

Appendices

The Appendix is divided into 10 sections. Appendix A presents a general ‘price-theory’ style model that illustrates how we measure the WTP for each group of beneficiaries of a policy and the net cost to the government. Appendix B presents the formal learning-by-doing model and derives the implications for the willingness-to-pay and government cost that enter the MVPF framework. Appendix C provides a detailed description of our measures of environmental externalities. Appendix D discusses how we incorporate rebound effects. Appendix E provides a detailed discussion of how we construct each MVPF in our sample. Appendix F provides details on our test of and correction for publication bias. Appendix G examines regulatory policy. It shows how we can use the MVPF framework to study whether tax or subsidy policies are more efficient at delivering environmental benefits than regulations targeting similar types of emissions. Appendix H discusses the distinction between the MVPF approach and more traditional benefit-cost metrics such as net social benefits and the benefit-cost ratio. Appendix I provides a detailed description of our construction of the resource cost per ton metrics for each policy in our sample. Finally, Appendix J discusses patterns of US environmental spending over the last 15 years. It compares spending under American Reinvestment and Recovery Act (ARRA) and Inflation Reduction Act (IRA).

Contents

A Model Appendix: Setup	126
B Model Appendix: Learning by Doing	129
B.1 Moving Forward in Time	132
B.2 Isoelastic Specification	134
B.3 Comparative Statics: Learning-by-Doing Effects Eventually Fade Out	136
C Externalities	139
C.1 Social Costs	139
C.2 Electric Grid Externalities	141
C.3 Natural Gas Externalities	145
C.4 Gasoline Externalities	146
D Rebound	155
E Policy Appendices	158
E.1 Wind Production Tax Credit	158
E.2 Solar Investment Tax Credit	167
E.3 Battery Electric Vehicles	178

E.4	Hybrid Electric Vehicles	202
E.5	Appliance Rebates	217
E.6	Weatherization	229
E.7	Vehicle Retirement	238
E.8	Home Energy Reports	254
E.9	Other Nudges	262
E.10	Gasoline Taxes	270
E.11	Other Fuel Taxes	292
E.12	Other Revenue Raisers	316
E.13	Cap and Trade	320
E.14	International Subsidies	331
E.15	International Rebates	340
E.16	International Nudges	342
E.17	Other Subsidies	343
F	Publication Bias	346
F.1	Estimating Publication Bias	346
F.2	Correcting for Publication Bias	347
G	Regulation	348
G.1	Corporate Average Fuel Economy Standards (CAFE)	349
G.2	Renewable Portfolio Standards (RPS)	352
H	Comparison to Net Benefits and Benefit-Cost Ratios	354
I	Resource Cost per Ton	355
I.1	Electric Vehicles	355
I.2	Wind	356
I.3	Solar	357
I.4	Appliance Rebates	358
I.5	Vehicle Retirement	361
I.6	Weatherization	363
I.7	Hybrid Vehicles	364
I.8	Home Energy Reports	365
I.9	Gasoline Taxes	366
I.10	Other Fuel Taxes	366

I.11	Other Revenue Raisers	367
I.12	Government Cost per Ton	367
I.13	Net Social Cost per Ton	368
J	Federal Energy Policy over the last 15 years	369
J.1	ARRA	369
J.2	IRA	370

A Model Appendix: Setup

The MVPF framework requires measuring the willingness-to-pay for each group in society along with the net cost to the government. In this Appendix, we develop a rich model structure that allows us to illustrate how straightforward applications of price theory allow us to measure the WTP of each individual for a policy change along with the net cost to the government. The model structure extends the discussion in Hendren & Sprung-Keyser (2020) to provide a general characterization of the MVPF in the presence of externalities. We include what one might call “traditional” externalities, such as pollution and congestion. In addition to these externalities, we also allow for imperfect competition, so that a marginal increase in demand can increase firm profits. We also include what one might call “production” externalities whereby the production of a good by one firm can induce learning by doing that lowers the marginal cost of production for all other firms. These learning-by-doing effects (Thompson 2012, Nagy et al. 2013, Farmer & Lafond 2016, Way et al. 2022) have often been cited as motivation for production subsidies for new technology that addresses climate change (Gillingham & Stock 2018).

Finally, we use the model to help think about how to move from partial equilibrium causal effects of a policy to general equilibrium impacts of policies through changes in prices. One particularly relevant channel in our setting is the so-called “rebound effect” whereby a policy that generates an increase (or decrease) in electricity demand will cause the price of energy to increase (or decrease), leading to further changes in the consumption of dirty and clean goods.⁹⁹

We assume each individual consumes a vector of goods, \mathbf{x} , which have consumer prices \mathbf{p} , producer prices \mathbf{q} , and consumer taxes \mathbf{t} (or subsidies), where $\mathbf{t}=\mathbf{p}-\mathbf{q}$. We assume goods are indexed by both type and time so that dimensions of the goods and prices differ over time and across goods. For example, $x(t) \in \mathbf{x}$ could be the consumption of electric vehicles at time t , where consumption in each time period is a separate element of \mathbf{x} .¹⁰⁰ For convenience, we use notation suggesting that \mathbf{x} is finite dimensional, but will find it convenient to allow time to be continuous in the section below that measures learning-by-doing externalities. The individual is also affected by a vector of externalities, \mathbf{e} , which impose a monetized harm of $\mathbf{v}_i * \mathbf{e}$ on individual i , where \mathbf{v}_i is a vector of valuations of the externality from individual i and “*” represents the dot product. For example, \mathbf{e} can contain measures of the quality of the

⁹⁹We focus on general equilibrium effects that arise from the causal effect of the policy on prices of the good. However, the changes we estimate will not typically include the full array of general equilibrium effects of a policy on all prices and quantities. Nonetheless, the framework illustrates that such effects would be important if they affect emissions (so that they affect aggregate WTP) or tax revenue (so that they affect government costs).

¹⁰⁰We do not directly discuss worker wages, they are incorporated by thinking of labor as a good with a negative price (i.e., paid by firms to workers instead of from individuals to firms).

climate in 2050, commute times in New York on a particular day in 2020, and the presence of PM2.5 in Beijing in 2030. Different individuals will naturally have different valuations, \mathbf{v}_i , of these externalities. These valuations come from the assumption that individuals maximize a well-behaved utility function $u_i(\mathbf{x}; \mathbf{e})$ subject to a budget constraint, $\mathbf{p}^* \mathbf{x} \leq m_i$, where m_i is unearned income of individual i . Given this maximization program, we define $\mathbf{v}_i = \frac{1}{\lambda_i} \nabla_{\mathbf{e}} u(x^*; e)$, where $\nabla_{\mathbf{e}}$ is the gradient of the utility function with respect to each externality, evaluated at the optimal bundle x^* and λ_i is the Lagrange multiplier on individual i 's budget constraint. The intuition is that $\nabla_{\mathbf{e}} u$ measures how the externalities affect utility and λ_i changes units from utils into dollars. Finally, we assume the vector of goods in the economy is produced by a composite firm that has pre-tax profits Π and faces a tax rate τ_c . Individual i owns a share s_i in after-tax profits, generating payments $(1 - \tau_c)s_i\Pi$. With these assumptions, the envelope theorem implies that the willingness-to-pay of individual i for a policy change is:

$$WTP_i = \mathbf{x}_i * \mathbf{d}\mathbf{p} + \mathbf{v}_i * \mathbf{d}\mathbf{e} + (1 - \tau_c)s_i d\Pi. \quad (31)$$

There are three reasons that individuals are willing to pay for a policy change: (a) it makes the goods they consume cheaper, $\mathbf{x}_i * \mathbf{d}\mathbf{p}$, where $\mathbf{d}\mathbf{p}$ is the causal effect of the policy on prices; (b) it changes the value of the externalities they experience, $\mathbf{v}_i * \mathbf{d}\mathbf{e}$, where $\mathbf{d}\mathbf{e}$ is the causal effect of the policy on the externalities; or (c) it changes the income they receive from firm profits, $(1 - \tau_c)s_i d\Pi$, where $d\Pi$ is the causal effect of the policy on firm profits.

We assume these profits arise from the production of goods consumed in the economy. Let $\mathbf{x} = \sum_i \mathbf{x}_i$ denote total production of goods in the economy. We assume there is a single representative firm with a marginal cost function, $\mathbf{c}(\mathbf{x})$, so that market profits are $\Pi(\mathbf{x}) = \mathbf{x}^*(\mathbf{q} - \mathbf{c}(\mathbf{x}))$. The policy impact on firm profits is:

$$d\Pi = \mathbf{d}\mathbf{x}^*(\mathbf{q} - \mathbf{c}(\mathbf{x})) + \mathbf{x}^*(\mathbf{d}\mathbf{q} - \nabla \mathbf{c}(\mathbf{x}) \cdot \mathbf{d}\mathbf{x}) \quad (32)$$

where we let “ \cdot ” denote the Hadamard product (element-wise multiplication), to contrast it with “ $*$ ” that denotes the standard dot product multiplication, and $\nabla \mathbf{c}(\mathbf{x})$ denotes the gradient of the cost function. The first term is the change in consumption multiplied by the firm markup. This sums across the change in production of each good multiplied by the markup for that good. This would be zero under perfect competition ($\mathbf{q} = \mathbf{c}(\mathbf{x})$), but under imperfect competition increasing firm demand leads to higher profits. The second term is the impact of the policy on producer markups (prices minus costs). If the policy increases (decreases) producer prices, $\mathbf{d}\mathbf{q}$, this increases (decreases) firm profits. If the policy increases firm costs, this reduces firm profits proportional to the Hadamard product of the derivatives of the cost function, $\nabla \mathbf{c}$, and the change in production of each good, $\mathbf{d}\mathbf{x}$. Note that the generality of the cost function notation means we are allowing for consumption of one good (e.g., EVs today) to affect marginal costs of another good (e.g., EVs in the future), a feature we discuss further in the next section.

We assume externalities arise from the production or consumption of the vector of goods, \mathbf{x} . For example, increased gasoline consumption has some impact on the vector of pollutants, such as CO_2 , SO_2 , and congestion. Producing electricity using solar or wind power instead of coal can reduce $PM_{2.5}$ in addition to CO_2 and other pollutants. We therefore model our vector of pollutants \mathbf{e} as a vector-valued function $\mathbf{e} = \mathbf{E}(\mathbf{x})$, and we let $\nabla \mathbf{E}$ denote its Jacobian matrix. Each individual i 's willingness-to-pay for the sum of the changes in pollution that arise from changes in consumption and production of \mathbf{x} is given by:

$$\mathbf{v}_i * \mathbf{d}\mathbf{e} = \mathbf{v}_i * \nabla \mathbf{E} * \mathbf{d}\mathbf{x} \quad (33)$$

where $\mathbf{v}_i * \nabla E$ is the vector of costs to individual i of the consumption of \mathbf{x} in the economy – the matrix sums across the externalities produced from each good in the economy and multiplies by each individual’s valuation, \mathbf{v}_i , of those externalities. It is important to note that \mathbf{E} is a vector and equation 33 is summing across all the possible externalities experienced by individual i . This means we allow for individuals to experience externalities very differently.¹⁰¹ In implementation, we often sum across many individuals when forming the environmental externalities, but will delineate amongst subgroups wherever possible (e.g., when an SCC model allows us to think about benefits to different regions/countries/generations).

We also will allow for environmental externalities to affect the government budget in addition to directly affecting individuals. For example, the DICE and RICE models report damages in GDP or GDP-equivalent units (Nordhaus 1993). If we consider these as impacting productivity, it suggests carbon decreases global economic output by \$SCC per ton of carbon. Globally, 15% of this incidence falls on the US. With a 30% tax rate,¹⁰² this suggests government tax revenue declines by 4.5% of the SCC per ton of carbon emitted today. Other models of carbon damages have different incidence: Rennert et al. (2022) suggests emissions lead to lost lives in the US and reductions in the productivity of agriculture, but no negative impact on US GDP (and thus no impact on tax revenue). Our approach will consider multiple models of carbon damages in our analysis and explore the robustness of our results; the key point here is that our framework asks us to think about not just the magnitude but also the incidence of the damages from carbon emissions.

Translating the impact of environmental harms on the government budget, we assume the government taxes goods and services, \mathbf{x} , and profits, Π , so that the net impact on the government budget of the policy is

$$Cost = \mathbf{x} * d\mathbf{t} + \mathbf{t} * d\mathbf{x} + \tau^c(d\mathbf{x}*(\mathbf{q}-\mathbf{c}(\mathbf{x})) + \mathbf{x}*(d\mathbf{q} - \nabla\mathbf{c}(\mathbf{x}) \cdot d\mathbf{x})) \quad (34)$$

This is equivalent to the sum of the mechanical cost of any change to the subsidies or taxes ($\mathbf{x} * d\mathbf{t}$), the impact of the behavioral response on the cost of subsidies ($\mathbf{t} * d\mathbf{x}$), and the impact of the policy on profits multiplied by the tax rate on capital income (τ^c), yielding revenue $\tau^c d\Pi$. The environmental impacts noted above are captured by the fact that taxed behavior (\mathbf{x}) changes in the future in response to carbon emissions today – a feature we discuss further in our implementation below. Equations 31 and 34 are the core components feeding into the construction of the WTP and cost components needed for our welfare analysis.

Causal Effects: Partial vs. General Equilibrium

Measuring the WTP and cost of the policy requires measuring the causal effects of the policy change on \mathbf{x} , which we have denoted by $d\mathbf{x}$. A casual glance at the equations might suggest that one can use “reduced form” evidence on the effect of the policy on allocations \mathbf{x} without worrying about the impact of subsequent general equilibrium effects or other changes in behavior. This interpretation, however, is not generally true because the “ $d\mathbf{x}$ ” term needs to reflect the full causal effect of the policy change. As discussed above, this includes the long-run impact of emissions today on future taxed behavior so that we accurately measure government costs. But even absent these dynamics, we also must include any spillover or “general equilibrium”

¹⁰¹For example, one element of \mathbf{e} could be commute times in NYC; another element can be the daily temperature in Kenya in 2050. New Yorkers may value their commute times but not care about the temperature in Nairobi in 2050. Farmers in Kenya in 2050 might care about their daily temperature, but not be as concerned with how long it takes an investment banker in NYC to get to work.

¹⁰²This further assumes government and private discount rates are equivalent.

impacts that are not captured by an RCT or quasi-experimental analysis.

For example, consider an electric vehicle subsidy that increases purchases today. This increase is readily measured in RCTs and quasi-experimental studies. However, purchasing more EVs can lead to a reduction in gasoline demand. This in turn can lead to a reduction in the price of gasoline which can increase driving of gasoline powered vehicles – a so-called “rebound effect” in the energy economics literature. Conversely, EV purchases may increase electricity demand causing electric prices to rise, reducing electricity consumption – a reverse rebound effect, so to speak. If we know how much an EV changes energy demand for electricity and gasoline, we can measure the size of these “rebound” effects using additional information on market supply and demand elasticities. These will be central components of our empirical analysis.

B Model Appendix: Learning by Doing

A more complicated way in which price changes can affect demand is via learning-by-doing externalities. This Appendix provides the mathematical details on our new sufficient statistics result (Theorem 1, introduced in the main text and stated precisely in a generalized form below) that translates learning-by-doing curves, demand curves, and an assumption about market equilibrium into a formal statement about society’s willingness-to-pay for the dynamic effects of policies that increase consumption of these goods today. Before delving into the analysis, it is useful to start by noting the model already allows for learning by doing through the general cost function $\mathbf{c}(\mathbf{x})$. The Jacobian of this cost function, $\nabla\mathbf{c}(\mathbf{x})$, specifies how changes in the production of one good (e.g., solar panels today) affects the cost of producing other goods (e.g., solar panels in the future).

The basic idea of our approach is to write out a cost function that follows the shape in Appendix Figure 1 and then solve for the impact on WTP and cost. Importantly, learning-by-doing effects means there will be indirect effects from the fact that a subsidy today can cause an increase in consumption of the good in the future (e.g., after that subsidy has ended). In other words, the causal effect \mathbf{dx} will have not only static components from when the policy operates but dynamic components from long run impacts on the cost of production.

We focus on a policy change that increases a subsidy for one particular good, which we call a ‘green good.’¹⁰³ We denote market-level consumption of this good at time t by $x(t) \in \mathbf{x}$, and an individual’s consumption by $x_i(t)$. We consider a policy change that increases the subsidy for this good from $\tau(t)$ to $\tau(t) + d\tau$ starting at some time $t^* - \eta$ that lasts for a length of time η and thus ends in period t^* . Without loss of generality, we normalize $t^* = 0$. Therefore, the subsidy change is in operation over the time window $[-\eta, 0]$. Later, it will be helpful to consider the limiting behavior as η and $d\tau$ become small.

The subsidy change, $d\tau$, over this time window $[-\eta, 0]$ has a causal effect on the market-level consumption of the green good at each time t , which we denote $dx(t)$.¹⁰⁴ Formally, $dx(t)$ is a (Fréchet) differential of the time path of consumption of $x(t)$ with respect to the subsidy

¹⁰³A similar derivation applies to other policies that increase consumption of a good today that has learning-by-doing effects.

¹⁰⁴We assume for now that the policy was unanticipated so that there is no causal effect prior to t^* but our approach can be generalized to include such anticipatory effects.

change, t^* , and η .¹⁰⁵ In addition, the subsidy also has an effect on consumer prices, $p(t)$, at each time t , which we denote by $dp(t)$.

First, consider the cost of this policy to the government if no one changed their behavior, also known as the “mechanical” cost of the policy. This is equal to the product of the size of the subsidy change, $d\tau$, the length of the change, η , and the flow of goods subject to the change, which for small η is equal to $x_i(0)$. Combining, this is $\eta d\tau x_i(0)$. We assume a pass through rate of γ of the subsidy to consumer prices $dp = -\gamma d\tau$. Absent direct estimates of pass through, we assume full pass through $\gamma = 1$; but we relax this assumption for alternative specifications where empirical evidence suggests incomplete pass through.

Next, consider the impact of the behavioral responses to the policy $dx(t)$ while the subsidy is in operation during $[-\eta, 0]$. This change generates environmental externalities that arise both because of the direct purchase of the new good but also because the purchase of the good offsets purchases of other goods (e.g., an EV purchase leads to lower gas consumption). To economize notation, let $\nu_i(t)$ denote the sum of the value of the externalities experienced by individual i per unit of change in the consumption of the good at time t , so that the environmental externality on individual i is given by $\nu_i(t)dx(t)$, and we let $\nu(t) = \sum_i \nu_i(t)$ denote the full externality.¹⁰⁶

In addition to environmental externalities, the subsidy also can affect firm profits in a non-competitive environment. To simplify the model and focus on learning-by-doing, we assume a parsimonious model of firm behavior. In particular, we assume that prices are set at a constant markup μ over marginal costs, so that $p = (\mu + 1)c$ in each time period, with our baseline case of $\mu = 0$ corresponding to perfect competition. Given that dynamically optimizing firms would partially internalize learning-by-doing externalities, we view our approach as an upper bound on the willingness-to-pay generated from learning-by-doing effects.

Finally, the change in consumption of the green good affects government costs proportional to the pre-existing subsidy, $\tau(t)dx(t)$. We assume for exposition this is the only fiscal externality, but relax this assumption in our empirical implementation (e.g., we account for lost gas tax revenue when people buy more EVs). We also assume for simplicity in the exposition that there is no subsidy in operation after $t = 0$, although we again relax this in our empirical implementations where relevant.¹⁰⁷

Static Benchmark Before turning to the dynamic components of the MVPF, it is helpful to establish a benchmark MVPF in the absence of dynamic cost curve effects. In this case, $dx(t) = 0$ for $t > 0$ and thus the terms discussed to this point allow us to construct the MVPF. The WTP is given by the sum of the mechanical benefit of the policy and the environmental externalities, while the government cost is the mechanical benefit plus the fiscal externality. These can be written as

$$MVPF = \frac{1 + SE}{1 + FE} \quad (35)$$

¹⁰⁵Note that as $\eta \rightarrow 0$, $dx(t) \rightarrow 0$. As in traditional calculus of variations approaches, the ratio of $dx(t)$ to η is what will matter for our analysis.

¹⁰⁶In the notation of our model, let $\mathbf{dx}(t)$ denote the vector of changes of time t variables but that has zeros everywhere else in \mathbf{x} . Then, $\nu_i(t) = \mathbf{v}_i \nabla \mathbf{E} \mathbf{dx}(t)$

¹⁰⁷Adding existing subsidies in place after t^* changes the structure of the differential equation governing our analysis such that there is no longer a closed form solution. In this case, we solve the ODE numerically out to a large time horizon.

where

$$SE = \frac{\nu(0)}{p(0)}\epsilon \quad (36)$$

is the static externality benefit from additional consumption of x and

$$FE = \frac{\tau(0)}{p(0)}\epsilon \quad (37)$$

is the fiscal externality impact of additional consumption of x . Here, the elasticity $\epsilon = \frac{dx(0)}{dp(0)} \frac{p(0)}{x(0)}$ is the ratio of the percent change in x relative to the percent change in prices due to the subsidy. Comparing the numerator and denominator of the MVPF, note that we have $MVPF = 1$ whenever the subsidy is at its Pigouvian optimal level, $\tau(0) = e(0)$. When the existing subsidy is less than this, the MVPF will exceed 1, indicating the value of a slightly higher subsidy exceeds its cost to the government.

Dynamic MVPF Having established this static benchmark, now suppose that the subsidy today has dynamic effects. This introduces two additional types of externalities. The first arises because the additional consumption of x today leads to lower marginal costs in the future. Motivated by Appendix Figure 1, we assume that the marginal cost of producing the good $x(t)$ is given by $c(X(t))$, where $X(t) = \int_0^t x(s)ds + x(0)$ is cumulative production at time t . With this expression, the causal effect of the subsidy near $t = 0$ of costs in period $t > 0$ per dollar of the mechanical cost of the policy, $x(0)\eta d\tau$, is given by $\gamma \frac{d[c(X(t))]}{x(0)\eta d\tau}$. This price reduction is valued depending on how much x individual i consumes in period t . Discounting using the real discount rate ρ yields a valuation from these price reductions of

$$DP_i = \int_0^\infty -x_i(t) \frac{d[p(X(t))]}{x(0)\eta d\tau} e^{-\rho(t)} dt \quad (38)$$

where “DP” stands for the dynamic price reduction generated from the response.¹⁰⁸ Meanwhile, firms have a willingness-to-pay of

$$D\pi = \int_0^\infty \frac{d[\pi(X(t), x(t))]}{x(0)\eta d\tau} e^{-\rho(t)} dt \quad (39)$$

where $\pi((X(t), x(t) = \mu c(X(t))x(t)$ are firm’s profits. In addition to the price response, the lower costs lead to greater consumption of x given by $\frac{dx(t)}{x(0)\eta d\tau}$ per dollar of mechanical spending on the policy. Individual i values this change in x at time t according to $\nu_i(t)$, leading to a PDV of benefits:

$$DE_i = \int_0^\infty \left[\frac{dx(t)}{x(0)\eta d\tau} \nu_i(t) \right] e^{-\rho(t)} dt \quad (40)$$

Turning to the government costs, just as the behavioral response at $t^* = 0$ affects government revenue, so does the behavioral response for $t > 0$. This effect depends on the size of the subsidies in place, $\tau(t)$, after $t = 0$ (which we assume for simplicity are zero in our baseline specification) and also any impacts from the future environmental quality on tax revenue received by the government, which we denote as $\nu_g(t)$.¹⁰⁹ The PDV impact on government costs

¹⁰⁸By the envelope theorem, the willingness-to-pay for future marginal consumption due to lower prices is zero.

¹⁰⁹This term is given by the impact of the policy today on future consumption of goods in the economy,

is then

$$DFE = \int_0^\infty \frac{dx(t)}{x(0)\eta d\tau} \nu_g(t) e^{-\rho(t)} dt \quad (41)$$

where we would replace $\nu_g(t)$ with $\tau(t) - \nu_g(t)$ in the presence of pre-existing subsidies.

Summing together, we arrive at the MVPF inclusive of these dynamic effects:

$$MVPF = \frac{1 + SE + DP + DE}{1 + FE + DFE} \quad (42)$$

which is equivalent to the above but now includes the impact of the policy today on future prices and environmental externalities. Now, the key question is: how do we measure these dynamic terms in the equation above?

B.1 Moving Forward in Time

In general, measuring the response of future prices and consumption is quite complex. However, in our model this task is simplified by the fact that the subsidy essentially “moves us forward in time.” To see, this, note that the subsidy in place over $[-\eta, 0]$ induces an increase in the initial stock of cumulative consumption ($X(0)$) and contemporaneous consumption ($x(0)$) to $X(0)'$ and $x(0)'$ in the post-subsidy period. Since cumulative production is continuous and strictly increasing over time, there exists a time $\bar{t} > 0$ such that $X(0)' = X(\bar{t})$. Because the ordinary differential equation (ODE) governing $X(t), x(t)$ is autonomous – depending on the time index only indirectly through X and x , this shift forward in the initial condition fully characterizes how the production paths change with a shock to the initial conditions.¹¹⁰

What remains then is to characterize how the initial conditions (the starting stock and flow of production in the post-subsidy period) change with an infinitesimal subsidy change over $[-\eta, 0]$. Formally, let $dX(0)$ denote the impact of the policy change on cumulative production at time 0. By definition of $X(t)$, we have that $dX(t) = X'(t)dt = x(t)dt$. Note that for small η , we can also write $dX(0)$ as

$$dX(0) \approx -\gamma\eta\epsilon \frac{d\tau}{p(0)} x(0) \quad (43)$$

where.¹¹¹ Intuitively, the change in cumulative consumption is given by the change in flow consumption from a change in prices, $\frac{dx(t)}{dp(t)} = \epsilon^p \frac{x(t)}{p(t)}$, multiplied by the subsidy change, $d\tau$, and

multiplied by the tax rate on those goods and services – i.e. the $\mathbf{t} * \mathbf{dx}$ term in our government cost equation but focusing on the components where $t > 0$.

¹¹⁰This is because autonomous ODEs exhibit “horizontal invariance”. That is, if $X(t)$ solves the autonomous ODE satisfying the initial condition $X(t_0) = X_0$, then $X(t + t_0)$ solves the same ODE with initial condition $X(0) = X_0$.

¹¹¹To see this, note that we can write

$$\begin{aligned} dX(0) &= \int_{-\eta}^0 dx(t) dt \\ &= \int_{-\eta}^0 -\gamma\epsilon x(t) \frac{d\tau}{p(t)} dt \end{aligned}$$

where $\epsilon \frac{x(t)}{p(t)} = \frac{dx(t)}{dp(t)}$

then cumulated over the length of the subsidy η . For small η , \approx holds exactly if we divide each side by η and take the limit as $\eta \rightarrow 0$ as we can approximate the flows using just the response measured at $t = 0$. This means we can think of the policy as moving us forward in time by

$$dt = -\gamma\eta\epsilon\frac{d\tau}{p(0)} \quad (44)$$

The subsidy today “pushes us down the cost curve” by an amount of time that is proportional to the elasticity of demand operating during the subsidy (ϵ), the length of time the subsidy is in place (η), and the size of the subsidy as a share of the price ($\frac{d\tau}{p(0)}$).

How does the increase in cumulative production affect costs (and thus prices) in future periods? Note that because marginal cost is given by $c(X(t))$, the derivative of marginal costs with respect to time is $\frac{d}{dt}c(X(t)) = c_X(X(t))X'(t) = c_X(X(t))x(t)$. So, moving costs forward by dt yields a reduction in costs that is given by $\frac{d}{dt}c(X(t))dt = c_X(X(t))x(t)dt$.

Plugging in $dt = -\gamma\eta\epsilon\frac{d\tau}{p(0)}$, we have:

$$\frac{d[p(X(t))]}{\eta d\tau} = (\mu + 1)\frac{d[c(X(t))]}{\eta d\tau} = -(\mu + 1)\gamma\frac{\epsilon}{p(0)}c_X(X(t))x(t) \quad (45)$$

The impact of the policy today of size $\eta d\tau$ on future prices depend on how much it increases consumption today, ϵ , multiplied by $x(t)$, and normalized by the ratio of marginal costs in the future to the present, $c_X(X(t))/p(0)$. The key insight here is that equation 45 measures how marginal costs change in all future periods, $t > 0$, as a result of the subsidy levied in $[-\eta, 0]$. So, we can now use this to plug back into our formulas for the dynamic price component of the MVPF:

$$DP = (\mu + 1) \int_0^\infty \frac{x(t)}{x(0)} \lim_{\eta \rightarrow 0} \left[\frac{-d[c(X(t))]}{\eta d\tau} \right] e^{-\rho t} dt \quad (46)$$

$$= (\mu + 1)\gamma \int_0^\infty \frac{x(t)}{x(0)} \left[\frac{\epsilon}{p(0)} c_X(X(t)) \right] e^{-\rho t} dt \quad (47)$$

$$= (\mu + 1)\gamma \int_0^\infty e^{-\rho t} \epsilon \left(\frac{c_X(X(t))}{c(X(t))} X(t) \right) \frac{c(X(t))}{c(X(0))} \frac{x(t)}{X(t)} \frac{x(t)}{x(0)} dt \quad (48)$$

The first line passes the limit to the variables that depend on η ($dp(t)$ and η), the second line plugs in equation 45, and the third line re-arranges terms and uses the fact that price is equal to marginal cost, $c(X(t)) = p(t)$, as recall we have assumed no subsidies for $t > 0$.

Turning next to firm profits, we have $\pi = \mu c(t)x(t)$ so that $d\pi = \mu (dc(t)x(t) + dx(t)c(t))$, therefore

$$D\pi = \mu \left(\int_0^\infty \frac{x(t)}{x(0)} \lim_{\eta \rightarrow 0} \left[\frac{dc(t)}{\eta d\tau} \right] e^{-\rho t} dt + \int_0^\infty \frac{c(t)}{x(0)} \lim_{\eta \rightarrow 0} \left[\frac{dx(t)}{\eta d\tau} \right] e^{-\rho t} dt \right) \quad (49)$$

$$= -\mu\gamma \left(\int_0^\infty \frac{x(t)}{x(0)} \frac{c_X(X(t))x(t)}{p(0)} \epsilon e^{-\rho t} dt + \int_0^\infty \frac{c(t)}{x(0)} \frac{\epsilon x'(t)}{p(0)} e^{-\rho t} dt \right) \quad (50)$$

Next, we turn to the dynamic externality term, DE . This is determined by how the subsidy

affects the time path of consumption of x , $dx(t)$. Recall that the policy change can be thought of as moving forward by $dt = -\gamma\eta\epsilon\frac{d\tau}{p(0)}$. So, we can think of the change in x at a point in time as following:

$$dx(t) = -\gamma\eta x'(t)\epsilon\frac{d\tau}{p(0)} \quad (51)$$

The intuition is that if x is increasing in time ($x'(t) > 0$) then moving down the cost curve leads to greater consumption at time t than in the world without the subsidy (in our setting, it is natural to envision that prices go down over time because marginal cost goes down over time, so the consumption of x increases over time, $x'(t) > 0$ for all t). The amount by which consumption goes up, $dx(t)$, is given by the slope of x multiplied by how far time moves forward as a result of the subsidy, $-\gamma\eta\epsilon\frac{d\tau}{p(0)}$. So, we can write DE as

$$DE = \int_0^\infty \nu(t) \lim_{\eta \rightarrow 0} \left[\frac{dx(t)}{x(0)\eta d\tau} \right] e^{-\rho t} dt \quad (52)$$

$$= \int_0^\infty \frac{\epsilon\nu(t)x'(t)}{x(0)c(X(0))} e^{-\rho t} dt \quad (53)$$

where the last line both substitutes equation 51 and uses the assumption that subsidies go away at $t = 0$ so that $c(X(0)) = p(0)$. Finally, replacing $\nu(t)$ with the government revenue component of the environmental externality yields DFE .

These equations for DP , DE , and DFE fully characterize the MVPF of environmental subsidies. We summarize the analysis above into the following Lemma.

Lemma 1. *Suppose there are no subsidies after $t = 0$ and price equals marginal cost for all periods $t > 0$. Then, the MVPF of a small subsidy ($d\tau \approx 0$ and $\eta \approx 0$) is given by*

$$MVPF = \frac{1 + SE + DP + DE}{1 + FE + DFE} \quad (54)$$

where the terms are defined as above.

B.2 Isoelastic Specification

So far, we have not imposed any functional forms on the structure of how cumulative production affects marginal costs or how prices affect demand. However, in order to estimate DP and DE , we need to be able to forecast the time path of future demand and costs, $x(t)$ and $c(X(t)) = p(t)$. To obtain our analytical solution for the future path of prices and consumption, we parameterize consumers' demand function by an isoelastic specification:

$$x(p(t)) = ap(t)^\epsilon \quad (55)$$

with $\epsilon < 0$. A one percent reduction in prices leads to an ϵ percent increase in demand.

For firms, we assume that each firm's marginal cost is also given by an isoelastic specification:

$$c(X) = \kappa X^\theta \quad (56)$$

A one percent increase in cumulative production leads to a θ percent decline in marginal costs. Under our assumption of constant markups, this in turn implies a θ percent reduction in prices.

The second key insight in our framework is that we can combine equations (1) and (2) to yield

$$\frac{d}{dt} \log(x) = \epsilon \frac{d}{dt} \log(p) = \frac{d}{dt} \frac{d \log(c(X))}{d \log(X)} \frac{d \log(X)}{dt} = \epsilon \theta \frac{x(t)}{X(t)} \quad (57)$$

where the second equality uses the fact that at the no-subsidy baseline, consumer prices are equal to marginal costs in each period. Recall $x(t) = X'(t)$, which means we can write the evolution of production as a 2nd order ordinary differential equation (ODE):

$$\frac{X''(t)}{X'(t)} = \epsilon \theta \frac{X'(t)}{X(t)} \quad (58)$$

Equation 58 characterizes how consumption of x evolves over time as a function of the demand and cost curve elasticities.

Recalling that $t = 0$ corresponds to the end of the hypothetical subsidy increase period, we impose the initial conditions $X(0) = X_0$, $x(0) = x_0$ for where X_0, x_0 are contemporaneous and cumulative production at the time at which we calculate the dynamic externalities (i.e., in context or in 2020), which we observe in the data.¹¹² This yields a general closed-form solution for $X(t)$ given by

$$X(t) = C_1(t + C_2)^{\frac{1}{1-\epsilon\theta}} \quad (59)$$

where $C_1, C_2 \in \mathbb{R}^+$ are pinned down by the initial conditions, and therefore

$$x(t) = \frac{C_1}{1 - \epsilon\theta} (t + C_2)^{\frac{-\epsilon\theta}{1-\epsilon\theta}} \quad (60)$$

Having solved this ODE, we now have a closed form expression for the MVPF.

Theorem 1 (Generalized Version) (Iso-elastic Specification). *Suppose demand is given by equation 55 and the marginal cost is given by equation 56. Then,*

$$DP = \gamma(\mu + 1) \frac{\theta\epsilon}{1 - \theta\epsilon} C_2^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}} \int_0^\infty e^{-\rho t} (t + C_2)^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}} dt \quad (61)$$

where the constant, C_2 , is identified from cumulative and flow production,

$$C_2 = \frac{X(0)}{x(0)(1 - \epsilon\theta)} \quad (62)$$

while

$$DE = -\frac{\gamma\theta\epsilon^2}{(1 - \epsilon\theta)c(X(0))} C_2^{-\frac{\epsilon\theta}{1-\epsilon\theta}} \int_0^\infty e^{-\rho t} (t + C_2)^{\frac{2\epsilon\theta-1}{1-\epsilon\theta}} \nu(t) dt \quad (63)$$

where $c(X(0))$ is marginal costs at the point at which we estimate the dynamic externalities, and DFE follows the same form as DE replacing $\nu(t)$ with the government budget externality per unit of $x(t)$.

Theorem 1 yields the MVPF for subsidies of a green good in the presence of learning-

¹¹²Given that we observe data in yearly increments, we define cumulative production to be lagged cumulative production, taking the sum of yearly production in all prior years. This is to capture the fact that we model learning by doing and not static economies of scale, such that contemporaneous production does not affect contemporaneous marginal costs. Put differently, this matches what we would obtain in the discrete time version of our model.

by-doing externalities. This generalized version corresponds to Theorem 1 in the main text when $\gamma = 1, \mu = 0$. We also express C_2 as the “starting time” t^* of the policy (rather than normalizing $t^* = 0$) to stress the interpretation of this parameter as how far along the cost curve a technology is at the time we consider a marginal subsidy.

The theorem shows that we need to know 3 key parameters (a) the elasticity of demand with respect to price, ϵ , (b) the elasticity of marginal cost with respect to cumulative production, θ , and the ratio of cumulative production to flow production at the time of the subsidy change adjusted by these first two parameters, $\frac{X(0)}{x(0)(1-\epsilon\theta)}$.

B.3 Comparative Statics: Learning-by-Doing Effects Eventually Fade Out

Because our model has a closed-form solution in our baseline case, it is possible to perform comparative statics to generate intuitions for what drives these effects. It is already straightforward to see that LBD effects are generally increasing in the magnitude of the demand elasticity and slope of the cost curve, ϵ and θ . Here, we provide one additional comparative static that is perhaps less ex-ante obvious from the formula but can be seen with some further math: learning-by-doing effects eventually fade out over time. This means LBD effects potentially provide a rationale for early subsidies that are limited in time or cumulative production of the good. We formalize this in Theorem 2.

Theorem 2 (Comparative statics) We have that $\lim_{C_2 \rightarrow \infty} DP = 0$, and there exists \bar{C} such that DP is strictly decreasing for C_2 greater than \bar{C} .

Proof For reference, we reproduce the expression for

$$DP = \gamma(\mu + 1) \frac{\epsilon\theta}{1 - \epsilon\theta} C_2^{-\theta \frac{1+\epsilon}{1-\epsilon\theta}} \int_0^\infty e^{-\rho t} (t + C_2)^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt.$$

Note that $\underbrace{\gamma}_{>0} \underbrace{(\mu + 1)}_{\geq 0} \underbrace{\frac{\epsilon\theta}{1 - \epsilon\theta}}_{\in(0,1)} > 0$ under our assumptions; therefore, it suffices to show

$$\frac{d}{dC_2} \left(C_2^{-\theta \frac{1+\epsilon}{1-\epsilon\theta}} \int_0^\infty e^{-\rho t} (t + C_2)^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt \right) < 0. \quad 113$$

The derivative of interest is

$$\begin{aligned} & -\theta \frac{(1 + \epsilon)}{1 - \epsilon\theta} C_2^{-1-\theta \frac{1+\epsilon}{1-\epsilon\theta}} \int_0^\infty e^{-\rho t} (t + C_2)^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt \\ & + C_2^{-\theta \frac{1+\epsilon}{1-\epsilon\theta}} \left(-1 + \theta \frac{1 + \epsilon}{1 - \epsilon\theta} \right) \int_0^\infty e^{-\rho t} (t + C_2)^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt, \end{aligned}$$

which has the same sign as

¹¹³Note that we can reexpress $\int_0^\infty e^{-\rho t} (t + C_2)^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt$ as $\int_{C_2}^\infty e^{-\rho(t-C_2)} t^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt$. This formulation will frequently prove useful.

$$- \underbrace{\int_0^\infty e^{-\rho t} (t + C_2)^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt}_A + \theta \frac{(1+\epsilon)}{1-\epsilon\theta} \underbrace{\left(\int_0^\infty e^{-\rho t} (t + C_2)^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt - C_2^{-1} \int_0^\infty e^{-\rho t} (t + C_2)^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt \right)}_B$$

The first term, A, is clearly negative. For the second term, B, note that $\int_0^\infty e^{-\rho t} (t + C_2)^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt \leq C_2^{-1} \int_0^\infty e^{-\rho t} (t + C_2)^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt$. Therefore, B has the same sign as $-\theta \frac{(1+\epsilon)}{1-\epsilon\theta}$. When $\theta \frac{(1+\epsilon)}{1-\epsilon\theta} > 0$, both terms are thus negative and we are done. When $\theta \frac{(1+\epsilon)}{1-\epsilon\theta} < 0$ so that B is negative, we show that it nevertheless becomes asymptotically negligible relative to A, implying there exists a cutoff \bar{C} above which the statement holds. To see this, consider

$$\begin{aligned} & \lim_{C_2 \rightarrow \infty} \frac{\int_0^\infty e^{-\rho t} (t + C_2)^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt - C_2^{-1} \int_0^\infty e^{-\rho t} (t + C_2)^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt}{-\int_0^\infty e^{-\rho t} (t + C_2)^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt} \\ &= \lim_{C_2 \rightarrow \infty} \frac{\int_{C_2}^\infty e^{-\rho t} t^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt - C_2^{-1} \int_{C_2}^\infty e^{-\rho t} t^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt}{-\int_{C_2}^\infty e^{-\rho t} t^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt} \end{aligned}$$

where the RHS follows after canceling out common terms $e^{\rho C_2}$. By the preceding analysis, both the numerator and denominator are negative for any $C_2 > 0$, so the limit is at least 0. Now we show it is at most 0. Direct substitution yields an indeterminate form of $\frac{0}{0}$. Applying l'Hopital's rule to the RHS above and using Leibniz' rule for differentiation under the integral sign yields

$$\begin{aligned} & \lim_{C_2 \rightarrow \infty} \frac{-e^{-\rho C_2} C_2^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}} + C_2^{-1} e^{-\rho C_2} C_2^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} + \frac{1}{C_2^2} \int_{C_2}^\infty e^{-\rho t} t^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt}{e^{-\rho C_2} C_2^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}}} \\ &= \lim_{C_2 \rightarrow \infty} \frac{\int_{C_2}^\infty e^{-\rho t} t^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt}{C_2^2 + e^{-\rho C_2} C_2^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}}} \end{aligned}$$

This again yields an indeterminate form of $\frac{0}{0}$, but, noting that $\int_{C_2}^\infty e^{-\rho t} t^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt \leq C_2^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} \int_{C_2}^\infty e^{-\rho t} dt = \frac{e^{-\rho C_2}}{\rho} C_2^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}}$ (because $t^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}}$ is decreasing for $\theta \frac{1+\epsilon}{1-\epsilon\theta} < 1$), we have that by monotonicity of limits, the above centered expression is bounded above by

$$\begin{aligned} & \lim_{C_2 \rightarrow \infty} \frac{\frac{e^{-\rho C_2}}{\rho} C_2^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}}}{C_2^2 e^{-\rho C_2} C_2^{-2+\theta \frac{1+\epsilon}{1-\epsilon\theta}}} \\ &= \lim_{C_2 \rightarrow \infty} \frac{1}{\rho C_2} = 0 \end{aligned}$$

This concludes the analysis of decreasingness. Note that decreasingness is not sufficient to establish that the limit is 0, since the derivative could become arbitrarily small. To show that

DP converges to 0 as C_2 grows large, clearly, it suffices to show that

$$\int_0^\infty e^{-\rho t} (t + C_2)^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt \leq KC_2^{\lceil -1+\theta \frac{1+\epsilon}{1-\epsilon\theta} \rceil}$$

for $t^* > e$ for some constant K independent of t^* . If this is the case, then

$$\frac{\theta\epsilon}{1-\theta\epsilon} C_2^{-\theta \frac{1+\epsilon}{1-\epsilon\theta}} \int_0^\infty e^{-\rho t} (t + C_2)^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt \leq KC_2^{\lceil -1+\theta \frac{1+\epsilon}{1-\epsilon\theta} \rceil - \theta \frac{1+\epsilon}{1-\epsilon\theta}}$$

Since $\lceil x \rceil < x + 1$, this term clearly converges to 0 as C_2 goes to infinity. Rewriting $\int_0^\infty e^{-\rho t} (t + C_2)^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt = e^{\rho C_2} \int_{C_2}^\infty e^{-\rho t} t^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt$, we note that if $\theta \frac{1+\epsilon}{1-\epsilon\theta} < 1$, then $\int_{C_2}^\infty e^{-\rho t} t^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt \leq C_2^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} \frac{1}{\rho} e^{-\rho C_2}$ by monotonicity of the integral and the fact that t^α is decreasing for $\alpha < 0$.

In the remaining case where $\theta \frac{1+\epsilon}{1-\epsilon\theta} > 1$, we rely on the following lemma: $\int_x^\infty e^{-z} z^{a-1} dz \leq ae^{-x} x^{a-1}, \forall a \in \mathbb{N} \setminus \{0\}$ and $x > a$. This is proven via induction on a . In the base case, start with $a = 1$. We have $\int_x^\infty e^{-z} z^{a-1} dz = \int_x^\infty e^{-z} dz = e^{-x} = ae^{-x} x^{a-1}$.

In the inductive step, consider $\int_x^\infty e^{-z} z^{a-1}$ for some $a \in \mathbb{N} \setminus \{0\}$ where the predicate holds for $1, \dots, a-1$. Integrating by parts with $u = z^{a-1}, dv = e^{-z}$, we get

$$\begin{aligned} \int_x^\infty e^{-z} z^{a-1} &= -e^{-z} z^{a-1} \Big|_x^\infty + \int_x^\infty (a-1) e^{-z} z^{a-2} dz. \\ &= e^{-x} x^{a-1} + (a-1) \int_x^\infty e^{-z} z^{a-2} dz \\ &\leq e^{-x} x^{a-1} + (a-1)(a-1) e^{-x} x^{a-2} \end{aligned}$$

by the inductive hypothesis.

For $x > a-1$, this obeys

$$\begin{aligned} &\leq e^{-x} x^{a-1} + (a-1) x e^{-x} x^{a-2} \\ &= e^{-x} x^{a-1} + (a-1) e^{-x} x^{a-1} \\ &= a e^{-x} x^{a-1}, \end{aligned}$$

confirming the inductive step.

Since $\int_{C_2}^\infty e^{-\rho t} t^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt = \frac{1}{\rho \theta \frac{1+\epsilon}{1-\epsilon\theta}} \int_{\rho C_2}^\infty e^{-t} t^{-1+\theta \frac{1+\epsilon}{1-\epsilon\theta}} dt \leq \frac{1}{\rho \theta \frac{1+\epsilon}{1-\epsilon\theta}} \int_{\rho C_2}^\infty e^{-t} t^{\lceil -1+\theta \frac{1+\epsilon}{1-\epsilon\theta} \rceil} dt$ for $t^* > e$, at which point we can apply the above lemma, this completes the proof.

Theorem 2 shows that eventually, the learning-by-doing externalities diminish over time as cumulative production increases. It also shows that higher demand elasticities and cost curve elasticities lead to greater dynamic price externalities. And, greater elasticities also lead to greater price externalities.

C Externalities

This appendix provides details on how we construct harmonized measures of externalities associated with electricity generation, natural gas production, and vehicles. We begin this appendix with a discussion of how we value global and local emissions.

C.1 Social Costs

C.1.1 Social Costs for Global Emissions

We value global damages from five pollutants: carbon dioxide (CO_2), methane (CH_4), nitrous oxide (N_2O), carbon monoxide (CO), and hydrocarbons (HC). Both CO and HC also impose local damages, and we defer the discussion of those damages to the next subsection.

As noted in Section 3.2, we construct MVPFs using three estimates of the social cost of carbon (SCC) in 2020 paired with the discount rate used to estimate the SCC: \$193 with a 2% discount rate (EPA 2023c), \$76 with a 2.5% discount rate (Interagency Working Group 2021), and \$337 with a 1.5% discount rate (EPA 2023c).¹¹⁴ All are expressed in 2020 dollars per metric ton of CO_2 . The \$193 and \$337 estimates come from the same report and differ only as a function of the discount rate. The \$76 SCC comes from an earlier report with different underlying inputs; inputting a 2.5% discount rate into the model that generates the \$193 and \$337 estimates would yield an SCC of \$117. While we refer to these estimates by the SCC in 2020, each reports a series of SCCs into the future. The \$193 and \$337 estimates contain annual SCC estimates until 2080, while the \$76 estimate contains annual estimates until 2050. We linearly extrapolate to obtain SCC estimates for years before 2020 and, for our \$76 SCC, for 2050 through 2080. We set the SCC to \$0 if the extrapolation yields a negative value.¹¹⁵

The reports from which we pull estimates of the SCC also contain annual estimates of the social costs of CH_4 and N_2O calculated using the same discount rate. For example, when using our \$193 SCC, we use a social cost of CH_4 of \$1,648 and a social cost of N_2O of \$54,139 in 2020, expressed in 2020 dollars per metric ton. We again linearly extrapolate to obtain estimates for years before 2020 and, when necessary, for 2050 through 2080 (setting estimates to zero in the rare cases where this extrapolation yields a negative value).

For global damages from CO and HC , we use global warming potential (GWP) factors to convert from tons of CO or HC to tons of CO_2 equivalent (CO_2e), to which we can then apply our preferred SCC. Both GWP factors come from Masnadi et al. (2018), who report GWP factors of 2.65 for CO and 4.5 for non-methane volatile organic compounds, which we apply to HC . With a SCC of \$193, these GWP factors imply a global social cost of \$511.45 for CO and a global social cost of \$868.5 for HC .

C.1.2 Social Costs for Local Pollutants

We value global damages from six pollutants: particulate matter 2.5 micrometers or less in diameter ($PM_{2.5}$), oxides of nitrogen (NO_x), sulfur dioxide (SO_2), hydrocarbons (HC), ammonia

¹¹⁴Recent work from Barrage & Nordhaus (2024) estimates an SCC in 2020 of around \$88.09 when using a 3% discount rate and \$32.40 with a 5% discount rate, adjusting to 2020 dollars.

¹¹⁵We note that few in-context MVPFs require social costs for years where extrapolations yield negative results.

(NH_3), and carbon monoxide (CO). Both CO and HC also impose local damages, which we describe above. We do not consider local damages from PM_{10} , toxic air pollutants (such as benzene), nor lead. We do not vary social costs for local damages over time. Importantly, our approach to measuring damages from local pollution accounts for the heterogeneity in these damages across areas of the US. Even a national policy has heterogeneous effects across the US because local pollution in more populated areas generates larger externalities. Hence, measuring the average damages requires accounting for where in the US the emissions are taking place. We account for the spatial distribution of power plants when considering electricity generation and spatial distribution of vehicle miles traveled when considering vehicle usage externalities.

We measure CO damages from Matthews & Lave (2000), who report an average local social cost per ton of CO of \$520 in 1992 dollars. Adjusting for inflation, this yields a local social cost of \$959.32 for CO . We do not construct different estimates of the local damages from CO depending on where the CO was released since we lack county-specific damage estimates.

The remaining social costs for local pollutants come from the AP3 model (Holland et al. 2016, Tschofen et al. 2019). For each pollutant released in a county, AP3 calculates social costs for emissions released from (in ascending height order) area sources, low stacks, medium stacks, and tall stacks, where each category is defined by the height at which emissions are released. To construct a national average social cost estimate for each pollutant, we construct two social costs per pollutant: one weighted by county-level electricity production data (our “baseline” local social cost estimate), and another weighted by county-level vehicle miles traveled (our “VMT-weighted” local social cost estimate). We subsequently use these when forming our estimates in cases where the damages stem from the electric grid and vehicles, respectively. We run AP3 using the EPA’s \$7.4 million VSL (in 2006 dollars), which corresponds to \$9.5 million in 2020 dollars (EPA 2010, 2023b). Finally, we use AP3’s estimated volatile organic compound (VOC) social cost to value local damages from HC .¹¹⁶

To calculate our baseline local social costs for electricity generation, we take AP3’s estimated county-level social costs for pollution released from low, medium, and tall stacks (e.g., all non area-source-level emissions) and weight by county-level electricity generation.¹¹⁷ We calculate average county-level social costs for each pollutant by weighting across height-dependent social costs using the quantity of emissions from that stack height. We identify power plant location using the EIA’s 2020 Form EIA-860 (EIA 2020a). We then match these location data to plant-level electricity generation data from the EIA’s Form EIA-923 using unique plant IDs (EIA 2020b).¹¹⁸ We focus on total fuel consumption (measured in MMBtu), assuming that emissions released are proportional to the quantity of polluting fuel consumed in the electricity generation process. We drop plants with zero fuel consumed as well as hydroelectric plants. We sum total fuel consumed in a county to calculate county-level weights. We then match county-level average social costs (calculated above) with the county-level electricity generation data. Weighting across counties, we obtain national average social cost estimates (in 2020 dollars) of \$64,190.32 for NH_3 , \$16,192.45 for NO_X , \$105,127.64 for $PM_{2.5}$, \$46,491.03 for SO_2 , and \$5,876.91 for HC . We note that this process only matches 952 counties; however, we obtain social cost estimates that are broadly similar to those reported by Tschofen et al. (2019).¹¹⁹

¹¹⁶When applying damages, we treat VOC and HC interchangeably. For clarity, we refer to whichever pollutant the author or data series reports.

¹¹⁷One motivation for weighting by county-level electricity production is that utilities may concentrate production in areas where few people live and where, in turn, social costs are low.

¹¹⁸Through this approach, we match 14,892 plants using their plant ID. We cannot match 2,264 plants.

¹¹⁹For reference, the emissions-weighted estimates reported by Tschofen et al. (2019) calculated using 2014

While we weight across counties using the share of fuel consumed by power plants, we also apply these baseline social costs to local pollution released during the upstream production of petroleum products (described below) and when valuing local pollutants abated by cap-and-trade schemes. We note that using social costs specific to refinery locations or social costs calculated only for counties covered by a given cap-and-trade scheme could change results if these locations had above average social costs in the AP3 model.

To calculate VMT-weighted local social costs that we apply to vehicle emissions, we focus on county-level social costs from AP3 calculated for area source emissions.¹²⁰ County-level VMT estimates come from EPA (2024b).¹²¹ We sum vehicle miles traveled (VMT) from passenger cars, passenger trucks, and light-duty trucks to calculate total VMT by county. We exclude VMT from buses. Excluding counties in Alaska and Hawaii (for which AP3 does not calculate social costs), only two counties cannot be matched. Weighting by total VMT, we obtain national average social cost estimates (in 2020 dollars) of \$186,992.41 for NH_3 , \$34,054.95 for NO_X , \$278,801.84 for $PM_{2.5}$, \$94,439.93 for SO_2 , and \$12,229.40 for HC .¹²² We note that we use an overall average for NH_3 since we lack on-road emission rates for NH_3 .

C.2 Electric Grid Externalities

A key input into our wind, solar, weatherization, rebates, EVs, and nudge policies is the electric grid. The externalities from the grid include global and local environmental externalities and externalities arising from imperfect competition of electricity providers such that the marginal change in electricity demand affects firm profits. In this section, we provide a detailed explanation of how we construct these externalities in each region of the country over time. We use our baseline local social costs to value local pollution from electricity production.

C.2.1 Marginal Emissions

To estimate the emissions from marginal changes in electricity demand and renewable energy supply, we use EPA’s Avoided Emissions and Generation Tool (AVERT). AVERT reports the marginal emissions factors for CO_2 , NO_X , $PM_{2.5}$, and SO_2 per KWh (EPA 2024b). This tool uses historical data on regional demand and generation to estimate the displaced emissions that would result from new energy programs including residential solar, wind, and energy efficiency programs. Since we know the ratio of CO_2 to other global pollutants (NH_4 and N_2O) from the EPA’s emissions factors, we also add these pollutants in the same proportion.

Since we are interested in the emissions from policy changes that affect electricity usage, we use AVERT’s estimates for the emissions associated with the additional, rather than the average, electricity usage. We note that a grid’s marginal emissions rate is often considerably higher than its average emissions rate (Holland et al. 2022). Using our estimates, the monetized

emissions quantities (the same data used in our AP3 run), converted from 2018 to 2020 dollars, are \$64,943.44 for NH_3 , \$21,647.81 for NO_X , \$134,010.27 for $PM_{2.5}$, \$43,295.63 for SO_2 , and \$6,906.68 for HC . With the exception of SO_2 , our baseline social costs are slightly smaller than the social costs weighted by attributed emissions reported in Tschofen et al. (2019).

¹²⁰One motivation for weighting by county-level VMT is that driving may be concentrated where many people live and where, in turn, social costs are high.

¹²¹Specifically, we pulled county-level VMT from AVERT v4.1, released April 2023.

¹²²Despite large VMT-weighted social costs, local pollution damages make up a small share of the total externality from gasoline in 2020.

externality from the marginal kWh is roughly two times higher than that for the average kWh in 2020. Regions of the grid that have a low average emissions rate due to renewable may still have a high marginal emissions rate if natural gas is the marginal generation source. Broadly, this distinction does not affect our results aside from the potential conclusions about nudge policies in the Northwest, which have average production from clean sources but marginal production from dirtier sources.

AVERT reports national and region-specific estimates. The heterogeneity in the monetized environmental externality per MWh across the US in 2020 is shown in Appendix Figure 3. AVERT splits the contiguous US into 14 electricity regions. Prior to 2019, AVERT used 10 regions. From 2007-2022, we construct state specific emissions factors by mapping each state-year pair with its corresponding AVERT region.¹²³ This is mostly trivial; a state is generally entirely contained within a region. For the instances in which a state is shared by multiple regions, there is still a single region that covers a significant majority of the state.

AVERT calculates emissions factors separately for programs that reduce energy consumption, increase solar installations, and increase wind adoption. For all policy categories besides wind and solar, we use the first set of estimates corresponding to reduced energy consumption. The monetized externalities using each of these three estimates are similar. Per kWh, the monetized environmental externality in 2020 that we apply to solar, wind, and energy efficiency programs is \$0.149, \$0.145, and \$0.159, respectively.¹²⁴

C.2.2 Forecasting the Grid

Many of the policies we study involve a change in electricity supply or demand that persists for multiple years. For example, we assume a wind turbine constructed as a result of the PTC will have a lifetime of 25 years (and 30 years in our robustness analyses). Therefore, to quantify the environmental impact of a wind turbine, we need to make assumptions about the time-path of the electric grid.

To forecast the grid after 2022, our baseline approach uses estimates from Princeton’s REPEAT Project (Jenkins & Mayfield 2023). We use their mid-range forecast that includes predicted changes to the electric grid from the Inflation Reduction Act. REPEAT forecasts the composition of the grid by generation source at various points in time until 2050. To obtain a complete time path we linearly interpolate between their estimates. REPEAT provides the electric grid mix, but does not report the mix of generation sources for the marginal unit of electricity.

To forecast the marginal emissions rate, we estimate the marginal emissions rate from a hypothetical 2020 grid that is entirely coal or natural gas, and we multiply by the estimated percent of the grid that is forecasted to be coal or natural gas using REPEAT estimates. The calculation is outlined below.

In the first step, we assume that the monetized 2020 environmental externality (r_{2020}) is entirely from coal and natural gas. We estimate the proportion of r_{2020} from coal versus natural gas by assuming that the proportion of these two generation sources in the average mix is equivalent to the proportion in the marginal mix. Using the emissions rates (e) and usage (u)

¹²³For in-context estimates that require earlier data, we apply the 2007 emissions rate to 2005 and 2006.

¹²⁴In 2007, the monetized externality for solar, wind, and energy efficiency programs was \$0.206, \$0.232, and \$0.241 in 2020 dollars.

of each generation source in 2020, we estimate that the ratio of natural gas to coal in r_{2020} is 1 to 1.157.

$$p_{coal} = \frac{e_{coal}}{e_{ng}} \cdot \frac{u_{coal}}{u_{ng}} = 2.429 \cdot 0.476 = 1.157$$

Using CO_2 output emissions rates from EPA’s eGRID, coal produces 2181 pounds of CO_2e per MWh and natural gas produces 898 lbs per MWh (EPA 2020). Therefore, one unit of coal produces 2.429 times the amount of emissions as one unit of natural gas. Natural gas makes up a larger share of the electricity mix in 2020 compared to coal. For every one unit of natural gas, there are 0.476 units of coal (EPA 2020). Since the ratio of natural gas to coal is 1:1.157, approximately 54% of the environmental externality in 2020, r_{2020} , is from coal and 46% is from natural gas.

Next, we calculate r_{2020} assuming that the entire electricity grid is made up of either coal or natural gas. For coal, this is given by:

$$c_{total} = \left(r_{2020} \cdot \frac{p_{coal}}{1 + p_{coal}} \right) \cdot \frac{1}{u_{coal}}$$

The first term gives the environmental externality from coal in 2020. The second term scales this to generate the environmental externality if the entire grid is made up of coal. An analogous calculation can be done for natural gas. The 2020 electricity mix is 19.28% coal and 40.47% natural gas. Using the 2020 externality from energy efficiency programs of \$0.16, c_{total} is \$0.44 and ng_{total} is \$0.18.

Finally, we arrive at the environmental externality per kWh by multiplying c_{total} and ng_{total} by the percent of the electricity mix made up of natural gas and coal in each year using the REPEAT forecasts. Appendix Figure 2 Panel B shows the evolution of the environmental externality over time.

We apply a similar process to construct the externality across time for individual states. REPEAT does not report the state or region-level grid mix over time. Instead, they report the combustion share of each state over time. Using linear interpolation, we construct a dataset of the combustion share for each state from 2022-2050. Instead of separately identifying natural gas and coal, we split generation into clean and dirty sources. We assume the entire 2020 environmental externality is from coal and natural gas sources. Analogous to the US-wide calculation, we construct the externality assuming the entire grid is dirty and multiply this by the forecasted dirty proportion using the combustion share estimates from Princeton.

For both in-context and US-wide estimates, we assume that the marginal emissions rate stays constant after 2050. Changes in the monetized environmental externality after 2050 are driven by changes in the social cost.

For robustness, we include MVPF estimates for a ‘dirty’ and ‘clean’ grid. Our dirty grid specification uses the state’s grid that has the highest monetized environmental externality, and the clean grid specification does the opposite. From 2005-2020, the cleanest state was California. The dirtiest state switches between the Mid-Atlantic (2005-2015) and the Midwest (2016-2020).

C.2.3 Measuring Electric Utility Profits

Electric utilities are a regulated industry with natural monopolies. To estimate the markup on electricity, we use the levelized cost of electricity (LCOE) and the retail price of electricity. We construct the total LCOE per MWh at the state and national level by taking an average of the LCOEs for each generation source weighted by the share of the grid each source represents. We use the total LCOE not including tax credits for new plants coming online in 2020. For wind and solar, we use the realized cost from projects installed in 2020 from the Department of Energy of \$32.99 and \$34.00, respectively (Wiser et al. 2023, Bolinger et al. 2021). The EIA’s 2018 Annual Energy Outlook provides the LCOE for natural gas plants coming online in 2020 of \$49.74 (EIA 2023a). For other sources, we use the EIA’s 2015 Annual Energy Outlook which provides the LCOEs for coal, nuclear, hydroelectric, biomass, and geothermal plants coming online in 2020.¹²⁵ For generation sources that do not have LCOE data, we exclude them and re-weight the included sources. We calculate the average LCOE per MWh for the US in 2020 of \$74.00 per MWh.

To account for the cost of delivering electricity from the source of generation to the point of use, we add distribution costs to the LCOE (EIA 2023a).¹²⁶ For the price of electricity, we use annual data on the retail price of electricity by state from the BLS (BLS 2024).¹²⁷

Markups generate externalities only when consumption is shifted from goods with low to those with high markups. As a result, the precise goal of our analysis is to measure the extent to which markups differ from the average economy-wide markup. De Loecker et al. (2020) find that the overall economy-wide markup is 8%. In our baseline specification, we also assume that 28% of utilities are publicly owned (EIA 2019) and that the effective corporate tax rate on private utilities is 10% (DOT 2016). Therefore, the producer WTP per additional kWh consumed is:

$$WTP_{prod} = (p - (\overline{LCOE} + c_{td}) \cdot (1 + m)) \cdot (1 - \tau) \cdot (1 - \alpha)$$

where p is the retail price, \overline{LCOE} is the generation-weighted average LCOE, c_{td} is the transmission and distribution cost, m is the economy-wide markup, τ is the tax rate, and α is the proportion of utilities that are publicly owned. For the US in 2020, the producer profit per MWh is \$11.03.

For the in-context version of these estimates, we use state-specific electricity prices and electricity generation mixes. m , c_{td} , τ , and α are constant across geography and time. The EIA does not report state-specific estimates of the LCOE, but they do report minimum and maximum values for the US for each generation source. To construct state-specific estimates of cost, we create 50 equally spaced bins from the minimum to the maximum LCOE for each generation source and assign states into each bin using their ranking in the BLS’ power generation industry wage index (BLS 2022).

A markup on utility profits affects government costs through profit tax revenue from utilities.

¹²⁵The LCOEs we use for coal, nuclear, hydroelectric, biomass, and geothermal are \$105.67, \$105.78, \$92.78, \$111.67, and \$53.11 (EIA 2023a).

¹²⁶The EIA reports distribution costs of \$32 per MWh in 2020, which are approximately 43% of the average 2020 LCOE.

¹²⁷The price of electricity per MWh in the US in 2020 is \$131.50. Among the 48 contiguous states, the most expensive state, Connecticut, and the least expensive state, Louisiana, had prices of \$227.10 and \$96.70, respectively.

Since we assume that 28% of utilities are publicly owned, effective corporate tax rates are 10%, and the effective tax rate on public utilities is 100%, the fiscal externality from utility profits is given by:

$$FE_{prod} = (P - (\overline{LCOE} + c_{td}) \cdot (1 + m)) \cdot (\alpha + (1 - \alpha) \cdot \tau)$$

In 2020, the fiscal externality from utility profits per MWh is \$5.99. The sum of the producer willingness to pay and government fiscal externality corresponds to a markup in excess of the economy-wide markup of 12.9%. For years prior to 2020, we assume the ratio of $(\overline{LCOE} + c_{td})/p$ is constant over time and use retail prices from the BLS.

C.3 Natural Gas Externalities

Some weatherization and appliance rebate policies induce changes in households' consumption of natural gas. These changes lead to environmental externalities as well as changes in producer profits arising from imperfect competition of natural gas distribution. In this section, we provide a detailed explanation of how we construct these externalities in each region of the country over time.

C.3.1 Environmental Externalities

We make the reasonable assumption that combustion emissions from one MMBtu of natural gas do not vary over place. We use emissions factors from the EPA's eGRID from 2011-2020 for CO_2 , CH_4 , and N_2O (EPA 2024c). eGRID does not report emissions factors for local pollutants associated with natural gas combustion. The emissions factors are constant over time for CH_4 and N_2O . For the CO_2 emissions factor, the pounds of CO_2 per MMBtu increased from 116.89 in 2011 to 116.98 in 2020. For years prior to 2011, we use the 2011 emissions factor. Applying our baseline social costs to these emissions factors result in a monetized 2020 environmental externality from natural gas of \$10.25.

C.3.2 Measuring Natural Gas Profits

Similar to electric utilities, we assume that natural gas utility companies experience profits arising from imperfect competition. To measure markups, we take the difference between the retail price of natural gas and the citygate price of natural gas. We take both of these prices from EIA (2023f) for each state from 2000-2022. In 2020, natural gas prices hit a record low partly as a result of the COVID-19 pandemic (EIA 2021a). Therefore, we use the markup in 2021 in our baseline 2020 MVPFs. Following our approach for electricity markups, we subtract the 8% economy-wide markup from our natural gas markup estimate. This results in a baseline markup for the US of 42.57%.

To construct the producer willingness to pay and fiscal externality in levels per MMBtu, we assume a 10% effective corporate tax rate on profits for private natural gas utilities and a 100% tax rate on public utilities (DOT 2016). Approximately 5% of natural gas utilities are publicly owned (EIA 2020c). Therefore, the producer profit and fiscal externality per MMBtu is \$4.40 and \$0.75, respectively.

C.4 Gasoline Externalities

A key input into our analysis of gasoline taxes, EV and HEV subsidies, and vehicle retirement programs is the dollar value of externalities generated by gasoline-powered, light-duty vehicles.¹²⁸ Appendix Table 12 contains the values of specific externalities in 2020, and Appendix Figure 2 shows how the values of vehicle externalities have varied over time (1990–2022). All externalities are reported in terms of dollars per gallon of gasoline, although we note below that some externalities arise per mile driven as opposed to per gallon of gasoline consumed. We factor this distinction into our externality measures for EVs, HEVs, and vehicle retirement programs. Unless otherwise noted, all dollar values are in 2020 dollars.

We consider two sources of pollution in developing monetized estimates of the externalities from light-duty vehicles. The first includes emissions released when vehicles use gasoline (“on-road emissions”). The second includes emissions that result from producing a gallon of gasoline (“upstream emissions”). When valuing on-road vehicle emissions, we use our VMT-weighted local social costs. When valuing upstream emissions, we use our baseline local social costs.

C.4.1 On-Road Pollution Externalities

Most emissions from gasoline are generated while vehicles are in operation. For each pollutant we consider (see Appendix Table 12), we proceed in three steps. First, we estimate the average emission rate (measured in grams per gallon) associated with a vehicle from a given model year. Second, we average emission rates and fuel economy across model years to measure the average per-gallon emission rate for the light-duty vehicle fleet in a given year. This fleet-wide emission rate reflects both the composition of the fleet in any given year as well as the driving behavior of cars of a particular age.¹²⁹ Finally, we translate annual emission rates for a particular pollutant into dollar terms using each pollutant’s corresponding social cost. This gives us the externality value for a given pollutant in a particular year in dollars per gallon.

The EPA emissions tests new vehicles to ensure compliance with regulatory standards at the time of production (EPA 2024a). For some pollutants, however, we must account for the fact that a vehicle’s emission control system may become less effective over time.¹³⁰ We begin with emissions that change as a vehicle ages, which consists of CO , HC , and NO_X . We follow Jacobsen et al. (2023), who pair comprehensive data on the initial emission rates of new light-duty vehicles from model years 1957 onward with smog check data from Colorado’s IM240 test to estimate how emissions increase with vehicle age. The authors calculate annual decay rates (e.g., the annual increase in emissions per mile) for CO , HC , and NO_X of 3.6%, 5.6%, and 4.0%, respectively. We follow the authors in assuming that vehicles do not decay after age 19. We also assume that vehicles from model years earlier than 1975 do not decay, as these vehicles predate contemporary emissions standards. For vehicles produced after 1975, $AgeFactor_p$ does not differ with model year (e.g., emissions control systems of newer vehicles do not decay at

¹²⁸The EPA’s definition of light-duty vehicles includes two regulatory classes: passenger cars and light trucks (EPA 2023d). Light trucks include minivans, pickups, and other vans, and passenger cars consist of coupes, sedans, and wagons. SUVs can be classified as either passenger cars or light trucks depending on the vehicle’s characteristics. Light-duty vehicles make up approximately 95% of vehicles on the road (DOE 2022). We consider differences between medium- and heavy-duty vehicles only when evaluating the externalities from diesel fuel.

¹²⁹For policies that displace a *new* vehicle, we omit this step but still consider changes in a vehicle’s emission rate over its lifetime. See Section C.4.4 below.

¹³⁰Catalytic converters, for example, deteriorate over a vehicle’s lifetime (Baronick et al. 2000).

different rates).

Combining $AgeFactor_p$ with data on initial emission rates and vehicle fuel economy, we approximate the emission rate of pollutant, p , measured in year y for a vehicle produced in model year m as

$$\underbrace{EmissionRate_{y,m,p}}_{\text{Grams per Gallon}} = \underbrace{EmissionRate_{m,p}}_{\text{Grams per Mile}} (1 + AgeFactor_p)^{y-m} \times \underbrace{FuelEconomy_m}_{\text{Miles per Gallon}} \quad (64)$$

where $EmissionRate_{m,p}$ is the initial emission rate of pollutant, p , for a vehicle from model year m ; $FuelEconomy_m$ is the average fuel economy of a vehicle from model year m ; and $AgeFactor_p$ is the annual rate of deterioration for pollutant, p , for a vehicle of age $y - m$. Initial per-mile emission rates for CO , NO_X , and HC for model years 1957 onward come from Jacobsen et al. (2023), who compile these data from a range of sources.¹³¹ We assume no vehicles from model years earlier than 1957 remain in use. Fuel economy data for model years 1957–1975 come from EPA (1973), and data for model years 1975 onward come from the EPA’s Automotive Trends Report (EPA 2023d).¹³² Both series are weighted by vehicle sales. We assume a vehicle’s fuel economy does not change with the vehicle’s age.

For NO_X , CO , and HC , we account for the fact that fuel containing ethanol burns differently than pure gasoline. To do so, we use emissions adjustment factors from Hubbard et al. (2014), who report emissions rates by ethanol content.¹³³ The authors find that vehicles running on fuel with 9.8% ethanol emit 13.2% less NO_X (authors’ Table S3), 6.8% more CO (authors’ Table S2), and 13.0% less HC (authors’ Table S1, referred to as “non-methane hydrocarbons (corr.)” by the authors) relative to a vehicle running on fuel without ethanol. Multiplying these percent differences by the ratio of the observed share of ethanol in gasoline in a given year to the share of ethanol used in these emissions tests (9.8%) allows us to account for differences in the ethanol content of the fuel used in the authors’ tests and the average gallon of gasoline, assuming a linear relationship between ethanol content and emission rates. These adjustments do not noticeably affect our externalities in 2020 given these pollutants’ low initial emission

¹³¹The authors calculate unweighted emission rates for model years 1957 through 2020 and sales-weighted emission rates for model years 1981 through 2015. The authors note that both series have similar levels and trends. We use unweighted emission rates to capture more model years. We apply a linear interpolation to account for model years with missing emission rates. Only 1994 and 1995 lack emission rates for all three measured pollutants, and 1973 is missing an emission rate for NO_X . We assume no further improvements to vehicle emissions have been made for model years later than 2020.

¹³²The earlier fuel economy series reports a national average fuel economy of 15.6 miles per gallon in 1975. The Automotive Trends Report, however, reports a national average fuel economy of 13.1 miles per gallon for 1975. So that each series has the same average fuel economy in 1975, we calculate the difference between each series’ estimate of the 1975 average fuel economy and add this difference to each estimate in the earlier series. After this transformation, each series has the same average fuel economy for 1975. When using data from the Automotive Trends Report, fuel economy data for 2022 was preliminary when reported.

¹³³We do not adjust the emission rates for CH_4 or N_2O because, as described below, estimates from Lee et al. (2021) include CH_4 and N_2O emissions from ethanol combustion. While we assume CO_2 from ethanol combustion is entirely offset, we cannot assume the same for CH_4 and N_2O . To avoid double counting damages from these two greenhouse gases, we do not adjust our emission rates for CH_4 and N_2O using adjustment coefficients from Hubbard et al. (2014). We scale down on-road CH_4 and N_2O emissions by the share of gasoline. We cannot isolate CH_4 and N_2O emissions from Lee et al. (2021) and therefore leave these damages as part of our reported upstream CO_2 damages even though these emissions are released during on-road operation. We note that CH_4 and N_2O emissions from ethanol combustion are the smallest contributors to ethanol’s life cycle carbon intensity estimated by Lee et al. (2021).

rates; in earlier years (when emission rates were larger), ethanol made up too small a share of gasoline for these adjustments to affect our conclusions.¹³⁴

Next, we consider pollutants for which it is reasonable to assume that the impact of vehicle age on emissions is negligible ($AgeFactor_p \approx 0$). This includes CO_2 , SO_2 , $PM_{2.5}$, CH_4 , N_2O . We do not consider differences in damages from SO_2 and $PM_{2.5}$ between gasoline and ethanol. Both CO_2 and SO_2 emissions proceed from the carbon and sulfur content of a gallon of gasoline, meaning per-gallon emission rates will not vary with model year. We calculate on-road CO_2 emissions using the emissions coefficient for motor gasoline (8,786 grams per gallon) from EIA (2023b). We adjust our externalities to account for the share of ethanol in finished motor gasoline. We assume the ethanol component of gasoline is non-emissive, as the carbon dioxide taken out of the atmosphere while growing the organic material needed for ethanol is offset by the carbon dioxide emitted when ethanol is burned (EIA 2023b, AFDC 2024c). We allow the share of ethanol in gasoline to vary over time. We calculate SO_2 emissions using the average sulfur content of a gallon of gasoline (EPA 2017).¹³⁵

Because catalytic converters do not affect $PM_{2.5}$, CH_4 , and N_2O emissions, we assume these pollutants are also unaffected by the deterioration of emission control systems (IPA 2024). Emission rates for these three pollutants for model years 1990 onward come from MOVES, a tool designed by the EPA to quantify pollution from mobile sources (EPA 2024d).¹³⁶ We use sales weights from the EPA (2023d) to average across vehicle classes included in MOVES’ definition of light-duty vehicles.¹³⁷ We assume emissions for vehicles released before 1990 emit at the same rate as the average new vehicle from 1990, and that vehicles produced after 2020 emit at the same rate as the average new vehicle from 2020.¹³⁸

Once we have emission rates for each pollutant by model year, we compute the average emission rate for the entire fleet in a given year using the distribution of model years on the road in a given year and data on vehicle usage by vehicle age. We use data on the age distribution of and miles traveled by light-duty vehicles from the 2017 National Household Transportation Survey (FHWA 2017).¹³⁹ This survey provides a snapshot of the vehicles on the road in 2017, which enables us to measure both the fraction of cars of a given age and the average annual vehicle miles traveled by vehicles of a given age from the sample of respondents who indicated their vehicle’s age and average annual VMT. We assume model years are distributed evenly within bins when reported as ranges. We assume VMT for vehicles older than 33 years equals the average VMT at age 33. We hold the age distribution of the fleet and the distribution of

¹³⁴For reference, in 2020, this adjustment moves per-gallon damages from NO_X from \$0.076 to \$0.071, from \$0.050 to \$0.052 for CO , and from \$0.039 to \$0.036 for HC .

¹³⁵EPA (2017) reports average annual sulfur contents for 1997–2016. For years before 1997, we assume the sulfur content equals the value observed in 1997. For 2017 onward, we set sulfur content equal to Tier 3 Motor Vehicle Emission and Fuel Standards (10 ppm) (EPA 2017). To convert from ppm to grams per gallon, we assume a density of 6.1 pounds per gallon (Hawley 2022). This results in a conversion rate from ppm to grams per gallon of 0.0028 (e.g., 30 ppm is equivalent to approximately 0.08 grams per gallon of gasoline).

¹³⁶We use emission rates derived from MOVES but reported by Cai et al. (2013).

¹³⁷MOVES includes three vehicle classes in its definition of light-duty vehicles (EPA 2016). These categories do not align with the vehicle classes used in the U.S. EPA’s fuel economy data set. To link these data sets, we assume “Passenger Cars” corresponds with the “All Cars” classification used in the Automotive Trends Report, “Passenger Truck” with the “Truck SUV” classification, and “Light-Duty Commercial Truck” with the “Minivan/Van” and “Pickup” classifications.

¹³⁸As described below, emission rates for $PM_{2.5}$ from tires and brakes also come from MOVES. All details described in this paragraph apply to our treatment of $PM_{2.5}$ from tires and brakes.

¹³⁹When using data from the NHTS, we exclude recreational vehicles and motorcycles, as these are not included in the Automotive Trends Report.

VMT by vehicle age constant over time.

We construct weights for each model year by multiplying the annual gallons of gasoline consumed by a vehicle that age and the share of vehicles on the road of that age. We calculate annual gallons of gasoline consumed by dividing the VMT by a vehicle of a given age by the vehicle’s fuel economy. We use this weight to calculate fleet-wide average emission rates for externalities that arise per-gallon of gasoline used. All fleet-wide emission rates considered thus far have been expressed in grams of pollution per gallon. We convert emission rates to metric tons per gallon and then multiply each emission rate by the corresponding social cost to monetize damages. Appendix Table 12 summarizes each pollutant’s contribution to the total per-gallon externality.

C.4.2 Per-Mile Driving Externalities

Many vehicle externalities are closely linked to gasoline consumption, and the value of these externalities is often estimated on a per-gallon basis. We assume all exhaust pollution arises per-gallon of gasoline burned.¹⁴⁰ However, some vehicle externalities arise on a per-mile basis and are most naturally measured per mile of driving. We consider three externalities that arise per-mile-traveled: $PM_{2.5}$ from tire and brake wear, accidents, and congestion.

The per-mile emission rate for $PM_{2.5}$ from tire and brake wear comes from MOVES (Cai et al. 2013), the same source from where we obtain per-mile emission rates for exhaust $PM_{2.5}$.¹⁴¹ To value accidents, we use the annual fatalities avoided from a 1% reduction in VMT (263 fatalities avoided) estimated in Jacobsen (2013b), apply the EPA’s VSL of \$9.5 million (EPA 2010), and divide the product of these terms by the number of miles reduced from a 1% reduction in total VMT in 2008 (30 billion miles), the year from which most of Jacobsen’s data come (AFDC 2024b).¹⁴² This calculation yields an average accident externality of \$0.08 per mile. To value congestion, we average per-mile externalities from three papers—Couture et al. (2018) (\$0.02), Parry & Small (2005) (\$0.05), and Parry et al. (2014) (\$0.03)—for an average congestion externality of \$0.03 per mile.¹⁴³ We assume vehicles of different model years and vehicle types impose the same per-mile accident and congestion externality. Accidents and congestion are local externalities, and we do not vary these values over time.

For externalities that arise per-mile traveled, we augment the weighting approach described above to assign greater weight to vehicles of a given age that travel more miles (rather than those that consume more gasoline). This approach does not affect the per-mile accidents and congestion rate, as these do not vary with model year, although $PM_{2.5}$ from tires and brakes does. We also use a weighting approach to calculate a fleet-wide average fuel economy that lets us express per-mile externalities in per-gallon terms. Multiplying per-mile emission rates by this VMT- and age-weighted fuel economy (23.1 MPG in 2020) yields per-gallon estimates for our three per-mile externalities.

When evaluating gasoline taxes, changes in gasoline consumption do not arise entirely from

¹⁴⁰Since local on-road pollution is a relatively small share of the total externality from a gallon of gasoline in 2020, our results are not sensitive to this assumption.

¹⁴¹As noted above, we handle emission rates for $PM_{2.5}$ from tires and brakes the same way we handle other emission rates from MOVES.

¹⁴²In particular, we use the fatalities avoided calculated in the author’s Appendix G of Jacobsen (2013b), where the author applies his main text findings to a gasoline tax.

¹⁴³For Parry et al. (2014), we use the author’s estimate constructed using more granular traffic delay data. This estimate is 41 percent smaller than their initial estimate but is more in line with previous findings.

changes in VMT. As a result, we must know how much of a change in gasoline consumption is due to changes in VMT. We follow Small & Van Dender (2007) in assuming that changes in VMT account for 52% of the price elasticity of gasoline. We refer to this share of the own price elasticity that arises from VMT changes as β . One could, in practice, multiply the price elasticity of gasoline or the per-gallon externality by β to account for the fact that changes in gasoline usage do not stem entirely from changes in VMT. In Appendix Table 12, we multiply accidents, congestion, and $PM_{2.5}$ from tires and brakes by our preferred value of β (0.52). This approach allows us to compare across externalities before applying an elasticity. We describe in Appendix E.10 an alternative approach where we apply β to the price elasticity of gasoline; each approach yields identical conclusions.

C.4.3 Upstream Pollution Externalities

Upstream emissions include the pollution released while extracting and refining crude oil. We decompose upstream emissions into well-to-refinery emissions and refinery emissions. Well-to-refinery emissions include emissions released while exploring for, extracting, processing, and transporting crude from the well to the refinery. We only consider greenhouse gases released during this process.¹⁴⁴ Refinery emissions include both the local air pollutants and greenhouse gases released by petroleum refineries. We ignore emissions generated while transporting gasoline from the refinery to the pump. We assume gas taxes do not affect vehicle production decisions and therefore exclude vehicle manufacturing emissions from these MVPFs. We also ignore the effects on vehicle scrappage and downstream effects on the used-vehicle market, as we assume the price elasticity of gasoline captures the total effect of the gasoline price on gas consumption.

For both processes, we estimate upstream emissions by dividing the pollution released per gallon of crude input by the gallons of petroleum product produced from one barrel of crude oil. Formally, for each source of pollution, s , equal to the sourcing of crude oil or the refining process, we write $Upstream_{y,p,s}$ of pollutant p as

$$Upstream_{y,p,s} = \frac{Pollution_{y,p,s}}{RefineryYield_y} \quad (65)$$

where $Pollution_{y,p,s}$ represents the metric tons of pollutant, p , released per barrel of crude oil from source s in year y , and $RefineryYield_y$ refers to the gallons of petroleum product generated from one barrel of crude. We calculate refinery yield for a given year by dividing the total gallons of output from refiners and blenders in that year by the total barrels of crude that entered refiners and blenders that year.¹⁴⁵ In 2020, one barrel of crude oil produced on average 44.3 gallons of petroleum product.¹⁴⁶ The following paragraphs explain how we obtain values for the pollution emitted from a barrel of crude, $Pollution_{y,p,s}$.

¹⁴⁴Since we consider petroleum extracted globally, valuing local air pollution from this process would require both information on where emissions are released and how to value local damages outside of the United States.

¹⁴⁵The EIA tracks inputs for three types of facilities (“refiners,” “blenders,” and “refiners and blenders”) in its “U.S. Refinery and Blender Net Input of Crude Oil and Petroleum Products (Thousand Barrels)” series (EIA 2023*i*). We look at refiners and blenders because data for these facilities are available for more years, and because these facilities tend to have greater outputs than the others.

¹⁴⁶Output data come from the EIA’s “U.S. Refinery and Blender Net Production of Crude Oil and Petroleum Products (Thousand Barrels)” series (EIA 2023*j*). One barrel of crude contains 42 U.S. gallons. Refiners and blenders have a “processing gain” (output outweighs input in a given period) due to the specific gravity of the petroleum products refined. If the products refined have a lower specific gravity than crude oil, refiners will

We begin with pollution generated during the production and transportation of crude oil from the well to the refinery. We use estimates from Masnadi et al. (2018) (authors’ Figure 1), who estimate well-to-refinery emissions to be 10.3 grams of CO_2 equivalent (CO_2e) per megajoule of crude produced.¹⁴⁷ One barrel of crude oil contains 6,119 megajoules (DOE 2020). Producing one barrel of crude thus yields 63,014 grams of CO_2e . We assume well-to-refinery emissions have remained constant over time. Using the 2020 refinery yield of 44.3 gallons of petroleum product per barrel of crude, sourcing the crude needed to produce one gallon of petroleum product releases 1,421.5 grams of CO_2e . This allocation method assigns pollution from crude to its downstream products (i.e., motor fuel and diesel fuel, among others) in proportion to the quantity produced. CO_2 and CH_4 make up 65% and 34% of total emissions, respectively, with VOC and N_2O making up the remaining one percent.¹⁴⁸ We then divide the share of total CO_2e attributable to CH_4 and N_2O by the GWP factors used by the authors to convert grams of non- CO_2 pollutant to grams of CO_2e . This gives us grams of CO_2 , CH_4 , and N_2O released during the well-to-refinery process. We apply each pollutant’s respective social cost to value well-to-refinery emissions in dollars per gallon of petroleum product.¹⁴⁹

We then consider pollution released by US petroleum refineries. From 1990 onward, the Inventory of U.S. Greenhouse Gas Emissions and Sinks (“the Inventory”) collects annual, facility-level emissions data from domestic refineries for three greenhouse gases (CO_2 , CH_4 , and N_2O) released during the “Crude Refining” activity (Inventory Tables 3-45, 3-47, and 3-49) (EPA 2021). The National Emissions Inventory (NEI) reports emissions of local air pollutants by source every three years (2008–2020), which we use to calculate emissions from refineries for six local pollutants (NH_3 , CO , HC , NO_X , $PM_{2.5}$, and SO_2) (EPA 2023a). We interpolate to estimate emissions for unobserved years between 2008 and 2020. For years before 2008, we assume petroleum refineries emitted the same amount of pollutant, p , that refineries emitted in 2008. We do the same for years after 2020. For all pollutants, we calculate emissions per gallon of petroleum product by dividing total emissions by the total barrels of crude oil that entered refiners and blenders that year (EIA 2023i). We then divide pollution released per barrel of crude oil by the refinery yield.¹⁵⁰ We again apply each pollutant’s corresponding social cost to value emissions in dollars, using our baseline social costs for local pollutants. We aggregate emissions for a pollutant, p , in year y from both upstream sources to construct an annual upstream emission rate for each pollutant, in dollars per gallon.

All upstream emission rates are calculated per gallon of petroleum product. However, gasoline purchased in the US contains ethanol. To account for the share of ethanol in gasoline, we scale down each upstream emission rate in year y by one minus the share of fuel ethanol in finished motor gasoline. We calculate this share using the approach outlined by the EIA

experience a processing gain and produce more than 42 gallons of product from one barrel of crude (EIA 2024k). The national refinery yield has remained roughly constant over time.

¹⁴⁷This reflects the authors’ global volume-weighted-average. We use this global value because the US continues to import a large volume of crude oil—8.33 million barrels per day in 2022 from 80 different countries (EIA 2023e). For policies that target crude oil production in specific countries, we rely on the authors’ country-specific carbon intensity measurements.

¹⁴⁸We assume N_2O and VOC each make up half of the remaining percent of pollution. Since we calculate global damages from VOC using the same GWP factors as the authors, we leave this pollutant in terms of CO_2e .

¹⁴⁹Since the social cost of a non- CO_2 GHG is roughly equal to the social cost of carbon scaled by the GWP factor of the pollutant, this approach generates approximately the same results if we were to directly apply our SCC to the gCO_2e estimate.

¹⁵⁰This is equivalent to dividing total emissions in a given year by total gallons of output from refiners and blenders that year.

(EIA 2023*b*). This approach assumes the ratio of the quantity of motor gasoline supplied to the quantity of fuel ethanol supplied (excluding denaturants, losses, and co-products) equals the percentage share of ethanol in finished motor gasoline. The quantity of motor gasoline supplied comes from the EIA’s “U.S. Product Supplied of Finished Motor Gasoline” series (EIA 2023*h*). The quantity of fuel ethanol supplied comes from the EIA’s Monthly Energy Review (Table 10.3) (EIA 2024*d*) (“Fuel Ethanol, Excluding Denaturant, Losses and Co-products”). For example, we multiply all upstream emissions for 2020 by 0.95 to account for the 4.9% of ethanol in gasoline. We describe above how we adjust on-road emissions for the share of ethanol in gasoline.

We also account for upstream emissions from ethanol production. We only consider greenhouse gas emissions from this process. We use estimates of the carbon intensity of ethanol production from Lee et al. (2021), who find a carbon intensity of 45 grams of CO_2e released upstream per MJ of ethanol produced in 2019. We allow the carbon intensity of ethanol production to vary over time (authors’ Figure 4).¹⁵¹ We add to this value an estimate of the carbon intensity of land-use change associated with ethanol production (7.4 grams of CO_2e per MJ) also from Lee et al. (2021). We hold this value constant overtime. We multiply the combined carbon intensity of ethanol production by the share of ethanol in gasoline, and then by the social cost of carbon in a given year to monetize these damages. Increased emissions from ethanol production are added to the upstream CO_2 estimate we present in Appendix Table 12. After adjusting for the ethanol content of gasoline, upstream CO_2 damages increase from \$0.18 to \$0.22 per gallon.

C.4.4 Lifetime Vehicle Externalities

We also estimate the total damages a given vehicle generates over its lifetime. Policy-specific appendices describe which values enter into our calculations. Here, we give a broad overview of how we move between average gasoline externalities to damages measured over a vehicle’s lifetime.

For policies that affect new vehicles, we perform the calculations described above but focus on emission rates specific to the model year of the affected vehicle. For example, if a subsidy induces the purchase of a vehicle in 2020, we consider the emission rates of a new vehicle purchased that year, rather than a fleet-wide average emission rate. We account for changes in the emission rate over the vehicle’s lifetime due to the decay of emissions abatement technologies and continue to assume vehicles do not decay after age nineteen. We also use the fuel economy associated with the vehicle’s model year rather than a fleet-average fuel economy. We hold upstream damages constant across new and fleet-average vehicles, as we assume these arise per gallon of petroleum product produced and should therefore not vary with either model year or fuel economy (although they do vary with the year being evaluated). We do not isolate car-specific emission rates when evaluating policies that specifically affect cars rather than an average light-duty vehicles, although we do use car-specific fuel economies.

¹⁵¹This estimate of the carbon intensity of ethanol includes emissions from activities such as increased farming, ethanol processing, and increased fertilizer and chemical usage. Lee et al. (2021) estimate carbon intensities (in grams of CO_2e per MJ) for 2005 through 2019. We assume ethanol production for years before 2005 had the same carbon intensity as estimated in 2005, and that years after 2019 had the same carbon intensity as estimated in 2019. We assume one gallon of pure ethanol contains approximately 89.2 MJ of energy (AFDC 2024*d*) when using the reported “higher heating value” and assuming there are 0.001055 MJ in a Btu (ENERGY STAR 2015).

We assume cars have a lifetime of 17 years and the average light-duty vehicle (which includes both cars and light-trucks) has an average lifetime of 19 years, both of which come from Greene & Leard (2023).¹⁵² For lifetime VMT, we again draw from the FHWA (2017). For cars, we use the annual VMT reported for automobiles, cars, and station wagons. For the average light-duty vehicle, we use the same annual VMT described above, which averages across all vehicle types (excluding RVs, motorcycles, and unspecified vehicle types) and weights by the samples of respondents who indicated that vehicle type.

Over the vehicle’s lifetime, we account for rising social costs for global emissions. In our baseline setting, our social costs rise more slowly than our discount rate. Social costs for local pollutants do not change over time, although damages from these pollutants rise as vehicles’ emissions abatement systems decay. For policies that specifically affect vehicle fuel economy, we assume that improvements in fuel economy do not also generate improvements in other vehicle emission rates.¹⁵³ We account for differences in emission rates between vehicles of different model years if the policy targets vehicles of different ages.

For most policies, we assume drivers maintain a given level of VMT regardless of what vehicle they select. In these instances, we ignore externalities that arise per mile traveled. This again assumes that per-mile externalities do not vary with vehicle type. When we incorporate the rebound in VMT due to improved vehicle fuel economy (which we include in our hybrid and vehicle retirement MVPFs), we account for accidents, congestion, and $PM_{2.5}$ from tires and brakes, as vehicles generate these externalities when they travel more miles (see Appendix D below). In these instances, per-mile accident and congestion externalities do not differ between vehicle types even though per-gallon pollution externalities vary as a function of fuel economy.¹⁵⁴ Accounting for increases in VMT can therefore more than offset initial benefits from improved fuel economy, since both vehicles—regardless of fuel economy—generate the same per-mile externalities.¹⁵⁵ For policies where we assume vehicles do not travel the average VMT reported by the FHWA (2017), we scale lifetime damages by the fraction of the annual average VMT we think the vehicle travels because VMT enters linearly into our calculations, assuming this fraction holds uniformly over the vehicle’s lifetime.

For EV and hybrid vehicle MVPFs, we forecast lifetime vehicle externalities until 2050. We hold upstream and on-road emission rates fixed but adjust for rising social costs as usual. To forecast vehicle fuel economies past 2021, we use information from the EIA’s 2023 Annual Energy Outlook (EIA 2023a). In the Annual Energy Outlook’s Table 40, the EIA projects the miles per gallon of light-duty cars and conventional light trucks from 2022 to 2050. In order to evolve smoothly from our historical estimates to our future projections, we start with our observed 2021 fuel economy estimate and apply the year-over-year percent change in the fuel

¹⁵²To calculate the lifetime of an average new light-duty vehicle, we take the authors’ calculated lifetimes of 17 for cars, 20 for SUVs, and 25 for pickup trucks and calculate a weighted average using the 2020 production shares of 0.44 (all cars), 0.42 (truck SUVs and minivans/vans), and 0.14 (pickups) from the EPA (2023d). This yields an average lifetime that rounds to 19 years.

¹⁵³In other words, while a policy might cause drivers to use fewer gallons of gas, we do not assume that the increase in vehicle fuel economy also comes with lower emission rates.

¹⁵⁴Per-mile $PM_{2.5}$ emissions from tires and brakes can differ between vehicles if policies target vehicles of different ages.

¹⁵⁵Phrased differently, more fuel-efficient vehicles do not impose smaller per-mile externalities than vehicles with lower fuel economies, meaning an increase in driving will always generate damages from driving externalities. In our MVPFs, since local pollution damages from gasoline consumption are a small component of the local externality (especially for new vehicles), we see that increases in accidents, congestion, and $PM_{2.5}$ from increased VMT more than offset the initial benefits from decreased local air pollution that arise from improved fuel economy.

economy of cars and light-duty trucks implied by the EIA’s forecast. We perform this exercise separately for all light-duty vehicles and for light-duty cars alone.¹⁵⁶ When forecasting the average light-duty vehicle fuel economy, we hold fixed the relative weighting of cars and trucks from 2021.

C.4.5 Measuring Gasoline Producer Profits

Imperfect competition among suppliers in three markets results in a markup on gasoline that is above the economy-wide average markup. We account for producers’ WTP for lost profits resulting from reduced gasoline consumption.

First, crude suppliers sell oil to refiners at a price (refiner acquisition cost) above the landed cost of producing a barrel of crude, both reported by the EIA (EIA 2024*g,e*). In 2020, moving one barrel of crude oil from well to refinery cost \$37.27 on average, while refiners purchased this barrel for, on average, \$40. We use the refinery yield (1 barrel of crude produces how many gallons of refined product) to convert barrels of crude to gallons of consumable petroleum product. This conversion allocates profits (as well as upstream emissions) to downstream products in proportion to the quantity produced. We set the per-gallon markup to \$0 if the difference between the landed cost and selling price of crude is negative.¹⁵⁷ In 2020, the average markup imposed by crude producers equaled \$0.06 per gallon or 2.6% of the price of gasoline.

Second, the EIA reports that 17.7% of the price of a gallon of gasoline arises from refining costs and profits, not including costs from crude production passed onto refiners (EIA 2024*b*). Favennec (2022) estimates that new refineries face a variable cost of \$10 per barrel of crude processed but notes that this cost could fall between \$3 and \$5 per barrel once capital investments fully depreciate. Combining the EIA’s estimate of the share of the price of gas owing to refining costs and profits with a \$4 (\$10) refining cost, we calculate a per-gallon markup of \$0.32 (\$0.20) in 2020, or 14% (9%) of the price of gas. We use a \$4 cost of refining as our baseline specification.¹⁵⁸

Third, we consider markups imposed by distributors, who purchase gasoline from refiners at the dealer tank wagon price and sell to consumers at the retail price of gasoline, both measured on a per-gallon basis. These data are reported in the EIA’s “U.S. Total Gasoline DTW Sales Price by Refiners” series (EIA 2022*c*). The markup from distributors is the difference between these prices. In 2020, distributors purchased gasoline from refineries at \$1.86 per gallon and sold the same gallon to consumers for \$2.27, implying a per-gallon markup of \$0.41 per gallon, or 18% of the per-gallon price of gasoline. We assume distributors face no variable costs other than the cost of purchasing refined gasoline.¹⁵⁹

Summing each producer’s markup yields a total per-gallon markup equal to 35% of the price of gasoline. We subtract from this gasoline markup the average, economy-wide markup (8%) estimated by De Loecker et al. (2020), resulting in a 27% average markup on a gallon of gas.

¹⁵⁶When evaluating hybrid and EV policies, we typically use externalities based on cars with higher-than-average fuel economies. In these instances, we simply adjust lifetime externalities by the ratio of the average light-duty car MPG to the higher-than-average counterfactual MPG, since fuel economy enters linearly into our calculations.

¹⁵⁷No monthly data reported a negative markup in 2020, and negative markups appear intermittently after January 1983.

¹⁵⁸Neither approach results in a negative markup in any period.

¹⁵⁹This approach generates a negative markup for one month in our data (October 2019). We set markups to \$0 if this approach yields a negative markup.

In 2020, the total markup on gasoline was \$0.61 per gallon, which we adjust by the corporate average tax rate (21%) to account for the share of profits producers keep (Watson 2022). We do not vary across time the effective corporate tax rate gasoline producers face.

D Rebound

When a policy causes people to consume more or less of a good such as energy, this can affect its price, leading to a “rebound” effect. This effect means that the standard “treatment vs control” comparison does not identify the ultimate causal effect of the policy, as the treatment and control group are generally both experiencing the price changes. In this Appendix, we discuss how we adjust estimates of the causal effect of policy changes estimated in reduced-form settings to account for rebound effects by using external estimates of the supply and demand curves of the market.

Let the total demand for energy be $Q(p)$ and supply be given by $S(p)$. Suppose we have a policy (e.g., EV subsidy) that increases the demand for energy by dE . In equilibrium, we require markets to clear so that

$$dE + Q'(p)dp = S'(p)dp$$

or

$$dp = \frac{-dE}{S' - Q'}$$

which means that the total change in energy consumption is given by

$$S'(p)dp = dE \frac{S'(p)}{S'(p) - Q'(p)} \tag{66}$$

$$= dE \frac{1}{1 - Q'(p)/S'(p)} \tag{67}$$

$$= dE \left(\frac{1}{1 - \epsilon^D/\epsilon^S} \right) \tag{68}$$

where $\epsilon^D = (Q'(p)/Q(p))p$ and $\epsilon^S = (S'(p)/S(p))p$. The last line follows from the fact that $S(p) = Q(p)$ in equilibrium. The causal effect estimated from reduced form approaches, dE , is offset by a ‘rebound’ effect given by $\frac{-\epsilon^D/\epsilon^S}{1 - \epsilon^D/\epsilon^S}$. Intuitively, if supply is perfectly elastic so that $\epsilon^S = \infty$, then there is no rebound effect; conversely, if supply is perfectly inelastic, then any policy that attempts to change energy consumption does not succeed in doing so: prices are lowered so that energy consumption remains constant.

We incorporate rebound effects into both the electricity generation markets and the market for natural gas. For the gasoline market, we assume that there is a flat global supply curve for gasoline so that there is no rebound effect on prices, however we do incorporate a rebound effect of changes in the price of driving on vehicle miles traveled.¹⁶⁰

¹⁶⁰We note that for gasoline taxes, any estimate of the causal effect of the tax would incorporate both the channel from changes in vehicle miles traveled and from changes in the cars people drive to have higher miles per gallon. In these cases, it is natural to assume that estimates of the causal effect of the gas tax on gasoline consumption already incorporate this rebound effect and hence we do not add an additional rebound effect.

Electricity Markets We account for supply curves that are locally upward sloping across markets in the US. In our baseline specification, we construct a demand elasticity for electricity using a weighted average of demand elasticities from residential, commercial, and industrial electricity demand. We use a commercial and residential demand elasticity from Serletis et al. (2010) of -0.134 and -0.287, respectively. We use an industrial demand elasticity of -0.125 from Jones (2014). These elasticities are weighted by their respective share of total electricity demand resulting in a demand elasticity of -0.19 (EIA 2023c). This estimate is similar to other estimates in the literature (EIA 2021b, Deryugina et al. 2020). For supply, we similarly construct a weighted average of the elasticities of each generation source. We follow the approach of the Department of Interior’s MarketSim model and use the supply elasticities by source which are derived from the EIA’s 2015 and 2020 Annual Energy Outlook (DOI 2021, EIA 2023a). The resulting supply elasticity is 0.78.

Using the demand elasticity of -0.19 and the supply elasticity of 0.78, we get our baseline estimate of the rebound effect of 20%.

Appendix Figure 8 explores the robustness of our rebound estimate to a range of supply and demand elasticities. The short to medium range electricity demand elasticity estimates generally hover around the 0 to 0.4 range. Deryugina et al. (2020) exploit exogenous shocks in retail electricity prices in Illinois to estimate a residential price elasticity of -0.27. The EIA’s 2020 Annual Energy Outlook (EIA 2023a) reports price elasticities for residential and commercial electricity demand. Weighting by each sector’s market share, this corresponds to a demand elasticity of -0.16 (EIA 2021b). These estimates are in the range of our baseline value of -0.19. The availability of electricity supply elasticities in the literature is limited. However, for a fixed -0.19 demand elasticity, the rebound effect is robust to a range of supply elasticities. Assuming the electricity supply elasticity is greater than 0.4, which is consistent with all of the solar and wind elasticities in our sample, the rebound has an upper bound of 33%.

Natural Gas Markets Following the approach in Appendix D, we apply a rebound effect to policy-induced changes in natural gas consumption. We use a natural gas supply elasticity of 1.50 from DOI (2021), which is the same natural gas elasticity as the one feeding into the average electricity supply elasticity. We use a natural gas demand elasticity of 0.20, which is the middle of the range of estimates from Rubin & Auffhammer (2024). These elasticities lead to a natural gas rebound effect of 11.76%.

Gasoline and Driving We also consider the potential rebound effects in the vehicle markets. For gasoline, we assume a flat global supply curve and hence no rebound effects. But, for policies that cause individuals to purchase more fuel efficient vehicles, we do account for the fact that this has the potential to cause people to drive more. We calculate this rebound using an elasticity of VMT with respect to the fuel cost per mile of -0.2221 from Small & Van Dender (2007) (authors’ Table 5). We define fuel cost per mile as the price of gasoline (dollars per gallon) divided by the vehicle’s fuel economy. We then multiply this elasticity by the policy’s induced percent change in the fuel cost per mile. For example, if a policy causes drivers to upgrade from a 25.38 MPG vehicle (the average fuel economy of a new light-duty vehicle in 2020) to a 26.38 MPG vehicle, and the price of gasoline was \$2.27 (the 2020 price of gasoline), the percent change in the fuel cost of driving was -3.79%, resulting in a rebound of 0.842%. This rebound offsets some of the initial benefits from driving a more fuel-efficient vehicle. We account for changes in per-mile externalities (accidents, congestion, and $PM_{2.5}$ from tires and

brakes) when accounting for the VMT rebound, as described in our policy appendices.

E Policy Appendices

This appendix outlines our approach to calculating MVPFs for all policies in our sample. Policy-specific appendices within the same category are often repetitive to ensure all appendices contain necessary details. We also note that, in some instances, specific components will not exactly align with those reported in tables, as components in tables have been normalized by each policy’s program cost. Throughout, we round each reported number noting that in some cases this can lead to sums that do not fully add up to do rounding.

E.1 Wind Production Tax Credit

Our category average MVPF for the wind production tax credit (PTC) is 5.87. This appendix describes the construction of the individual MVPFs that feed into this category average.

Production tax credits incentivize the production of wind energy by paying producers a fixed amount per kilowatt hour of production for the first 10 years of a wind turbine’s lifetime. The PTC was first enacted as part of the Energy Policy Act of 1992 at a rate of 1.5 cents per kWh. Since 1992, the credit has lapsed and been retroactively reinstated over a dozen times, often at varying rates. Most recently, the Inflation Reduction Act extended the PTC through 2024 at its full level of 2.6 cents per kWh. Depending on a wind project’s adherence to various IRA provisions including prevailing wage and apprenticeship standards, domestic content shares, and placement in energy communities, certain developments may receive bonuses or deductions from the 2.6 cent baseline. The 2020 pre-IRA PTC level is 1.5 cents, which is the PTC level used in our 2020 baseline specification. Our baseline 2020 MVPF uses a levelized cost of wind of 3.3 cents per kWh, and we present robustness to higher levelized costs (Wiser et al. 2023).

To construct the MVPF, we estimate the individual components of the WTP and government cost. The WTP consists of the mechanical inframarginal transfer to wind developers, local and global environmental externalities, and learning-by-doing effects. The government cost consists of the mechanical transfer, the fiscal externality from induced turbine construction, and a climate fiscal externality described in Section 4. Work by Hitaj (2013), Metcalf (2010), and Shrimali et al. (2015) provide the primary causal estimates for this analysis. They report the behavioral response of wind turbine investment with respect to the production tax credit.

In order to calculate the initial upfront externalities from induced changes in the PTC, we calculate the behavioral response of wind installations with respect to a change in the PTC level. For the learning-by-doing benefits induced by the PTC, we calculate the behavioral response of wind installations with respect to changes in the levelized cost of wind generation. Since wind PTCs do not last for the lifetime of the turbine, these two elasticities are distinct. For example, a one cent increase in the wind PTC level per kWh does not correspond to a one cent decrease in the net cost of generation per kWh. We transform the causal estimate provided by each of the papers in our sample into a semi-elasticity of wind installations with respect to a one cent change in the wind PTC as well as an elasticity of wind installations with respect to a 1% change in the levelized cost of wind generation.

In order to compute the LCOE of wind net of the PTC, we discount the flow of LCOE costs and PTC benefits to present day using a weighted average cost of capital of 2.80%, which is consistent with the rate used in the construction of the LCOE (Wiser et al. 2023). The present

discounted value (PDV) of the 3.3 cent LCOE per kWh over 25 years is \$0.6040 and the PDV of the 2020 PTC over 10 years is \$0.1329. Therefore, the PDV of the LCOE in 2020 net of the PTC is \$0.4711. A one cent expansion of the PTC corresponds to a further reduction in the PDV of the LCOE of \$0.0886. We present analogous calculations for the in-context estimates in the MVPF constructions outlined below.

Since the PTC is provided at a per-kWh level, both the government costs and environmental benefits scale proportional to increases in production. Therefore, variables related to the amount of energy output from a wind turbine such as average turbine size and capacity factor only affect the MVPF to the extent in which they change the levelized cost of wind generation.

For ease of interpretation, we will estimate the externalities and cost of a turbine that produces one MWh a year. Our baseline MVPFs use a wind turbine lifetime of 25 years, which is consistent with the lifetime used in the 2020 LCOE calculation from Wiser et al. (2023). While we don't need the level of the capacity factor for our analysis, we do use results from Kay & Ricks (2023) that suggest that capacity reduces by 5% once the PTC is removed. Therefore, for our externality assumptions, we will assume there is one MWh produced for the first ten years followed by 0.95 MWh produced for the next fifteen years. To crosswalk the WTP and cost component values calculated below with those in Table 2, one can divide each component value by the mechanical cost to recover the cost per \$1 of government spending on the policy.

As explained in Appendix C.2, we use EPA's AVERT model to estimate the emissions saved from the marginal kWh displaced by wind energy and monetize those local and global environmental benefits using our social cost and marginal damage estimates. We also assume that the positive supply shock as a result of the PTC lowers the equilibrium price of electricity which increases the quantity demanded. This rebound effect, as described in Appendix D, leads to an offsetting 20% increase in electricity consumption. We incorporate life-cycle emissions from the manufacturing, maintenance, and decommissioning of wind turbines, which the DOE estimates to be 11 grams per kWh generated (DOE 2023c).

Using our baseline assumptions, the monetized local and global environmental benefits from a wind turbine producing one MWh per year (and adjusting for the 5% capacity reduction) is \$202.96 and \$1,556.71 over the 25 years of the turbine's lifetime, respectively. The total rebound effect is \$344.66 and the lifecycle emissions costs are \$51.48.

Our wind MVPFs include learning by doing effects explained in Section 2.3. We apply a learning rate of 0.194 and use annual and cumulative production data from the Land Based Wind Market Report (Way et al. (2022); Wiser et al. (2023)). We do not incorporate changes in producer profits in our WTP. We assume producers of wind turbines also have the ability to produce electricity through other generation sources. If they are optimizing at the margin, there would be no additional profits from the marginal kWh generated from wind relative to other sources.

We construct both in-context MVPFs from the last year of each paper's sample as well as baseline MVPFs that consider a national expansion of the PTC in 2020.

Federal Wind PTC - Hitaj (2013)

Our MVPF for wind PTCs using estimates from Hitaj (2013) are 4.63 [1.28, 11.63] in 2020 and 5.39 in-context. Hitaj (2013) uses a panel dataset of annual wind capacity additions by county

across the U.S. from 1998-2007. Using a linear probability model with county fixed effects and a vector of controls, the paper estimates the change in the probability of installation in response to the subsidy. The direction and significance of the linear probability model coefficient is consistent with the results from the Tobit, Probit, and IV models also presented in the paper. The paper estimates a 0.317% change in the probability of installation in a county for a 1 cent increase in the PTC. By dividing this by the baseline probability of installation, we recover the semi-elasticity of installation with respect to a 1 cent change in the PTC. There are 20,908 counties with zero installed capacity and 612 counties with positive installed capacity during the sample period, leading to a probability of installation of 2.84%. Therefore, a 1 cent increase in the PTC leads to an 11% increase in wind capacity additions ($0.00317/0.0284$).

We can scale the in-context semi-elasticity by the ratio of the costs in-context and the costs in 2020. From the present discounted value exercise above, we know that the LCOE over the lifetime of the wind turbine net of the PTC is \$0.4711. To get the in-context costs, we take the PTC and LCOE estimates from the years in the papers sample (1998-2007). We compute the average LCOE and the average PTC during this timeframe weighted by the amount of capacity additions in each year. Since we don't have capacity additions data prior to 2000, we only compute this average for 2000-2007. The resulting average LCOE and average PTC in-context in nominal dollars is 5.87 and 1.98 cents, respectively. The PDV of the net LCOE minus PTC over the lifetime of the turbine is \$0.8991. The ratio of the in-context and 2020 costs is 1.91. Scaling the in-context semi-elasticity of 0.11 by this ratio results in a 2020 semi-elasticity of -0.21. Therefore, in 2020, a one cent increase in the PTC leads to a 21% increase in wind capacity installations.

Next, we calculate the elasticity with respect to a change in the LCOE. Using the numbers from above, a one cent change in the PTC has a PDV of 0.0886 and the LCOE net of the PTC in-context has a PDV of 0.8991. Therefore, a one cent change in the PTC corresponds to a 9.85% change in the cost. Dividing the percent change in quantity (11.14%) by the percent change in cost (9.85%) results in an elasticity of 1.13.

For learning-by-doing benefits, we use the elasticity with respect to the cost of generation (1.13). For the upfront externalities induced by the PTC, we use the semie elasticity with respect to a one cent change in the PTC (0.21).

Cost The cost is made up of three components: a mechanical transfer, PTC spending from induced demand, and a climate fiscal externality from increased GDP. As explained above, for ease of interpretation, all the components will be estimated for a wind turbine that produces one MWh a year for each year of its lifetime. We imagine the MVPF of a policy that expands the PTC by one cent per kWh. Since there are 1000 kWh in a MWh, the mechanical transfer per year is \$10. Discounting over the ten years that the turbine is eligible for the PTC results in a transfer of \$82.62.

The baseline PTC in 2020 is 1.5 cents per kWh. Since the PTC induces further wind construction, there is a fiscal externality from implementing the PTC. The fiscal externality per induced MWh is \$15. Multiplying this value by the semi-elasticity (0.21 in 2020) and discounting over the first ten years of the wind turbine results in a fiscal externality of \$29.24 in 2020 and \$15.32 in-context.

As described in Section 4, the climate fiscal externality is 1.9% of the global environmental externality, including learning by doing effects which are described in the following section. The climate fiscal externality lowers the government cost by \$6.40 in 2020 and \$5.58 in-context.

Summing together these components, we estimate a total government cost of \$105.45 in 2020 and \$92.36 in-context.

WTP The WTP is made up of four components: a mechanical transfer, environmental externalities, a rebound effect, and learning-by-doing effects. The transfer to inframarginal producers is fully valued and therefore is the same as the mechanical cost of \$82.62. The remaining components are only affected by the induced installations as a result of the policy expansion.

To calculate the environmental externalities arising from wind turbines, we use estimates of the marginal emissions avoided from wind energy from AVERT and forecast the grid using REPEAT as described in Appendix C.2. We monetize these estimates and discount them over the 25 years of the wind turbine’s lifetime. The monetized environmental externality per kWh is 2.65 and 13.28 cents for local and global pollutants for the U.S. in 2020, respectively. The corresponding values in 2007 are 10.51 and 8.76 cents. Our grid forecasting predicts that these values will decrease over time. Converting each year’s externalities to a per MWh level, multiplying by the semi-elasticity, and discounting over the 25 year lifetime of the turbine leads to a global and local environmental externality in 2020 of \$324.81 and \$43.17, respectively. For the in-context specification, the values are \$193.17 and \$87.50.

The rebound effect, explained in Appendix D, is roughly 20% of the local and global externality. This corresponds to a rebound of \$72.08 in 2020 and \$54.97 in-context, respectively. We also include lifecycle emissions costs of 11 grams per kWh. The resulting lifecycle costs are \$10.74 in 2020 and \$3.34 in-context.

We incorporate learning by doing effects using the modeling approach discussed in Section 2.3. The learning by doing effects for wind use a learning rate of 0.194 (Way et al. 2022). The cumulative wind production in 2020 is 742,689 MW and in 2007 is 93,924 MW. The static wind production in 2020 is 92,490 MW and in 2007 is 19,967 MW. The resulting environmental and price reduction benefits from learning by doing are \$82.42 and \$37.62 in 2020, respectively. For 2007, they are \$142.69 and \$49.76. Summing across all WTP components, we arrive at a 2020 WTP of \$487.83 and an in-context WTP of \$497.42. The resulting MVPF is 4.63 in 2020 and 5.39 in-context.

Federal Wind PTC - Metcalf (2010)

Our MVPF for wind PTCs using estimates from Metcalf (2010) are 5.30 [2.65, 9.28] in 2020 and 6.43 in-context. Metcalf (2010) uses data on wind generation investments in the U.S. between 1990 and 2007. For each wind turbine, the paper constructs a user cost of capital measure that takes into account regional variation in costs and corporate tax rates. Metcalf estimates a Tobit regression with state and year fixed effects. The coefficient on the user cost of capital implies an elasticity of 1.3 as reported in the paper. In the standard user cost model, this elasticity corresponds to the elasticity of investment in turbines with respect their price. We feed this elasticity into our learning by doing model.

To get the upfront externalities, we convert the 1.3 elasticity with respect to the cost of generation into a semi-elasticity with respect to a one cent change in the PTC. To do so, we divide the elasticity by the 2020 cost per kWh. Using the values from above, we know that a one cent change in the PTC corresponds to a 14.67% change in the LCOE in 2020 ($0.0886 / 0.6040$). Since the 2020 LCOE is 3.299 cents and the 2020 PTC is 1.5 cents, the resulting LCOE net of the PTC is 2.573 cents ($3.299 \cdot (1 - (1.5 \cdot 0.1412))$). Dividing the 1.3 elasticity by 2.573 gives the semi-elasticity with respect to a one cent change in the cost of generation.

Since we are interested in the semi-elasticity with respect to a one cent change in the PTC, we translate a one cent change in the PTC to a one cent change in the LCOE. The ratio of the PDV of a one cent change in the PTC to the PDV of a one cent change in the LCOE is 0.4839. This results in a 2020 semi-elasticity of -0.24. Therefore, a one cent increase in the PTC leads to a 24% increase in wind capacity installations. Using the corresponding values in-context results in a semi-elasticity of -0.12.

Cost The cost is made up of three components: a mechanical transfer, PTC spending from induced demand, and a climate fiscal externality from increased GDP. As explained above, for ease of interpretation, all the components will be estimated for a wind turbine that produces one MWh a year for each year of its lifetime. We imagine the MVPF of a policy that expands the PTC by one cent per kWh. Since there are 1000 kWh in a MWh, the mechanical transfer per year is \$10. Discounting over the ten years that the turbine is eligible for the PTC results in a transfer of \$82.62.

The baseline PTC in 2020 is 1.5 cents per kWh. Since the PTC induces further wind construction, there is a fiscal externality from implementing the PTC. The fiscal externality per induced MWh is \$15. Multiplying this value by the semi-elasticity (0.24 in 2020) and discounting over the first ten years of the wind turbine results in a fiscal externality of \$33.60 in 2020 and \$16.51 in-context.

As described in Section 4, the climate fiscal externality is 1.9% of the global environmental externality, including learning by doing effects which are described in the following section. The climate fiscal externality lowers the government cost by \$7.79 in 2020 and \$7.01 in-context.

Summing together these components, we estimate a total government cost of \$108.43 in 2020 and \$92.12 in-context.

WTP The WTP is made up of four components: a mechanical transfer, environmental externalities, a rebound effect, and learning-by-doing effects. The transfer to inframarginal producers is fully valued and therefore is the same as the mechanical cost of \$82.62. The remaining components are only affected by the induced installations as a result of the policy expansion.

To calculate the environmental externalities arising from wind turbines, we use estimates of the marginal emissions avoided from wind energy from AVERT and forecast the grid using REPEAT as described in Appendix C.2. We monetize these estimates and discount them over the 25 years of the wind turbine's lifetime. The monetized environmental externality per kWh is 2.65 and 13.28 cents for local and global pollutants for the U.S. in 2020, respectively. The corresponding values in 2007 are 10.51 and 8.76 cents. Our grid forecasting predicts that these values will decrease over time. Converting each year's externalities to a per MWh level, multiplying by the semi-elasticity, and discounting over the 25 year lifetime of the turbine leads to a global and local environmental externality in 2020 of \$373.26 and \$49.62, respectively. For the in-context specification, the values are \$208.16 and \$94.29.

The rebound effect, explained in Appendix D, is roughly 20% of the local and global externality. This corresponds to a rebound of \$82.83 in 2020 and \$59.24 in-context, respectively. We also include lifecycle emissions costs of 11 grams per kWh. The resulting lifecycle costs are \$12.34 in 2020 and \$3.59 in-context.

We incorporate learning by doing effects using the modeling approach discussed in Section 2.3. The learning by doing effects for wind use a learning rate of 0.194 (Way et al. 2022). The cumulative wind production in 2020 is 742,689 MW and in 2007 is 93,924 MW. The static wind production in 2020 is 92,490 MW and in 2007 is 19,967 MW. The resulting environmental and

price reduction benefits from learning by doing are \$117.89 and \$46.26 in 2020, respectively. For 2007, they are \$207.93 and \$62.08. Summing across all WTP components, we arrive at a 2020 WTP of \$574.48 and an in-context WTP of \$592.25. The resulting MVPF is 5.30 in 2020 and 6.43 in-context.

Federal Wind PTC - Shrimali et al. (2015)

Our MVPF for wind PTCs using estimates from Shrimali et al. (2015) is 7.55 [1.74, ∞] in 2020 and 8.04 in-context. The paper uses data on wind generation investments in the U.S. between 1990 and 2011. They estimate a state fixed effects model with a vector of social, economic, and policy control variables. The coefficient on the PTC dummy variable in their regression suggests that the tax credit is responsible for adding 28.15 MW of capacity annually per state, equivalent to 1407.5 MW across the U.S.

To convert this to an elasticity with respect to a change in the cost of wind generation, we calculate the percent change in price and quantity from the PTC. To get the average PTC and average annual capacity additions in-context, we weight each year's PTC and capacity additions by the level of annual capacity additions. An upper bound on the percent change on quantity would be to take a simple average and a lower bound would be to place all the weight on the largest capacity addition change. Our approach to weight by annual capacity additions results in an average capacity addition of 6,328.33 MW and the average PTC in 2020 dollars is 2.375 cents. The average in-context LCOE in 2020 dollars was 8.59 cents.

First, we calculate the percent change in price. As calculated in the MVPF above, the ratio of the PDV of a one cent change in the PTC to the PDV of a one cent change in the LCOE is 0.4839. We use an arc elasticity construction since there is a large non-marginal price and quantity change. The price change, given by $(2.375 * 0.4839)/(8.59 - (2.375 * 0.4839 * 0.50))$, results in a 14.34% change in price. An analogous calculation can be done for quantity $1407.5/(6328.33 - 1407.5 * 0.5)$ and results in a 25.02% change in quantity. Dividing the percent change in quantity by the percent change in price results in an elasticity of 1.746.

Next, we convert the elasticity with respect to the cost of generation to a semi-elasticity with respect to a one cent change in the PTC. We use the same calculation as the one outlined in the MVPF using estimates from Metcalf (2010). Since the 2020 LCOE is 3.299 cents and the 2020 PTC is 1.5 cents, the resulting LCOE net of the PTC is 2.573 cents $(3.299 * (1 - (1.5 * 0.1412)))$. Dividing the 1.746 elasticity by 2.573 gives the semi-elasticity with respect to a one cent change in the cost of generation. Since we are interested in the semi-elasticity with respect to a one cent change in the PTC, we translate a one cent change in the PTC to a one cent change in the LCOE. The ratio of the PDV of a one cent change in the PTC to the PDV of a one cent change in the LCOE is 0.4839. This results in a 2020 semi-elasticity of -0.33. Therefore, a one cent increase in the PTC leads to a 33% increase in wind capacity installations. Using the corresponding values in-context results in a semi-elasticity of -0.11.

For the in-context specification, we use externality values for the US in 2011, the last year of the sample.

Cost The cost is made up of three components: a mechanical transfer, PTC spending from induced demand, and a climate fiscal externality from increased GDP. As explained above, for ease of interpretation, all the components will be estimated for a wind turbine that produces one MWh a year for each year of its lifetime. We imagine the MVPF of a policy that expands the PTC by one cent per kWh. Since there are 1000 kWh in a MWh, the mechanical transfer

per year is \$10. Discounting over the ten years that the turbine is eligible for the PTC results in a transfer of \$82.62.

The baseline PTC in 2020 is 1.5 cents per kWh. Since the PTC induces further wind construction, there is a fiscal externality from implementing the PTC. The fiscal externality per induced MWh is \$15. Multiplying this value by the semi-elasticity (0.33 in 2020) and discounting over the first ten years of the wind turbine results in a fiscal externality of \$45.11 in 2020 and \$15.60 in-context.

As described in Section 4, the climate fiscal externality is 1.9% of the global environmental externality, including learning by doing effects which are described in the following section. The climate fiscal externality lowers the government cost by \$12.54 in 2020 and \$9.26 in-context.

Summing together these components, we estimate a total government cost of \$115.19 in 2020 and \$88.97 in-context.

WTP The WTP is made up of four components: a mechanical transfer, environmental externalities, a rebound effect, and learning-by-doing effects. The transfer to inframarginal producers is fully valued and therefore is the same as the mechanical cost of \$82.62. The remaining components are only affected by the induced installations as a result of the policy expansion.

To calculate the environmental externalities arising from wind turbines, we use estimates of the marginal emissions avoided from wind energy from AVERT and forecast the grid using REPEAT as described in Appendix C.2. We monetize these estimates and discount them over the 25 years of the wind turbine's lifetime. The monetized environmental externality per kWh is 2.65 and 13.28 cents for local and global pollutants for the U.S. in 2020, respectively. The corresponding values in 2007 are 10.51 and 8.76 cents. Our grid forecasting predicts that these values will decrease over time. Converting each year's externalities to a per MWh level, multiplying by the semi-elasticity, and discounting over the 25 year lifetime of the turbine leads to a global and local environmental externality in 2020 of \$501.18 and \$66.62, respectively. For the in-context specification, the values are \$198.99 and \$58.96.

The rebound effect, explained in Appendix D, is roughly 20% of the local and global externality. This corresponds to a rebound of \$111.21 in 2020 and \$50.52 in-context, respectively. We also include lifecycle emissions costs of 11 grams per kWh. The resulting lifecycle costs are \$16.57 in 2020 and \$4.08 in-context.

We incorporate learning by doing effects using the modeling approach discussed in Section 2.3. The learning by doing effects for wind use a learning rate of 0.194 (Way et al. 2022). The cumulative wind production in 2020 is 742,689 MW and in 2007 is 93,924 MW. The static wind production in 2020 is 92,490 MW and in 2007 is 19,967 MW. The resulting environmental and price reduction benefits from learning by doing are \$270.71 and \$75.98 in 2020, respectively. For 2011, they are \$337.10 and \$92.21. Summing across all WTP components, we arrive at a 2020 WTP of \$869.33 and an in-context WTP of \$715.28. The resulting MVPF is 7.55 in 2020 and 8.04 in-context.

Wind Feed in Tariffs (FIT)

To supplement our analysis of wind production tax credits, we draw from international estimates of wind elasticities that use variation in the wind feed in tariff. We imagine these elasticities apply in the US context and estimate the MVPF of a production tax credit using the implied elasticity from the FIT. We take elasticities for Germany, Spain, UK, and France from Bolkesj 

et al. (2014), the EU from Nicolini & Tavoni (2017), and a second estimate for Germany from Hitaj & Löschel (2019).

The feed in tariff is the price that wind installers are paid for energy generation. Since the installation market is competitive, the FIT closely tracks the LCOE in each country. The FIT may also price in other benefits such as price stability. We assume that the response from US wind installers to changes in the LCOE is analogous to the response of European wind installers to changes in the FIT. Therefore, we estimate the elasticity with respect to the FIT and use this value in our MVPF construction for wind PTCs.

The elasticity construction for each MVPF is outlined below. After calculating the implied elasticity, the 2020 MVPF is constructed using the same approach and externality values as the US PTC estimates. We do not estimate in-context MVPFs for these papers.

Germany FIT using estimates from Bolkesjø et al. (2014)

Bolkesjø et al. (2014) uses a panel dataset with fixed effects to estimate the impact of the FIT on installed wind capacity. The paper estimates that cumulative capacity in Germany increases by 6,449 MW in response to a one euro-cent change in the FIT per kWh. To estimate the MVPF, we divide the percent change in quantity installed by the percent change in the FIT. The baseline cumulative capacity in the year of the estimation (2012) in Germany was 31,308 MW (EWEA 2013). The percent change in capacity installed is 20.60%. The average feed-in tariff value in 2012 in Germany is \$0.115 per kWh (OECD 2022). Converting this to Euros using the average exchange rate during the sample period (1.20) leads to a mean FIT of \$0.096 euro cents. The percent change in price from a one euro-cent change is 10.42%. Therefore, the elasticity is 1.97.

Following the approach of our wind PTC MVPFs, the elasticity of 1.97 leads to an MVPF of 9.15 in 2020.

Spain FIT using estimates from Bolkesjø et al. (2014)

Bolkesjø et al. (2014) uses a panel dataset with fixed effects to estimate the impact of the FIT on installed wind capacity. The paper estimates that cumulative capacity in Spain increases by 4,424 MW in response to a one euro-cent change in the FIT per kWh. To estimate the MVPF, we divide the percent change in quantity installed and the percent change in the FIT. The baseline cumulative capacity in the year of the estimation (2011) in Spain was 22,796 MW (EWEA 2013). The percent change in capacity installed is 19.41%. The average feed-in tariff value in 2011 in Spain is \$0.108 per kWh (OECD 2022). Converting this to Euros using the average exchange rate during the sample period (1.20) leads to a mean FIT of \$0.09 euro cents. The percent change in price from a one euro-cent change is 11.11%. Therefore, the elasticity is 1.75.

Following the approach of our wind PTC MVPFs, the elasticity of 1.75 leads to an MVPF of 7.55 in 2020.

France FIT using estimates from Bolkesjø et al. (2014)

Bolkesjø et al. (2014) uses a panel dataset with fixed effects to estimate the impact of the FIT on installed wind capacity. The paper estimates that cumulative capacity in France increases by 1,245 MW in response to a one euro-cent change in the FIT per kWh. To estimate the MVPF, we divide the percent change in quantity installed by the percent change in the FIT.

The baseline cumulative capacity in the year of the estimation (2012) in France was 7,564 MW (EWEA 2013). The percent change in capacity installed is 16.46%. The average feed-in tariff value in 2012 in France is \$0.105 per kWh (OECD 2022). Converting this to Euros using the average exchange rate during the sample period (1.20) leads to a mean FIT of \$0.088 euro cents. The percent change in price from a one euro-cent change is 11.43%. Therefore, the elasticity is 1.44.

Following the approach of our wind PTC MVPFs, the elasticity of 1.44 leads to an MVPF of 5.91 in 2020.

UK FIT using estimates from Bolkesjø et al. (2014)

Bolkesjø et al. (2014) uses a panel dataset with fixed effects to estimate the impact of the FIT on installed wind capacity. The paper estimates that cumulative capacity in the UK increases by 704 MW in response to a one euro-cent change in the FIT per kWh. To estimate the MVPF, we divide the percent change in quantity installed by the percent change in the FIT. The baseline cumulative capacity in the year of the estimation (2012) in the UK was 8,445 MW (EWEA 2013). The percent change in capacity installed is 8.34%. The average feed-in tariff value in 2012 in the UK is \$0.086 per kWh (OECD 2022). Converting this to Euros using the average exchange rate during the sample period (1.20) leads to a mean FIT of \$0.072 euro cents. The percent change in price from a one euro-cent change is 13.95%. Therefore, the elasticity is 0.597.

Following the approach of our wind PTC MVPFs, the elasticity of 0.597 leads to an MVPF of 2.82 in 2020.

Germany FIT using estimates from Hitaj & Löschel (2019)

Hitaj & Löschel (2019) use variation in the PTC level over time and regions to estimate the impact of the FIT on wind capacity additions in Germany. The paper estimates that annual capacity additions in Germany increased by 2,247.6 KW in response to a one euro-cent change in the FIT per kWh. To estimate the MVPF, we divide the percent change in quantity installed by the percent change in the FIT. Average capacity additions, in KW, among observations with positive capacity additions was 11,891 KW. The percent change in capacity installed is 18.90%. The average feed-in tariff value during the sample period in Germany is 8.81 euro cents per kWh. The percent change in price from a one euro-cent change is 11.35%. Therefore, the elasticity is 1.665.

Following the approach of our wind PTC MVPFs, the elasticity of 1.665 leads to an MVPF of 7.07 in 2020.

EU Renewable Energy Incentives using estimates from Nicolini & Tavoni (2017)

Nicolini & Tavoni (2017) use variation in renewable energy incentives across five EU countries to estimate the impact of the FIT on renewable energy capacity. Since the paper pools together five renewable energy sources (biomass, geothermal, hydroelectric, solar, wind), we exclude it from our main wind analysis in Figure 2.

The paper employs four models (OLS, fixed effects, random effects, and Hausman-Taylor) to estimate the impact of a one euro-cent increase in the the lagged FIT on the total installed renewable energy capacity. A precision weighted average of these four estimates results in a semi-elasticity with respect to total renewable capacity of 3.5%. The mean tariff over the five

countries and ten years of the sample is \$0.058 per kWh or \$0.048 euro-cents (OECD 2022). A one euro-cent change in the FIT corresponds to a 20.69% change. Dividing the percent change in quantity (3.5%) by the percent change in FIT (20.69%) results in an elasticity of 0.17.

This is the elasticity for all renewable energy generation. However, if one were to imagine this as just the wind elasticity in the US, the MVPF of the 2020 PTC would be 1.50.

E.2 Solar Investment Tax Credit

Our category average MVPF for the residential solar investment tax credit in 2020 is 3.86. The investment tax credit (ITC) was created by the Energy Policy Act of 2005. The ITC is a non-refundable federal tax credit that applies to the total installation cost of residential and utility scale solar. Since our literature review produced no estimates of utility scale solar’s response to the ITC, we focus our analysis on rooftop solar. In certain states, homeowners are eligible for state rebates in addition to the federal tax credit. State rebates do not reduce the amount of federal ITC a homeowner can claim (and vice versa). In 2020, the ITC was 26%. Most recently, the Inflation Reduction Act increased the ITC to 30%. Our baseline MVPF will represent a marginal expansion to the 26% ITC level using 2020 externality values.

To construct the MVPF, we start by estimating the individual components of the WTP. The WTP consists of a mechanical inframarginal transfer to residential solar consumers and installers, local and global environmental externalities, learning by doing effects, and utility profit losses. To estimate the behavioral change induced by the ITC, we use price elasticity estimates from Crago & Chernyakhovskiy (2017), Gillingham & Tsvetanov (2019), Pless & van Benthem (2019), and Hughes & Podolefsky (2015). We multiply these elasticities by the externality per dollar of spending on residential solar, V/p , to get the externality components of the MVPF.

The output per watt of a solar panel varies across the US. For our baseline estimates, we take the approach of the Department of Energy’s National Renewable Energy Laboratory (Ramaswamy et al. 2022). The NREL uses estimates of solar panel efficiency from the geographic center of the contiguous US - Fredonia, Kansas. Fredonia is also roughly in the center of the distribution of solar panel efficiency across the US. Using a tilt of 20 degrees, azimuth of 214 degrees, and inverter efficiency of 96%, the annual kWh production per watt is 1.44 (Barbose et al. (2020); Ramaswamy et al. (2022)). We use the NREL’s PVWatts calculator to estimate this value for individual states for our in-context MVPF estimates (NREL 2022a). We also assume that solar panels have a lifetime of 25 years, which is in the middle of existing estimates that range from 20 to 30 years.

We use the model in Section 2.3 to account for potential learning by doing externalities. The cost of solar has been declining over time. In 2020, the NREL estimates that the cost per watt of residential solar was \$3.13 in 2022 dollars (NREL 2022b). Ten years earlier, the cost was estimated to be \$8.70 per watt. For the learning by doing externalities, we apply a learning rate of 0.319 and use annual and cumulative production data from the International Renewable Energy Agency (Way et al. (2022); IRENA (2023a)).

As explained in Appendix C.2, we use EPA’s AVERT model to estimate the emissions saved from the marginal kWh displaced by solar energy and monetize those local and global environmental benefits using marginal damage estimates. We also assume that the positive supply shock as a result of the ITC lowers the equilibrium price of electricity which increases

the quantity demanded. This rebound effect, as described in Appendix D, leads to a 20% increase in electricity consumption. Our environmental externality includes life-cycle emissions from the manufacturing, maintenance, and decommissioning of solar panels which the NREL estimates to be approximately 40 grams per kWh generated (NREL 2013).

The utility market has imperfect competition due to regulation and cost characteristics of the industry. As described in Appendix C.2, we estimate utility profits per kWh using the difference between the levelized cost of electricity (LCOE) and the retail price of electricity. Since solar panels reduce the quantity of electricity purchased from utilities, we incorporate a negative WTP for utility producers. The profit loss per kWh is 1.10 cents in 2020.

Some of the papers in our sample also report pass-through estimates. Since our elasticities are with respect to the price of installation, not with respect to the subsidy level, we scale the elasticity by the pass through rate. A pass-through rate below one lowers the magnitude of each externality and a rate greater than one increases the magnitude. For the in-context MVPFs, we use the pass-through rate estimated in each paper.¹⁶¹ For the baseline 2020 MVPF, we take the average of the pass through estimates for home-owned solar from Gillingham & Tsvetanov (2019) and Pless & van Benthem (2019). This results in a pass-through of 81.1%. The MVPF for third party owned residential solar from Pless & van Benthem (2019) uses a pass-through of 152.8%.

We construct in-context MVPFs using the geography and time period from each paper as well as national MVPFs using harmonized 2020 assumptions. For ease of interpretation, the calculations below imagine a \$1 per watt increase in the rebate level. For consistency, we transform all externalities to be at a per watt level.

Connecticut Residential Solar Investment Program

Using estimates from Gillingham & Tsvetanov (2019), we estimate a baseline MVPF of 1.63 [1.10, 3.55] in 2020 and 0.65 in-context for the Connecticut Residential Solar Investment Program. The paper estimates a price elasticity of demand of -0.65. They use a panel dataset from 2008 to 2014 with the number of annual installations in each census block group. To account for excess zeroes in their count dataset, they employ a Poisson hurdle model using instrumental variables and fixed effects. To address endogeneity in rebate variation, they use local roofing wages and state subsidies as instruments. Our in-context MVPF uses externality values for Connecticut in 2014, the last year in their sample.

To estimate the monetized externalities that feed into the MVPF, we multiply the elasticity by the externality per dollar of spending (V/p). For the subsample with positive installations, the paper reports an average state incentive of \$3.04 and a cost per watt net of the state rebate of \$3.89. The federal ITC is applied on the entire installation cost, not just the post state rebate cost. Therefore, after applying the 30% ITC that was in effect in 2014, the resulting cost per watt (p) in 2014 dollars is \$2.05. For the baseline MVPF, the cost per watt of \$3.13 in 2022 dollars and the 26% ITC in 2020 imply a cost per watt net of the subsidy of \$2.05 in 2020 dollars (NREL 2022b).

For the baseline 2020 MVPF, we use the average solar pass-through rates across the sample

¹⁶¹If a paper does not estimate pass through, we take the pass through estimate from the paper with the most similar context as the sample studied in the paper. The pass-through rate used for each in-context estimate is explained further in the MVPF calculations below.

of 81.1%. For the in-context MVPF, we use the pass-through from the paper of 84%. The in-context MVPF studies the policy in Connecticut in 2014.

Cost The government cost is made up of five components: a mechanical transfer, ITC spending from induced demand, state subsidy spending from induced demand, tax revenue loss from reduced utility profits, and a climate fiscal externality. There is a \$1 mechanical transfer to inframarginal households that would have installed solar panels in the absence of the subsidy. Since there is a federal ITC of 26% of installation costs already in place in 2020, there is a fiscal externality from induced demand. The federal subsidy per watt in 2020 is \$0.72. Since the cost per watt net of subsidies is \$2.05, the fiscal externality per dollar of spending (V/p) is \$0.35 in 2020. To get the fiscal externality per dollar of government spending on the policy, we multiply this value by the product of the elasticity and pass-through rate. The resulting externality from the federal subsidy is \$0.185 in 2020 and \$0.340 in-context.

The fiscal externality from state level rebates is zero in 2020, since we assume there are no state rebates in our baseline MVPF. For our in-context externality, the state rebate in Connecticut as estimated in the paper is \$3.04 per watt. To get the fiscal externality, we divide this by the in-context cost per watt of \$1.81 and multiply by the product of the elasticity and pass through to get an in-context state fiscal externality of \$0.886.

We assume that utilities lose profit from customers switching to rooftop solar and therefore the government loses profit tax revenue. As explained above, one watt produces 1.44 kWh in 2020. Since there is a rebound effect of 20%, from the perspective of a utility company, it is as if they only lose 1.15 kWh per watt worth of revenue. From Appendix C.2, we know that the government revenue from utility profit per kWh is \$0.006 in 2020 and \$0.029 in-context. To get the externality per dollar (V/p), we multiply the externality by 1.15 and take the discounted sum over the solar panels 25 year lifetime. Dividing this by the price per watt gives us V/p and multiplying by the product of the elasticity and pass through results in a fiscal externality from lost government tax revenue of \$0.036 in 2020 and \$0.174 in-context.

Using the method described in Section 2.3, the climate fiscal externality represents 1.9% of the global environmental externality, which includes learning by doing and is net of the rebound effect and life cycle emissions which are explained in the following section. The resulting externality will reduce the government cost by \$0.01 in 2020 and \$0.01 in-context. Summing together these components, we estimate a total government cost of \$1.21 in 2020 and \$2.39 in-context.

WTP The WTP is made up of five components: a mechanical transfer, environmental externalities, a rebound effect, learning-by-doing, and producer profit loss. The \$1 transfer to inframarginal consumers is split by the solar installers and the homeowners. In our baseline MVPF, we use a pass through rate of 81.1%. Therefore, 18.9 cents of the transfer flows to installers and the rest flows to homeowners. Using the estimation strategy described in Appendix C.2, the monetized environmental externality per kWh is \$0.125 and \$0.025 cents for global and local pollutants for the U.S. in 2020, respectively. The corresponding values for Connecticut in 2014 are \$0.082 and \$0.027 cents. While these are the point in time externalities for 2020 and 2014, we allow for the grid to change over the course of the solar panel's 25 year lifetime as described in Appendix C.2. On a per watt basis, solar generates 1.44 kWh per year in 2020 and 1.29 kWh per year in-context. To get the global and local environmental externality per dollar of spending (V/p), we discount the stream of environmental benefits per watt over the 25 year lifetime of solar panels and divide by the cost per watt. To get the value fed into the MVPF we multiply this ratio by the product of the elasticity and pass through. The global

environmental externality also includes the lifecycle cost of solar which is 40 grams of CO₂e per kWh generated. Putting all this together, we arrive at an environmental externality per mechanical dollar of government spending of \$0.078 and \$0.533 for local and global benefits in 2020, respectively. The corresponding values in-context are \$0.043 and \$0.372. The rebound effect is calculated as 20% of the local and global externalities, excluding the lifecycle emissions. The local and global rebound in 2020 is \$0.015 and \$0.115, respectively. The in-context rebound is \$0.008 and \$0.082.

Section 2.3 explains how we calculate learning-by-doing benefits for solar subsidies. Cumulative global production of solar panels was at 176,111 MW in 2014 and 713,918 MW in 2020. Static production of solar panels was at 39,541 MW in 2014 and 128,050 MW in 2020 (IRENA 2023a). Using a learning rate of 0.319, we arrive at a learning by doing effect on prices of \$0.346 and on the environment of \$0.216 in 2020 (Way et al. 2022). The corresponding values for 2014 are \$0.381 and \$0.160.

As described above and in Appendix C.2, the per kWh utility profit loss is \$0.011 in 2020 and \$0.053 cents in Connecticut in 2014. Using our estimate for the kWh generated per watt, discounting over the 25 year solar panel lifetime, and dividing by the cost per watt, we get the externality per dollar of spending on solar. Scaling this value down by the rebound rate (20%) and multiplying by the product of the elasticity and pass through, we arrive at a utility WTP of -\$0.066 in 2020 and -\$0.321 in 2014. Summing up all of the individual WTP components, the resulting WTP is \$1.976 in 2020 and \$1.545 in-context. Dividing by the total government cost, the MVPF is 1.63 in 2020 and 0.65 in-context.

Northeast Solar Rebates

Using estimates from Crago & Chernyakhovskiy (2017), we estimate a baseline MVPF of 4.68 [2.16, 91.72] in 2020 and 4.13 in-context. The authors use a panel dataset of county level installations from 2005 to 2012 for 13 states in the northeast. They exploit inter-temporal variation in the timing of state subsidies for residential solar to estimate the relationship between rebate levels and adoption in the northeast. To account for the endogeneity of the rebate levels and timing, they use year fixed effects and control for time-varying indicators of environmental preferences. They estimate that a \$1 increase in the rebate level increases installations by 47.2%. We apply an in-context pass through rate of 84% from Gillingham & Tsvetanov (2019) who estimate pass through on a sample of residential solar installations in Connecticut. Therefore, to get the increase in installations from a \$1 decrease in the price per watt, we multiply the 47.2% by 1/0.844 to get a increase of 55.9%.

To convert this semi-elasticity to a price elasticity, we multiply the semi-elasticity by the cost per watt during the sample period. To get the cost per watt before the state and federal subsidies, we take an average of the cost per watt from NREL (2022b) weighted by annual solar installations from 2010 to 2012.¹⁶² The resulting cost per watt in-context is \$5.42 in 2008 dollars. The paper reports an average rebate of \$1.13. After including the 30% federal subsidy in place during the sample period, the implied cost per watt net of the subsidy is \$2.66 in 2008 dollars. The resulting price elasticity is 1.49.

To estimate the monetized externalities that feed into the MVPF, we multiply the elasticity by the externality per dollar of spending (V/p). For the baseline MVPF, the cost per watt of

¹⁶²The paper's sample begins in 2005, but the National Renewable Energy Laboratory (NREL) only reports the cost per watt starting in 2010.

\$3.13 in 2022 dollars and the 26% ITC in 2020 imply a cost per watt net of the subsidy (p) of \$2.05. The in-context MVPF is estimated for 2012, the last year of the sample. The in-context MVPF is localized to the 13 northeast states the paper studies. For all geography specific externalities, we take a population weighted average in 2012 to estimate the externality.

Cost The government cost is made up of five components: a mechanical transfer, ITC spending from induced demand, state subsidy spending from induced demand, tax revenue loss from reduced utility profits, and a climate fiscal externality. There is a \$1 mechanical transfer to inframarginal households that would have installed solar panels in the absence of the subsidy. Since there is a federal ITC of 26% of installation costs already in place in 2020, there is a fiscal externality from induced demand. The federal subsidy per watt in 2020 is \$0.72. Since the cost per watt net of subsidies is \$2.05, the fiscal externality per dollar of spending (V/p) is \$0.35 in 2020. To get the fiscal externality per dollar of government spending on the policy, we multiply this value by the product of the elasticity and pass-through rate.¹⁶³ The resulting externality from the federal subsidy is \$0.424 in 2020 and \$0.607 in-context.

The fiscal externality from state level rebates is zero in 2020, since we assume there are no state rebates in our baseline MVPF. For our 2012 in-context externality, the average state rebate as reported in the paper is \$1.21 per watt in 2012 dollars. To get the fiscal externality, we divide this by the in-context cost per watt of \$2.84 in 2012 dollars and multiply by the product of the elasticity and pass through to get an in-context state fiscal externality of \$0.533.

We assume that utilities lose profit from customers switching to rooftop solar and therefore the government loses profit tax revenue. As explained above, one watt produces 1.44 kWh in 2020. Since there is a rebound effect of 20%, from the perspective of a utility company, it is as if they only lose 1.15 kWh per watt worth of revenue. From Appendix C.2, we know that the government revenue from utility profit per kWh is \$0.006 in 2020 and \$0.010 in-context. To get the externality per dollar (V/p), we multiply the externality by 1.15 and take the discounted sum over the solar panels 25 year lifetime. Dividing this by the price per watt gives us V/p and multiplying by the product of the elasticity and pass through results in a fiscal externality from lost government tax revenue of \$0.082 in 2020 and \$0.085 in-context.

Using the method described in Section 4, the climate fiscal externality represents 1.9% of the global environmental externality, which includes learning by doing and is net of the rebound effect and life cycle emissions which are explained in the following section. The resulting externality will reduce the government cost by \$0.076 in 2020 and \$0.079 in-context. Summing together these components, we estimate a total government cost of \$1.431 in 2020 and \$2.144 in-context.

WTP The WTP is made up of five components: a mechanical transfer, environmental externalities, a rebound effect, learning-by-doing, and producer profit loss. The \$1 transfer to inframarginal consumers is split by the solar installers and the homeowners. In our baseline MVPF, we use a pass through rate of 81.1%. Therefore, 18.9 cents of the transfer flows to installers and the rest flows to homeowners. Using the estimation strategy described in Appendix C.2, the monetized environmental externality per kWh is \$0.125 and \$0.025 cents for global and local pollutants for the U.S. in 2020, respectively. For the in-context MVPF, we take the population weighted average of each state's environmental externality per kWh. While these are the point in time externalities, we allow for the grid to change over the course of the solar panel's 25 year lifetime as described in Appendix C.2. On a per watt basis, solar generates

¹⁶³For the baseline MVPF, we use the sample average 81.1% pass-through rate. For the in-context estimate, we use the pass-through of 84% from Gillingham & Tsvetanov (2019).

1.44 kWh per year in 2020 and 1.20 kWh per year in-context. To get the global and local environmental externality per dollar of spending (V/p), we discount the stream of environmental benefits per watt over the 25 year lifetime of solar panels and divide by the cost per watt. To get the value fed into the MVPF we multiply this ratio by the product of the elasticity and pass through. The global environmental externality also includes the lifecycle cost of solar which is 40 grams of CO_{2e} per kWh generated. Putting all this together, we arrive at an environmental externality per mechanical dollar of government spending of \$0.179 and \$1.220 for local and global benefits in 2020, respectively. The corresponding values in-context are \$0.232 and \$0.700. The rebound effect is calculated as 20% of the local and global externalities, excluding the lifecycle emissions. The local and global rebound in 2020 is \$0.035 and \$0.264, respectively. The in-context rebound is \$0.045 and \$0.149.

Appendix 2.3 explains how we apply learning by doing to solar. Cumulative global production of solar panels was at 176,111 MW in 2014 and 713,918 MW in 2020. Static production of solar panels was at 39,541 MW in 2014 and 128,050 MW in 2020 (IRENA, 2023). Using a learning rate of 0.319, we arrive at a learning by doing effect on prices of \$1.610 and on the environment of \$3.132 in 2020 (Way et al. 2022). The corresponding values for 2014 are \$2.365 and \$4.906.

As described above and in Appendix C.2, the per kWh utility profit loss is \$0.011 in 2020 and \$0.018 cents in-context. Using our estimate for the kWh generated per watt, discounting over the 25 year solar panel lifetime, and dividing by the cost per watt, we get the externality per dollar of spending on solar. Scaling this value down by the rebound rate (20%) and multiplying by the product of the elasticity and pass through, we arrive at a utility WTP of -\$0.152 in 2020 and -\$0.157 in 2014. Summing up all of the individual WTP components, the resulting WTP is \$6.690 in 2020 and \$8.852 in-context. Dividing by the total government cost, the MVPF is 4.68 in 2020 and 4.129 in-context.

California Solar Initiative (HO) - Pless & van Benthem (2019)

Using estimates from Pless & van Benthem (2019), we estimate a baseline MVPF of 2.71 in 2020 and 1.79 in-context for the California Solar Initiative. The paper estimates an elasticity for home owned (HO) and third party owned (TPO) residential solar. This section will focus on HO solar and the following section will outline the MVPF calculation for TPO solar.

The California Solar Initiative was enacted in 2007 and is the largest state rebate program for solar in the US. It provides homeowners a lump sum payment for residential solar in addition to the federal ITC. The authors use a panel dataset of solar installations in California from 2010 to 2013. They use price variations caused by sharp changes in the incentive schedule over time and between IOUs to estimate a price elasticity. To calculate the elasticity, we take the implied derivative at the mean price of \$3.70 from their IV demand estimation and multiply by the ratio of the mean price to the mean installations (0.461). The resulting price elasticity for HO residential solar is -1.14.

Our in-context MVPF uses the pass through rate for HO solar estimated in the paper of 77.8%. Our baseline MVPF uses the average pass through rate in our solar sample of 81.1%.

To estimate the monetized externalities that feed into the MVPF, we multiply the elasticity by the externality per dollar of spending (V/p). For the baseline MVPF, the cost per watt of \$3.13 in 2022 dollars and the 26% ITC in 2020 imply a cost per watt net of the subsidy (p) of \$2.05 in 2020 dollars. The in-context MVPF is estimated for 2013, the last year of the sample.

The average HO cost per watt used in the demand estimation net of the state subsidy is \$3.89. The average state subsidy is reported as \$0.42. The 30% federal ITC in effect in 2013 applies to the total installation cost before the state rebate. The cost to the homeowner after the state and federal incentives is \$2.64 per watt.

Cost The government cost is made up of five components: a mechanical transfer, ITC spending from induced demand, state subsidy spending from induced demand, tax revenue loss from reduced utility profits, and a climate fiscal externality. There is a \$1 mechanical transfer to inframarginal households that would have installed solar panels in the absence of the subsidy. Since there is a federal ITC of 26% of installation costs already in place in 2020, there is a fiscal externality from induced demand. The federal subsidy per watt in 2020 is \$0.72. Since the cost per watt net of subsidies is \$2.05, the fiscal externality per dollar of spending (V/p) is \$0.35 in 2020. To get the fiscal externality per dollar of government spending on the policy, we multiply this value by the product of the elasticity and pass-through rate. The resulting externality from the federal subsidy is \$0.324 in 2020 and \$0.398 in-context.

The fiscal externality from state level rebates is zero in 2020, since we assume there are no state rebates in our baseline MVPF. For our in-context externality, the average state rebate as reported in the paper is \$1.21 per watt. To get the fiscal externality, we divide this by the in-context cost per watt of \$2.66 and multiply by the product of the elasticity and pass through to get an in-context state fiscal externality of \$0.143.

We assume that utilities lose profit from customers switching to rooftop solar and therefore the government loses profit tax revenue. As explained above, one watt produces 1.44 kWh in 2020. Since there is a rebound effect of 20%, from the perspective of a utility company, it is as if they only lose 1.15 kWh per watt worth of revenue. From Appendix C.2, we know that the government revenue from utility profit per kWh is \$0.006 in 2020 and \$0.024 in-context. To get the externality per dollar (V/p), we multiply the externality by 1.15 and take the discounted sum over the solar panels 25 year lifetime. Dividing this by the price per watt gives us V/p and multiplying by the product of the elasticity and pass through results in a fiscal externality from lost government tax revenue of \$0.063 in 2020 and \$0.211 in-context.

Using the method described in Section 4, the climate fiscal externality represents 1.9% of the global environmental externality, which includes learning by doing and is net of the rebound effect and life cycle emissions which are explained in the following section. The resulting externality will reduce the government cost by \$0.034 in 2020 and \$0.027 in-context. Summing together these components, we estimate a total government cost of \$1.353 in 2020 and \$1.725 in-context.

WTP The WTP is made up of five components: a mechanical transfer, environmental externalities, a rebound effect, learning-by-doing, and producer profit loss. The \$1 transfer to inframarginal consumers is split by the solar installers and the homeowners. In our baseline MVPF, we use a pass through rate of 81.1%. Therefore, 18.9 cents of the transfer flows to installers and the rest flows to homeowners. Using the estimation strategy described in Appendix C.2, the monetized environmental externality per kWh is \$0.125 and \$0.025 cents for global and local pollutants for the U.S. in 2020, respectively. The corresponding values for California in 2013 are \$0.070 and \$0.008 cents. While these are the point in time externalities for 2020 and 2012, we allow for the grid to change over the course of the solar panel's 25 year lifetime as described in Appendix C.2. On a per watt basis, solar generates 1.44 kWh per year in 2020 and 1.63 kWh per year in-context. To get the global and local environmental externality per dollar of spending (V/p), we discount the stream of environmental benefits per watt over the

25 year lifetime of solar panels and divide by the cost per watt. To get the value fed into the MVPF we multiply this ratio by the product of the elasticity and pass through. The global environmental externality also includes the lifecycle cost of solar which is 40 grams of CO₂e per kWh generated. Putting all this together, we arrive at an environmental externality per mechanical dollar of government spending of \$0.137 and \$0.932 for local and global benefits in 2020, respectively. The corresponding values in-context are \$0.038 and \$0.514. The rebound effect is calculated as 20% of the local and global externalities, excluding the lifecycle emissions. The local and global rebound in 2020 is \$0.027 and \$0.202, respectively. The in-context rebound is \$0.007 and \$0.113.

Appendix 2.3 explains how we apply learning by doing to solar. Cumulative global production of solar panels was at 176,111 MW in 2014 and 713,918 MW in 2020. Static production of solar panels was at 39,541 MW in 2014 and 128,050 MW in 2020 (IRENA, 2023). Using a learning rate of 0.319, we arrive at a learning by doing effect on prices of \$0.864 and on the environment of \$1.081 in 2020 (Way et al. 2022). The corresponding values for 2014 are \$1.011 and \$1.036.

As described above and in Appendix C.2, the per kWh utility profit loss is \$0.011 in 2020 and \$0.044 cents in-context. Using our estimate for the kWh generated per watt, discounting over the 25 year solar panel lifetime, and dividing by the cost per watt, we get the externality per dollar of spending on solar. Scaling this value down by the rebound rate (20%) and multiplying by the product of the elasticity and pass through, we arrive at a utility WTP of -\$0.116 in 2020 and -\$0.388 in 2014. Summing up all of the individual WTP components, the resulting WTP is \$3.670 in 2020 and \$3.090 in-context. Dividing by the total government cost, the MVPF is 2.71 in 2020 and 1.82 in-context.

California Solar Initiative (TPO) - Pless & van Benthem (2019)

Using estimates from Pless & van Benthem (2019), we estimate a baseline MVPF for third party owned solar of 3.82 in 2020 and 3.28 in-context. The paper estimates an elasticity for home owned (HO) and third party owned (TPO) residential solar. This section will focus on the MVPF calculation for TPO. As described in the previous section, we calculate the elasticity by taking the implied derivative at a price of \$3.70 from their IV demand estimation for TPO and multiply by the ratio of the mean price (\$3.70) to the mean installations (0.518). The resulting price elasticity for TPO residential solar is -1.04.

The paper estimates a pass through of 153% for TPO solar. The other papers in our solar sample use a pass through rate of 81.1% for the baseline MVPF. However, for our TPO MVPF, we use the 153% pass through rate for both the in-context and baseline since this is the only estimate of TPO pass through in our sample.

To estimate the monetized externalities that feed into the MVPF, we multiply the elasticity by the externality per dollar of spending (V/p). For the baseline MVPF, the cost per watt of \$3.13 in 2022 dollars and the 26% ITC in 2020 imply a cost per watt net of the subsidy (p) of \$2.05 in 2020 dollars. The in-context MVPF is estimated for 2013, the last year of the sample. The average TPO cost per watt used in the demand estimation net of the state subsidy is \$3.43. The average state subsidy is reported as \$0.41. The 30% federal ITC in effect in 2013 applies to the total installation cost before the rebate. Therefore, the cost to the homeowner after the state and federal incentives is \$2.31 per watt.

Cost The government cost is made up of five components: a mechanical transfer, ITC spending

from induced demand, state subsidy spending from induced demand, tax revenue loss from reduced utility profits, and a climate fiscal externality. There is a \$1 mechanical transfer to inframarginal households that would have installed solar panels in the absence of the subsidy. Since there is a federal ITC of 26% of installation costs already in place in 2020, there is a fiscal externality from induced demand. The federal subsidy per watt in 2020 is \$0.72. Since the cost per watt net of subsidies is \$2.05, the fiscal externality per dollar of spending (V/p) is \$0.35 in 2020. To get the fiscal externality per dollar of government spending on the policy, we multiply this value by the product of the elasticity and pass-through rate. The resulting externality from the federal subsidy is \$0.558 in 2020 and \$0.718 in-context.

The fiscal externality from state level rebates is zero in 2020, since we assume there are no state rebates in our baseline MVPF. For our in-context externality, the average state rebate as reported in the paper is \$1.21 per watt. To get the fiscal externality, we divide this by the in-context cost per watt of \$2.66 and multiply by the product of the elasticity and pass through to get an in-context state fiscal externality of \$0.286.

We assume that utilities lose profit from customers switching to rooftop solar and therefore the government loses profit tax revenue. As explained above, one watt produces 1.44 kWh in 2020. Since there is a rebound effect of 20%, from the perspective of a utility company, it is as if they only lose 1.15 kWh per watt worth of revenue. From Appendix C.2, we know that the government revenue from utility profit per kWh is \$0.006 in 2020 and \$0.024 in-context. To get the externality per dollar (V/p), we multiply the externality by 1.15 and take the discounted sum over the solar panels 25 year lifetime. Dividing this by the price per watt gives us V/p and multiplying by the product of the elasticity and pass through results in a fiscal externality from lost government tax revenue of \$0.108 in 2020 and \$0.432 in-context.

Using the method described in Section 4, the climate fiscal externality represents 1.9% of the global environmental externality, which includes learning by doing and is net of the rebound effect and life cycle emissions which are explained in the following section. The resulting externality will reduce the government cost by \$0.061 in 2020 and \$0.086 in-context. Summing together these components, we estimate a total government cost of \$1.606 in 2020 and \$2.349 in-context.

WTP The WTP is made up of five components: a mechanical transfer, environmental externalities, a rebound effect, learning-by-doing, and producer profit loss. The \$1 transfer to inframarginal consumers is split by the solar installers and the homeowners. In our baseline MVPF, we use a pass through rate of 81.1%. Therefore, 18.9 cents of the transfer flows to installers and the rest flows to homeowners. Using the estimation strategy described in Appendix C.2, the monetized environmental externality per kWh is \$0.125 and \$0.025 cents for global and local pollutants for the U.S. in 2020, respectively. The corresponding values for California in 2013 are \$0.070 and \$0.008 cents. While these are the point in time externalities for 2020 and 2012, we allow for the grid to change over the course of the solar panel's 25 year lifetime as described in Appendix C.2. On a per watt basis, solar generates 1.44 kWh per year in 2020 and 1.63 kWh per year in-context. To get the global and local environmental externality per dollar of spending (V/p), we discount the stream of environmental benefits per watt over the 25 year lifetime of solar panels and divide by the cost per watt. To get the value fed into the MVPF we multiply this ratio by the product of the elasticity and pass through. The global environmental externality also includes the lifecycle cost of solar which is 40 grams of CO₂e per kWh generated. Putting all this together, we arrive at an environmental externality per mechanical dollar of government spending of \$0.235 and \$0.347 for local and global benefits in

2020, respectively. The corresponding values in-context are \$0.077 and \$1.053. The rebound effect is calculated as 20% of the local and global externalities, excluding the lifecycle emissions. The local and global rebound in 2020 is \$0.046 and \$0.347, respectively. The in-context rebound is \$0.015 and \$0.232.

Appendix 2.3 explains how we apply learning by doing to solar. Cumulative global production of solar panels was at 176,111 MW in 2014 and 713,918 MW in 2020. Static production of solar panels was at 39,541 MW in 2014 and 128,050 MW in 2020 (IRENA, 2023). Using a learning rate of 0.319, we arrive at a learning by doing effect on prices of \$1.371 and on the environment of \$1.982 in 2020 (Way et al. 2022). The corresponding values for 2014 are \$2.200 and \$3.878.

As described above and in Appendix C.2, the per kWh utility profit loss is \$0.011 in 2020 and \$0.044 cents in-context. Using our estimate for the kWh generated per watt, discounting over the 25 year solar panel lifetime, and dividing by the cost per watt, we get the externality per dollar of spending on solar. Scaling this value down by the rebound rate (20%) and multiplying by the product of the elasticity and pass through, we arrive at a utility WTP of -\$0.200 in 2020 and -\$0.795 in 2014. Summing up all of the individual WTP components, the resulting WTP is \$6.128 in 2020 and \$7.694 in-context. Dividing by the total government cost, the MVPF is 3.82 in 2020 and 3.28 in-context.

California Solar Initiative - Hughes & Podolefsky (2015)

Using estimates from Hughes & Podolefsky (2015), we estimate a baseline MVPF of 5.06 in 2020 and 1.87 in-context for the California Solar Initiative. This paper focuses on home owned solar rebates offered through the California Solar Initiative, the largest state rebate program in the US. The paper uses a panel dataset of installations at the zip code level from 2007 to 2012. The authors exploit geographic variation in incentive levels to measure the relationship between rebates and adoption. The paper finds that a \$1 increase in the rebate level per watt leads to a 101.09% increase in solar adoption. We apply an in-context pass through rate of 77.8% from Pless & van Benthem (2019). To be consistent with the approach in other papers in the sample, we first convert this semi-elasticity with respect to the rebate level into an elasticity with respect to price. To get the percent change in installations from a \$1 change in the price, we multiply the 101.09% by $1/0.778$, resulting in an increase in installations of 129.93%. Next, we estimate the percent change in price that corresponds to a \$1 increase in the rebate level.

As reported in the paper, the average cost per watt before the state and federal incentives are applied across the three IOUs in California is \$8.09 in 2012 dollars. A weighted average of the state rebate is \$1.62 per watt. After applying the 30% federal ITC and state rebate, the in-context cost per watt was \$4.04. Taking the product of the cost per watt and the semi-elasticity results in an elasticity of -5.25. We use the -5.25 elasticity to calculate the static environmental and fiscal externalities from the policy. For the dynamic learning-by-doing benefits, the solution to the differential equation is undefined when the product of the learning rate and elasticity exceeds one. Intuitively, this suggests that the learning-by-doing benefits are infinite in these cases. For this policy, we decide to use the elasticity estimate of -1.13 from Pless & van Benthem (2019) to estimate the learning-by-doing benefits since the -1.13 estimate comes from the same CSI solar subsidy program.

Our in-context MVPF uses the pass through rate for HO solar estimated in Pless & van Benthem (2019) of 77.8%. Our baseline MVPF uses the average pass through rate in our solar

sample of 81.1%.

To estimate the monetized externalities that feed into the MVPF, we multiply the elasticity by the externality per dollar of spending (V/p). For the value of p in the baseline MVPF, the cost per watt of \$3.13 in 2022 dollars and the 26% ITC in 2020 imply a cost per watt net of the subsidy (p) of \$2.05 in 2020 dollars. The in-context MVPF is estimated for California in 2012, the last year of the sample. For the in-context MVPF, we use the semi-elasticity (ϵ/p) as reported in the paper of 1.01.

Cost The government cost is made up of five components: a mechanical transfer, ITC spending from induced demand, state subsidy spending from induced demand, tax revenue loss from reduced utility profits, and a climate fiscal externality. There is a \$1 mechanical transfer to inframarginal households that would have installed solar panels in the absence of the subsidy. Since there is a federal ITC of 26% of installation costs already in place in 2020, there is a fiscal externality from induced demand. The federal subsidy per watt in 2020 is \$0.72. Since the cost per watt net of subsidies is \$2.05, the fiscal externality per dollar of spending (V/p) is \$0.35 in 2020. To get the fiscal externality per dollar of government spending on the policy, we multiply this value by the product of the elasticity and pass-through rate. The resulting externality from the federal subsidy is \$1.496 in 2020 and \$1.909 in-context.

The fiscal externality from state level rebates is zero in 2020, since we assume there are no state rebates in our baseline MVPF. For our in-context externality, the average state rebate as reported in the paper is \$1.62 per watt. To get the fiscal externality, we multiply this by the product of the semi-elasticity of 1.01 and the pass through to get an in-context state fiscal externality of \$1.273.

We assume that utilities lose profit from customers switching to rooftop solar and therefore the government loses profit tax revenue. As explained above, one watt produces 1.44 kWh in 2020. Since there is a rebound effect of 20%, from the perspective of a utility company, it is as if each watt is only producing 1.15 kWh per watt. From Appendix C.2, we know that the government revenue from utility profit per kWh is \$0.006 in 2020 and \$0.023 in-context. To get the externality per dollar (V/p), we multiply the externality by 1.15 and take the discounted sum over the solar panels 25 year lifetime. Dividing this by the price per watt gives us V/p and multiplying by the product of the elasticity and pass through results in a fiscal externality from lost government tax revenue of \$0.291 in 2020 and \$0.399 in-context.

Using the method described in Section 4, the climate fiscal externality represents 1.9% of the global environmental externality, which includes learning by doing and is net of the rebound effect and life cycle emissions which are explained in the following section. The resulting externality will reduce the government cost by \$0.157 in 2020 and \$0.058 in-context. Summing together these components, we estimate a total government cost of \$2.630 in 2020 and \$4.527 in-context.

WTP The WTP is made up of five components: a mechanical transfer, environmental externalities, a rebound effect, learning-by-doing, and producer profit loss. The \$1 transfer to inframarginal consumers is split by the solar installers and the homeowners. In our baseline MVPF, we use a pass through rate of 81.1%. Therefore, 18.9 cents of the transfer flows to installers and the rest flows to homeowners. Using the estimation strategy described in Appendix C.2, the monetized environmental externality per kWh is \$0.125 and \$0.025 cents for global and local pollutants for the U.S. in 2020, respectively. The corresponding values for California in 2013 are \$0.070 and \$0.008 cents. While these are the point in time externalities for 2020 and 2012, we allow for the grid to change over the course of the solar panel's 25 year lifetime

as described in Appendix C.2. On a per watt basis, solar generates 1.44 kWh per year in 2020 and 1.39 kWh per year in-context. To get the global and local environmental externality per dollar of spending (V/p), we discount the stream of environmental benefits per watt over the 25 year lifetime of solar panels and divide by the cost per watt. To get the value fed into the MVPF we multiply this ratio by the product of the elasticity and pass through. The global environmental externality also includes the lifecycle cost of solar which is 40 grams of CO_{2e} per kWh generated. Putting all this together, we arrive at an environmental externality per mechanical dollar of government spending of \$0.631 and \$4.299 for local and global benefits in 2020, respectively. The corresponding values in-context are \$0.081 and \$1.059. The rebound effect is calculated as 20% of the local and global externalities, excluding the lifecycle emissions. The local and global rebound in 2020 is \$0.124 and \$0.930, respectively. The in-context rebound is \$0.016 and \$0.231.

Appendix 2.3 explains how we apply learning by doing to solar. Cumulative global production of solar panels was at 176,111 MW in 2014 and 713,918 MW in 2020. Static production of solar panels was at 39,541 MW in 2014 and 128,050 MW in 2020 (IRENA, 2023). As explained above, we use the -1.13 elasticity from Pless & van Benthem (2019) to estimate the learning-by-doing benefits. Using a learning rate of 0.319, we arrive at a learning by doing effect on prices of \$3.987 and on the environment of \$4.988 in 2020 (Way et al. 2022). The corresponding values for 2014 are \$5.001 and \$2.315.

As described above and in Appendix C.2, the per kWh utility profit loss is \$0.011 in 2020 and \$0.044 cents in-context. Using our estimate for the kWh generated per watt, discounting over the 25 year solar panel lifetime, and dividing by the cost per watt, we get the externality per dollar of spending on solar. Scaling this value down by the rebound rate (20%) and multiplying by the product of the elasticity and pass through, we arrive at a utility WTP of -\$0.535 in 2020 and -\$0.734 in 2014. Summing up all of the individual WTP components, the resulting WTP is \$13.316 in 2020 and \$8.476 in-context. Dividing by the total government cost, the MVPF is 5.06 in 2020 and 1.87 in-context.

E.3 Battery Electric Vehicles

E.3.1 State-level Rebates for Battery Electric Vehicles

Clinton & Steinberg (2019) analyze seven state-level direct vehicle rebates offered between 2011 and 2014. The programs each varied in the value of the incentive, the time they were in effect, and eligibility requirements. Using a fixed-effect specification, they estimate the effect of a \$1,000 increase in financial incentives to be a 7.8% increase in per capita BEV registrations. We translate this into an elasticity below.

Throughout this section, the “in-context” specification will mean the seven states (CA, HI, IL, MA, PA, TN, TX) from 2011 to 2014, which is the time and geography analyzed in the paper.

WTP The WTP for an expansion of BEV subsidies is the sum of the transfer (which we normalize to 1) plus the environmental and other market externalities discussed in Section 4. These broadly have the form of the elasticity (-2.931) times the societal willingness to pay for one additional dollar of spending on the BEV, V/p . We estimate V/p separately by focusing on the per-car externalities, V , and the consumer price, p . To measure consumer prices, the Office of Energy Efficiency and Renewable Energy (EERE) lists the manufacturer’s suggested

retail price (MSRP) for most BEV models. We compute a sales-weighted average using data from Kelley Blue Book’s Electrified Light-Vehicle Sales Report for Q4 2021, which includes estimates of 2020 sales of each BEV model. For models with multiple trims, EERE would often report a range of MSRPs, in which case we use the mean of the MSRP for that model. For 2020, we have an average MSRP of \$54,025. The total subsidy is the sum of the average federal and state-level subsidies in 2020, which we describe in more detail in the Fiscal Externalities section, and is \$8104.27. To compute the elasticity, we take the semi-elasticity reported in the paper 0.078, divide it by \$1,000, and then multiply it by the MSRP net of subsidies for the in-context specification, \$36,248. This gives us an elasticity of -2.931.

For the in-context specification, the sales data by model comes from the Transportation Research Center at Argonne National Laboratory and covers all models up until 2019 that are full-sized and capable of 60 mph. This gives us an average MSRP of \$46,006, which net of the average federal subsidy over 2011-2014 and the average subsidy among the fourteen states leaves us with a net MSRP of \$36,248.

Transfer We consider a \$1 increase in the BEV subsidy with no pass-through rate.

Global Environmental Externalities As described in the main text of the paper, the global environmental benefits are calculated as the difference in the WTP for avoiding the global pollutants from driving the counterfactual internal combustion engine (ICE) vehicle and the WTP for facing the global pollutants from electricity generation needed to drive a battery electric vehicle (BEV). We calculate the two WTPs as follows:

BEV Global Externalities To determine the damages from one electric vehicle, we need to know how many kilowatt-hours (kWh) of electricity the average BEV uses each year and the average vehicle miles traveled (VMT) driven. EERE reports the kWh of electricity needed for each BEV model to travel 100 miles. Combining this data with the sales data mentioned above, we calculate a sales-weighted average of the kWh used by a BEV per mile. For 2020, this average energy consumption is 0.293 kWh per mile, and for in-context, it is 0.326 kWh per mile.

For VMT we use the 2017 results from the Federal Highway Administration’s National Household Travel Survey (NHTS). They report VMT by vehicle age and vehicle type, where the eight types of vehicles include options such as “Automobile/Car/Station Wagon”, “Van (Mini/Cargo/Passenger)”, “SUV (Santa Fe, Tahoe, Jeep, etc.)”, “Pickup Truck”, and “Other Truck”. We use for our main specification the VMT reported in the Automobile/Car/Station Wagon category as is. Then, the VMT is multiplied by 0.6154 following Zhao et al. (2023)’s analysis finding that BEVs accumulate fewer annual miles than ICE vehicles: 7,165 versus 11,642.

Throughout our analysis, we assume the counterfactual ICE vehicles and BEVs have the same VMT, although our primary results are robust to reasonable variations in this assumption. This is because it is natural to assume that the difference in VMT found in Zhao et al. is due to selection in the types of drivers that purchase BEVs.¹⁶⁴ We assume a 17-year lifespan for

¹⁶⁴It would be straightforward to adjust our approach to allow for differential mileage driven, noting that one would need to distinguish between total reduction in miles and substitution to other ICE vehicles.

both ICE vehicles and BEVs, so we use the VMT numbers corresponding to a car’s lifespan within that range. The survey reports an average VMT of 12,245 miles for the first year, with the remaining values of VMT ranging between 5,885 (in the last year) and 13,078 (in the fifth year). For in-context, we impute age-state-level VMT data using the NHTS VMT by household state and vehicle type. We calculate a percent difference between the sample-weighted average VMT across all ages and each age’s VMT. Then, we assume this percent difference holds within states. Thus, we use the EV registration-weighted average of the nine sample states’ imputed VMTs for the in-context specification.

Now that we can calculate the energy consumption for a BEV in each year of its lifetime, we estimate the global damages from electricity consumption from the grid using information from AVERT and forecasts of the cleanliness of the grid from Jenkins & Mayfield (2023). Since the VMT changes for each year of the car’s lifetime, the global damages each year change as well. For exposition, the first year’s energy consumption will be 2425.525 kWh in 2020 and 2573.102 in-context, which leads to \$347.623 and \$218.546 respectively, of global damages from grid pollution. See a more detailed description of the grid externalities calculations in Appendix C.2. We then estimate the global damages for each year of the vehicle’s lifetime, which totals \$3183.512 in global damages in 2020 and \$2769.536 in context. Normalized by the net MSRP, we have \$0.069 for 2020 and \$0.076 for in-context.

Finally, we take this amount and multiply it by the elasticity of -2.931 and we subtract out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, which is 1.9% of the total amount (see Section 4). Written out this is $0.069 \cdot -2.931 \cdot (1 - 0.15 \cdot 0.2554 \cdot 0.5)$. Thus, our final values for the damages from using a BEV are -0.199 (2020) and -0.212 (in-context).

ICE Global Externalities For the counterfactual ICE vehicle, we use the same VMT as for the BEVs. We need the average fuel economy to obtain the final gas consumption in each year of the ICE vehicle’s life. Holland et al. (2016) report in their appendix a substitute gas vehicle for each popular BEV in 2014. Most of these are the ICE versions of the BEV (i.e., the gas-powered Ford Focus as the counterfactual for the BEV Ford Focus), but some are not. They test the reasonableness of their choices based on market research data from MaritzCX. Using the sales data in 2014 for the BEVs for which they find counterfactuals, we compute an average counterfactual MPG for 2014, which is 37.16 miles per gallon. This is higher than the average new car fuel economy in 2014 of 27.63 MPG. We then extrapolate to an average counterfactual MPG for other years by calculating the ratio of the counterfactual MPG to the car fleet’s average MPG in 2014 and assuming that the ratio of 1.34 holds for other years. This gives us a counterfactual MPG of 41.23 in 2020 and 39.37 in context. We also explore the robustness of our results to assuming BEVs displace an average fleet light duty vehicle, which slightly raises the MVPF to 1.664.

In the first year of the counterfactual ICE vehicle’s life, we estimate it consumed 326.599 gallons of gas in 2020 and 365.79 in context. Using the gas consumed by an ICE vehicle in each year of its life, we estimate the total damages from an ICE with the estimated pollutants as described in Appendix C.4. Again, we take this amount and normalize it by the net MSRP of a BEV to get 0.187 for 2020 and 0.130 in-context and multiply it by the elasticity to get 0.549 and 0.368 and we subtract out the portion of the benefits that will accrue to the US government via

increased GDP from avoiding carbon emissions, which is 1.9% of the total amount (see Section 4). Thus, a \$1 increase in the subsidy for a BEV leads to a reduction in damages from driving the counterfactual ICE of 0.538 in 2020 and 0.361 in context.

Upstream Battery Externalities We also include emissions from the production of the batteries used in electric vehicles. Battery production is a unique source of emissions from BEVs compared to ICE vehicles. Winjobi et al. (2022) note that “batteries in electric vehicles can account for one-third of their production greenhouse gas (GHG) emissions” and find in their analysis of different battery chemistries that lithium nickel manganese cobalt oxide batteries with equal proportions of nickel, manganese, and cobalt (NMC11) have life-cycle GHG emissions of 59.5 kg CO₂-eq/kWh. Using the average battery capacity for 2020 (73.004) and in-context (39.777), we have 4343.744 for 2020 and 2366.736 in-context. After converting the kg of emissions to tons, multiplying by the SCC, and normalizing by the MSRP, we have 0.018 for 2020 and 0.011 in context. Multiplying by the elasticity and pass-through rate and subtracting out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, we have 0.052 for 2020 and 0.030 for in-context.

The sum of the reduction in ICE emissions, increase in grid-related emissions, and upstream battery production emissions gives our final externality measure of 0.287 (2020) and 0.119 (in-context) for the global environmental externality for a \$1 mechanical increase in subsidies to BEVs.

Local Externalities Local externalities are estimated in nearly the same way as global externalities. We use the same energy and gas consumption values. The difference is the marginal damage per kWh of electricity or gallon of gasoline. We describe the estimation of these local damages in Appendices C.2 and C.4. We calculate -0.010 in 2020 and -0.019 in-context in local damages from BEVs and 0.016 in 2020 and 0.011 in-context from the counterfactual ICE. Taking the difference gives 0.006 and -0.008, respectively. After multiplying by the elasticity, we have 0.018 and -0.024.

Rebound The previous analysis assumes that the increased consumption of BEVs does not affect grid-wide electricity prices. We now consider some general equilibrium effects where the increased use of BEVs increases the price of electricity and thus decreases electricity consumption. We assume this effect does not lead to a secondary rebound in using ICE vehicles because while we allow for local production of electricity with upward-sloping supply curves, we assume a flat global supply curve for gasoline so that there are few if any rebound effects in the gasoline market. The rebound effect is calculated as $\frac{1}{1-\varepsilon_D/\varepsilon_S}$ where ε_D is the demand elasticity for electricity, which is -0.19 and comes from DOI (2021) and ε_S is the supply elasticity for electricity, which is 0.78 (see Appendix D for calculation). With this rebound effect of about 20% multiplied by the previously calculated global and local damages from electricity consumption for BEVs, we now have a decrease in damages of 0.045 in 2020 and 0.052 in-context from less electricity consumption.

Learning-by-Doing Our model of learning-by-doing is described in Appendix B. Here, we describe any necessary preliminary calculations as well as the data sources for the model inputs. There are nine inputs into the model: the demand elasticity for BEVs, the discount

rate, the learning rate, the fraction of a BEV's price that is from non-battery components (we refer to this as the fixed cost ratio), the marginal sales in a given year, the cumulative sales up until the year of interest, the net MSRP, the environmental damage per EV, and the social cost of carbon (SCC).

The demand elasticity for this policy of -2.931 is calculated as described above. Our baseline discount rate is 2%. The learning rate for batteries of 0.421 comes from Way et al. (2022). We show in Appendix B how we adjust for the fact that batteries comprise only a fraction of the total cost of the car. Environmental damages are as above for the global and local externalities but converted back to the per-car level. However, we allow for the environmental externality to vary over time to account for the SCC increasing over time. Our method of extrapolating to future SCC values is described in Section 4. The net MSRP is as described at the top of the WTP section. The last input is the SCC, which we use a baseline value of 193 for 2020 and allow to vary over time following projections of a rising SCC from the Environmental Protection Agency's recent guidance regarding the social cost of carbon at a 2% discount rate.

This leaves the fixed cost ratio and marginal and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries, as well as the cost per kWh, come from Ziegler & Trancik (2021). They report a representative series of the price of all types of lithium-ion cells and one for the market size of all types of lithium-ion cells measured in energy capacity. However, the price data only goes until 2018, and the sales data goes until 2016, so we append price data from the Department of Energy and sales data from the IEA. For 2020, the marginal sales are 167700 MWh of batteries, and the cumulative sales are 917708 starting from 1991. For in-context, the average marginal sales are 29,680 and the average cumulative sales are 147359. The price per kWh in 2020 is \$181.978 and over 2011-14 on average is \$248.777. To calculate the fixed cost ratio, we take the price per kWh and multiply it by the sales-weighted average battery capacity for EVs in 2020, 73.004 kWh (battery capacity for each model comes from Edmunds), or 2011-14, 39.777 kWh. We then divide this by the MSRP (\$54,025 or \$46,006) to get the proportion of the cost due to the battery of 0.246 for 2020 and 0.187 for in-context. One minus that proportion gives us our fixed cost ratio of 0.754 for 2020 and 0.813 for in-context. Recall that for our learning-by-doing model, we only think the learning is affecting the battery price and not the price of other parts of the vehicle, so any price and environmental benefits will be a result of only those battery prices coming down. Following the steps outlined in Appendix B, we obtain a dynamic environmental component for 2020 (in-context) of 0.119 (-0.323) and a dynamic price component of 0.564 (0.403).

Profits Lastly, for WTP, we estimate the gasoline producers' and utilities' WTP for the subsidy. Since the gasoline market has higher markups than the average market and electricity markets are heavily regulated to have fixed levels of markups, we believe this is an important component to consider. Using the previously estimated increase in electricity consumption and decrease in gas consumption, we calculate the producers' WTP as the markup multiplied by the change in electricity or gas, normalized by the net MSRP, and then multiplied by the elasticity. For each year of the car's lifetime, the annual profits will be discounted. For gasoline, this is 4,642.3/3,199.9 gallons of gasoline multiplied by 0.613/0.933 markup (discounted each year), normalized by the net MSRP, multiplied by the elasticity, and finally we subtract out 21% of the surplus due to the effective corporate tax rate to get 0.125/0.142. For utilities, we have

34,476/36,574 kWh of electricity multiplied by 0.011/0.022 markup, normalized, and multiplied by the same elasticity to get 0.017/0.044.

With the various components, we can now calculate the total WTP of 1.925 for 2020 and 1.130 for in-context.

Cost The cost of \$1 mechanical increase in BEV subsidies is equal to the \$1 plus the fiscal externalities induced by the demand response to the subsidies. These include additional costs to state and federal subsidies, reductions in gas tax revenue, changes in profits tax collected, and the climate fiscal externality from changes in CO_2 emissions. We discuss each in turn.

State and Federal FEs In 2020, and in context, there are some existing federal and state subsidies for BEVs; an additional dollar of spending will spur further behavioral effects that will increase the spending from the preexisting programs. Federally, in 2020, most BEVs no longer qualified for the \$7,500 subsidy from the Qualified Plug-In Electric Drive Motor Vehicle Credit. Only eight models qualified, and some of those for less than the full amount. EERE reports the subsidy each model qualified for and when it stopped qualifying for the maximum amount and subsequent smaller amounts. Similar to how we estimate the sales-weighted average for other parameters, we estimate a sales-weighted average federal subsidy of \$7500.00 in 2020 and \$7107.55 in-context. Normalizing that value by the net MSRP gives us \$0.163 in 2020 and \$0.196 in context. Finally, multiplying by the elasticity gives us the federal fiscal externality of 0.479 for 2020 and 0.554 for in-context.

For state subsidies, we use the Alternative Fuels Data Center's (AFDC) database on incentives and laws related to alternative fuels and advanced vehicles. Nine states in 2020 have subsidies for BEVs, with varying levels of subsidy size, MSRP eligibility rules, and income eligibility rules. For example, in Oregon, there was a \$7,500 subsidy that applied if the battery was greater than 10 kWh, the MSRP was less than \$50,000, and the income for a household of one was below \$54,360. If subsidies differ by driver's income, we scale the subsidy based on the approximate proportion of EV drivers within that income constraint. Muehlegger & Rapson (2019) report the proportion of EV drivers within income categories and find that 14% have household incomes between 0 and \$49,000 and 30% have incomes between 50 and \$99,000. This implies (assuming a uniform distribution of EV purchasers within each bucket) that 16.7% of Oregon's EV purchasers were eligible for the subsidy. In 2020, only 63% of BEVs had an MSRP of less than \$50,000. Thus, we can estimate that the percent of EV purchasers who received the subsidy (assuming independence) is $0.63 \cdot 0.167 = 0.1056$. The average subsidy received in Oregon was \$792.33. We repeat these steps for the other eight states, and then we use the AFDC's data on how many EV registrations occur in each state to compute a final EV registration-weighted average state subsidy of \$604.27. This value normalized by the net MSRP and multiplied by the elasticity gives us the state fiscal externality of 0.039.

Since the in-context MVPF is looking at a specific state's subsidy, we take the subsidy amount to be \$2,650.3 as reported in Clinton & Steinberg (2019). When normalized and multiplied by the elasticity, this gives us a state fiscal externality of 0.207.

Gas Tax FE We also calculate a gasoline tax fiscal externality using the average state and federal gas tax rates as described in Appendix E.10. The FE is the elasticity multiplied by the tax rate of 0.465/0.271 multiplied by the decrease in gas consumption from ICE vehicles of 4,642.3/3,199.9 and normalized by the net MSRP to get 0.120/0.056.

Profits Tax FE Similarly to gasoline taxes, we have an average combined revenue rate of 0.006/0.012 that accounts for profits to publicly-owned utilities and corporate taxes on privately owned utilities as described in Appendix C.2.3. The FE is calculated in the same way as the gas tax FE to get 0.009/0.024.

Climate FE Finally, the climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Our baseline social cost of carbon comes from EPA (2023*c*) based 50% on GDP and estimates that 15% of that avoided GDP loss would flow to the US. Since the average tax rate in the US is 25.54%, we estimate the climate FE as $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the static global environmental externality and the dynamic global environmental externality that results from learning by doing. Taking those two pieces and multiplying their sum by 1.9% gives us -0.011/0.002.

Thus, our final cost is 1.650/1.833, which gives us our MVPFs of 1.167 and 0.617.

E.3.2 Qualified Plug-In Electric Drive Motor Vehicle Credit

Li et al. (2017) studied the Qualified Plug-In Electric Drive Motor Vehicle Credit (PEDVC). This is a credit for electric vehicles purchased beginning in 2009. Li et al. (2017) simulate the effect of the PEDVC on battery electric vehicle (BEV) sales from 2011-13 using a model of indirect network effects between BEV sales and the availability of public charging stations. They find that 40.4% of the total BEV sales during the three years were a result of the subsidy program. We translate this into an elasticity below. Throughout this section, the “in-context” specification will mean the US from 2011 to 2013, which is the time and geography analyzed in the paper.

WTP The WTP for an expansion of BEV subsidies is the sum of the transfer (which we normalize to 1) plus the environmental and other market externalities discussed in Section 4. These broadly have the form of the elasticity (-2.611) times the societal willingness to pay for one additional dollar of spending on the BEV, V/p . We estimate V/p separately by focusing on the per-car externalities, V , and the relevant consumer price, p . To measure consumer prices, the Office of Energy Efficiency and Renewable Energy (EERE) lists the manufacturer’s suggested retail price (MSRP) for most BEV models. We compute a sales-weighted average using data from Kelley Blue Book’s Electrified Light-Vehicle Sales Report for Q4 2021, which includes estimates of 2020 sales of each BEV model. For models with multiple trims, EERE would often report a range of MSRPs, in which case we use the mean of the MSRP for that model. For 2020, we have an average MSRP of \$54,025. The total subsidy is the sum of the average federal and state-level subsidies in 2020, which we describe in more detail in the Fiscal Externalities section, and is \$8104.27.

For the in-context specification, the sales data by model comes from the Transportation Research Center at Argonne National Laboratory and covers all models up until 2019 that are full-sized and capable of 60 mph. This gives us an average MSRP of \$47,436, which net of the average state subsidy over 2011-2013 leaves us with a net MSRP of \$39,269.

We use in-context MSRPs and subsidy amounts to compute the elasticity. We take the BEV sales increase reported in the paper, 0.404, and divide it by the percent change in price of a BEV that corresponds to the subsidy, which is \$6592.25 (the average subsidy from the PEDVC as reported in the paper) divided by the average net MSRP over 2011-13, \$42,565. This gives us $0.404 / -0.144$ and our final elasticity of -2.611.

Transfer We consider a \$1 increase in the BEV subsidy with no pass-through rate.

Global Environmental Externalities As described in the main text of the paper, the global environmental benefits are calculated as the difference in the WTP for avoiding the global pollutants from driving the counterfactual internal combustion engine (ICE) vehicle and the WTP for facing the global pollutants from electricity generation needed to drive a battery electric vehicle (BEV). We calculate the two WTPs as follows: **BEV Global Externalities** To determine the damages from one electric vehicle, we need to know how many kilowatt-hours (kWh) of electricity the average BEV uses each year and the average vehicle miles traveled (VMT). EERE reports the kWh of electricity needed for each BEV model to travel 100 miles. Combining this data with the sales data mentioned above, we calculate a sales-weighted average of the kWh used by a BEV per mile. For 2020, this average energy consumption is 0.293 kWh per mile, and for in-context, it is 0.331 kWh per mile.

For VMT we use the 2017 results from the Federal Highway Administration’s National Household Travel Survey (NHTS). They report VMT by vehicle age and vehicle type, where the eight types of vehicles include options such as “Automobile/Car/Station Wagon”, “Van (Mini/Cargo/Passenger)”, “SUV (Santa Fe, Tahoe, Jeep, etc.)”, “Pickup Truck”, and “Other Truck”. We use for our main specification the VMT reported in the Automobile/Car/Station Wagon category. Then, the VMT is multiplied by 0.6154 following Zhao et al. (2023)’s analysis finding that BEVs accumulate fewer annual miles than ICE vehicles: 7,165 versus 11,642.

Throughout our analysis, we assume the counterfactual ICE vehicles and BEVs have the same VMT, although our primary results are robust to reasonable variations in this assumption. Then, the VMT is multiplied by 0.6154 following Zhao et al. (2023) analysis found that BEVs accumulate fewer annual miles than ICE vehicles: 7,165 versus 11,642.¹⁶⁵ We assume a 17-year lifespan for both ICE vehicles and BEVs, so we use the VMT numbers corresponding to each year of a car’s lifespan within that range. The survey reports an average VMT of 12,245 miles for the first year, with the remaining values of VMT ranging between 5,885 (in the last year) and 13,078 (in the fifth year). This input does not change between in-context and 2020 because the NHTS is only run every 5-8 years.

¹⁶⁵It would be straightforward to adjust our approach to allow for differential mileage driven, noting that one would need to distinguish between total reduction in miles and substitution to other ICE vehicles.

Now that we can calculate the energy consumption for a BEV in each year of its lifetime, we estimate the damages from electricity consumption from the grid using information from AVERT and forecasts of the cleanliness of the grid from Jenkins & Mayfield (2023). Since the VMT changes for each year of the car’s lifetime, the damages change as well. For exposition, the first year’s energy consumption will be 2425.525 kWh in 2020 and 2747.203 in-context which leads to 347.623 and 303.581 respectively, of damages from grid pollution. See a more detailed description of the grid externalities calculations in Appendix C.2. We then estimate the damages for each year of the vehicle’s lifetime, which totals 3183.512 in damages in 2020 and 3993.640 in context. Normalized by the net MSRP, we have 0.069 for 2020 and 0.102 for in-context.

Finally, we take this amount and multiply it by the elasticity of -2.611 and we subtract out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, which is 1.9% of the total amount (see Section 4). Written out this is $0.069 \cdot -2.611 \cdot (1 - 0.15 \cdot 0.3)$. Thus, our final values for the damages from using a BEV are -0.178 (2020) and -0.209 (in-context).

ICE Global Externalities For the counterfactual ICE vehicle, we use the same VMT as for the BEVs. We need the average fuel economy to obtain the final gas consumption in each year of the ICE vehicle’s life. Holland et al. (2016) report in their appendix a substitute gas vehicle for each popular BEV in 2014. Most of these are the ICE versions of the BEV (i.e., the gas-powered Ford Focus as the counterfactual for the BEV Ford Focus), but some are not. They test the reasonableness of their choices based on market research data from MaritzCX. Using the sales data in 2014 for the BEVs for which they find counterfactuals, we compute an average counterfactual MPG for 2014, which is 37.16 miles per gallon. This is higher than the average new car fuel economy in 2014 of 27.63 MPG. We then extrapolate to an average counterfactual MPG for other years by calculating the ratio of the counterfactual MPG to the car fleet’s average MPG in 2014 and assuming that the ratio of 1.34 holds for other years. This gives us a counterfactual MPG of 41.233 in 2020 and 36.556 in context. We also explore the robustness of our results to assuming EVs displace an average fleet light duty vehicle, which slightly raises the MVPF to 1.559.

In the first year of the counterfactual ICE vehicle’s life, we estimate it consumed 326.599 gallons of gas in 2020 and 368.71 in context. Now, with the gas consumed by an ICE vehicle in each year of its life, we estimate the total damages from an ICE with the estimated pollutants as described in Appendix C.4. Again, we take this amount and normalize it by the net MSRP of a BEV to get 0.187 for 2020 and 0.118 in-context and multiply it by the elasticity to get 0.489 and 0.309 and we subtract out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, which is 1.9% of the total amount (see Section 4). Thus, \$1 of spending on a BEV generates \$0.187 of environmental savings from reduced ICE emissions. Multiplying by the elasticity, this suggests that a \$1 increase in the subsidy for a BEV leads to a reduction in damages from driving the counterfactual ICE is 0.480 in 2020 and 0.303 in context.

Upstream Battery Externalities We also include emissions from the production of the batteries used in electric vehicles. Winjobi et al. (2022) note that “batteries in electric vehicles can account for one-third of their production greenhouse gas (GHG) emissions” and find in

their analysis of different battery chemistries that lithium nickel manganese cobalt oxide batteries with equal proportions of nickel, manganese, and cobalt (NMC11) have life-cycle GHG emissions of 59.5 kg CO₂-eq/kWh. Using the average battery capacity for 2020 (73.004) and in-context (40.008), we have 4343.744 for 2020 and 2380.470 in-context. After converting the kg of emissions to tons, multiplying by the SCC, and normalizing by the MSRP, we have 0.018 for 2020 and 0.010 in context. Multiplying by the elasticity and pass-through rate and subtracting out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, we have 0.047 for 2020 and 0.025 for in-context.

The sum of the reduction in ICE emissions, increase in grid-related emissions, and upstream battery production emissions gives our final externality measure of 0.255 (2020) and 0.068 (in-context) for the global environmental externality for a \$1 mechanical increase in subsidies to BEVs.

Local Externalities Local externalities are estimated in nearly the same way as global externalities. We use the same energy and gas consumption values. The difference is the marginal damage per kWh of electricity or gallon of gasoline. We describe the estimation of these local damages in Appendices C.2 and C.4. We calculate local damages from increased grid usage from BEVs of -0.010 in 2020 and -0.029 in context. We calculate savings from reduced gasoline consumption of 0.016 in 2020 and 0.010 in context from the counterfactual ICE. Taking the difference between the grid usage and gas consumption yields total benefits of 0.006 and -0.019, respectively. After multiplying by the price elasticity, it suggests \$1 of mechanical spending on the subsidy delivers local environmental benefits of \$0.016 in 2020 and \$-0.034 in context.

Rebound The previous analysis assumes that the increased consumption of BEVs does not affect grid-wide electricity prices. We now consider some general equilibrium effects where the increased use of BEVs increases the price of electricity and thus decreases electricity consumption. We assume this effect does not lead to a secondary rebound in using ICE vehicles because while we allow for local production of electricity with upward-sloping supply curves, we assume a flat global supply curve for gasoline so that there are few if any rebound effects in the gasoline market. The rebound effect is calculated as $\frac{1}{1-\varepsilon_D/\varepsilon_S}$ where ε_D is the demand elasticity for electricity, which is -0.19 and comes from DOI (2021) and ε_S is the supply elasticity for electricity, which is 0.78 (see Appendix D for calculation). With this rebound effect of about 20% multiplied by the previously calculated global and local damages from electricity consumption for BEVs, we now have a decrease in damages of 0.040 in 2020 and 0.053 in-context from less electricity consumption.

Learning-by-Doing Our model of learning by doing is described in Appendix B. Here, we describe any necessary preliminary calculations as well as the data sources for the model inputs. There are nine inputs into the model: the demand elasticity for BEVs, the discount rate, the learning rate, the fraction of a BEV's price that is from non-battery components (we refer to this as the fixed cost ratio), the marginal sales in a given year, the cumulative sales up until the year of interest, the net MSRP, the environmental damage per EV, and the social cost of carbon (SCC).

The demand elasticity for this policy of -2.611 is calculated as described above. Our baseline discount rate is 2%. The learning rate for batteries of 0.421 comes from Way et al. (2022). We show in Appendix B how we adjust for the fact that batteries comprise only a fraction of the total cost of the car. Environmental damages are as above for the global and local externalities but converted back to the per-car level. However, we allow for the environmental externality to vary over time to account for the SCC increasing over time. Our method of extrapolating to future SCC values is described in Section 3.2. The net MSRP is as described at the top of the WTP section. The last input is the SCC, which we use a baseline value of 193 for 2020 and allow to vary over time following projections of a rising SCC from the Environmental Protection Agency's recent guidance regarding the social cost of carbon at a 2% discount rate.

This leaves the fixed cost ratio and marginal and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries, as well as the cost per kWh, come from Ziegler & Trancik (2021). They report a representative series of the price of all types of lithium-ion cells and one for the market size of all types of lithium-ion cells measured in energy capacity. However, the price data only goes until 2018, and the sales data goes until 2016, so we append price data from the Department of Energy and sales data from the IEA. For 2020, the marginal sales are 167700 MWh of batteries, and the cumulative sales are 917708 starting from 1991. For in-context, the average marginal sales are 29,680 and the average cumulative sales are 147359. The price per kWh in 2020 is \$181.978 and over 2011-13 on average is \$259.153. To calculate the fixed cost ratio, we take the price per kWh and multiply it by the sales-weighted average battery capacity for EVs in 2020, 73.004 kWh (battery capacity for each model comes from Edmunds), or 2011-13, 40.008 kWh. We then divide this by the MSRP (\$54,025 or \$47,436) to get the proportion of the cost due to the battery of 0.246 for 2020 and 0.190 for in-context. One minus that proportion gives us our fixed cost ratio of 0.754 for 2020 and 0.810 for in-context. Recall that for our learning-by-doing model, we only think the learning is affecting the battery price and not the price of other parts of the vehicle, so any price and environmental benefits will be a result of only those battery prices coming down. Following the steps outlined in Appendix B, we obtain a dynamic environmental component for 2020 (in-context) of 0.090 (0.050) and a dynamic price component of 0.482 (0.356).

Profits Lastly, for WTP, we estimate the gasoline producers' and utilities' WTP for the subsidy. Since the gasoline market has higher markups than the average market and electricity markets are heavily regulated to have fixed levels of markups, we believe this is an important component to consider. Using the previously estimated increase in electricity consumption and decrease in gas consumption, we calculate the producers' WTP as the markup multiplied by the change in electricity or gas, normalized by the net MSRP, and then multiplied by the elasticity. For each year of the car's lifetime, the annual profits will be discounted. For gasoline, this is 4,642.3/3,225.4 gallons of gasoline multiplied by 0.613/0.921 markup (discounted each year), normalized by the net MSRP, multiplied by the elasticity, and finally we subtract out 21% of the surplus due to the effective corporate tax rate to get 0.111/0.128. For utilities, we have 34,476/39,049 kWh of electricity multiplied by 0.011/0.010 markup, normalized, and multiplied by the same elasticity to get 0.015/0.018.

With the various components, we can now calculate the total WTP of 1.788 for 2020 and 1.383 for in-context.

Cost The cost of \$1 mechanical increase in BEV subsidies is equal to the \$1 plus the fiscal externalities induced by the demand response to the subsidies. These include additional costs to state and federal subsidies, reductions in gas tax revenue, changes in profits tax collected, and the climate fiscal externality from changes in CO_2 emissions. We discuss each in turn.

State and Federal FEs In 2020, and in context, there are some existing federal and state subsidies for BEVs; an additional dollar of spending will spur further behavioral effects that will increase the spending from the preexisting programs. Federally, in 2020, most BEVs no longer qualified for the \$7,500 subsidy from the Qualified Plug-In Electric Drive Motor Vehicle Credit. Only eight models qualified, and some of those for less than the full amount. EERE reports the subsidy each model qualified for and when it stopped qualifying for the maximum amount and subsequent smaller amounts. Similar to how we estimate the sales-weighted average for other parameters, we estimated a sales-weighted average federal subsidy of \$7,500 in 2020. For in-context, we use the average federal subsidy reported by the paper, which is \$6,592.2 in-context. Normalizing that value by the net MSRP gives us \$0.163 in 2020 and \$0.168 in context. Finally, multiplying by the elasticity gives us the federal fiscal externality of 0.426 for 2020 and 0.438 for in-context.

For state subsidies, we use the Alternative Fuels Data Center’s (AFDC) database on incentives and laws related to alternative fuels and advanced vehicles. Nine states in 2020 have subsidies for BEVs, with varying levels of subsidy size, MSRP eligibility rules, and income eligibility rules. For example, in Oregon, there was a \$7,500 subsidy that applied if the battery was greater than 10 kWh, the MSRP was less than \$50,000, and the income for a household of one was below \$54,360. If subsidies differ by driver’s income, we scale the subsidy based on the approximate proportion of EV drivers within that income constraint. Muehlegger & Rapson (2019) report the proportion of EV drivers within income categories and find that 14% have household incomes between 0 and \$49,000 and 30% have incomes between 50 and \$99,000. This implies (assuming a uniform distribution of EV purchasers within each bucket) that 16.7% of Oregon’s EV purchasers were eligible for the subsidy. In 2020, only 63% of BEVs had an MSRP of less than \$50,000. Thus, we can estimate that the percent of EV purchasers who received the subsidy (assuming independence) is $0.63 \cdot 0.167 = 0.1056$. The average subsidy received in Oregon was \$792.33. We repeat these steps for the other eight states, and then we use the AFDC’s data on how many EV registrations occur in each state to compute a final EV registration-weighted average state subsidy of \$7,500. This value normalized by the net MSRP and multiplied by the elasticity gives us the state fiscal externality of 0.426.

For the in-context MVPF, we take the subsidy amount reported in Table 1 of Li et al. (2017), which is \$1,575. When normalized and multiplied by the elasticity, this gives us a state fiscal externality of 0.105.

Gas Tax FE We also calculate a gasoline tax fiscal externality using the average state and federal gas tax rates as described in Appendix E.10. The FE is the elasticity multiplied by the tax rate of $0.465/0.398$ multiplied by the decrease in gas consumption from ICE vehicles of $4,642.3/3,225.4$ and normalized by the net MSRP to get $0.107/0.074$.

Profits Tax FE Similarly to gasoline taxes, we have an average combined revenue rate of 0.006/0.005 that accounts for profits to publicly-owned utilities and corporate taxes on privately owned utilities as described in Appendix C.2.3. The FE is calculated in the same way as the gas tax FE to get 0.008/0.010.

Climate FE Finally, the climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Our baseline social cost of carbon comes from EPA (2023*c*) based on 50% GDP and estimates that 15% of that avoided GDP loss would flow to the US. Since the average tax rate in the US is 25.54%, we estimate the climate FE as $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the static global environmental externality and the dynamic global environmental externality that results from learning by doing. Taking those two pieces and multiplying their sum by 1.9% gives us -0.009/-0.004.

Thus, our final cost is 1.580/1.637, which gives us our MVPFs of 1.132 and 0.844.

E.3.3 Enhanced Fleet Modernization Program

Muehlegger & Rapson (2022) study the Enhanced Fleet Modernization Program. This is a voluntary vehicle scrappage program that promotes the purchase of new battery electric vehicles (BEVs) for California residents who have low incomes. The program was evaluated by exploiting exogenous variation in large EV subsidies within “disadvantaged” zip codes across pilot and control regions. Results suggested a consumer price elasticity of EV demand of -2.1 and an average subsidy pass-through rate of 85 percent.

Throughout this section, the “in-context” specification will mean California from 2015 to 2018, which is the time and geography analyzed in the paper.

WTP The WTP for an expansion of BEV subsidies is the sum of the transfer (which we normalize to 1) plus the environmental and other market externalities discussed in Section 4. These broadly have the form of the product of the elasticity (-2.1) times the pass-thru rate of subsidies to prices, times the societal willingness to pay for one additional dollar of spending on the BEV, V/p . We estimate V/p separately by focusing on the per-car externalities, V , and the relevant consumer price, p . To measure consumer prices, the Office of Energy Efficiency and Renewable Energy (EERE) lists the manufacturer’s suggested retail price (MSRP) for most BEV models. We compute a sales-weighted average using data from Kelley Blue Book’s Electrified Light-Vehicle Sales Report for Q4 2021, which includes estimates of 2020 sales of each BEV model. For models with multiple trims, EERE would often report a range of MSRPs, in which case we use the mean of the MSRP for that model. For 2020, we have an average MSRP of \$54,025. The total subsidy is the sum of the average federal and state-level subsidies in 2020, which we describe in more detail in the Fiscal Externalities section, and is \$8104.27.

For the in-context specification, the sales data by model comes from the Transportation Research Center at Argonne National Laboratory and covers all models up until 2019 that are full-sized and capable of 60 mph. This gives us an average MSRP of \$61,678, which net of the average federal subsidy over 2015-2018 and the subsidy specific to this EFMP program leaves us with a net MSRP of \$45,656.

Transfer We consider a \$1 increase in the BEV subsidy, where 85% flows to the consumers and 15% to the dealers as shown in Table 4 of Muehlegger & Rapson (2022) implies.

Global Environmental Externalities As described in the main text of the paper, the global environmental benefits are calculated as the difference in the WTP for carbon dioxide emissions with an internal combustion engine (ICE) vehicle and a battery electric vehicle (BEV).

We calculate this difference as follows: **BEV Global Externalities** To determine the damages from one electric vehicle, we need to know how many kilowatt-hours (kWh) of electricity the average BEV uses each year and the average vehicle miles traveled (VMT) driven. EERE reports the kWh of electricity needed for each BEV model to travel 100 miles. Combining this data with the sales data mentioned above, we calculate a sales-weighted average of the kWh used by a BEV per mile. For 2020, this average energy consumption is 0.293 kWh per mile, and for in-context is 0.314 kWh per mile.

For VMT we use the 2017 results from the Federal Highway Administration’s National Household Travel Survey (NHTS). They report VMT by vehicle age and vehicle type, where the eight types of vehicles include options such as “Automobile/Car/Station Wagon”, “Van (Mini/Cargo/Passenger)”, “SUV (Santa Fe, Tahoe, Jeep, etc.)”, “Pickup Truck”, and “Other Truck”. We use for our main specification the VMT reported in the Automobile/Car/Station Wagon category as is. Then, the VMT is multiplied by 0.6154 following Zhao et al. (2023)’s analysis finding that BEVs accumulate fewer annual miles than ICE vehicles: 7,165 versus 11,642.

Throughout our analysis, we assume the counterfactual ICE vehicles and BEVs have the same VMT, although our primary results are robust to reasonable variations in this assumption. Then, the VMT is multiplied by 0.6154 following Zhao et al. (2023) analysis found that BEVs accumulate fewer annual miles than ICE vehicles: 7,165 versus 11,642.¹⁶⁶ We assume a 17-year lifespan for both ICE vehicles and BEVs, so we use the VMT numbers corresponding to each year of a car’s lifespan within that range. The survey reports an average VMT of 12,245 miles for the first year, with the remaining values of VMT ranging between 5,885 (in the last year) and 13,078 (in the fifth year). For in-context, we impute age-state-level VMT data using the NHTS VMT by household state and vehicle type. We calculate a percent difference between the sample-weighted average VMT across all ages and each age’s VMT. Then, we assume this percent difference holds within states. Thus, we use the California-specific imputed VMT for the in-context specification.

Once we calculate the energy consumption for a BEV in each year of its lifetime, we estimate the damages from electricity consumption from the grid using information from AVERT and forecasts of the cleanliness of the grid from Jenkins & Mayfield (2023). Since the VMT changes for each year of the car’s lifetime, the damages change as well. For exposition, the first year’s energy consumption will be 2425.525 kWh in 2020 and 2424.229 in-context which leads to 347.623 and 200.897 respectively, of damages from grid pollution. See a more detailed de-

¹⁶⁶It would be straightforward to adjust our approach to allow for differential mileage driven, noting that one would need to distinguish between total reduction in miles and substitution to other ICE vehicles.

scription of the grid externalities calculations in Appendix C.2. We then estimate the damages for each year of the vehicle’s lifetime, which totals 3183.512 in damages in 2020 and 2136.125 in context. Normalized by the net MSRP, we have 0.069 for 2020 and 0.047 for in-context.

Finally, we take this amount and multiply it by the elasticity of -2.1 and the pass-through rate of 0.85 and we subtract out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, which is 1.9% of the total amount (see Section 4). Written out for 2020 this is $0.069 \cdot -2.1 \cdot 0.85 \cdot (1 - 0.15 \cdot 0.2554 \cdot 0.5)$. Thus, our final values for the damages from using a BEV are -0.121 (2020) and -0.082 (in-context).

ICE Global Externalities For the counterfactual ICE vehicle, we use the same VMT as for the BEVs. We need the average fuel economy to obtain the final gas consumption in each year of the ICE vehicle’s life. Holland et al. (2016) report in their appendix a substitute gas vehicle for each popular BEV in 2014. Most of these are the ICE versions of the BEV (i.e., the gas-powered Ford Focus as the counterfactual for the BEV Ford Focus), but some are not. They test the reasonableness of their choices based on market research data from MaritzCX. Using the sales data in 2014 for the BEVs for which they find counterfactuals, we compute an average counterfactual MPG for 2014, which is 37.16 miles per gallon. This is higher than the average new car fuel economy in 2014 of 27.63 MPG. We then extrapolate to an average counterfactual MPG for other years by calculating the ratio of the counterfactual MPG to the car fleet’s average MPG in 2014 and assuming that the ratio of 1.34 holds for other years. This gives us a counterfactual MPG of 41.23 in 2020 and 39.37 in context. We also explore the robustness of our results to assuming EVs displace an average fleet light duty vehicle, which slightly raises the MVPF to 1.350.

In the first year of the counterfactual ICE vehicle’s life, we estimate it consumed 326.599 gallons of gas in 2020 and 342.247 in context. Now, with the gas consumed by an ICE vehicle in each year of its life, we estimate the total damages from an ICE with the estimated pollutants as described in Appendix C.4. Again, we take this amount and normalize it by the net MSRP of a BEV to get 0.187 for 2020 and 0.109 in-context and multiply it by the elasticity and the pass-through rate to get 0.334 and 0.195 and we subtract out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, which is 1.9% of the total amount (see Appendix 4). Thus, \$1 of spending on an EV generates \$0.187 of environmental savings from reduced ICE emissions. Multiplying by the elasticity and pass-through rate, this suggests that a \$1 increase in the subsidy for an EV leads to a reduction in damages from driving the counterfactual ICE of 0.328 in 2020 and 0.191 in context.

Upstream Battery Externalities We also include emissions from the production of the batteries used in electric vehicles. Winjobi et al. (2022) note that “batteries in electric vehicles can account for one-third of their production greenhouse gas (GHG) emissions” and find in their analysis of different battery chemistries that lithium nickel manganese cobalt oxide batteries with equal proportions of nickel, manganese, and cobalt (NMC11) have life-cycle GHG emissions of 59.5 kg CO₂-eq/kWh. Using the average battery capacity for 2020 (73.004) and in-context (63.497), we have 4343.744 for 2020 and 3778.063 in-context. After converting the kg of emissions to tons, multiplying by the SCC, and normalizing by the MSRP, we have 0.018 for 2020 and 0.015 in context. Multiplying by the elasticity and pass-through rate and subtracting out the portion of the benefits that will accrue to the US government via increased GDP from

avoiding carbon emissions, we have 0.032 for 2020 and 0.026 for in-context.

The sum of the reduction in ICE emissions, increase in grid-related emissions, and upstream battery production emissions gives our final externality measure of 0.175 (2020) and 0.083 (in-context) for the global environmental externality for a \$1 mechanical increase in subsidies to BEVs.

Local Externalities Local externalities are estimated in nearly the same way as global externalities. We use the same energy and gas consumption values. The difference is the marginal damage per kWh of electricity or gallon of gasoline. We describe the estimation of these local damages in Appendix C.4. We calculate local damages from increased grid usage from BEVs of \$-0.010 in 2020 and \$-0.003 in context. We calculate savings from reduced gasoline consumption of \$0.016 in 2020 and \$0.009 in-context from the counterfactual ICE. Taking the difference between the grid usage and gas consumption yields total benefits of \$0.006 and \$0.006, respectively. After multiplying by the price elasticity and pass-through rate, it suggests \$1 of mechanical spending on the subsidy delivers local environmental benefits of \$0.011 in 2020 and \$0.010 in context.

Rebound The previous analysis assumes that the increased consumption of BEVs does not affect grid-wide electricity prices. We now consider some general equilibrium effects where the increased use of BEVs increases the price of electricity and thus decreases electricity consumption. We assume this effect does not lead to a secondary rebound in using ICE vehicles because while we allow for local production of electricity with upward-sloping supply curves, we assume a flat global supply curve for gasoline so that there are few if any rebound effects in the gasoline market. The rebound effect is calculated as $\frac{1}{1-\varepsilon_D/\varepsilon_S}$ where ε_D is the demand elasticity for electricity, which is -0.19 and comes from DOI (2021) and ε_S is the supply elasticity for electricity, which is 0.78 (see Appendix D for calculation). With this rebound effect of about 20% multiplied by the previously calculated global and local damages from electricity consumption for BEVs, we now have a decrease in damages of 0.027 in 2020 and 0.017 in-context from less electricity consumption.

Learning-by-Doing Our model of learning-by-doing is described in Appendix B. Here, we describe any necessary preliminary calculations as well as the data sources for the model inputs. There are nine inputs into the model: the demand elasticity for BEVs, the discount rate, the learning rate, the fraction of a BEV's price that is from non-battery components (we refer to this as the fixed cost ratio), the marginal sales in a given year, the cumulative sales up until the year of interest, the net MSRP, the environmental damage per EV, and the social cost of carbon (SCC).

The demand elasticity for this policy of -2.1 comes directly from the paper. Our baseline discount rate is 2%. The learning rate for batteries of 0.421 comes from Way et al. (2022). We show in Appendix B how we adjust for the fact that batteries comprise only a fraction of the total cost of the car. Environmental damages are as above for the global and local externalities but converted back to the per-car level. However, we allow for the environmental externality to vary over time to account for the SCC increasing over time. Our method of extrapolating to future SCC values is described in Section 3.2. The net MSRP is as described at the top of the

WTP section. The last input is the SCC, for which we use our baseline specification that has a value of 193 for 2020 and allow it to vary over time following projections of a rising SCC from the Environmental Protection Agency's recent guidance regarding the social cost of carbon at a 2% discount rate.

This leaves the fixed cost ratio, marginal sales, and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries, as well as the cost per kWh, come from Ziegler & Trancik (2021). They report a representative series of the price of all types of lithium-ion cells and one for the market size of all types of lithium-ion cells measured in energy capacity. However, the price data only goes until 2018, and the sales data goes until 2016, so we append price data from the Department of Energy and sales data from the IEA. For 2020, the marginal sales are 167700 MWh of batteries, and the cumulative sales are 917708 starting from 1991. For in-context, the average marginal sales are 62,906 and the average cumulative sales are 337250. The price per kWh in 2020 is \$181.978 and over 2015-18 on average is \$226.306. To calculate the fixed cost ratio, we take the price per kWh and multiply it by the sales-weighted average battery capacity for EVs in 2020, 73.004 kWh (battery capacity for each model comes from Edmunds), or 2015-18, 63.497 kWh. We then divide this by the MSRP (\$54,025 or \$61,678) to get the proportion of the cost due to the battery of 0.246 for 2020 and 0.213 for in-context. One minus that proportion gives us our fixed cost ratio of 0.754 for 2020 and 0.787 for in-context. Recall that for our learning-by-doing model, we only think the learning is affecting the battery price and not the price of other parts of the vehicle, so any price and environmental benefits will be a result of only those battery prices coming down. Following the steps outlined in Appendix B, we obtain a dynamic environmental component for 2020 (in-context) of 0.046 (0.156) and a dynamic price component of 0.309 (0.261).

Profits Lastly, for WTP, we estimate the gasoline producers' and utilities' WTP for the subsidy. Since the gasoline market has higher markups than the average market and electricity markets are heavily regulated to have fixed levels of markups, we believe this is an important component to consider. Using the previously estimated increase in electricity consumption and decrease in gas consumption, we calculate the producers' WTP as the markup multiplied by the change in electricity or gas, normalized by the net MSRP, and then multiplied by the elasticity and the pass-through rate. For each year of the car's lifetime, the annual profits will be discounted. For gasoline, this is 4,642.3/2,993.9 gallons of gasoline multiplied by 0.613/0.682 markup (discounted each year), normalized by the net MSRP, multiplied by the elasticity and the 85% pass-through rate, and finally we subtract out 21% of the surplus due to the effective corporate tax rate to get 0.076/0.045. For utilities, we have 34,476/34,458 kWh of electricity multiplied by 0.011/0.048 markup, normalized, and multiplied by the same elasticity and pass-through rate to get 0.010/0.045.

With the various components, we can now calculate the total WTP of 1.503 for 2020 and 1.527 for in-context.

Cost The cost of a \$1 mechanical increase in EV subsidies is equal to the \$1 plus the fiscal externalities induced by the demand response to the subsidies. These include additional costs to state and federal subsidies, reductions in gas tax revenue, changes in profits taxes collected,

and the climate fiscal externality from changes in CO_2 emissions. We discuss each in turn.

State and Federal FEs In 2020, and in context, there are some existing federal and state subsidies for BEVs; an additional dollar of spending will spur further behavioral effects that will increase the spending from the preexisting programs. Federally, in 2020, most BEVs no longer qualified for the \$7,500 subsidy from the Qualified Plug-In Electric Drive Motor Vehicle Credit. Only eight models qualified, and some of those for less than the full amount. EERE reports the subsidy each model qualified for and when it stopped qualifying for the maximum amount and subsequent smaller amounts. Similar to how we estimate the sales-weighted average for other parameters, we estimate a sales-weighted average federal subsidy of \$7500.00 in 2020 and \$7021.44 in-context. Normalizing that value by the net MSRP gives us \$0.163 in 2020 and \$0.154 in-context. Finally, multiplying by the elasticity gives us the federal fiscal externality of 0.292 for 2020 and 0.275 for in-context.

For state subsidies, we use the Alternative Fuels Data Center's (AFDC) database on incentives and laws related to alternative fuels and advanced vehicles. Nine states in 2020 have subsidies for BEVs, with varying levels of subsidy size, MSRP eligibility rules, and income eligibility rules. For example, in Oregon, there was a \$7,500 subsidy that applied if the battery was greater than 10 kWh, the MSRP was less than \$50,000, and the income for a household of one was below \$54,360. If subsidies differ by driver's income, we scale the subsidy based on the approximate proportion of EV drivers within that income constraint. Muehlegger & Rapson (2019) report the proportion of EV drivers within income categories and find that 14% have household incomes between 0 and \$49,000 and 30% have incomes between 50 and \$99,000. This implies (assuming a uniform distribution of EV purchasers within each bucket) that 16.7% of Oregon's EV purchasers were eligible for the subsidy. In 2020, only 63% of BEVs had an MSRP of less than \$50,000. Thus, we can estimate that the percent of EV purchasers who received the subsidy (assuming independence) is $0.63 \cdot 0.167 = 0.1056$. The average subsidy received in Oregon was \$792.33. We repeat these steps for the other eight states, and then we use the AFDC's data on how many EV registrations occur in each state to compute a final EV registration-weighted average state subsidy of \$604.27. This value normalized by the net MSRP and multiplied by the elasticity gives us the state fiscal externality of 0.028.

Since the in-context MVPF is looking at a specific state's subsidy, we take the subsidy amount to be \$9,000 as reported in Muehlegger & Rapson (2022). When normalized and multiplied by the elasticity, this gives us a state fiscal externality of 0.414.

Gas Tax FE We also calculate a gasoline tax fiscal externality using the average state and federal gas tax rates as described in Appendix E.10. The FE is the elasticity multiplied by the tax rate of 0.465/0.434 multiplied by the decrease in gas consumption from ICE vehicles of 4,642.3/2,993.9 and normalized by the net MSRP to get 0.073/0.041.

Profits Tax FE Similarly to gasoline taxes, we have an average combined revenue rate of 0.006/0.026 that accounts for profits to publicly-owned utilities and corporate taxes on privately owned utilities as described in Appendix C.2.3. The FE is calculated in the same way as the gas tax FE to get 0.006/0.025.

Climate FE Finally, the climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Our baseline social cost of carbon comes from EPA (2023c) and estimates that 15% of that avoided GDP loss (which is 50% of the SCC) would flow to the US. Since the average tax rate in the US is 25.54%, we estimate the climate FE as $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the static global environmental externality and the dynamic global environmental externality that results from learning by doing. Taking those two pieces and multiplying their sum by 1.9% gives us -0.006/-0.006.

Thus, our final cost is 1.401/1.711, which gives us our MVPFs of 1.073 and 0.893.

E.3.4 State-level Income Tax Credits for Battery Electric Vehicles

Clinton & Steinberg (2019) analyze eight state-level income tax credits for battery electric vehicles (BEVs) offered between 2011 and 2014. The programs each varied in the value of the incentive, the time they were in effect, and eligibility requirements. Using a fixed-effect specification, they estimate the effect of a \$1,000 increase in financial incentives to be a 5.5% decrease in per capita BEV registrations. We translate this into an elasticity below.

Throughout this section, the “in-context” specification will mean the eight states (CO, GA, LA, MD, OR, SC, UT, WV) from 2011 to 2014, which is the time and geography analyzed in the paper.

WTP The WTP for an expansion of BEV subsidies is the sum of the transfer (which we normalize to 1) plus the environmental and other market externalities discussed in Section 4. These broadly have the form of the elasticity (-0.001) times the societal willingness to pay for one additional dollar of spending on the BEV, V/p . We estimate V/p separately by focusing on the per-car externalities, V , and the relevant consumer price, p . To measure consumer prices, the Office of Energy Efficiency and Renewable Energy (EERE) lists the manufacturer’s suggested retail price (MSRP) for most BEV models. We compute a sales-weighted average using data from Kelley Blue Book’s Electrified Light-Vehicle Sales Report for Q4 2021, which includes estimates of 2020 sales of each BEV model. For models with multiple trims, EERE would often report a range of MSRPs, in which case we use the mean of the MSRP for that model. For 2020, we have an average MSRP of \$54,025. The total subsidy is the sum of the average federal and state-level subsidies in 2020, which we describe in more detail in the Fiscal Externalities section, and is \$647.25. To compute the elasticity, we take the semi-elasticity reported in the paper -0.055, divide it by \$1,000, and then multiply it by the MSRP net of subsidies for the in-context specification, \$35,866. This gives us an elasticity of -0.001.

For the in-context specification, the sales data by model comes from the Transportation Research Center at Argonne National Laboratory and covers all models up until 2019 that are full-sized and capable of 60 mph. This gives us an average MSRP of \$46,006, which net of the average federal subsidy over 2011-2014 and the average subsidy among the fourteen states leaves us with a net MSRP of \$35,866.

Transfer We consider a \$1 increase in the BEV subsidy with no pass-through rate.

Global Environmental Externalities As described in the main text of the paper, the global environmental benefits are calculated as the difference in the WTP for avoiding the global pollutants from driving the counterfactual internal combustion engine (ICE) vehicle and the WTP for facing the global pollutants from electricity generation needed to drive a battery electric vehicle (BEV). We calculate the two WTPs as follows:

BEV Global Externalities To determine the damages from one electric vehicle, we need to know how many kilowatt-hours (kWh) of electricity the average BEV uses each year and the average vehicle miles traveled (VMT) driven. EERE reports the kWh of electricity needed for each BEV model to travel 100 miles. Combining this data with the sales data mentioned above, we calculate a sales-weighted average of the kWh used by a BEV per mile. For 2020, this average energy consumption is 0.293 kWh per mile, and for in-context, it is 0.326 kWh per mile.

For VMT we use the 2017 results from the Federal Highway Administration’s National Household Travel Survey (NHTS). They report VMT by vehicle age and vehicle type, where the eight types of vehicles include options such as “Automobile/Car/Station Wagon”, “Van (Mini/Cargo/Passenger)”, “SUV (Santa Fe, Tahoe, Jeep, etc.)”, “Pickup Truck”, and “Other Truck”. We use for our main specification the VMT reported in the Automobile/Car/Station Wagon category as is. Then, the VMT is multiplied by 0.6154 following Zhao et al. (2023)’s analysis finding that BEVs accumulate fewer annual miles than ICE vehicles: 7,165 versus 11,642.

Throughout our analysis, we assume the counterfactual ICE vehicles and BEVs have the same VMT, although our primary results are robust to reasonable variations in this assumption. This is because we believe the difference in VMT found in Zhao et al. (2023) is due to selection in the types of drivers that purchase BEVs independent of vehicle type.¹⁶⁷ We assume a 17-year lifespan for both ICE vehicles and BEVs, so we use the VMT numbers corresponding to a car’s lifespan within that range. The survey reports an average VMT of 12,245 miles for the first year, with the remaining values of VMT ranging between 5,885 (in the last year) and 13,078 (in the fifth year). For in-context, we impute age-state-level VMT data using the NHTS VMT by household state and vehicle type. We calculate a percent difference between the sample-weighted average VMT across all ages and each age’s VMT. Then, we assume this percent difference holds within states. Thus, we use the EV registration-weighted average of the nine sample states’ imputed VMTs for the in-context specification.

Once we calculate the energy consumption for a BEV in each year of its lifetime, we estimate the global damages from electricity consumption from the grid using information from AVERT and forecasts of the cleanliness of the grid from Jenkins & Mayfield (2023). Since the VMT changes for each year of the car’s lifetime, the global damages each year change as well. For example, the first year’s energy consumption will be 2,425.525 kWh in 2020 and 2,589.740 in-context, which leads to \$347.623 and \$296.455 respectively, of global damages from grid pollution. See a more detailed description of the grid externalities calculations in Appendix C.2. We then estimate the global damages for each year of the vehicle’s lifetime, which totals \$3,183.512 in global damages in 2020 and \$3,607.115 in context. Normalized by the net MSRP, we have \$0.060 for 2020 and \$0.101 for in-context.

¹⁶⁷It would be straightforward to adjust our approach to allow for differential mileage driven, noting that one would need to distinguish between total reduction in miles and substitution to other ICE vehicles.

We take this amount and multiply it by the elasticity of -0.001 and we subtract out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, which is 1.9% of the total amount (see Section 4). Written out this is $0.060 \cdot -0.001 \cdot (1 - 0.15 \cdot 0.2554 \cdot 0.5)$. Thus, our final values for the damages from using a BEV are 0.120 (2020) and 0.195 (in-context).

ICE Global Externalities For the counterfactual ICE vehicle, we use the same VMT as for the BEVs. We need the average fuel economy to obtain the final gas consumption in each year of the ICE vehicle’s life. Holland et al. (2016) report in their appendix a substitute gas vehicle for each popular BEV in 2014. Most of these are the ICE versions of the BEV (i.e., the gas-powered Ford Focus as the counterfactual for the BEV Ford Focus), but some are not. They test the reasonableness of their choices based on market research data from MaritzCX. Using the sales data in 2014 for the BEVs for which they find counterfactuals, we compute an average counterfactual MPG for 2014, which is 37.16 miles per gallon. This is higher than the average new car fuel economy in 2014 of 27.63 MPG. We then extrapolate to an average counterfactual MPG for other years by calculating the ratio of the counterfactual MPG to the car fleet’s average MPG in 2014 and assuming that the ratio of 1.34 holds for other years. This gives us a counterfactual MPG of 41.23 in 2020 and 39.37 in context. We also explore the robustness of our results to assuming BEVs displace an average fleet light duty vehicle, which slightly raises the MVPF to 0.993.

In the first year of the counterfactual ICE vehicle’s life, we estimate it consumed 326.599 gallons of gas in 2020 and 365.79 in context. With the gas consumed by an ICE vehicle in each year of its life, we estimate the total damages from an ICE with the estimated pollutants as described in Appendix C.4. We take this amount and normalize it by the net MSRP of a BEV to get 0.099 for 2020 and 0.132 in-context and multiply it by the elasticity to get -0.204 and -0.260 and we subtract out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, which is 1.9% of the total amount (see Section 4). Thus, a \$1 increase in the subsidy for a BEV leads to a reduction in damages from driving the counterfactual ICE of -0.200 in 2020 and -0.255 in context.

Upstream Battery Externalities We also include emissions from the production of the batteries used in electric vehicles. Battery production is a unique source of emissions from BEVs compared to ICE vehicles. Winjobi et al. (2022) note that “batteries in electric vehicles can account for one-third of their production greenhouse gas (GHG) emissions” and find in their analysis of different battery chemistries that lithium nickel manganese cobalt oxide batteries with equal proportions of nickel, manganese, and cobalt (NMC11) have life-cycle GHG emissions of 59.5 kg CO₂-eq/kWh. Using the average battery capacity for 2020 (73.004) and in-context (39.777), we have 4,343.744 for 2020 and 2,366.736 in-context. After converting the kg of emissions to tons, multiplying by the SCC, and normalizing by the MSRP, we have 0.016 for 2020 and 0.011 in context. Multiplying by the elasticity and subtracting out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, we have -0.032 for 2020 and -0.021 for in-context.

The sum of the reduction in ICE emissions, increase in grid-related emissions, and upstream battery production emissions gives our final externality measure of -0.048 (2020) and -0.039 (in-context) for the global environmental externality for a \$1 mechanical increase in subsidies to

BEVs.

Local Externalities Local externalities are estimated in nearly the same way as global externalities. We use the same energy and gas consumption values. The difference is the marginal damage per kWh of electricity or gallon of gasoline. We describe the estimation of these local damages in Appendices C.2 and C.4. We calculate -0.009 in 2020 and 0.033 in-context in local damages from BEVs and 0.008 in 2020 and 0.011 in-context from the counterfactual ICE. Taking the difference gives -0.000 and 0.044, respectively. After multiplying by the elasticity, we have 0.000 and 0.043.

Rebound The previous analysis assumes that the increased consumption of BEVs does not affect grid-wide electricity prices. We now consider some general equilibrium effects where the increased use of BEVs increases the price of electricity and thus decreases electricity consumption. We assume this effect does not lead to a secondary rebound in using ICE vehicles because while we allow for local production of electricity with upward-sloping supply curves, we assume a flat global supply curve for gasoline so that there are few if any rebound effects in the gasoline market. The rebound effect is calculated as $\frac{1}{1-\varepsilon_D/\varepsilon_S}$ where ε_D is the demand elasticity for electricity, which is -0.19 and comes from DOI (2021) and ε_S is the supply elasticity for electricity, which is 0.78 (see Appendix D for calculation). With this rebound effect of about 20% multiplied by the previously calculated global and local damages from electricity consumption for BEVs, we estimate a decrease in damages of -0.027 in 2020 and -0.051 in-context from less electricity consumption.

Learning-by-Doing Our model of learning-by-doing is described in Appendix B. Here, we describe preliminary calculations and the data sources for the model inputs. There are nine inputs into the model: the demand elasticity for BEVs, the discount rate, the learning rate, the fraction of a BEV's price that is from non-battery components (we refer to this as the fixed cost ratio), the marginal sales in a given year, the cumulative sales up until the year of interest, the net MSRP, and the environmental damage per EV.

The demand elasticity for this policy of -0.001 is calculated as described above. Our baseline discount rate is 2%. The learning rate for batteries of 0.421 comes from Way et al. (2022). We show in Appendix B how we adjust for the fact that batteries comprise only a fraction of the total cost of the car. Environmental damages are as above for the global and local externalities but converted to a per-car level. However, we allow for the environmental externality to vary over time to account for the SCC increasing over time. Our method of extrapolating to future SCC values is described in Section 4. The net MSRP is as described at the top of the WTP section.

This leaves the fixed cost ratio and marginal and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries, as well as the cost per kWh, come from Ziegler & Trancik (2021). They report a representative series of the price of all types of lithium-ion cells and one for the market size of all types of lithium-ion cells measured in energy capacity. However, the price data only goes until 2018, and the sales data goes until 2016, so we append price data from the Department of Energy and sales data from the IEA. For 2020, the marginal sales are 167,700 MWh of batteries, and the cumulative sales are 917,708 starting from 1991. For

in-context, the average marginal sales are 29,680 and the average cumulative sales are 147,359. The price per kWh in 2020 is \$181.978 and over 2011-14 on average is \$248.777. To calculate the fixed cost ratio, we take the price per kWh and multiply it by the sales-weighted average battery capacity for EVs in 2020, 73.004 kWh (battery capacity for each model comes from Edmunds), or 2011-14, 39.777 kWh. We then divide this by the MSRP (\$54,025 or \$46,006) to get the proportion of the cost due to the battery of 0.246 for 2020 and 0.187 for in-context. One minus that proportion gives us our fixed cost ratio of 0.754 for 2020 and 0.813 for in-context. Recall that for our learning-by-doing model, we only think the learning is affecting the battery price and not the price of other parts of the vehicle, so any price and environmental benefits will be a result of only those battery prices coming down. Following the steps outlined in Appendix B, we obtain a dynamic environmental component for 2020 (in-context) of 0.000 (0.000) and a dynamic price component of 0.000 (0.000).

Profits We estimate the gasoline producers' and utilities' WTP for the subsidy. We note that the gasoline market has higher markups than the average market and electricity markets are heavily regulated to have fixed levels of markups. Using the previously estimated increase in electricity consumption and decrease in gas consumption, we calculate the producers' WTP as the markup multiplied by the change in electricity or gas, normalized by the net MSRP, and then multiplied by the elasticity. For each year of the car's lifetime, the annual profits are discounted. For gasoline, this is 2,857.1/3,199.9 gallons of gasoline multiplied by 0.613/0.933 markup (discounted each year), normalized by the net MSRP, multiplied by the elasticity, and finally we subtract out 21% of the surplus due to the effective corporate tax rate to get -0.046/-0.101. For utilities, we have 34,476/36,810 kWh of electricity multiplied by 0.011/0.001 markup, normalized, and multiplied by the same elasticity to get -0.010/-0.002.

With the various components, we can now calculate the total WTP of 0.961 for 2020 and 1.053 for in-context.

Cost The net government cost of a \$1 mechanical increase in BEV subsidies is equal to the \$1 plus the fiscal externalities induced by the demand response to the subsidies. These include additional costs to state and federal subsidies, reductions in gas tax revenue, changes in profits tax collected, and the climate fiscal externality from changes in CO_2 emissions. We discuss each in turn.

State and Federal FEs In 2020, and in context, there are some existing federal and state subsidies for BEVs; an additional dollar of spending will spur further behavioral effects that will increase the spending from the preexisting programs. Federally, in 2020, most BEVs no longer qualified for the \$7,500 subsidy from the Qualified Plug-In Electric Drive Motor Vehicle Credit. Only eight models qualified, and some of those for less than the full amount. EERE reports the subsidy each model qualified for and when it stopped qualifying for the maximum amount and subsequent smaller amounts. Similar to how we estimate the sales-weighted average for other parameters, we estimate a sales-weighted average federal subsidy of \$42.98 in 2020 and \$7,107.55 in-context. Normalizing that value by the net MSRP gives us \$0.001 in 2020 and \$0.198 in context. Finally, multiplying by the elasticity gives us the federal fiscal externality of -0.002 for 2020 and -0.391 for in-context.

For state subsidies, we use the Alternative Fuels Data Center’s (AFDC) database on incentives and laws related to alternative fuels and advanced vehicles. Nine states in 2020 have subsidies for BEVs, with varying levels of subsidy size, MSRP eligibility rules, and income eligibility rules. For example, in Oregon, there was a \$7,500 subsidy that applied if the battery was greater than 10 kWh, the MSRP was less than \$50,000, and the income for a household of one was below \$54,360. If subsidies differ by driver’s income, we scale the subsidy based on the approximate proportion of EV drivers within that income constraint. Muehlegger & Rapson (2019) report the proportion of EV drivers within income categories and find that 14% have household incomes between 0 and \$49,000 and 30% have incomes between 50 and \$99,000. This implies (assuming a uniform distribution of EV purchasers within each bucket) that 16.7% of Oregon’s EV purchasers were eligible for the subsidy. In 2020, only 63% of BEVs had an MSRP of less than \$50,000. Thus, we can estimate that the percent of EV purchasers who received the subsidy (assuming independence) is $0.63 \cdot 0.167 = 0.1056$. The average subsidy received in Oregon was \$792.33. We repeat these steps for the other eight states, and then we use the AFDC’s data on how many EV registrations occur in each state to compute a final EV registration-weighted average state subsidy of \$604.27. This value normalized by the net MSRP and multiplied by the elasticity gives us the state fiscal externality of -0.023.

Since the in-context MVPF is looking at a specific state’s subsidy, we take the subsidy amount to be \$3,032 as reported in Clinton & Steinberg (2019). When normalized and multiplied by the elasticity, this gives us a state fiscal externality of -0.167.

Gas Tax FE We also calculate a gasoline tax fiscal externality using the average state and federal gas tax rates as described in Appendix E.10. The FE is the elasticity multiplied by the tax rate of 0.465/0.398 multiplied by the decrease in gas consumption from ICE vehicles of 2,857.1/3,199.9 and normalized by the net MSRP to get -0.044/-0.058.

Profits Tax FE Similar to gasoline taxes, we have an average combined revenue rate of 0.006/0.001 that accounts for profits to publicly-owned utilities and corporate taxes on privately owned utilities as described in Appendix C.2.3. The FE is calculated in the same way as the gas tax FE to get -0.006/-0.001.

Climate FE The climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Our baseline social cost of carbon comes from EPA (2023c) based 50% on GDP and estimates that 15% of that avoided GDP loss would flow to the US. Since the average tax rate in the US is 25.54%, we estimate the climate FE as $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the static global environmental externality and the dynamic global environmental externality that results from learning by doing. Taking those two pieces and multiplying their sum by 1.9% gives us 0.003/0.002. Thus, our final cost is 0.927/0.361, which gives us our MVPFs of 1.037 and 2.919.

E.4 Hybrid Electric Vehicles

E.4.1 HEV USA - Income Tax Credit

Gallagher & Muehlegger (2011) analyze eight state-level income tax credits for hybrid vehicles offered between 2000 and 2006. The programs each varied in the value of the incentive, the time they were in effect, and eligibility requirements. Using a fixed-effect specification, they estimate the effect of a \$1,000 increase in financial incentives to be a 2.39% increase in per capita HEV sales. We translate this into an elasticity below. Throughout this section, the “in-context” specification will mean the twelve states (CO, MD, NY, OR, PA, SC, UT, and WV) from 2000 to 2006, which is the time and geography analyzed by the paper.

WTP The WTP for an expansion of HEV subsidies is the sum of the transfer (which we normalize to 1) plus the environmental and other market externalities discussed in Section 4. These broadly have the form of the elasticity (-0.430) times the societal willingness to pay for one additional dollar of spending on the HEV, V/p . We estimate V/p separately by focusing on the per-car externalities V and the relevant consumer price net of subsidies, p . To measure consumer prices, we use the manufacturer’s suggested retail price (MSRP) for each HEV model year from Edmunds and Kelley Blue Book, and we compute a sales-weighted average using data from Kelley Blue Book’s Electrified Light-Vehicle Sales Report for Q4 2021, which includes estimates of 2020 sales of each HEV model. For 2020, we have an average MSRP of \$33,464. The total subsidy is the sum of the average federal and state-level subsidies in 2020, which we describe in more detail in the Fiscal Externalities section, and is \$0. To compute the elasticity, we take the semi-elasticity reported in the paper 0.024, divide it by \$1,000, and then multiply it by the MSRP net of subsidies for the in-context specification, \$17,000. This gives us an elasticity of -0.430.

For the in-context specification, the sales data by model comes from the Transportation Research Center at Argonne National Laboratory and covers all models up until 2019 that are full-sized and capable of 60 mph. This gives us an average MSRP of \$20,084, which net of the average federal subsidy over 2000-2006 and the average subsidy among the twelve states leaves us with a net MSRP of \$17,000.

Transfer We consider a \$1 increase in the HEV subsidy so that inframarginal purchasers value this transfer at \$1.

Global Environmental Externalities As described in the main text of the paper, the global environmental benefits are calculated as the difference in the societal WTP for avoiding the global pollutants from driving the counterfactual internal combustion engine (ICE) vehicle and the WTP for facing the global pollutants from driving a hybrid electric vehicle (HEV). We calculate the two societal WTPs in almost identical ways, only varying the MPG for each vehicle.

HEV Global Externalities To determine the damages from one hybrid vehicle, we need to know how many gallons of gas the average HEV uses each year and the average vehicle miles traveled (VMT) driven. EERE reports the gas mileage of each HEV model. Combining

this data with the sales data mentioned above, we can calculate a sales-weighted average fuel economy of an HEV for a given year. For 2020, this average fuel economy is 42.520 and for in-context is 40.842.

For VMT we use the 2017 results from the Federal Highway Administration’s National Household Travel Survey (NHTS). They report VMT by vehicle age and vehicle type, where the eight types of vehicles include options such as “Automobile/Car/Station Wagon”, “Van (Mini/Cargo/Passenger)”, “SUV (Santa Fe, Tahoe, Jeep, etc.)”, “Pickup Truck”, and “Other Truck”. We use for our main specification the VMT reported in the Automobile/Car/Station Wagon category.

Throughout our analysis, we assume ICE vehicles and HEVs have the same VMT, although our primary results are robust to reasonable variations in this assumption.¹⁶⁸ We assume a 17-year lifespan for both ICE vehicles and HEVs, so we use the VMT numbers corresponding to each year of a car’s lifespan within that range. The survey reports an average VMT of 12,245 miles for the first year, with the remaining values of VMT ranging between 5,885 (in the last year) and 13,078 (in the fifth year). For in-context, we impute age-state-level VMT data using the NHTS VMT by household state and vehicle type. We calculate a percent difference between the sample-weighted average VMT across all ages and each age’s VMT. Then, we assume this percent difference holds within states. Thus, we use the population-weighted average of the twelve sample states’ imputed VMTs for the in-context specification.

Using the gas consumption for an HEV in each year of its lifetime, we estimate the damages from an HEV with the estimated pollutants as described in Appendix C.4. Since the VMT changes for each year of the car’s lifetime, the damages change as well. For exposition, the first year’s gasoline consumption will be 316.721 gallons in 2020 and 302.455 in-context, which leads to \$600.913 and \$445.380 respectively, of damages from gasoline. See a more detailed description of the gasoline externalities calculations in Appendix C.4. We then estimate the damages for each year of the vehicle’s lifetime, which totals 8341.459 in damages in 2020 and -4.9e+03 in context. Normalized by the net MSRP, we have 0.249 for 2020 and -0.291 for in-context.

Finally, we take this amount and multiply it by the elasticity of -0.430 and we subtract out the portion of the benefits that will accrue to the US government via the climate FE from increased GDP from avoiding carbon dioxide emissions, which is 1.9% of the total environmental externality in our baseline specification (see Section 4). Writing this out is $0.249 \cdot -0.430 \cdot (1 - 0.15 \cdot 0.2554 \cdot 0.5)$. Thus, our final values for the damages from using an HEV are -0.105 (2020) and -0.116 (in-context).

ICE Global Externalities For the counterfactual ICE vehicle, we use the same VMT as for the HEVs and estimate a counterfactual average fuel economy based on Muehlegger & Rapson (2023). They report in Table 3 of their paper an estimated effect of hybrid vehicle incentives on fleet gallons per mile of -0.000011. Assuming this result holds over time, we apply the formula suggested by the paper $1 / ((1/\text{MPG}) - -0.000011 * 100)$, where MPG is the average HEV MPG in a given year, to calculate the counterfactual ICE MPG for each year. This gives us a counterfactual MPG of 40.620 in 2020 and 39.081 in context. We also explore the robustness of our results to assuming HEVs displace an average fleet light duty vehicle, which

¹⁶⁸It would be straightforward to adjust our approach to allow for differential mileage driven, noting that one would need to distinguish between total reduction in miles and substitution to other ICE vehicles.

slightly raises the MVPF to 0.970.

In the first year of the counterfactual ICE vehicle's life, we estimate it consumed 331.535 gallons of gas in 2020 and 316.085 in context. With the gas consumed by an ICE vehicle in each year of its life, we estimate the total damages from an ICE with the estimated pollutants as described in Appendix C.4. We take this amount and normalize it by the net MSRP of an HEV to get 0.261 for 2020 and 0.301 in-context and multiply it by the elasticity to get 0.112 and 0.122 and we subtract out the portion of the benefits that will accrue to the US government via the climate FE from increased GDP from avoiding carbon dioxide emissions, which is 1.9% of the total amount (see Section 4). Thus, \$1 of spending on an HEV generates \$0.261 of environmental savings from reduced ICE emissions. Multiplying by the elasticity, this suggests that a \$1 increase in the subsidy for an HEV leads to a reduction in damages from driving the counterfactual ICE is 0.110 in 2020 and 0.120 in context.

Upstream Battery Externalities We also include emissions from the production of the batteries used in hybrid vehicles. Using GREET 2020.NET from the Argonne National Laboratory, Pipitone et al. (2021) find that the production and materials of average HEV battery lead to 234 kg CO_2 eq emissions. After converting the kg of emissions to tons, multiplying by the SCC, and normalizing by the MSRP, we have 0.001 for 2020 and 0.002 in context. Multiplying by the elasticity and subtracting out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon dioxide emissions, we have 0.001 for 2020 and 0.001 for in-context.

The difference in ICE vs. HEV emissions plus upstream battery production emissions gives our final externality measure of 0.004 (2020) and 0.003 (in-context) for the global environmental externality for a \$1 mechanical increase in subsidies to HEVs.

Local Externalities Local externalities are estimated in nearly the same way as global externalities. We use the same gas consumption values. The difference is the marginal damage per gallon of gasoline. We describe the estimation of these local damages in Appendices C.4.1 and C.2. We calculate 0.027 in 2020 and -884.134 in-context in local damages from HEVs and 0.022 in 2020 and 0.045 in-context from the counterfactual ICE. The difference between these yields 0.050 and -884.089, respectively for our 2020 and in-context specifications. After multiplying by the elasticity, we have 0.000 and 0.001.

Rebound We assume there is a rebound effect due to the lowered cost of driving for HEVs due to their higher fuel economy. Small & Van Dender (2007) estimate an elasticity of vehicle miles traveled (VMT) with respect to fuel costs per mile of -0.221. This rebound effect imposes additional damages on society that counteract the benefits of higher fuel economy. To calculate the rebound effect, we first calculate the percent difference in the cost of driving one mile in an HEV compared to the cost in our counterfactual ICE vehicle, which is -0.045 in 2020 and -0.043 in-context. After multiplying by the Small and Van Dender elasticity, we have 0.010 in 2020 and 0.010 in context. We then take the local and global damages and multiply them by the rebound % to arrive at the total externality of -0.004 for 2020 and -0.005 in-context. This same 1.1% increase is applied to the gas consumption values used for gasoline producer profits and the gasoline tax fiscal externality discussed below.

Learning-by-Doing We incorporate potential externalities arising from learning-by-doing in the production of batteries. Relative to the baseline model described in the text, we need to account for the fact that the battery is a small fraction of the total cost of the car; we discuss how we incorporate this in Appendix B. Here, we describe preliminary calculations and data sources used for the model inputs. There are nine inputs into the model: the demand elasticity for HEVs, the discount rate, the learning rate, the fraction of an HEV’s price that is from non-battery components (we refer to this as the fixed cost ratio), the flow of sales in a given year, the cumulative sales up until the year of interest, the net MSRP, the environmental damage per HEV, and the social cost of carbon (SCC).

The demand elasticity for this policy of -0.430 is calculated as described above. Our baseline discount rate is 2%. The learning rate for batteries of 0.421 comes from Way et al. (2022). We show in Appendix B how we adjust for the fact that batteries comprise only a fraction of the total cost of the car. Environmental damages are as above for the global and local externalities but converted to a per-car level. However, we allow for the environmental externality to vary over time to account for the SCC increasing over time. Our method of extrapolating to future SCC values is described in Section 3.2. The net MSRP is as described at the top of the WTP section. The last input is the SCC, which we use a baseline value of \$193 for 2020 and allow to vary over time following projections of a rising SCC from the Environmental Protection Agency’s recent guidance regarding the social cost of carbon at a 2% discount rate.

This leaves the fixed cost ratio and marginal and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries, as well as the cost per kWh, come from Ziegler & Trancik (2021). They report a representative series of the price of all types of lithium-ion cells and one for the market size of all types of lithium-ion cells measured in energy capacity. However, the price data only goes until 2018, and the sales data goes until 2016, so we append price data from the Department of Energy and sales data from the IEA. For 2020, the marginal sales are 167700 MWh of batteries, and the cumulative sales are 917708 starting from 1991. For in-context, the average marginal sales are 2,029.4 and the average cumulative sales are 5,940.7. The price per kWh in 2020 is \$176.532 and over 2000-06 on average is \$710.256. To calculate the fixed cost ratio, we take the price per kWh and multiply it by the sales-weighted average battery capacity for EVs in 2020, 1.469 kWh (battery capacity for each model comes from Edmunds), or 2000-06, 1.354 kWh. We then divide this by the MSRP (\$33,464 or \$20,084) to get the proportion of the cost due to the battery of 0.008 for 2020 and 0.033 for in-context. One minus that proportion gives us our fixed cost ratio of 0.992 for 2020 and 0.967 for in-context. Following the steps outlined in Appendix B, we obtain a dynamic environmental component for 2020 (in-context) of 0.000 (0.000) and a dynamic price component of 0.002 (0.009).

Profits We estimate the gasoline producers’ WTP for the subsidy. If gasoline has higher markups than other goods in the economy, a change in gasoline demand can cause an externality on producers from the change in profits in the economy. Appendix C.4.5 discusses our approach to estimating the markups in the gasoline market, which we estimate to be 35% in 2020, relative to a national average of 8%, for a net markup of 27%. Using the previously estimated gas consumption for the counterfactual ICE and HEVs, we calculate the producers’ WTP as the markup multiplied by the present discounted value of the change in gasoline consumption, normalized by the net MSRP, and then multiplied by the elasticity. This yields

165.883/165.804 gallons of gasoline multiplied by 0.613/0.671 markup (discounted each year), normalized by the net MSRP, multiplied by the elasticity, and finally we subtract out 21% of the surplus due to the effective corporate tax rate to get -0.001/0.002.

With all of these components, we calculate a total WTP of 1.002 for 2020 and 1.011 for in-context.

Cost The net government cost of a \$1 mechanical increase in HEV subsidies is equal to the \$1 plus the fiscal externalities induced by the demand response to the subsidies. These include additional state and federal subsidies, reductions in gas tax revenue, changes in profits tax collected, and the climate fiscal externality from changes in CO_2 emissions. We discuss each in turn.

State and Federal FEs For the in-context specification, there are some existing federal and state subsidies for HEVs that an additional dollar of spending will spur further behavioral effects that will increase the spending from the preexisting programs. By 2020, though, there were no more subsidies available for non-plug-in hybrid vehicles. In Table 3 of Gallagher & Muehlegger (2011), they report an average federal tax incentive of \$1,073. Normalizing that value by the net MSRP gives us \$0.063 in-context. Finally, multiplying by the elasticity gives us the federal fiscal externality of 0.026 for in-context. This quantity is zero for 2020 because there was no subsidy in place.

There were also no state-level subsidies available for HEVs in 2020. For in-context, we take the subsidy amount to be \$2,011, which is the average of the state income tax credit and sales tax incentives reported in Gallagher & Muehlegger (2011). When normalized and multiplied by the elasticity, this gives a state fiscal externality of 0.048.

Gas Tax FE We also calculate a gasoline tax fiscal externality using the average state and federal gas tax rates as described in Appendix E.10. The FE is the elasticity multiplied by the tax rate, which is 0.465 in 2020 and 0.377 in context multiplied by the difference in gas consumption between ICE and HEV vehicles of 165.883 for 2020 and 165.804 for in-context and normalized by the consumer price net of subsidies to get 0.001/0.002.

Profits Tax FE Since we estimate gasoline producers' profits, we account for corporate profits taxation as a fiscal externality. The corporate tax rate is 21%. Multiplying the gasoline producers' WTP by 0.21 gives -0.0002 for the 2020 specification and 0.0005 for the in-context specification.

Climate FE The climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Our baseline social cost of carbon comes from EPA (2023c) estimates that 15% of that avoided GDP loss (which is 50% of the SCC) would flow to the US. Since the average tax rate in the US is 25.54%, we estimate the climate FE as $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the static global environmental externality and the dynamic global environmental externality

that results from learning by doing. Taking those two pieces and multiplying their sum by 1.9% gives us -0.000/-0.000.

Our final cost is 1.001/1.076, which gives us our MVPFs of 1.002 and 0.940.

E.4.2 Federal Income Tax Credit for Hybrid Vehicles

Beresteanu & Li (2011) analyze the federal income tax credit for hybrid vehicles offered between 2000 and 2006. The program authorized a credit of up to \$3,400, depending on the model and the improvement in fuel economy relative to the nonhybrid counterpart. Using a market equilibrium model with both demand and supply sides in the spirit of Berry et al. (1995), they estimate the effect of a \$2,276 increase in financial incentives to be a 19.75% increase in per capita HEV sales. We translate this into an elasticity below. Throughout this section, the “in-context” specification will mean the eighteen states the authors have data from (AR, AZ, CA, CO, CT, FL, GA, IA, MO, NM, NV, NY, OH, PA, TN, TX, WA, and WI) in 2006, which is the time analyzed by the paper.

WTP The WTP for an expansion of HEV subsidies is the sum of the transfer (which we normalize to 1) plus the environmental and other market externalities discussed in Section 4. These broadly have the form of the elasticity (-1.985) times the societal willingness to pay for one additional dollar of spending on the HEV, V/p . We estimate V/p separately by focusing on the per-car externalities V and the relevant consumer price net of subsidies, p . To measure consumer prices, we use the manufacturer’s suggested retail price (MSRP) for each HEV model year from Edmunds and Kelley Blue Book, and we compute a sales-weighted average using data from Kelley Blue Book’s Electrified Light-Vehicle Sales Report for Q4 2021, which includes estimates of 2020 sales of each HEV model. For 2020, we have an average MSRP of \$33,464. The total subsidy is the sum of the average federal and state-level subsidies in 2020, which we describe in more detail in the Fiscal Externalities section, and is \$0. To compute the elasticity, we take the semi-elasticity reported in the paper 0.198, divide it by \$1,000, and then multiply it by the MSRP net of subsidies for the in-context specification, \$21,736. This gives us an elasticity of -1.985.

For the in-context specification, the sales data by model comes from the Transportation Research Center at Argonne National Laboratory and covers all models up until 2019 that are full-sized and capable of 60 mph. This gives us an average MSRP of \$25,758, which is net of the average state subsidy in 2006, and the average federal subsidy leaves us with a net MSRP of \$21,736.

Transfer We consider a \$1 increase in the HEV subsidy so that inframarginal purchasers value this transfer at \$1.

Global Environmental Externalities As described in the main text of the paper, the global environmental benefits are calculated as the difference in the societal WTP for avoiding the global pollutants from driving the counterfactual internal combustion engine (ICE) vehicle and the WTP for facing the global pollutants from driving a hybrid electric vehicle (HEV).

We calculate the two societal WTPs in almost identical ways, only varying the MPG for each vehicle.

HEV Global Externalities To determine the damages from one hybrid vehicle, we need to know how many gallons of gas the average HEV uses each year and the average vehicle miles traveled (VMT) driven. EERE reports the gas mileage of each HEV model. Combining this data with the sales data mentioned above, we can calculate a sales-weighted average fuel economy of an HEV for a given year. For 2020, this average fuel economy is 42.52 and for in-context is 38.44.

For VMT, we use the 2017 results from the Federal Highway Administration’s National Household Travel Survey (NHTS). They report VMT by vehicle age and vehicle type, where the eight types of vehicles include options such as “Automobile/Car/Station Wagon,” “Van (Mini/Cargo/Passenger),” “SUV (Santa Fe, Tahoe, Jeep, etc.),” “Pickup Truck,” and “Other Truck.” For our main specification, we use the VMT reported in the Automobile/Car/Station Wagon category as is. Throughout our analysis, we assume ICE vehicles and HEVs have the same VMT, although our primary results are robust to reasonable variations in this assumption.¹⁶⁹ We assume a 20-year lifespan for both ICE vehicles and HEVs, so we use the VMT numbers corresponding to each year of a car’s lifespan within that range. The survey reports an average VMT of 12,245 miles for the first year, with the remaining values of VMT ranging between 5,885 (in the last year) and 13,078 (in the fifth year). For in-context, we impute age-state-level VMT data using the NHTS VMT by household state and vehicle type. We calculate a percent difference between the sample-weighted average VMT across all ages and each age’s VMT. Then, we assume this percent difference holds within states. Thus, we use the population-weighted average of the twelve sample states’ imputed VMTs for the in-context specification.

Using the gas consumption for an HEV in each year of its lifetime, we estimate the damages from an HEV with the estimated pollutants as described in Appendix C.4. Since the VMT changes for each year of the car’s lifetime, the damages change as well. For exposition, the first year’s gasoline consumption will be 316.721 gallons in 2020 and 333.344 in-context, which leads to 600.913 and 600.913 respectively, of damages from gasoline. See a more detailed description of the gasoline externalities calculations in Appendix C.4. We then estimate the damages for each year of the vehicle’s lifetime, which totals 8341.459 in damages in 2020 and 5510.344 in context. Normalized by the net MSRP, we have 0.249 for 2020 and 0.254 for in-context.

We take this amount and multiply it by the elasticity of -1.985 and we subtract out the portion of the benefits that will accrue to the US government via the climate FE from increased GDP from avoiding carbon emissions, which is 1.9% of the total environmental externality in our baseline specification (see Section 4). Writing this out is $0.249 \cdot -1.985 \cdot (1 - 0.15 \cdot 0.2554 \cdot 0.5)$. Thus, our final values for the damages from using an HEV are -0.485 (2020) and -0.469 (in-context).

ICE Global Externalities For the counterfactual ICE vehicle, we use the same VMT as for the HEVs and estimate a counterfactual average fuel economy based on Muehlegger & Rapson (2023). They report in Table 3 of their paper an estimated effect of hybrid vehicle

¹⁶⁹It would be straightforward to adjust our approach to allow for differential mileage driven, noting that one would need to distinguish between total reduction in miles and substitution to other ICE vehicles.

incentives on fleet gallons per mile of -0.000011. Assuming this result holds over time, we apply the formula suggested by the paper $1 / ((1/\text{MPG}) - -0.000011 * 100)$, where MPG is the average HEV MPG in a given year, to calculate the counterfactual ICE MPG for each year. This gives us a counterfactual MPG of 40.62 in 2020 and 36.88 in context. We also explore the robustness of our results to assuming HEVs displace an average fleet light duty vehicle, which slightly raises the MVPF to 0.859.

In the first year of the counterfactual ICE vehicle's life, we estimate it consumed 331.535 gallons of gas in 2020 and 347.44 in context. With the gas consumed by an ICE vehicle in each year of its life, we estimate the total damages from an ICE with the estimated pollutants as described in Appendix C.4. We take this amount and normalize it by the net MSRP of an HEV to get 0.261 for 2020 and 0.264 in-context and multiply it by the elasticity to get 0.518 and 0.498 and we subtract out the portion of the benefits that will accrue to the US government via the climate FE from increased GDP from avoiding carbon emissions, which is 1.9% of the total amount (see Section 4). Thus, \$1 of spending on an HEV generates \$0.261 of environmental savings from reduced ICE emissions. Multiplying by the elasticity, this suggests that a \$1 increase in the subsidy for an HEV leads to a reduction in damages from driving the counterfactual ICE is 0.508 in 2020 and 0.489 in context.

Upstream Battery Externalities We also include emissions from the production of the batteries used in hybrid vehicles. Using GREET 2020.NET from the Argonne National Laboratory, Pipitone et al. (2021) find that the production and materials of average HEV battery lead to 234 kg CO_2 eq emissions. After converting the kg of emissions to tons, multiplying by the SCC, and normalizing by the MSRP, we have 0.001 for 2020 and 0.001 in context. Multiplying by the elasticity and subtracting out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, we have 0.003 for 2020 and 0.003 for in-context.

The difference in ICE vs. HEV emissions plus upstream battery production emissions gives our final externality measure of 0.020 (2020) and 0.017 (in-context) for the global environmental externality for a \$1 mechanical increase in subsidies to HEVs.

Local Externalities Local externalities are estimated in nearly the same way as global externalities. We use the same gas consumption values. The difference is the marginal damage per gallon of gasoline. We describe the estimation of these local damages in Appendices C.2 and C.4. We calculate 0.027 in 2020 and 0.037 in-context in local damages from HEVs and 0.022 in 2020 and 0.032 in-context from the counterfactual ICE. The difference between these yields 0.050 and 0.070, respectively for our 2020 and in-context specifications. After multiplying by the elasticity, we have 0.002 and 0.002.

Rebound We assume there is a rebound effect due to the lowered cost of driving for HEVs due to their higher fuel economy. Small & Van Dender (2007) estimate an elasticity of vehicle miles traveled (VMT) with respect to fuel costs per mile of -0.221. This rebound effect imposes additional damages on society that counteract the benefits of higher fuel economy. To calculate the rebound effect, we first calculate the percent difference in the cost of driving one mile in an HEV compared to the cost in our counterfactual ICE vehicle, which is -0.045 in 2020 and

-0.041 in-context. After multiplying by the Small and Van Dender elasticity, we have 0.010 in 2020 and 0.009 in context. We then take the local and global damages and multiply them by the rebound % to arrive at the total externality of -0.017 for 2020 and -0.017 in-context. This same 1.1% increase is applied to the gas consumption values used for gasoline producer profits and the gasoline tax fiscal externality discussed below.

Learning-by-Doing We incorporate potential externalities arising from learning-by-doing in the production of batteries. Relative to the baseline model described in the text, we need to account for the fact that the battery is a small fraction of the total cost of the car; we discuss how we incorporate this in Appendix B. Here, we describe preliminary calculations and data sources used for the model inputs. There are nine inputs into the model: the demand elasticity for HEVs, the discount rate, the learning rate, the fraction of an HEV's price that is from non-battery components (we refer to this as the fixed cost ratio), the flow of sales in a given year, the cumulative sales up until the year of interest, the net MSRP, the environmental damage per HEV, and the social cost of carbon (SCC).

The demand elasticity for this policy of -1.985 is calculated as described above. Our baseline discount rate is 2%. The learning rate for batteries of 0.421 comes from Way et al. (2022). We show in Appendix B how we adjust for the fact that batteries comprise only a fraction of the total cost of the car. Environmental damages are as above for the global and local externalities but converted back to the per-car level. However, we allow for the environmental externality to vary over time to account for the SCC increasing over time. Our method of extrapolating to future SCC values is described in Section 3.2. The net MSRP is as described at the top of the WTP section. The last input is the SCC, which we use a baseline value of 193 for 2020 and allow to vary over time following projections of a rising SCC from the Environmental Protection Agency's recent guidance regarding the social cost of carbon at a 2% discount rate.

This leaves the fixed cost ratio and marginal and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries, as well as the cost per kWh, come from Ziegler & Trancik (2021). They report a representative series of the price of all types of lithium-ion cells and one for the market size of all types of lithium-ion cells measured in energy capacity. However, the price data only goes until 2018, and the sales data goes until 2016, so we append price data from the Department of Energy and sales data from the IEA. For 2020, the marginal sales are 167700 MWh of batteries, and the cumulative sales are 917708 starting from 1991. For in-context, the average marginal sales are 11,343 and the average cumulative sales are 43,801. The price per kWh in 2020 is \$176.532 and over 2000-06 on average is \$561.411. To calculate the fixed cost ratio, we take the price per kWh and multiply it by the sales-weighted average battery capacity for EVs in 2020, 1.469 kWh (battery capacity for each model comes from Edmunds), or 2000-06, 1.436 kWh. We then divide this by the MSRP (\$33,464 or \$25,758) to get the proportion of the cost due to the battery of 0.008 for 2020 and 0.025 for in-context. One minus that proportion gives us our fixed cost ratio of 0.992 for 2020 and 0.975 for in-context. Following the steps outlined in Appendix B, we obtain a dynamic environmental component for 2020 (in-context) of 0.0001 (0.0002) and a dynamic price component of 0.009 (0.031).

Profits We estimate the gasoline producers' WTP for the subsidy. If gasoline has higher markups than other goods in the economy, a change in gasoline demand can cause an externality on producers from the change in profits in the economy. Appendix C.4.5 discusses our approach to estimating the markups in the gasoline market, which we estimate to be 35% in 2020, relative to a national average of 8%, for a net markup of 27%. Using the previously estimated gas consumptions for the counterfactual ICE and HEVs, we calculate the producers' WTP as the markup multiplied by the present discounted value of the change in gasoline consumption, normalized by the net MSRP, and then multiplied by the elasticity. This yields 165.883/165.691 gallons of gasoline multiplied by 0.613/0.725 markup (discounted each year), normalized by the net MSRP, multiplied by the elasticity, and finally we subtract out 21% of the surplus due to the effective corporate tax rate to get -0.004/0.009.

With all of these components, we can now calculate the total WTP of 1.010 for 2020 and 1.043 for in-context.

Cost The net government cost of a \$1 mechanical increase in HEV subsidies is equal to the \$1 plus the fiscal externalities induced by the demand response to the subsidies. These include additional costs to state and federal subsidies, reductions in gas tax revenue, changes in profits tax collected, and the climate fiscal externality from changes in CO_2 emissions. We discuss each in turn.

State and Federal FEs For the in-context specification, there are some existing federal and state subsidies for HEVs that an additional dollar of spending will spur further behavioral effects that will increase the spending from the preexisting programs. By 2020, though, there were no more subsidies available for non-plug-in hybrid vehicles. In Table 6 of Beresteanu & Li (2011), they report an average federal credit in 2006 of \$2,276. Normalizing that value by the net MSRP gives us \$0.105 in-context. Finally, multiplying by the elasticity gives us the federal fiscal externality of 0.197 for in-context. This quantity is zero for 2020 because there was no subsidy in place.

There were also no state-level subsidies available for HEVs in 2020. For in-context, we take the subsidy amount to be \$1746.89, which is the average of the state income tax credit and sales tax incentives reported in Gallagher & Muehlegger (2011). When normalized and multiplied by the elasticity, this gives us a state fiscal externality of 0.197.

Gas Tax FE We calculate a gasoline tax fiscal externality using the average state and federal gas tax rates as described in Appendix E.10. The FE is the elasticity multiplied by the tax rate, which is 0.465 in 2020 and 0.387 in context multiplied by the difference in gas consumption between ICE and HEV vehicles of 165.883 for 2020 and 165.691 for in-context and normalized by the consumer price net of subsidies to get 0.004/0.006.

Profits Tax FE Since we estimate gasoline producers' profits, we account for corporate profits taxation as a fiscal externality. The corporate tax rate is 21%, so multiplying the

gasoline producers' WTP by 0.21, we have -0.0011 for the 2020 specification and 0.0023 for the in-context specification.

Climate FE Finally, the climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Our baseline social cost of carbon comes from EPA (2023c) and estimates that 15% of that avoided GDP loss (which is 50% of the SCC) would flow to the US. Since the average tax rate in the US is 25.54%, we estimate the climate FE as $0.15 \cdot 0.5 \cdot 0.2554 = 1.9\%$ of the static global environmental externality and the dynamic global environmental externality that results from learning by doing. Taking those two pieces and multiplying their sum by 1.9% gives us -0.000/-0.000.

Thus our final cost is 1.002/1.357, which gives us our MVPFs of 1.008 and 0.769.

E.4.3 HEV USA - Sales Tax Waiver

Gallagher & Muehlegger (2011) analyze four state-level income tax credits for hybrid vehicles offered between 2000 and 2006. The programs each varied in the value of the incentive, the time they were in effect, and eligibility requirements. Using a fixed-effect specification, they estimate the effect of a \$1,000 increase in financial incentives to be a 2.39% increase in per capita HEV sales. We translate this into an elasticity below. Throughout this section, the “in-context” specification will mean the twelve states (CT, DC, ME, NM) from 2000 to 2006, which is the time and geography analyzed by the paper.

WTP The WTP for an expansion of HEV subsidies is the sum of the transfer (which we normalize to 1) plus the environmental and other market externalities discussed in Section 4. These broadly have the form of the elasticity (-6.916) times the societal willingness to pay for one additional dollar of spending on the HEV, V/p . We estimate V/p separately by focusing on the per-car externalities V and the relevant consumer price net of subsidies, p . To measure consumer prices, we use the manufacturer's suggested retail price (MSRP) for each HEV model year from Edmunds and Kelley Blue Book, and we compute a sales-weighted average using data from Kelley Blue Book's Electrified Light-Vehicle Sales Report for Q4 2021, which includes estimates of 2020 sales of each HEV model. For 2020, we have an average MSRP of \$33,464. The total subsidy is the sum of the average federal and state-level subsidies in 2020, which we describe in more detail in the Fiscal Externalities section, and is \$0. To compute the elasticity, we take the semi-elasticity reported in the paper 0.374, divide it by \$1,000, and then multiply it by the MSRP net of subsidies for the in-context specification, \$17,974. This gives us an elasticity of -6.916.

For the in-context specification, the sales data by model comes from the Transportation Research Center at Argonne National Laboratory and covers all models up until 2019 that are full-sized and capable of 60 mph . This gives us an average MSRP of \$20,084, which net of the average federal subsidy over 2000-2006 and the average subsidy among the twelve states leaves us with a net MSRP of \$17,974.

Transfer We consider a \$1 increase in the HEV subsidy so that inframarginal purchasers value this transfer at \$1.

Global Environmental Externalities As described in the main text of the paper, the global environmental benefits are calculated as the difference in the societal WTP for avoiding the global pollutants from driving the counterfactual internal combustion engine (ICE) vehicle and the WTP for facing the global pollutants from driving a hybrid electric vehicle (HEV). We calculate the two societal WTPs in almost identical ways, only varying the MPG for each vehicle.

HEV Global Externalities To determine the damages from one hybrid vehicle, we need to know how many gallons of gas the average HEV uses each year and the average vehicle miles traveled (VMT) driven. EERE reports the gas mileage of each HEV model. Combining this data with the sales data mentioned above, we can calculate a sales-weighted average fuel economy of an HEV for a given year. For 2020, this average fuel economy is 42.520 and for in-context is 40.842.

For VMT we use the 2017 results from the Federal Highway Administration’s National Household Travel Survey (NHTS). They report VMT by vehicle age and vehicle type, where the eight types of vehicles include options such as “Automobile/Car/Station Wagon”, “Van (Mini/Cargo/Passenger)”, “SUV (Santa Fe, Tahoe, Jeep, etc.)”, “Pickup Truck”, and “Other Truck”. We use for our main specification the VMT reported in the Automobile/Car/Station Wagon category as is.

Throughout our analysis, we assume ICE vehicles and HEVs have the same VMT, although our primary results are robust to reasonable variations in this assumption.¹⁷⁰ We assume a 17-year lifespan for both ICE vehicles and HEVs, so we use the VMT numbers corresponding to each year of a car’s lifespan within that range. The survey reports an average VMT of 12,245 miles for the first year, with the remaining values of VMT ranging between 5,885 (in the last year) and 13,078 (in the fifth year). For in-context, we impute age-state-level VMT data using the NHTS VMT by household state and vehicle type. We calculate a percent difference between the sample-weighted average VMT across all ages and each age’s VMT. Then, we assume this percent difference holds within states. Thus, we use the population-weighted average of the twelve sample states’ imputed VMTs for the in-context specification.

Using the gas consumption for an HEV in each year of its lifetime, we estimate the damages from an HEV with the estimated pollutants as described in Appendix C.4. Since the VMT changes for each year of the car’s lifetime, the damages change as well. For exposition, the first year’s gasoline consumption will be 316.721 gallons in 2020 and 307.184 in-context, which leads to 600.913 and 445.380 respectively, of damages from gasoline. See a more detailed description of the gasoline externalities calculations in Appendix C.4. We then estimate the damages for each year of the vehicle’s lifetime, which totals 8341.459 in damages in 2020 and -4.9e+03 in context. Normalized by the net MSRP, we have 0.249 for 2020 and -0.275 for in-context.

Finally, we take this amount and multiply it by the elasticity of -6.916 and we subtract

¹⁷⁰It would be straightforward to adjust our approach to allow for differential mileage driven, noting that one would need to distinguish between total reduction in miles and substitution to other ICE vehicles.

out the portion of the benefits that will accrue to the US government via the climate FE from increased GDP from avoiding carbon emissions, which is 1.9% of the total environmental externality in our baseline specification (see Section 4). Writing this out is $0.249 \cdot -6.916 \cdot (1 - 0.15 \cdot 0.2554 \cdot 0.5)$. Thus, our final values for the damages from using an HEV are -1.691 (2020) and -1.813 (in-context).

ICE Global Externalities For the counterfactual ICE vehicle, we use the same VMT as for the HEVs and estimate a counterfactual average fuel economy based on Muehlegger & Rapson (2023). They report in Table 3 of their paper an estimated effect of hybrid vehicle incentives on fleet gallons per mile of -0.000011. Assuming this result holds over time, we apply the formula suggested by the paper $1 / ((1/\text{MPG}) - -0.000011 * 100)$, where MPG is the average HEV MPG in a given year, to calculate the counterfactual ICE MPG for each year. This gives us a counterfactual MPG of 40.620 in 2020 and 39.081 in context. We also explore the robustness of our results to assuming HEVs displace an average fleet light duty vehicle, which slightly raises the MVPF to 0.476.

In the first year of the counterfactual ICE vehicle’s life, we estimate it consumed 331.535 gallons of gas in 2020 and 321.026 in context. With the gas consumed by an ICE vehicle in each year of its life, we estimate the total damages from an ICE with the estimated pollutants as described in Appendix C.4. Again, we take this amount and normalize it by the net MSRP of an HEV to get 0.261 for 2020 and 0.285 in-context and multiply it by the elasticity to get 1.805 and 1.913 and we subtract out the portion of the benefits that will accrue to the US government via the climate FE from increased GDP from avoiding carbon emissions, which is 1.9% of the total amount (see Section 4). Thus, \$1 of spending on an HEV generates \$0.261 of environmental savings from reduced ICE emissions. Multiplying by the elasticity, this suggests that a \$1 increase in the subsidy for an HEV leads to a reduction in damages from driving the counterfactual ICE is 1.770 in 2020 and 1.876 in context.

Upstream Battery Externalities We also include emissions from the production of the batteries used in hybrid vehicles. Using GREET 2020.NET from the Argonne National Laboratory, Pipitone et al. (2021) find that the production and materials of average HEV battery lead to 234 kg CO_2 eq emissions. After converting the kg of emissions to tons, multiplying by the SCC, and normalizing by the MSRP, we have 0.001 for 2020 and 0.002 in context. Multiplying by the elasticity and subtracting out the portion of the benefits that will accrue to the US government via increased GDP from avoiding carbon emissions, we have 0.009 for 2020 and 0.012 for in-context.

The difference in ICE vs. HEV emissions plus upstream battery production emissions gives our final externality measure of 0.070 (2020) and 0.052 (in-context) for the global environmental externality for a \$1 mechanical increase in subsidies to HEVs.

Local Externalities Local externalities are estimated in nearly the same way as global externalities. We use the same gas consumption values. The difference is the marginal damage per gallon of gasoline. We describe the estimation of these local damages in Appendices C.2 and C.4. We calculate 0.027 in 2020 and -884.134 in-context in local damages from HEVs and 0.022 in 2020 and 0.043 in-context from the counterfactual ICE. The difference between these yields 0.050 and -884.092, respectively for our 2020 and in-context specifications. After multiplying

by the elasticity, we have 0.007 and 0.012.

Rebound We assume there is a rebound effect due to the lowered cost of driving for HEVs due to their higher fuel economy. Small & Van Dender (2007) estimate an elasticity of vehicle miles traveled (VMT) with respect to fuel costs per mile of -0.221. This rebound effect imposes additional damages on society that counteract the benefits of higher fuel economy. To calculate the rebound effect, we first calculate the percent difference in the cost of driving one mile in an HEV compared to the cost in our counterfactual ICE vehicle, which is -0.045 in 2020 and -0.043 in-context. After multiplying by the Small and Van Dender elasticity, we have 0.010 in 2020 and 0.010 in context. We then take the local and global damages and multiply them by the rebound % to arrive at the total externality of -0.059 for 2020 and -0.072 in-context. This same 1.1% increase is applied to the gas consumption values used for gasoline producer profits and the gasoline tax fiscal externality discussed below.

Learning-by-Doing Next, we incorporate potential externalities arising from learning-by-doing in the production of batteries. Relative to the baseline model described in the text, we need to account for the fact that the battery is a small fraction of the total cost of the car; we discuss how we incorporate this in Appendix B. Here, we describe preliminary calculations and data sources used for the model inputs. There are nine inputs into the model: the demand elasticity for HEVs, the discount rate, the learning rate, the fraction of an HEV's price that is from non-battery components (we refer to this as the fixed cost ratio), the flow of sales in a given year, the cumulative sales up until the year of interest, the net MSRP, the environmental damage per HEV, and the social cost of carbon (SCC).

The demand elasticity for this policy of -6.916 is calculated as described above. Our baseline discount rate is 2%. The learning rate for batteries of 0.421 comes from Way et al. (2022). We show in Appendix B how we adjust for the fact that batteries comprise only a fraction of the total cost of the car. Environmental damages are as above for the global and local externalities but converted back to the per-car level. However, we allow for the environmental externality to vary over time to account for the SCC increasing over time. Our method of extrapolating to future SCC values is described in Section 3.2. The net MSRP is as described at the top of the WTP section. The last input is the SCC, which we use a baseline value of 193 for 2020 and allow to vary over time following projections of a rising SCC from the Environmental Protection Agency's recent guidance regarding the social cost of carbon at a 2% discount rate.

This leaves the fixed cost ratio and marginal and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries, as well as the cost per kWh, come from Ziegler & Trancik (2021). They report a representative series of the price of all types of lithium-ion cells and one for the market size of all types of lithium-ion cells measured in energy capacity. However, the price data only goes until 2018, and the sales data goes until 2016, so we append price data from the Department of Energy and sales data from the IEA. For 2020, the marginal sales are 167700 MWh of batteries, and the cumulative sales are 917708 starting from 1991. For in-context, the average marginal sales are 2,029.4 and the average cumulative sales are 5,940.7. The price per kWh in 2020 is \$176.532 and over 2000-06 on average is \$710.256. To calculate the fixed cost ratio, we take the price per kWh and multiply it by the sales-weighted average battery capacity for EVs in 2020, 1.469 kWh (battery capacity for each model comes from

Edmunds), or 2000-06, 1.354 kWh. We then divide this by the MSRP (\$33,464 or \$20,084) to get the proportion of the cost due to the battery of 0.008 for 2020 and 0.033 for in-context. One minus that proportion gives us our fixed cost ratio of 0.992 for 2020 and 0.967 for in-context. Following the steps outlined in Appendix B, we obtain a dynamic environmental component for 2020 (in-context) of 0.001 (0.002) and a dynamic price component of 0.031 (0.167).

Profits We estimate the gasoline producers' WTP for the subsidy. If gasoline has higher markups than other goods in the economy, a change in gasoline demand can cause an externality on producers from the change in profits in the economy. Appendix C.4.5 discusses our approach to estimating the markups in the gasoline market, which we estimate to be 35% in 2020, relative to a national average of 8%, for a net markup of 27%. Using the previously estimated gas consumptions for the counterfactual ICE and HEVs, we calculate the producers' WTP as the markup multiplied by the present discounted value of the change in gasoline consumption, normalized by the net MSRP, and then multiplied by the elasticity. This yields 165.883/165.804 gallons of gasoline multiplied by 0.613/0.671 markup (discounted each year), normalized by the net MSRP, multiplied by the elasticity, and finally we subtract out 21% of the surplus due to the effective corporate tax rate to get -0.014/0.028.

With all of these components, we calculate a total WTP of 1.036 for 2020 and 1.188 for in-context.

Cost The net government cost of a \$1 mechanical increase in HEV subsidies is equal to the \$1 plus the fiscal externalities induced by the demand response to the subsidies. These include additional costs to state and federal subsidies, reductions in gas tax revenue, changes in profits tax collected, and the climate fiscal externality from changes in CO_2 emissions. We discuss each in turn.

State and Federal FEs For the in-context specification, there are some existing federal and state subsidies for HEVs that an additional dollar of spending will spur further behavioral effects that will increase the spending from the preexisting programs. By 2020, though, there were no more subsidies available for non-plug-in hybrid vehicles. In Table 3 of Gallagher & Muehlegger (2011), they report an average federal tax incentive of \$1,073. Normalizing that value by the net MSRP gives us \$0.060 in-context. Finally, multiplying by the elasticity gives us the federal fiscal externality of 0.401 for in-context. This quantity is zero for 2020 because there was no subsidy in place.

There were also no state-level subsidies available for HEVs in 2020. For in-context, we take the subsidy amount to be \$1,037, which is the average of the state income tax credit and sales tax incentives reported in Gallagher & Muehlegger (2011). When normalized and multiplied by the elasticity, this gives us a state fiscal externality of 0.388.

Gas Tax FE We calculate a gasoline tax fiscal externality using the average state and federal gas tax rates as described in Appendix E.10. The FE is the elasticity multiplied by

the tax rate, which is 0.465 in 2020 and 0.377 in context multiplied by the difference in gas consumption between ICE and HEV vehicles of 165.883 for 2020 and 165.804 for in-context and normalized by the consumer price net of subsidies to get 0.014/0.024.

Profits Tax FE Since we estimate gasoline producers' profits, we account for corporate profits taxation as a fiscal externality. The corporate tax rate is 21%, so multiplying the gasoline producers' WTP by 0.21, we have -0.0038 for the 2020 specification and 0.0076 for the in-context specification.

Climate FE The climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Our baseline social cost of carbon comes from EPA (2023*c*) and estimates that 15% of that avoided GDP loss (which is 50% of the SCC) would flow to the US. Since the average tax rate in the US is 25.54%, we estimate the climate FE as $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the static global environmental externality and the dynamic global environmental externality that results from learning by doing. Taking those two pieces and multiplying their sum by 1.9% gives us -0.002/-0.002.

Thus our final cost is 1.008/1.819, which gives us our MVPFs of 1.028 and 0.653.

E.5 Appliance Rebates

Our category average MVPF for appliance rebates is 1.16. This appendix describes the construction of the individual MVPFs that feed into this category average.

Appliance rebates provide financial incentives to individuals and businesses to adopt energy-saving technologies. Such incentives can take various forms, including tax credits, cash rebates, or discounts on energy-efficient products. While weatherization programs typically involve comprehensive changes to household energy infrastructure (e.g., HVAC), energy rebates focus on individual technologies (e.g., dish washer).

The appliance rebates in our sample are for dishwashers, refrigerators, clothes washers, and water heaters. Most of the papers focus on rebates for Energy Star rated appliances. Energy Star is administered by the US Environmental Protection Agency and provides a set of energy efficiency criteria that companies can voluntarily meet. If an appliance meets this criteria, it receives the Energy Star label. Americans purchased over 300 million Energy Star certified products in 2021 (DOE 2023*b*).

The MVPF construction for appliance rebates is similar to that of weatherization. We allow for differences across MVPFs in subsidy levels, appliance costs, baseline energy usage, and treatment effects, but we harmonize the underlying electricity and natural gas externalities as described in Appendix C.2 and C.3, respectively. The willingness to pay consists of the mechanical transfer to households, environmental externality, rebound effect, and effect on producer profits. The total cost for each program is the sum of the average subsidy level, fiscal externality from the change in utility profit tax revenue, and the climate fiscal externality.

All papers in our sample report the average treatment effect of the appliance rebate on energy usage, but they do not all report the percent of beneficiaries that are inframarginal. For papers that do not estimate this, we apply the same assumption as we do for weatherization policies and assume that 50% of households are marginal. We also assume a uniform distribution

over the potential threshold subsidy at which people would do the retrofit, resulting in marginal households valuing the subsidy, on average, at 50%. Therefore, a \$1 mechanical transfer will lead to \$0.50 of benefits for inframarginal households and \$0.25 for marginal households.

For marginal beneficiaries, the environmental externality is calculated as the product of the treatment effect, baseline usage, and monetized damage per kwh. The effect on producer profits has an analogous calculation. The environmental externality per kWh in 2020 is \$0.159 and the producer profit level for electricity in 2020 is \$0.011. We allow for the electricity grid to change over the lifetime of each appliance as described in Appendix C.2. Some appliance rebate programs also effect natural gas usage. The externalities for natural gas are explained in Appendix C.3 and result in an environmental and producer profits externality per MMBtu of \$10.247 and \$4.396, respectively.

Some papers in this category estimate a price elasticity for energy efficient appliances. For those MVPFs, we calculate the externality per dollar spent (V/p) and multiply this ratio by the elasticity to calculate the externality components of the WTP and cost.

We construct both in-context MVPFs using externalities from the year and state the policy was implemented in as well as baseline MVPFs for the US in 2020.

E.5.1 Cash for Appliances - Houde & Aldy (2017)

Houde & Aldy (2017) estimate take-up of cash for appliance rebates for clothes washers, dishwashers, and refrigerators. The corresponding MVPFs in 2020 are 1.405, 1.132, and 1.042. In context, they are 1.460, 1.153, and 1.053.

The State Energy Efficient Appliance Rebate Program (SEEARP), more commonly known as “Cash for Appliances” (C4A), was funded through the 2009 Recovery Act. The goal of the program was to incentivize the purchase of energy efficient residential appliances. State governments received \$300 million to subsidize purchases of appliances that had an ENERGYSTAR (ES) rating. State governments had significant discretion over the roll-out of the program which created geographic variation in the program’s timing, generosity, and appliance eligibility. Houde & Aldy (2017) use this variation to estimate the impact of C4A on household purchases of energy efficient refrigerators, clothes washers, and dishwashers. Analyzing transaction-level sales data from a large national retailer, they estimate the proportion of individuals that were induced to change their behavior as a result of the policy.

Houde & Aldy (2017) find that a vast majority of individuals would have bought the ES-rated appliance in the absence of the rebate or simply delayed/accelerated their consumption by a few weeks. We consider people who made short-term changes in their purchase timings inframarginal. In total, they find that the percent of inframarginal beneficiaries of clothes washers, dishwashers, and refrigerators rebates was 90.5%, 85.9%, and 92.0%, respectively.

Consistent with Houde & Aldy (2017), we assume that these appliances have a 15-year lifetime. Using data on the manufacturing year of scrapped appliances, Houde & Aldy (2017) estimate that the marginal beneficiary accelerates their purchase of an appliance by 5 years. For the in-context MVPF, we estimate the MVPF of the policy in 2010, the year it was implemented. For the first 5 years of the appliance lifetime, we assume the counterfactual appliance is a non-ES rated appliance purchased 10 years prior (2001). For the next 10 years of the appliance lifetime, we assume the counterfactual is a non-ES rated appliance purchased at the time of the

policy (2010) ¹⁷¹. Houde & Aldy (2017) provide estimates of the difference in energy efficiency for ES and non-ES rated appliances for 2001 and 2010. For the baseline MVPF, we assume the difference between ES and non-ES appliances during the first five years and next ten years is the same as the in-context difference.

The clothes washer, dish washer, and refrigerator rebates were implemented across 43, 37, and 44 states, respectively. Therefore, we will use values for the entire U.S. for both the in-context and baseline specifications. The in-context MVPF will use externalities from 2010 and the baseline MVPF will use 2020 values ¹⁷².

Clothes Washers

Our MVPF for clothes washers using estimates from Houde & Aldy (2017) is 1.41 and 1.46 in-context. The average rebate for clothes washers was \$107 (in 2010 dollars). There were 43 states that offered rebates for clothes washers and a total of 580,863 rebate claims. Houde & Aldy (2017) estimate that the percent of inframarginal rebate recipients for clothes washers (including those that only slightly delayed/accelerated consumption) was 90.5%.

To estimate the energy savings from a marginal beneficiary purchasing an ES-rated clothes washer, the authors report the difference between an ES and non-ES rated clothes washer in 2010 as 201 kWh per year. We use this number for the kWh reduction in years 6-15 of the clothes washers lifetime. For years 1-5, we compare the 2010 ES-rated clothes washer with a 2001 non-ES rated clothes washer. The paper does not directly report this number. It does report the ratio of the rebate amount to the total lifetime reduction. Using this ratio, we calculate the kWh difference for the first five years of the ES-rated appliance to be 668 kWh per year. We multiply the annual reductions by the percent of marginal rebate recipients (9.5%) to get the kWh change each year as a result of the subsidy.

Cost The total cost is comprised of the direct rebate cost and fiscal externalities. The average rebate for clothes washers was \$107 in 2010 dollars.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2. The fiscal externality for electricity per kWh is \$0.006 in 2020 and \$0.005 in 2010. Using the annual kWh reduction per rebate, discounting over the 15 years of the appliance lifetime, we arrive at a total fiscal externality per rebate in the US in 2020 of \$2.67 and in 2010 of \$1.960.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government cost by \$1.10 in 2020 and \$0.776 in-context. The resulting total cost is \$128.58 in the baseline and \$108.18 in-context.

WTP The willingness to pay is comprised of the marginal and inframarginal benefits, environmental externality, and loss in producer profits.

90.5% of the beneficiaries of the policy are inframarginal. Since we assume that inframarginal households value 100% of the \$107 rebate and marginal households value 50% of the \$107 rebate, the willingness to pay for the transfer is \$101.92 in-context. Inflation adjusting to 2020 dollars, the baseline willingness to pay is \$120.97.

¹⁷¹In theory, years 6-15 of the ES appliance lifetime should be compared to a non-ES appliance bought 5 years after the policy (2015). However, the paper only reports energy efficiency data for the difference in 2010.

¹⁷²Houde & Aldy (2017) use price data from 2008 to 2012. We assume that all the prices reported in their paper are in 2010 dollars.

The local and global environmental externality per kWh of electricity is explained in Appendix C.2. The resulting local and global per-kWh environmental externality for the US in 2020 is \$0.133 and \$0.026, respectively. The corresponding externality in 2010 is \$0.100 and \$0.081. While these are point in time estimates, we allow the electricity grid and social costs to change over the 15 years of the appliance lifetime. Using the annual changes in energy consumption and discounting over 15 years, the global environmental externality is \$71.20 in 2020 and \$50.35 in-context. The local environmental externality is \$10.53 in 2020 and \$24.83 in-context. The rebound effect, as explained in Appendix D, offsets 20% of the local and global environmental benefits. The resulting rebound effect is \$16.01 in 2020 and \$14.73 in-context.

Reduced energy consumption as a result of the ES-rated clothes washer leads to lower profits for electric utilities, as explained in Appendix C.2. The loss in profits per kWh of electricity is \$0.011 in 2020 and \$0.010 in 2010. Discounting over the lifetime of the appliance, we arrive at a total producer willingness to pay of -\$4.92 in the baseline specification and -\$3.61 in-context. Summing across these components, the total willingness to pay in 2020 is \$180.67 and in-context is \$157.99. This results in a baseline MVPF of 1.41 and in-context MVPF of 1.46.

Dish Washers

Our MVPF for dish washers using estimates from Houde & Aldy (2017) is 1.13 and 1.15 in-context. The average rebate for dish washers was \$84 (in 2010 dollars). There were 37 states that offered rebates for dish washers and a total of 316,117 rebate claims. Houde & Aldy (2017) estimate that the percent of inframarginal rebate recipients for dish washers (including those that only slightly delayed/accelerated consumption) was 85.9%.

To estimate the energy savings from a marginal beneficiary purchasing an ES-rated dish washer, the authors report the difference between an ES and non-ES rated dish washer in 2010 as 34 kWh per year. We use this number for the kWh reduction in years 6-15 of the dish washers lifetime. For years 1-5, we compare the 2010 ES-rated clothes washer with a 2001 non-ES rated clothes washer. The paper does not directly report this number. It does report the ratio of the rebate amount to the total lifetime reduction. Using this ratio, we calculate the kWh difference for the first five years of the ES-rated appliance to be 234.5 kWh per year. We multiply the annual reductions by the percent of marginal rebate recipients (14.1%) to get the kWh change each year as a result of the subsidy.

Cost The total cost is comprised of the direct rebate cost and fiscal externalities. The average rebate for clothes washers was \$84 in 2010 dollars.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2. The fiscal externality for electricity per kWh is \$0.006 in 2020 and \$0.005 in 2010. Using the annual kWh reduction per rebate, discounting over the 15 years of the appliance lifetime, we arrive at a total fiscal externality in the US in 2020 of \$0.900 and in 2010 of \$0.661.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government cost by \$0.381 in 2020 and \$0.259 in-context. The resulting total cost is \$100.22 in the baseline and \$84.40 in-context.

WTP The willingness to pay is comprised of the marginal and inframarginal benefits, environmental externality, and loss in producer profits.

85.9% of the beneficiaries of the policy are inframarginal. Since we assume that inframarginal households value 100% of the \$84 rebate and marginal households value 50% of the \$84 rebate,

the willingness to pay for the transfer is \$78.08 in-context. Inflation adjusting to 2020 dollars, the baseline willingness to pay is \$92.67.

The local and global environmental externality per kWh of electricity is explained in Appendix C.2. The resulting local and global per-kWh environmental externality for the US in 2020 is \$0.133 and \$0.026, respectively. The corresponding externality in 2010 is \$0.100 and \$0.081. While these are point in time estimates, we allow the electricity grid and social costs to change over the 15 years of the appliance lifetime. Using the annual changes in energy consumption and discounting over 15 years, the global environmental externality is \$24.74 in 2020 and \$16.82 in-context. The local environmental externality is \$3.68 in 2020 and \$8.93 in-context. The rebound effect, as explained in Appendix D, offsets 20% of the local and global environmental benefits. The resulting rebound effect is \$5.57 in 2020 and \$5.04 in-context.

Reduced energy consumption as a result of the ES-rated dish washer leads to lower profits for electric utilities, as explained in Appendix C.2. The loss in profits per kWh of electricity is \$0.011 in 2020 and \$0.010 in 2010. Discounting over the lifetime of the appliance, we arrive at a total producer willingness to pay of -\$1.66 in the baseline specification and \$1.22 in-context. Summing across these components, the total willingness to pay in 2020 is \$113.49 and in-context is \$97.31. This results in a baseline MVPF of 1.132 and in-context MVPF of 1.153.

Refrigerators

Our MVPF for refrigerators using estimates from Houde & Aldy (2017) is 1.04 and 1.05 in-context. The average rebate for refrigerators was \$128 (in 2010 dollars). There were 44 states that offered rebates for refrigerators and a total of 613,561 rebate claims. Houde & Aldy (2017) estimate that the percent of inframarginal rebate recipients for refrigerators (including those that only slightly delayed/accelerated consumption) was 92.0%.

To estimate the energy savings from a marginal beneficiary purchasing an ES-rated refrigerator, the authors report the difference between an ES and non-ES rated refrigerator in 2010 as 65 kWh per year. We use this number for the kWh reduction in years 6-15 of the refrigerators lifetime. For years 1-5, we compare the 2010 ES-rated refrigerator with a 2001 non-ES rated refrigerator. The paper does not directly report this number. It does report the ratio of the rebate amount to the total lifetime reduction. Using this ratio, we calculate the kWh difference for the first five years of the ES-rated appliance to be 207.6 kWh per year. We multiply the annual reductions by the percent of marginal rebate recipients (8%) to get the kWh change each year as a result of the subsidy.

Cost The total cost is comprised of the direct rebate cost and fiscal externalities. The average rebate for refrigerators was \$128 in 2010 dollars.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2. The fiscal externality for electricity per kWh is \$0.006 in 2020 and \$0.005 in 2010. Using the annual kWh reduction per rebate, discounting over the 15 years of the appliance lifetime, we arrive at a total fiscal externality in the US in 2020 of \$0.593 and in 2010 of \$0.434.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government cost by \$0.236 in 2020 and \$0.173 in-context. The resulting total cost is \$152.29 in the baseline and \$128.26 in-context.

WTP The willingness to pay is comprised of the marginal and inframarginal benefits, environmental externality, and loss in producer profits.

92.0% of the beneficiaries of the policy are inframarginal. Since we assume that inframarginal households value 100% of the \$128 rebate and marginal households value 50% of the \$128 rebate, the willingness to pay for the transfer is \$122.88 in-context. Inflation adjusting to 2020 dollars, the baseline willingness to pay is \$145.85.

The local and global environmental externality per kWh of electricity is explained in Appendix C.2. The resulting local and global per-kWh environmental externality for the US in 2020 is \$0.133 and \$0.026, respectively. The corresponding externality in 2010 is \$0.100 and \$0.081. While these are point in time estimates, we allow the electricity grid and social costs to change over the 15 years of the appliance lifetime. Using the annual changes in energy consumption and discounting over 15 years, the global environmental externality is \$15.32 in 2020 and \$11.24 in-context. The local environmental externality is \$2.25 in 2020 and \$5.17 in-context. The rebound effect, as explained in Appendix D, offsets 20% of the local and global environmental benefits. The resulting rebound effect is \$3.44 in 2020 and \$3.21 in-context.

Reduced energy consumption as a result of the ES-rated refrigerator leads to lower profits for electric utilities, as explained in Appendix C.2. The loss in profits per kWh of electricity is \$0.011 in 2020 and \$0.010 in 2010. Discounting over the lifetime of the appliance, we arrive at a total producer willingness to pay of -\$1.09 in the baseline specification and -\$0.80 in-context. Summing across these components, the total willingness to pay in 2020 is \$158.66 and in-context is \$135.10. This results in a baseline MVPF of 1.04 and in-context MVPF of 1.05.

E.5.2 Energy Star Rebates - Datta & Gulati (2014)

Datta & Gulati (2014) estimate the impact of rebates for Energy Star appliances on consumer demand. They separately identify the impact of a \$1 increase in utility rebate levels on the demand for clothes washers, dishwashers, and refrigerators. While Houde & Aldy (2017) study the cash for appliances program enacted in 2009, Datta & Gulati (2014) evaluate energy star rebates in place from 2001 to 2006. They leverage the variation in size and timing of the rebates across states to estimate the effect of the rebate on consumption.

Datta & Gulati (2014) report the percentage change in the consumption of each of the three appliances with respect to a \$1 increase in each of the appliances' rebate levels. This value allows us to calculate a price elasticity for each appliance. For the other appliance rebates in our sample, we estimate the MVPF of the average dollar spent on the rebate using estimates of the share of inframarginal beneficiaries and assuming that marginal beneficiaries value 50% of the rebate. For the rebates evaluated in Datta & Gulati (2014), we estimate the MVPF of a \$1 expansion of existing rebate.

We take the standard MVPF approach as outlined in equation 9 and as implemented in wind, solar, EVs, and other policy categories. We convert the semi-elasticity reported in the paper to an elasticity by multiplying by the retail cost of each appliance. For each externality V , we divide by the cost of each appliance p , to get the externality per dollar spent on the good (V/p). We multiply this by the elasticity to get the externality per dollar of government spending. To harmonize with Houde & Aldy (2017), we assume that these appliances have a 15-year lifetime.

The clothes washer, dish washer, and refrigerator rebates were implemented across 19, 12, and 14 states, respectively. Therefore, we will use values for the entire U.S. for both the in-context and baseline specifications. The rebates were implemented from 2001 to 2006. Due to data limitations for certain externality values in 2001, we will use externality values from 2006

for the in-context MVPF and 2020 for the baseline MVPF.

Clothes Washers

Our MVPF for Energy Star rebates for clothes washers using estimates from Datta & Gulati (2014) is 1.310 [1.134 , 1.440] and 2.126 in-context. They report that the average rebate for clothes washers in their sample was \$68.65. They also report that the cost of an Energy Star clothes washer is \$966. We assume that these values are in nominal dollars from the middle of the paper's sample (2004). To get the in-context levels, we inflation-adjust these value to 2006 dollars. To get the baseline levels, we inflation adjust to 2020 dollars. Therefore, the in-context cost of the appliance net of the rebate is \$950.

Datta & Gulati (2014) report that a \$1 increase in rebate levels leads to a 0.395% increase in the share of clothes washers. Multiplying by the net cost of a clothes washer in-context, we arrive at a price elasticity of -3.78.

For environmental externalities, we use the annual difference in kWh usage between Energy Star and non-Energy Star clothes washers. Datta & Gulati (2014) report that ES-rated clothes washers in 2006 use 531 kWh compared to 234 for the average non ES-rated clothes washer. This leads to an annual in-context difference of 297 kWh. For the baseline, we use the kWh difference from Houde & Aldy (2017) of 201 kWh since this value is estimated closer to 2020.

Cost The total cost is comprised of the mechanical transfer to consumers and fiscal externalities. There is a \$1 mechanical transfer cost to inframarginal consumers as a result of expanding the rebate level by \$1.

Since there is a pre-existing subsidy of \$68.65 (2004 dollars), there will be a fiscal externality as a result of induced demand from the rebate expansion. The fiscal externality divided by the cost of the appliance is the externality per dollar of spending on clothes washers. This value multiplied by the elasticity of 3.78 results in a fiscal externality of \$0.289 in 2020 and \$0.289 in-context.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2. The fiscal externality for electricity per kWh is \$0.005 in 2020 and \$0.005 in 2006. Using the annual kWh reduction per rebate, discounting over the 15 years of the appliance lifetime, and dividing by the cost of a clothes washer, we get the fiscal externality per dollar of spending on the good (V/p). We multiply this value by the elasticity to get the fiscal externality per dollar of government spending on the rebate. The fiscal externality is \$0.059 in-context and \$0.039 in 2020.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government cost by \$0.014 in 2020 and \$0.023 in-context. The resulting total cost is \$1.315 in the baseline and \$1.325 in-context.

WTP The willingness to pay is comprised of the inframarginal benefits, environmental externality, and loss in producer profits. The mechanical \$1 transfer is fully valued by consumers.

The local and global environmental externality per kWh of electricity is explained in Appendix C.2. The resulting local and global per-kWh environmental externality for the US in 2020 is \$0.026 and \$0.133, respectively. The corresponding externality in 2006 is \$0.102 and \$0.083. While these are point in time estimates, we allow the electricity grid and social costs to change over the 15 years of the appliance lifetime. Using the annual changes in energy consumption, discounting over 15 years, and dividing by the appliance cost, we get the local

and global environmental externality per dollar of spending on the good (V/p). We multiply these values by the elasticity to get the environmental externality per dollar of government spending. We exclude 1.9% of the global environmental benefits that flow to the government. The resulting global environmental externality is \$0.861 in 2020 and \$1.458 in-context. The local environmental externality is \$0.126 in 2020 and \$0.935 in-context. The rebound effect, as explained in Appendix D, offsets 20% of the local and global environmental benefits. The rebound effect is \$0.193 in 2020 and \$0.469 in-context.

Reduced energy consumption as a result of the ES-rated clothes washer leads to lower profits for electric utilities, as explained in Appendix C.2. The loss in profits per kWh of electricity is \$0.011 in 2020 and \$0.010 in 2006. Discounting over the lifetime of the appliance and dividing by the appliance cost, we get the profit loss per dollar of spending on the good. To convert this to the profit loss per dollar of government spending, we multiply by the elasticity. The producer willingness to pay is -\$0.072 in the baseline specification and -\$0.108 in-context. Summing across these components, the total willingness to pay in 2020 is \$1.722 and in-context is \$2.816. This results in the baseline MVPF of 2.13 and in-context MVPF of 1.31.

Dishwashers

Our MVPF for Energy Star rebates for dishwashers using estimates from Datta & Gulati (2014) is 1.053 [.988,1.200] and 0.883 in-context. They report that the average rebate for dishwashers in their sample was \$34.35 (2004 dollars). They do not report the cost of an energy star dish washer. We get the appliance cost of \$627 in 2010 dollars from Houde & Aldy (2017) and inflation adjust to 2006 dollars for the in-context cost and 2020 dollars for the baseline cost. Therefore, the in-context cost of the appliance net of the rebate is \$543.

Datta & Gulati (2014) report that a \$1 increase in rebate levels leads to a 0.6% decrease in the share of dishwashers. While this result is not statistically significant, it suggests a positive price elasticity. Multiplying by the net cost of a dishwasher in-context, we arrive at a price elasticity of 3.28. Since consumption decreases as a result of the rebate, the environmental externality will be negative and the producer willingness to pay will be positive.

For environmental externalities, we use the annual difference in kWh usage between Energy Star and non-Energy Star dishwashers. Datta & Gulati (2014) do not report this difference. We use the kWh difference from Houde & Aldy (2017) of 34 kWh for both the in-context and baseline specification.

Cost The total cost is comprised of the mechanical transfer to consumers and fiscal externalities. There is a \$1 mechanical transfer cost to inframarginal consumers as a result of expanding the rebate level by \$1.

Since there is a pre-existing subsidy of \$34.35 (2004 dollars), there will be a fiscal externality as a result of induced demand from the rebate expansion. The fiscal externality divided by the cost of the appliance is the externality per dollar of spending on dishwashers. This value multiplied by the elasticity of 3.28 results in a fiscal externality of -\$0.221 in 2020 and -\$0.221 in-context. This value is negative because the rebate reduces demand.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2. The fiscal externality for electricity per kWh is \$0.005 in 2020 and \$0.005 in 2006. Using the annual kWh reduction per rebate, discounting over the 15 years of the appliance lifetime, and dividing by the cost of a clothes washer, we get the fiscal externality per dollar of spending on the good (V/p). We multiply this value by

the elasticity to get the fiscal externality per dollar of government spending on the rebate. The fiscal externality is $-\$0.010$ in-context and $-\$0.010$ in 2020.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality increases government cost (since the environmental benefits are negative) by $\$0.003$ in 2020 and $\$0.004$ in-context. The resulting total cost is $\$0.772$ in the baseline and $\$0.772$ in-context.

WTP The willingness to pay is comprised of the inframarginal benefits, environmental externality, and loss in producer profits. The mechanical $\$1$ transfer is fully valued by consumers.

The local and global environmental externality per kWh of electricity is explained in Appendix C.2. The resulting local and global per-kWh environmental externality for the US in 2020 is $\$0.026$ and $\$0.133$, respectively. The corresponding externality in 2006 is $\$0.102$ and $\$0.083$. While these are point in time estimates, we allow the electricity grid and social costs to change over the 15 years of the appliance lifetime. Using the annual changes in energy consumption, discounting over 15 years, and dividing by the appliance cost, we get the local and global environmental externality per dollar of spending on the good (V/p). We multiply these values by the elasticity to get the environmental externality per dollar of government spending. We exclude 1.9% of the global environmental benefits that flow to the government. The resulting global environmental externality is $-\$0.223$ in 2020 and $-\$0.255$ in-context. The local environmental externality is $-\$0.033$ in 2020 and $-\$0.164$ in-context. These values are negative since the rebates decreases the consumption of energy efficient dishwashers. The rebound effect, as explained in Appendix D, offsets 20% of the local and global environmental benefits. The rebound effect is $\$0.050$ in 2020 and $\$0.082$ in-context.

Reduced energy consumption as a result of the ES-rated dishwasher leads to lower profits for electric utilities, as explained in Appendix C.2. However, in the case of dishwasher rebates, profits increase for utilities due to the decrease in efficient dishwasher purchases. The increase in profits per kWh of electricity is $\$0.011$ in 2020 and $\$0.010$ in 2006. Discounting over the lifetime of the appliance and dividing by the appliance cost, we get the profit loss per dollar of spending on the good. To convert this to the profit gain per dollar of government spending, we multiply by the elasticity. The producer willingness to pay is $\$0.019$ in the baseline specification and $\$0.019$ in-context. Summing across these components, the total willingness to pay in 2020 is $\$0.813$ and in-context is $\$0.682$. This results in the baseline MVPF of 1.05 and in-context MVPF of 0.88.

Refrigerators

Our MVPF for Energy Star rebates for refrigerators using estimates from Datta & Gulati (2014) is 1.011 [1.000,1.020] and 1.113 in-context. They report that the average rebate for refrigerators in their sample was $\$49.06$ (2004 dollars). They do not report the cost of an energy star dishwasher. We get the appliance cost of $\$1,240$ in 2010 dollars from Houde & Aldy (2017) and inflation adjust to 2006 dollars for the in-context cost and 2020 dollars for the baseline cost. Therefore, the in-context cost of the appliance net of the rebate is $\$1,094$.

Datta & Gulati (2014) report that a $\$1$ increase in rebate levels leads to a 0.282% increase in the share of ES-rated refrigerators. Multiplying by the net cost of a refrigerators in-context, we arrive at a price elasticity of -3.08 .

For environmental externalities, we use the annual difference in kWh usage between Energy Star and non-Energy Star refrigerators. Datta & Gulati (2014) do not report this difference.

We use the kWh difference from Houde & Aldy (2017) of 65 kWh for both the in-context and baseline specification.

Cost The total cost is comprised of the mechanical transfer to consumers and fiscal externalities. There is a \$1 mechanical transfer cost to inframarginal consumers as a result of expanding the rebate level by \$1.

Since there is a pre-existing subsidy of \$49.06 (2004 dollars), there will be a fiscal externality as a result of induced demand from the rebate expansion. The fiscal externality divided by the cost of the appliance is the externality per dollar of spending on refrigerators. This value multiplied by the elasticity of 3.08 results in a fiscal externality of \$0.148 in 2020 and \$0.148 in-context.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2. The fiscal externality for electricity per kWh is \$0.005 in 2020 and \$0.005 in 2006. Using the annual kWh reduction per rebate, discounting over the 15 years of the appliance lifetime, and dividing by the cost of a refrigerators, we get the fiscal externality per dollar of spending on the good (V/p). We multiply this value by the elasticity to get the fiscal externality per dollar of government spending on the rebate. The fiscal externality is \$0.009 in-context and \$0.009 in 2020.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government cost by \$0.003 in 2020 and \$0.004 in-context. The resulting total cost is \$1.154 in the baseline and \$1.153 in-context.

WTP The willingness to pay is comprised of the inframarginal benefits, environmental externality, and loss in producer profits. The mechanical \$1 transfer is fully valued by consumers.

The local and global environmental externality per kWh of electricity is explained in Appendix C.2. The resulting local and global per-kWh environmental externality for the US in 2020 is \$0.026 and \$0.133, respectively. The corresponding externality in 2006 is \$0.102 and \$0.083. While these are point in time estimates, we allow the electricity grid and social costs to change over the 15 years of the appliance lifetime. Using the annual changes in energy consumption, discounting over 15 years, and dividing by the appliance cost, we get the local and global environmental externality per dollar of spending on the good (V/p). We multiply these values by the elasticity to get the environmental externality per dollar of government spending. We exclude 1.9% of the global environmental benefits that flow to the government. The resulting global environmental externality is \$0.199 in 2020 and \$0.228 in-context. The local environmental externality is \$0.029 in 2020 and \$0.146 in-context. The rebound effect, as explained in Appendix D, offsets 20% of the local and global environmental benefits. The rebound effect is \$0.045 in 2020 and \$0.073 in-context.

Reduced energy consumption as a result of the ES-rated refrigerator leads to lower profits for electric utilities, as explained in Appendix C.2. The loss in profits per kWh of electricity is \$0.011 in 2020 and \$0.010 in 2006. Discounting over the lifetime of the appliance and dividing by the appliance cost, we get the profit loss per dollar of spending on the good. To convert this to the profit loss per dollar of government spending, we multiply by the elasticity. The producer willingness to pay is -\$0.017 in the baseline specification and -\$0.017 in-context. Summing across these components, the total willingness to pay in 2020 is \$1.167 and in-context is \$1.284. This results in the baseline MVPF of 1.01 and in-context MVPF of 1.11.

E.5.3 California Energy Savings Assistance Program - Refrigerators

Our MVPF for refrigerator replacements using estimates from Blonz (2023) is 0.96 [0.93 , 0.99] in 2020 and 0.57 in-context. Blonz (2023) estimates the change in energy consumption from refrigerator replacements in California.

Blonz (2023) uses data from 2009-2012 from the California Energy Savings Assistance (ESA) program. Among other appliances and retrofits, the ESA provides low-income households energy efficient refrigerator replacements if their existing fridge meets the eligibility criteria. The paper finds that the contractors installing the new fridge intentionally misreported the percentage of fridges that met the eligibility criteria in order to increase their compensation.

Blonz (2023) finds that 3,715 replacements were for qualified refrigerators compared to 1,261 for unqualified refrigerators. Therefore, about 75% of the replacements were for qualified fridges. The paper also finds that the people who qualified accelerated their replacement decisions by five years and those who should not have qualified accelerated their replacement decision six years. During this window, the paper estimates that the qualified refrigerators saved 73.45 kWh per month and the unqualified refrigerators saved 38.02 kWh per month. Since the paper estimates the average change in purchase timing across all the beneficiaries, we assume that everyone is marginal to the policy and changes their decision by either 5 or 6 years depending on whether they should have qualified for the replacement. Consistent with the other appliance rebate policies, we assume that these appliances have a 15-year lifetime.

The program paid contractors \$850 per fridge replacement. We assume that all values reported in the paper are in 2010 dollars (the middle of the sample). For the in-context MVPF, we use externality values from California in 2009, the first year of the program. For the baseline MVPF, we use values from the US in 2020.

Cost The total cost is comprised of the direct subsidy and fiscal externalities. The program paid \$850 (2010 dollars) per refrigerator replacement. We inflation-adjust this cost to 2009 dollars for the in-context MVPF and 2020 dollars for the baseline MVPF.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2. The fiscal externality for electricity per kWh is \$0.006 in 2020 and \$0.022 in California in 2009. For the qualified refrigerators, we multiply the annual kWh reduction per qualified fridge by the fiscal externality per kWh for each of the five years, and discount over those years to get the total fiscal externality for the qualified fridges. We repeat this calculation for the unqualified fridges. We take an average of the fiscal externality for qualified and unqualified fridges weighted by their respective shares to get the fiscal externality per replacement. The fiscal externality is \$18.92 in 2020 and \$68.36 in-context.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government cost by \$8.706 in 2020 and \$3.99 in-context. The resulting total cost is \$1035.63 in the baseline and \$914.37 in-context.

WTP The willingness to pay is comprised of the marginal and inframarginal benefits, environmental externality, and loss in producer profits.

We assume that everyone is marginal to the policy and delays their replacement decision by five or six years depending on whether they qualify for the replacement. Consistent with the other appliance rebate policies, we assume that marginal beneficiaries value 50% of the rebate. Therefore the willingness to pay for the rebate is \$425.00 in-context and \$512.71 in 2020. The

in-context value is in 2009 dollars and the baseline value is in 2020 dollars.

The local and global environmental externality per kWh of electricity is explained in Appendix C.2. The resulting local and global per-kWh environmental externality for the US in 2020 is \$0.026 and \$0.133, respectively. The corresponding externality for California in 2009 is \$0.006 and \$0.064. While these are point in time estimates, we allow the electricity grid and social costs to change over time. For the qualified refrigerators, we multiply the annual kWh reduction per qualified fridge by the environmental externality per kWh for each of the five years, and discount over those years to get the total environmental externality for the qualified fridges. We repeat this calculation for the unqualified fridges. We take an average of the environmental externality for qualified and unqualified fridges weighted by their respective shares to get the local and global environmental externality. The global environmental externality is \$554.37 in 2020 and \$253.89 in-context. The local environmental externality is \$85.56 in 2020 and \$24.97 in-context. The rebound effect, as explained in Appendix D, offsets 20% of the local and global environmental benefits. The resulting rebound effect is \$125.34 in 2020 and \$54.62 in-context.

Reduced energy consumption as a result of the energy efficient refrigerator leads to lower profits for electric utilities, as explained in Appendix C.2. The loss in profits per kWh of electricity is \$0.011 in 2020 and \$0.040 in California in 2009. Using the same calculation as the environmental externality, we arrive at a total producer willingness to pay of -\$34.83 in the baseline and -\$125.84 in-context. Summing across these components, the total willingness to pay in 2020 is \$992.47 and in-context is \$523.41. This results in a baseline MVPF of 0.958 and in-context MVPF of 0.572.

E.5.4 Energy Star Water Heater Rebates

Our MVPF for water heater rebates using estimates from Allcott & Sweeney (2017) is 1.340 [1.250 , 1.367] in 2020 and 0.998 in-context.

Partnering with a water heater retailer, Allcott & Sweeney (2017) run a natural field experiment where they randomize customers into multiple treatment arms. The MVPF explained in this section focuses on the treatment group that receives a \$100 rebate. In the nudge section, we also construct an MVPF for the treatment arm in which sales agents receive a \$25 incentive for each energy efficient water heater they sell. In the field experiment, sales agents called potential customers who were randomized into a control or treatment group. Those in the treatment group were offered a \$100 subsidy for Energy Star water heaters.

Among those in the control group, there was a 0.9% chance they purchased an Energy Star water heater. The \$100 rebate increased the purchase probability by 3.7 percentage points. Therefore, roughly 20% of the beneficiaries of the rebate are inframarginal ($0.9/(0.9 + 3.7)$) For the marginal beneficiaries, the water heater rebate leads to reductions in energy usage.

While water heaters can be either natural gas or electric, the water heaters offered during the field experiment were all natural gas. The EIA estimates that the average water heater in a four person household uses 22.7 MMBtu of natural gas (EIA 2018). An Energy Star water heater uses 8% less energy than a standard model (DOE 2024b). Therefore, we estimate that an ES-rated water heater saves 1.816 MMBtu per year. Multiplying this by the proportion of marginal recipients, we get a reduction of 1.46 MMBtu per rebate per year. Consistent with the other appliance rebate MVPFs in our sample, we assume a lifetime of 15 years.

The experiment ran from 2012 to 2014. We assume that all values reported in the paper are in 2013 dollars (the middle of the sample). For the in-context MVPF, we use externality values for the US in 2012, the first year of the program. For the baseline MVPF, we use values from the US in 2020.

Cost The total cost is comprised of the direct rebate cost and fiscal externalities. The rebate for water heaters was \$100 in 2012 dollars.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2. The fiscal externality for natural gas per MMBtu is \$0.75 in 2020 and \$0.76 in 2012. Using the annual MMBtu reduction per rebate, discounting over the 15 years of the appliance lifetime, we arrive at a total fiscal externality in the US in 2020 of \$12.59 and in 2010 of \$12.88.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government cost by \$3.76 in 2020 and \$2.79 in-context. The resulting total cost is \$121.58 in the baseline and \$110.09 in-context.

WTP The willingness to pay is comprised of the marginal and inframarginal benefits, environmental externality, and loss in producer profits.

The 20% of inframarginal beneficiaries value the entire \$100 rebate. Consistent with the other appliance rebate MVPFs, we assume that the 80% of marginal beneficiaries value 50% of the \$100 rebate. Therefore, the total consumer willingness to pay is \$59.78 in-context and \$67.41 in 2020 (\$100 rebate inflation adjusted to 2012 and 2020 dollars, respectively).

The environmental externality per MMBtu of natural gas is explained in Appendix C.3. The externality per MMBtu in 2020 is \$10.25 and in 2012 is \$7.61. Using the annual changes in energy consumption and discounting over 15 years, the environmental externality per rebate is \$192.43 in 2020 and \$142.93 in-context. The rebound effect, as explained in Appendix D, offsets 20% of the environmental benefits. The resulting rebound effect is \$22.64 in 2020 and \$16.82 in-context.

Reduced energy consumption as a result of the ES-rated water heater leads to lower profits for utilities, as explained in Appendix C.3. The loss in profits per MMBtu of natural gas is \$4.40 in 2020 and \$4.50 in 2012. Discounting over the lifetime of the appliance, we arrive at a total producer willingness to pay of -\$74.26 in the baseline specification and -\$75.98 in-context. Summing across these components, the total willingness to pay in 2020 is \$162.92 and in-context is \$109.92. This results in a baseline MVPF of 1.340 and in-context MVPF of 0.998.

E.6 Weatherization

Our category average MVPF for weatherization programs is 0.98. This appendix describes the construction of the individual MVPFs that feed into this category average.

Weatherization programs are intended to improve the energy efficiency of residential, commercial, and industrial buildings. These programs typically involve measures such as insulation, air sealing, HVAC system upgrades, and window and door improvements. Such programs are implemented by governmental agencies, nonprofit organizations, and utility companies and often target low-income households. The Inflation Reduction Act includes \$8.8 billion for weatherization programs of which approximately 50% is for whole-home energy upgrades and the other 50% is allocated for appliance and efficient electric technology rebates (DOE n.d.).

We create MVPFs for state-level weatherization policies implemented in Michigan, Illinois, Arizona, Wisconsin, and Florida. Each policy focuses on different types of retrofits for households with varying baseline energy usage. We take the average treatment effect on energy usage, retrofit cost, subsidy level, and baseline energy usage from each paper. We do not harmonize these measures across policies because we believe the papers' treatment effect is dependent on the retrofit cost and baseline energy usage. Our baseline MVPFs use environmental externalities and producer profit values corresponding to the US in 2020.

The willingness to pay for weatherization consists of the mechanical transfer to households, environmental externality, rebound effect, and effect on producer profits. The total cost for each program is the sum of the average subsidy level, fiscal externality from the change in utility profit tax revenue, and the climate fiscal externality.

The papers in our sample do not observe the counterfactual take-up of weatherization in the absence of weatherization subsidies. Therefore, we do not have an empirical estimate of the share of marginal beneficiaries. Our baseline MVPF assumes that 50% of households are marginal to the subsidy. For the marginal households, some are convinced to take up the subsidy by the first few dollars and some are only convinced by the last dollar. We assume a uniform distribution over the potential threshold subsidy at which people would do the retrofit, resulting in marginal households valuing the subsidy at 50%. Inframarginal households value the entire subsidy. Therefore, a \$1 mechanical transfer will lead to \$0.50 of benefits for inframarginal households and \$0.25 for marginal households. If we assumed a marginal fraction of 0% the MVPF is 1 by construction and with an assumed marginal fraction of 100% the category average MVPF is 0.97.

The externalities included in the MVPF are only from the 50% of households that are induced to take up weatherization. The environmental externality is calculated as the product of the treatment effect, baseline usage, proportion of marginal households, and the monetized environmental externality per kWh. The effect on producer profits has an analogous calculation. The environmental externality per kWh in 2020 is \$0.16 and the producer profit level for electricity in 2020 is \$0.01. Some weatherization programs also affect natural gas usage. The externalities for natural gas are explained in Appendix C.3 and result in an environmental and producer profits externality per MMBtu of \$10.25 and \$4.40, respectively.

We construct both in-context MVPFs using externalities from the year and state the policy was implemented in as well as baseline MVPFs for the US in 2020. For ease of interpretation, the numbers referenced in each policy are in terms of the cost reported in the paper (generally per household). To crosswalk the MVPF component numbers with those in Table 2, one can divide each component by the mechanical spending on weatherization.

E.6.1 Michigan Weatherization Assistance Program

Our MVPF for weatherization using estimates from Fowlie et al. (2018) is 0.92 [0.82, 1.05] in 2020 and 0.96 in-context. Fowlie et al. (2018) conducts a large-scale randomized control experiment of the Weatherization Assistance Program (WAP) on 30,000 households in Michigan. WAP is a federal program run by the US Department of Energy. It is the largest energy efficiency program in the country, assisting over 7 million households since it began in 1976. WAP does not provide funding for energy efficiency upgrades unless it passes a cost-benefit analysis from engineering predictions.

This paper studies energy efficiency investments in Michigan between 2011-2014 - a period

in which WAP funding was significantly increased as a result of the American Recovery and Reinvestment Act. All owner-occupied households at or below 200% of the poverty line were eligible to apply for assistance. The most common energy upgrades included furnace replacement, attic and wall insulation, and infiltration reduction.

The paper uses a randomized encouragement treatment in which treated households are encouraged to apply for the weatherization program through intensive communication and marketing. Using treatment status as an instrument, Fowlie et al. (2018) estimate the per household energy reduction caused by the weatherization program.

The average household in the paper's sample uses 76.68 MMBtu of natural gas and 7490.90 kWh of electricity annually. The paper's main specification estimates that weatherization reduces natural gas consumption by 18.9% and electricity consumption by 9.5%. This translates into an annual 712.85 kWh and 14.52 MMBtu reduction. Fowlie et al. (2018) presents their results for weatherization lifetimes of 10, 16, and 20 years. Our baseline MVPF uses a 20-year lifetime. The in-context MVPF studies the policy in 2011, the first year of the paper's sample.

Cost The total cost is comprised of the direct program cost and fiscal externalities. Since the MVPF measures the effectiveness of the weatherization program and not the effectiveness of the encouragement nudge, the program cost does not include the cost of the encouragement treatment. We do provide an MVPF for the encouragement nudge (0.29) explained in a forthcoming policy appendix. Fowlie et al. (2018) find that the average cost of the energy upgrade per household was \$5,150 in 2011 dollars.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2 (electricity) and Appendix C.3 (natural gas). The fiscal externality for electricity per kWh is \$0.006 in 2020 and \$0.01 in Michigan in 2011. The fiscal externality for natural gas per MMBtu is \$0.75 in 2020 and \$0.52 in Michigan in 2011. Using the annual 15 MMBtu reduction in natural gas and annual 717.9 kWh reduction in electricity, discounting over the lifetime of the weatherization, we arrive at a total fiscal externality in the US in 2020 of \$108.32 and for Michigan in 2011 of \$111.38.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government cost by \$29.53 in 2020 and \$26.07 in-context. The resulting total cost is \$6,005.52 in the baseline and \$5,235.31 in-context.

WTP The willingness to pay is comprised of the marginal and inframarginal benefits, environmental externality, and loss in producer profits.

We assume that half of the beneficiaries are marginal and the other half are inframarginal. Since we assume that inframarginal households value 100% of the subsidy and marginal households value 50% of the subsidy, the willingness to pay for the transfer is \$3,862.50 in-context. Inflation adjusting to 2020 dollars, the baseline willingness to pay is \$4,445.04 or 75% of the subsidy.

Consistent with other policy categories, we split the environmental externality into a global and local component. The environmental externality per kWh of electricity and per MMBtu of natural gas, and their local and global sub-components, are explained in Appendices C.2 and C.3, respectively. The resulting per-kWh environmental externality for the US in 2020 is \$0.16 and for Michigan in 2011 is \$0.24. While these are point-in-time estimates, we allow the electricity grid and social costs to change over the 20 years of weatherization benefits. For natural gas, the externalities per MMBtu are \$10.25 in 2020 and \$7.28 in context. Using the annual change in energy consumption and discounting over 20 years, the global environmental

externality is \$1,761.88 in 2020 and \$1,575.96 in context. The local environmental externality is \$76.42 in 2020 and \$291.74 in context. The rebound effect offsets approximately 20% of environmental benefits from electricity and 12% of the environmental benefits from natural gas. The resulting rebound effect is -\$264.84 in 2020 and -\$298.15 in context.

Reduced energy consumption as a result of weatherization leads to lower profits for electric and natural gas utilities. The construction of the producer profits externality is explained in Appendix C.2 (electricity) and Appendix C.3 (natural gas). The loss in profits per kWh of electricity is \$0.01 in 2020 and \$0.02 in Michigan in 2011. The loss in profits per MMBtu of natural gas is \$4.40 in 2020 and \$3.06 in Michigan in 2011. Using the annual reduction in electricity and natural gas, discounting over the lifetime of the weatherization, we arrive at a total producer willingness to pay of -\$522.57 in the baseline specification and -\$430.26 in-context. Summing across these components, the total willingness to pay in 2020 is \$5,495.94 and in-context is \$5,001.78. This results in a baseline MVPF of \$0.92 and in-context MVPF of \$0.96.¹⁷³

E.6.2 Illinois Home Weatherization Assistance Program

Our MVPF for weatherization using estimates from Christensen, Francisco & Myers (2023) is 0.98 [0.96, 1.00] in 2020 and 1.05 in context. Christensen, Francisco & Myers (2023) studies the Illinois Home Weatherization Assistance Program (IHWAP). IHWAP uses funding from the federal Weatherization Assistance Program which provides rebates to low-income households for dwelling upgrades (e.g., insulation, appliance replacements) and repairs aimed at boosting energy efficiency. Households were eligible provided their incomes were less than 200 percent of the national poverty line. Households qualifying for other social assistance programs (e.g., Low Income Home Energy Assistance Program (LIHEAP), households with members receiving Security Disability (SSD), Supplemental Security Income (SSI) or Temporary Assistance for Needy Families (TANF)) were also eligible.

Christensen, Francisco & Myers (2023) use data from households who received upgrades from 2018 to 2019 through IHWAP. They use an event study fixed effects model to estimate the impact of weatherization on energy usage. The paper also studies the impact of performance incentives for contractors who are performing the weatherization. The MVPF for these incentives is 1.07 for the high incentive and 1.06 for the low incentive. These MVPFs are further explained in the nudge and marketing policy appendix. The IHWAP MVPF focuses exclusively on weatherization and excludes the benefits and costs from the performance incentive.

Following the approach in Christensen, Francisco & Myers (2023), we use a 34-year lifetime for the weatherization benefits. Our estimate of the MVPF is within the range of MVPFs reported in the paper.¹⁷⁴

The paper estimates the monthly change in electricity and natural gas consumption. Converting these estimates to annual changes, the average household in their sample reduces annual electricity consumption by 1656.44 kWh and annual natural gas consumption by 19.48 MMBtu. The in-context MVPF studies the policy in 2018, the first year of the paper's sample.

¹⁷³If we were to instead assume that 100% of the beneficiaries are marginal, then the baseline MVPF would be 0.83

¹⁷⁴Christensen, Francisco & Myers (2023) estimate MVPFs for weatherization of 0.72, 0.95, and 1.14 corresponding to SCCs of \$51, \$125, and \$185. The main difference between our calculation and theirs is that they assume all beneficiaries are marginal and do not include a rebound effect for electricity and natural gas.

Cost The total cost is comprised of the direct program cost and fiscal externalities. Christensen, Francisco & Myers (2023) reports that the average cost of the energy upgrade per household was \$9,655 in 2017 dollars.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2 (electricity) and Appendix C.3 (natural gas). The fiscal externality for electricity per kWh is \$0.006 in 2020 and is zero in Illinois in 2009. The fiscal externality for natural gas per MMBtu is \$0.52 in 2020 and \$0.38 in Illinois in 2018. Using the annual reduction in natural gas and electricity, discounting over the lifetime of the weatherization, we arrive at a total fiscal externality in the US in 2020 of \$259.79 and for Illinois in 2018 of \$123.06.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government costs by \$68.42 in 2020 and \$65.23 in context. The resulting total cost is \$10,386.98 in the baseline and \$9,948.34 in context.

WTP The willingness to pay is comprised of the marginal and inframarginal benefits, environmental externality, and loss in producer profits.

We assume that half of the beneficiaries are marginal and the other half are inframarginal. Since we assume that inframarginal households value 100% of the subsidy and marginal households value 50% of the subsidy, the willingness to pay for the transfer is \$7,417.88 in-context. Inflation adjusting to 2020 dollars, the baseline willingness to pay is \$7,646.71 or 75% of the subsidy amount.

Consistent with other policy categories, we split the environmental externality into a global and local component. The environmental externality per kWh of electricity and per MMBtu of natural gas, and their local and global sub-components, are explained in Appendix C.2 and C.3, respectively. The resulting per-kWh environmental externality for the US in 2020 is \$0.16 and for Illinois in 2018 is \$0.21. While these are point-in-time estimates, we allow the electricity grid and social costs to change over the 34 years of weatherization benefits. For natural gas, the externalities per MMBtu are \$10.25 in 2020 and \$9.51 in context. Using the annual change in energy consumption and discounting over 34 years, the global environmental externality is \$4,119.10 in 2020 and \$3,932.84 in context. The local environmental externality is \$196.86 in 2020 and \$475.10 in context. The rebound effect, as explained in Section D, offsets approximately 20% of environmental benefits from electricity and 12% of the environmental benefits from natural gas. The resulting rebound effect is -\$654.04 in 2020 and -\$685.90 in context.

Reduced energy consumption as a result of weatherization leads to lower profits for electric and natural gas utilities. The construction of the producer profits externality is explained in Appendix C.2 (electricity) and Appendix C.3 (natural gas). The loss in profits per kWh of electricity is \$0.01 in 2020 and zero in Illinois in 2018. The loss in profits per MMBtu of natural gas is \$4.40 in 2020 and \$3.38 in Illinois in 2018. Using the annual reduction in natural gas and electricity, discounting over the lifetime of the weatherization, we arrive at a total producer willingness to pay of -\$1,127.49 in the baseline specification and -\$725.63 in context. Summing across these components, the total willingness to pay in 2020 is \$10,181.14 and in context is \$10,414.29. This results in a baseline MVPF of 0.98 and an in-context MVPF of 1.05.¹⁷⁵

¹⁷⁵If we were to instead assume that 100% of the beneficiaries are marginal, then the baseline MVPF would be 0.96

E.6.3 Gainesville Regional Utility LEEP Plus Program

Our MVPF for weatherization using estimates from Hancevic & Sandoval (2022) is 0.86 [0.80, 0.92] in 2020 and 0.87 in context. Hancevic & Sandoval (2022) studies Gainesville, Florida’s Low-income Energy Efficiency Program Plus (LEEP Plus). Gainesville Regional Utilities (GRU), the fifth largest municipal electric utility company in Florida, established the LEEP Plus in 2007. This program helps low-income households in Gainesville, Florida, with home improvements to reduce electricity consumption. To be eligible, households must live in homes built before 1997 and have a family income lower than 80% of the metro area’s median income.

Hancevic & Sandoval (2022) use panel data from 2012 through 2018 for households that received an energy upgrade through GRU’s LEEP Plus. To estimate the causal impact of participation, the paper compares treated households that received an energy upgrade to control households that applied but were not selected to receive an upgrade. Households were untreated for a variety of reasons such as incomplete applications and incomes above the eligible cap. LEEP Plus focuses on retrofits that affect electricity usage and the paper finds that the program did not affect natural gas. Therefore, the MVPF focuses on the treatment effect on electricity consumption.

Using household and time fixed effects, the paper finds that treated households reduce electricity consumption relative to control households by 7.4% following the weatherization. The average electricity usage of the households in their sample was 9,965.5 kWh per year, implying a reduction of 706.9 kWh. The paper reports that the energy efficiency upgrades have a lifetime of 10-20 years. We assume a lifetime of 20 years in our MVPF calculations. The in-context MVPF studies the policy in 2012, the first year of the paper’s sample.

Cost The total cost is comprised of the direct program cost and fiscal externalities. Hancevic & Sandoval (2022) reports that the average cost of the energy upgrade per household was \$3,783.60 in 2018 dollars.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2. The fiscal externality for electricity per kWh is \$0.006 in 2020 and also \$0.006 in Florida in 2012. Using the annual 706.9 kWh reduction in electricity and discounting over the lifetime of the weatherization, we arrive at a total fiscal externality in the US in 2020 of \$28.40 and for Florida in 2012 of \$26.81.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government cost by \$8.48 in 2020 and \$7.54 in-context. The resulting total cost is \$3,920.24 in the baseline and \$3,478.70 in-context.

WTP The willingness to pay is comprised of the marginal and inframarginal benefits, environmental externality, and loss in producer profits.

We assume that half of the beneficiaries are marginal and the other half are inframarginal. Since we assume that inframarginal households value 100% of the subsidy and marginal households value 50% of the subsidy, the willingness to pay for the transfer is \$2,594.57 in-context. Inflation adjusting to 2020 dollars, the baseline willingness to pay is \$2,952.24.

Consistent with other policy categories, we split the environmental externality into a global and local component. The local and global environmental externality per kWh of electricity is explained in Appendix C.2. The resulting per-kWh environmental externality for the US in 2020 is \$0.16 and for Florida in 2012 is \$0.17. While these are point-in-time estimates, we allow the electricity grid and social costs to change over the 20 years of weatherization benefits. Using the

annual change in energy consumption and discounting over 20 years, the global environmental externality is \$539.92 in 2020 and \$480.26 in context. The local environmental externality is \$75.78 in 2020 and \$133.81 in context. The rebound effect, as explained in Appendix D, offsets 20% of the environmental benefits from electricity and 12% from natural gas. The resulting rebound effect is -\$120.60 in 2020 and -\$120.28 in context.

Reduced energy consumption as a result of weatherization leads to lower profits for electric utilities, as explained in Appendix C.2. The loss in profits per kWh of electricity is \$0.01 in 2020 and \$0.01 in Florida in 2012. Using the annual 706.9 kWh reduction in electricity, discounting over the lifetime of the weatherization, we arrive at a total producer willingness to pay of -\$52.28 in the baseline specification and -\$49.36 in context. Summing across these components, the total willingness to pay in 2020 is \$3,368.06 and in context is \$3,039.00. This results in a baseline MVPF of 0.86 and an in-context MVPF of 0.87.¹⁷⁶

E.6.4 Energize Phoenix Program - Residential Buildings

Our MVPF for weatherization using estimates from Liang et al. (2018) is 1.21 [0.93,1.43] in 2020 and 1.33 in-context. Liang et al. (2018) studies Energize Phoenix, a weatherization program that targeted buildings within a 10-mile radius of downtown Phoenix, Arizona. The program was in operation from 2010 to 2013 and had a goal of reducing energy consumption by 30% for residential buildings.

There were three subsidy programs for residential buildings that depended on household income: Energy Assist 60/40, Energy Assist 100%, and Rebate Match. Since cost data is only available for the 60/40 program, the MVPF for residential buildings is limited to this subsidy. Households were eligible for the Energy Assist 60/40 program if they had an annual income of less than 400% of the federal poverty level. This group received a subsidy that covers 60% of the upgrade costs.

Liang et al. (2018) estimate the average treatment effect of Energize Phoenix on residential electricity consumption using month and household fixed effects. In addition to the primary event study design, they validate their results with a difference-in-difference approach that compares treated households to those that applied for but did not receive the subsidy. They find that the program reduces electricity consumption by 26%. The average baseline annual electricity usage for the households in the 60/40 program before receiving energy upgrades was 14,349.60 kWh. This results in an annual reduction of approximately 3,740.39 kWh. The authors do not observe natural gas data, so they do not report changes in natural gas consumption. The in-context MVPF studies the policy in 2010, the first year of the paper's sample.

Cost The total cost is comprised of the direct program cost and fiscal externalities. Liang et al. (2018) reports the total retrofit cost of the program. Converting the total cost to a per household cost and accounting for the fact that the government is only subsidizing 60% of retrofit costs, the resulting per household subsidy is \$4,333.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2. The fiscal externality for electricity per kWh is \$0.006 in 2020 and \$0.003 in Arizona in 2010. Using the annual kWh reduction in electricity and discounting over the lifetime of the weatherization, we arrive at a total fiscal

¹⁷⁶If we were to instead assume that 100% of the beneficiaries are marginal, then the baseline MVPF would be 0.72

externality in the US in 2020 of \$150.28 and for Arizona in 2010 of \$79.16.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government cost by \$44.86 in 2020 and \$42.92 in-context. The resulting total cost is \$4,920.16 in the baseline and \$4,092.63 in-context.

WTP The willingness to pay is comprised of the marginal and inframarginal benefits, environmental externality, and loss in producer profits.

We assume that half of the beneficiaries are marginal and the other half are inframarginal. Since we assume that inframarginal households value 100% of the subsidy and marginal households value 50% of the subsidy, the willingness to pay for the transfer is \$3,042.30 in-context. Inflation adjusting to 2020 dollars, the baseline willingness to pay is \$3,611.06 or 75% of the subsidy.

Consistent with other policy categories, we split the environmental externality into a global and local component. The local and global environmental externality per kWh of electricity is explained in Appendix C.2. The resulting per-kWh environmental externality for the US in 2020 is \$0.16 and for Arizona in 2010 is \$0.10. While these are point-in-time estimates, we allow the electricity grid and social costs to change over the 20 years of weatherization benefits. Using the annual change in energy consumption and discounting over 20 years, the global environmental externality is \$2,856.83 in 2020 and \$2,733.04 in context. The local environmental externality is \$400.98 in 2020 and \$430.27 in context. The rebound effect, as explained in Appendix D, offsets 20% of the environmental benefits from electricity and 12% from natural gas. The resulting rebound effect is -\$638.09 in 2020 and -\$619.58 in context.

Reduced energy consumption as a result of weatherization leads to lower profits for electric utilities, as explained in Appendix C.2. The loss in profits per kWh of electricity is \$0.01 in 2020 and \$0.01 in Arizona in 2010. Using the annual 3,609.84 kWh reduction in electricity, discounting over the lifetime of the weatherization, we arrive at a total producer willingness to pay of -\$276.64 in the baseline specification and -\$145.72 in context. Summing across these components, the total willingness to pay in 2020 is \$5,954.13 and in context is \$5,440.30. This results in a baseline MVPF of 1.21 and an in-context MVPF of 1.33.¹⁷⁷

E.6.5 Wisconsin Energy Efficiency Retrofit Program

Our MVPF for weatherization using estimates from Allcott & Greenstone (2024) is 0.92 in 2020 and 0.93 in-context. Allcott & Greenstone (2024) study two home retrofit programs in Wisconsin: Green Madison and Milwaukee Energy Efficiency. They are both funded through the federal Better Buildings Neighborhood Program as part of the initial 2009 economic stimulus bill. The program took place from 2010 to 2013.

Households were randomized into two treatment groups and a control group. The treatment group received additional subsidies for home energy audits, the first stage of the weatherization process. Allcott & Greenstone (2024) find that while the audit subsidies increased take-up of audits, it had a small insignificant impact on households' decisions to invest in weatherization. Allcott & Greenstone (2024) combine this experimental variation with observational variation in household energy use to determine the energy savings from both the audit and retrofit stages of the weatherization.

¹⁷⁷If we were to instead assume that 100% of the beneficiaries are marginal, then the baseline MVPF would be 1.41

To construct the other weatherization MVPFs in our sample, we had to make relatively strong assumptions about the share of marginal beneficiaries and the valuation by marginal people. As an alternative, Allcott & Greenstone (2024) estimate a structural model of weatherization take-up to measure consumer surplus from the subsidy.

In their paper, Allcott and Greenstone estimate an MVPF of 0.93 using a 2020 social cost of carbon of \$190. For our in-context MVPF, we take this number directly. We harmonize this to a 2020 national MVPF using ratios of in-context externalities (Wisconsin in 2013) to 2020 US externalities. The explanation of the construction of this MVPF will therefore focus on the 2020 baseline MVPF. Following the approach in the paper, the cost and willingness to pay components are normalized per household in the population rather than per program participant.

The weatherization subsidy led to a 15% change in audit take-up and a 2% change in retrofit investment take-up. To construct the ratio of the in-context externalities to our 2020 externalities, we need to determine the percent of the environmental externality and producer profit loss that are from electricity versus natural gas. The paper finds that the audit reduced electricity consumption by 0.949 KWh per day and increased natural gas consumption by 0.064 therms per day. It also finds that the weatherization investment decreased electricity consumption by 0.193 KWh per day and decreased natural gas consumption by 0.46 therms per day. Combining this with the 15% and 2% changes in audit and investment probabilities results in a relative weighting of 109% on electricity and -9% on natural gas.

For ease of interpretation, we will start by constructing the WTP and then construct the Cost.

WTP The willingness to pay is comprised of the transfer benefits, environmental externality, and loss in producer profits.

To get the amount that beneficiaries value the transfer, we take the sum of the investment distortion (-0.91) and consumer surplus (10.79) and inflation adjust these values from 2013 to 2020 dollars. The resulting willingness to pay is \$10.98.

Consistent with other policy categories, we split the environmental externality into a global and local component. We begin by constructing the global and local externality components implied by the paper. The paper reports the monetized global and local externality values they use for electricity and natural gas. They report global damages of \$15.3 per MMBtu and \$0.11 per kWh. Similarly, they report \$1.00 per MMBtu and \$0.07 per kWh in local damages. The total environmental externality value that feeds into their MVPF construction is \$0.87 per household. Using the 109% and -9% weights from above, we find that the implied local and global split in the paper is \$0.35 and \$0.52, respectively. We can construct these values using our 2020 baseline externality values per MMBtu of natural gas and per kWh of electricity. The local environmental externality implied by our estimates in 2020 is 39% of that in the paper and our global environmental externality in 2020 is 127%. Scaling these numbers, and removing the 1.9% of the global benefits that flow to the government, results in a local environmental externality of \$0.65 and a global environmental externality of \$0.14. The rebound effect, as explained in Appendix D, offsets roughly 20% of the environmental benefits from electricity and 12% from natural gas. The resulting rebound effect is \$0.16.

Reduced energy consumption as a result of weatherization leads to lower profits for electric and natural gas utilities, as explained in Appendices C.2 and C.3. Following our approach for the environmental externalities, we take their producer profit component and scale it by the

ratio of our 2020 markup to their in-context markup. The paper uses a markup of \$2.75 per MMBtu of natural gas and \$0.10 per kWh of electricity. Our 2020 estimates of the natural gas markup is \$4.40 and of the electricity markup is \$0.01. The producer willingness to pay component reported in the paper is 0.21. We scale this by 0.10, which is the ratio of our implied markups to theirs. This leads to a producer willingness to pay component of \$0.02.

Summing across these components, the total willingness to pay in 2020 is \$11.60.

Cost The total cost is comprised of the direct program cost and fiscal externalities. For the direct program cost, we take the paper’s reported cost per household in the population of \$11.35 and inflation adjust this to 2020 dollars. This results in a cost of \$12.61.

The construction of the fiscal externality from the loss in government profit tax revenue from utility companies is explained in Appendix C.2. We take the total profit loss for producers and assume that the government loses tax revenue from 72% of private utilities and loses total profit from 28% of public utilities. Assuming a profit tax on private utilities of 10%, this results in a fiscal externality of \$0.007.

As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality reduces government costs by \$0.01 in 2020. The resulting total cost is \$12.61 in 2020. Dividing the WTP by the total cost, we arrive at the baseline MVPF of 0.92.

E.7 Vehicle Retirement

Vehicle retirement programs subsidize the scrapping of older vehicles, conditional on retiring a vehicle with specific qualities (e.g., the retired vehicle is above a specified age) and/or purchasing a vehicle that meets certain requirements (e.g., the purchased vehicle meets a stated fuel economy requirement). These subsidies can generate externalities through three channels. First, consumers accelerate the retirement of older, typically dirtier vehicles and the purchase of newer, cleaner vehicles, decreasing the usage of more polluting vehicles. Second, consumers purchase cleaner vehicles that qualify for the subsidy than they otherwise would have. Third, if the subsidy requires vehicles to be scrapped (rather than sold to the used-vehicle market), vehicles that would have otherwise stayed on the road are no longer used.¹⁷⁸ To abstract from the third channel, we assume all vehicles retired under these programs would have still been retired the next time the consumer purchased a vehicle.

We consider externalities from the first and second channels when forming two MVPFs for the “Cash for Clunkers” program (Li et al. 2013, Hoekstra et al. 2017). We only include externalities from acceleration when evaluating the Bay Area Air Quality Management District’s (BAAQMD) Vehicle Buyback Program (Sandler 2012). We account for the rebound in vehicle miles traveled due to owning a more fuel-efficient vehicle when evaluating both sources of externalities. In each instance, we quantify the change in the cost of driving one mile due to the fuel-economy improvement relative to the fuel economy of the vehicle that would have been used during that period.

¹⁷⁸We note that, while the third channel has the potential to generate large environmental externalities if a large number of years are taken off a vehicle’s life, valuing the effects of scrapping rather than selling an older vehicle requires an understanding of how the introduction of an additional vehicle into the used-vehicle market displaces which used vehicles are consumed and how this additional vehicle affects market prices.

E.7.1 Cash for Clunkers (Hoekstra et al. 2017)

Hoekstra et al. (2017) rely on variation in whether a household's vehicle just barely qualified for the federal Cash for Clunkers subsidy (which only subsidized the retirement of vehicles that received less than 18 miles per gallon) to estimate for a subset of participants in Texas the effects of this program on the acceleration of vehicle consumption and fuel economy improvements. They find that the Cash for Clunkers program accelerated vehicle purchases by, at most, eight months, and that the policy caused consumers to purchase vehicles that were, on average, 3.12 MPG more efficient than the vehicle they would have purchased.¹⁷⁹ We form a confidence interval for this MVPF using only the standard error (0.0929) reported for fuel economy improvement.

To form an MVPF in 2020, we pair the reported MPG improvement with the average fuel economy of a new vehicle released in 2009 (the year the policy change occurred) to calculate a percent improvement in vehicle fuel economy of 13.9% (3.12 MPG/22.40 MPG). We hold this percent improvement and the months accelerated constant in 2020. When determining how consumers value the subsidy, we assume everyone is inframarginal, as consumers did not vary their decision to buy a vehicle. However, when valuing the policy's externalities, we assume everyone is marginal, as consumers accelerated their purchase and varied what type of vehicle they purchased. We assume that, absent the policy, consumers would have purchased an average new light-duty vehicle released the year in which we evaluate the policy. Our in-context specification is set in 2009.

Transfer The authors find that the Cash for Clunkers program primarily shifted when consumers purchased a new vehicle (rather than generating additional consumption of new vehicles). As a result, we assume a 100% inframarginal share when valuing the transfer, as consumers did not change their decision to purchase a vehicle but rather when to purchase the vehicle. A 100% inframarginal share means consumers value the entire subsidy, which was on average \$4,210 in 2009 (in nominal dollars) (Hoekstra et al. 2017). We adjust this value for inflation (\$5,078.84 in 2020 dollars) to form an MVPF in 2020.

Retirement Acceleration Accelerating vehicle consumption decisions also accelerates vehicle retirement decisions. We assume all retired vehicles would have been scrapped at the time the new vehicle is purchased regardless of the subsidy. The Cash for Clunker programs therefore changes how long older, dirtier vehicles remain on the road.

Hoekstra et al. (2017) report that the average retired vehicle had a lifetime mileage of 160,155 miles. This is nearly identical to the total mileage reported after age 13 by the average light-duty vehicle in the FHWA (2017) (160,186.7 miles).¹⁸⁰ We assume the retired vehicle is 14 years old when scrapped and use this age to infer the vehicle's fuel economy and emission rates. For example, a 14-year-old vehicle in 2009 corresponded to a vehicle released in 1996, which had a fuel economy of 20.43 MPG. In 2020, a 14-year-old vehicle had a fuel economy of 20.60 MPG. We account for emission system decay (e.g., increases in emission rates) before this date.

¹⁷⁹The authors report this figure in the text (page 30), but to obtain a standard error for this result, we take the fuel economy improvement reported in Table 3, Column 6 (0.6578) and divide by the percent subsidized reported in Table 3, Column 6 (0.2107), which mirrors the authors' calculation.

¹⁸⁰See Appendix C.4.1 for details on how we use data from the NHTS and how we calculate vehicle externalities.

In 2009, the average retired vehicle imposed externalities of \$1.25 per gallon in global damages, \$0.41 per gallon in local pollution damages, and \$0.0981 per mile in driving damages from accidents, congestion, and $PM_{2.5}$ from tires and brakes, all expressed in nominal dollars. The more fuel efficient average new light-duty vehicle purchased in 2009 as a result of the subsidy (which had a fuel economy of 25.52 miles per gallon) imposed externalities of \$1.21 per gallon in global damages, \$0.12 per gallon in local pollution damages, and \$0.0980 per mile in driving damages from accidents, congestion, and $PM_{2.5}$ from tires and brakes, all expressed in nominal dollars. Driving damages vary entirely from differences in $PM_{2.5}$ from tires and brakes, while global and local pollution damages vary with differences in model-year-specific emission rates.

In 2020, the average retired vehicle imposed externalities of \$1.88 per gallon in global damages, \$0.15 per gallon in local pollution damages, and \$0.1183 per mile in driving damages from accidents, congestion, and $PM_{2.5}$ from tires and brakes, all expressed in nominal dollars. The more fuel efficient average new light-duty vehicle purchased in 2020 as a result of the subsidy (which had a fuel economy of 28.92 miles per gallon) imposed externalities of \$1.87 per gallon in global damages, \$0.14 per gallon in local pollution damages, and \$0.1184 per mile in driving damages from accidents, congestion, and $PM_{2.5}$ from tires and brakes, all expressed in nominal dollars. Driving damages vary entirely from differences in $PM_{2.5}$ from tires and brakes, while global and local pollution damages vary with differences in model-year-specific emission rates.

We assume that both the new and retired vehicles would have traveled 6,512.97 miles before accounting for the rebound in VMT from improved fuel economy. We calculate this mileage by taking the annual VMT for a 14-year-old vehicle (9,769.45 miles) and assuming that VMT is evenly distributed across the eight acceleration months, which yields a total VMT over the acceleration period of 6,512.97 miles. The fuel economy improvement induced by the policy caused a 19.95% reduction in the cost of driving one mile in 2009 and 28.76% in 2020, which we calculate by dividing the price of gasoline (\$2.40 per gallon in 2009 and \$2.27 per gallon in 2020) by the vehicle's fuel economy. We calculate this rebound relative to the cost of driving one mile with the average retired vehicle. Using an elasticity of VMT with respect to the cost of driving one mile of -0.2221 (Small & Van Dender 2007), we find that the fuel economy improvement caused drivers to travel an additional 288.63 miles (4.43% increase) over the acceleration period in 2009 and 415.98 miles (6.39% increase) in 2020. In 2009, over the eight month acceleration period, the retired vehicle would have traveled 6,512.97 miles while the the new vehicle would have traveled 6,801.60 miles. In 2020, over the eight month acceleration period, the retired vehicle would have traveled 6,512.97 miles while the the new vehicle would have traveled 6,928.94 miles. Dividing VMT by fuel economy gives us total gallons of gasoline consumed over this period.

Combining our externalities with vehicle usage, we find that Cash for Clunkers generated in 2009 \$76.13 (\$399.88 - \$310.01 - \$13.74) in global benefits: while the retired vehicle generated a total of \$399.88 in global damages from consuming 318.77 gallons of gas, the new vehicle generated \$310.01 in global damages from consuming 255.16 gallons of gas initially plus an additional \$13.74 in damages from consuming 11.31 gallons from the rebound. Applying the same approach, we find that Cash for Clunkers generated \$100.10 in local pollution benefits (\$130.94 - \$29.53 - \$1.31) in 2009. Pairing miles traveled with the per-mile driving externality noted above, we find -\$28.17 in local driving benefits (\$638.63 - \$638.50 - \$28.30) in 2009. The policy generates driving *damages* because of the increase in damages from accidents and congestion as a result of the policy, and improvements in vehicle fuel economy do not generate corresponding reductions in these driving externalities. The small difference between driving damages before accounting for the rebound arises entirely from differences in emissions of $PM_{2.5}$

from tires and brakes.

In 2020, we find that Cash for Clunkers generated in 2020 \$146.88 in global benefits (\$595.78 - \$421.95 - \$26.95): while the retired vehicle generated a total of \$595.78 in global damages from consuming 316.10 gallons of gas, the new vehicle generated \$421.95 in global damages from consuming 225.20 gallons of gas initially plus an additional \$26.95 in damages from consuming 14.38 gallons from the rebound. Applying the same approach, we find that Cash for Clunkers generated \$13.61 in local pollution benefits (\$47.71 - \$32.05 - \$2.05) in 2020. Pairing miles traveled with the per-mile driving externality noted above, we find -\$49.83 in local driving benefits (\$770.54 - \$771.12 - \$49.25) in 2020. The policy generates driving *damages* because of the increase in damages from accidents and congestion as a result of the policy, and improvements in vehicle fuel economy do not generate corresponding reductions in these driving externalities. The small difference between driving damages before accounting for the rebound arises entirely from differences in emissions of $PM_{2.5}$ from tires and brakes.

Each gallon of gasoline sold in 2009 provided gasoline producers with \$0.660 (in nominal dollars) in pre-tax profits. In 2020, the per-gallon markup was \$0.613 (in nominal dollars). Appendix C.4.5 describes how we calculate this per-gallon markup. Combining the estimated markup with the gallons consumed by each vehicle, we find that Cash for Clunkers generated -\$34.51 in pre-tax benefits (-\$210.35 + \$168.37 + \$7.46) for producers in 2009 and -\$46.87 (-\$193.64 + \$137.95 + \$8.81) for producers in 2020. Producers face *damages* because gas consumption falls. With a tax rate of 21% (Watson 2022), accelerating vehicle retirements cost gasoline producers \$27.26 in after-tax profits in 2009 and \$37.03 in 2020. We add the lost corporate tax revenue to the denominator of the MVPF, as described below.

Since consumers accelerated their purchase by less than one year, we need not discount nor account for rising social costs when valuing groups' WTP for accelerated retirement.

Fuel Economy Improvement In addition to accelerating vehicle retirements, the Cash for Clunkers program caused consumers to purchase more fuel efficient vehicles. We assume that, absent the policy, consumers would have purchased an average new light-duty vehicle released the year in which we evaluate the policy. Although we typically assume that the average light-duty vehicle drives for 19 years (see Appendix C.4.4), we assume here that new light-duty vehicles last 19 years plus the number of months accelerated. In this example, the vehicle lasts 19 years and eight months.

We assume the acceleration period takes up the remainder of the year in which the policy is being evaluated. For example, in our 2020 specification, the policy would go into effect on May 1, 2020, so that the eight months over which you accelerate your purchase take up the rest of 2020. This lets us begin the vehicle's normal 19 year lifespan and VMT schedule in January 2021. This changes which social costs we use to value the vehicle's damages. We discount to the baseline year being evaluated (2009 or 2020). All components noted below have been discounted. If the acceleration period is longer than six months, we assume the vehicle enters the first year of its lifetime with one year of emissions abatement system decay.

The average new light-duty vehicle purchased in 2009 generated \$12,337.97 (in nominal dollars) in global damages over its lifetime, assuming the vehicle begins its 19 year, 215,521.3 mileage lifetime on January 1, 2010. It also generated \$1,062.11 in local pollution damages, \$17,719.03 in local driving damages, and \$5,323.71 in profits for gasoline producers. For externalities that vary as a function of gasoline usage (global and local pollution damages and gasoline producer profits), we calculate the externalities generated by the more fuel efficient

vehicle by dividing the baseline externality by one plus the percent improvement in vehicle fuel economy. We assume externalities generated per mile traveled do not vary with vehicle fuel economy. With a 13.9% improvement in vehicle fuel economy, the more fuel efficient vehicle in 2009 generated \$10,828.89 in global damages, \$932.20 in local pollution damages, \$17,719.03 in local driving damages, and \$4,672.56 in profits for gasoline producers.

The average new light-duty vehicle purchased in 2020 generated \$15,917.92 (in nominal dollars) in global damages over its lifetime, assuming the vehicle begins its 19 year, 215,521.3 mileage lifetime on January 1, 2021. It also generated \$1,125.36 in local pollution damages, \$21,399.21 in local driving damages, and \$4,361.78 in profits for gasoline producers. For externalities that vary as a function of gasoline usage (global and local pollution damages and gasoline producer profits), we calculate the externalities generated by the more fuel efficient vehicle by dividing the baseline externality by one plus the percent improvement in vehicle fuel economy. We assume externalities generated per mile traveled do not vary with vehicle fuel economy. With a 13.9% improvement in vehicle fuel economy, the more fuel efficient vehicle in 2020 generated \$13,970.98 in global damages, \$987.72 in local pollution damages, \$21,399.21 in local driving damages, and \$3,828.29 in profits for gasoline producers.

We again account for the rebound in VMT that follows from driving a more fuel efficient vehicle. The fuel economy improvement induced by the policy caused a 12.23% reduction in the cost of driving, which we calculate by dividing the price of gasoline (\$2.40 per gallon in 2009 and \$2.27 per gallon in 2020) by the vehicle's fuel economy. The percent change in the cost of traveling one mile does not vary over time because we hold the percent improvement in fuel economy fixed across specifications. Using an elasticity of VMT with respect to the cost of driving one mile of -0.2221 (Small & Van Dender 2007), we find that the fuel economy improvement caused drivers to use their vehicle 2.72% more over the vehicle's lifetime.

Accounting for the rebound effect, we find that Cash for Clunkers generated in 2009 \$1,214.90 (\$12,337.97 - \$10,828.89 - \$294.17) in global benefits over the vehicle's lifetime: while the original vehicle generated a total of \$12,337.97 in global damages, the more fuel efficient vehicle generated \$10,828.89 in global damages initially plus an additional \$294.17 in damages from the 2.72% VMT rebound. Applying the same approach, we find that Cash for Clunkers generated \$104.59 in local pollution benefits (\$1,062.11 - \$932.20 - \$25.32) in 2009. Pairing miles traveled with the per-mile driving externality noted above, we find -\$481.34 in local driving benefits (\$17,719.03 - \$17,719.03 - \$481.34) in 2009. Local driving damages follow entirely from the rebound effect since driving externalities do not vary as a function of fuel economy. Pre-tax producer profits fell by \$524.22 (\$5,323.71 - \$4,672.56 - \$126.93) in 2009 due to the policy change, implying post-tax benefits of -\$414.13.

We find that Cash for Clunkers generated in 2020 \$1,567.42 (\$15,917.92 - \$13,970.98 - \$379.53) in global benefits over the vehicle's lifetime: while the original vehicle generated a total of \$15,917.92 in global damages, the more fuel efficient vehicle generated \$13,970.98 in global damages initially plus an additional \$379.53 in damages from the 2.72% rebound. Applying the same approach, we find that Cash for Clunkers generated \$110.81 in local pollution benefits (\$1,125.36 - \$987.72 - \$26.83) in 2020. Pairing miles traveled with the per-mile driving externality noted above, we find -\$581.32 in local driving benefits (\$21,399.21 - \$21,399.21 - \$581.32) in 2020. Local driving damages follow entirely from the rebound effect since driving externalities do not vary as a function of fuel economy. Pre-tax producer profits fell by \$429.50 (\$4,361.78 - \$3,828.29 - \$104.00) in 2020 due to the policy change, implying post-tax benefits of -\$339.31.

Total WTP Assuming all consumers are inframarginal to the subsidy, consumers are willing to pay \$4,210 in 2009 and \$5,078.84 in 2020, both in nominal dollars. We assume everyone who received the subsidy changed their behavior by both accelerating the retirement of their old vehicle and opting for a more fuel efficient vehicle, so scaling the sum of benefits from acceleration and fuel economy improvements by the share of consumers that are marginal (100%) returns the sum of benefits from acceleration and fuel economy improvements. Producers have a negative willingness to pay for lost profits, and society has a negative willingness to pay for increases in driving damages (e.g., accidents, congestion, and $PM_{2.5}$ from tires and brakes).

In 2009, combining benefits from retirement acceleration and fuel economy improvements and multiplying by the share of consumers that were marginal to the subsidy, society is willing to pay \$1,291.03 ($\$76.13 + \$1,214.90$) for global damages, \$204.69 ($\$100.10 + \104.59) for local pollution damages, and $-\$509.52$ ($-\$28.17 + -\481.34) for local driving damages. Producers are willing to pay $-\$558.73$ ($-\$34.51 + -\524.22) for changes in pre-tax profits, or \$441.40 in post-tax profits. We adjust society's willingness to pay for the share of global benefits that do not flow to the US government (0.981), which yields an adjusted willingness to pay of \$1,266.50. Society's willingness to pay for all local damages ($-\$304.83$) is the sum of society's willingness to pay for pollution damages and driving damages. Society's willingness to pay for the rebound in pollution and driving damages is $-\$838.33$ ($-\$13.74 \times 0.981 + -\$1.31 + -\$28.30$ from acceleration, and $-\$294.17 \times 0.981 + -\$25.32 + -\$481.34$) in 2009.

In 2020, combining benefits from retirement acceleration and fuel economy improvements and multiplying by the share of consumers that were marginal to the subsidy, society is willing to pay \$1,714.30 ($\$146.88 + \$1,567.42$) for global damages, \$124.42 ($\$13.61 + \110.81) for local pollution damages, and $-\$631.14$ ($-\$49.83 + -\581.32) for local driving damages. Producers are willing to pay $-\$476.37$ ($-\$46.87 + -\429.50) for changes in pre-tax profits, or \$376.33 in post-tax profits. We adjust society's willingness to pay for the share of global benefits that do not flow to the US government (0.981), which yields an adjusted willingness to pay of \$1,681.46. Society's willingness to pay for all local damages ($-\$506.72$) is the sum of society's willingness to pay for pollution damages and driving damages. Society's willingness to pay for the rebound in pollution and driving damages is $-\$1,058.21$ ($-\$26.95 \times 0.981 + -\$2.05 + -\$49.25$ from acceleration, and $-\$379.53 \times 0.981 + -\$26.83 + -\$581.32$) in 2020.

Summing all willingness to pay components (consumers, producers, global benefits, and local damages), we calculate a total willingness to pay of \$4,730.27 in 2009 ($\$4,210 + -\$441.40 + \$1,266.50 + -\304.83) and \$5,877.25 in 2020 ($\$5,078.84 + -\$376.33 + \$1,681.46 + -\506.72).

Cost The total cost is comprised of the direct program cost and fiscal externalities. The program cost is equal to the size of the subsidy, or \$4,210 in 2009 and \$5,078.84 in 2020 (Hoekstra et al. 2017) in nominal dollars, adjusting for inflation. All fiscal externalities are scaled by the share of consumers that are marginal, but we abstract from this step in this section since we assume the share of consumers that are marginal is 100%.

The government loses gas tax revenue as a result of the acceleration of vehicle retirements and improvements in fuel economy. In 2009, the average total gas tax levied by the federal government and states was \$0.39 per gallon (FHWA 2021, 2020). In 2020, the average total gas tax was \$0.46 per gallon. In 2009, accelerating vehicle retirement cost the government \$20.49 ($\$124.89 - \$99.97 - \4.43) in lost gas tax revenue, while fuel economy improvements cost the government \$311.26 ($\$3,160.96 - \$2,774.34 - \75.37), resulting in a total fiscal externality from lost gas tax revenue of \$331.75 (in nominal dollars). In 2020, accelerating vehicle retirement cost

the government \$35.57 (\$146.95 - \$104.69 - \$6.69) in lost gas tax revenue, while fuel economy improvements cost the government \$325.94 (\$3,310.10 - \$2,905.24 - \$78.92), resulting in a total fiscal externality from lost gas tax revenue of \$361.51 (in nominal dollars).

The government also loses corporate tax revenue when gasoline producers lose profits. As noted above, in 2009, gasoline producers faced pre-tax lost profits of \$34.51 as a result of retirement acceleration and \$524.22 from fuel economy improvements. With a 21% effective corporate tax rate (Watson 2022), the government lost \$117.33 in corporate tax revenue in 2009 as a result of the Cash for Clunkers program. In 2020, gasoline producers lost \$46.87 in pre-tax profits from acceleration and \$429.50 from fuel economy improvements for a total fiscal externality of \$100.04.

We place the share of global benefits that flow to the US government in the denominator of the MVPF. The Cash for Clunkers program generated \$1,291.04 in global benefits in 2009 before adjusting for the share of benefits that flow to the US government. With 1.92% of global benefits flowing to the US government as long-run revenue, Cash for Clunkers cost the government -\$24.73 (in nominal dollars) in 2009 by abating greenhouse gases. In 2020, Cash for Clunkers generated \$1,714.30 in total global benefits, for a climate fiscal externality of -\$32.84 (in nominal dollars). This fiscal externality is negative because it raises revenue for the government.

Summing the program cost and the fiscal externalities, Cash for Clunkers imposed a net cost on the government of \$4,634.35 in 2009 and \$5,507.56 in 2020. The fiscal externality from lost tax revenue was \$449.08 in 2009 and \$461.55 in 2020, while the climate fiscal externality cost the government -\$24.73 in 2009 and -\$32.84 in 2020.

MVPF Dividing the total WTP calculated above (\$4,730.27) by the total cost (\$4,634.35), we form an MVPF of 1.021 in 2009. Using our 2020 estimates of a total WTP of \$5,877.25 and a net government cost of \$5,507.56, we obtain an MVPF of 1.067 [1.052, 1.082].

E.7.2 Cash for Clunkers (Li et al. 2013)

Li et al. (2013) use a difference-in-difference to compare changes in the United States vehicle market (treatment group) as a result of the Cash for Clunkers program to changes in the Canadian vehicle market (control group) over the same period. They find that the Cash for Clunkers program accelerated vehicle purchases by at most seven months (from December to June) and that consumers purchased vehicles that were 1.94 MPG more efficient than the vehicle they would have purchased.¹⁸¹ To calculate the MPG improvement, we take the average improvement in vehicle fuel economy for all vehicles purchased during the sample period (0.21, s.e. 0.04, authors' Table 5, Panel 2) and divide by the share of vehicle transactions in the full sample that were eligible for the subsidy (678,359 / 6,270,967).

To form an MVPF in 2020, we pair the reported MPG improvement with the average fuel economy of a new vehicle released in 2009 (the year the policy change occurred) to calculate

¹⁸¹We use seven months of acceleration as the behavioral response to the policy since the authors find increases in sales in the summer of 2009 but effectively no net increase in sales between June and December 2009, meaning consumers could have, at most, accelerated by seven months by moving their purchase from December to June. Using the greatest possible acceleration allows us also to be consistent with the acceleration duration reported by Hoekstra et al. (2017), who report that the program accelerated vehicle purchases by no more than eight months.

a percent improvement in vehicle fuel economy of 8.67% (1.94 MPG/22.40 MPG). We hold this percent improvement and the months accelerated constant in 2020. When determining how consumers value the subsidy, we assume everyone is inframarginal, as consumers did not vary their decision to buy a vehicle. However, when valuing the policy’s externalities, we assume everyone is marginal, as consumers accelerated their purchase and varied what type of vehicle they purchased. We assume that, absent the policy, consumers would have purchased an average new light-duty vehicle released the year in which we evaluate the policy. Our in-context specification is set in 2009.

Transfer The authors find that the Cash for Clunkers program primarily shifted when consumers purchased a new vehicle (rather than generating additional consumption of new vehicles). As a result, we assume a 100% inframarginal share when valuing the transfer, as consumers did not change their decision to purchase a vehicle but rather when to purchase the vehicle. A 100% inframarginal share means consumers value the entire subsidy, which was on average \$4,210 in 2009 (in nominal dollars) (Hoekstra et al. 2017). We adjust this value for inflation (\$5,078.84 in 2020 dollars) to form an MVPF in 2020.

Retirement Acceleration Accelerating vehicle consumption decisions also accelerates vehicle retirement decisions. We assume all retired vehicles would have been scrapped at the time the new vehicle is purchased regardless of the subsidy. The Cash for Clunker programs therefore changes how long older, dirtier vehicles remain on the road.

Hoekstra et al. (2017) report that the average retired vehicle had a lifetime mileage of 160,155 miles. This is nearly identical to the total mileage reported after age 13 by the average light-duty vehicle in the FHWA (2017) (160,186.7 miles).¹⁸² We assume the retired vehicle is 14 years old when scrapped and use this age to infer the vehicle’s fuel economy and emission rates. For example, a 14-year-old vehicle in 2009 corresponded to a vehicle released in 1996, which had a fuel economy of 20.43 MPG. In 2020, a 14-year-old vehicle had a fuel economy of 20.60 MPG. We account for emission system decay (e.g., increases in emission rates) before this date.

In 2009, the average retired vehicle imposed externalities of \$1.25 per gallon in global damages, \$0.41 per gallon in local pollution damages, and \$0.0981 per mile in driving damages from accidents, congestion, and $PM_{2.5}$ from tires and brakes, all expressed in nominal dollars. The more fuel efficient average new light-duty vehicle purchased in 2009 as a result of the subsidy (which had a fuel economy of 24.34 miles per gallon) imposed externalities of \$1.21 per gallon in global damages, \$0.12 per gallon in local pollution damages, and \$0.0980 in driving damages from accidents, congestion, and $PM_{2.5}$ from tires and brakes, all expressed in nominal dollars. Driving damages vary entirely from differences in $PM_{2.5}$ from tires and brakes, while global and local pollution damages vary with differences in model-year-specific emission rates.

In 2020, the average retired vehicle imposed externalities of \$1.88 per gallon in global damages, \$0.15 per gallon in local pollution damages, and \$0.1183 per mile in driving damages from accidents, congestion, and $PM_{2.5}$ from tires and brakes, all expressed in nominal dollars. The more fuel efficient average new light-duty vehicle purchased in 2020 as a result of the subsidy (which had a fuel economy of 27.58 miles per gallon) imposed externalities of \$1.87 per gallon in global damages, \$0.14 per gallon in local pollution damages, and \$0.1184 in driving damages

¹⁸²See Appendix C.4.1 for details on how we use data from the NHTS and how we calculate vehicle externalities.

from accidents, congestion, and $PM_{2.5}$ from tires and brakes, all expressed in nominal dollars. Driving damages vary entirely from differences in $PM_{2.5}$ from tires and brakes, while global and local pollution damages vary with differences in model-year-specific emission rates.

We assume that both the new and retired vehicles would have traveled 5,698.85 miles before accounting for the rebound in VMT from improved fuel economy. We calculate this mileage by taking the annual VMT for a 14-year-old vehicle (9,769.45 miles) and assuming that VMT is evenly distributed across the seven acceleration months, which yields a total VMT over the acceleration period of 5,698.85 miles. The fuel economy improvement induced by the policy caused a 16.07% reduction in the cost of driving one mile in 2009 and 25.30% in 2020, which we calculate by dividing the price of gasoline (\$2.40 per gallon in 2009 and \$2.27 per gallon in 2020) by the vehicle's fuel economy. We calculate this rebound relative to the cost of driving one mile with the average retired vehicle. Using an elasticity of VMT with respect to the cost of driving one mile of -0.2221 (Small & Van Dender 2007), we find that the fuel economy improvement caused drivers to travel an additional 203.42 miles (3.57% increase) over the acceleration period in 2009 and 320.25 miles (5.62% increase) in 2020. In 2009, over the seven month acceleration period, the retired vehicle would have traveled 5,698.85 miles while the the new vehicle would have traveled 5,902.26 miles. In 2020, over the seven month acceleration period, the retired vehicle would have traveled 5,698.85 miles while the the new vehicle would have traveled 6,019.09 miles. Dividing VMT by fuel economy gives us total gallons of gasoline consumed over this period.

Combining our externalities with vehicle usage, we find that Cash for Clunkers generated in 2009 \$55.33 (\$349.89 - \$284.41 - \$10.15) in global benefits: while the retired vehicle generated a total of \$349.89 in global damages from consuming 278.92 gallons of gas, the new vehicle generated \$284.41 in global damages from consuming 234.10 gallons of gas initially plus an additional \$10.15 in damages from consuming 8.36 gallons from the rebound. Applying the same approach, we find that Cash for Clunkers generated \$86.52 in local pollution benefits (\$114.58 - \$27.09 - \$0.97) in 2009. Pairing miles traveled with the per-mile driving externality noted above, we find -\$19.83 in local driving benefits (\$558.80 - \$558.69 - \$19.94) in 2009. The policy generates driving *damages* because of the increase in damages from accidents and congestion as a result of the policy, and improvements in vehicle fuel economy do not generate corresponding reductions in these driving externalities. The small difference between driving damages before accounting for the rebound arises entirely from differences in emissions of $PM_{2.5}$ from tires and brakes.

We find that Cash for Clunkers generated in 2020 \$112.44 (\$521.31 - \$387.11 - \$21.75) in global benefits: while the retired vehicle generated a total of \$521.31 in global damages from consuming 276.59 gallons of gas, the new vehicle generated \$387.11 in global damages from consuming 206.61 gallons of gas initially plus an additional \$21.75 in damages from consuming 11.61 gallons from the rebound. Applying the same approach, we find that Cash for Clunkers generated \$10.69 in local pollution benefits (\$41.74 - \$29.40 - \$1.65) in 2020. Pairing miles traveled with the per-mile driving externality noted above, we find -\$38.42 in local driving benefits (\$674.22 - \$674.73 - \$37.92) in 2020. The policy generates driving *damages* because of the increase in damages from accidents and congestion as a result of the policy, and improvements in vehicle fuel economy do not generate corresponding reductions in these driving externalities. The small difference between driving damages before accounting for the rebound arises entirely from differences in emissions of $PM_{2.5}$ from tires and brakes.

Each gallon of gasoline sold in 2009 provided gasoline producers with \$0.660 (in nominal

dollars) in pre-tax profits. In 2020, the per-gallon markup was \$0.613 (in nominal dollars). Appendix C.4.5 describes how we calculate this per-gallon markup. Combining the estimated markup with the gallons consumed by each vehicle, we find that Cash for Clunkers generated -\$24.07 in pre-tax benefits (-\$184.05 + \$154.47 + \$5.51) for producers in 2009 and -\$35.76 (-\$169.43 + \$126.56 + \$7.11) for producers in 2020. Producers face *damages* because gas consumption falls. With a tax rate of 21% (Watson 2022), accelerating vehicle retirements cost gasoline producers \$19.01 in after-tax profits in 2009 and \$28.25 in 2020. We add the lost corporate tax revenue to the denominator of the MVPF, as described below.

Since consumers accelerated their purchase by less than one year, we need not discount nor account for rising social costs when valuing groups' WTP for accelerated retirement.

Fuel Economy Improvement In addition to accelerating vehicle retirements, the Cash for Clunkers program caused consumers to purchase more fuel efficient vehicles. We assume that, absent the policy, consumers would have purchased an average new light-duty vehicle released the year in which we evaluate the policy. Although we typically assume that the average light-duty vehicle drives for 19 years (see Appendix C.4.4), we assume here that new light-duty vehicles last 19 years plus the number of months accelerated. In this example, the vehicle lasts 19 years and seven months.

We assume the acceleration period takes up the remainder of the year in which the policy is being evaluated. For example, in our 2020 specification, the policy would go into effect on June 1, 2020, so that the seven months over which you accelerate your purchase take up the rest of 2020. This lets us begin the vehicle's normal 19 year lifespan and VMT schedule in January 2021. This changes which social costs we use to value the vehicle's damages. We discount to the baseline year being evaluated (2009 or 2020). All components noted below have been discounted. If the acceleration period is longer than six months, we assume the vehicle enters the first year of its lifetime with one year of emissions abatement system decay.

The average new light-duty vehicle purchased in 2009 generated \$12,337.97 (in nominal dollars) in global damages over its lifetime, assuming the vehicle begins its 19 year, 215,521.3 mileage lifetime on January 1, 2010. It also generated \$1,062.11 in local pollution damages, \$17,719.03 in local driving damages, and \$5,323.71 in profits for gasoline producers. For externalities that vary as a function of gasoline usage (global and local pollution damages and gasoline producer profits), we calculate the externalities generated by the more fuel efficient vehicle by dividing the baseline externality by one plus the percent improvement in vehicle fuel economy. We assume externalities generated per mile traveled do not vary with vehicle fuel economy. With an 8.67% improvement in vehicle fuel economy, the more fuel efficient vehicle in 2009 generated \$11,354.09 in global damages, \$977.41 in local pollution damages, \$17,719.03 in local driving damages, and \$4,899.17 in profits for gasoline producers.

The average new light-duty vehicle purchased in 2020 generated \$15,917.92 (in nominal dollars) in global damages over its lifetime, assuming the vehicle begins its 19 year, 215,521.3 mileage lifetime on January 1, 2021. It also generated \$1,125.36 in local pollution damages, \$21,399.21 in local driving damages, and \$4,361.78 in profits for gasoline producers. For externalities that vary as a function of gasoline usage (global and local pollution damages and gasoline producer profits), we calculate the externalities generated by the more fuel efficient vehicle by dividing the baseline externality by one plus the percent improvement in vehicle fuel economy. We assume externalities generated per mile traveled do not vary with vehicle fuel economy. With a 8.67% improvement in vehicle fuel economy, the more fuel efficient vehicle in

2020 generated \$14,648.56 in global damages, \$ 1,035.62 in local pollution damages, \$21,399.21 in local driving damages, and \$4,013.95 in profits for gasoline producers.

We again account for the rebound in VMT that follows from driving a more fuel efficient vehicle. The fuel economy improvement induced by the policy caused a 7.97% reduction in the cost of driving, which we calculate by dividing the price of gasoline (\$2.40 per gallon in 2009 and \$2.27 per gallon in 2020) by the vehicle's fuel economy. The percent change in the cost of traveling one mile does not vary over time because we hold the percent improvement in fuel economy fixed across specifications. Using an elasticity of VMT with respect to the cost of driving one mile of -0.2221 (Small & Van Dender 2007), we find that the fuel economy improvement caused drivers to use their vehicle 1.77% more over the vehicle's lifetime.

Accounting for the rebound effect, we find that Cash for Clunkers generated in 2009 \$782.79 (\$12,337.97 - \$11,354.09 - \$201.09) in global benefits over the vehicle's lifetime: while the original vehicle generated a total of \$12,337.97 in global damages, the more fuel efficient vehicle generated \$11,354.09 in global damages initially plus an additional \$201.09 in damages from the 1.77% VMT rebound. Applying the same approach, we find that Cash for Clunkers generated \$67.39 in local pollution benefits (\$1,062.11 - \$977.41 - \$17.31) in 2009. Pairing miles traveled with the per-mile driving externality noted above, we find -\$313.83 in local driving benefits (\$17,719.03 - \$17,719.03 - \$313.83) in 2009. Local driving damages follow entirely from the rebound effect since driving externalities do not vary as a function of fuel economy. Pre-tax producer profits fell by \$337.77 (\$5,323.71 - \$4,899.17 - \$86.77) in 2009 due to the policy change, implying post-tax benefits of -\$266.83.

We find that Cash for Clunkers generated in 2020 \$1,009.92 (\$15,917.92 - \$14,648.56 - \$259.44) in global benefits over the vehicle's lifetime: while the original vehicle generated a total of \$15,917.92 in global damages, the more fuel efficient vehicle generated \$14,648.56 in global damages initially plus an additional \$259.44 in damages from the 1.77% rebound. Applying the same approach, we find that Cash for Clunkers generated \$71.40 in local pollution benefits (\$1,125.36 - \$1,035.62 - \$18.34) in 2020. Pairing miles traveled with the per-mile driving externality noted above, we find -\$379.01 in local driving benefits (\$21,399.21 - \$21,399.21 - \$379.01) in 2020. Local driving damages follow entirely from the rebound effect since driving externalities do not vary as a function of fuel economy. Pre-tax producer profits fell by \$276.74 (\$4,361.78 - \$4,013.95 - \$71.09) in 2020 due to the policy change, implying post-tax benefits of -\$218.62.

Total WTP Assuming all consumers are inframarginal to the subsidy, consumers are willing to pay \$4,210 in 2009 and \$5,078.84 in 2020, both in nominal dollars. We assume everyone who received the subsidy changed their behavior by both accelerating the retirement of their old vehicle and opting for a more fuel efficient vehicle, so scaling the sum of benefits from acceleration and fuel economy improvements by the share of consumers that are marginal (100%) returns the sum of benefits from acceleration and fuel economy improvements. Producers have a negative willingness to pay for lost profits, and society has a negative willingness to pay for increases in driving damages (e.g., accidents, congestion, and $PM_{2.5}$ from tires and brakes).

In 2009, combining benefits from retirement acceleration and fuel economy improvements and multiplying by the share of consumers that were marginal to the subsidy, society is willing to pay \$838.12 (\$55.33 + \$782.79) for global damages, \$153.91 (\$86.52 + \$67.39) for local pollution damages, and -\$333.66 (-\$19.83 + -\$313.83) for local driving damages. Producers are willing to pay -\$361.83 (-\$24.07 + -\$337.77) for changes in pre-tax profits, or \$285.85 in

post-tax profits. We adjust society's willingness to pay for the share of global benefits that do not flow to the US government (0.981), which yields an adjusted willingness to pay of \$822.07. Society's willingness to pay for all local damages (-\$179.75) is the sum of society's willingness to pay for pollution damages and driving damages. Society's willingness to pay for the rebound in pollution and driving damages is -\$559.28 ($-\$10.15 \times 0.981 + -\$0.97 + -\19.94 from acceleration, and $-\$201.09 \times 0.981 + -\$17.31 + -\$313.83$) in 2009.

In 2020, combining benefits from retirement acceleration and fuel economy improvements and multiplying by the share of consumers that were marginal to the subsidy, society is willing to pay \$1,122.36 ($\$112.44 + \$1,009.92$) for global damages, \$82.09 ($\$10.69 + \71.40) for local pollution damages, and -\$417.43 ($-\$38.42 + -\379.01) for local driving damages. Producers are willing to pay -\$312.49 ($-\$35.76 + -\276.74) for changes in pre-tax profits, or \$246.87 in post-tax profits. We adjust society's willingness to pay for the share of global benefits that do not flow to the US government (0.981), which yields an adjusted willingness to pay of \$1,100.86. Society's willingness to pay for all local damages (-\$335.34) is the sum of society's willingness to pay for pollution damages and driving damages. Society's willingness to pay for the rebound in pollution and driving damages is -\$712.77 ($-\$21.75 \times 0.981 + -\$1.65 + -\37.92 from acceleration, and $-\$259.44 \times 0.981 + -\$18.34 + -\$379.01$) in 2020.

Summing all willingness to pay components (consumers, producers, global benefits, and local damages), we calculate a total willingness to pay of \$4,566.47 in 2009 ($\$4,210 + -\$285.85 + \$822.07 + -\179.75) and \$5,597.49 in 2020 ($\$5,078.84 + -\$246.87 + \$1,100.86 + -\335.34).

Cost The total cost is comprised of the direct program cost and fiscal externalities. The program cost is equal to the size of the subsidy, or \$4,210 in 2009 and \$5,078.84 in 2020 (Hoekstra et al. 2017) in nominal dollars, adjusting for inflation. All fiscal externalities are scaled by the share of consumers that are marginal, but we abstract from this step in this section since we assume the share of consumers that are marginal is 100%.

The government loses gas tax revenue as a result of the acceleration of vehicle retirements and improvements in fuel economy. In 2009, the average total gas tax levied by the federal government and states was \$0.39 per gallon (FHWA 2021, 2020). In 2020, the average total gas tax was \$0.46 per gallon. In 2009, accelerating vehicle retirement cost the government \$14.29 ($\$109.28 - \$91.72 - \3.27) in lost gas tax revenue, while fuel economy improvements cost the government \$200.55 ($\$3,160.96 - \$2,908.89 - \51.52), resulting in a total fiscal externality from lost gas tax revenue of \$214.84 (in nominal dollars). In 2020, accelerating vehicle retirement cost the government \$27.14 ($\$128.58 - \$96.05 - \5.40) in lost gas tax revenue, while fuel economy improvements cost the government \$210.01 ($\$3,310.10 - \$3,046.14 - \53.95), resulting in a total fiscal externality from lost gas tax revenue of \$237.15 (in nominal dollars).

The government also loses corporate tax revenue when gasoline producers lose profits. As noted above, in 2009, gasoline producers faced pre-tax lost profits of \$24.07 as a result of retirement acceleration and \$337.77 from fuel economy improvements. With a 21% effective corporate tax rate (Watson 2022), the government lost \$75.98 in corporate tax revenue in 2009 as a result of the Cash for Clunkers program. In 2020, gasoline producers lost \$35.76 in pre-tax profits from acceleration and \$276.74 from fuel economy improvements for a total fiscal externality of \$65.62.

We place the share of global benefits that flow to the US government in the denominator of the MVPF. The Cash for Clunkers program generated \$838.12 in global benefits in 2009 before adjusting for the share of benefits that flow to the US government. With 1.92% of

global benefits flowing to the US government as long-run revenue, Cash for Clunkers cost the government -\$16.05 (in nominal dollars) in 2009 by abating greenhouse gases. In 2020, Cash for Clunkers generated \$1,122.36 in total global benefits, for a climate fiscal externality of -\$21.50 (in nominal dollars). This fiscal externality is negative because it raises revenue for the government.

Summing the program cost and the fiscal externalities, Cash for Clunkers imposed a net cost on the government of \$4,484.77 in 2009 and \$5,360.11 in 2020. The fiscal externality from lost tax revenue was \$290.82 in 2009 and \$302.77 in 2020, while the climate fiscal externality cost the government -\$16.05 in 2009 and -\$21.50 in 2020.

MVPF Dividing the total WTP calculated above (\$4,566.47) by the total cost (\$4,484.77), we form an MVPF of 1.018 in 2009. Using our 2020 estimates of a total WTP of \$5,597.49 and a net government cost of \$5,360.11, we obtain an MVPF of 1.044 [1.030, 1.058].

E.7.3 BAAQMD Vehicle Buyback Program

Sandler (2012) uses data from California vehicle inspections to compare vehicles purchased and scrapped under the Bay Area Air Quality Management District's Vehicle Buyback Program to similar vehicles that were not retired as a result of the program. Specifically, the author constructs a control group by propensity-scoring vehicles depending on vehicle and driver characteristics and then uses these scores to calculate counterfactual estimates of vehicle usage. We focus on the estimated VMT remaining (16,027.6 miles, s.e. 161.0) and days of driving remaining (1,394.2 days, s.e. 9.509) for the control group vehicles that qualified for the program (authors' Table 2, rows 1 and 2). We form a confidence interval for this MVPF using only uncertainty in the paper's estimate of miles abated.

We only consider benefits from accelerated vehicle retirement. We assume the retired vehicle is 26 years old, as the Buyback Program required vehicles to be from model year 1998 or older in 2023 (BAAQMD 2023).¹⁸³ Both fuel economy and emission rates vary as a function of the vehicle's model year. We follow the author in assuming that consumers replaced their retired vehicle with a fleet average vehicle. When determining how consumers value the subsidy, we assume everyone is inframarginal, as consumers did not vary their decision to buy a vehicle. However, when valuing the policy's externalities, we assume everyone is marginal, as consumers accelerated their purchase. Our in-context specification is set in 1996, the earliest year of the program.

Transfer The author finds that the BAAQMD Vehicle Buyback Program shifted when consumers purchased a new vehicle. As a result, we assume a 100% inframarginal share when valuing the transfer, as consumers did not change their decision to purchase a vehicle but rather when to purchase the vehicle. A 100% inframarginal share means consumers value the entire subsidy, which was on average \$500 (in 2000 dollars) before 2004 (Sandler 2012), or \$455.47 in 1996 dollars. For our 2020 specification, we use the post-2004 average subsidy of \$650 (in 2000 dollars) (Sandler 2012), or \$977.11 in 2020 dollars.

¹⁸³This is consistent with our assumption that a new vehicle purchased in 2020 would be one year old in 2020. Although vehicles older than this qualified for the program as well, we take the youngest possible vehicle age as a conservative estimate.

Retirement Acceleration Accelerating vehicle consumption decisions also accelerates vehicle retirement decisions. We assume all retired vehicles would have been scrapped at the time the new vehicle is purchased regardless of the subsidy. The BAAQMD Vehicle Buyback Program therefore changes how long older, dirtier vehicles remain on the road. We assume the abated VMT (16,027.6 miles) is evenly distributed over the accelerated years (1,394.2 days, or 3.82 years), so that vehicles travel 4,196 miles annually for the first three accelerated years and 3,439.58 miles in the final 299.2 days.

In 1996, a 26-year-old light-duty vehicle corresponded to a 1971 model year vehicle. In the first accelerated year (1996), this vehicle imposed \$1.533 per gallon in local pollution damages, \$0.673 per gallon in global pollution damages, and \$0.072 per mile in local driving damages, all expressed in 1996 dollars. We account for rising the social costs of greenhouse gases over the 3.82 years accelerated. Since the vehicle is from before 1975, we do not account for vehicle decay, as we assume these vehicles predate modern emissions abatement technologies; as a result, local pollution and driving damages do not increase over the accelerated period. In 1997, the vehicle generated \$0.699 per gallon in global pollution damages, \$0.724 per gallon in 1998, and \$0.750 per gallon in 1999. With a fuel economy of 9.67 miles per gallon, the vehicle consumed 433.93 gallons annually for the first three years and 355.71 gallons in the final 299.2 days. Discounting after the first year, the retired vehicle would have generated \$1,143.08 in global damages, \$2,470.72 in local pollution damages, and \$1,116.18 in local driving damages. With a pre-tax nominal per-gallon markup of \$0.519 in 1996, the vehicle generated \$836.31 in discounted pre-tax profits for gasoline producers, or \$660.68 in post-tax profits with an assumed 21% corporate tax rate (Watson 2022).

In 2020, a 26-year-old light-duty vehicle corresponded to a 1995 model year vehicle. In the first accelerated year (2020), this vehicle imposed \$0.641 per gallon in local pollution damages, \$1.957 per gallon in global pollution damages, and \$0.1183 per mile in local driving damages, all expressed in 2020 dollars and accounting for vehicle decay. We account for rising the social costs of greenhouse gases over the 3.82 years accelerated. Since the vehicle is 26 years old, we need not account for further vehicle decay over the acceleration period, as we assume vehicle decay stops after age 19 (Jacobsen et al. 2023); as a result, local pollution and driving damages do not increase over the accelerated period. In 2021, the vehicle generated \$1.998 per gallon in global pollution damages, \$2.030 per gallon in 2022, and \$2.071 per gallon in 2023. With a fuel economy of 20.486 miles per gallon, the vehicle consumed 204.83 gallons annually for the first three years and 167.90 gallons in the final 299.2 days. Discounting after the first year, the retired vehicle would have generated \$1,529.50 in global damages, \$487.86 in local pollution damages, and \$1,843.44 in local driving damages. With a pre-tax nominal per-gallon markup of \$0.613 in 2020, the vehicle generated \$466.00 in discounted pre-tax profits for gasoline producers, or \$368.14 in post-tax profits with an assumed 21% corporate tax rate (Watson 2022).

When calculating externalities from the replacement vehicle, we account for the rebound in VMT from driving a more fuel efficient vehicle.¹⁸⁴ The fuel economy improvement induced by the policy caused a 52.14% reduction in the cost of driving one mile in 1996 and 11.36% in 2020, which we calculate by dividing the price of gasoline (\$1.25 per gallon in 1996 and \$2.27 per gallon in 2020) by the vehicle's fuel economy.¹⁸⁵ We calculate this rebound relative to the cost of driving one mile with the average retired vehicle. Using an elasticity of VMT with respect to the cost of driving one mile of -0.2221 (Small & Van Dender 2007), we find that the fuel

¹⁸⁴Although BAAQMD identifies dirty vehicles based on the vehicle's age, older vehicles will typically have lower fuel economies.

¹⁸⁵The fleet average fuel economy in 1996 was 20.204 miles per gallon and 23.112 miles per gallon in 2020.

economy improvement caused drivers to travel an additional 1,856.06 miles (11.58% increase) over the acceleration period in 1996 and 404.55 miles (2.52% increase) in 2020. Including the rebound in 1996, the replacement vehicle traveled 4,681.92 miles annually for the first three years and 3,837.89 miles in the final 299.2 days, for a total VMT of 17,883.66 miles. In 2020, the replacement vehicle traveled 4,301.92 miles annually for the first three years and 3,526.39 miles in the final 299.2 days, for a total VMT of 16,432.15 miles.

In 1996, the fleet average vehicle that replaces the 26-year-old vehicle imposed \$0.692 per gallon in local pollution damages, \$0.634 per gallon in global damages, and \$0.071 per mile in local driving damages, all expressed in 1996 dollars and accounting for vehicle decay. We account for rising the social costs of greenhouse gases over the 3.82 years accelerated. Since the fleet average vehicle is approximately 10 years old, we account for vehicle decay over this period; as a result, local pollution damages rise each year. Local driving damages do not change annually. See Appendix C.4.1 for details on how we account for vehicle decay. In 1997, the vehicle generated \$0.712 per gallon in local pollution damages and \$0.659 per gallon in global pollution damages. In 1998, the vehicle generated \$0.736 per gallon in local pollution damages and \$0.684 per gallon in global pollution damages. In 1999, the vehicle generated \$0.755 per gallon in local pollution damages and \$0.709 per gallon in global pollution damages. With a fuel economy of 20.204 miles per gallon, the vehicle consumed 231.73 gallons annually for the first three years and 189.95 gallons in the final 299.2 days, accounting for the rebound in VMT. Discounting after the first year, the replacement vehicle would have generated \$575.89 (\$516.12 + \$59.77) in discounted global damages, with \$516.12 in damages coming from the initial VMT and \$59.77 from the rebound effect. The replacement vehicle also generated \$621.50 (\$557.00 + \$64.50) in discounted local pollution damages and \$1,238.68 (\$1,110.12 + \$128.56) in discounted local driving damages. With a pre-tax nominal per-gallon markup of \$0.519 in 1996, the vehicle generated \$446.61 (\$400.25 + \$46.35) in discounted pre-tax profits for gasoline producers, or \$352.82 in post-tax profits with an assumed 21% corporate tax rate (Watson 2022).

In 2020, the fleet average vehicle that replaces the 26-year-old vehicle imposed \$0.226 per gallon in local pollution damages, \$1.891 per gallon in global damages, and \$0.118 per mile in local driving damages, all expressed in 2020 dollars and accounting for vehicle decay. We account for rising the social costs of greenhouse gases over the 3.82 years accelerated. Since the fleet average vehicle is approximately 10 years old, we account for vehicle decay over this period; as a result, local pollution damages rise each year. Local driving damages do not change annually. See Appendix C.4.1 for details on how we account for vehicle decay. In 2021, the vehicle generated \$0.232 per gallon in local pollution damages and \$1.931 per gallon in global pollution damages. In 2022, the vehicle generated \$0.238 per gallon in local pollution damages and \$1.962 per gallon in global pollution damages. In 2023, the vehicle generated \$0.243 per gallon in local pollution damages and \$2.003 per gallon in global pollution damages. With a fuel economy of 23.112 miles per gallon, the vehicle consumed 186.13 gallons annually for the first three years and 152.58 gallons in the final 299.2 days, accounting for the rebound in VMT. Discounting after the first year, the replacement vehicle would have generated \$1,343.39 (\$1,310.32 + \$33.07) in discounted global damages, with \$1,310.32 in damages coming from the initial VMT and \$33.07 from the rebound effect. The replacement vehicle also generated \$161.81 (\$157.82 + \$3.98) in discounted local pollution damages and \$1,878.84 (\$1,832.59 + \$46.26) in discounted local driving damages. With a pre-tax nominal per-gallon markup of \$0.613 in 2020, the vehicle generated \$423.47 (\$413.04 + \$10.43) in discounted pre-tax profits for gasoline producers, or \$334.54 in post-tax profits with an assumed 21% corporate tax rate (Watson 2022).

Total WTP Assuming all consumers are inframarginal to the subsidy, consumers are willing to pay \$455.47 in 1996 and \$977.11 in 2020, both in nominal dollars. We assume everyone who received the subsidy changed their behavior by accelerating the retirement of their old vehicle, so scaling the benefits from acceleration by the share of consumers that are marginal (100%) returns the benefits from acceleration, or the difference between damages from the retired and replacement vehicles.

In 1996, the old, retired vehicle generated \$1,143.08 in global damages, \$2,470.72 in local pollution damages, \$1,116.18 in local driving damages, and \$660.68 in post-tax profits for gasoline producers. The fleet average replacement vehicle generated \$575.89 in global damages, \$621.50 in local pollution damages, \$1,238.68 in local driving damages, and \$352.82 in post-tax profits for gasoline producers, all accounting for the rebound effect. Taking the difference between these vehicles, we find that the BAAQMD Vehicle Buyback Program resulted in \$567.19 in global benefits, \$1,849.22 in local pollution benefits, -\$122.50 in local driving benefits, and -\$307.86 in post-tax profits. Producers have a negative WTP for the policy since accelerating the retirement of the less fuel efficient vehicle decreases their profits. We adjust society's willingness to pay for the share of global benefits that do not flow to the US government (0.981), which yields an adjusted willingness to pay of \$556.33. Society's willingness to pay for all local damages (\$1,726.72) is the sum of society's willingness to pay for pollution damages and driving damages. Society's willingness to pay for the rebound in pollution and driving damages is -\$251.69 ($-\$59.77 \times 0.981 + -\$64.50 + -\128.56) in 1996. Combining components, we calculate a total WTP of \$2,430.65 ($\$455.47 + -\$307.87 + \$556.33 + \$1,726.72$) in 1996.

In 2020, the old, retired vehicle generated \$1,529.50 in global damages, \$487.86 in local pollution damages, \$1,843.44 in local driving damages, and \$368.14 in post-tax profits for gasoline producers. The fleet average replacement vehicle generated \$1,343.39 in global damages, \$161.81 in local pollution damages, \$1,878.84 in local driving damages, and \$334.54 in post-tax profits for gasoline producers, all accounting for the rebound effect. Taking the difference between these vehicles, we find that the BAAQMD Vehicle Buyback Program resulted in \$186.11 in global benefits, \$326.05 in local pollution benefits, -\$35.40 in local driving benefits, and -\$33.60 in post-tax profits. Producers have a negative WTP for the policy since accelerating the retirement of the less fuel efficient vehicle decreases their profits. We adjust society's willingness to pay for the share of global benefits that do not flow to the US government (0.981), which yields an adjusted willingness to pay of \$182.55. Society's willingness to pay for all local damages (\$290.65) is the sum of society's willingness to pay for pollution damages and driving damages. Society's willingness to pay for the rebound in pollution and driving damages is -\$82.68 ($-\$33.07 \times 0.981 + -\$3.98 + -\46.26) in 2020. Combining components, we calculate a total WTP of \$1,416.71 ($\$977.11 + -\$33.60 + \$182.55 + \290.65) in 2020.

Cost The total cost is comprised of the transfer, administrative costs, and fiscal externalities. The transfer is equal to the size of the subsidy, or \$455.47 in 1996 and \$977.11 in 2020 (Sandler 2012) in nominal dollars, both in nominal dollars. The author reports that administrative costs were \$240 (in 2000 dollars). We adjust this cost for inflation and combine with the transfer to calculate the total program cost, or \$674.10 in 1996 and \$1,337.88 in 2020. All fiscal externalities are scaled by the share of consumers that are marginal, but we abstract from this step in this section since we assume the share of consumers that are marginal is 100%.

Accelerating the retirement of the older, less fuel efficient vehicle results in less gas tax revenue for the government. In 1996, the average total gas tax levied by the federal government

and states was \$0.37 per gallon (FHWA 2021, 2020). In 2020, the average total gas tax was \$0.46 per gallon. In 1996, the old, retired vehicle generated \$595.82 in tax revenue while the new, replacement vehicle generated \$318.18 (\$285.16 + \$33.02), for a change in gas tax revenue of \$277.64. In 2020, the old, retired vehicle generated \$353.64 in tax revenue while the new, replacement vehicle generated \$321.36 (\$313.45 + \$7.91), for a change in gas tax revenue of \$32.28. This fiscal externality is positive since accelerating vehicle retirement and collecting less in gas tax revenue imposes a cost on the government.

The government also loses corporate tax revenue when gasoline producers lose profits. In 1996, gasoline producers faced pre-tax lost profits of \$389.70 (\$836.31 - \$446.61) as a result of retirement acceleration. With a 21% effective corporate tax rate (Watson 2022), the government lost \$81.84 in corporate tax revenue in 1996 as a result of the BAAQMD Vehicle Buyback Program. In 2020, gasoline producers lost \$42.53 (\$466.00 - \$423.47) in pre-tax profits from acceleration for a fiscal externality of \$8.93.

We place the share of global benefits that flow to the US government in the denominator of the MVPF. The BAAQMD Vehicle Buyback Program generated \$567.19 in global benefits in 1996 before adjusting for the share of benefits that flow to the US government. With 1.92% of global benefits flowing to the US government as long-run revenue, the program cost the government -\$10.86 (in nominal dollars) in 1996 by abating greenhouse gases. In 2020, Cash for Clunkers generated \$186.11 in total global benefits, for a climate fiscal externality of -\$3.56 (in nominal dollars). This fiscal externality is negative because it raises revenue for the government.

Summing the program cost and the fiscal externalities, the BAAQMD Vehicle Buyback Program imposed a net cost on the government of \$1,022.72 in 1996 and \$1,375.53 in 2020. The fiscal externality from lost tax revenue was \$359.48 in 1996 and \$41.21 in 2020, while the climate fiscal externality cost the government -\$10.86 in 1996 and -\$3.56 in 2020.

MVPF Dividing the total WTP calculated above (\$2,430.65) by the total cost (\$1,022.72), we form an MVPF of 2.38 in 1996. Using our 2020 estimates of a total WTP of \$1,416.71 and a net government cost of \$1,375.53, we obtain an MVPF of 1.030 [1.025, 1.036].

We note that the in-context MVPF is sensitive to our assumptions about vehicle quality. If we instead assumed that vehicles released before 1975 could decay, we obtain an MVPF of 5.14. If we allow for vehicle decay and assume the average vehicle retired was 30 years old, we obtain an MVPF of 6.80. Since we allow for no vehicle decay for vehicles released before 1975 and assume retired vehicles are the youngest age possible, our average MVPF is likely a lower bound.

E.8 Home Energy Reports

E.8.1 Home Energy Reports (17 RCTs)

Allcott (2011) evaluates randomized natural field experiments on 600,000 households run by the company Opower. The treatment households were sent Home Energy Report (HER) letters comparing their electricity use to that of their neighbors. The control group were not sent any letters. Using 17 separate RCTs around the U.S., the author finds treatment effects on energy consumption ranging from 1.4 to 3.3%. We take these results and estimate an average

treatment effect across the RCTs below.

Throughout this section, the “in-context” specification will mean the US (since the RCTs occur in multiple states) in 2009, the last year of any of the experiments. However, the experiments do not constitute a representative sample of the US. There are no experiments in the South for example. We cannot make our in-context specification any more precise though since the experiments’ geographies are just labeled by region and not by state.

WTP The WTP for a HER is the sum of the environmental and other market externalities discussed in Section 5. To compute each externality, we first calculate a weighted average number of kWh of electricity reduced annually across the experiments. The ATEs are weighted by the number of treated households in each experiment reported in Table 1 of Allcott (2011) and those weights are divided evenly among experiments with multiple treatment groups when applicable. Figure 5 reports the baseline daily electricity usage of households in each experiment, which we use to convert the ATEs into level reductions in electricity. Our average kWh reduction per year is 243.26 kWh. In general, each externality has the form of this average reduction times the societal willingness to pay for one additional kWh of electricity normalized by the average program cost (see calculation of program cost below).

Transfer No transfer is included in this MVPF calculation. One reason a nudge might change a person’s behavior is because they were close to indifferent about the choice to begin with. Thus, we model the willingness to pay for a nudge to be zero for those who are directly nudged.

Environmental Externalities Using AVERT’s reported marginal emissions coefficients (see Appendix C for more details), we estimate the global pollutants and local pollutants from the average HER. The global damages are \$34.20 in 2020 and \$23.17 in context and the local damages are \$3.88 in 2020 and \$22.81 in context. After normalizing by the program cost, the global damages are \$3.87 in 2020 and \$3.17 in context while the local damages are \$0.44 in 2020 and \$3.12 in context.

The previous calculation assumes that the decreased consumption of electricity does not affect grid-wide electricity prices. We now consider some general equilibrium effects where the decreased use of electricity decreases the price of electricity and thus increases electricity consumption (see Appendix D for calculation). With this rebound effect of about 20% multiplied by the previously calculated normalized global and local damages from electricity consumption, we now have a decrease in damages avoided due to the HERs of -0.84 in 2020 and -1.23 in context from more electricity consumption.

Thus, the WTP for avoiding the environmental damages from the reduced electricity consumption is 3.53 in 2020 and 5.10 in context.

Profits Lastly for WTP, we estimate the utilities’ WTP for the HERs. Since the electricity markets are heavily regulated to have fixed levels of markups, we believe this is an important component to consider. Using the previously estimated decrease in electricity consumption,

we calculate the producers' WTP as the markup multiplied by the change in electricity and normalized by the program cost. We have 243.26 kWh of electricity multiplied by a markup of 0.01 in 2020 and 0.01 in context, and normalized to get -0.24 in 2020 and -0.26 in context.

With the various components, we can now calculate the total WTP of 3.22 for 2020 and 4.79 for in-context.

Cost The net government cost of a \$1 mechanical increase in HER spending is equal to the \$1 plus the fiscal externalities induced by the response to the subsidies. These include changes in profits tax collected and the climate fiscal externality from changes in CO_2 emissions. We discuss each in turn.

Program Cost Similar to how we calculate the average kWh reduction per year, we also estimate a treated households-weighted average of the costs of the experiments. Figure 5 of Allcott (2011) shows the cost-effectiveness of each experiment in units of cents per kWh. We inflation adjust the costs for experiments implemented in 2008 and then take the weighted average. The final program cost is \$8.83 in 2020 and \$7.32 in context. This cost is per household per year not per nudge as experiments varied in the number of nudges sent between 4 and 12. The program cost is normalized to 1 before being added to the other costs.

Profits Tax FE We calculate a profits tax fiscal externality that accounts for profits to publicly-owned utilities and corporate taxes on privately-owned utilities as described in Appendix C.2.3. The FE is the annual kWh reduction (including the rebound effect) multiplied by the tax rate of 0.01 in 2020 and 0.01 in context, and normalized by the net MSRP to get 0.13 in 2020 and 0.14 in context.

Climate FE Finally, the climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Consistent with the GIVE model, we assume that 50% of the social cost of carbon is from lost productivity. Thus, by abating CO_2 , a program increases productivity and thus tax revenue to the government. We use DICE's estimate that 15% of that avoided GDP loss would flow to the US and an average tax rate in the US of 25.54%. The climate FE as $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the global environmental externality including the rebound. Taking that piece and multiplying by 4.5% gives us -0.06 in 2020 and -0.05 in context.

Thus, our final cost is 1.07 in 2020 and 1.09 in context, which gives us our MVPFs of 3.01 and 4.40.

E.8.2 Opower Electricity Program Evaluations (166 RCTs)

Oracle compiled Opower reports between 2010 and 2018 evaluating randomized natural field experiments on residential utility customers. The treatment households were sent Home Energy Report letters comparing their electricity use to that of their neighbors. Using 166 RCTs covered in these reports, we estimate an average treatment effect across the RCTs below.

Throughout this section, the “in-context” specification will mean the US (since the RCTs occur in multiple states) in 2012, the average year from the evaluations.

WTP The WTP for a Home Energy Report (HER) is the sum of the environmental and other market externalities discussed in Section 5. To compute each externality, we first calculate a weighted average number of kWh of electricity reduced annually across the experiments. The ATEs are weighted by the treatment group size and the population of the Census region that the experiment occurred in. All of the reports include the baseline daily electricity usage of households in each experiment, which we use to convert the ATEs into level reductions in electricity. Our average kWh reduction per year is 160.73 kWh. In general, each externality has the form of this average reduction times the societal willingness to pay for one additional kWh of electricity normalized by the average program cost (see calculation of program cost below).

Transfer No transfer is included in this MVPF calculation. One reason a nudge might change a person’s behavior is because they were close to indifferent about the choice to begin with. Thus, we model the willingness to pay for a nudge to be zero for those who are directly nudged.

Environmental Externalities Using AVERT’s reported marginal emissions coefficients (see Appendix C for more details), we estimate the global pollutants and local pollutants from the average HER. The global damages are \$22.59 in 2020 and \$17.46 in context and the local damages are \$2.56 in 2020 and \$10.44 in context. After normalizing by the program cost, the global damages are \$3.25 in 2020 and \$2.83 in context while the local damages are \$0.37 in 2020 and \$1.69 in context.

The previous calculation assumes that the decreased consumption of electricity does not affect grid-wide electricity prices. We now consider some general equilibrium effects where the decreased use of electricity decreases the price of electricity and thus increases electricity consumption (see Appendix D for calculation). With this rebound effect of about 20% multiplied by the previously calculated normalized global and local damages from electricity consumption, we now have a decrease in damages avoided due to the HERs of -0.71 in 2020 and -0.89 in context from more electricity consumption.

Thus, the WTP for avoiding the environmental damages from the reduced electricity consumption is 2.96 in 2020 and 3.68 in context.

Profits Lastly for WTP, we estimate the utilities’ WTP for the HERs. Since the electricity markets are heavily regulated to have fixed levels of markups, we believe this is an important component to consider. Using the previously estimated decrease in electricity consumption, we calculate the producers’ WTP as the markup multiplied by the change in electricity and normalized by the net MSRP. We have 160.73 kWh of electricity multiplied by a markup of 0.01 in 2020 and 0.01 in context, and normalized to get -0.20 in 2020 and -0.21 in context.

With the various components, we can now calculate the total WTP of 2.70 for 2020 and 3.43 for in-context.

Cost The net government cost of a \$1 mechanical increase in HER spending is equal to the \$1 plus the fiscal externalities induced by the response to the subsidies. These include changes in profits tax collected and the climate fiscal externality from changes in CO_2 emissions. We discuss each in turn.

Program Cost Allcott (2011) estimates that the cost of mailing and printing one HER was approximately \$1 in 2009. Each report mentions how many nudges each household received per year and we take a treatment size and region population-weighted average of the annual number of nudges. We take this, multiply it by \$1, and inflation adjust it to estimate our program cost. The final program cost is \$6.96 in 2020 and \$6.17 in context. The program cost is normalized to 1 before being added to the other costs.

Profits Tax FE We calculate a profits tax fiscal externality that accounts for profits to publicly-owned utilities and corporate taxes on privately-owned utilities as described in Appendix C.2.3. The FE is the annual kWh reduction (including the rebound effect) multiplied by the tax rate of 0.01 in 2020 and 0.01 in context, and normalized by the net MSRP to get 0.11 in 2020 and 0.11 in context.

Climate FE Finally, the climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Consistent with the GIVE model, we assume that 50% of the social cost of carbon is from lost productivity. Thus, by abating CO_2 , a program increases productivity and thus tax revenue to the government. We use DICE's estimate that 15% of that avoided GDP loss would flow to the US and an average tax rate in the US of 25.54%. The climate FE as $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the global environmental externality including the rebound. Taking that piece and multiplying by 4.5% gives us -0.05 in 2020 and -0.04 in context.

Thus, our final cost is 1.06 in 2020 and 1.07 in context, which gives us our MVPFs of 2.55 and 3.20.

E.8.3 Peak Energy Reports

Brandon et al. (2019) examine the response of household electricity consumption to social nudges during peak load events. They use data from a natural field experiment with 42,100 households and find that peak energy reports (PERs) reduce peak load electricity consumption by 3.8% when implemented in isolation.

Throughout this section, the “in-context” specification will mean California in 2014, which is the time and geography analyzed in the paper.

We compute three MVPFs for each specification: one with a marginal cost of electricity during peak load events of \$1 per kWh (high marginal cost), one with a marginal cost of \$0.50

per kWh (low marginal cost), and one where we assume that peak load events lead to blackouts with 100% probability so the marginal kWh of electricity that a PER reduces one person from using is shifted to another person (value of lost load marginal cost). That shifted kWh is valued at \$4.29 which is the WTP to avoid a blackout.

WTP The WTP for a PER is the sum of the environmental and other market externalities as discussed in Section 5. To compute each externality, we take the average treatment effect of 0.038 and turn this into a level reduction in electricity using the baseline consumption of the PER treatment group from Figure 5 of Brandon et al. (2019). With a baseline electricity use per hour of 0.65 kWh, the average kWh reduction in each hour is 0.02. Peak load events were called during five-hour periods on three different days for the experiment, so the kWh reduction per PER was 0.12 kWh. In general, each externality has the form of this average reduction times the societal willingness to pay for one additional kWh of electricity normalized by the average program cost (see calculation below).

Transfer No transfer is included in this MVPF calculation. One reason a nudge might change a person’s behavior is because they were close to indifferent about the choice to begin with. Thus, we model the willingness to pay for a nudge to be zero for those who are directly nudged.

However, in the value of lost load marginal cost MVPF, consumers have a WTP of \$4.29 multiplied by the kWh reduction per PER and normalized by the program cost which is 5.30 in 2020 and 4.85 in context.

Environmental Externalities During peak load events, we assume the marginal kWh of electricity is being provided through burning coal since coal possesses “a ‘flexible’ characteristic that enables rapid adjustments in output to accommodate fluctuations in demand” (Jacob et al. 2023). Using eGrid’s emissions factors in 2020, we calculate the local emissions per kWh of electricity to be 0.05 and the global emissions per kWh to be 0.05. We multiply these emissions numbers by the kWh reduction to get \$0.01 in local damages and \$0.02 in global damages. After removing the portion of the global damages that will accrue to the US government via increased GDP from avoiding carbon emissions, which is 1.9% of the total amount, and dividing both amounts by the program cost, we have a global environment component of 0.23 and a local component of 0.06. Note that there are no environmental externalities when the marginal cost is the value of lost load.

We can’t compute a rebound effect for this policy because during peak load events supply is inelastic (see Appendix D).

Profits Lastly for WTP, we estimate the utilities’ WTP for the PERs. During a peak load event, we assume the marginal cost of providing one kWh of electricity is much higher than the price utilities receive. Our low (\$0.50) and high (\$1) marginal costs come from Department of Market Monitoring (2021). The average price received per kWh in 2020 across the US is \$0.13 so assuming that 28% of utilities are private and that 10% of private utility profits

are taxed, the WTP of producers is 0.70 in 2020 with the high marginal cost and 0.70 in context.

When we assume the low marginal cost, the WTP of producers is 0.29 in 2020 and 0.29 in context. There is no producers' WTP when the marginal cost is the value of lost load.

With the various components, we can now calculate the total WTP with a high marginal cost of 0.99 for 2020 and 0.94 for in-context. With a low marginal cost, the 2020 WTP is 0.59 and the in-context WTP is 0.54. Using the value of lost load marginal cost gives a 2020 WTP of 5.30 and an in-context WTP of 4.85.

Cost The net government cost of a \$1 mechanical increase in PER spending is equal to the \$1 plus the fiscal externalities induced by the response to the nudges. These include changes in profits tax collected and the climate fiscal externality from changes in greenhouse gas emissions. We discuss each in turn.

Program Cost We assume each PER costs \$0.10. This is normalized to 1 before being added to the other costs.

Profits Tax FE We calculate a profits tax fiscal externality that accounts for the taxation of profits arising from imperfect competition and the revenues to public utilities as described in Appendix C.2.3. Using the same \$1 per kWh marginal cost as in the utilities' WTP calculation, the public portion of the FE is $(0.13 - 1) \cdot 0.28 \cdot 0.12 = -0.03$ and the private portion is $(0.13 - 1) \cdot (1 - 0.28) \cdot 0.1 \cdot 0.12 = -0.01$. Adding these together and dividing by the program cost, gives us -0.38 in 2020 and -0.38 in context.

When we assume the low marginal cost, the profits tax FE is -0.16 in 2020 and -0.16 in context. There is no profits tax FE when the marginal cost is the value of lost load.

Climate FE Finally, the climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Consistent with the GIVE model, we assume that 50% of the social cost of carbon is from lost productivity. Thus, by abating CO_2 , a program increases productivity and thus tax revenue to the government. We use DICE's estimate that 15% of that avoided GDP loss would flow to the US and an average tax rate in the US of 25.54%. The climate FE as $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the global environmental externality. Taking that piece and multiplying by 1.9% gives us -0.00 in 2020 and -0.00 in context. Note that there is no climate fiscal externality when the marginal cost is the value of lost load.

Thus, our final total cost when using a high marginal cost is 0.62 in 2020 and 0.62 in context, which gives us our MVPFs of 1.60 and 1.52. With a low marginal cost, the total cost is 0.84 in 2020 and 0.84 in context with MVPFs of 0.70 in 2020 and 0.64 in context. In the value of lost load specification the total cost in 2020 is 1.00 and in context is 1.00, so the MVPF in 2020 is 5.30 and in context is 4.85.

E.8.4 Opower Natural Gas Program Evaluations (52 RCTs)

Oracle compiled Opower reports between 2010 and 2018 evaluating randomized natural field experiments on residential utility customers. The treatment households were sent Home Energy Report letters comparing their natural gas use to that of their neighbors. Using 52 RCTs covered in these reports, we estimate an average treatment effect across the RCTs below.

Throughout this section, the “in-context” specification will mean the US (since the RCTs occur in multiple states) in 2012, the average year from the evaluations.

WTP The WTP for a Home Energy Report (HER) is the sum of the environmental and other market externalities discussed in Section 5. To compute each externality, we first calculate the average number of MMBtu of natural gas reduced annually across the experiments. The ATEs are weighted by the treatment group size and the population of the census region in which the experiment occurred. All of the reports include households’ baseline daily natural gas usage in each experiment, which we use to convert the ATEs into level reductions in natural gas. Our average MMBtu reduction per year is 0.94 MMBtu. In general, each externality has the form of this average reduction times the societal willingness to pay for one additional MMBtu of natural gas normalized by the average program cost (see calculation of program cost below).

Transfer No transfer is included in this MVPF calculation. One reason a nudge might change a person’s behavior is because they were close to indifferent about the choice to begin with. Thus, we model the willingness to pay for a nudge to be zero for those who are directly nudged.

Environmental Externalities Using emissions factors from the EPA’s eGRID from 2011-2020 for CO_2 , CH_4 , and N_2O (see Appendix C for more details), we estimate the global pollutants from the average HER. The global damages are \$9.46 in 2020 and \$7.03 in context. eGRID does not report emissions factors for local pollutants associated with natural gas combustion. After normalizing by the program cost, the global damages are \$0.95 in 2020 and \$0.80 in context.

The previous calculation assumes that the decreased consumption of natural gas does not affect market-wide natural gas prices. We now consider some general equilibrium effects where the decreased use of natural gas decreases the price of natural gas and thus increases natural gas consumption (see Appendix D for calculation). With this natural gas rebound effect of about 11.76% multiplied by the previously calculated normalized global damages from natural gas consumption, we now have a decrease in damages avoided due to the HERs of -0.11 in 2020 and -0.09 in context from more natural gas consumption.

Thus, the WTP for avoiding the environmental damages from the reduced natural gas consumption is in 2020 and in context.

Profits Lastly for WTP, we estimate the utilities’ WTP for the HERs. Due to imperfect competition in the natural gas market, there exist markups above the average economy-wide

markups. Using the previously estimated decrease in natural gas consumption, we calculate the producers' WTP as the markup multiplied by the change in natural gas and normalized by the program cost. We have 0.94 MMBtu of natural gas multiplied by a markup of 4.40 in 2020 and 4.50 in context, and normalized to get -0.37 in 2020 and -0.42 in context.

With the various components, we can now calculate the total WTP of 0.47 for 2020 and 0.28 for in-context.

Cost The net government cost of a \$1 mechanical increase in HER spending is equal to the \$1 plus the fiscal externalities induced by the response to the subsidies. These include changes in profits tax collected and the climate fiscal externality from changes in greenhouse gas emissions. We discuss each in turn.

Program Cost Allcott (2011) estimates that the cost of mailing and printing one HER was approximately \$1 in 2009. Each Opower report mentions how many nudges each household received per year and we take a treatment size and region population-weighted average of the annual number of nudges. We take this, multiply it by \$1, and inflation adjust it to estimate our program cost. The final program cost is \$9.96 in 2020 and \$8.83 in context. The program cost is normalized to 1 before being added to the other costs.

Profits Tax FE We calculate a profits tax fiscal externality that accounts for profits arising from imperfect competition as described in Appendix C.2.3. The FE is the annual MMBtu reduction (including the rebound effect) multiplied by the tax rate of 0.75 in 2020 and 0.76 in context, and normalized by the net MSRP to get 0.06 in 2020 and 0.07 in context.

Climate FE Finally, the climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Consistent with the GIVE model, we assume that 50% of the social cost of carbon is from lost productivity. Thus, by abating CO_2 , a program increases productivity and thus tax revenue to the government. We use DICE's estimate that 15% of that avoided GDP loss would flow to the US and an average tax rate in the US of 25.54%. The climate FE as $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the global environmental externality including the rebound. Taking that piece and multiplying by 4.5% gives us -0.02 in 2020 and -0.01 in context.

Thus, our final cost is 1.05 in 2020 and 1.06 in context, which gives us our MVPFs of 0.45 and 0.26.

E.9 Other Nudges

E.9.1 Solarize Connecticut

Gillingham & Bollinger (2021) examine intensive community-level campaigns in Connecticut to increase the adoption of solar photovoltaic (PV) installations. The "Solarize" program involves a competitive bidding process to choose a single installer for a campaign, volunteer promoters to provide information about solar PV, community-based recruitment, a group pricing discount,

and a limited time frame for the campaign. In total, the effect of these packaged components in treated municipalities is an increase of 37.128 solar PV installations per municipality and a decrease in price by \$0.46 per watt. In spillover municipalities, the effect is a decrease of 0.616 solar PV installations per municipality and a decrease in price by \$0.05 per watt.

Throughout this section, the “in-context” specification will mean Connecticut in 2012, the first year of the campaigns.

WTP The WTP for the Solarize campaign is the sum of the environmental and other market externalities as discussed in Section 4. To compute each externality, we take the average treatment effect on the treated municipalities of 37.128 and on the spillover municipalities of -0.616 and turn this into an annual reduction in electricity using the average kWh per install. We assume an average system capacity of 7.15 kW for 2020 (Ramaswamy et al. 2022) and 6.97 kW for in-context (Gillingham & Tsvetanov 2019) and an average annual output of 10,296 kWh for 2020 and 8,997 kWh for in-context for a system of the average capacity. The average annual output for 2020 assumes a 20/214 tilt/azimuth from Fredonia, Kansas, inverter efficiency of 96%, and system losses of 14%. The in-context value comes from Gillingham & Tsvetanov (2019).

For 2020 this translates into 375,928 kWh of energy switched to solar each year ($(37.128 - 0.616) \cdot 10296$) and in context, it is 328,486 kWh of energy each year. In general, each externality has the form of the present discounted value of this average reduction times the lifetime of a solar panel times the societal willingness to pay for one solar system’s worth of electricity normalized by the average program cost (see calculation below).

Consumers Since this policy is not a direct subsidy to consumers, we don’t consider the policy spending a mechanical transfer. We assume that consumers do not value the actual marketing spending from the nudge because we think they are indifferent. However, consumers still receive private benefits from the price reductions resulting from the group pricing discount. Gillingham & Bollinger (2021) found a \$0.46 per watt decrease in solar prices. Using the control group average price of \$4.63, this translates into a 9.94% decrease in prices. For 2020, we use an average cost per watt of \$3.13 from the National Renewable Energy Laboratory, and thus, consumers save \$0.31 per watt. We also assume that consumers value the price discount at 50%. The WTP for consumers in the treated municipalities is then 50% times the average system capacity of 7.15 kW (which converted to watts is 7,150 W) times the installation treatment effect of 37.1 times the price reduction, which equals 41,276. Using the \$0.46 price savings and an average system capacity of 6.972 kW, the WTP for consumers in the treated municipalities in context is 57,744. Normalized by the program cost, the WTP for consumers in the treated municipalities is 1.15 in 2020 and 1.81 in context.

In addition, the authors estimate spillover effects on prices and installations in neighboring municipalities. They estimate a decrease in price by \$0.05 per watt and a decrease of 0.616 solar PV installations per municipality. This translates to a normalized consumer WTP of -0.002 in 2020 and -0.003 in context. Thus, the total normalized consumer WTP for both types of municipalities is 1.14 in 2020 and 1.81 in context.

Environmental Externalities Using the annual kWh of energy switched to solar from above; we calculate the global and local pollutants over the 25-year lifetime of a solar system using AVERT's reported marginal emissions coefficients (see Appendix C.2 for more details). In 2020, the global damages are \$608,377 and the local damages are \$79,215. In context, the global damages are \$391,589 and the local damages are \$51,503. After removing the portion of the global damages that will accrue to the US government via increased GDP from avoiding carbon emissions, which is 1.9% of the total amount, and dividing both amounts by the program cost, we have a global environment component of 16.58 and a local component of 2.20 in 2020 and a global environment component of 12.03 and a local component of 1.61 in context.

Additionally, we estimate a 19.59% rebound effect from increased electricity supply (see Appendix D). This counteracts the abatement of global and local pollutants and leads to a social WTP of -3.25 for the global rebound in 2020 and -0.43 for the local rebound in 2020. In context, these rebounds are -2.36 for global and -0.32 for local.

Profits Lastly for WTP, we estimate the utilities' WTP for the lost profits caused by the solar panels. We use a markup per kWh of 0.01 in 2020 and 0.05 in context. Taking the present discounted value of the stream of profits lost over the lifetime of a solar panel and accounting for the rebound effect, we have a producers' WTP of -1.84 in 2020 and -7.62 in context.

With the various components, we can now calculate the total WTP of 12.82 for 2020 and 4.00 for in-context.

Cost The net government cost of a \$1 mechanical increase in Solarize program spending is equal to the \$1 plus the fiscal externalities induced by the response to the program. These include changes in profits tax collected and the climate fiscal externality from changes in greenhouse gas emissions. We discuss each in turn.

Program Cost According to Gillingham and Bollinger, the Solarize program costs \$860 per new installation. Inflation-adjusted to 2020\$ and multiplied by the 37.1 installs gives us a program cost of \$35,999 in 2020. For in-context, this value is \$31,930. This is normalized to \$1 before being added to the other costs.

State and Federal FEs Each solar system that the Solarize program induces people to purchase leads to more spending from the state and federal government subsidy programs. In 2020, there was a 26% subsidy for solar panels from the federal government, and in context, there was a 30% federal subsidy and a Connecticut rebate of \$1.25 per watt.

For 2020, the fiscal externality is just from federal spending and is 26% of the price treated consumers pay multiplied by the number of installations in their municipalities plus 26% of the price spillover consumers pay multiplied by the number of installations in their municipalities. We assume a baseline cost per watt of \$3.13 (NREL 2022b) and account for the price decrease seen in treated municipalities to get \$2.67 per watt. We convert this to a per-system cost by multiplying by 1,000 and the average system capacity, which gives \$20,156 as the total cost a treated consumer pays for an average solar system. Finally, we multiply by the treatment effect and 26% to get \$194,572 of federal spending for the treated municipalities. We apply the

same logic to the spillover municipalities using the relevant treatment effects on installations and price to get -3,546. Adding the two pieces together and normalizing by the program cost gives us 5.31 for the federal fiscal externality in 2020.

In context, there is both a federal and a state fiscal externality. The state FE is \$1.25 multiplied by the system capacity in watts multiplied by the number of installations in both treated and spillover municipalities. This value when normalized by the program cost is 9.97. The federal FE is calculated the same way in context as 2020 except using 30% as the subsidy rate and \$4.63 (converted from 2014\$ to 2012\$) as the original cost per watt before the group discount. This value when normalized by the program cost is 9.66.

Profits Tax FE We calculate a profits tax fiscal externality that accounts for the taxation of profits arising from imperfect competition and the revenues to public utilities as described in Appendix C.2.3. Using the same annual kWh reduced as in the environmental externalities calculation and a government revenue per kWh value of 0.01 in 2020 and 0.03 in-context, the present discounted value of the profits tax FE over the lifetime of a solar system is 36,508 in 2020 and 133,804 in context. Normalizing by the program cost gives 1.01 in 2020 and 4.19 in context.

Climate FE Finally, the climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Consistent with the GIVE model, we assume that 50% of the social cost of carbon is from lost productivity. Thus, by abating CO_2 , a program increases productivity and thus tax revenue to the government. We use DICE's estimate that 15% of that avoided GDP loss would flow to the US and an average tax rate in the US of 25.54%. The climate FE is $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the global environmental externality net of the rebound effect. Taking that piece and multiplying by 1.9% gives us -0.23 in 2020 and -0.17 in context.

Thus, our final total cost is 7.09 in 2020 and 24.65 in context, which gives us our MVPFs of 1.81 and 0.16.

E.9.2 Energize CT Home Energy Solutions Program Energy Audit

Gillingham & Tsvetanov (2018) evaluated a randomized trial in Connecticut that looked at the effect of personalized notecards sent 14 days before a scheduled residential energy audit on the probability of completing the audit. The goal of the notecards is to use social norms, salience, and a personal touch to nudge individuals toward following through on the energy audit. These audits involve professional home assessments to identify energy efficiency investments that a household can make to reduce energy consumption and save money on energy bills. The authors find that the households that received a notecard had a 1.1 percentage point higher likelihood of completing its audit on a given day, conditional on not yet having completed it. They converted this result to the effect of the treatment on the overall uptake during the study period and found that the notecards led to a 6.4 percentage points higher overall uptake for an average treated household in this experiment.

Throughout this section, the “in-context” specification will mean Connecticut in 2014, the place and year of the experiment.

WTP The WTP for the personalized notecards is the sum of the environmental and other market externalities as discussed in Section 4. To compute each externality, we calculate the number of kWh reduced annually on average by an audit, which is $0.05 \cdot 10,566 = 528.3$ kWh in 2020 and $0.05 \cdot 7,794 = 389.7$ in context. The paper assumes that an audit leads to a 5% energy reduction (section 4.3) and we get 2020 US annual energy use is 10,566 kWh from the EIA’s 2020 Residential Energy Consumption Survey. Energy consumption is fairly stable over time, so for the in-context specification, we use the 2020 value for Connecticut from the same survey, which is 7,794 kWh. In general, each externality has the form of the present discounted value of this average reduction times five (the paper assumes the impacts of an audit last for five years) times the societal willingness to pay for one audit’s worth of electricity savings normalized by the average program cost (see calculation below).

Transfer No transfer is included in this MVPF calculation. One reason a nudge might change a person’s behavior is because they were close to indifferent about the choice to begin with. Thus, we model the willingness to pay for a nudge to be zero for those who are directly nudged.

Environmental Externalities Using AVERT’s reported marginal emissions coefficients (see Appendix C.2 for more details), we estimate the global pollutants and local pollutants from one audit. The global damages are \$360.39 in 2020 and \$157.96 in context, and the local damages are \$55 in 2020 and \$37 in context. After normalizing by the program cost, the global damages are \$8.68 in 2020 and \$4.23 in context while the local damages are \$1.33 in 2020 and \$0.99 in context.

The previous calculation assumes that the decreased consumption of electricity does not affect grid-wide electricity prices. We now consider some general equilibrium effects where the decreased use of electricity decreases the price of electricity and thus increases electricity consumption. The rebound effect is calculated as $\frac{1}{1-\varepsilon_D/\varepsilon_S}$ where ε_D is the demand elasticity for electricity, which is -0.19 and comes from EIA (2023c) and ε_S is the supply elasticity for electricity, which is 0.78 (see Appendix D for calculation). With this rebound effect of about 20% multiplied by the previously calculated normalized global and local damages from electricity consumption, we now have a decrease in damages avoided due to the audits of -1.96 in 2020 and -1.02 in context from more electricity consumption.

Profits Lastly, for WTP, we estimate the utilities’ WTP for the audits. Since the electricity markets are heavily regulated to have fixed levels of markups, we believe this is an important component to consider. Using the previously estimated decrease in electricity consumption, we calculate the producers’ WTP as the markup multiplied by the change in electricity and normalized by the program cost. We have 528.30 kWh of electricity multiplied by a markup of 0.01 in 2020 and 0.05 in context and normalized to get -0.54 in 2020 and -1.89 in context.

With the various components, we can now calculate the total WTP of 7.51 for 2020 and 2.31 for in-context.

Cost The cost of \$1 mechanical increase in audit nudge spending is equal to the \$1 plus the fiscal externalities induced by the response to the subsidies. These include changes in profits tax collected and the climate fiscal externality from changes in CO_2 emissions. We discuss each in turn.

Program Cost Each notecard cost \$2.40 as reported in Section 4.3 of Gillingham & Tsvetanov (2018). We calculate the program cost as the cost to achieve one successful audit, so we need to determine how many cards need to be sent to get one audit. This is $1/0.064 = 15.57$ cards. Thus, the program cost in context is \$37.38 and when inflation-adjusted to 2020 is \$41.53. The program cost is normalized to 1 before being added to the other costs.

Profits Tax FE We calculate a profits tax fiscal externality that accounts for profits to publicly-owned utilities and corporate taxes on privately-owned utilities as described in Appendix C.2.3. The FE is the annual kWh reduction (including the rebound effect) multiplied by the tax rate of 0.01 in 2020 and 0.03 in context, and normalized by the program cost to get 0.31 in 2020 and 1.08 in context.

Climate FE Finally, the climate fiscal externality comes from the increased GDP due to decreasing carbon emissions. Consistent with the GIVE model, we assume that 50% of the social cost of carbon is from lost productivity. Thus, by decreasing CO_2 , a program increases productivity and thus tax revenue to the government. We use DICE's estimate that 15% of that avoided GDP loss would flow to the US and an average tax rate in the US of 25.54%. The climate FE is calculated as $0.15 \cdot 0.2554 \cdot 0.5 = 1.9\%$ of the global environmental externality, including the rebound. Taking those global damages and multiplying by 1.9% gives us -0.14 in 2020 and -0.07 in context.

Thus, our final cost is 3.55 in 2020 and 4.38 in context, which gives us our MVPFs of 2.12 and 0.53.

E.9.3 ENERGY STAR Rebate - Water Heaters (w/ Sales Agent Incentive)

Our MVPF for the combined sales agent incentive payment and water heater subsidy using estimates from Allcott & Sweeney (2017) is 1.14 [1.08, 1.15] in 2020 and 0.82 in-context. This MVPF follows the construction of the MVPF for water heater rebates also from Allcott & Sweeney (2017). The difference is that this MVPF includes the impact of a \$25 incentive payment to sales agents.

The sales agent incentive payment changes the original water heater MVPF in two ways. First, it increases the cost of the program by the size of the sales incentive (\$25) and the change in fiscal externalities. Second, the improved sales performance increases take-up of the water heater subsidy. Water heater take-up in the control group that were not offered the subsidy was 0.9%. The rebate increased take-up by 3.7 percentage points. The combination of the rebate and the sales incentive increased take-up by 21.9 percentage points. Therefore, the percent

of marginal recipients of the subsidy increased from 80.43% to 96.05%. The increase in the percent of marginal recipients increases the environmental impact of the subsidy and reduces the amount the recipients value the subsidy.

Cost Following the approach outlined above, the total cost of the sales incentive and rebate combination increases from \$121.58 to \$151.48 in 2020 and from \$110.09 to \$137.05 in-context.

WTP Since inframarginal recipients value the entire subsidy and marginal recipients value half of the subsidy, the increase in the share of marginal recipients from the sales incentive lowers the willingness to pay from program recipients. Since there are more additional recipients, the environmental externality and loss of producer profits also increases proportional to change in the marginal share. The WTP increases from \$162.92 to \$172.67 in 2020 and from \$109.92 to \$111.85 in-context. The MVPF in 2020 declines from 1.34 to 1.14 after including the sales incentive.

E.9.4 Illinois Home Weatherization Assistance Program Bonus Payments

The high and low bonus weatherization MVPF studies the same program as our weatherization MVPF for Illinois (IHWAP) and uses the same treatment effects on energy consumption. The difference is that these MVPFs include the cost and benefits of the contractor bonus payment in addition to the weatherization. Christensen, Francisco & Myers (2023) use data from households who received upgrades from 2018 to 2019 through IHWAP. They use an event study fixed effects model to estimate the impact of weatherization on energy usage. The paper also randomizes some contractors into bonus treatment groups. These contractors would receive incentive payments based on how well they implement the air sealing aspect of the home upgrade measured by CFM50 levels. Contractors received either low bonus payments (\$114) or high bonus payments (\$283).

This changes the original IHWAP weatherization MVPF in three ways. First, we add the cost of the bonus payment in the denominator, increasing the cost of the program. Second, we take estimates from Christensen, Francisco & Myers (2023) who measure the contractors willingness to pay (producer surplus) from the bonus payments which are \$13 for the low payments and \$99 for the high payments. We add these to the numerator. Third, we include the additional environmental benefits from the contractors improved performance. The baseline weatherization program reduces household annual electricity consumption by 1656.44 kWh and annual natural gas consumption by 19.48 MMBtu. The incentive payments lead to a further reduction in electricity of 230.71 kWh for high payments and 118.87 kWh for low payments. For natural gas, they reduce annual consumption by 4.82 MMBtu for high payments and 4.24 MMBtu for low payments.

High Bonus Payment

Our MVPF for high bonus payments for weatherization using estimates from Christensen, Francisco & Myers (2023) is 1.07 [0.90, 1.24] in 2020 and 1.15 in-context. This MVPF is calculated by making the three adjustments outlined above to the MVPF for the Illinois weatherization calculation. After making these adjustments, in 2020, the cost increases from \$10,386.98 to \$10,651.67 and the WTP increases from \$10,181.14 to \$11,391.18.

Low Bonus Payment

Our MVPF for low bonus payments for weatherization using estimates from Christensen, Francisco & Myers (2023) is 1.06 [0.99, 1.14] in 2020 and 1.14 in-context. This MVPF is calculated by making the three adjustments outlined above to the MVPF for the Illinois weatherization calculation. After making these adjustments, in 2020, the cost increases from \$10,386.98 to \$10,462.16 and the WTP increases from \$10,181.14 to \$11,115.70.

E.9.5 Michigan Weatherization Assistance Program (Marketing)

Our MVPF for marketing spending on weatherization using estimates from Fowlie et al. (2018) and Fowlie et al. (2015) is 0.28 [0.10, 0.51] in 2020 and 0.35 in-context. This MVPF studies the same program as our weatherization MVPF for Michigan and uses the same treatment effects on energy consumption. The difference is that this MVPF focuses on the welfare gain from spending on marketing weatherization to households whereas the weatherization MVPF focuses only on spending on weatherization itself. The marketing MVPF takes into account that it is difficult and costly to convince households to take up weatherization subsidies.

Fowlie et al. (2015) estimate that the total marketing spend in their quasi-random encouragement treatment was \$450,000. In the control group who received no marketing, one percent of households took up the weatherization. In the treatment group, six percent took up weatherization. Therefore, of the 435 households that received weatherization subsidies in the treatment, 362.5 households were induced by the marketing. The adjusted cost of the marketing per induced household is \$1,241.37. We assume that each household that is induced to participate in the weatherization program as a result of the marketing is indifferent to receiving the weatherization and therefore does not value the subsidy.

The treatment effect on energy consumption and the cost of the weatherization are identical to that in the Michigan weatherization MVPF. The average household in the paper's sample uses 76.68 MMBtu of natural gas and 7490.90 kWh of electricity annually. The paper's main specification estimates that weatherization reduces natural gas consumption by 18.9% and electricity consumption by 9.5%. This translates into an annual 712.85 kWh and 14.52 MMBtu reduction. Fowlie et al. (2018) presents their results for weatherization lifetimes of 10, 16, and 20 years. Our baseline MVPF uses a 20-year lifetime. The in-context MVPF studies the policy in 2011, the first year of the paper's sample. The average cost of the weatherization upgrades is \$5,150 per home in 2011 dollars.

Cost The total cost is comprised of the direct marketing/weatherization cost and fiscal externalities. The cost of the nudge per induced household is \$1,241.37 and the cost of the energy upgrade per household is \$5,150 in 2011 dollars.

The construction of fiscal externality from the loss in government profit tax revenue from utility companies and the climate fiscal externality is similar to that in the Michigan weatherization MVPF. The only difference is that the households induced by the marketing are all marginal whereas we applied a 50% marginal assumption to the weatherization MVPFs. Therefore, the fiscal externality components for this MVPF are double that of the Michigan weatherization MVPF. The fiscal externality from lost tax revenue increases the cost by \$216.64 in 2020 and \$222.09 in-context. The climate externality reduces government cost by \$59.06 in 2020 and \$52.14 in-context. The total cost is \$7512.92 in 2020 and \$6561.33 in-context.

WTP The willingness to pay is comprised of the environmental externality and the loss in producer profits.

As explained above, we assume that the people who are induced to take-up weatherization from the marketing nudge are indifferent on the margin and therefore do not value the subsidy. However, there are still environmental externalities and producer profits that are affected by the subsidy. Following the logic above, these components will be exactly double the component values from the Michigan weatherization MVPF.

The global environmental externality is \$3,523.76 in 2020 and \$3,151.91 in context. The local environmental externality is \$152.84 in 2020 and \$583.48 in context. The rebound effect offsets approximately 20% of environmental benefits from electricity and 12% of the environmental benefits from natural gas. The resulting rebound effect is -\$529.67 in 2020 and -\$596.30 in context. The total producer willingness to pay is -\$1,045.14 in the baseline specification and -\$859.30 in-context. Summing across these components, the total willingness to pay in 2020 is \$2,101.79 and in-context is \$2,279.79. This results in a baseline MVPF of 0.28 and in-context MVPF of 0.35.

E.9.6 Carbon Footprint Food Label Field Experiment

Our MVPF for placing carbon footprint labels on meal choices is 7.02, using estimates from Lohmann et al. (2022). Lohmann et al. (2022) run a large-scale field experiment at five college cafeterias at the University of Cambridge, implementing a difference-in-differences identification strategy to determine the causal effect of carbon footprint labels on individual food choices. They find that carbon footprint labels caused a 4.3% reduction in average carbon emissions per meal.

WTP Using a representative university cafeteria which serves 1000 meals per day, each with an average carbon footprint of 2 kg, carbon footprint labels avert 2.84 tons of carbon emissions per month. With an SCC of 193, this results in a WTP of 538.03.

Cost The total cost is comprised of the program cost and climate fiscal externality. Lohmann et al. (2022) use pricing estimates by Footsteps Inc. to determine that the cost for implementing the label program on all meals is \$87.2 per month. The climate fiscal externality reduces the cost by \$10.5, resulting in a total cost of \$76.7. Dividing WTP by total cost results in a MVPF of 7.02.

E.10 Gasoline Taxes

Gasoline taxes reduce the quantity of fuel consumed while generating revenue for the government. The MVPF for a gasoline tax combines the price elasticity of gasoline with a measure of the value of the externalities generated per dollar of spending on gasoline. We form MVPFs using 12 estimates of the price elasticity of gasoline and then harmonize the externalities (V/p) and tax rates (τ/p) across MVFPs in our 2020 baseline specification. We discuss differences in the in-context specification at the end of this section.

Consumers' WTP The envelope theorem implies that consumers value the policy change at the value of the price increase. In other words, consumers are willing to pay \$1 for a \$1 change in their cost of gasoline, holding their consumption of gasoline constant. Following Marion & Muehlegger (2011)—who use variation in changes to state-level fuel taxes to show that suppliers fully and immediately pass gasoline and diesel taxes on to consumers—we assume

the \$1 increase in the price of gas is completely passed onto consumers.

Society’s WTP In response to higher fuel prices, drivers (a) reduce the number of miles traveled, and (b) substitute toward more fuel-efficient (higher MPG) vehicles. Each response reduces the total quantity of gasoline consumed. The price elasticity of gasoline (ϵ_{Gas}) is the sum of these behavioral responses. Although purchasing a more fuel-efficient vehicle lowers the cost of driving one mile, we need not account for increases in driving due to improved fuel economy, as estimates of the total change in gas consumption already include this rebound effect.

Burning fewer gallons of gasoline benefits society through less global and local air pollution. In 2020, we estimate that each gallon of gasoline imposed \$2.116 in damages from pollution; \$1.891 of these damages came from global pollutants, while the remaining \$0.226 came from local pollution. Appendix C explains how we estimate these externalities. Using the average retail price of gasoline for all grades and formulations reported by the EIA (\$2.27 in 2020), burning one gallon of gas imposed \$0.93 of damages per dollar of spending on gas in 2020.¹⁸⁶ Multiplying this value by the price elasticity of gasoline gives us society’s WTP for reduced pollution. Using the price elasticity (-0.334) reported by Small & Van Dender (2007), society was willing to pay \$0.311 for a \$1 increase in the gas tax rate, with \$0.2776 for greenhouse gases that contribute to global warming and \$0.0331 for air pollution with adverse health effects.

We scale global benefits by the share of the social cost of carbon that does *not* flow to the US government as increased revenue. We assume that 50% of the social cost of carbon imposes damages on society by affecting GDP, that 15% of global benefits flow to future US residents, and that the government imposes an effective 25.54% tax on economic activity, implying that society captures 98.08% ($1 - 0.15 \times 0.2554 \times 0.5$) of benefits from abating greenhouse gases today and the US government the remaining 1.92% ($\sigma_{US Govt}$). We discuss below how to integrate the remaining 1.92% of global benefits into the MVPF’s denominator as increased long-run revenue. We multiply the \$0.2776 society is willing to pay to avoid global damages (calculated using a price elasticity of -0.334) by 0.9808, resulting in a WTP for avoided greenhouse gases of \$0.2723 in 2020.

Driving fewer miles also benefits society through fewer accidents, less congestion, and reduced pollution ($PM_{2.5}$) from tire and brake wear, which we refer to collectively as “driving externalities.” In 2020, driving externalities imposed \$2.73 of damages per gallon of gas consumed, or \$1.20 per dollar spent on gasoline. All driving externalities impose local damages on society. Appendix C.4.2 explains how we calculate these externalities. Since these externalities arise per mile traveled, we care only about the decline in gasoline consumption owing to reductions in miles traveled. Following Small & Van Dender (2007), we assume 52% of the

¹⁸⁶To calculate the annual average price of gasoline, we average monthly price data from the EIA’s “U.S. All Grades All Formulations Retail Gasoline Prices” series (EIA 2023g) and weight by monthly data on the quantity of gasoline supplied from the EIA’s “U.S. Product Supplied of Finished Motor Gasoline” series (EIA 2023h), which approximates the total quantity of reformulated and convention gasoline consumed in a given month. We construct annual averages rather than using the reported annual average price to account for changes in the federal gas tax rate that went into effect in specific months. Our annual averages are nearly identical to those reported by the EIA. For years not included in the EIA’s price series (earlier than 1994), we impose on each month the average annual historic gas price reported in the DOE’s “Historical Gasoline Prices, 1929-2011” series (DOE 2016), although only one in-context MVPF from our extended sample (Gas (Hughes - Ext)) requires price data from before 1994. So that each series has the same average gas price in 1994, we calculate the difference between each series’ estimate of the 1994 average price and add this difference to each estimate in the earlier series. After this transformation, each series had the same average fuel economy in 1994.

change in gasoline consumption arises from reduced driving. Multiplying the price elasticity of gasoline by this parameter isolates the reduction in gasoline consumption that follows from reduced driving. We refer to the share of the change in gasoline consumption from changes in VMT as β . Using our 2020 driving externality estimate of \$1.20 in damages per gallon of gasoline consumed, a price elasticity of -0.334, and a β of 0.52, society was willing to pay \$0.209 ($-0.334 \times 0.52 \times \1.20) for avoided damages from driving.

Changes in fleet composition typically arise from consumers substituting for more fuel-efficient gasoline-powered vehicles, but purchasing an EV allows drivers to consume fewer gallons of gas while traveling the same number of miles. We account for benefits and costs from charging more EVs. Because the own price elasticity of gasoline measures the total change in gas consumption, we assume the price elasticities used to construct MVPFs account for reductions in gas consumption due to consumers switching to electric vehicles. As a result, any WTPs arising from gasoline usage (gasoline producer profits, environmental benefits of reduced gas consumption, and gas tax revenue) need not be adjusted. The cross-price elasticity between gasoline and EVs governs the amount of substitution toward EVs due to higher gas prices. Formally, let $\eta_{EV, Gas}$ represent the cross-price elasticity between gasoline and EVs:

$$\frac{dQ_{EV}}{dP_{Gas}} \times \frac{P_{Gas}}{Q_{EV}} = \eta_{EV, Gas} \quad (69)$$

We assume consumers choose between purchasing either an EV or a gas-powered vehicle. Under this discrete choice framework, Slutsky symmetry implies that the relationship between a change in the price of an EV and the consumption of gas-powered vehicles is identical to the relationship between a change in the price of a gas-powered vehicle and the consumption of EVs. Moreover, the magnitude of the shifting of consumption from gas-powered cars to EVs is equivalent to the opposite direction of the change in own-good consumption: the increase in EV consumption is equal to the decrease in gas-powered vehicle consumption. We can express this relationship as

$$\frac{dQ_{EV}}{dP_{Gas Car}} = \frac{dQ_{Gas Car}}{dP_{EV}} \quad (70)$$

$$\frac{dQ_{Gas Car}}{dP_{EV}} = \frac{-dQ_{EV}}{dP_{EV}} \quad (71)$$

where the price of owning a gas-powered vehicle ($P_{Gas Car}$) is the present discounted value of gas consumed over the vehicle's lifetime.¹⁸⁷ Changes in the price of gasoline enter linearly into the price of owning a gas-powered vehicle such that

$$\frac{dQ_{EV}}{dP_{Gas Car}} = \frac{dQ_{EV}}{dP_{Gas}} \quad (72)$$

Combining equations, Slutsky symmetry implies that the relationship between a change in the price of gasoline and EV consumption is negatively proportional to the relationship between a change in the price of an EV and EV consumption:

¹⁸⁷We calculate the present discounted value of gasoline consumption for a given year by holding the average annual price of gas for that year fixed over the vehicle's lifetime, discounting with our selected discount rate. We assume EVs and gas-powered cars travel the same number of miles over their 17 year lifetimes.

$$\frac{dQ_{EV}}{dP_{Gas}} = \frac{dQ_{EV}}{dP_{Gas Car}} = \frac{-dQ_{EV}}{dP_{EV}} \quad (73)$$

Eq. 69 can now be expressed in terms of the own-price elasticity of EVs and the price ratio between the present discounted value of gasoline expenses and the price of an EV.

$$\frac{-dQ_{EV}}{dP_{EV}} \times \frac{P_{Gas}}{Q_{EV}} = \eta_{EV, Gas} \quad (74)$$

$$\left(\frac{-dQ_{EV}}{dP_{EV}} \times \frac{P_{Gas}}{Q_{EV}} \right) \times \frac{P_{EV}}{P_{EV}} = \eta_{EV, Gas} \quad (75)$$

$$\underbrace{\left(\frac{-dQ_{EV}}{dP_{EV}} \times \frac{P_{EV}}{Q_{EV}} \right)}_{\text{Price Elasticity of EVs } (\epsilon_{EV})} \times \frac{P_{Gas}}{P_{EV}} = \eta_{EV, Gas} \quad (76)$$

To calculate the cross-price elasticity implied by Eq. 76, we use the own price elasticity (-2.1) estimated by Muehlegger & Rapson (2022).¹⁸⁸ As described in Section 4, we assume an EV displaces a cleaner-than-average gas-powered car. Fueling a vehicle with this counterfactual fuel economy (41.2 MPG) in 2020 would cost \$5,643.48 over its lifetime, using a 2% discount rate, the average annual price of gas in 2020 (\$2.27), and an average annual VMT that is 61% of the VMT of an average car. An EV purchased in 2020 sold for, on average, \$53,378.23, net of the average total subsidy in 2020 (\$647.25). Together, these parameters imply a cross-price elasticity of 0.22 in 2020.

Increased EV consumption generates environmental damages from increased electricity usage and dirtier manufacturing processes. The lifetime damages from EVs are expressed in dollars per EV purchased. To convert from dollars per EV purchased to dollars of spending on gasoline, we divide the lifetime damages from an EV by the price of gasoline and multiply by the ratio of EV consumption to gas consumed by light-duty vehicles. We then multiply this term by the behavioral response (the cross-price elasticity) or

$$\eta_{EV, Gas} \frac{V_{EV} Q_{EV}}{P_{Gas} Q_{Gas}} \quad (77)$$

The average EV purchased in 2020 imposed \$3,398.31 in global damages ($V_{EV, Global}$) and \$366.02 in local damages ($V_{EV, Local}$) over its lifetime.¹⁸⁹ In 2020, US consumers purchased 238,540 EVs (DOE 2024a, Cox Automotive 2023), generating a total of \$810.63 million and \$87.31 million in global and local damages, respectively.¹⁹⁰ Dividing monetized damages from EV consumption by the product of the price of gasoline (\$2.27) and the number of gallons of gasoline consumed

¹⁸⁸We need not account for the pass-through rate we apply when calculating the MVPF of an EV subsidy using the behavioral response estimated by Muehlegger & Rapson (2022), as here we examine a change in the gas tax rate, not a change in the EV rebate level.

¹⁸⁹As shown below, we scale $V_{EV, Global}$ by $\sigma_{US Govt}$ to isolate society's WTP for global benefits from the US government's added revenues from abating carbon today. Global damages come from EV charging and battery production. Local damages come from EV charging alone. We account for the rebound in electricity usage due to higher prices in both global and local damages.

¹⁹⁰For 2022, we use data from Cox Automotive (2023) as DOE (2024a) did not report 2022 data at the time we accessed it. Only one in-context MVPF uses these 2022 data ("Gas (Kilian)").

by light-duty vehicles in 2020 (1.127 billion) expresses the effects of induced EV substitution in levels of gasoline spending: in 2020, EVs imposed \$0.003 in global damages and \$0.0003 in local damages per dollar spent on gasoline.¹⁹¹ Multiplying by the cross-price elasticity (0.22) then provides society’s WTP to avoid the environmental damages associated with charging more EVs. We add these terms to society’s WTP for the local and global benefits of reduced gasoline consumption. Accounting for damages from EVs decreases society’s WTP for global damages by \$0.0007 and local damages by \$0.00008.

Collecting the environmental benefits and damages from reduced gas usage and driving and increased EV manufacturing and charging, society’s WTP for reduced damages from gasoline consumption can be expressed as

$$\begin{aligned}
 WTP_{Society} = & \underbrace{\left(\epsilon_{Gas} \frac{(1 - \sigma_{US Govt}) V_{Gas, Global}}{P_{Gas}} + \eta_{EV, Gas} \frac{(1 - \sigma_{US Govt}) V_{EV, Global} Q_{EV}}{P_{Gas} Q_{Gas}} \right)}_{\text{Global Env (-\$0.272)}} \\
 & + \underbrace{\left(\epsilon_{Gas} \frac{V_{Gas, Local}}{P_{Gas}} + \eta_{EV, Gas} \frac{V_{EV, Local} Q_{EV}}{P_{Gas} Q_{Gas}} \right)}_{\text{Local Env (-\$0.033)}} + \underbrace{\beta \epsilon_{Gas} \frac{V_{Driving}}{P_{Gas}}}_{\text{Driving (-\$0.209)}}
 \end{aligned} \tag{78}$$

Using an ϵ_{Gas} of -0.334, our 2020 values, and our preferred parameters outlined above, we estimate a WTP for global pollution of -\$0.2716 (with -\$0.2723 for reduced gas consumption and \$0.0007 for increased EV usage), a WTP for local pollution of -\$0.0331 (with -\$0.03314 for reduced gas consumption and \$0.00008 for increased EV usage), and a WTP for driving damages of \$0.209. Each term’s label corresponds to the components displayed in Figure 7. Summing by damage type, society was willing to pay \$0.2716 for global benefits (Global Env) and \$0.2419 for local benefits (Local Env + Driving) when using the price elasticity estimated by Small & Van Dender (2007). As described below, we sign each component depending on society’s WTP to remove the tax: society has a negative WTP to remove the tax. Within this component, society has a negative WTP for the benefits associated with reduced gas usage and driving but a positive WTP for the damages induced by greater EV adoption.

As noted, we do not calculate a rebound effect for gas tax MVPFs, since estimates of the total change in gas consumption should account for increases in VMT in response to substitution toward more fuel-efficient vehicles. We do not isolate this rebound in VMT. We do, however, account for the rebound in electricity prices due to increased EV charging: a greater number of EVs drawing from the grid increases electricity demand, resulting in higher electricity prices and, in turn, less electricity consumption. See Section 4 for more on this calculation. This rebound in electricity prices is accounted for and included in the environmental damages from increased EV charging today and the dynamic environmental benefits of increased EV consumption tomorrow. Additionally, we note that, when we present results without learning by doing, we also remove the static effects from EV substitution.

¹⁹¹We calculate total gallons of gasoline consumed by light-duty vehicles by aggregating monthly supply data from the EIA’s “U.S. Product Supplied of Finished Motor Gasoline” series (described above) (EIA 2023*h*). To those annual values, we add the total annual quantity of aviation gasoline supplied (EIA 2024*i*) to replicate the EIA’s approach to measuring the total quantity of gasoline consumed in a year. Lastly, we again follow the EIA by assuming that light-duty vehicles consume 91% of all gasoline sold in a year (EIA 2023*d*).

Learning-by-Doing Benefits We augment our learning-by-doing framework to allow a change in the gas tax rate (rather than a change in subsidy amount) to induce greater EV adoption.¹⁹² Specifically, let $V_{Dynamic}$ be some benefits from future EV consumption induced by a \$1 change in the subsidy for an EV. $V_{Dynamic}$ is calculated using an own price elasticity of ϵ_{EV} and is measured per dollar of spending on an EV.

A change in the subsidy for an EV generates $V_{Dynamic}$ through consumers behavioral response to the price of an EV (ϵ_{EV}). However, we do not care about a change in the subsidy amount; rather, we care about a change in the price of gasoline. We, therefore, multiply $V_{Dynamic}$ by the price of an EV and divide by the behavioral response used to calculate $V_{Dynamic}$ to return to dollars of benefits per EV. From there, we can apply the same conversion used in equation (77) to move from dollars of spending on EVs to dollars of spending on gasoline and multiply by the cross-price elasticity to calculate society’s WTP for the learning-by-doing benefits generated by increased EV substitution. Specifically,

$$\eta_{EV, Gas} \left(V_{Dynamic} \frac{P_{EV}}{\epsilon_{EV}} \frac{Q_{EV}}{Q_{Gas}} \right) \quad (79)$$

Put differently, equation (79) scales $V_{Dynamic}$ by the ratio of the cross-price elasticity between EVs and gasoline to the own price elasticity of EV consumption used to calculate these benefits and converts from dollars of spending on EVs to dollars of spending on gasoline. $V_{Dynamic}$ includes both the environmental and price benefits from learning-by-doing. In 2020, society was willing to pay \$0.002 for future EV price reductions and less than \$0.001 for the environmental benefits from learning-by-doing.¹⁹³

$$\begin{aligned} WTP_{LBD} = & \underbrace{\eta_{EV, Gas} \left(((1 - \sigma_{US Govt}) V_{Dynamic, Env. Global} + V_{Dynamic, Env. Local}) \frac{P_{EV}}{\epsilon_{EV}} \frac{Q_{EV}}{P_{Gas} Q_{Gas}} \right)}_{\text{Dynamic Env (-\$0.00024)}} \\ & + \underbrace{\eta_{EV, Gas} \left(V_{Dynamic, Price} \frac{P_{EV}}{\epsilon_{EV}} \frac{Q_{EV}}{P_{Gas} Q_{Gas}} \right)}_{\text{Dynamic Price (-\$0.0019)}} \end{aligned} \quad (80)$$

Like the components associated with substitution toward EV, these learning-by-doing benefits are common across all gas tax MVPFs calculated in 2020, as our derived cross-price elasticity is independent of the own-price elasticity of gasoline unique to each MVPF calculation.

¹⁹²We do not consider substitution toward EVs for in-context MVPFs calculated for years before 2011. When calculating in-context learning-by-doing and EV substitution benefits for 2022 (for the in-context MVPF of “Gas (Kilian)”), we use externality estimates calculated for 2022. For the fiscal externality from EV subsidies, we inflation adjust the 2020 state-average subsidy value to 2022 dollars (\$683.12) and use a federal average subsidy of \$1,064.95.

¹⁹³For the price benefits, we take the \$0.368 of learning-by-doing benefits calculated using price elasticity of -2.1, multiply by the ratio of behavioral responses (0.22/-2.1), and multiply again by the ratio of spending on EVs to spending on gasoline (0.05). This yields a -\$0.0019 learning-by-doing price benefit per dollar of spending on gasoline. We do the same for the learning-by-doing environmental benefits (\$0.043 in global benefits and \$0.004 in local benefits), which yields a learning-by-doing environmental benefit of \$0.00024 per dollar of spending on gasoline.

Producers' WTP As described in Appendix C.4.5, we calculate a total per-gallon markup equal to 35% of the price of gasoline. We subtract from this gasoline markup the average, economy-wide markup (8%) estimated by De Loecker et al. (2020), resulting in a 27% average markup on a gallon of gas. In 2020, the total markup on gasoline was \$0.61 per gallon, which we adjust by the corporate average tax rate (21%) to account for the share of profits producers keep (Watson 2022).¹⁹⁴ This results in a post-tax externality borne by producers of \$0.21 per dollar of spending on gasoline. With a price elasticity of -0.334, producers were willing to pay \$0.071 in 2020 for the policy change.

To account for utilities' WTP, we perform the same calculations described above to move from WTP per EV to WTP per dollar of spending on gasoline. In 2020, the average EV generated \$265.56 in post-tax profits for utilities over its lifetime. From the 238,540 EVs purchased in 2020, utilities would earn a total of \$63.35 million over these vehicles' lifetimes.¹⁹⁵ Dividing by total spending on gasoline in 2020 (1.127 billion times \$2.27 per gallon) denotes utility profits in dollars of spending on gasoline, and multiplying by the cross-price elasticity yields utilities' WTP for a \$1 increase in the gas tax rate. In 2020, utilities were WTP -\$0.00005 for the policy change. We sign this component as a negative since utilities are WTP to keep the policy change.

Producers' total WTP can be expressed as

$$WTP_{Producers} = \underbrace{\epsilon_{Gas} \frac{(1 - \tau_{Corp, Gas}) \mu_{Gas}}{P_{Gas}}}_{\text{Gasoline Producers } (\$0.071)} + \underbrace{\eta_{EV, Gas} \frac{(1 - \tau_{Corp, Utilities}) \mu_{EV} Q_{EV}}{P_{Gas} Q_{Gas}}}_{\text{Utilities } (-\$0.00005)} \quad (81)$$

where μ is the pre-tax profit producers earn per unit of good sold, and τ_{Corp} , is the effective corporate tax rate producers face. Each term's label corresponds to the components included in Figure 7.

Total WTP Summing across components, a \$1 change in the gas tax rate results in a total WTP of \$0.555 when using a price elasticity of gasoline of -0.334. Consumers (\$1) and producers (\$0.071) are both willing to pay to avoid the tax increase, while society (-\$0.514) and future consumers (-\$0.002) are willing to pay to keep the tax increase. We sign each component depending on the group's willingness to pay to remove the tax. Consumers and producers are both willing to pay to remove the tax since these groups are made worse off through higher prices and reduced profits, respectively. On the other hand, society is willing to pay to keep the tax on the books, as they are made better off through reduced environmental and driving externalities. Future consumers also have a negative willingness to pay to remove the tax, although these future consumers' WTP does not offset contemporary consumers' WTP to avoid higher gas prices. Considering the removal of a tax (rather than the addition of one) also allows us to normalize both taxes and subsidies to a positive \$1 mechanical cost.

Cost A \$1 increase in the gas tax rate mechanically raises \$1 of revenue for the government. However, the accompanying decrease in gas consumption reduces the amount of revenue the government collects by the size of the behavioral response times the tax collected per dollar of

¹⁹⁴We do not vary across time the effective corporate tax rate gasoline producers face.

¹⁹⁵We hold the price of electricity in 2020 constant over the vehicle's lifetime and discount using our preferred discount rate of 2%.

gasoline spending. In 2020, the federal gas tax rate was \$0.184 per gallon (FHWA 2021) while the average state tax on gasoline (weighted by gross gallons of gasoline taxed) was \$0.281 per gallon (FHWA 2022, 2020). Accounting for federal and state gas taxes, the government collected \$0.20 per dollar spent on gas. Multiplying by a price elasticity of -0.334, the government faced a \$0.068 loss in revenue from decreased gasoline consumption.

We also account for four other fiscal externalities that impact the revenue raised from a \$1 change in the gas tax rate. Decreases in producer profits reduce government revenue in the form of lost corporate taxes (assuming a 21% tax rate). The pre-tax markup on gasoline was \$0.27 per dollar of gas spending in 2020, meaning the government collected \$0.06 in corporate tax revenue for each dollar spent on gas.¹⁹⁶ With the price elasticity (-0.334) estimated by Small & Van Dender (2007), we calculate a \$0.019 fiscal externality from lost corporate tax revenue.

Second, public and private utilities generate revenue for the government, meaning substitution toward EVs should increase revenue collected through increased vehicle charging. Charging the average EV purchased in 2020 generated \$144.25 in profits for private utilities over the vehicle's lifetime (\$34.41 million across the lifetimes of all EVs sold in 2020), or \$0.0001 per dollar of spending on gasoline. Applying the cross-price elasticity yields a fiscal externality of \$0.00003 from increased revenue collected from utilities. Through this component, the government raises revenue by inducing substitution toward EVs.

Third, EVs qualified for \$647.25 in federal and state subsidies in 2020 (\$154.4 million across all EVs sold in 2020). Applying the same transformation described above, the federal government lost \$0.0006 in revenue per dollar spent on gasoline, implying a fiscal externality of \$0.0001 from increased spending on EV subsidies after applying a cross-price elasticity of 0.22. Through this component, the government loses revenue by having to subsidize more EV purchases.

Lastly, abating greenhouse gas emissions through a gas tax raises revenue for the government in the long run. When calculating society's WTP for global pollution benefits, we scale the WTP component by the share of global benefits that do not flow to the US government as revenue (98.1%). We put the remaining 1.92% in the denominator. With a price elasticity of -0.334, the total, unadjusted WTP in 2020 for *all* global benefits was \$0.2776, implying the government generated \$0.005 ($\0.2776×0.0192) in revenue by abating carbon emissions today and promoting economic output tomorrow.¹⁹⁷ Abating carbon emissions generates a positive fiscal externality, as the US government generates additional revenue through the policy change.

Summing the mechanical \$1 of revenue raised and the five fiscal externalities, we obtain a total "cost" of \$0.918 when using a price elasticity of -0.334: a \$1 increase in the gas tax rate raises \$0.918 in revenue for the government.¹⁹⁸ Collecting the mechanical revenue raised and the fiscal externalities, we can express the denominator of the MVPF as

¹⁹⁶We assume all gasoline producer profits are subject to the American tax schedule but note that the geographic variation in crude oil production could render our calculation of this fiscal externality an overestimate.

¹⁹⁷This total WTP for global damages includes global benefits and damages from EV substitution and learning-by-doing.

¹⁹⁸To match our approach with subsidies, we again consider the effects of removing the gasoline tax: this would mechanically lower revenue by \$1 but would positively impact government revenue by increasing gas consumption. This allows us to treat both taxes and subsidies as having a \$1 mechanical cost.

$$\begin{aligned}
\text{Cost} = & \underbrace{\epsilon_{Gas} \frac{\tau_{Gas, Federal} + \tau_{Gas, State}}{P_{Gas}} + \epsilon_{Gas} \frac{\tau_{Corp, Gas} \mu_{Gas}}{P_{Gas}} + \eta_{EV, Gas} \frac{\tau_{Corp, Utilities} \mu_{EV} Q_{EV}}{P_{Gas} Q_{Gas}}}_{\text{Taxes } (-\$0.087)} \\
& + \underbrace{\eta_{EV, Gas} \frac{\tau_{EV} Q_{EV}}{P_{Gas} Q_{Gas}}}_{\text{Subsidies } (-\$0.0001)} \\
& + \underbrace{\left(\epsilon_{Gas} \frac{\sigma_{US Govt} V_{Gas, Global}}{P_{Gas}} + \eta_{EV, Gas} \frac{\sigma_{US Govt} V_{EV, Global} Q_{EV}}{P_{Gas} Q_{Gas}} + \eta_{EV, Gas} \frac{\sigma_{US Govt} V_{Dynamic, Env. Global} P_{EV} Q_{EV}}{\epsilon_{EV} P_{Gas} Q_{Gas}} \right)}_{\text{Climate FE } (\$0.005)}
\end{aligned} \tag{82}$$

Each component is labeled using the corresponding label from Figure 7. “Taxes” is the sum of fiscal externalities from changes in gas taxes collected and profit taxes paid by gas producers and utilities.

MVPF Dividing the total WTP calculated above (\$0.555) by the total cost (\$0.918), both calculated with a price elasticity of -0.334, we form an MVPF of 0.605 in 2020.

Estimates in Sample The following paragraphs explain how each paper in our sample estimates the price elasticity of gasoline and how we form MVPFs using each paper’s estimate. All papers in our sample estimate an elasticity (rather than a semi-elasticity). For all estimates, we evaluate the policy change at the national level. Our baseline estimate focuses on 2020, and our in-context estimates are set in the last year within each paper’s sample.

To further broaden our treatment of gas taxes, we include in our extended sample additional elasticity estimates from papers already in our main sample as well as elasticities from papers that use more structural methods. All papers in our extended sample find elasticities that fall within the range of elasticities in our main sample. We do not report confidence intervals for extended sample MVPFs. Table 1 notes which policies belong to our extended sample, but we also add an “*” to extended sample entries below for reference.

Given the relatively low price of gasoline in 2020 (\$2.27 per gallon), we note for reference that our category average gasoline tax MVPF would increase to 0.765 were one to use the average price of gasoline in 2021 (\$3.11 per gallon) in our calculations and to 0.823 using 2022 price (\$4.08 per gallon).

State-level Gas Tax Variation (Davis & Kilian 2011)

Davis & Kilian (2011) use variation in state-level fuel prices between 1989 and 2008 to estimate a long-run own price elasticity of gasoline of -0.46 (s.e. 0.23). The authors’ Table 4 (Column 4, Row 1) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2008.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases (-\$0.374), reduced local air pollution (-\$0.046), and reduced driving externalities (-\$0.288), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of -0.46. Society’s WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of -\$0.002 and added

environmental benefits smaller than $-\$0.001$. Gasoline producers are willing to pay $\$0.098$ for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of $\$0.388$ for the policy change in 2020.

A $\$1$ increase in the gas tax rate raised the government $\$0.887$ in 2020. In addition to the $\$1$ of revenue mechanically raised, the government lost $\$0.12$ in tax revenue ($\$0.094$ from lost gas tax revenue, $\$0.026$ from lost corporate tax revenue collected from gas producers, and less than $\$0.001$ gained from utility profits). The government also spent less than $\$0.001$ in EV subsidies and raised $\$0.007$ by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a $\$1$ increase in the gas tax rate generated $\$0.887$ in government revenue in 2020.

With a price elasticity of -0.46 , dividing the total WTP of $\$0.388$ by the total cost of $\$0.887$ yields an MVPF of 0.437 $[-0.208, 0.997]$ in 2020.

In 2008, the nominal gas price was $\$3.310$. Society was WTP $-\$0.166$ for reduced greenhouse gases, $-\$0.051$ for reduced local air pollution, and $-\$0.143$ for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP $\$0.099$ for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of $\$0.739$. A $\$1$ increase in the gas tax costs the government $\$0.054$ in lost gas tax revenue and $\$0.026$ in lost corporate tax revenue. The policy change also earned the government $\$0.003$ in revenue from abated greenhouse gases, for a total cost of $\$0.923$. Dividing total WTP by total cost results in an MVPF of 0.801 in the context (2008) in which the authors estimated the price elasticity of gasoline.

Urban Area-level Gas Price Variation (Su 2011)

Su (2011) uses variation at the urban area level in 2001 to estimate an own price elasticity of gasoline of -0.397 (t-statistic -2.52). The authors' Table 3 (Column 1, Row 16) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2001.

In 2020, consumers were WTP $\$1$ for a $\$1$ increase in the gas tax rate. Society was WTP for reduced greenhouse gases ($-\$0.323$), reduced local air pollution ($-\$0.039$), and reduced driving externalities ($-\$0.248$), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 ($\$2.27$) and multiplying by the price elasticity of -0.397 . Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of $-\$0.002$ and added environmental benefits smaller than $-\$0.001$. Gasoline producers are willing to pay $\$0.084$ for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of $\$0.472$ for the policy change in 2020.

A $\$1$ increase in the gas tax rate raised the government $\$0.903$ in 2020. In addition to the $\$1$ of revenue mechanically raised, the government lost $\$0.104$ in tax revenue ($-\$0.081$ from lost gas tax revenue, $\$0.022$ from lost corporate tax revenue collected from gas producers, and less than $\$0.001$ gained from utility profits). The government also spent less than $\$0.0001$ in EV subsidies and raised $\$0.006$ by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a $\$1$ increase in the gas tax rate generated $\$0.903$ in government revenue in 2020.

With a price elasticity of -0.397 , dividing the total WTP of $\$0.472$ by the total cost of $\$0.903$ yields an MVPF of 0.523 $[0.113, 0.907]$ in 2020.

In 2001, the nominal gas price was \$1.466. Society was WTP -\$0.222 for reduced greenhouse gases, -\$0.150 for reduced local air pollution, and -\$0.230 for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP \$0.122 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.519. A \$1 increase in the gas tax costs the government \$0.102 in lost gas tax revenue and \$0.032 in lost corporate tax revenue. The policy change also earned the government \$0.004 in revenue from abated greenhouse gases, for a total cost of \$0.870. Dividing total WTP by total cost results in an MVPF of 0.596 in the context (2001) in which the authors estimated the price elasticity of gasoline.

State-level Gas Tax Variation (Coglianese et al. 2017)

Coglianese et al. (2017) use state-level variation in gas taxes over 1989-2008 and include leads and lags to estimate a long-run own price elasticity of gasoline of -0.368 (s.e. 0.239). The authors' Table 2 (Column 5, row labeled "Cumulative Effect") reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters, and another in the context of 2008.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases (-\$0.299), reduced local air pollution (-\$0.036), and reduced driving externalities (-\$0.230), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of -0.368. Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of -\$0.002 and added environmental benefits smaller than -\$0.001. Gasoline producers are willing to pay \$0.078 for lost profits, and utilities smaller than -\$0.001 for increased profits. Summing these components yields a total WTP of \$0.510 for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government \$0.910 in 2020. In addition to the \$1 of revenue mechanically raised, the government lost \$0.096 in tax revenue (\$0.075 from lost gas tax revenue, \$0.021 from lost corporate tax revenue collected from gas producers, and less than \$0.001 gained from utility profits). The government also spent an additional \$0.0001 in EV subsidies and raised \$0.006 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.910 in government revenue in 2020.

With a price elasticity of -0.368, dividing the total WTP of \$0.510 by the total cost of \$0.910 yields an MVPF of 0.561 [-0.079, 1.113] in 2020.

In 2008, the nominal gas price was \$3.310. Society was WTP -\$0.133 for reduced greenhouse gases, -\$0.041 for reduced local air pollution, and -\$0.114 for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP \$0.079 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.792. A \$1 increase in the gas tax costs the government -\$0.043 in lost gas tax revenue and \$0.021 in lost corporate tax revenue. The policy change also earned the government \$0.003 in revenue from abated greenhouse gases, for a total cost of \$0.938. Dividing total WTP by total cost results in an MVPF of 0.844 in the context (2008) in which the authors estimated the price elasticity of gasoline.

Regional Gas Price Variation (Manzan & Zerom 2010)

Manzan & Zerom (2010) uses household-level survey data from 1991 and 1994 to estimate an own price elasticity of gasoline of -0.355 (s.e. $.117$). The authors' Table 4 (Column 4, Row 1) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 1994.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases ($-\$0.289$), reduced local air pollution ($-\$0.035$), and reduced driving externalities ($-\$0.222$), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 ($\$2.27$) and multiplying by the price elasticity of -0.355 . Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of $-\$0.002$ and added environmental benefits smaller than $-\$0.001$. Gasoline producers are willing to pay $\$0.076$ for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of $\$0.527$ for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government $\$0.913$ in 2020. In addition to the \$1 of revenue mechanically raised, the government lost $\$0.093$ in tax revenue ($-\$0.073$ from lost gas tax revenue, $\$0.020$ from lost corporate tax revenue collected from gas producers, and less than $\$0.001$ gained from utility profits). The government also spent less than $\$0.0001$ in EV subsidies and raised $\$0.006$ by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated $\$0.913$ in government revenue in 2020.

With a price elasticity of -0.335 , dividing the total WTP of $\$0.527$ by the total cost of $\$0.913$ yields an MVPF of 0.578 [$0.287, 0.863$] in 2020.

In 1994, the nominal gas price was $\$1.080$. Society was WTP $-\$0.179$ for reduced greenhouse gases, $-\$0.242$ for reduced local air pollution, and $-\$0.231$ for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP $\$0.118$ for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of $\$0.466$. A \$1 increase in the gas tax costs the government $\$0.121$ in lost gas tax revenue and $\$0.031$ in lost corporate tax revenue. The policy change also earned the government $\$0.004$ in revenue from abated greenhouse gases, for a total cost of $\$0.851$. Dividing total WTP by total cost results in an MVPF of 0.548 in the context (1994) in which the authors estimated the price elasticity of gasoline.

State-level Gas Price Variation (Small & Van Dender 2007)

Small & Van Dender (2007) use variation in state-level fuel prices between 1997 and 2001 to estimate a long-run own price elasticity of gasoline of -0.334 (s.e. 0.045). The authors' Table 5 (Column 2, Row 7) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2001.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases ($-\$0.272$), reduced local air pollution ($-\$0.033$), and reduced driving externalities ($-\$0.209$), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 ($\$2.27$) and multiplying by the price elasticity of -0.3340 . Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of $-\$0.002$ and added environmental benefits smaller than $-\$0.001$. Gasoline producers are willing to pay $\$0.071$ for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of $\$0.555$ for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government \$0.918 in 2020. In addition to the \$1 of revenue mechanically raised, the government lost \$0.087 in tax revenue (\$0.068 from lost gas tax revenue, \$0.019 from lost corporate tax revenue collected from gas producers, and less than \$0.001 gained from utility profits). The government also spent an additional \$0.0001 in EV subsidies and raised \$0.005 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.918 in government revenue in 2020.

With a price elasticity of -0.334, dividing the total WTP of \$0.555 by the total cost of \$0.918 yields an MVPF of 0.605 [0.498, 0.717] in 2020.

In 2001, the nominal gas price was \$1.47. Society was WTP -\$0.187 for reduced greenhouse gases, -\$0.126 for reduced local air pollution, and -\$0.194 for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP \$0.102 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.595. A \$1 increase in the gas tax costs the government \$0.085 in lost gas tax revenue and \$0.027 in lost corporate tax revenue. The policy change also earned the government \$0.004 in revenue from abated greenhouse gases, for a total cost of \$0.891. Dividing total WTP by total cost results in an MVPF of 0.668 in the context (2001) in which the authors estimated the price elasticity of gasoline.

National Crude Price Variation (Li et al. 2014)

Li et al. (2014) use variation in the price of imported crude oil from 1968 to 2008 to estimate an own price-elasticity of gasoline of -0.323 (s.e. 0.083). The authors' Table 4 (Panel A, Column 6, Row 2) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2008.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases (-\$0.263), reduced local air pollution (-\$0.032), and reduced driving externalities (-\$0.202), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of -0.323. Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of -\$0.002 and added environmental benefits smaller than -\$0.001. Gasoline producers are willing to pay \$0.069 for lost profits, and utilities less than -\$0.001 for increased profits. Summing these components yields a total WTP of \$0.570 for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government \$0.921 in 2020. In addition to the \$1 of revenue mechanically raised, the government lost \$0.084 in tax revenue (\$0.066 from lost gas tax revenue, \$0.018 from lost corporate tax revenue collected from gas producers, and less than \$0.001 gained from utility profits). The government also spent an additional \$0.0001 in EV subsidies and raised \$0.005 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.921 in government revenue in 2020.

With a price elasticity of -0.323, dividing the total WTP of \$0.570 by the total cost of \$0.921 yields an MVPF of 0.619 [0.420, 0.821] in 2020.

In 2008, the nominal gas price was \$3.310. Society was WTP -\$0.116 for reduced greenhouse gases, -\$0.036 for reduced local air pollution, and -\$0.100 for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP \$0.069 for the

policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.817. A \$1 increase in the gas tax costs the government \$0.038 in lost gas tax revenue and \$0.018 in lost corporate tax revenue. The policy change also earned the government \$0.002 in revenue from abated greenhouse gases, for a total cost of \$0.946. Dividing total WTP by total cost results in an MVPF of 0.864 in the context (2008) in which the authors estimated the price elasticity of gasoline.

City-level Gas Price Variation (Levin et al. 2017)

Levin et al. (2017) use variation in city-level fuel prices between 2006 and 2009 to estimate an own price elasticity of gasoline of -0.295 (s.e. .031). The authors' Table 1 (Column 1, Row 1) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2009.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases (-\$0.240), reduced local air pollution (-\$0.029), and reduced driving externalities (-\$0.184), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of -.295. Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of -\$0.002 and added environmental benefits smaller than -\$0.001. Gasoline producers are willing to pay \$0.063 for lost profits, and utilities less than -\$0.001 for increased profits. Summing these components yields a total WTP of \$0.607 for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government \$0.928 in 2020. In addition to the \$1 of revenue mechanically raised, the government lost \$0.077 in tax revenue (\$0.060 from lost gas tax revenue, \$0.017 from lost corporate tax revenue collected from gas producers, and less than \$0.001 gained from utility profits). The government also spent an additional \$0.0001 in EV subsidies and raised \$0.005 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.928 in government revenue in 2020.

With a price elasticity of -.295, dividing the total WTP of \$0.607 by the total cost of \$0.928 yields an MVPF of 0.654 [0.583, 0.731] in 2020.

In 2009, the nominal gas price was \$2.403. Society was WTP -\$0.149 for reduced greenhouse gases, -\$0.042 for reduced local air pollution, and -\$0.127 for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP \$0.064 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.746. A \$1 increase in the gas tax costs the government \$0.048 in lost gas tax revenue and \$0.017 in lost corporate tax revenue. The policy change also earned the government \$0.003 in revenue from abated greenhouse gases, for a total cost of \$0.938. Dividing total WTP by total cost results in an MVPF of 0.796 in the context (2009) in which the authors estimated the price elasticity of gasoline.

National Gas Tax Variation (Sentenac-Chemin 2012)

Sentenac-Chemin (2012) use variation in national fuel prices between 1978 and 2005 to estimate a long-run own price elasticity of gasoline of -0.28 (t-statistic -5.28). The authors report this elasticity in their Table 1. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2005.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for

reduced greenhouse gases (-\$0.228), reduced local air pollution (-\$0.028), and reduced driving externalities (-\$0.175), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of -0.28. Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of -\$0.002 and added environmental benefits smaller than -\$0.001. Gasoline producers are willing to pay \$0.60 for lost profits, and utilities less than -\$0.001 for increased profits. Summing these components yields a total WTP of \$0.627 for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government \$0.931 in 2020. In addition to the \$1 of revenue mechanically raised, the government lost \$0.073 in tax revenue (\$0.057 from lost gas tax revenue, \$0.016 from lost corporate tax revenue collected from gas producers, and less than \$0.001 gained from utility profits). The government also spent less than \$0.0001 in EV subsidies and raised \$0.004 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.931 in government revenue in 2020.

With a price elasticity of -0.28, dividing the total WTP of \$0.627 by the total cost of \$0.931 yields an MVPF of 0.673 [0.550, 0.801] in 2020.

In 2005, the nominal gas price was \$2.315. Society was WTP -\$0.122 for reduced greenhouse gases, -\$0.052 for reduced local air pollution, and -\$0.112 for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP \$0.067 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.781. A \$1 increase in the gas tax costs the government \$0.046 in lost gas tax revenue and \$0.018 in lost corporate tax revenue. The policy change also earned the government \$0.002 in revenue from abated greenhouse gases, for a total cost of \$0.939. Dividing total WTP by total cost results in an MVPF of 0.831 in the context (2005) in which the authors estimated the price elasticity of gasoline.

State-level Crude Price Pass-through Variation (Kilian & Zhou 2024)

Kilian & Zhou (2024) uses state-level variation in the pass-through from oil shocks to retail gasoline prices from 1989 to 2022 to estimate an own price elasticity of gasoline of -.198 (s.e. .053). The authors' Table 7 (Column 4, Row 1) reports this elasticity. We construct two MVPFs with this elasticity: one using our harmonized 2020 parameters and another in the context of 2022.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases (-\$0.161), reduced local air pollution (-\$0.020), and reduced driving externalities (-\$0.124), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of -.198. Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of -\$0.002 and added environmental benefits smaller than -\$0.001. Gasoline producers are willing to pay \$0.042 for lost profits, and utilities less than -\$0.001 for increased profits. Summing these components yields a total WTP of \$0.736 for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government \$0.951 in 2020. In addition to the \$1 of revenue mechanically raised, the government lost \$0.052 in tax revenue (-\$0.040 from lost gas tax revenue, \$0.011 from lost corporate tax revenue collected from gas producers, and less than \$0.001 gained from utility profits). The government also spent less than \$0.0001 in EV subsidies

and raised \$0.003 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.951 in government revenue in 2020.

With a price elasticity of -0.198 , dividing the total WTP of \$0.736 by the total cost of \$0.951 yields an MVPF of 0.773 [0.656, 0.896] in 2020.

In 2022, the nominal gas price was \$3.443. Society was WTP $-\$0.104$ for reduced greenhouse gases, $-\$0.012$ for reduced local air pollution, and $-\$0.080$ for reduced driving externalities. We do account for EV substitution in years after 2011. Learning-by-doing generated a WTP by future consumers of $-\$0.004$ and added environmental benefits smaller than $-\$0.001$. Producers were WTP \$0.033 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.832. A \$1 increase in the gas tax costs the government \$0.022 in lost gas tax revenue and \$0.009 in lost corporate tax revenue. The policy change also earned the government \$0.002 in revenue from abated greenhouse gases, for a total cost of \$0.970. Dividing total WTP by total cost results in an MVPF of 0.858 in the context (2022) in which the authors estimated the price elasticity of gasoline.

National Crude Price Shock Variation (Gelman et al. 2023)

Gelman et al. (2023) uses cross-sectional variation in gasoline spending interacted with the 2014 price shock to crude oil to estimate an own price elasticity of gasoline of -0.164 (s.e. $.024$) with data spanning from 2013 to 2016. The authors' Table 5 (Column 2, Row 1) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2016.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases ($-\$0.133$), reduced local air pollution ($-\$0.016$), and reduced driving externalities ($-\$0.103$), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of -0.164 . Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of $-\$0.002$ and added environmental benefits smaller than $-\$0.001$. Gasoline producers are willing to pay \$0.035 for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of \$0.781 for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government \$0.960 in 2020. In addition to the \$1 of revenue mechanically raised, the government lost \$0.043 in tax revenue ($-\0.034 from lost gas tax revenue, \$0.009 from lost corporate tax revenue collected from gas producers, and less than \$0.001 gained from utility profits). The government also spent less than \$0.001 in EV subsidies and raised \$0.003 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.960 in government revenue in 2020.

With a price elasticity of -0.164 , dividing the total WTP of \$0.781 by the total cost of \$0.960 yields an MVPF of 0.814 [0.762, 0.869] in 2020.

In 2016, the nominal gas price was \$2.256. Society was WTP $-\$0.114$ for reduced greenhouse gases, $-\$0.018$ for reduced local air pollution, and $-\$0.091$ for reduced driving externalities. We do account for EV substitution for years after 2011. Producers were WTP \$0.040 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.816. A \$1 increase in the gas tax costs the government \$0.032 in lost gas tax revenue and \$0.011 in

lost corporate tax revenue. The policy change also earned the government \$0.002 in revenue from abated greenhouse gases, for a total cost of \$0.960. Dividing total WTP by total cost results in an MVPF of 0.850 in the context (2016) in which the authors estimated the price elasticity of gasoline.

National Gas Price Variation (Park & Zhao 2010)

Park & Zhao (2010) use variation in national gas prices between 1976 and 2008 to estimate an own price elasticity of gasoline of -0.161 (0.015). This elasticity was scraped from the authors' Figure 2. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2008.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases (-\$0.130), reduced local air pollution (-\$0.016), and reduced driving externalities (-\$0.100), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of -0.161. Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of -\$0.002 and added environmental benefits smaller than -\$0.001. Gasoline producers are willing to pay \$0.034 for lost profits, and utilities less than -\$0.001 for increased profits. Summing these components yields a total WTP of \$0.786 for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government \$0.961 in 2020. In addition to the \$1 of revenue mechanically raised, the government lost \$0.042 in tax revenue (\$0.033 from lost gas tax revenue, \$0.009 from lost corporate tax revenue collected from gas producers, and less than \$0.001 gained from utility profits). The government also spent less than \$0.001 in EV subsidies and raised \$0.003 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.961 in government revenue in 2020.

With a price elasticity of -0.161, dividing the total WTP of \$0.786 by the total cost of \$0.961 yields an MVPF of 0.818 [0.786, 0.852] in 2020.

In 2008, the nominal gas price was \$3.310. Society was WTP -\$0.058 for reduced greenhouse gases, -\$0.018 for reduced local air pollution, and -\$0.050 for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP \$0.035 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.909. A \$1 increase in the gas tax costs the government \$0.019 in lost gas tax revenue and \$0.009 in lost corporate tax revenue. The policy change also earned the government \$0.001 in revenue from abated greenhouse gases, for a total cost of \$0.973. Dividing total WTP by total cost results in an MVPF of 0.934 in the context (2008) in which the authors estimated the price elasticity of gasoline.

National Gas Price Variation (Hughes et al. 2008)

Hughes et al. (2008) uses monthly variation in gasoline prices at the national level from 2001 to 2006 to estimate an own price elasticity of gasoline of -0.042 (s.e. 0.009). The authors' Table 1 (Column 2, Row 2) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2006.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases (-\$0.034), reduced local air pollution (-\$0.004), and reduced driving externalities (-\$0.026), all calculated by dividing the per-gallon externality by the price of

gasoline in 2020 (\$2.27) and multiplying by the price elasticity of $-.042$. Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of $-\$0.002$ and added environmental benefits smaller than $-\$0.001$. Gasoline producers are willing to pay $\$0.009$ for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of $\$0.943$ for the policy change in 2020.

A $\$1$ increase in the gas tax rate raised the government $\$0.990$ in 2020. In addition to the $\$1$ of revenue mechanically raised, the government lost $\$0.011$ in tax revenue ($-\$0.009$ from lost gas tax revenue, $\$0.002$ from lost corporate tax revenue collected from gas producers, and less than $\$0.001$ gained from utility profits). The government also spent around $\$0.0001$ in EV subsidies and raised $\$0.001$ by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a $\$1$ increase in the gas tax rate generated $\$0.990$ in government revenue in 2020.

With a price elasticity of $-.042$, dividing the total WTP of $\$0.943$ by the total cost of $\$0.990$ yields an MVPF of 0.953 [$0.939, 0.968$] in 2020.

In 2006, the nominal gas price was $\$2.622$. Society was WTP $-\$0.017$ for reduced greenhouse gases, $-\$0.006$ for reduced local air pollution, and $-\$0.015$ for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP $\$0.009$ for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of $\$0.970$. A $\$1$ increase in the gas tax costs the government $\$0.006$ in lost gas tax revenue and $\$0.002$ in lost corporate tax revenue. The policy change also earned the government less than $\$0.001$ in revenue from abated greenhouse gases, for a total cost of $\$0.992$. Dividing total WTP by total cost results in an MVPF of 0.978 in the context (2006) in which the authors estimated the price elasticity of gasoline.

Almost Ideal Demand System (West & Williams 2007)*

West & Williams (2007) uses an Almost Ideal Demand System model estimated on data from 1996 to 1998 and finds an own price elasticity of gasoline of $-.46$. The authors report separate elasticities for one-adult and two-adult households in Table 4 (Panel 2, Row 1, Column 1) and Table 5 (Panel 2, Row 1, Column 1) respectively. The authors also report the average gasoline consumption per week for one-adult and two-adult households in Table 1 (Row 1). Our final elasticity of $-.46$ comes from taking an average of the reported elasticities for one and two adult households weighted by their average gasoline consumption. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 1998.

In 2020, consumers were WTP $\$1$ for a $\$1$ increase in the gas tax rate. Society was WTP for reduced greenhouse gases ($-\$0.373$), reduced local air pollution ($-\$0.045$), and reduced driving externalities ($-\$0.286$), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 ($\$2.27$) and multiplying by the price elasticity of $-.46$. Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of $-\$0.002$ and added environmental benefits smaller than $-\$0.001$. Gasoline producers are willing to pay $\$0.097$ for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of $\$0.391$ for the policy change in 2020.

A $\$1$ increase in the gas tax rate raised the government $\$0.888$ in 2020. In addition to the $\$1$ of revenue mechanically raised, the government lost $\$0.119$ in tax revenue ($-\$0.094$ from lost gas tax revenue, $\$0.026$ from lost corporate tax revenue collected from gas producers, and less than

\$0.001 gained from utility profits). The government also spent around \$0.0001 in EV subsidies and raised \$0.007 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.888 in government revenue in 2020.

With a price elasticity of -0.46 , dividing the total WTP of \$0.391 by the total cost of \$0.888 yields an MVPF of 0.440 in 2020.

In 1998, the nominal gas price was \$1.072. Society was WTP $-\$0.295$ for reduced greenhouse gases, $-\$0.269$ for reduced local air pollution, and $-\$0.336$ for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP \$0.170 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.270. A \$1 increase in the gas tax costs the government \$0.160 in lost gas tax revenue and \$0.045 in lost corporate tax revenue. The policy change also earned the government \$0.006 in revenue from abated greenhouse gases, for a total cost of \$0.800. Dividing total WTP by total cost results in an MVPF of 0.337 in the context (1998) in which the authors estimated the price elasticity of gasoline.

Quadratic Almost Ideal Demand System (Tiezzi & Verde 2016)*

Tiezzi & Verde (2016) uses a Quadratic Almost Ideal Demand System model estimated on data from 2007 to 2010 and finds an own price elasticity of gasoline of -0.435 (s.e. 0.027). The authors' Table 4 (Column 5, Row 7) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2010.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases ($-\$0.354$), reduced local air pollution ($-\$0.043$), and reduced driving externalities ($-\$0.272$), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of -0.435 . Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of $-\$0.002$ and added environmental benefits smaller than $-\$0.001$. Gasoline producers are willing to pay \$0.093 for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of \$0.421 for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government \$0.893 in 2020. In addition to the \$1 of revenue mechanically raised, the government lost \$0.113 in tax revenue ($-\0.089 from lost gas tax revenue, \$0.025 from lost corporate tax revenue collected from gas producers, and less than \$0.001 gained from utility profits). The government also spent less than \$0.0001 in EV subsidies and raised \$0.007 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.893 in government revenue in 2020.

With a price elasticity of -0.435 , dividing the total WTP of \$0.421 by the total cost of \$0.893 yields an MVPF of 0.472 in 2020.

In 2010, the nominal gas price was \$2.835. Society was WTP $-\$0.193$ for reduced greenhouse gases, $-\$0.049$ for reduced local air pollution, and $-\$0.162$ for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP \$0.092 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.687. A \$1 increase in the gas tax costs the government \$0.062 in lost gas tax revenue and \$0.024 in lost corporate tax revenue. The policy change also earned the government \$0.004 in

revenue from abated greenhouse gases, for a total cost of \$0.918. Dividing total WTP by total cost results in an MVPF of 0.749 in the context (2010) in which the authors estimated the price elasticity of gasoline.

Multimarket Simulation Model (Bento et al. 2009)*

Bento et al. (2009) uses a multimarket simulation model with 2002 regional gas price data to estimate an own price elasticity of gasoline of $-.35$. The authors' Table 4 (Column 1, Row 1) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2002.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases ($-\$0.285$), reduced local air pollution ($-\$0.035$), and reduced driving externalities ($-\$0.219$), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of $-.35$. Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of $-\$0.002$ and added environmental benefits smaller than $-\$0.001$. Gasoline producers are willing to pay \$0.074 for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of \$0.534 for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government \$0.914 in 2020. In addition to the \$1 of revenue mechanically raised, the government lost \$0.091 in tax revenue ($-\0.072 from lost gas tax revenue, \$0.020 from lost corporate tax revenue collected from gas producers, and less than \$0.001 gained from utility profits). The government also spent less than \$0.0001 in EV subsidies and raised \$0.006 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.914 in government revenue in 2020.

With a price elasticity of $-.35$, dividing the total WTP of \$0.534 by the total cost of \$0.914 yields an MVPF of 0.584 in 2020.

In 2002, the nominal gas price was \$1.386. Society was WTP $-\$0.216$ for reduced greenhouse gases, $-\$0.132$ for reduced local air pollution, and $-\$0.218$ for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP \$0.109 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.542. A \$1 increase in the gas tax costs the government \$0.095 in lost gas tax revenue and \$0.029 in lost corporate tax revenue. The policy change also earned the government \$0.004 in revenue from abated greenhouse gases, for a total cost of \$0.881. Dividing total WTP by total cost results in an MVPF of 0.616 in the context (2002) in which the authors estimated the price elasticity of gasoline.

National Gas Price Variation (Hughes et al. 2008)*

Hughes et al. (2008) uses monthly variation in gasoline prices at the national level from 1975 to 1980 to estimate an own price elasticity of gasoline of -0.335 (s.e. 0.024). The authors' Table 1 (Column 1, Row 2) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 1980.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases ($-\$0.272$), reduced local air pollution ($-\$0.033$), and reduced driving externalities ($-\$0.209$), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of $-.335$. Society's WTPs for

global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of $-\$0.002$ and added environmental benefits smaller than $-\$0.001$. Gasoline producers are willing to pay $\$0.071$ for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of $\$0.554$ for the policy change in 2020.

A $\$1$ increase in the gas tax rate raised the government $\$0.918$ in 2020. In addition to the $\$1$ of revenue mechanically raised, the government lost $\$0.087$ in tax revenue ($-\$0.068$ from lost gas tax revenue, $\$0.019$ from lost corporate tax revenue collected from gas producers, and less than $\$0.001$ gained from utility profits). The government also spent less than $\$0.0001$ in EV subsidies and raised $\$0.005$ by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a $\$1$ increase in the gas tax rate generated $\$0.918$ in government revenue in 2020.

With a price elasticity of -0.335 , dividing the total WTP of $\$0.554$ by the total cost of $\$0.918$ yields an MVPF of 0.604 in 2020.

In 1980, the nominal gas price was $\$0.83$. Society was WTP $-\$0.141$ for reduced greenhouse gases, $-\$0.271$ for reduced local air pollution, and $-\$0.201$ for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP $\$0.115$ for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of $\$0.503$. A $\$1$ increase in the gas tax costs the government $\$0.086$ in lost gas tax revenue and $\$0.031$ in lost corporate tax revenue. The policy change also earned the government $\$0.003$ in revenue from abated greenhouse gases, for a total cost of $\$0.886$. Dividing total WTP by total cost results in an MVPF of 0.567 in the context (1980) in which the authors estimated the price elasticity of gasoline.

State-level Crude Price Pass-through Variation (Kilian & Zhou 2024)*

Kilian & Zhou (2024) uses state-level variation in the pass-through from oil shocks to retail gasoline prices from 1989 to 2014 to estimate an own price elasticity of gasoline of -0.314 (s.e. $.066$). The authors' Table 4 (Column 4, Row 1) reports this elasticity. Although this precise point estimate is based on 1989 to 2008 data, the author also uses an extended data sample (through 2022) and reports that the elasticity remains stable around -0.3 until the end of 2014. Thus, we construct two MVPFs with this elasticity: one using our harmonized 2020 parameters and another in the context of 2014.

In 2020, consumers were WTP $\$1$ for a $\$1$ increase in the gas tax rate. Society was WTP for reduced greenhouse gases ($-\$0.255$), reduced local air pollution ($-\$0.031$), and reduced driving externalities ($-\$0.196$), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 ($\$2.27$) and multiplying by the price elasticity of -0.314 . Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of $-\$0.002$ and added environmental benefits smaller than $-\$0.001$. Gasoline producers are willing to pay $\$0.067$ for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of $\$0.582$ for the policy change in 2020.

A $\$1$ increase in the gas tax rate raised the government $\$0.923$ in 2020. In addition to the $\$1$ of revenue mechanically raised, the government lost $\$0.082$ in tax revenue ($\$0.064$ from lost gas tax revenue, $\$0.018$ from lost corporate tax revenue collected from gas producers, and less than $\$0.001$ gained from utility profits). The government also spent less than $\$0.0001$ in EV subsidies and raised $\$0.005$ by abating carbon emissions. Combining the mechanical revenue raised with

the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.923 in government revenue in 2020.

With a price elasticity of -0.314 , dividing the total WTP of \$0.582 by the total cost of \$0.923 yields an MVPF of 0.630 in 2020.

In 2014, the nominal gas price was \$3.443. Society was WTP $-\$0.136$ for reduced greenhouse gases, $-\$0.024$ for reduced local air pollution, and $-\$0.110$ for reduced driving externalities. We do account for EV substitution in years after 2011, although society was only willing to pay less than $-\$0.0001$. Producers were WTP \$0.072 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.801. A \$1 increase in the gas tax costs the government \$0.037 in lost gas tax revenue and \$0.019 in lost corporate tax revenue. The policy change also earned the government \$0.003 in revenue from abated greenhouse gases, for a total cost of \$0.946. Dividing total WTP by total cost results in an MVPF of 0.847 in the context (2014) in which the authors estimated the price elasticity of gasoline.

State-level Gas Price Variation (Small & Van Dender 2007)*

Small & Van Dender (2007) use variation in state-level fuel prices between 1997 and 2001 to estimate a short-run own price elasticity of gasoline of -0.067 (s.e. 0.009). The authors' Table 5 (Column 1, Row 7) reports this elasticity. We construct two MVPFs: one using our harmonized 2020 parameters and another in the context of 2001.

In 2020, consumers were WTP \$1 for a \$1 increase in the gas tax rate. Society was WTP for reduced greenhouse gases ($-\$0.054$), reduced local air pollution ($-\$0.007$), and reduced driving externalities (-0.042), all calculated by dividing the per-gallon externality by the price of gasoline in 2020 (\$2.27) and multiplying by the price elasticity of -0.067 . Society's WTPs for global and local pollution are inclusive of added damages from increased EV charging and manufacturing. Learning-by-doing generated a WTP by future consumers of $-\$0.002$ and added environmental benefits smaller than \$0.001. Gasoline producers are willing to pay \$0.014 for lost profits, and utilities less than $-\$0.001$ for increased profits. Summing these components yields a total WTP of \$0.910 for the policy change in 2020.

A \$1 increase in the gas tax rate raised the government \$0.984 in 2020. In addition to the \$1 of revenue mechanically raised, the government lost \$0.017 in tax revenue (\$0.014 from lost gas tax revenue, \$0.004 from lost corporate tax revenue collected from gas producers, and less than \$0.001 gained from utility profits). The government also spent an additional \$0.0001 in EV subsidies and raised \$0.001 by abating carbon emissions. Combining the mechanical revenue raised with the fiscal externalities, a \$1 increase in the gas tax rate generated \$0.984 in government revenue in 2020.

With a price elasticity of -0.067 , dividing the total WTP of \$0.910 by the total cost of \$0.984 yields an MVPF of 0.925 in 2020.

In 2001, the nominal gas price was \$1.47. Society was WTP $-\$0.037$ for reduced greenhouse gases, $-\$0.025$ for reduced local air pollution, and \$0.039 for reduced driving externalities. We do not account for EV substitution for years before 2011. Producers were WTP \$0.020 for the policy change. Summing consumers', producers', and society's WTPs yields a total WTP of \$0.919. A \$1 increase in the gas tax costs the government \$0.017 in lost gas tax revenue and \$0.005 in lost corporate tax revenue. The policy change also earned the government \$0.001 in revenue from abated greenhouse gases, for a total cost of \$0.978. Dividing total WTP by total cost results in an MVPF of 0.940 in the context (2001) in which the authors estimated the price

elasticity of gasoline.

E.11 Other Fuel Taxes

In this section, we describe how we form MVPFs for six taxes on fuels other than gasoline. Our main sample includes MVPFs for fuel taxes on jet fuel and diesel fuel. Our extended sample includes four additional MVPFs (for fuel taxes on heavy fuel oil, flex fuel, and crude oil). We highlight in those sections why we include these policies in our extended sample.

E.11.1 Tax on Jet Fuel

Taxing jet fuel reduce the quantity of fuel consumed while generating revenue for the government. The MVPF for a tax on jet fuel combines the price elasticity of jet fuel with a measure of the value of the externalities generated per dollar of spending on jet fuel. We use a price elasticity of jet fuel from Fukui & Miyoshi (2017), who rely on historical variation in jet fuel price and consumption data between 1995 to 2013 to estimate the responsiveness of US airlines' fuel use to fuel prices. The authors employ a quantile regression to estimate the long-run price elasticity (authors' Table 7). We take the coefficient (-0.166, s.e. 0.0836) estimated for airlines at the median of the fuel consumption distribution but note that our results would vary if one used estimates of the responsiveness of airlines at different ends of the fuel consumption distribution.

Consumers' WTP The envelope theorem implies that consumers value the policy change at the value of the price increase. In other words, consumers are willing to pay \$1 for a \$1 change in their cost of jet fuel, holding their consumption of jet fuel constant. Following our treatment of gasoline taxes, we assume the \$1 increase in the price of jet fuel is completely passed onto consumers.

Society's WTP In response to higher fuel prices, airlines reduce the number of miles flown, either by reducing the number of trips flown or the average distance of trips. This reduction in fuel usage benefits society through less global and local air pollution. CO_2 emissions from jet fuel use are a function of the carbon content of jet fuel (9,752.236 grams per gallon), which we take from EIA (2023b). Converting the carbon content from grams to tons and multiplying by the social cost of carbon gives us society's WTP for airlines to burn one fewer gallon of jet fuel, or \$1.88 per gallon in 2020 when using an SCC of \$193 (in 2020 dollars). In our in-context specification (2013), carbon damages from jet fuel were \$1.45 per gallon (in 2013 dollars).

We also account for global and local emissions from jet fuel production. We explain how we account for upstream emissions in Appendix C.4.3. These externalities are equivalent to the upstream externalities that enter our diesel and heavy fuel oil MVPFs. They differ from our upstream gasoline externalities as we need not adjust for ethanol production. In 2020, upstream emissions generated an additional \$0.231 (in 2020 dollars) per gallon of petroleum product produced. In 2013, upstream emissions generated an additional \$0.177 (in 2013 dollars) per gallon of petroleum product produced.

Combining damages from CO_2 and upstream fuel production, we estimate that jet fuel imposed \$2.11 of damages per gallon burned in 2020, with the majority of damages (\$2.09)

coming from greenhouse gases (both upstream and in air emissions) and the remaining \$0.02 from upstream air pollution. We adjust damages from greenhouse gases by the share of global damages that do not flow to the US government (0.981), resulting in a total externality of \$2.07 per gallon in 2020 and \$1.60 in 2013 (both in nominal dollars).¹⁹⁹ Using price data from the EIA’s “U.S. Gulf Coast Kerosene-Type Jet Fuel Spot Price FOB” data series (EIA 2024f), our externality calculations imply that society faced \$1.89 in total damages per dollar of spending on jet fuel in 2020, when the price of jet fuel was \$1.10 per gallon. In 2013, when the price of jet fuel was \$2.92 per gallon, society faced \$0.55 (in 2013 dollars) in total damages per dollar of spending on jet fuel.

Multiplying the externality generated per dollar of spending on jet fuel by the own-price elasticity of jet fuel gives us society’s WTP for a \$1 change in the tax on jet fuel. With a price elasticity of -0.166 from Fukui & Miyoshi (2017), society was willing to pay \$0.313 ($-0.166 \times \1.89, in 2020 dollars) in 2020 for pollution abated due to reduced jet fuel consumption. In 2013, society was willing to pay \$0.09 ($-0.166 \times \0.55, in 2013 dollars) for pollution abated due to reduced jet fuel consumption. We do not consider any rebound effects when evaluating a tax on jet fuel.

Assessing local damages from aviation requires information on where air pollution is released, as estimates of the social costs of local air pollutants (such as those from AP3) take as an input the location of emissions. For example, if aviation generates large quantities of emissions while planes are taxiing, and airports are located near large population centers, we would need to assign large social costs to the pollution from aviation to account for the large number of people exposed to the pollution (see Schlenker & Walker 2015). Additionally, since AP3 takes as an input the height at which emissions are released, we would also require information about where along a plane’s flight path it releases air pollution to account for the dispersion of pollution along air currents (see Taylor & Du 2024). Despite these concerns, we note that incorporating benefits from abating local air pollution would make a tax on jet fuel a more efficient way to raise revenue.

We attempt to quantify potential local air pollution benefits by assuming the total quantity of emissions assigned to aircraft according to the National Emissions Inventory (EPA 2023a). Since aircraft can use jet fuel or aviation gasoline, we determine what share of local emissions reported by the NEI aircraft that use jet fuel are responsible for by assuming total emissions are proportional to the share of aviation fuel consumed that is jet fuel. Data on aviation gasoline and jet fuel consumption come from the “U.S. Product Supplied of Aviation Gasoline (Thousand Barrels)” (EIA 2024i) and “U.S. Product Supplied of Kerosene-Type Jet Fuel (Thousand Barrels)” (EIA 2024j), both annual series published by the EIA.²⁰⁰ After multiplying the total quantity of emissions by the share of aviation fuel consumed that is jet fuel (99% in 2020), we convert reported emissions from short to metric tons and divide by the total quantity of jet fuel consumed in the year of analysis, where quantity of jet fuel consumed is expressed in gallons. This calculation gives us an average quantity of pollution released per gallon of jet fuel consumed, assuming emissions from all aircraft are proportional to the quantity of each fuel type consumed. We apply this process to four local pollutants documented in the NEI: CO , NO_X , $PM_{2.5}$, and VOC . For our in-context specification, we use NEI data from 2014, the

¹⁹⁹Although upstream emissions vary over time, the rising social cost of carbon is largely responsible for the rise in the externality from jet fuel. In 2020 dollars, the externality from CO_2 rose from \$1.62 in 2013 to \$1.88, accounting for roughly 90% of the real change in the per-gallon externality associated with jet fuel consumption.

²⁰⁰Given the relatively small quantity of aviation gasoline consumed annually, this method assigns nearly all pollution to aircraft that use jet fuel.

closest year available to 2013. We apply to these quantities our baseline social cost estimates but note that this application ignores spatial variation in aviation emissions.

As an example of this calculation, the NEI reports that aircraft released 7,360 short tons of primary $PM_{2.5}$ in 2020, or 6,676.88 metric tons. Assuming aircraft that use jet fuel are responsible for 98.99% of these emissions (since 98.99% of all aviation fuel consumed is jet fuel, jet fuel usage resulted in 6,609.74 metric tons of $PM_{2.5}$ pollution in 2020, or 3.99×10^{-7} metric tons (0.399 grams) per gallon, since 16.55 billion gallons (393,976 barrels) of jet fuel were consumed in 2020. Multiplying by a social cost of \$105,127.64 in 2020, we obtain a WTP for $PM_{2.5}$ damages of \$0.042 per gallon of jet fuel consumed.

In addition to applying the above process to all four pollutants reported by the NEI, we account for sulfur released from burning jet fuel, which follows from the fuel's sulfur content. We take this to be the midpoint (600 ppm) of the range of sulfur contents reported by PARTNER (2011), or 1.83 grams of SO_2 per gallon, assuming a density of jet fuel of 807.5 kilograms per cubic meter (ExxonMobil n.d.). Sulfur released from burning jet fuel adds an additional \$0.085 per gallon to jet fuel's externality in 2020, using our baseline social cost of SO_2 of \$46,491.03.

If we were to include local damages from SO_2 , CO , NO_X , $PM_{2.5}$, and VOC , the total per-gallon externality associated with jet fuel consumption in 2020 rises from \$2.11 per gallon to \$2.36 per gallon. Local damages rise from \$0.02 to \$0.25 per gallon, while global damages rise from \$2.09 to \$2.11 per gallon. In 2013, the total externality rises from \$1.63 to \$1.86 (in 2013 dollars), with local damages increasing by \$0.22 and global damages by \$0.01. We report below what the MVPF would be both in 2020 and 2013 if one were to use the adjusted externalities calculated above. However, we do not include these additional local damages in our baseline specification due to uncertainty about which social costs to apply.

Producers' WTP We account for producers' WTP for lost profits resulting from reduced jet fuel consumption. Specifically, since jet fuel production involves the same three processes required to produce gasoline, we assume the percent markup imposed on gasoline by crude oil producers, refiners, and distributors holds for other fuels. We describe how we calculate this 27% markup in Appendix C.4.5. This percent markup is net of the assumed 8% economy-wide markup estimated by De Loecker et al. (2020).

Applying the 27% net markup to the price of jet fuel in 2020 (\$1.10) implies a per-gallon markup of \$0.298. We adjust by the corporate average tax rate (21%) to account for the share of profits producers keep (Watson 2022).²⁰¹ This results in a post-tax externality borne by producers of \$0.235 per dollar of spending on jet fuel. With a price elasticity of -0.166, producers were willing to pay \$0.036 in 2020 for the policy change. In 2013, when the percent markup was 27.6%, producers were willing to pay \$0.036 (in 2013 dollars), using the 2013 price of jet fuel and holding the corporate tax rate fixed.

Total WTP Summing across components, a \$1 change in the jet fuel tax rate results in a total WTP of \$0.722 in 2020 when using a price elasticity of jet fuel of -0.166. Consumers (\$1) and producers (\$0.036) are both willing to pay to avoid the tax increase, while society (-\$0.313) is willing to pay to keep the tax increase. We sign each component depending on the group's willingness to pay to remove the tax. Consumers and producers are both willing to pay to remove the tax since these groups are made worse off through higher prices and reduced profits,

²⁰¹We do not vary across time the effective corporate tax rate gasoline producers face.

respectively. On the other hand, society is willing to pay to keep the tax on the books, as they are made better off through reduced environmental externalities. Considering the removal of a tax (rather than the addition of one) also allows us to normalize both taxes and subsidies to a positive \$1 mechanical cost. Using the externality values calculated for 2013, we obtain a total WTP of \$0.945 (in 2013 dollars), with society willing to pay \$0.09 for increased pollution and producers \$0.036 for increased profits.

Cost A \$1 increase in the jet fuel tax rate mechanically raises \$1 of revenue for the government. However, the accompanying decrease in jet fuel consumption reduces the amount of revenue the government collects by the size of the behavioral response times the tax collected per dollar of jet fuel spending. In 2020, the federal jet fuel tax rate was \$0.219 per gallon while the average state tax on jet fuel was \$0.036 (EIA 2024c), meaning the government collected \$0.232 per dollar spent on jet fuel.²⁰² We hold these tax rates fixed over time and do not adjust for inflation. Multiplying by a price elasticity of -0.166, the government faced a \$0.039 loss in revenue from decreased jet fuel consumption. Applying the same method to the in-context prices and taxes (when the average state tax was \$0.036 per gallon and the federal rate was \$0.219 per gallon) results in a \$0.0145 fiscal externality from lost jet fuel tax revenue.

Decreases in producer profits reduce government revenue in the form of lost corporate taxes (assuming a 21% tax rate). The pre-tax markup on jet fuel was \$0.298 per gallon, meaning the government collected \$0.057 in corporate tax revenue for each dollar spent on jet fuel in 2020, when the price was \$1.10 per gallon.²⁰³ With our price elasticity of -0.166, we calculate a \$0.009 fiscal externality from lost corporate tax revenue. Applying the same method to the in-context prices and markups (noted above) results in a \$0.010 fiscal externality from lost corporate tax revenue.

Finally, abating greenhouse gas emissions through a jet fuel tax raises revenue for the government in the long run. When calculating society's WTP for global pollution benefits, we scale the WTP component by the share of global benefits that do not flow to the US government as revenue (98.1%). We put the remaining 1.92% in the denominator. With a price elasticity of -0.166, the total, unadjusted WTP in 2020 for *all* global benefits was \$0.316, implying the government generated \$0.006 ($\0.316×0.0192) in revenue by abating carbon emissions today and promoting economic output tomorrow. Abating carbon emissions generates a positive fiscal externality, as the US government generates additional revenue through the policy change. In context, the climate FE equaled \$0.0018.

Summing the mechanical \$1 of revenue raised and the three fiscal externalities, we obtain a total "cost" of \$0.958 when using a price elasticity of -0.166: a \$1 increase in the gas tax rate raises \$0.958 in revenue for the government. To match our approach with subsidies and other revenue raisers, we again consider the effects of removing the tax. This allows us to treat both taxes and subsidies as having a \$1 mechanical cost. The government loses \$1 in revenue by decreasing the tax rate by \$1, raises \$0.048 in revenue from jet fuel taxes (\$0.039) and corporate taxes (\$0.009) by encouraging jet fuel consumption, and loses \$0.006 by increasing carbon emissions and decreasing long-run revenue, for a net government cost of \$0.958 in 2020.

²⁰²To calculate the average state tax on jet fuel, we take the state-specific tax rates reported by EIA (2024c) and weight by the quantity of jet fuel consumed in that state as reported by EIA (2022a). We set a state's tax on jet fuel equal to \$0 per gallon if a tax rate is not reported or if the state lacks a tax specifically levied on jet fuel. We do not adjust this tax rate for inflation since fuel taxes are typically not indexed for inflation.

²⁰³We assume all producer profits are subject to the American tax schedule but note that the geographic variation in crude oil production could render our calculation of this fiscal externality an overestimate.

Using our 2013 specifications and applying the same calculations, we obtain a net government cost of \$0.978, with the government gaining \$0.024 in jet fuel tax and corporate tax revenue and losing \$0.002 in long-run revenue due to increased carbon emissions.

MVPF Dividing the total WTP calculated above (\$0.722) by the total cost (\$0.958), both calculated with a price elasticity of -0.166, we form an MVPF of 0.754 with a 95 percent confidence interval of [0.563, 0.936] in 2020. Using our in-context estimates of a total WTP of \$0.945 and a net government cost of \$0.978, we obtain an MVPF of 0.967.

If we were to use our estimates of society's WTP that include increased local pollution from aircraft, we obtain an MVPF of 0.715 in 2020 (\$0.685/\$0.958) and 0.954 in 2013 (\$0.932/\$0.978).

E.11.2 Tax on Diesel Fuel

Taxing diesel fuel reduce the quantity of fuel consumed while generating revenue for the government. The MVPF for a tax on diesel fuel combines the price elasticity of diesel fuel with a measure of the value of the externalities generated per dollar of spending on diesel fuel. We use a price elasticity of diesel fuel from Dahl (2012), who reports the price and income elasticities for gasoline and diesel by country. We use their US price elasticity of diesel of -0.07 (author's Table 1). We do not form a confidence interval for this MVPF because the paper does not include standard errors. Our in-context specification corresponds to 2006.

Consumers' WTP The envelope theorem implies that consumers value the policy change at the value of the price increase. In other words, consumers are willing to pay \$1 for a \$1 change in their cost of diesel fuel, holding their consumption of diesel fuel constant. Following our treatment of gasoline taxes, we assume the \$1 increase in the price of diesel fuel is completely passed onto consumers.

Society's WTP Our valuation of the per-gallon externality from diesel fuel closely mirrors our approach to estimating the externality from gasoline. Upstream externalities are common across both fuel types, as we calculate upstream externalities per gallon of petroleum product produced (see Appendix C.4.3). In 2006, upstream emissions generated \$0.138 in damages per gallon of petroleum product produced and in 2020 generated \$0.264 in damages per gallon, both in nominal dollars. We do not consider biodiesel blends. We use the same VMT-weighted social costs used to value all on-road emissions but note that social costs could differ if VMT from diesel-powered vehicles follows a different spatial distribution than gas-powered vehicles.

Regardless of vehicle type, burning one gallon of diesel fuel releases 10,183.15 grams of CO_2 (EIA 2023b), resulting in a willingness to pay of \$1.102 per gallon in 2006 and \$1.965 per gallon in 2020, both in nominal dollars. Similarly, one gallon of diesel fuel released 403 ppm of SO_2 in 2006 and 15 ppm in 2020. To calculate the sulfur content of diesel, we assume that sulfur content regulations were perfectly binding, meaning one gallon of diesel contained 5,000 ppm of sulfur in years before 1994, 500 ppm of sulfur between 1994 and 2005, and 15 ppm of sulfur for 2010 onward (EIA 2015). We take a linear interpolation to find specific sulfur contents between 2006 and 2009, assuming the sulfur content of diesel fell at a constant rate in this period. We convert from ppm to grams per gallon using a density of diesel of 0.85 kilograms

per liter (Speight 2011). We find a willingness to pay of \$0.095 per gallon in 2006 and \$0.0046 per gallon in 2020, both in nominal dollars.

Other emissions from diesel vehicles vary with the type of vehicle (light-, medium-, or heavy-duty) consuming the fuel. We outline how we calculate emission rates for each vehicle type and then explain how we weight across vehicle types to construct a single per-gallon externality from diesel fuel. We consider a fleet of diesel vehicles where the oldest possible model year is 1975.

For light-duty, diesel-powered vehicles, we separately consider emissions from light-duty cars and trucks. To calculate emission rates for all diesel-powered light-duty vehicles, we use the reported share of light-duty cars and trucks that contained diesel power trains reported in the EPA’s Automotive Trends Report (EPA 2023*d*) to weight across light-duty vehicle types and within a given model year. For CH_4 and N_2O emissions, we take diesel-specific emission rates from Cai et al. (2013). We crosswalk these emission rates with the EPA’s production shares by assuming that GREET’s “Passenger Car” classification corresponds to the EPA’s “All Car” classification, GREET’s “Passenger Truck” classification corresponds to the EPA’s “Truck SUV” classification, and GREET’s “Light-Duty Commercial Truck” corresponds to the EPA’s “Light Truck” classification. We assume vehicles from model years released before the series began emit at the rate of vehicles released in the first year reported. As with gas-powered vehicles, we do not consider changes in CH_4 or N_2O emission rates throughout the vehicle’s lifetime.

For all other emissions from light-duty vehicles (HC , CO , NO_x , and $PM_{2/5}$ from exhaust and tires and brakes), we calculate the percent difference between gas- and diesel-powered vehicles reported by (DOT 2024) and apply these percent differences to our preferred light-duty, gas-vehicle emission rates from Jacobsen et al. (2023) and Cai et al. (2013) (see Appendix C.4.1). We calculate percent differences for model years 2000 onward. We assume vehicles from model years released before the series began emit at the rate of vehicles released in the first year reported. We focus on light-duty trucks and cars here. This approach tells us how much more or less polluting a diesel vehicle is than a gas vehicle of the same vehicle type and model year.²⁰⁴ We convert from per-mile to per-gallon emission rates by adjusting the fuel economies reported in the EPA’s Automotive Trends Report (EPA 2023*d*) by the percent difference between gas- and diesel-powered light-duty cars (25.5%) and light-duty trucks (24.5%) in the fuel economy data reported in the 2020 Annual Energy Outlook (Table 40) (EIA 2023*a*).²⁰⁵ These adjustments imply light-duty, diesel-powered vehicles are more fuel efficient than their gasoline counterparts. While we hold the adjustment factor fixed, the fuel economy of gas-powered vehicles varies with vehicle model year. We assume gas- and diesel-powered light-duty vehicles face the same rate and duration of emissions system decay (e.g., annual increase in emissions) as gas-powered vehicles.

For diesel-powered, light-duty vehicles, we use the same per-mile accident and congestion externalities described in Appendix C.4.2. We assume the entire change in diesel consumption arises from changes in VMT (e.g., $\beta = 1$). We assume light-duty vehicle generate no road damage externalities.

²⁰⁴DOT (2024) does not separately isolate emissions from light-duty cars, so we assume light-duty vehicles reflects emissions from light-duty cars.

²⁰⁵We specifically calculate the percent difference between vehicles released in model year 2019, the earliest available year in the 2020 Annual Energy Outlook time series. Our results would be similar had we calculated the percent difference using the 2020 model year data. We hold this percent difference fixed over time. This percent difference is in line with the fuel economy difference reported by [fueleconomy.gov](https://www.fueleconomy.gov).

Our approach to valuing externalities from medium- and heavy-duty vehicles closely mirrors the approach described above for light-duty vehicles, but we need not consider differences in emissions between different types of medium- and heavy-duty vehicles. Whereas there were light-duty cars and trucks that consumed diesel, diesel-powered medium- and heavy-duty vehicles are almost always trucks.²⁰⁶ In most instances, emissions rates for medium-duty vehicles are not reported. We assume that medium-duty vehicles have the same emission rates as heavy-duty vehicles but note that this is an upper-bound on our externality calculation. We use the same emission rates data from the DOT (2024) to calculate the percent difference between heavy-duty, diesel-powered trucks and gasoline-powered, light-duty vehicles.²⁰⁷ CH_4 and N_2O emission rates for diesel both come from Cai et al. (2013). While we do not separately identify per-mile emission rates for medium-duty vehicles, we apply type-specific fuel economies to medium- and heavy-duty vehicles. We take fuel economy distributions for medium- and heavy-duty vehicles from the 2021 Vehicle Inventory and Use Survey (VIUS) (DOT 2023), taking the midpoint of the reported fuel economy range and weighting using the reported sample sizes.²⁰⁸ We hold these fuel economies fixed over time. In our calculations, medium-duty vehicles received 12.02 miles per gallon while heavy-duty vehicles received 6.08 miles per gallon.

For diesel-powered, medium- and heavy-duty vehicles, we use the same accident externality for light-duty vehicles.²⁰⁹ We scale up the per-mile congestion externality by 95.98%, the percent difference between the congestion externality estimated for automobiles and combination trucks in the 1997 Federal Highway Cost Allocation Study (authors' Table V-23, middle estimate for "All Highways"). We also assume medium- and heavy-duty vehicles impose a road damage externality of \$0.069 per mile, also from the 1997 Federal Highway Cost Allocation Study (authors' Table V-19, calculated for "All Combinations").²¹⁰ We again treat driving externalities from medium-duty vehicles as equal to those generated by heavy-duty vehicles.

Once we have calculated average emissions rates for light-, medium-, and heavy-duty vehicles for all pollutants, we weight across model years and vehicle types by either the quantity of diesel consumed or the number of miles traveled, depending on whether the externality arises on a per-gallon or per-mile basis (see Appendix C.4 for more). This mirrors how we weigh across model years when calculating our per-gallon externality for gasoline, but it differs in that we consider variation in diesel usage and VMT across light-, medium-, and heavy-duty vehicles. For light-duty vehicles, we assume diesel vehicles travel 4.31% more miles, which comes from the reported average difference between diesel and gasoline vehicles in the FHWA (2017).²¹¹ VMT for medium- and heavy-duty vehicles by vehicle age comes from the VIUS (DOT 2023).²¹²

²⁰⁶Following this fact, we assume all medium- and heavy-duty vehicles are trucks, and that all medium- and heavy-duty vehicles use diesel.

²⁰⁷As noted above, we weight across light-duty cars and trucks using production shares from EPA (2023*d*).

²⁰⁸Specifically, we use data reported in the Table 22, "Miles per Gallon by Registration State and Vehicle Size." When using data from the VIUS, we look at values calculated for the entire United States. We ignore data points without reported samples or without reported fuel economies. For medium-duty vehicles, we weight across the two reported classes, "Class 3, 4, 5, and 6" and "Class 7."

²⁰⁹While heavier vehicles may impose larger risks to other drivers when involved in accidents, differences in both driver quality and where these larger vehicles travel push the externality in the other direction. See Muehlenbachs et al. (2017) for more. We hold the per-mile accident externality fixed across vehicle types.

²¹⁰We assume this cost is reported in 1997 dollars. In 2020 dollars, the road damage externality from medium- and heavy-duty vehicles was \$0.11 per mile.

²¹¹The 2017 NHTS reports that all gas vehicles travel on average 10,069.91 miles annually, while diesel vehicles travel 10,503.85 annually. We calculate the percent difference in VMT between gas- and diesel-powered light-duty vehicles using these figures. We assume this difference is the same throughout the vehicle's lifetime. We do not separately consider VMT differences for light-duty cars and trucks.

²¹²VMT and age data come from VIUS Table 2, "Model Year by Registration State and Vehicle Size."

We assume vehicles older than the oldest vehicle reported in the VIUS travel the same VMT as the oldest vehicle in the survey. We note above how we handle data from the VIUS. We use the reported samples to construct age distributions for medium- and heavy-duty vehicles, assuming vehicles older than the oldest vehicle reported are evenly distributed.²¹³ Data on the diesel fleet composition come from the VIUS (DOT 2023).²¹⁴

After weighting across vehicle types and model years according to whether they arise per-mile or per-gallon, we have damages in dollar per gallon terms for the average vehicle in the diesel fleet in a given year.²¹⁵ In 2006, the total diesel externality was \$5.87 per gallon (in nominal dollars), with global damages contributing \$1.21, local pollution \$2.84, and driving damages \$1.81. In 2006, $PM_{2.5}$ from exhaust and NO_X together made up over 90% of local pollution damages. In 2020, the total diesel externality was \$6.90 per gallon (in nominal dollars), with global damages contributing \$2.19, local pollution \$2.36, and driving damages \$2.36. $PM_{2.5}$ from exhaust and NO_X together again made up over 90% of local pollution damages. We again note that we assume all changes in diesel consumption arise from changes in VMT ($\beta = 1$).

In our in-context specification, we use the author’s reported diesel price of \$2.61 per gallon. In 2020, we use a diesel price of \$2.55 per gallon from the EIA’s “U.S. No 2 Diesel Retail Prices” data series (EIA 2024h). Dividing the total externality by the price per gallon, we find each dollar spent on diesel generated \$2.247 and \$2.705 in damages in 2006 and 2020, respectively. Multiplying by our elasticity of -0.07, we find a willingness to pay by society for a change in environmental and driving externalities of \$0.16 and \$0.19 in 2006 and 2020, respectively. We adjust global damages for the share that flows to the US government as increased long-run revenue, which results in a WTP for global damages of \$0.032 and \$0.059 in 2006 and 2020, respectively.

Producers’ WTP We account for producers’ WTP for lost profits resulting from reduced diesel consumption. Specifically, since diesel production involves the same three processes required to produce gasoline, we assume the percent markup imposed on gasoline by crude oil producers, refiners, and distributors holds for other fuels. We describe how we calculate this 27% markup in Appendix C.4.5. This percent markup is net of the assumed 8% economy-wide markup estimated by De Loecker et al. (2020).

Applying the 27% net markup to the price of diesel in 2020 (\$2.55) implies a per-gallon markup of \$0.692. We adjust by the corporate average tax rate (21%) to account for the share of profits producers keep (Watson 2022).²¹⁶ This results in a post-tax externality borne by producers of \$0.213 per dollar of spending on diesel. With a price elasticity of -0.07, producers were willing to pay \$0.015 in 2020 for the policy change. In 2006, when the percent markup was 27.9%, producers were willing to pay \$0.015 (in 2006 dollars), using the 2006 price of diesel and holding the corporate tax rate fixed.

²¹³We assume the oldest model year in our sample is 1975 for consistency with emission rate data.

²¹⁴Specifically, fleet composition data come from Table 22 “Fuel Type and Cubic Inch Displacement by Registration State and Vehicle Size.” We assume the number of vehicles in sample is coded as “S” for not meeting publication standards.

²¹⁵As with gasoline externalities, we convert driving externalities to dollars per gallon using the average VMT-weighted fuel economy in the fleet.

²¹⁶We do not vary across time the effective corporate tax rate gasoline producers face.

Total WTP Summing across components, a \$1 change in the diesel tax rate results in a total WTP of \$0.827 in 2020 when using a price elasticity of -0.07. Consumers (\$1) and producers (\$0.015) are both willing to pay to avoid the tax increase, while society (-\$0.19) is willing to pay to keep the tax increase. We sign each component depending on the group’s willingness to pay to remove the tax. Consumers and producers are both willing to pay to remove the tax since these groups are made worse off through higher prices and reduced profits, respectively. On the other hand, society is willing to pay to keep the tax on the books, as they are made better off through reduced environmental externalities. Considering the removal of a tax (rather than the addition of one) also allows us to normalize both taxes and subsidies to a positive \$1 mechanical cost. Using the externality values calculated for 2006, we obtain a total WTP of \$0.859 (in 2006 dollars), with society willing to pay \$0.16 for increased pollution and producers \$0.015 for increased profits.

Cost A \$1 increase in the diesel tax rate mechanically raises \$1 of revenue for the government. However, the accompanying decrease in diesel consumption reduces the amount of revenue the government collects by the size of the behavioral response times the tax collected per dollar of diesel spending. In 2020, the federal diesel tax rate was \$0.244 per gallon (FHWA 2021) while the average state tax on diesel (weighted by gross gallons of diesel taxed) was \$0.296 per gallon (FHWA 2020). Accounting for federal and state diesel taxes, the government collected \$0.212 per dollar spent on diesel in 2020. Multiplying by a price elasticity of -0.07, the government faced a \$0.0148 loss in revenue from decreased diesel consumption. In 2006, the federal diesel tax was still \$0.244 per gallon (FHWA 2021) while the average state tax on diesel was \$0.2047 per gallon (FHWA 2020), resulting in a fiscal externality of \$0.012 from lost diesel consumption.

Decreases in producer profits reduce government revenue in the form of lost corporate taxes (assuming a 21% tax rate). The pre-tax markup on diesel was \$0.692 per gallon, meaning the government collected \$0.057 in corporate tax revenue for each dollar spent on diesel in 2020, when the price was \$2.55 per gallon.²¹⁷ With our price elasticity of -0.07, we calculate a \$0.004 fiscal externality from lost corporate tax revenue. Applying the same method to the in-context prices and markups (noted above) results in a \$0.004 fiscal externality from lost corporate tax revenue.

Finally, abating greenhouse gas emissions through a diesel tax raises revenue for the government in the long run. When calculating society’s WTP for global pollution benefits, we scale the WTP component by the share of global benefits that do not flow to the US government as revenue (98.1%). We put the remaining 1.92% in the denominator. With a price elasticity of -0.07, the total, unadjusted WTP in 2020 for *all* global benefits was \$0.060, implying the government generated \$0.0011 ($\0.060×0.0192) in revenue by abating carbon emissions today and promoting economic output tomorrow. Abating carbon emissions generates a positive fiscal externality, as the US government generates additional revenue through the policy change. In context, the climate FE equaled \$0.0006.

Summing the mechanical \$1 of revenue raised and the three fiscal externalities, we obtain a total “cost” of \$0.982 when using a price elasticity of -0.07: a \$1 increase in the diesel tax rate raises \$0.982 in revenue for the government. To match our approach with subsidies and other revenue raisers, we again consider the effects of removing the tax. This allows us to treat both taxes and subsidies as having a \$1 mechanical cost. The government loses \$1 in revenue

²¹⁷We assume all producer profits are subject to the American tax schedule but note that the geographic variation in crude oil production could render our calculation of this fiscal externality an overestimate.

by decreasing the tax rate by \$1, raises \$0.0188 in revenue from diesel taxes (\$0.0148) and corporate taxes (\$0.004) by encouraging diesel consumption, and loses \$0.0011 by increasing carbon emissions and decreasing long-run revenue, for a net government cost of \$0.982 in 2020. Using our 2006 specifications and applying the same calculations, we obtain a net government cost of \$0.984, with the government gaining \$0.016 in diesel tax and corporate tax revenue and losing \$0.0006 in long-run revenue due to increased carbon emissions.

MVPF Dividing the total WTP calculated above (\$0.827) by the total cost (\$0.982), both calculated with a price elasticity of -0.07, we form an MVPF of 0.842 in 2020. Using our in-context estimates of a total WTP of \$0.859 and a net government cost of \$0.984, we obtain an MVPF of 0.872.

E.11.3 Tax on Heavy Fuel Oil

Taxing heavy fuel oil (also referred to as residual fuel oil and bunker fuel) reduces the quantity of fuel consumed while generating revenue for the government. The MVPF for a tax on heavy fuel oil combines the price elasticity of heavy fuel oil for maritime vessels with a measure of the value of the externalities generated per dollar of spending on heavy fuel oil. We use a price elasticity of heavy fuel oil for maritime vessels from Mundaca et al. (2021), who focus on how vessels respond to higher fuel prices by reducing the weight of product shipped and the distance vessels travel, both of which relate linearly to the quantity of fuel consumed. Specifically, the authors rely on variation in the global average price of heavy fuel oil to estimate vessel's responsiveness to fuel prices along the intensive margin, reporting elasticities by the type of cargo.

We include this policy in our extended sample for two reasons. First, firms may respond to a tax on bunker fuel by not only reducing the weight and distance of shipments, but also by reducing the total number of shipments they charter. Additionally, the degree of leakage from this policy is unclear. If, for example, the United States imposed a tax on heavy fuel oil, vessels may respond by strategically refueling in countries without a tax. While a tax implemented globally (as discussed by the authors) would sidestep this concern, we focus on policies implemented by the United States. Out of concern for leakage, we use the authors' smallest reported elasticity (-0.032229), which corresponds to the price elasticity of ships carrying furniture. For reference, we report the MVPF at the end of this section if one were to use the author's largest reported elasticity (-0.416961), which applies to ships carrying fossil fuels. We do not consider nor discuss whether a reduction in the movement of polluting goods (such as fossil fuels) has added benefits to the environment. Both elasticities come from the authors' Table A.3 and correspond to the range of elasticities reported in the paper's abstract. We do not form a confidence interval for this MVPF since we include it in our extended sample.

As described in Appendix E.11.1, the location of local air pollution emissions matters when determine the social cost applied to these emissions; if the majority of emissions are released while ships are stationed in densely populated ports, the social cost we apply should be higher. The ambiguity about where ships release local pollutants further complicates our MVPF of a tax on heavy fuel oil. We discuss below how we attempt to address this issue.

Consumers' WTP The envelope theorem implies that consumers value the policy change at the value of the price increase. In other words, consumers are willing to pay \$1 for a \$1 change in their cost of heavy fuel oil, holding their consumption of heavy fuel oil constant.

Following our treatment of gasoline taxes, we assume the \$1 increase in the price of heavy fuel oil is completely passed onto consumers.

Society’s WTP In response to higher fuel prices, vessels reduce the number of miles traveled and the weight of their cargo, both of which reduce fuel consumption. This reduction in fuel usage benefits society through less global and local air pollution. We take emission rates (reported in grams per kWh and by engine tiers) for six pollutants (NO_X , $PM_{2.5}$, CO , CO_2 , SO_2 , and VOC) from ERG (2022) (author’s Table 5. C1C2).²¹⁸ Assuming a density of 905 kilograms per cubic meter (Live Bunkers 2024), or 3,425.8 grams per gallon, and a specific fuel consumption of 180.4 grams per kWh (Sustainable Ships 2024), we convert emission rates from grams per kWh to grams per gallon by multiplying each emission rate by 18.99 kWh per gallon (3,425.8 grams per gallon divided by 180.4 grams per kWh). We weight across engine tiers using the number of unique vessels belonging to each tier reported by ICCT (2023).

As noted above and discussed in Appendix E.11.1, assessing local damages from the maritime sector requires information on where air pollution is released, as estimates of the social costs of local air pollution (such as those from AP3) take as an input the location of emissions. For example, if ships generate large quantities of emissions while in port, and ports are located near large population centers, we would need to assign large social costs to the pollution from ships to account for the large number of people exposed to the pollution. Alternatively, if most local air pollution is released while ships are at sea, then fewer people would be exposed to the damages, meaning a smaller social cost should be applied. Despite this concern we note that incorporating benefits from abating local air pollution would make a tax on heavy fuel oil a more efficient way to raise revenue.

We also account for global and local emissions from heavy fuel oil production. We explain how we account for upstream emissions in Appendix C.4.3. These externalities are equivalent to the upstream externalities that enter our diesel and jet fuel MVPFs. They differ from our upstream gasoline externalities as we need not adjust for ethanol production. In 2020, upstream emissions generated an additional \$0.231 (in 2020 dollars) per gallon of petroleum product produced. In 2004, upstream emissions generated an additional \$0.122 (in 2004 dollars) per gallon of petroleum product produced.

As our baseline specification, we account for global damages released by burning heavy fuel oil and global and local damages from producing petroleum products. Pairing the damages from burning heavy fuel oil with upstream damages, we calculate a total WTP in 2020 of \$2.74 per gallon, with \$2.72 coming from global damages and \$0.02 from local damages. We adjust damages from greenhouse gases by the share of global damages that do not flow to the US government (0.981), resulting in a total externality of \$2.69 per gallon. Carbon dioxide released when burning heavy fuel oil imposes the largest externality (\$2.49 per gallon). In our in-context specification (2004), the total WTP was \$1.34 per gallon (in 2004 dollars), with global damages (adjusted for the share that does not flow to the US government) contributing \$1.31 and local damages the remaining \$0.03. We only adjust for the rising social costs of greenhouse gases over time. If we were to value the additional local damages reported by ERG (2022) using the baseline social cost estimates reported in the main text (which we use in our local damage extension discussed in Appendix E.11.1), the total per-gallon externality would rise to \$5.93 in 2020, with all of the increase coming from local damages (namely, NO_X and $PM_{2.5}$. In

²¹⁸Since CO_2 and SO_2 emissions follow from the fuel’s carbon content, these emissions do not differ by engine tier. We ignore reported PM_{10} emissions, as we do not have a social cost with which to value PM_{10} damages.

2004, the total externality rises to \$3.71 per gallon (in 2004 dollars) when incorporating local damages from burning heavy fuel oil. We note below what the MVPFs would be if one used these higher values.

Using price data from the EIA’s “U.S. Residual Fuel Oil Wholesale/Resale Price by Refiners (Dollars per Gallon)” (EIA 2022*b*), our externality calculations imply that society faced \$2.35 in total damages per dollar of spending on heavy fuel oil in 2020, when the price of heavy fuel oil was \$1.143 per gallon. In 2004, when the price of heavy fuel oil was \$0.681, society faced \$1.97 (in 2004 dollars) in total damages per dollar of spending on heavy fuel oil. Multiplying the externality generated per dollar of spending on heavy fuel oil by the own-price elasticity of heavy fuel oil gives us society’s WTP for a \$1 change in the tax on heavy fuel oil. With a price elasticity of -0.032 from Mundaca et al. (2021), society was willing to pay \$0.076 (-0.032 *times* \$2.35, in 2020 dollars) in 2020 for pollution abated due to reduced heavy fuel oil consumption. In 2004, society was willing to pay \$0.064 (-0.03 \times \$1.97, in 2004 dollars) for pollution abated due to reduced heavy fuel oil consumption. We do not consider any rebound effects when evaluating a tax on heavy fuel oil.

Producers’ WTP We account for producers’ WTP for lost profits resulting from reduced heavy fuel oil consumption. Specifically, since heavy fuel oil production involves the same three processes required to produce gasoline, we assume the percent markup imposed on gasoline by crude oil producers, refiners, and distributors holds for other fuels. We describe how we calculate this 27% markup in Appendix C.4.5. This percent markup is net of the assumed 8% economy-wide markup estimated by De Loecker et al. (2020).

Applying the 27% net markup to the price of heavy fuel oil in 2020 (\$1.143) implies a per-gallon markup of \$0.31. We adjust by the corporate average tax rate (21%) to account for the share of profits producers keep (Watson 2022).²¹⁹ This results in a post-tax externality borne by producers of \$0.214 per dollar of spending on heavy fuel oil. With a price elasticity of -0.032, producers were willing to pay \$0.007 in 2020 for the policy change. In 2004, when the percent markup was 31.4%, producers were willing to pay \$0.008 (in 2004 dollars), using the 2004 price of heavy fuel oil and holding the corporate tax rate fixed.

Total WTP Summing across components, a \$1 change in the heavy fuel oil tax rate results in a total WTP of \$0.931 in 2020 when using a price elasticity of heavy fuel oil of -0.032. Consumers (\$1) and producers (\$0.007) are both willing to pay to avoid the tax increase, while society (-\$0.076) is willing to pay to keep the tax increase. We sign each component depending on the group’s willingness to pay to remove the tax. Consumers and producers are both willing to pay to remove the tax since these groups are made worse off through higher prices and reduced profits, respectively. On the other hand, society is willing to pay to keep the tax on the books, as they are made better off through reduced environmental externalities. Considering the removal of a tax (rather than the addition of one) also allows us to normalize both taxes and subsidies to a positive \$1 mechanical cost. Using the externality values calculated for 2004, we obtain a total WTP of \$0.944 (in 2004 dollars), with society willing to pay \$0.064 for increased pollution and producers \$0.008 for increased profits.

²¹⁹We do not vary across time the effective corporate tax rate gasoline producers face.

Cost A \$1 increase in the heavy fuel tax rate mechanically raises \$1 of revenue for the government. While fuel consumed by maritime vessels is taxed under certain circumstances (IRS 2023), we simplify our MVPF by assuming no preexisting tax on heavy fuel oil purchased by maritime vessels exists, meaning we have no baseline fiscal externality from changes in heavy fuel oil consumption. However, we account for changes in corporate taxes paid and changes to the government budget from effects on the climate.

Decreases in producer profits reduce government revenue in the form of lost corporate taxes (assuming a 21% tax rate). The pre-tax markup on heavy fuel oil was \$0.31 per gallon, meaning the government collected \$0.057 in corporate tax revenue for each dollar spent on heavy fuel oil in 2020, when the price was \$1.143 per gallon.²²⁰ With our price elasticity of -0.032, we calculate a \$0.0018 fiscal externality from lost corporate tax revenue.

Abating greenhouse gas emissions through a heavy fuel oil tax raises revenue for the government in the long run. When calculating society's WTP for global pollution benefits, we scale the WTP component by the share of global benefits that do not flow to the US government as revenue (98.1%). We put the remaining 1.92% in the denominator. With a price elasticity of -0.032, the total, unadjusted WTP in 2020 for *all* global benefits was \$0.077, implying the government generated \$0.0015 ($\0.077×0.0192) in revenue by abating carbon emissions today and promoting economic output tomorrow. Abating carbon emissions generates a positive fiscal externality, as the US government generates additional revenue through the policy change.

Summing the mechanical \$1 of revenue raised and the three fiscal externalities, we obtain a total "cost" of \$0.9996 when using a price elasticity of -0.032: a \$1 increase in the gas tax rate raises \$0.9996 in revenue for the government. To match our approach with subsidies and other revenue raisers, we again consider the effects of removing the tax. This allows us to treat both taxes and subsidies as having a \$1 mechanical cost. The government loses \$1 in revenue by decreasing the tax rate by \$1, raises \$0.009 from corporate taxes by encouraging heavy fuel oil consumption, and loses \$0.0015 by increasing carbon emissions and decreasing long-run revenue, for a net government cost of \$0.9996 in 2020. Using our 2004 specifications and applying the same calculations, we obtain a net government cost of \$0.9991, with the government gaining \$0.0021 in corporate tax revenue and losing \$0.0012 in long-run revenue due to increased carbon emissions.

MVPF Dividing the total WTP calculated above (\$0.931) by the total cost (\$0.9996), both calculated with a price elasticity of -0.032, we form an MVPF of 0.931 in 2020. Using our in-context estimates of a total WTP of \$0.944 and a net government cost of \$0.9991, we obtain an MVPF of 0.945.

If we were to use our estimates of society's WTP that include increased local pollution from marine vessels, we obtain an MVPF of 0.84 in 2020 ($\$0.840/\0.9996) and 0.83 in 2004 ($\$0.832/\0.9991). If we were to use the largest elasticity estimated by the authors (-0.416961) and our baseline externalities, we obtain an MVPF of 0.109 in 2020 ($\$0.109/\0.9953) and 0.285 in 2004 ($\$0.281/\0.9882). If we were to implement both changes, we obtain an MVPF of -1.08 in 2020 ($-\$1.07/\0.9953) and -1.18 in 2004 ($-\$1.17/\0.9882).

²²⁰We assume all producer profits are subject to the American tax schedule but note that the geographic variation in crude oil production could render our calculation of this fiscal externality an overestimate.

E.11.4 Windfall Profit Tax on Crude Oil

Taxing the extraction of crude oil by US firms reduces the quantity of domestic crude oil produced. However, whether this change has environmental benefits depends on the slope of the global crude oil supply curve. We assume that the long-run global supply of crude oil is perfectly elastic, meaning a tax on US crude oil production would not affect the total quantity of crude oil consumed but would change the location of crude oil extraction.

One approach to taxing crude oil production involves taxing the profits producers earn from crude extraction. The Crude Oil Windfall Profit Tax Act of 1980 (WPT) imposed a tax on most domestic crude oil wells and subjected both less costly and higher value wells to higher tax rates. Rao (2018) uses variation in the after-tax price of crude oil introduced by the 1980 WPT to estimate how the policy affected the quantity of crude oil extracted, using a sample that includes monthly production data for wells in California. The author finds an after-tax price elasticity of oil production of 0.295 (s.e. 0.038, author’s Table 4, column 3, row 4). We do not form a confidence interval for this MVPF since we include it in our extended sample.

We follow the author’s approach to calculating the after-tax price of crude oil (see Rao (2018), page 274).

$$ATP = (1 - \tau^{Corporate})(P - \tau^{WPT}(P - B)) \quad (83)$$

where ATP is the after-tax price, P is the real selling price of crude oil, B is a statutory base price of crude oil against which the selling price is measured, $\tau^{Corporate}$ is the prevailing corporate tax rate, and τ^{WPT} is the WPT rate.²²¹ In our in-context specification (which is set in 1985), we use the author’s reported average after-tax price of \$18.30 (author’s Table 3). We solve for the real base price of \$7.14 using the reported average real selling price (\$41, author’s Table 3), the WTP rate (0.21, author’s Table 3), and the corporate tax rate (0.46, author’s replication package). We assume all values reported in the author’s Table 3 are in 1985 dollars (the final year of the author’s sample). In 2020, we calculate the average after-tax price by adjusting the real base price for inflation (\$17.17 in 2020 dollars) and updating the real selling price to be the 2020 refiner acquisition cost (\$40 in 2020 dollars), which measures the price at which crude oil suppliers sell oil to refiners (EIA 2024e).²²² We also substitute the author’s 46% corporate tax rate for the 21% corporate tax rate faced by petroleum producers in 2020 (Watson 2022), and we set the WPT rate to 0% in 2020 since the initial tax expired in 1988 (Lazzari 1990). Combining these pieces, the after-tax price would be \$31.60 in 2020.

We evaluate a one percentage point increase in the WPT rate. In 1985 (our in-context specification), we envision increasing the WPT rate from 21% to 22%. In 2020, we envision implementing a 1% tax on windfall profits accrued by domestic crude producers.

Producers’ WTP The envelope theorem implies that producers value the policy change at the value of the price increase. In other words, producers are willing to pay \$1 for a \$1 change in the after-tax price, holding production of crude constant. However, since we consider a one percentage point increase in the WPT rate—rather than a \$1 increase taxes paid—producers

²²¹The author notes that a different formula is applied should the real selling price fall below the statutory base price. We do not encounter this in our applications and therefore ignore this feature of the policy.

²²²This price closely tracks crude oil spot prices, such as those reported in the EIA’s “Cushing, OK WTI Spot Price FOB (Dollars per Barrel)” series (EIA 2024a).

subject to the tax have a willingness to pay equal to the change in the after-tax price.

We calculate the change in the after-tax price by taking the difference between the after-tax price calculated using the new WPT rate (22% in 1985 and 1% in 2020) and the after-tax price calculated with the real WPT rate (21% in 1985 and 0% in 2020). In 1985, increasing the WPT rate from 21% to 22% decreased the after-tax price by \$0.183. In 2020, increasing the WPT rate from 0% to 1% decreased the after-tax price by \$0.180.

Society’s WTP Since we assume a perfectly elastic long-run global supply curve for crude oil, the price of crude oil is fixed in the long-run. Without shifting the demand for crude oil, there will not be any change in the quantity of crude oil consumed. However, because the windfall profit tax reduces domestic crude oil production while not affecting the global supply of crude oil, the location of crude oil production must adjust.

The environmental externality from relocating crude oil production requires estimates of the carbon intensity of crude oil production by country.²²³ We use estimates from Masnadi et al. (2018), who measure the carbon intensity (grams of CO_2e generated per MJ of crude oil produced) of crude oil production by country. The carbon intensity of crude oil includes emissions released during exploration, extraction, processing, and transporting crude, or all “well-to-refinery” emissions. We abstract from changes in emissions resulting from differences in the cleanliness of refining.

Masnadi et al. (2018) find a global volume-weighted carbon intensity of 10.3 grams of CO_2e per MJ of crude produced, and a US-specific carbon intensity of 11.3 grams of CO_2e per MJ. Shifting crude oil production from the US to outside the US therefore generates environmental benefits equal to 1 gram of CO_2e per MJ of crude produced. One barrel of crude oil contains 6,119 megajoules (DOE 2020), so shifting one barrel of crude produced from the US to the rest of the world abates 6,119 grams of CO_2e . We assume well-to-refinery emissions have remained constant over time. CO_2 and CH_4 make up 65% and 34% of total emissions, respectively, with VOC and N_2O making up the remaining one percent.²²⁴ We then divide the share of total CO_2e attributable to CH_4 and N_2O by the GWP factors used by the authors to convert grams of a non- CO_2 pollutant to grams of CO_2e . This gives us grams of CO_2 , CH_4 , and N_2O released during the well-to-refinery process.

We apply each pollutant’s respective social cost to value well-to-refinery emissions in dollars per barrel of crude oil shifted overseas.²²⁵ We linearly extrapolate to obtain social costs for years before 2020, setting the social cost of a given pollutant equal to \$0 if the extrapolation yields a negative value. In 1985, the social cost of carbon was \$24.61, the social cost of methane was \$0, and the social cost of nitrous oxide was \$3,602.78, all expressed in 1985 dollars. In 2020, we use our baseline social costs of \$193 for carbon dioxide, \$1,648 for methane, and \$54,139 for nitrous oxide, all expressed in 2020 dollars.

Combining the change in the quantity of greenhouse gases produced with our social costs,

²²³Since our estimates of the social costs of local air pollutants are specific to the United States, we abstract from changes in local pollution but note that our approach could be applied to changes in local pollution given country-specific social costs and country-specific estimates of air pollution released from crude production.

²²⁴We assume N_2O and VOC each make up half of the remaining percent of the pollution. Since we calculate global damages from VOC using the same GWP factor as the authors we leave this pollutant in terms of CO_2e .

²²⁵Since the social cost of non- CO_2 greenhouse gases are roughly equal to the social cost of carbon scaled by the pollutant’s GWP factor, this approach generates approximately the same results if we were to apply our preferred social cost of carbon to the grams of CO_2e estimate.

we find a \$0.099 reduction in global damages per barrel of crude oil relocated from the US to the rest of the world in 1985. In 2020, each relocated barrel of crude oil generated \$0.894 in global benefits, in 2020 dollars. We adjust each by the share of the social cost of carbon that does not flow to the US government (0.981), which results in an adjusted externality of \$0.0972 in 1985 and \$0.877 in 2020, both expressed in nominal dollars. With an after-tax nominal price of \$18.30 and \$31.60 in 1985 and 2020, respectively, we obtain an externality of \$0.0053 per after-tax dollar generated in 1985 from crude oil production and \$0.0278 per after-tax dollar generated in 2020 from crude oil production.

In 1985, the \$0.183 decrease in the after-tax price is equivalent to a 1% reduction in the after-tax price. We multiply this observed percent change in the after-tax price by the after-tax price elasticity of oil production of 0.295 from Rao (2018) to obtain the percent change in crude oil supplied, or -0.295%. In 2020, we find a percent change in the after-tax price of -0.168% by applying the same approach. Combining the behavioral response to a one percentage point change in the WTP rate and the per-dollar externality estimated above gives us society's WTP for environmental benefits induced by the policy change. In 1985, society was willing to pay \$0.000016 ($0.00295 \times \0.0053) in nominal dollars for a 1 percentage point increase in the WTP rate. In 2020, society was willing to pay \$0.000047 ($0.00168 \times \0.0278) in nominal dollars for a 1 percentage point increase in the WTP rate.

Total WTP Summing across components, a one percentage point change in the WTP rate results in a total WTP of \$0.1803 in 2020 when using an after-tax elasticity of 0.295. Producers (\$0.1804) are willing to pay to avoid the tax increase, while society (-\$0.00005) is willing to pay to keep the tax increase. We sign each component depending on the group's willingness to pay to remove the tax. Producers are willing to pay to remove the tax since they are made worse off through lower after-tax prices. On the other hand, society is willing to pay to keep the tax on the books, as they are made better off through reduced environmental externalities. Considering the removal of a tax (rather than the addition of one) also allows us to normalize both taxes and subsidies to a positive mechanical cost. Using the externality values calculated for 1985, we obtain a total WTP of \$0.1828 (in 1985 dollars), with society willing to pay -\$0.00002 for increased pollution and producers \$0.1829 for the change in the after-tax price.

Cost Increasing the WTP rate by one percentage point mechanically changes how much producers pay in taxes. The change in the amount of taxes paid equals the change in the after-tax price, meaning the mechanical change in government revenue equals producers' WTP for the policy change. In 1985, increasing the WTP rate from 21% to 22% decreased the after-tax price by \$0.183. In 2020, increasing the WTP rate from 0% to 1% decreased the after-tax price by \$0.180.

The induced decrease in crude oil production reduces government revenue collected in the form of windfall profit taxes and corporate taxes. In 1985, prior to the policy change, the government collected \$22.70 per barrel of crude oil produced: the difference between the selling price (\$41) and the after-tax price (\$18.30) implies that the government collects \$22.70 per barrel of crude oil sold.²²⁶ With an after-tax price of \$18.30, the government collected \$1.24

²²⁶We can also calculate the tax revenue paid as

$$Revenue = \tau^{Corporate}(P - \tau^{WPT}(P - B)) + \tau^{WPT}(P - B) \quad (84)$$

where the first term isolates revenue collected as corporate taxes and the second identifies revenue collected

per after-tax dollar generated from extracting crude oil. We multiply the tax collected per dollar generated by the percent change in crude oil production that arises from a one percentage point increase in the WPT rate. The \$0.183 decrease in the after-tax price is equivalent to a 1% reduction in the after-tax price. We then multiply this observed percent change in the after-tax price by the after-tax price elasticity of oil production of 0.295 from Rao (2018) to obtain the percent change in crude oil supplied, or -0.295%. Combining the percent change in crude oil supplied and the revenue collected per dollar generated from crude oil produced provides the change in government revenue arising from changes in crude oil production: \$0.004 in 1985.

Applying the same approach to our 2020 components, we obtain a fiscal externality of \$0.0004 (in 2020 dollars). The government collected \$0.266 per dollar generated from crude oil production, and raising the WPT rate from 0% to 1% changes the after-tax price by only 0.6%. Holding the after-tax supply elasticity fixed over time, the government loses \$0.0004 due to the decline in crude oil production.

Abating greenhouse gas emissions through a tax on crude oil production raises revenue for the government in the long run. When calculating society's WTP for global pollution benefits, we scale the WTP component by the share of global benefits that do not flow to the US government as revenue (98.1%). We put the remaining 1.92% in the denominator. With an after-tax elasticity of 0.295, the total, unadjusted WTP in 1985 for *all* global benefits was \$0.00001596, implying the government generated \$0.00000031 ($\0.00001596×0.0192) in nominal dollars in revenue by abating carbon emissions today and promoting economic output tomorrow. Abating carbon emissions generates a positive fiscal externality, as the US government generates additional revenue through the policy change. Applying the same approach to our 2020 values results in a fiscal externality from abated greenhouse gases of \$0.000000913 (in nominal dollars).

Summing the cost components, we obtain a total "cost" of \$0.9996: in 1985, a one percentage point increase in the WPT rate raises \$0.1792 in revenue for the government. To match our approach with subsidies and other revenue raisers, we again consider the effects of removing the tax. This allows us to treat both taxes and subsidies as having a \$1 mechanical cost. The government loses \$0.1829 in revenue by removing the higher tax, raises \$0.0037 by encouraging crude oil production, and loses \$0.00000031 by increasing carbon emissions and decreasing long-run revenue, for a net government cost of \$0.1792 in 1985. Using our 2020 specifications and applying the same calculations, we obtain a net government cost of \$0.1799, with the government gaining \$0.00045 from increased crude production and losing \$0.000000913 in long-run revenue due to increased carbon emissions.

MVPF Dividing the total WTP calculated above (\$0.1804) by the total cost (\$0.1799), both calculated with a price elasticity of 0.295, we form an MVPF of 1.002 in 2020. Using our in-context estimates of a total WTP of \$0.1828 and a net government cost of \$0.1792, we obtain an MVPF of 1.020.

through the windfall profit tax. This approach is equivalent to the method described above but allows one to decompose revenue by source.

E.11.5 State-level Crude Oil Taxes

Taxing the extraction of crude oil by US firms reduces the quantity of domestic crude oil produced. However, whether this change has environmental benefits depends on the slope of the global crude oil supply curve. We assume that the long-run global supply of crude oil is perfectly elastic, meaning a tax on US crude oil production would not affect the total quantity of crude oil consumed but would change the location of crude oil extraction.

One approach to taxing crude oil production involves imposing a severance tax on crude oil extraction, often expressed in dollars per barrel of crude oil drilled. Brown et al. (2020) use cross-state variation in severance taxes and changes in the price of oil over time to estimate the responsiveness of crude producers to taxes on oil drilling. The authors find an elasticity of crude oil extraction with respect to the severance tax rate of -0.315 (s.e. 0.124, authors' table 3, column 4). We do not form a confidence interval for this MVPF since we include it in our extended sample.

We evaluate a \$1 increase in the severance tax on crude oil in 2015 (the final year of the authors' sample) and 2020. For our in-context specification, we use the baseline severance tax rate of \$3.266 (in 2012 dollars) reported by the paper (authors' Table 1) and convert to 2015 dollars. We use the same severance tax rate when evaluating the policy in 2020 but again adjust for inflation.

Producers' WTP The envelope theorem implies that producers value the policy change at the value of the tax increase. In other words, producers are willing to pay \$1 for a \$1 change in tax paid on extracted crude, holding their production of crude oil fixed. We assume domestic crude producers bear the entire tax burden since the global price of oil is fixed in the long-run as a result of the perfectly elastic global supply curve.

Society's WTP Since we assume a perfectly elastic long-run global supply curve for crude oil, the price of crude oil is fixed in the long-run. Without shifting the demand for crude oil, there will not be any change in the quantity of crude oil consumed. However, because the windfall profit tax reduces domestic crude oil production while not affecting the global supply of crude oil, the location of crude oil production must adjust.

The environmental externality from relocating crude oil production requires estimates of the carbon intensity of crude oil production by country.²²⁷ We use estimates from Masnadi et al. (2018), who measure the carbon intensity (grams of CO_2e generated per MJ of crude oil produced) of crude oil production by country. The carbon intensity of crude oil includes emissions released during exploration, extraction, processing, and transporting crude, or all "well-to-refinery" emissions. We abstract from changes in emissions resulting from differences in the cleanliness of refining.

Masnadi et al. (2018) find a global volume-weighted carbon intensity of 10.3 grams of CO_2e per MJ of crude produced, and a US-specific carbon intensity of 11.3 grams of CO_2e per MJ. Shifting crude oil production from the US to outside the US therefore generates environmental benefits equal to 1 gram of CO_2e per MJ of crude produced. One barrel of crude oil contains

²²⁷Since our estimates of the social costs of local air pollutants are specific to the United States, we abstract from changes in local pollution but note that our approach could be applied to changes in local pollution given country-specific social costs and country-specific estimates of air pollution released from crude production.

6,119 megajoules (DOE 2020), so shifting one barrel of crude produced from the US to the rest of the world abates 6,119 grams of CO_2e . We assume well-to-refinery emissions have remained constant over time. CO_2 and CH_4 make up 65% and 34% of total emissions, respectively, with VOC and N_2O making up the remaining one percent.²²⁸ We then divide the share of total CO_2e attributable to CH_4 and N_2O by the GWP factors used by the authors to convert grams of a non- CO_2 pollutant to grams of CO_2e . This gives us grams of CO_2 , CH_4 , and N_2O released during the well-to-refinery process.

We apply each pollutant's respective social cost to value well-to-refinery emissions in dollars per barrel of crude oil shifted overseas.²²⁹ We linearly extrapolate to obtain social costs for years before 2020, setting the social cost of a given pollutant equal to \$0 if the extrapolation yields a negative value. In 2015, the social cost of carbon was \$158.55, the social cost of methane was \$1,050.51, and the social cost of nitrous oxide was \$43,262.49, all expressed in 2015 dollars. In 2020, we use our baseline social costs of \$193 for carbon dioxide, \$1,648 for methane, and \$54,139 for nitrous oxide, all expressed in 2020 dollars.

Combining the change in the quantity of greenhouse gases produced with our social costs, we find a \$0.713 reduction in global damages per barrel of crude oil relocated from the US to the rest of the world in 2015. In 2020, each relocated barrel of crude oil generated \$0.894 in global benefits, in 2020 dollars. We adjust each by the share of the social cost of carbon that does not flow to the US government (0.981), which results in an adjusted externality of \$0.70 in 2015 and \$0.877 in 2020, both expressed in nominal dollars. With a severance tax of \$3.37 per barrel in 2015 and \$3.68 per barrel in 2020 (both in nominal dollars), we obtain an externality of \$0.208 per dollar of severance tax levied in 2015 and \$0.238 per dollar of severance tax levied in 2020.

Multiplying by the elasticity of supply with respect to the tax rate (-0.315) from Brown et al. (2020) gives us society's WTP for the environmental benefits from a \$1 change in the severance tax rate. In 2015, society was willing to pay \$0.065 for the environmental benefits from a \$1 increase in the severance tax rate. In 2020, society was willing to pay \$0.075 for the environmental benefits from a \$1 increase in the severance tax rate.

Total WTP Summing across components, a \$1 change in the severance tax rate results in a total WTP of \$0.925 in 2020 when using the elasticity of production with respect to the tax rate of -0.315 from Brown et al. (2020). Producers (\$1) are willing to pay to avoid the tax increase, while society (-\$0.075) is willing to pay to keep the tax increase. We sign each component depending on the group's willingness to pay to remove the tax. Producers are willing to pay to remove the tax since they are made worse off by paying more in taxes on each barrel of crude produced. On the other hand, society is willing to pay to keep the tax on the books, as they are made better off through reduced environmental externalities. Considering the removal of a tax (rather than the addition of one) also allows us to normalize both taxes and subsidies to a positive mechanical cost. Using the externality values calculated for 2015, we obtain a total WTP of \$0.935 (in 2015 dollars), with society willing to pay -\$0.065 for increased pollution and producers \$1 for the change in the tax rate.

²²⁸We assume N_2O and VOC each make up half of the remaining percent of the pollution. Since we calculate global damages from VOC using the same GWP factors as the authors we leave this pollutant in terms of CO_2e .

²²⁹Since the social cost of non- CO_2 greenhouse gases are roughly equal to the social cost of carbon scaled by the pollutant's GWP factor, this approach generates approximately the same results if we were to apply our preferred social cost of carbon to the grams of CO_2e estimate.

Cost A \$1 increase in the severance tax rate mechanically raises \$1 of revenue for the government. However, the accompanying decrease in crude extraction reduces the amount of revenue the government collects. Since the elasticity of production is with respect to the severance tax, the fiscal externality from the decreased crude oil extraction induced by a \$1 increase in the severance tax rate equals the elasticity, or \$0.315.²³⁰ The fiscal externality from decreased crude extraction is the same in our 2020 and in-context specifications.

Although the envelope theorem implies we need not account for lost producer profits in the numerator of the MVPF, we account for the effects of lost profits on corporate tax revenue. As described in Appendix C.4.5, crude suppliers sell oil to refiners at a price (refiner acquisition cost) above the landed cost of producing a barrel of crude, both reported by the EIA (EIA 2024*g,e*). In 2020, moving one barrel of crude oil from well to refinery cost \$37.27 on average, while refiners purchased this barrel for, on average, \$40. We set the markup to \$0 if the difference between the landed cost and selling price of crude is negative.²³¹ In 2020, this approach yields a per-barrel markup of \$2.73, and in 2015, the markup was \$2.99 per barrel, both expressed in nominal dollars. The government then collects 21% of this markup in the form of corporate taxes (Watson 2022). In 2020, the government collected \$0.57 in corporate taxes per barrel of crude, or \$0.16 per dollar of severance tax levied (all in nominal dollars). In 2015, the government collected \$0.63 in corporate taxes per barrel of crude, or \$0.19 per dollar of severance tax levied (all in nominal dollars). Multiplying by the elasticity of production with respect to the tax rate, we estimate a fiscal externality from changes in corporate tax revenue of \$0.049 in 2020 and \$0.059 in 2015.

Abating greenhouse gas emissions through a tax on crude oil production raises revenue for the government in the long run. When calculating society’s WTP for global pollution benefits, we scale the WTP component by the share of global benefits that do not flow to the US government as revenue (98.1%). We put the remaining 1.92% in the denominator. With an elasticity of -0.315, the total, unadjusted WTP in 2020 for *all* global benefits was \$0.0765, implying the government generated \$0.0015 ($\0.0765×0.0192) in nominal dollars in revenue by abating carbon emissions today and promoting economic output tomorrow. Abating carbon emissions generates a positive fiscal externality, as the US government generates additional revenue through the policy change. Applying the same approach to our 2015 values results in a fiscal externality from abated greenhouse gases of \$0.0013 (in nominal dollars).

Summing the mechanical \$1 of revenue raised and the three fiscal externalities, we obtain a total “cost” of \$0.637 when using a price elasticity of -0.315: a \$1 increase in the severance tax rate raises \$0.637 in revenue for the government. To match our approach with subsidies and other revenue raisers, we again consider the effects of removing the tax. This allows us to treat both taxes and subsidies as having a \$1 mechanical cost. The government loses \$1 in revenue by decreasing the tax rate by \$1, raises \$0.364 in revenue from increased crude production (-\$0.315) and corporate taxes paid (-\$0.049), and loses \$0.0015 by increasing carbon emissions and decreasing long-run revenue, for a net government cost of \$0.637 in 2020. Using our 2015 specifications and applying the same calculations, we obtain a net government cost of \$0.628, with the government gaining \$0.3736 in severance tax and corporate tax revenue and losing \$0.00128 in long-run revenue due to increased carbon emissions.

²³⁰In other words, the government collects \$1 in severance tax revenue per dollar of severance tax levied. For a \$1 change in the tax rate, multiplying by the elasticity of crude oil extraction with respect to the severance tax rate gives us a fiscal externality equal to the elasticity.

²³¹No monthly data reported a negative markup in 2020, and negative markups appear intermittently after January 1983.

MVPF Dividing the total WTP calculated above (\$0.925) by the total cost (\$0.637), both calculated with an elasticity of -0.315, we form an MVPF of 1.451 in 2020. Using our in-context estimates of a total WTP of \$0.935 and a net government cost of \$0.628, we obtain an MVPF of 1.489.

E.11.6 Tax on E85 (Flex Fuel)

“Flex fuel,” or E85, is a blend of gasoline and ethanol that can be consumed by flexible fuel vehicles. It contains between 51% to 83% ethanol; we focus on E85 that contains the maximum possible share of ethanol (83% ethanol, 17% gasoline) (DOE 2023a). Flexible fuel vehicles can consume gasoline or E85 (DOE 2023a). We assume E85 and the average gallon of gasoline are substitutes, meaning raising the tax on E85 induces drivers to consume more gasoline.

We use an own-price elasticity of ethanol from Anderson (2012), who estimates the price elasticity using monthly variation in ethanol prices in Minnesota between 1997 and 2006. The author finds that a \$0.10 increase in the per-gallon price of E85 (relative to the price of gasoline) leads to a 16.22% decrease (s.e. 0.217, author’s Table 2) in the quantity of E85 demanded. With an average E85 retail price of \$2.37 in the author’s sample (for a percent change in the E85 price of 4.219%), we calculate an elasticity of -3.844 for E85. We assume changes in the price of E85 from a tax on ethanol do not affect the price of gasoline, meaning the change in the price of E85 relative to the price of gasoline is driven entirely by the change in the E85 price from the tax change.

Our in-context specification is set in 2006. We specifically consider a \$1 increase in the tax on E85.

Consumers’ WTP The envelope theorem implies that consumers value the policy change at the value of the price increase. In other words, consumers are willing to pay \$1 for a \$1 change in their cost of E85, holding their consumption of E85. Following our treatment of gasoline taxes, we assume the \$1 increase in the price of E85 is completely passed onto consumers.

Society’s WTP Switching from E85 to gasoline generates environmental benefits because E85 is cleaner on a per-gallon basis than the average gallon of gasoline. However, given E85’s lower energy content, we must account for the fact that drivers must consume more E85 than gasoline to meet a target mileage. We assume flex-fuel vehicles operating on E85 achieve 27% lower fuel economies (AFDC 2024d); if the flex-fuel vehicle needed 1 gallon of gasoline to achieve its target mileage, the same vehicle would need to consume 1.37 gallons of E85.²³² This adjustment implies that, in 2006, while a flex-fuel vehicle using the average gallon of gasoline received 19.96 miles per gallon, a vehicle operating on E85 received only 14.57 miles per gallon. In 2020, while the average gallon of gasoline gave flex-fuel vehicles 23.11 miles per gallon, operating on E85 let the same vehicles drive 16.87 miles per gallon. This assumes flex-fuel vehicles achieve the fleet-average fuel economy when operating on the average gallon of gasoline.

²³²The EPA’s reported penalty to fuel economy is calculated between flex-fuel vehicles operating on ethanol-free gasoline and those operating on E85 (see the source’s footnotes). As a result, this fuel economy penalty is an overestimate. However, given only around 4.9% of the average gallon of gasoline was ethanol (see Appendix C.4), comparing E85 to unblended gasoline is the closest available estimate reported by AFDC (2024d).

We use our average per-gallon gasoline externality to estimate the damages from consuming the average gallon of gasoline. These externalities account for the small share (4.9%) of ethanol in the average gallon of gasoline. Since we assume the flex-fuel vehicle will travel the same number of miles, we do not include any driving externalities (accidents, congestion, and $PM_{2.5}$ from tires and brakes). In 2006, burning one gallon of gasoline generated, on average, \$1.091 in global damages and \$0.404 in local damages, both in nominal dollars. In 2020, burning one gallon of gasoline generated, on average, \$1.891 in global damages and \$0.226 in local pollution damages, both in nominal dollars.

To estimate the externality from one gallon of E85, we extend the framework outlined in Appendix C.4. First, we assume ethanol is non-emissive, meaning burning one gallon of E85 will only generate carbon emissions from the share of volume that is gasoline. Since we focus on E85 (which is 83% ethanol), we multiply the damages from burning one gallon of pure gasoline (\$1.69 per gallon) by 0.17, resulting in an externality from on-road CO_2 of \$0.162 per gallon in 2006 and \$0.288 per gallon in 2020.²³³ We note below how relaxing this assumption affects our conclusions.

Second, all upstream emission rates are calculated per gallon of petroleum product. However, a gallon of gasoline purchased in the US is not made up of only gasoline. To account for the share of ethanol in gasoline, we scale down each upstream emission rate by the share of gasoline in E85 (17%). After scaling down upstream emissions to account for the share of gasoline not derived from petroleum, we add upstream emissions from ethanol production. We only consider greenhouse gas emissions from this process. We use estimates of the carbon intensity of ethanol production from Lee et al. (2021), who find a carbon intensity of 45 grams of CO_2e released upstream per MJ of ethanol produced in 2019 and a carbon intensity of 56 grams of CO_2e released upstream per MJ of ethanol produced in 2006 (see authors' Figure 4).²³⁴ We add to this value an estimate of the carbon intensity of land-use change associated with ethanol production (7.4 grams of CO_2e per MJ) also from Lee et al. (2021). We hold this value constant overtime. We multiply the combined carbon intensity of ethanol production by the share of ethanol in E85 (83%), and then by the social cost of carbon in a given year to monetize these damages. Increased emissions from ethanol production are added to the upstream CO_2 estimate we present in Appendix Table 12. After adjusting for the ethanol content of gasoline, upstream carbon dioxide emissions increase from \$0.183 to \$0.778 per gallon in 2020. In 2006, upstream CO_2 emissions increase from \$0.102 to \$0.525 per gallon.

Third, for NO_X , CO , and HC , we account for the fact that fuel containing ethanol burns differently than pure gasoline. To do so, we use emissions adjustment factors from Hubbard et al. (2014), who report emissions rates from vehicles using fuel containing different amounts of ethanol.²³⁵ The authors find that vehicles running on fuel with 80.1% ethanol emit 52.0% less

²³³As discussed in Appendix C.4, converting the EIA's reported carbon content of pure gasoline (8.78 kilograms per gallon) to tons per gallon and multiplying by our preferred social cost of carbon in 2020 (\$193) gives us the unadjusted CO_2 externality of \$1.69 per gallon. Repeating this calculation using the (nominal) social cost of carbon in 2006, \$108.22, gives us the externality in 2006.

²³⁴This estimate of the carbon intensity of ethanol includes emissions from activities such as increased farming, ethanol processing, and increased fertilizer and chemical usage. Lee et al. (2021) estimate carbon intensities (in grams of CO_2e per MJ) for 2005 through 2019. We assume ethanol production for years before 2005 had the same carbon intensity as estimated in 2005, and that years after 2019 had the same carbon intensity as estimated in 2019. We assume one gallon of pure ethanol contains approximately 89.2 MJ of energy (AFDC 2024d) when using the reported "higher heating value" and assuming there are 0.001055 MJ in a Btu (ENERGY STAR 2015).

²³⁵We do not adjust the emission rates for CH_4 or N_2O because estimates from Lee et al. (2021) include CH_4

NO_X (authors' Table S3), 2.79% more CO (authors' Table S2), and 57.5% less HC (authors' Table S1, referred to as "non-methane hydrocarbons (corr.*)" by the authors) relative to a vehicle running on fuel without ethanol. Multiplying these percent differences by the ratio of the observed share of ethanol in E85 (83%) to the share of ethanol used in these emissions tests (80.1%) allows us to account for differences in the ethanol content of the fuel used in the authors' tests and E85, assuming a linear relationship between ethanol content and emission rates. In 2006, per-gallon damages from NO_X rose from \$0.191 to \$0.095, from \$0.06 to \$0.064 for CO , and \$0.064 to \$0.028 for HC . In 2020, per-gallon damages from NO_X rose from \$0.076 to \$0.039, from \$0.050 to \$0.053 for CO , and \$0.039 to \$0.017 for HC . We do not consider differences in damages from SO_2 and $PM_{2.5}$ between gasoline and ethanol.

Making these adjustments to pollutant-specific emission rates results in a total per-gallon externality from E85 of \$0.967 in 2006 and \$1.255 in 2020, in nominal dollars. In 2006, one gallon of E85 generated \$0.710 per gallon in global damages and \$0.257 in local pollution damages. In 2020, one gallon of E85 generated \$1.095 per gallon in global damages and \$0.160 in local pollution damages. Since flex-fuel vehicles must consume more E85 to travel a specified number of miles, we multiply these externalities by the ratio of the vehicle's fuel economy when using the average gallon of gasoline to the vehicle's fuel economy when using E85, which is equivalent to dividing the E85 per-gallon externalities by one minus the fuel economy penalty (27%) from using E85. This results in total per-gallon externality from E85 of \$1.325 in 2006 and \$1.719 in 2020.

Taking the difference of the per-gallon externalities from the average gallon of gasoline and E85, we find that fueling a vehicle with E85 generates \$0.118 in global environmental benefits and \$0.052 in local pollution benefits in 2006, for a total environmental benefit of \$0.17 (in nominal dollars). With an E85 price of \$2.91 per gallon in 2006 and \$2.62 per gallon in 2020 (AFDC 2024a), each dollar spent on ethanol generated \$0.058 in benefits in 2006 and \$0.152 in benefits in 2020.²³⁶ With an elasticity of -3.844, society has a WTP of \$0.221 for a \$1 increase in the tax on E85 in 2006, and a WTP of \$0.572 for the policy change in 2020. We adjust society's WTP for global damages by the share of the social cost of carbon that does not flow to the US government (0.981), resulting in an adjusted WTP for global damages of \$0.153 in 2006 and \$0.562 in 2020. Society has a positive willingness to pay since taxing E85 causes consumers to substitute toward dirtier fuel.

Producers' WTP We assume producers impose the same percent markup on a gallon of E85 and the average gallon of gasoline. We explain our markup calculation in Appendix C.4.5. The percent markup was 27.67% and 26.93% in 2006 and 2020 respectively. Both reported markups are net of the 8% economy-wide markup reported by De Loecker et al. (2020). With an E85 price of \$2.91 per gallon and a gasoline price of \$2.62 per gallon in 2006, consumers spent \$3.98 to fuel their vehicle with E85 and \$2.91 to fuel their vehicle with the average gallon

and N_2O emissions from ethanol combustion. While we assume CO_2 from ethanol combustion is entirely offset, we cannot assume the same for CH_4 and N_2O . To avoid double counting damages from these two greenhouse gases, we do not adjust our emission rates for CH_4 and N_2O using adjustment coefficients from Hubbard et al. (2014). We scale down on-road CH_4 and N_2O emissions by the share of gasoline in E85. We cannot isolate CH_4 and N_2O emissions from Lee et al. (2021) and therefore leave these damages as part of our reported upstream CO_2 damages even though these emissions are released during on-road operation. We note that CH_4 and N_2O emissions from ethanol combustion are the smallest contributors to ethanol's life cycle carbon intensity estimated by Lee et al. (2021).

²³⁶We weight by monthly gas consumption to construct annual average gas prices. This approach resembles the approach described in Appendix E.10 but restricts our sample to months with E85 prices.

of gasoline, meaning producers collected \$1.10 in profits when consumers used E85 and \$0.72 when consumers used gasoline. This results in a loss of \$0.376 when consumers switch from E85 to gasoline, or \$0.129 per dollar spent on E85. In 2020, with an E85 price of \$2.62 per gallon and a gasoline price of \$2.27 per gallon, consumers spent \$3.59 to fuel their vehicle with E85 and \$2.27 to fuel their vehicle with the average gallon of gasoline, meaning producers collected \$0.97 in profits when consumers used E85 and \$0.61 when consumers used gasoline. This results in a loss of \$0.354 when consumers switch from E85 to gasoline, or \$0.135 per dollar spent on E85.

With an elasticity of -3.844, producers have a WTP of \$0.497 for a \$1 increase in the tax on E85 in 2006 and \$0.519 for the policy change in 2020. Producers have a positive willingness to pay for the tax since taxing E85 causes consumers to spend less on fuel, leading to less collected as profits. We adjust both by a profit tax rate of 21% (Watson 2022) for a WTP of \$0.393 in 2006 and \$0.411 in 2020.

Total WTP We sum consumers', producers', and society's WTP for the tax change to calculate the total WTP for the policy change. In 2006, we calculate a total WTP of \$1.614 (\$1 + \$0.221 + \$0.393). In 2020, we calculate a total WTP of \$1.982 (\$1 + \$0.572 + \$0.411). We sign each component depending on the group's willingness to pay to remove the tax. All parties have a positive willingness to pay since all are made worse off by the tax on E85.

Cost A \$1 increase in the severance tax rate mechanically raises \$1 of revenue for the government. However, the accompanying decrease in E85 consumption reduces the amount of revenue the government collects. Although certain states tax ethanol and gasoline differently, we simplify the MVPF by assuming that E85 and gasoline are taxed at the same rate, which results in a fiscal externality only from differences in the quantity of ethanol needed to drive a given distance versus the quantity of gasoline needed to drive the same distance. We note, however, that taxing ethanol at a lower rate would decrease the size of the fiscal externality and, at very low tax rates, could result in a positive fiscal externality. See Appendix E.10 for more information on gas tax rates. With a nominal tax rate of \$0.387 in 2006, the government generated \$0.14 more in tax revenue when consumers fueled their vehicle with E85 rather than gasoline, or \$0.049 per dollar spent on E85. In 2020, with a nominal tax rate of \$0.465, the government generated \$0.17 in tax revenue when consumers fueled their vehicle with E85 rather than gasoline, or \$0.066 per dollar spent on E85. With an elasticity of -3.844, a \$1 increase in the tax on E85 generated a fiscal externality of -\$0.189 in 2006 and -\$0.252 in 2020. The fiscal externality is negative since the government loses money by increasing the tax on ethanol.

Since producers lose revenue from taxing E85 at a higher rate, the government also loses corporate tax revenue. Before adjusting for corporate taxes, producers had a pre-tax WTP of \$0.497 in 2006 and \$0.519 in 2020. Multiplying by an effective tax rate of 21% (Watson 2022), we calculate a fiscal externality from lost corporate profit taxes of -\$0.104 in 2006 and -\$0.109 in 2020. The fiscal externality is negative since the government loses money by reducing producer profits.

Abating greenhouse gas emissions through a tax on E85 raises revenue for the government in the long run. When calculating society's WTP for global pollution benefits, we scale the WTP component by the share of global benefits that do not flow to the US government as revenue (98.1%). We put the remaining 1.92% in the denominator. The total, unadjusted WTP in 2006 for global benefits was \$0.156, implying the government generated \$0.0030 (in nominal

dollars) in revenue by abating carbon emissions today and promoting economic output tomorrow. Abating carbon emissions generates a positive fiscal externality, as the US government generates additional revenue through the policy change. Applying the same approach to our 2020 values results in a fiscal externality from abated greenhouse gases of \$0.01098 ($\0.573×0.981) in nominal dollars. The fiscal externality is positive since the government raises money by abating carbon emissions.

We sum the program cost and the fiscal externalities from the tax change to calculate the total cost of the policy change. In 2006, we calculate a total cost of \$0.709 ($\$1 + -\$0.189 + -\$0.104 + \0.0030). In 2020, we calculate a total WTP of \$0.650 ($\$1 + -\$0.252 + -\$0.109 + \0.01098). We sign each component depending on when removing the tax raises or loses revenue for the government.

MVPF Dividing the total WTP calculated above (\$1.982) by the total cost (\$0.650), we form an MVPF of 3.051 in 2020. Using our in-context estimates of a total WTP of \$1.614 and a net government cost of \$0.709, we obtain an MVPF of 2.276.

If one assumed that ethanol was emissive when consumed by the vehicle (e.g., that the biomass would have been grown anyway but is now being refined and burned in a vehicle's engine), we find an MVPF of 1.13 in 2020, with society's WTP for global pollution falling to $-\$0.714$ from \$0.562 and the climate FE falling to $-\$0.0139$ from \$0.01098. This difference arises from our assumption that the ethanol component would release 5,769.6949 grams of carbon per gallon of E85 if it were emissive (EIA 2014) and that upstream emissions would only come from the ethanol refining process (45% of GHG emissions released upstream, according to Lee et al. (2021)).²³⁷

E.12 Other Revenue Raisers

In this section, we describe the calculation of three MVPFs that do not explicitly fit into the other tax categories. Of the three MVPFs, two are for critical peak pricing policies, and one evaluates the California Alternate Rates for Energy (CARE) program.

E.12.1 Critical Peak Pricing

Fowlie et al. (2021) study the effects of time-varying pricing on residential electricity demand. The paper implements a field experiment in California through the Sacramento Municipal Utility District (SMUD) from 2011-2013. SMUD customers were randomized into a control group and two treatment groups. One treatment group was signed up for a time-varying pricing plan while the other was given the option to opt-in. 90% of customers that were signed up stuck with the program whereas only 20% of people chose to opt-in. The paper finds that

²³⁷This scenario assumes that all the emissions from growing the biomass and from land-use change would have been released even if the biomass was not used for ethanol. If, for example, the biomass would simply have been stored and never consumed if it were not used for ethanol, then refining the biomass into ethanol and consuming this fuel would generate emissions. We isolate the upstream emissions from ethanol refining by taking the reported net quantity of emissions from Lee et al. (2021) (45 grams per MJ in 2019, which we hold constant), adding back in the share of the carbon offsets generated (total of 57 grams per MJ), and applying the share (45%) of total upstream emissions attributable to ethanol refining reported in the text.

the roughly 70% of passive joiners who chose not to opt out but would not have opted-in still respond to electricity price increases. However, their response is half as large.

We estimate the MVPF of a tax that increases the price of electricity during peak periods when marginal costs are high. The assumptions made for our baseline MVPF are similar to those made in our MVPF for peak energy reports using estimates from Brandon et al. (2019). First, we assume that the marginal cost of the next unit of electricity generation during peak periods is \$1,000 MWh. This is near the highest marginal cost reported in the sample from Fowlie et al. (2021). We present robustness to a marginal cost of \$500 per MWh. Second, we assume that the marginal generation source during peak events has emissions levels on par with coal. While many peaker plants are natural gas, using the natural gas emissions rate may underestimate the emissions involved with starting up the plant during peak times. We use coal's emissions rates from EPA's eGRID.

We also present a version of the MVPF assuming that the peak pricing allows for one customer to essentially transfer their electricity consumption to another customer. This could happen in the event that the peak pricing helps delay a blackout. In this case, an electricity customer would value the marginal MWh at the "value of lost load" (VLL). We use a VLL of \$4,292 per MWh from Brown & Muehlenbachs (2024).

Fowlie et al. (2021) estimate the responsiveness of electricity consumption to price changes separately for the active and passive joiners. We create two MVPFs that correspond to the treatment effects for passive and active joiners. We do not estimate an in-context MVPF since we do not have a good estimate of how the marginal cost of generation during peak times varies by year and energy market.

Active Joiners

Using the treatment effect on active joiners from Fowlie et al. (2021), we estimate an MVPF of 0.459 [0.393,0.529]. The paper estimates that active joiners reduced their hourly consumption during peak periods by 0.658 kW off of a mean of 2.49 kW. The price increased by roughly 350% during this period. Therefore, the elasticity during the peak period is -0.076. The retail price per kWh in 2020 was \$0.13. Dividing the elasticity by the price gives us that a one-cent change in the peak price leads to a 0.57% change in consumption. Equivalently, a one-dollar change in the peak price leads to a 57% change in consumption.

Cost We imagine a policy that imposes a tax of one dollar during peak periods. This tax raises one dollar in revenue. There are also two fiscal externalities associated with the tax. Utility companies lose profit when they have to supply electricity during high marginal cost periods. Since the government collects profit tax revenue from utility companies, reducing electricity consumption during peak times increases government profit tax revenue. From Appendix C.2, we estimate that 28% of utilities are public. For the 72% that are private, we apply a 10% effective tax rate. We do not use our standard profit per kWh value because the marginal cost during peak times is much higher than that during normal times. To get the profit loss per additional kWh of consumption during peak times, we take the difference between the retail electricity price and the marginal cost. The retail cost per kWh is \$0.13, and the marginal cost per kWh is \$1, resulting in a loss of \$0.87. Therefore, the increase in government revenue per kWh reduced is given by $0.87 * 0.28 + 0.87 * 0.72 * 0.1$ and is \$0.306. We multiply this by the semi-elasticity of 0.57 to get the change in government revenue for a one-dollar tax. The resulting fiscal externality is \$0.18. If one were to use the \$500 marginal cost assumption, the fiscal externality would be \$0.07.

The second fiscal externality to consider is the climate fiscal externality from reduced electricity consumption. As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality is \$0.002, and the total cost is \$1.18. For the marginal cost of \$500, the total cost is \$1.08.

WTP Consumers are willing to pay one dollar to avoid a one-dollar tax. The climate benefits from the tax bring down the willingness to pay to avoid the tax. The monetized local and global environmental externality from one reduced kWh using coal's emissions factor is \$0.052 and \$0.190. Multiplying these values by the semi-elasticity results in a WTP of -\$0.11 and -\$0.03. There is no rebound effect on market prices and demand because of the assumption that we are on the inelastic portion of the electricity supply curve during peak periods.

Utilities are also made better off by the tax because consumers reduce consumption during high marginal cost periods. Following the approach used to calculate the fiscal externality, we take the net of tax profits (assuming a 10% tax rate) per kWh from the 72% of private utilities and multiply by the semi ($0.306 * 0.72 * 0.90 * 0.57$) to get a WTP of -\$0.32. The \$500 marginal cost version has a WTP of utility companies of -\$0.14. Summing these components, we estimate a total WTP of \$0.54. Dividing by the cost, the MVPF is 0.46 for the baseline and 0.67 using the \$500 marginal cost assumption.

Passive Joiners

Using the treatment effect on passive joiners from Fowlie et al. (2021), we estimate an MVPF of 0.780 [0.697,0.869]. The paper estimates that passive joiners reduced their hourly consumption during peak periods by 0.242 kW off of a mean of 2.49 kW. The price increased by roughly 350% during this period. Therefore, the elasticity during the peak period is -0.028. The retail price per kWh in 2020 was \$0.13. Dividing the elasticity by the price gives us that a one-cent change in the peak price leads to a 0.21% change in consumption. Equivalently, a one-dollar change in the peak price leads to a 21% change in consumption.

Cost We imagine a policy that imposes a tax of one dollar during peak periods. This tax raises one dollar in revenue. There are also two fiscal externalities associated with the tax. Utility companies lose profit when they have to supply electricity during high marginal cost periods. Since the government collects profit tax revenue from utility companies, reducing electricity consumption during peak times increases government profit tax revenue. From Appendix C.2, we estimate that 28% of utilities are public. For the 72% that are private, we apply a 10% effective tax rate. We do not use our standard profit per kWh value because the marginal cost during peak times is much higher than that during normal times. To get the profit loss per additional kWh of consumption during peak times, we take the difference between the retail electricity price and the marginal cost. The retail cost per kWh is \$0.13, and the marginal cost per kWh is \$1, resulting in a loss of \$0.87. Therefore, the increase in government revenue per kWh reduced is given by $0.87 * 0.28 + 0.87 * 0.72 * 0.1$ and is \$0.306. We multiply this by the semi-elasticity of 0.21 to get the change in government revenue for a one-dollar tax. The resulting fiscal externality is \$0.06. If one were to use the \$500 marginal cost assumption, the fiscal externality would be \$0.03.

The second fiscal externality to consider is the climate fiscal externality from reduced electricity consumption. As explained in Section 4, the climate fiscal externality is 1.9% of the global environmental benefits. This externality is essentially zero, and the total cost is \$1.07. For the marginal cost of \$500, the total cost is \$1.03.

WTP Consumers are willing to pay one dollar to avoid a one-dollar tax. The climate benefits

from the tax bring down the willingness to pay to avoid the tax. The monetized local and global environmental externality from one reduced kWh using coal's emissions factor is \$0.052 and \$0.190. Multiplying these values by the semi-elasticity results in a WTP of -\$0.01 and -\$0.04 for local and global pollutants, respectively. There is no rebound effect on market prices and demand because of the assumption that we are on the inelastic portion of the electricity supply curve during peak periods.

Utilities are also made better off by the tax because consumers reduce consumption during high marginal cost periods. Following the approach used to calculate the fiscal externality, we take the net of tax profits (assuming a 10% tax rate) per kWh from the 72% of private utilities and multiply by the semi ($0.306 * 0.72 * 0.90 * 0.21$) to get a WTP of -\$0.12. The \$500 marginal cost version has a WTP of utility companies of -\$0.05. Summing these components, we estimate a total WTP of \$0.83. Dividing by the cost, the MVPF is 0.78 for the baseline and 0.87 using the \$500 marginal cost assumption.

E.12.2 California Alternate Rates for Energy (CARE)

Our MVPF for CARE using estimates from Hahn & Metcalfe (2021) is 0.719 [0.562, 0.914] in 2020 and 0.763 in context. CARE provides a 20% subsidy for natural gas and a 30-35% subsidy for electricity for low-income residents in California. Hahn & Metcalfe (2021) use an encouragement design field experiment in which they randomize which households are encouraged to enroll in CARE. They exploit this differential take-up rate to estimate a price elasticity of demand for natural gas of -0.35. CARE, as a subsidy, increases GHG emissions because it lowers the price of energy. We treat CARE as a tax and imagine the MVPF of reducing one dollar of spending on CARE. Reducing CARE spending is costly for the households who receive it but leads to positive environmental benefits for society. Since Hahn & Metcalfe (2021) only estimates the elasticity for natural gas, our MVPF is focused on reducing CARE subsidies for natural gas and not for electricity.

The subsidy for natural gas is 20% of the marginal price of natural gas. California has a price schedule in which the average market price and the marginal price are different. The paper reports that the average price per therm is \$0.90 for non-CARE recipients and \$0.70 for CARE recipients. From this, we can infer that the subsidy corresponds to a 22% subsidy on the average market price ($0.20/0.90$). Since the elasticity is -0.35, the subsidy corresponds to a 7.73% change in demand.

For ease of interpretation, we report all components of the MVPF at the per therm level. Our baseline MVPF uses externality values for the US in 2020, and our in-context MVPF uses values for California in 2014.

Cost The cost consists of the sum of the mechanical revenue raised from reducing spending on CARE, administrative costs, and the fiscal externalities from changes in tax revenue and changes in greenhouse gas emissions.

The retail price per therm in the US in 2020 is \$1.08. Therefore, the CARE subsidy per therm is \$0.24 ($0.22 * 1.08$). Reducing spending on CARE by \$0.24 raises \$0.24 worth of revenue. In the case of CARE, administrative costs have a significant impact on welfare, as found in Hahn & Metcalfe (2021). Since it is likely that these costs scale with program size, we include them in our analysis. The paper reports that the administrative costs for the natural gas portion of CARE administered by SoCalGas in 2015 (their sample population) is seven million dollars. The total program cost for SoCalGas in 2015 was 109 million. Therefore, roughly 6%

of the cost is administrative. We scale the \$0.240 worth of revenue raised by the additional saved administrative cost. This increases the revenue raised per therm to \$0.256.

The decrease in the subsidy also induces decreased consumption of natural gas, which generates a fiscal externality from further reduced CARE subsidies. This externality is equal to the product of the subsidy per therm (\$0.24) and the change in demand calculated above (7.73%). The fiscal externality from decreased natural gas consumption is \$0.016 in 2020 and \$0.014 in context.

From Appendix C.3, we calculate that there is a \$0.075 profit tax fiscal externality per therm of natural gas in the baseline and \$0.0859 in California in 2014. This externality arises from a change in profit tax revenue when natural gas consumption decreases. Since the subsidy corresponds to a 7.73% change in demand, the fiscal externality per induced therm is \$0.006 in 2020 and \$0.007 in context.

The climate externality, explained in Section 4, corresponds to 1.9% of the global environmental externalities outlined below. The climate externality is \$0.002 in 2020 and \$0.001 in context. The corresponding total cost is \$0.280 in 2020 and \$0.235 in context.

WTP Consumers value the policy change at the value of the price increase. At a per therm level, consumers are willing to pay \$0.240 to avoid the government reducing subsidy spending by \$0.240.

In response to higher natural gas prices, consumers will decrease natural gas consumption and the corresponding emissions from natural gas. To calculate the environmental benefits, we multiply the monetized externality per therm of natural gas by the change in demand (7.73%) from the price change. The environmental benefit per therm of natural gas in the US in 2020 is \$1.025, and in 2014 is \$0.822. Taking out the 1.9% that flows to the Treasury through the climate fiscal externality, the resulting environmental externality is -\$0.078 in 2020 and -\$0.062 in context. This is negative to reflect the fact that society is willing to pay less to avoid this policy change since there are environmental benefits. The rebound effect for natural gas, explained in Appendix C.3, offsets 12% of these benefits. The rebound effect is \$0.009 in 2020 and \$0.007 in context.

From Appendix C.3, we calculate that there is a loss in natural gas profits of \$0.440 for each one therm of reduced natural gas consumption in 2020 and \$0.507 in context. Taking into account the rebound and the change in demand from the policy change, the willingness to pay to avoid the policy change for natural gas utilities is \$0.030 in 2020 and \$0.035 in context. Summing together all of the WTP components results in a total WTP of \$0.201 in 2020 and \$0.180 in context. The resulting MVPF is 0.719 in 2020 and 0.763 in context.

E.13 Cap and Trade

In this section, we describe how we form MVPFs for cap-and-trade auctions. Section 6 describes our generalized approach to evaluating cap-and-trade auctions using the MVPF framework. Here, we focus on specific inputs for each policy's MVPF. We begin each subsection with an overview of how we collect the necessary data and estimated effects of the policy before explaining how we calculate each component of the MVPF. The in-context MVPF for a cap-and-trade auction evaluates a one-unit change in all permits auctioned in the set of years over which the policy change is evaluated, weighting each year by the number of allowances sold. When components are reported in tables, we normalize by program cost, which, for cap-and-

trade policies, we treat as the change in revenue from the change in permit prices. This ensures that firms' WTP for the change in permit price (listed in the "transfer" column of MVPF tables) equals the reported program cost.

Our main sample includes MVPFs of two domestic cap-and-trade systems: the Regional Greenhouse Gas Initiative (RGGI) and the California Cap-and-Trade Program. Our extended sample includes two evaluations of the European Union's Emissions Trading System (ETS).

Here, we abstract from fiscal externalities induced by take-up of other pre-existing subsidies. Similarly, our analysis of other subsidies assumes that there is no quantity cap which binds on producers. Those interaction effects between can be quantitatively meaningful. Perino et al. (2023), for example, note the presence of "waterbed effects" where by changes in subsidies can induce little or no change in emissions quantities if caps bind. Those subsidies simply reduce the price of permits.

E.13.1 Regional Greenhouse Gas Initiative

The Regional Greenhouse Gas Initiative (RGGI) is a cap-and-trade system that covers CO_2 emissions from regulated power plants in 11 states in the Northeastern US. We collect data on RGGI allowances and quarterly auction clearing prices from RGGI (2024). When averaging across periods, we weigh by the total number of allowances sold each quarter. One allowance authorizes a firm to emit one short ton of CO_2 .

We use results from Chan & Morrow (2019) to estimate the marginal abatement cost curve for power plants covered by RGGI. The authors use a difference-in-difference approach that leverages differences in the locations of power plants to quantify RGGI's impact on tons of pollution emitted between 2009 and 2016. They estimate RGGI's impact on both tons of pollution emitted and damages generated. We focus solely on changes in quantities and apply our preferred social costs to harmonize with other policies in our sample, but note that focusing on changes in quantities overlooks the redistribution of pollution to areas with higher social costs considered by the authors.²³⁸ When discussing RGGI, we denote all quantities as short tons and adjust our social costs accordingly, as the EPA's Air Markets Program database (used by the authors) and RGGI's allowances are both expressed quantities in short tons.

Chan & Morrow (2019) report changes in CO_2 , SO_2 , and NO_X emissions from power plants in RGGI states (authors' Tables A.1, A.2, and A.3) between 2009 and 2016.²³⁹ For CO_2 and SO_2 , the authors isolate changes in emissions from coal versus combined cycle natural gas power plants located in RGGI states and in neighboring states (Pennsylvania and Ohio) whose power plants are not covered by RGGI (authors' Tables A.2 and A.3). We apply the reported change in emissions to the quantities of emissions released in 2008 by power plants in states that joined RGGI, which we take from the authors' replication package.

The authors' replication package reports that power plants in states that joined RGGI collectively emitted 85,043,072 short tons of CO_2 in 2008, with 49,345,056 tons coming from

²³⁸Even if RGGI induced a net reduction in tons of pollution emitted, these benefits could be partially (or entirely) offset if leaked emissions reemerge in locations where marginal damages are higher (e.g., in places with larger populations). We do not consider spatial differences in social costs in this paper; instead, we apply our national average social costs to the change in tons estimated by the authors.

²³⁹The authors' dependent variables are logged. We exponentiate the reported coefficients to find the quantity of emissions released in 2016 and calculate the change in emissions. We report the coefficients included in the authors' tables to allow readers to cross-walk our approach with the paper's estimates.

coal power plants and the remaining 35,698,016 from combined cycle natural gas power plants. Using the estimated change in CO_2 emissions from coal power plants in RGGI states (-0.79, s.e. 0.17 – Table A.2, Column 2) and natural gas power plants in RGGI states (0.13, s.e. 0.11 – Table A.2, Column 3), we calculate a change in CO_2 emissions from power plants covered by RGGI of -21,994,162 tons of CO_2 between 2009 and 2016.

We follow the same approach for SO_2 and NO_X emissions. In 2008, coal power plants in states that joined RGGI emitted 220,302.31 tons of SO_2 while combined cycle natural gas power plants emitted 1,110.33 tons. Using the estimated change in SO_2 emissions from coal power plants in RGGI states (-0.89, s.e. 0.39 – Table A.3, Column 2) and natural gas power plants in RGGI states (-0.15, s.e. 0.13 – Table A.3, Column 3), we calculate a change in SO_2 emissions from power plants covered by RGGI of -129,988.56 tons between 2009 and 2016. The authors do not report changes in NO_X emissions by energy source, so we apply the total change in NO_X emissions from power plants in RGGI states (-0.19, s.e. 0.14 – Table 3, Column 4) to the sum of NO_X emissions released in 2008 from coal (72,341.55) and natural gas power plants (9,309.23) in RGGI states. This implies a total change in tons of NO_X emitted of -14,128.92. Since RGGI allocates allowances per short ton of CO_2 , we pair the estimated change in CO_2 emissions with the estimated change in SO_2 and NO_X to calculate the co-benefits generated per ton of carbon abated. For every one ton of CO_2 abated, RGGI abated 0.006 (-129,988.56/-21,994,162) short tons of SO_2 and 0.0006 (-14,128.92/-21,994,162) short tons of NO_X .

While the reported changes in emissions in nearby non-RGGI states allow us to quantify leakage from RGGI, we use an alternative approach that assumes a constant leakage share, as the authors' results imply RGGI had positive spillovers on neighboring states by encouraging greater natural gas usage and decreased coal usage.²⁴⁰ We also follow this alternative approach since we do not consider shifting social costs for local pollutants. This alternative approach also allows us to consider NO_X leakage, as the authors' appendix tables only report changes in CO_2 and SO_2 in Pennsylvania and Ohio. Fell & Maniloff (2018) estimates that RGGI states decreased CO_2 annual emissions by 8.8 million tons while neighboring states increased annual emissions by 4.5 million tons, implying a 51.1% leakage rate. We assume this leakage rate applies to all pollutants and does not vary over time.

To calculate the slope of the marginal abatement cost curve, dp/dq , we pair the total reduction in short tons of CO_2 emitted between 2009 and 2016 ($dq = -21,994,162$ short tons) with the average clearing price per allowance from this time period ($dp = \$3.19$ in 2016 dollars). To calculate the average clearing price, we weigh the clearing price of each auction conducted between 2009 and 2016 by the number of allowances sold. Multiplying dp/dq by the total number of allowances sold between 2009 and 2016 (q) tells us the impact of auctioning one additional permit on the auction price; this term captures both firms' WTP for the policy change and the effect on government revenue from the change in permit price. 816,224,961 million allowances were sold between 2009 and 2016 (RGGI 2024). Multiplying 816.2 million by dp/dq (3.19/-21,994,162) indicates that auctioning one additional permit decreased spending on all permits in circulation by \$118.48 (in 2016 dollars).

Next we detail each component that goes into our MVPF calculation. We focus on our

²⁴⁰If we were to apply the same approach described above to the reported change in emissions in Pennsylvania and Ohio, we would find that these neighboring states emitted 6,101,105.3 and 76,776.39 fewer tons of CO_2 and SO_2 , respectively. This reduction in emissions in neighboring states is consistent with the increased natural gas usage in nearby states as observed by the authors. Following this approach would imply that RGGI generated co-benefits by encouraging a shift to natural gas energy production in states not covered by RGGI and would make auctioning one less permit an even more efficient way to raise revenue.

in-context MVPFs for cap-and-trade programs but include what the MVPF would be in 2020 if one assumed $q \times dp/dq$ remained constant over time, only adjusting dp for inflation.

Firms' WTP Firms covered by RGGI have a WTP for the increased spending on permits that arises from the change in the quantity of allowances auctioned. As described in Section 6, marginal firms that no longer obtain permits are indifferent to doing so by the envelope theorem. Firms' WTP for the change in the price of all permits in circulation equals $q \times dp/dq$, or \$118.48 (in 2016 dollars). In 2020, holding q and dq constant but adjusting dp from 2016 to 2020 dollars (\$3.44 in 2020 dollars), firms would be willing to pay \$127.78 (in 2020 dollars) for a one unit change in the number of allowances auctioned.

Society's WTP Issuing one fewer allowance abates one short ton of CO_2 before accounting for leakage. We use the average social cost of carbon between 2009 and 2016 (all expressed in 2016 dollars) when calculating the in-context MVPF, weighting by the total quantity of allowances sold each year. This yields an SCC of \$136.04 (in 2016 dollars) per short ton. We multiply this value by the share of the SCC that does not flow to the US government (0.981). Before accounting for leakage, society is willing to pay \$133.44 (in 2016 dollars) for abated CO_2 .

For each abated short ton of CO_2 , RGGI also abated 0.006 short tons of SO_2 and 0.0006 short tons of NO_X . Multiplying these estimates by the social cost of one short ton of pollution yields society's WTP for local pollution. The social costs of local pollutants do not rise over time, but we again convert damages from metric to short tons. Using the quantity of SO_2 and NO_X abated per short ton of CO_2 abated calculated above and social costs of \$39,106.08 per short ton of SO_2 and \$13,620.33 per short ton of NO_X (both in 2016 dollars), abating one short ton of CO_2 yields an additional \$239.87 in local benefits, with \$231.12 from abated SO_2 and \$8.75 from abated NO_X .

To account for leakage, we sum the total willingness to pay for pollution (\$373.31 = \$133.44 + \$239.87) and multiply by one minus the leakage rate (1 - 0.511). This yields a total WTP of \$182.41 (in 2016 dollars) after accounting for the 51.1% of emissions that reemerge outside of RGGI, assuming that all three pollutants have the same leakage rate. We do not consider any rebound effect when evaluating RGGI, as the total change in CO_2 emissions estimated between 2009 and 2016 should include changes in emissions that arise due to changing electricity prices.

For our 2020 specification, we augment our results by adjusting for the rising SCC per short ton. All other components remain the same. This adjustment implies a WTP of \$210.33 (in 2020 dollars) for abated pollution in 2020.

Total WTP The total willingness to pay for a one-unit change in allowances auctioned in RGGI is the sum of firms' WTP for the change in permit spending (\$118.48) and society's WTP for the net change in pollution (\$182.41). For consistency with other revenue raisers, we consider auctioning one additional allowance, as this results in a negative WTP for society (more CO_2 is emitted) and a positive WTP for firms (spending on permits decreases). Auctioning one additional permit results in a total WTP of -\$63.93 (in 2016 dollars). Total WTP is negative since society's WTP for abated pollution outweighs firms' WTP for higher permit prices.

In 2020, the total WTP for a one unit increase in allowances auctioned is -\$82.55, with firms willing to pay \$127.78 and society willing to pay -\$210.33.

Cost The denominator of RGGI’s MVPF contains three terms. First, the government generates revenue by selling one additional permit at a given price. In our in-context specification, we use the price used to calculate dp (\$3.19 in 2016 dollars) to quantify the revenue generated from selling one additional permit. In our 2020 specification, we use the average clearing price for all auctions held in 2020, weighted by the quantity of allowances sold (\$6.41 in 2020 dollars).

The change in permit prices also affects the government’s budget, as all permits in circulation are sold at a different price. When auctioning one more allowance, the government loses revenue from all permits selling at a lower price, or $q \times dp/dq$. This is equal to firms’ WTP for higher permit prices. In our in context specification, the government loses \$118.48 (in 2016 dollars) from lower permit prices. In 2020, the government loses \$127.78 (in 2020 dollars) from lower permit prices.

We also account for the fiscal externality on government revenue from reduced CO_2 emissions, or the share of global benefits that do not enter the numerator. In our in-context specification, auctioning one additional permit costs the government \$1.27 due to increased carbon emissions. In 2020, the climate fiscal externality costs the government \$1.64. This externality accounts for CO_2 leakage.

As with WTP, we consider auctioning one additional allowance, which lowers permit prices and results in greater CO_2 emissions. Auctioning one more allowance generates revenue since one extra permit is sold at a given price (-\$3.19 in 2016 dollars). However, this additional permit costs the government \$118.48 (in 2016 dollars) from lower permit prices and \$1.27 due to increased CO_2 emissions. This yields a net government cost of \$116.56 (in 2016 dollars) in context. In 2020, auctioning one more permit cost the government \$123.01 (in 2020 dollars).

MVPF Combining our total WTP and net government cost, we obtain an MVPF of -0.55 in context (-63.93/116.56). In our 2020 specification, we obtain an MVPF of -0.67, again assuming the inflation-adjusted $q \times dp/dq$ holds over time. The MVPF is negative since society’s WTP for changes in pollution outweighs firms’ WTP for the change in permit spending, and since the government loses revenue by auctioning one more permit.

E.13.2 California Cap-and-Trade Program

The California Cap-and-Trade Program covers emissions from around 450 firms that collectively produce approximately 85% of greenhouse emissions released in the state (C2ES 2024). We collect data on allowances sold and quarterly auction settlement prices from CARB (2024). When averaging across periods, we weight by the total number of allowances auctioned each quarter. As described below, since California shares an allowance auction with Québec, we isolate the number of allowances sold to Californian firms using data from CARB (2023). One allowances authorizes a firm to emit one metric ton of CO_2 . Unlike in our discussion of RGGI, all quantities are denoted in metric tons.

We take estimates of the effect of the cap-and-trade system on emissions from Hernandez-Cortes & Meng (2023), who use a differential trend-break model and the introduction of the cap-and-trade system to estimate the total change in CO_2 emissions, as well as changes in local air pollution emissions. The authors report the total change in tons of CO_2e , $PM_{2.5}$, PM_{10} , NO_X , and SO_X between 2012 and 2017 (authors’ Table 1).²⁴¹ The authors find that the

²⁴¹We ignore reductions in PM_{10} since AP3 does not calculate marginal damages for this pollutant. We apply

California Cap-and-Trade Program abated 3,200,000 tons of CO_2 emissions between 2012 and 2017. Since the cap-and-trade system allocates allowances per metric ton of CO_2 , we pair the estimated change in CO_2 emissions with the estimated change in $PM_{2.5}$, NO_X , and SO_X to calculate the co-benefits generated per ton of carbon abated. For every one ton of CO_2 abated, the California Cap-and-Trade program abated 0.00003 tons of $PM_{2.5}$, 0.0002 tons of NO_X , and 0.00002 tons of SO_2 . Because we rely on the total change in tons of pollution emitted (which do not have corresponding standard errors), we cannot form confidence intervals for this MVPF. We do not consider leakage from this cap-and-trade program.

To calculate the slope of the marginal abatement cost curve, dp/dq , we pair the total reduction in tons of CO_2 emitted between 2012 and 2017 ($dq = -3,200,000$ metric tons) with the average settlement price per allowance from this time period ($dp = \$13.24$ in 2017 dollars). To calculate the average settlement price, we weight the settlement price of each auction conducted between 2012 and 2017 by the number of allowances sold.²⁴² Multiplying dp/dq by the total number of allowances sold between 2012 and 2017 (q) tells us the impact of auctioning one additional permit on the auction price; this term captures both firms' WTP for the policy change and the effect on government revenue from the change in permit price. 709,985,605 allowances were sold between 2012 and 2017 (CARB 2023). While firms in both California and Québec faced the same settlement price, we focus solely on changes in emissions in California. We isolate the allowances held by Californian firms using the annual "Total Allocations" reported by CARB (2023).²⁴³ Multiplying 710 million by dp/dq ($13.24/-3,200,000$) indicates that auctioning one fewer permit increased the price of all permits in circulation by \$2,936.95 (in 2017 dollars).

Hernandez-Cortes & Meng (2023) focus on changes in emissions from firms responsible for only 5% of California's greenhouse gas emissions. As a conservative assumption, we assume the remaining 95% of firms did not change their behavior, meaning dq equals 3,200,000 tons of CO_2 abated. We also explore an alternative approach where we assume the unobserved firms responded identically to the firms included in the authors' sample, which implies a dq of 64,000,000. Under this assumption, we obtain a $q \times dp/dq$ of \$146.85 (in 2017 dollars).

Next we detail each component that enters into our MVPF calculation. We focus on the in-context MVPFs for cap and trade but include what the MVPF would be in 2020 if one assumed $q \times dp/dq$ remained constant over time, only adjusting dp for inflation.

Firms' WTP Firms covered by California's Cap-and-Trade Program have a WTP for the increased spending on permits that arises from the change in the quantity of allowances auctioned. As described in Section 6, marginal firms that no longer obtain permits are indifferent to doing so by the envelope theorem. Firms' WTP for the change in the price of all permits in circulation equals $q \times dp/dq$, or \$2,936.95 (in 2017 dollars). In 2020, holding q and dq constant but adjusting dp from 2017 to 2020 dollars (\$13.98 in 2020 dollars), firms would be willing to pay \$3,101.40 (in 2020 dollars) for a one unit change in the quantity of allowances auctioned.

the social cost of SO_2 to reported SO_X emissions.

²⁴²We use the total allowances sold to firms in both California and Québec when weighting across quarters.

²⁴³CARB (2023) reports total annual allocations for 2013 onward. Since only one auction was held in 2012 and Québec and California had not yet joined auctions, the total allocations for that year equals the allowances sold at that single auction.

Society’s WTP Issuing one fewer allowance abates one metric ton of CO_2 . We use the average social cost of carbon between 2012 and 2017 (all expressed in 2017 dollars) when calculating the in-context MVPF, weighting by the total quantity of allowances sold to Californian firms each year. This yields an SCC of \$165.65 (in 2017 dollars) per ton. We multiply this value by the share of the SCC that does not flow to the US government (0.981). Society is therefore willing to pay \$162.48 (in 2017 dollars) for abated CO_2 .

For each abated ton of CO_2 , California’s Cap-and-Trade Program also abated 0.00003 tons of $PM_{2.5}$, 0.0002 tons of NO_X , and 0.00002 tons of SO_2 . Multiplying these estimates by the social cost of each pollutant yields society’s WTP for local pollution. The social costs of local pollutants do not rise over time. Using the quantity of NO_X , $PM_{2.5}$, and SO_2 abated per ton of CO_2 abated calculated above and the social costs reported in text, abating one ton of CO_2 yields an additional \$4.69 (in 2017 dollars) in local benefits, with \$1.35 from $PM_{2.5}$, \$2.49 from NO_X , and \$0.85 from SO_2 .

For our 2020 specification, we augment our results by adjusting for the rising social cost of carbon (\$193 in 2020, in 2020 dollars). All other components remain the same but are adjusted for inflation. This adjustment implies a WTP of \$194.26 (in 2020 dollars) for abated pollution in 2020, with \$189.30 coming as global benefits and the remaining \$4.96 coming from local pollutants.

We do not consider leakage from the California Cap-and-Trade Program, nor do we consider any rebound effect, as the total change in CO_2 emissions estimated between 2012 and 2017 should include changes in emissions that arise due to changing electricity prices.

Total WTP The total willingness to pay for a one unit change in allowances auctioned in the California Cap-and-Trade Program is the sum of firms’ WTP for the change in permit spending (\$2,936.95) and society’s WTP for the net change in pollution (\$167.17). For consistency with other revenue raisers, we consider auctioning one additional allowance, as this results in a negative WTP for society (more CO_2 is emitted) and a positive WTP for firms (spending on permits decreases). Auctioning one additional permit results in a total WTP of \$2,769.78 (in 2017 dollars). Total WTP is positive since society’s WTP for abated pollution is less than firms’ WTP for higher permit prices.

In 2020, the total WTP for a one unit increase in allowances auctioned is \$2,907.14 (in 2020 dollars), with firms willing to pay \$3,101.40 and society willing to pay -\$194.26.

Cost The denominator of the California Cap-and-Trade Program’s MVPF contains three terms. First, the government generates revenue by selling one additional permit at a given price. In our in-context specification, we use the price used to calculate dp (\$13.24 in 2017 dollars) to quantify the revenue generated from selling one additional permit. In our 2020 specification, we use the average clearing price for all auctions held in 2020, weighted by the quantity of allowances sold (\$17.12 in 2020 dollars).

The change in permit prices also affects the government’s budget, as all permits in circulation are sold at a different price. When auctioning one more allowance, the government loses revenue from all permits selling at a lower price, or $q \times dp/dq$. This is equal to firms’ WTP for higher permit prices. In our in context specification, the government loses \$2,936.95 (in 2017 dollars) from lower permit prices. In 2020, the government loses \$3,101.40 (in 2020 dollars) from lower permit prices.

We also account for the fiscal externality on government revenue from reduced CO_2 emissions, or the share of global benefits that do not enter the numerator. In our in-context specification, auctioning one additional permit costs the government \$3.17 (in 2017 dollars) due to increased carbon emissions. In 2020, the climate fiscal externality costs the government \$3.70 (in 2020 dollars).

As with WTP, we consider auctioning one additional allowance, which lowers permit prices and results in greater CO_2 emissions. Auctioning one more allowance generates revenue since one extra permit is sold at a given price (-\$13.24 in 2017 dollars). However, this additional permit costs the government \$2,936.95 from lower permit prices and \$3.17 due to increased CO_2 emissions. This yields a net government cost of \$2,926.89 (in 2017 dollars) in context. In 2020, auctioning one more permit cost the government \$3,087.98 (in 2020 dollars).

MVPF Combining our total WTP and net government cost, we obtain an MVPF of 0.95 (2,769.78/2,926.89) in context. In our 2020 specification, we obtain an MVPF of 0.94, again assuming the inflation-adjusted $q \times dp/dq$ holds over time. The MVPF is positive since society’s WTP for changes in pollution is less than firms’ WTP for the change in permit spending, and since the government loses revenue by auctioning one more permit.

E.13.3 EU Emissions Trading System (Bayer & Aklin 2020)

The European Union’s Emissions Trading System is the largest and oldest cap-and-trade system that targets greenhouse gases, covering around 5% of global carbon emissions (Colmer et al. 2024). We collect data on allowances auctioned or sold to all stationary installations in a given year from EEA (2024).²⁴⁴ Annual settlement prices come from World Bank (2024), who report the annual settlement price (in US dollars per ton of CO_2e) as measured on or around April 1 of each year. One allowances authorizes a firm to emit one metric ton of CO_2 . Unlike in our discussion of RGGI, all quantities are denoted in metric tons.

We use estimates from Bayer & Aklin (2020), who rely on a generalized synthetic control approach to structurally estimate counterfactual CO_2 emissions from 25 EU member states. The authors estimate that ETS abated 1,219.5 million tons of CO_2 between 2008 and 2016 (authors’ Table S4). We do not form a confidence interval for this MVPF since we include it in our extended sample.

To calculate the slope of the marginal abatement cost curve, dp/dq , we pair the total reduction in tons of CO_2 emitted between 2008 and 2016 ($dq = -1,219.5$ million metric tons) with the average settlement price per allowance from this time period ($dp = \$7.69$ in 2016 dollars). To calculate the average settlement price, we weight the annual allowance price reported by World Bank (2024) by the total number of allowances auctioned or sold that year, as reported by EEA (2024). Multiplying dp/dq by the total number of allowances sold between 2008 and 2016 (q) tells us the impact of auctioning one additional permit on the auction price; this term captures both firms’ WTP for the policy change and the effect on government revenue from the change in permit price. Bayer & Aklin (2020) focus their analysis on five sectors whose emissions are stationary (energy, metals, minerals, chemicals, and paper – authors’ Table S2), which aligns with the quantity of allowances we use in our calculations. 3,029,915,982 allowances were “auctioned or sold” according to EEA (2024) between 2008 and 2016. Multiplying 3,029,915,982

²⁴⁴Both papers in our extended sample that evaluate ETS focus on changes in emissions released by stationary polluters. Although ETS now covers emissions from aviation, our analysis does not address these emissions.

by dp/dq (7.69/-1,219.5 million) indicates that auctioning one additional permit decreased the price of all allowances in circulation by \$19.09 (in 2016 dollars).

Next we detail each component that enters into our MVPF calculation. We focus on the in-context MVPFs for cap and trade but include what the MVPF would be in 2020 if one assumed $q \times dp/dq$ remained constant over time, only adjusting dp for inflation.

Firms' WTP Firms covered by ETS have a WTP for the increased spending on permits that arises from the change in the quantity of allowances auctioned. As described in Section 6, marginal firms that no longer obtain permits are indifferent to doing so by the envelope theorem. Firms' WTP for the change in the price of all permits in circulation equals $q \times dp/dq$, or \$19.09 (in 2016 dollars). In 2020, holding q and dq constant but adjusting dp from 2016 to 2020 dollars (\$8.29 in 2020 dollars), firms would be willing to pay \$20.59 (in 2020 dollars) for a one unit change in the quantity of allowances auctioned.

Society's WTP Issuing one fewer allowance abates one metric ton of CO_2 . We use the average social cost of carbon between 2008 and 2016 (all expressed in 2016 dollars) when calculating the in-context MVPF, weighting by the total quantity of allowances sold to firms each year. This yields an SCC of \$156.78 (in 2016 dollars) per ton. We multiply this value by the share of the SCC that does not flow to the US government (0.981). Society is therefore willing to pay \$153.78 (in 2016 dollars) for abated CO_2 .

For our 2020 specification, we augment our results by adjusting for the rising social cost of carbon (\$193 in 2020, in 2020 dollars). This adjustment implies a WTP of \$189.30 (in 2020 dollars) after accounting for the share of the SCC that flows to the US government as revenue.

We do not consider leakage from ETS, nor do we consider any rebound effect, as the total change in CO_2 emissions estimated between 2008 and 2016 should include changes in emissions that arise due to changing electricity prices. We focus our analysis on CO_2 emissions but note that incorporating any local benefits would reinforce our conclusions by further increasing society's WTP for the policy change (Basaglia et al. 2024).

Total WTP The total willingness to pay for a one unit change in allowances auctioned in ETS is the sum of firms' WTP for the change in permit spending (\$19.09) and society's WTP for the net change in CO_2 damages (\$153.78). For consistency with other revenue raisers, we consider auctioning one additional allowance, as this results in a negative WTP for society (more CO_2 is emitted) and a positive WTP for firms (spending on permits decreases). Auctioning one additional permit results in a total WTP of -\$134.68 (in 2016 dollars). Total WTP is negative since society's WTP for abated pollution outweighs firms' WTP for higher permit prices.

In 2020, the total WTP for a one unit increase in allowances auctioned is -\$168.71 (in 2020 dollars), with firms willing to pay \$20.59 and society willing to pay -\$189.30.

Cost The denominator of the ETS MVPF contains three terms. First, the government generates revenue by selling one additional permit at a given price. In our in-context specification, we use the price used to calculate dp (\$7.69 in 2016 dollars) to quantify the revenue generated from selling one additional permit. In our 2020 specification, we use the average clearing price for all auctions held in 2020, weighted by the quantity of allowances sold (\$18.53 in 2020 dollars).

The change in permit prices also affects the government’s budget, as all permits in circulation are sold at a difference price. When auctioning one more allowance, the government loses revenue from all permits selling at a lower price, or $q \times dp/dq$. This is equal to firms’ WTP for higher permit prices. In our in-context specification, the government loses \$19.09 (in 2016 dollars) from lower permit prices. In 2020, the government loses \$20.59 (in 2020 dollars) from lower permit prices.

We also account for the fiscal externality on government revenue from reduced CO_2 emissions, or the share of global benefits that do not enter the numerator. In our in-context specification, auctioning one additional permit costs the government \$3.00 (in 2016 dollars) due to increased carbon emissions. In 2020, the climate fiscal externality costs the government \$3.70 (in 2020 dollars).

As with WTP, we consider auctioning one additional allowance, which lowers permit prices and results in greater CO_2 emissions. Auctioning one more allowance generates revenue since one extra permit is sold at a given price (-\$7.69 in 2016 dollars). However, this additional permit costs the government \$19.09 from lower permit prices and \$3.00 due to increased CO_2 emissions. This yields a net government cost of \$14.41 (in 2016 dollars) in context. In 2020, auctioning one more permit cost the government \$5.76 (in 2020 dollars).

MVPF Combining our total WTP and net government cost, we obtain an MVPF of -9.35 (-134.68/14.41) in context. In our 2020 specification, we obtain an MVPF of -29.29, again assuming the inflation-adjusted $q \times dp/dq$ holds over time. The MVPF is negative since society’s WTP for changes in pollution outweighs firms’ WTP for the change in permit spending, and since the government loses revenue by auctioning one more permit.

E.13.4 EU Emissions Trading System (Colmer et al. 2024)

The European Union’s Emissions Trading System is the largest and oldest cap-and-trade system that targets greenhouse gases, covering around 5% of global carbon emissions (Colmer et al. 2024). Annual settlement prices come from World Bank (2024), who report the annual settlement price (in US dollars per ton of CO_2e) as measured on or around April 1 of each year. For our MVPF constructed using estimates from (Colmer et al. 2024), we only require data on the quantity of allowances auctioned to weight price data across years; we take allowance data from EEA (2024). One allowances authorizes a firm to emit one metric ton of CO_2 . Unlike in our discussion of RGGI, all quantities are denoted in metric tons.

We use estimates from (Colmer et al. 2024), who pair a matched difference-in-differences approach with variation in exposure to ETS arising from the roll out of ETS and participation criteria to estimate the change in CO_2 emissions from regulated firms relative to unaffected unregulated firms. The authors report the percent change in carbon emissions from ETS, or dq/q , for two time periods (authors’ Table 2, column 1, rows 3 and 4). We use an average (-0.1515) of the estimate for trading phase I (-0.140, authors’ Table 2, column 1, row 3) and the estimate for trading phase II (-0.163, authors’ Table 2, column 1, row 4). We evaluate ETS between 2005 and 2012 to account for changes in permit price during both trading phase. We do not form a confidence interval for this MVPF since we include it in our extended sample.

Since the authors report the percent change in emissions, or dq/q , we only require the change in permit price (dp) to form an MVPF. We use the average settlement price ($dp = \$19.90$ in 2012

dollars) per allowance from the time period studied by the authors (2005–2012). To calculate the average settlement price, we weight the annual allowance price reported by World Bank (2024) by the total number of allowances auctioned or sold that year, as reported by EEA (2024). Dividing dp by the authors' dq/q gives us $q \times dp/dq$, or \$131.32 (in 2012 dollars). As in our other cap-and-trade MVPFs, this value indicates that auctioning one additional permit decreased the price of all allowances in circulation by \$131.32 (in 2012 dollars).

Next we detail each component that enters into our MVPF calculation. We focus on the in-context MVPFs for cap and trade but include what the MVPF would be in 2020 if one assumed $q \times dp/dq$ remained constant over time, only adjusting dp for inflation.

Firms' WTP Firms covered by ETS have a WTP for the increased spending on permits that arises from the change in the quantity of allowances auctioned. As described in Section 6, marginal firms that no longer obtain permits are indifferent to doing so by the envelope theorem. Firms' WTP for the change in the price of all permits in circulation equals $q \times dp/dq$, or \$131.32 (in 2012 dollars). In 2020, holding q and dq constant but adjusting dp from 2012 to 2020 dollars (\$22.43 in 2020 dollars), firms would be willing to pay \$148.06 (in 2020 dollars) for a one unit change in the quantity of allowances auctioned.

Society's WTP Issuing one fewer allowance abates one metric ton of CO_2 . We use the average social cost of carbon between 2005 and 2012 (all expressed in 2012 dollars) when calculating the in-context MVPF, weighting by the total quantity of allowances sold to firms each year. This yields an SCC of \$137.43 (in 2012 dollars) per ton. We multiply this value by the share of the SCC that does not flow to the US government (0.981). Society is therefore willing to pay \$134.79 (in 2012 dollars) for abated CO_2 .

For our 2020 specification, we augment our results by adjusting for the rising social cost of carbon (\$193 in 2020, in 2020 dollars). This adjustment implies a WTP of \$189.30 (in 2020 dollars) after accounting for the share of the SCC that flows to the US government as revenue.

We do not consider leakage from ETS, nor do we consider any rebound effect, as the total change in CO_2 emissions estimated between 2005 and 2012 should include changes in emissions that arise due to changing electricity prices. We focus our analysis on CO_2 emissions but note that incorporating any local benefits would reinforce our conclusions by further increasing society's WTP for the policy change (Basaglia et al. 2024).

Total WTP The total willingness to pay for a one unit change in allowances auctioned in ETS is the sum of firms' WTP for the change in permit spending (\$131.32) and society's WTP for the net change in CO_2 damages (\$134.79). For consistency with other revenue raisers, we consider auctioning one additional allowance, as this results in a negative WTP for society (more CO_2 is emitted) and a positive WTP for firms (spending on permits decreases). Auctioning one additional permit results in a total WTP of -\$3.47 (in 2012 dollars). Total WTP is negative since society's WTP for abated pollution outweighs firms' WTP for higher permit prices.

In 2020, the total WTP for a one unit increase in allowances auctioned is -\$41.24 (in 2020 dollars), with firms willing to pay \$148.06 and society willing to pay -\$189.30.

Cost The denominator of the ETS MVPF contains three terms. First, the government generates revenue by selling one additional permit at a given price. In our in-context specification, we use the price used to calculate dp (\$19.90 in 2012 dollars) to quantify the revenue generated from selling one additional permit. In our 2020 specification, we use the average clearing price for all auctions held in 2020, weighted by the quantity of allowances sold (\$18.53 in 2020 dollars).

The change in permit prices also affects the government’s budget, as all permits in circulation are sold at a difference price. When auctioning one more allowance, the government loses revenue from all permits selling at a lower price, or $q \times dp/dq$. This is equal to firms’ WTP for higher permit prices. In our in-context specification, the government loses \$131.32 (in 2012 dollars) from lower permit prices. In 2020, the government loses \$148.06 (in 2020 dollars) from lower permit prices.

We also account for the fiscal externality on government revenue from reduced CO_2 emissions, or the share of global benefits that do not enter the numerator. In our in-context specification, auctioning one additional permit costs the government \$2.63 (in 2012 dollars) due to increased carbon emissions. In 2020, the climate fiscal externality costs the government \$3.70 (in 2020 dollars).

As with WTP, we consider auctioning one additional allowance, which lowers permit prices and results in greater CO_2 emissions. Auctioning one more allowance generates revenue since one extra permit is sold at a given price (-\$19.90 in 2012 dollars). However, this additional permit costs the government \$131.32 from lower permit prices and \$2.63 due to increased CO_2 emissions. This yields a net government cost of \$114.06 (in 2012 dollars) in context. In 2020, auctioning one more permit cost the government \$133.23 (in 2020 dollars).

MVPF Combining our total WTP and net government cost, we obtain an MVPF of -0.03 (-3.47/114.06) in context. In our 2020 specification, we obtain an MVPF of -0.31, again assuming the inflation-adjusted $q \times dp/dq$ holds over time. The MVPF is negative since society’s WTP for changes in pollution outweighs firms’ WTP for the change in permit spending, and since the government loses revenue by auctioning one more permit.

E.14 International Subsidies

We compile an illustrative set of international policies and construct MVPFs as if they were to be implemented by the US Government. We do not construct in-context estimates for these policies. For each subsidy, we construct two MVPFs, one considering only benefits accruing to the US and one including non-US benefits. The MVPFs including non-US benefits are likely underestimates given that we are not able to monetize all of the local benefits. In particular, many of these policies provide health and productivity benefits to local communities that we do not account for.

E.14.1 Subsidies for Cookstoves

We include two estimates of cookstove subsidies which lead to divergent MVPFs. Using estimates from Berkouwer & Dean (2022) for improved cookstove subsidies in Kenya, we find an MVPF of 323.5 (36.6 for US-only benefits). Using estimates from Hanna et al. (2016) for

cookstove subsidies in India, we find an MVPF of -2.28 (-0.37 for US-only benefits). The construction of both MVPFs are discussed below.

Cookstoves in Kenya using Estimates from Berkouwer & Dean (2022)

Cookstoves can provide both environmental benefits from reduced charcoal usage as well as private benefits from lower charcoal spending. Berkouwer & Dean (2022) test whether easing credit restrictions and providing information on the amount of savings increases take-up of improved cookstoves. Berkouwer & Dean (2022) run a randomized experiment with 1,000 households in Nairobi, Kenya in which they assign each household a credit treatment (delayed repayment) and an attention treatment (information on cost savings). Exploiting this randomization, the paper is able to estimate the causal impact of the cookstove on charcoal usage.

Berkouwer & Dean (2022) use the take-up from each treatment arm to estimate a demand curve of households' willingness to pay for the cookstove. The paper estimates that a \$30 subsidy would increase adoption from 0.6 percent to 54.5 percent. Therefore, only 1.1% of beneficiaries are inframarginal (0.6/54.5). For marginal households, Berkouwer & Dean (2022) estimate that a cookstove provides 3.5 tons of annual carbon benefits over two years.

Since there are credit constraints that prevent otherwise interested buyers from purchasing cookstoves, we assume that recipients of the subsidy value the reduced charcoal spending from the improved cookstove. Berkouwer & Dean (2022) estimate that households that take-up the cookstove reduce charcoal expenditures by \$2.28 per week compared to the control group.

For ease of interpretation, the components of the cost and WTP below are reported at a per rebate level. To match these values to those in Table 2, one can divide each component by the program cost (\$30.37 in 2020 dollars).

Cost The total cost per subsidy is the sum of the subsidy amount and the climate fiscal externality. The subsidy per cookstove, adjusted to 2020 dollars, is \$30.37. The climate fiscal externality is 1.9% of the total global environmental benefits as explained in Section 7. For the roughly 99% of marginal households, the cookstove provides seven tons of carbon benefits. After discounting and applying a \$193 social cost, the climate fiscal externality reduces the total cost by \$25.60. The total cost per subsidy net of the climate fiscal externality is \$4.77.

WTP The WTP consists of private benefits to households receiving the cookstoves as well as climate benefits from the reduced carbon emissions. The 1.1% of inframarginal households value the entire subsidy as a transfer. The remaining marginal households value the subsidy to the extent to which it provides them with savings on charcoal spending. The cookstove lasts for two years and provides \$2.28 in savings benefits per week. After discounting the second year of savings at the baseline 2% level, the subsidy provides \$232.21 in savings for marginal households. Adding the inframarginal benefits, the total private benefits are \$232.54. The climate benefits from the 7 tons of abated carbon amounts to \$1,311.00 in benefits per cookstove. Therefore, the total WTP is \$1,543.54. To calculate the US benefits, we ignore the private benefits and only consider the 15% of global benefits that flow to the US. The resulting US MVPF is 36.65.

Cookstoves in India using Estimates from Hanna et al. (2016)

Hanna et al. (2016) run a large scale randomized experiment in India in which treated households received subsidized cookstoves. Unlike the results in Berkouwer & Dean (2022), this paper finds that households receiving the improved cookstove did not reduce their wood consumption

and experienced no health benefits. This is due in part to the fact that households continued to use their old stoves for a significant portion of their meals and did not invest the necessary time and money to maintain their stoves.

Take-up of the stoves in the control group was 6.18% and take-up in the treatment group was 68.23%. Therefore, roughly 9% of the beneficiaries were inframarginal. We assume that the inframarginal beneficiaries value the entire transfer. Assuming a linear demand curve and a uniform distribution over the threshold at which one chooses to take-up the subsidy, we estimate that marginal beneficiaries value 50% of the transfer. We do not value the energy savings accrued from this cookstove because there was no reduction on household wood expenditure and these cookstoves are sufficiently inexpensive such that there are no significant credit constraints.

To value the climate benefits from the improved cookstove, we first estimate the climate damages from burning wood used for the standard cookstove. The paper finds that the control group uses 3.73 kg of wood per meal and cooks, on average, 12.61 meals per week on their stoves. The emissions factor for wood burning is 1590 grams per kilogram (Bhattacharya et al. 2002). This results in \$678.72 in carbon damages from a standard cookstove each year.

The paper has a four year follow-up period. It finds that the amount of wood used per meal (in kg) in the first four years relative to the control group is -0.022, 0.156, 0.235, and -0.180. While these estimates are statistically insignificant, the paper finds that total wood consumption actually increased in the treatment group relative to the control group. Dividing these level changes by the control group mean of 3.73 kg of wood per meal gives us the percent change per year. We use these yearly percent changes to calculate the externality components of the MVPF.

Cost The total cost per subsidy is the sum of the subsidy amount and the climate fiscal externality. The subsidy per cookstove is \$12.50. The climate fiscal externality is 1.9% of the total global environmental benefits as explained in Section 7. Since the cookstoves actually increased wood consumption, the climate fiscal externality increases the cost by \$0.72 resulting in a total cost of \$13.22.

WTP The WTP consists of transfer benefits to households receiving the cookstoves as well as climate benefits from the reduced carbon emissions. The 9% of inframarginal households value the entire subsidy as a transfer. The remaining marginal households value 50% of the subsidy. Therefore, the total private benefits to households per subsidy is \$6.82. For the climate benefits we scale the \$687.72 in damages per cookstove by the annual percent change in wood consumption. The discounted sum of the four years of the cookstove is -\$36.95. Therefore, the total WTP is -\$30.13. resulting in an MVPF of -2.28. To calculate the US benefits, we ignore the private benefits and only consider the 15% of global benefits that flow to the US. The resulting US MVPF is -0.37.

E.14.2 Rice Burning PES

We evaluate the MVPF of two payments for ecosystem services (PES) policies targeted at preventing crop residue burning in India from Jack et al. (2022). While crop residue burning is technically illegal in India, limited enforcement has led to burning becoming the primary method for clearing crop residue from rice harvesting. This paper conducts a randomized experiment in which farmers are placed into one of three groups: control, upfront payment, and standard payment. We construct an MVPF for both the standard and upfront payment from the perspective of the US government.

Approximately 15,119,030 tons of carbon are emitted annually from rice burning in the Indian state of Punjab, the location of the experiment (Deshpande et al. 2023). Given that there are two million hectares of burned land in Punjab, there are 3.06 tons of carbon emitted per acre of burned land.

For both the upfront and standard payment MVPFs, the cost and WTP components are reported per unburned acre. Therefore, the environmental benefits are the monetized 3.06 tons of avoided carbon damages.

Standard Payment

For villages that were randomized into the standard payments treatment, the farmers only received the PES after it was verified that they did not burn the crop residue on their land.

Cost The average standard payment per acre was 102.50 rupees. Scaling this cost by the proportion of marginal acres unburned results in a cost per unburned acre of 5,280 rupees or \$71.25. The climate fiscal externality is 1.9% of the total global environmental benefits as explained in Section 7. This externality lowers the cost by \$11.31 resulting in a total cost of \$59.94.

WTP The WTP consists of transfer benefits to households receiving the cookstoves as well as climate benefits from the reduced carbon emissions. Using the paper's max accuracy model, an additional 2% of farmers in the standard payment group did not burn their plots relative to a mean of 9.8%. Therefore, the percent of inframarginal beneficiaries is 83%. For marginal beneficiaries, we assume that the threshold subsidy amount that convinces farmers to take-up the PES is uniformly distributed between zero and the subsidy amount. As a result, we estimate that marginal farmers value 50% of the subsidy leading to a total farmer willingness to pay of \$65.21.

Monetizing the 3.06 tons of carbon benefits at a social cost of carbon of \$193 gives us an environmental externality of \$579.12. The total WTP is \$644.34 and the MVPF is 10.75. To calculate the US benefits, we ignore the private benefits and only consider the 15% of global benefits that flow to the US. The resulting US MVPF is 1.29.

Upfront Payment

For villages that were randomized into the upfront payments treatment, the farmers received a portion of the PES unconditionally prior to complying with the contract. The total amount was fixed, so the conditional portion of the PES in the upfront group was smaller than the conditional payment in the standard payment group. Even though their incentive to comply was lower, the farmers had higher compliance and lower inframarginality in the upfront group.

Cost The average upfront payment per acre was 310.50 rupees. Scaling this cost by the proportion of marginal acres unburned results in a cost per unburned acre of 4,032.07 rupees or \$54.42. The climate fiscal externality is 1.9% of the total global environmental benefits as explained in Section 7. This externality lowers the cost by \$11.31 resulting in a total cost of \$43.11.

WTP The WTP consists of transfer benefits to households receiving the cookstoves as well as climate benefits from the reduced carbon emissions. All recipients fully value the entire upfront unconditional payment since it is just a transfer to the farmers. For the conditional payment, the inframarginal farmers value it fully and the marginal farmers value 50% of the payment. The paper reports that 18.3% of people complied with the upfront contract and received the

conditional payment. Using the paper's max accuracy model, an additional 7.7% of farmers in the upfront payment group did not burn their plots relative to a mean of 16.1%. Therefore, of the 18.3% of farmers that complied, the percent of inframarginal beneficiaries is 67.6%.

We can imagine an average payment in which 18.3% of people receive the full 800 rupees and the remaining 82.7% of people receive only the 375 rupee unconditional payment. This results in a average payment of 456.53 rupee. There are three portions of this payment that are fully valued. The 82.7% of non-compliers who receive the 375 rupee unconditional payment value it fully. Secondly, the 12.4% ($0.676 * 0.183$) of inframarginal beneficiaries value the entire 800 rupee payment. Lastly, the 5.9% of marginal beneficiaries ($0.323 * 0.183$) value the entire 375 rupee unconditional payment. Summing up these groups and dividing by the average 56.53 rupee payment results in 94.44% of the transfer being valued entirely. The remaining portion, which consists of the conditional payment for marginal farmers, is valued at 50%. Applying these percentages to the cost per unburned acre of \$54.42 results in a private WTP of \$52.91.

Monetizing the 3.06 tons of carbon benefits at a social cost of carbon of \$193 gives us an environmental externality of \$579.12. The total WTP is \$632.03 and the MVPF is 14.66. To calculate the US benefits, we ignore the private benefits and only consider the 15% of global benefits that flow to the US. The resulting US MVPF is 1.79.

E.14.3 Deforestation Payments

REDD+ Offsets (Sierra Leone)

The Paris Climate Agreement established REDD+, which stands for reducing emissions from deforestation and forest degradation in developing countries. Through REDD+, governments and other groups can pay foresters to preserve forests and in turn reduce greenhouse gas emissions from deforestation. The effectiveness of these programs is dependent on the additionality of the conservation efforts. If foresters would have preserved their forests even in the absence of the financial incentive, then the policy would simply act as a transfer to foresters.

Malan et al. (2024) studies the effectiveness of carbon offsets using evidence from the Gola Rainforest National Park - a REDD+ project in Sierra Leone. The paper uses satellite images of the areas within the REDD+ zone and the areas directly outside the zone to measure the causal impact of the credits on deforestation. They find that the REDD+ program decreased yearly deforestation rates by 30%.

To calculate the environmental externality, we take the carbon reduction estimate of 342,954 tons per year directly from the paper. Malan et al. (2024) finds that there are 929 hectares of avoided forest loss per year and the carbon stock per hectare is 369.30 tons. Monetizing these damages using a \$193 social cost of carbon results in \$66,190,149 in climate benefits.

Cost The cost consists of the mechanical cost of the policy as well as the climate fiscal externality. Since we are constructing an MVPF for the 2020 version of this policy, we apply a cost per permit of \$4.70 in 2021 dollars reported in UN (2012). There were 403,458 permits awarded each year for the Gola project resulting in a total cost of the permits of \$1,811,435.40. The climate fiscal externality is 1.9% of the total global environmental benefits as explained in Section 7. This externality lowers the cost by \$1,267,872.30 resulting in a total cost of \$543,563.08.

WTP We conservatively assume that the foresters are indifferent on the margin to receiving the credit and deforesting their land. In theory, there are likely some inframarginal recipients that value the permit. This assumption will not affect the US-only MVPF. For the climate

benefits, we monetize the avoided carbon damages using a \$193 social cost of carbon and take out the 1.9% that flows to the US treasury. The total WTP is \$64,922,277. The MVPF is 119.44 and excluding non-US benefits, the MVPF is 15.93.

REDD+ Carbon Offsets (Mix)

West et al. (2023) studies 26 REDD+ projects, similar to the Sierra Leone project evaluated in Malan et al. (2024), across six countries and three continents. To evaluate the causal impact of the carbon credits on deforestation, the paper constructs a synthetic control for each of the 26 projects. The synthetic controls are made from a pool of control areas that have similar levels of forest cover and deforestation pressures. They find that eight of the 26 projects showed evidence of additional reductions in deforestation. Therefore, approximately 69% of the projects were inframarginal. The marginal projects also did not achieve the full level of avoided deforestation as expected based on the number of credits they received. Overall, the paper finds that the projects achieved 7% of the expected carbon benefits - suggesting that 7% of the credits were given to marginal foresters.

We report WTP and cost components at a per credit level. To match the values in the table, one can divide each component by the cost per credit (4.49 in 2020 dollars).

Cost The cost consists of the sum of the cost per credit and the climate fiscal externality. We use a cost per permit of \$4.70 in 2021 dollars reported in UN (2012). The climate fiscal externality is 1.9% of the total global environmental benefits as explained in Section 7. This externality lowers the cost by \$0.26 resulting in a total cost of \$4.23.

WTP The WTP consists of transfer benefits to foresters and climate benefits from avoided deforestation. We assume that the 93% of credits that go to inframarginal recipients are valued fully as a transfer. Consistent with our approach in Malan et al. (2024), we assume that the marginal recipients are indifferent between receiving the credit and deforesting their land and therefore do not value the credit. The climate benefits per marginal credit are equivalent to avoiding one ton of carbon. Since 7% of the benefits are marginal, the total climate benefit, after removing the share that flows to the US treasury, is \$13.25. The total WTP is \$17.58 and the MVPF is 4.16 (and 0.42 excluding non-US benefits).

Deforestation PES (Uganda)

Jayachandran et al. (2017) evaluate a program of ecosystem services in Uganda that offered forest-owning households annual payments for conserving their forest. Payments are approximately \$40.69 per hectare in 2020 USD and the program lasted for two years. The treatment group deforested 4.2% of their land which is about half of the control group's deforestation of 9.1% of their land.

WTP There are two components to this policy's WTP: the value private forest owners (PFOs) put on the transfer and the global environmental benefits from the reduced CO2 emissions.

The transfer is calculated as 100% of the subsidy the inframarginal PFOs receive plus 50% of the subsidy the marginal PFOs receive. The proportion inframarginal equals the percent of forests not deforested in the control group divided by the percent of forests not deforested in the treatment group. This is $0.91/0.96 = 0.95$, which means 95% of hectares would have been

conserved without the payments. The average subsidy per PFO was \$36.09, which converted to 2020\$ is \$40.69. Thus the transfer is $36.09 \cdot (0.5 \cdot 0.05 + 0.95) = 39.69$.

The authors estimate the averted CO₂ per PFO to be 183.5 tons if the hectares conserved by the program are permanently conserved. We align with their assumptions about the delay of CO₂ emissions, i.e., that status quo deforestation is uniform over time and that after the program ends, treated PFOs deforest at a 50% higher rate than usual. The authors also assume that CO₂ is emitted ten years after trees are cut down in their base case. Since the payments end after two years, it would take four years to “undo” the two years of conservation. This means the conservation occurs on average at $t = 4$, while in the control group, the two years of conservation occur on average at $t = 1$, so on average the program delayed deforestation by three years. Thus, we value the 183.5 tons per PFO at the difference in the present discounted value of the SCC in 2030 vs. in 2033. The SCC in 2030 is \$230.00 and in 2033 is \$241.00 so the difference in their present discounted values is \$2.38. Thus, the global environmental benefits are $183.5 \cdot 2.38 \cdot 0.981 = \428.23 , where the 0.981 is to account for the global environmental benefits that flow to the US government through increased tax revenue as discussed in Section 7.

The overall WTP is then 467.93.

Cost There are also two components to this policy’s total cost: the program cost and the climate FE.

The program cost includes the subsidies paid out to PFOs (\$40.69 as above) and administrative costs. Section 4 of the Supplementary Material reports a \$14 monitoring cost per eligible PFO when two spot checks happen per day, a \$30 marketing and program management cost per eligible PFO, and a 10% transaction fee for PES payments. Totaling this up and converting to 2020\$ gives \$53.68 in administrative costs, so the total program cost is \$94.37.

The climate FE is approximately 1.92% of the global environmental benefits, which is -8.36.

Thus, the overall cost is 86.00 and the MVPF is 5.44. If we only consider the WTP of the US (15% of the global environmental benefits excluding the climate FE), then the MVPF is 0.66.

Deforestation PES (Mexico)

Izquierdo-Tort et al. (2024) evaluated a randomized trial in Mexico that compared the effects of a standard payment for ecosystems services (PES) contract to one requiring enrollees to enroll all of their forests instead of just some. The new agreements aimed to prevent landowners from only enrolling parcels that they would have conserved anyway, thus reducing inframarginal payments. Payments are MX\$1,000 per conserved hectare for both full-enrollment and standard PES and the program with new contracts lasted approximately one year. The treatment group deforested 5.7 percentage points less of their land than the control group which is equivalent to 39% less deforestation. We calculate the MVPF using no PES as the counterfactual though, using the paper’s citation that standard PES leads to 1.1 percentage points less deforestation from Costedoat et al. (2015). Thus, the relevant treatment effect for our purposes is preventing

deforestation of 6.8% of forest area compared to no PES.

WTP There are two components to this policy's WTP: the value landowners put on the transfer and the global environmental benefits from the reduced CO₂ emissions.

The transfer is calculated as 100% of the subsidy the inframarginal landowners receive plus 50% of the subsidy the marginal landowners receive. In Section 5.3 of Izquierdo-Tort et al. (2024), the authors write that they paid a total of MX\$591,000 to the treatment group, which implies that landowners conserved 591 hectares. Above that, they note that 65.8 hectares were marginal conservation compared to a baseline of no PES. Thus, the proportion marginal is $65.8/591 = 0.11$. The authors also note that the payment per marginal hectare of conservation was US\$448.29 for full-enrollment PES, so the WTP for the transfer can be calculated as $0.11 \cdot 428.24 \cdot 0.5 + (1 - 0.11) \cdot 428.24 = \404.40 .

The authors use prior estimates that the Lacandona forest stores 550 metric tons of CO₂ per hectare. We align with their assumptions about the delay of CO₂ emissions, i.e., that status quo deforestation is uniform over time and that after the program ends, treated PFOs deforest at the same rate as usual. The authors also assume that CO₂ is emitted immediately after trees are cut down in their base case. Since the standard PES landholders deforest 14.2% of their land per year (Table 2 of Izquierdo-Tort et al. (2024)) and we're assuming that is a reduction of 1.1% compared to no PES, it'll take $1/0.153 = 6.54$ years to deplete their remaining forest when PES is over. Thus, to value the delay, we compare the SCC in 2020 to the PDV of the SCC in 2027 (approximately 6.54 years from 2020). The SCC in 2020 is 193 and in 2027 is 219 in 2020\$. The PDV of the 2027 SCC is then 192.41. We can now value the 550 tons of CO₂ at $193 - 192.41 = \$0.59$ per ton. Thus, the global environmental benefits are $550 \cdot 0.59 \cdot 0.981 = \316.80 , where the 0.981 is to account for the global environmental benefits that flow to the US government through increased tax revenue as discussed in Section 7.

The overall WTP is then 721.20.

Cost There are also two components to this policy's total cost: the program cost and the climate FE. The program cost is \$448.29 per marginal hectare as noted above. The climate FE is approximately 1.92% of the global environmental benefits, which is -6.19.

Thus, the overall cost is \$422.05 and the MVPF is 1.71. If we only consider the WTP of the US (15% of the global environmental benefits excluding the climate FE), then the MVPF is 0.10.

E.14.4 Wind Offsets in India

Calel et al. (Forthcoming) study the effectiveness of carbon offsets implemented through the Clean Development Mechanism (CDM), the world's largest carbon offset program. Established in the Kyoto Protocol and run by the UN, the CDM allows countries and firms to fund carbon offsets in other countries and count those reductions towards their own goals. Similar to REDD+, the effectiveness of this program is dependent on the additionality of the offsets.

Calel et al. (Forthcoming) use evidence from wind farms built in India through the CDM and estimate that at least 52% of the projects were inframarginal. They identify blatantly inframarginal projects (BLIMPs) by checking if there was another unsubsidized wind project built in the same state and year that had strictly lower returns. They use three factors to consider the returns to a wind project: the capacity factor, windiness of the location, and proximity to a connection point.

We convert the 52% inframarginal percentage into an elasticity with respect to the cost of wind installation, similar to the elasticities calculated for the wind PTC MVPFs in our sample. We calculate the percent change in price and quantity of wind installations as a result of the CDM credit. We take the CDM credit prices and capacity additions per year from Figure 1 in Calel et al. (Forthcoming). The average credit price in 2020 dollars during the sample period, weighted by the capacity additions per year, is \$14.08 per ton of carbon. Using a CO_2 emissions factor for India in 2013 of 0.81, the CDM credit reduces costs per MWh by \$11.40 (India Ministry of Power 2018). We take annual wind LCOE data from IRENA (2023b) and find a weighted average LCOE during the same period of \$96.06 per MWh. The percent change in price, using the midpoint approach, is 12.62% ($11.40/(96.06 - 11.40 * 0.5)$). To get the percent change in quantity from CDM, we take the installations from CDM projects and multiply by the percent of those projects that are marginal (48%). For the non-CDM projects we take the sum of the inframarginal projects funded through CDM (52%) and the projects not funded through CDM. We find that there is an average of 3.79 GW of wind added each year during the sample from projects marginal to CDM and 11.79 GW added from other projects. The percent change in quantity installed from CDM is 27.67% ($3.79/(11.79 + 3.79 * 0.5)$). Therefore, the elasticity implied by the 52% inframarginal share is -2.19.

We estimate the MVPF for these wind offsets from the perspective of the US government in 2020. Assuming the elasticity estimated in-context applies to 2020, we can calculate the percent change in installations from a \$1 decrease in cost per MWh by dividing by the elasticity by the 2020 cost. The 2020 LCOE for wind in India is \$35.74. The average CDM credit price is 5.3 per ton and the emissions factor in 2020 is 0.71 tons of CO_2 per MWh resulting in a credit price per MWh of \$3.76. Therefore, the cost net of credit is \$31.97 (in 2020 dollars) and a \$1 decrease in cost per MWh leads to a 6.86% change ($2.19/31.97$) in wind installations. We use this semi-elasticity to construct the externalities per mechanical dollar of government spending on CDM credits.

In our MVPF construction, we assume that emissions in the US do not change in response to the US government purchasing carbon credits in India. Alternatively, one could imagine that the US chooses to emit more because they are able to offset these emissions through the CDM. This assumption would lead to a lower MVPF. We also use the 52% inframarginal share in our baseline estimate, which is likely an underestimate of the true inframarginal share. As the inframarginal share increases, the MVPF would converge to one.

To calculate the costs and WTP, we follow the same approach as the cost of the wind PTCs. For ease of interpretation, we will imagine a wind turbine that produces 1 MWh of energy each year for the 25 year lifetime of the turbine. Since the costs accrue at a per unit of generation level, this capacity factor assumption has no impact on the MVPF ratio. The MVPF in 2020 imagines increasing the subsidy level per MWh by \$1.

Cost The cost consists of the sum of the \$1 mechanical transfer per MWh and fiscal externalities. The \$1 per year transfer over the 25 year lifetime discounted at 2% results in a mechanical cost of the subsidy of \$19.91. The 2020 credit, at the per MWh level, is \$3.76. To get the

total fiscal externality from induced spending on marginal projects, we take the product of the semi-elasticity and 3.76 and discount this flow of costs over the 25 year lifetime of the turbine. The resulting fiscal externality is \$5.14.

We also include a climate fiscal externality which is 1.9% of the total global environmental benefits as explained in Section 7. This externality lowers the cost by \$2.91 resulting in a total cost of \$22.14.

WTP The WTP consists of the mechanical \$19.91 transfer and the environmental externality from the clean energy generation. We do not have a comprehensive forecast of the Indian electricity grid over time. Instead, we simply use the carbon emissions factor of 0.71 tons per MWh in 2020 to construct the environmental externality. Each year of the turbine's lifetime, the environmental externality is the product of 0.71, the semi-elasticity, and the social cost of carbon. We also apply a 11 grams per kWh life cycle emissions cost from DOE (2023c). The resulting environmental externality is \$186.30. To be consistent with the wind PTC estimates, we apply a rebound that offsets 20% of the environmental externality. The environmental externality is an underestimate because we do not include local pollutants and do not include learning-by-doing benefits. The total WTP is \$169.16. Dividing the WTP by the cost, we arrive at an MVPF of 7.64 (and 0.90 excluding non-US benefits).

E.15 International Rebates

We estimate the MVPFs of two appliance rebate policies and one weatherization policy in Mexico. We apply a carbon emissions factor for the Mexico grid of 431 grams per kWh from Climate Transparency (2021). We do not have reliable estimates of local pollution from the grid, so our environmental externality is likely an underestimate of local benefits. To be consistent with the other appliance rebates and weatherization policies in our sample, we apply a 20% rebound to electricity consumption in response to a negative shock in electricity demand that lowers prices. We construct MVPFs for 2020 from the perspective of the US government funding these rebates.

Cash for Coolers Appliance Rebate - Refrigerators

Davis et al. (2014) study the impact of an appliance rebate program in Mexico that replaced old refrigerators and air conditioners with energy-efficient models. We construct the MVPF separately for the fridge and air conditioner rebates. The paper finds that a fridge replacement through the program reduces household electricity consumption by 11.2 kWh per month. In work by Boomhower & Davis (2014), they find that roughly half of the program participants in the Mexico appliance rebate program are additional. Based on this estimate, we use a 50% marginal share in our MVPF calculation.

Cost The cost consists of the mechanical cost of the policy as well as the climate fiscal externality. The paper reports that the program costs \$129,400,000 in 2010 dollars. There were 858,962 refrigerators replaced through the program leading to a rebate per fridge of \$169.85 in 2020 dollars. The climate fiscal externality is 1.9% of the total global environmental benefits as explained in Section 7. This externality lowers the cost by \$0.33 resulting in a total cost of \$169.51.

WTP The willingness to pay consists of the transfer benefits to marginal and inframarginal households as well the climate benefits from the reduced electricity consumption. The in-

framarginal households value the entire subsidy as a transfer. For marginal households, we assume a uniform distribution over the potential threshold subsidy at which people would do the retrofit, resulting in the average marginal households valuing the subsidy at 50%. Therefore, the \$169.85 subsidy will lead to \$84.92 of benefits for inframarginal households and \$42.46 for marginal households. The paper assumes that the subsidy accelerates a household's decision to buy a new refrigerator by five years. Therefore, we assume that the 11.2 monthly kWh reduction in electricity usage persists for five years. Using a \$193 social cost of carbon and a carbon emissions factor of 431 grams per kWh, we arrive at a global environmental externality net of the climate fiscal externality of \$21.20. We also account for a 20% rebound effect of \$4.15 as described in Section D. The resulting total WTP is \$144.43 and MVPF is 0.85 (excluding non-US benefits is 0.01).

Cash for Coolers Appliance Rebate - Air Conditioners

Davis et al. (2014) study the impact of an appliance rebate program in Mexico that replaced old refrigerators and air conditioners with energy-efficient models. We construct the MVPF separately for the fridge and air conditioner rebates. Energy efficient air conditioners make it cheaper for households to use their air conditioning. The paper finds that the AC has near zero impact on household energy usage in the winter months but significantly increases energy usage in the summer months. They estimate that during the six winter months, households receiving the AC units increase consumption by 0.2 kWh a month and during the summer months households increase consumption by 15 kWh a month. As a result, the program actually increases environmental damages. In work by Boomhower & Davis (2014), they find that roughly half of the program participants in the Mexico appliance rebate program are additional. Based on this estimate, we use a 50% marginal share in our MVPF calculation.

Cost The cost consists of the mechanical cost of the policy as well as the climate fiscal externality. The paper reports that the program costs \$13,400,000 in 2010 dollars. There were 98,604 air conditioners replaced through the program leading to a rebate per AC of \$153.22 in 2020 dollars. The climate fiscal externality is 1.9% of the total global environmental benefits as explained in Section 7. Since the program increases environmental damages, this externality increases the cost by \$0.23 resulting in a total cost per rebate of \$153.44.

WTP The willingness to pay consists of the transfer benefits to marginal and inframarginal households as well the climate damages from the increased electricity consumption. The inframarginal households value the entire subsidy as a transfer. For marginal households, we assume a uniform distribution over the potential threshold subsidy at which people would do the retrofit, resulting in the average marginal households valuing the subsidy at 50%. Therefore, the \$153.22 subsidy will lead to \$76.61 of benefits for inframarginal households and \$38.30 for marginal households. The paper assumes that the subsidy accelerates a household's decision to buy a new refrigerator by five years. Therefore, we assume that the monthly kWh increase in electricity usage persists for five years. Using a \$193 social cost of carbon and a carbon emissions factor of 431 grams per kWh, we arrive at a global environmental externality net of the climate fiscal externality of -\$14.38. We also account for a 20% rebound effect of \$2.82 as described in Section D. The resulting total WTP is \$103.35 and MVPF is 0.67 (excluding non-US benefits is -0.01).

Weatherization Subsidies in Mexico

Davis et al. (2020) study the impact of energy efficient housing upgrades in Mexico. They

implement a field trial in which some new homes were provided with energy efficient upgrades prior to residents moving in. They deploy data loggers to track energy consumption in the treated and control units to evaluate the change in energy consumption following the upgrades. While the results are not statistically significant, the paper finds that treated households used a cumulative 199 kWh more than control units from October 2016 to November 2017. Therefore, the difference per year is 183.69 kWh. This increase is conservative relative to the (statistically insignificant) increase of 16% found in the paper's regression specification on log electricity consumption. We assume a lifetime of seven years which is the average of the three year and 11 year persistence of energy efficiency upgrades found in Kotchen (2017) for electricity and natural gas, respectively.

Cost The cost consists of the cost of the energy upgrade and the climate fiscal externality. The paper reports that the upgrades cost between \$650 and \$850 per home. Taking the average and converting to 2020 dollars leads to a cost per home of \$808.88. The climate fiscal externality is 1.9% of the global environmental externality. Since the program increased environmental damages due to increased energy consumption, the climate fiscal externality increases the total cost. The resulting climate fiscal externality is \$1.22 and the total cost per upgrade is \$810.10.

WTP The willingness to pay consists of the transfer benefits to marginal households as well the climate damages from increased electricity consumption. Since the energy upgrades were quasi-random and households didn't elect to have them done, it is unlikely that there are inframarginal beneficiaries. We assume all households value 50% of the cost of the upgrade. Therefore, the \$808.88 subsidy will lead to \$404.44 of benefits for households. Since households who received the upgrades use 183.69 more kWh per year, the rebates have a negative environmental externality. Using a seven year lifetime, \$193 social cost of carbon, and a carbon emissions factor of 431 grams per kWh, we arrive at a global environmental externality net of the climate fiscal externality of -\$78.00. We also account for a 20% rebound effect of \$15.28 as described in Section D. The resulting total WTP is \$341.72 and MVPF is 0.42 (excluding non-US benefits is -0.01).

E.16 International Nudges

We estimate the MVPFs of two electricity reduction nudge programs in Qatar and Germany. The construction of these MVPFs are similar to the construction of the Home Energy Report MVPFs in the U.S.

Nudge (Qatar)

Our MVPF for electricity nudges in Qatar using estimates from Al-Ubaydli et al. (2023) is 6.53. While the externalities and causal effects are specified to the Qatar context, the estimation approach is identical to Home Energy Reports. Al-Ubaydli et al. (2023) send 12 nudges to households and estimate the cumulative annual treatment effect to be 3.8%. The baseline electricity consumption for households in Qatar is reported to be 2,780.67 kWh per month. Therefore, households reduce electricity by 1267.99 kWh a year in response to the nudge. There are 489.97 grams of carbon emitted per kWh of electricity in Qatar. For consistency with our HER MVPFs, we assume a cost per nudge of \$1 (Allcott & Kessler 2019) and a rebound effect of 20%.

Nudge (Germany)

Our MVPF for electricity nudges in Germany using estimates from Andor et al. (2020) is 0.33. Andor et al. (2020) send 4 nudges to households and estimate the cumulative annual treatment effect to be 0.719%. The baseline electricity consumption for households in Germany is reported to be 3,304 kWh. Therefore, households reduce electricity by 23.76 kWh a year in response to the nudge. There are 486 grams of carbon emitted per kWh of electricity in Germany. For consistency with our HER MVPFs, we assume a cost per nudge of \$1 (Allcott & Kessler 2019) and a rebound effect of 20%.

E.17 Other Subsidies

We construct MVPFs for two additional subsidy policies that do not cleanly fit into the other categories we included.

E.17.1 California 20/20 Electricity Rebate Program

Our MVPF for California's electricity rebate program using estimates from Ito (2015) is 2.57 [1.90, 3.26] in 2020 and 1.00 in-context. Ito (2015) study an electricity rebate program in California in which households received a 20% discount on their summer monthly electricity bills if they reduced consumption by 20% or more relative to the previous year. The program was implemented in the summer of 2005 and auto-enrolled everyone who had an electricity account prior to a cutoff date in order to avoid manipulation and self-selection. Ito (2015) estimate the program's impact on electricity consumption by leveraging the eligibility requirements in their regression discontinuity design.

The paper finds that the program reduced electricity consumption by 4% in inland California and had persistent impacts for the following three years. The find that the program had no impact on energy consumption in coastal areas. The difference is driven by higher temperatures and usage of air conditioning in inland areas.

In order to calculate the WTP in the MVPF, we need to determine how much households receiving the 20% subsidy value it. We assume that inframarginal households - people who would have reduced their consumption by at least 20% in the absence of the subsidy - value the entire subsidy. For marginal households, we assume that they are indifferent between consuming more electricity and receiving the subsidy on the margin, so they do not value the subsidy. We do not know the exact breakdown of inframarginal to marginal households. Since we know that coastal group has no decline in their energy consumption, we take the ratio of program spending on coastal to total spending to calculate the inframarginal share. This is likely an underestimate of the inframarginal share because it assumes all of the inland beneficiaries are marginal. The program spending on the coastal and inland group was \$9,358,919 and \$1,250,621, respectively. This results in an inframarginal share of roughly 88%.

The baseline summer electricity consumption for all households in the coastal and inland region is approximately 8.26 and 1.20 billion kWh, respectively. The policy reduced electricity consumption by 0.12% in the coastal area and 4.11% in the inland area in the summer it was implemented. While there were no persistent impacts on the coastal area, the inland area saw reduced electricity consumption by 3.92%, 4.69%, and 4.21% in the three summers following the program's implementation.

We compute a baseline MVPF for this policy in the US in 2020 and an in-context estimate of the MVPF for the policy in California in 2005.

Cost The cost consists of the sum of the total cost of the electricity rebates, the fiscal externality from lost tax revenue from electric utilities, and the climate fiscal externality. The sum of the coastal and inland rebates in 2020 dollars is \$14,064,010 and in 2005 dollars is \$10,609,540.

The lost government revenue per kWh, as explained in Appendix C.2, is \$0.006 in the US in 2020 and \$0.018 in CA in 2005. Applying this externality to the reduced electricity consumption outlined above results in a fiscal externality from lost utility profit tax revenue of \$998,863.26 in the baseline and \$3,062,978.60 in-context. The climate fiscal externality, as explained in Section 4, reduces the total cost of the policy by 1.9% of the global environmental benefits. The climate fiscal externality is -\$461,553.70 in 2020 and -\$177,044.07 in-context. The resulting total cost is \$14,601,319 in 2020 and \$13,495,475 in-context.

WTP The willingness to pay is comprised of the inframarginal benefits, environmental externality, and loss in producer profits. As described above, we use an inframarginal share of 88%. Since we assume that marginal households are indifferent between consuming more electricity and receiving the subsidy on the margin, they do not value the subsidy. The transfer value of the subsidy is 88% of the total subsidy cost which is \$12,406,186 in 2020 and \$9,358,919 in-context.

The environmental externality per kWh for the US in 2020 is \$0.16 and for California in 2005 is \$0.06. Applying these monetary damages to the change in electricity consumption in coastal and inland areas leads to a global and local environmental externality in 2020 of \$29,390,844 and \$4,178,802.50, respectively. For the in-context specification, these are \$11,273,823 and \$859,505.85. The rebound effect offsets approximately 20% of environmental benefits resulting in a reduction in the environmental externality of \$6,575,153.20 in 2020 and \$2,376,506.90 in context.

Reduced energy consumption as a result of these incentive rebates leads to lower profits for electric utilities. The construction of the producer profits externality is explained in Appendix C.2. Applying the reduction in electricity consumption in coastal and inland areas, we arrive at a total producer willingness to pay of -\$1,838,816.50 in the baseline specification and -\$5,638,665.20 in-context. Summing across these components, the total willingness to pay in 2020 is \$37,561,863 and in-context is \$13,477,076. This results in a baseline MVPF of 2.57 and in-context MVPF of 1.00.

E.17.2 USDA Conservation Reserve Program

Our MVPF for the USDA's Conservation Reserve Program (CRP) using estimates from Aspelund & Russo (2024) is 2.41 [2.15, 2.66] in 2020. The CRP is one of the largest payment for ecosystem services mechanisms in the world. The program auctions conservation contracts in which landowners receive payments to take cropland out of production and instead conserve the land by planting grass mixes, trees, or establish habitats for a duration of ten years. Aspelund & Russo (2024) investigate the extent to which these payments are additional. In other words, are these payments paying farmers who would not have cropped their land even absent the payment or are they actually changing the behavior of farmers.

Aspelund & Russo (2024) estimate that the additionality under the current CRP auction ranges between 21% and 31%. We conservatively assume that 21% of the permits are addi-

tional. They also model an optimal auction market which would result in 55% of permits being additional. The MVPF under this assumption jumps to 4.72 in 2020.

The average bid price per acre per year as reported in the paper is \$83. In order to calculate the MVPF, we estimate the environmental externalities from one additional acre that is not cropped as a result of receiving a conservation contract. Consistent with Aspelund & Russo (2024), we take the carbon abatement estimate from the USDA (2017). In 2017, there were 23.4 million acres enrolled and 44 million metric tons of carbon reduced through the CRP. The carbon reduction per additional acre is 1.88 metric tons, roughly 25% of which is from reduced fertilizer use. The monetized carbon benefit per additional acre per year is \$362.84.

To estimate the non-carbon benefits of the program, we follow the approach of Aspelund & Russo (2024). They take the average welfare gains from three papers in the literature (Feather et al. (1999); Hansen (2007); Johnson et al. (2016)) and monetize those benefits. For carbon benefits, they apply a social cost of carbon of \$43. From the three papers, they compute the average of four values of the social benefits per acre enrolled in the program: \$98.34, \$255.70, \$367.96, and \$456.04. To compute the average local benefits, we subtract the carbon benefits from each of these four estimates and compute the adjusted average. A social cost per ton of \$43 and 1.88 tons of carbon avoided per acre implies carbon benefits of \$80.84. Subtracting \$80.84 of benefits from each estimate and computing the new average leads to local benefits of \$213.66 per additional acre per year.

We estimate the MVPF for this policy from the perspective of the US government purchasing a permit in the 2020 CRP auction market. Since the program covers the entire US and the paper uses data until 2021, we treat the in-context and 2020 MVPF as the same. The WTP and cost components reported below calculated for one permit per year. To match the values from Table 2, one can divide each component value by the mechanical cost of \$83 per permit.

Cost The cost consists of the mechanical cost of the policy as well as the climate fiscal externality. The cost per permit per year is \$83. The climate fiscal externality is 1.9% of the total global environmental benefits as explained in Section 7. This externality lowers the cost by \$1.49 resulting in a total cost of \$81.51.

WTP The WTP consists of the transfer benefits to landowners as well as the local and global environmental benefits. Since 21% of the permits are additional, 79% of the payments are flowing to landowners that are not changing their behavior in response to receiving the permits. These inframarginal landowners value the entire payment. For marginal landowners, we assume a uniform distribution over the permit price at which they would enroll in the CRP, resulting in the marginal landowner valuing the subsidy at 50%. Therefore, the \$83 subsidy will lead to \$65.21 of benefits for inframarginal households and \$8.89 for marginal households.

As explained above, the monetized carbon benefits from each additional acre per year are \$355.95, excluding the 1.9% that flow to the US Treasury and the local benefits are \$213.66 per additional acre per year. Taking the sum of these and multiplying by the percent additional (21.4%) results in a total externality of \$122.06. The total WTP is \$196.17 and the MVPF is 2.41.

F Publication Bias

In this section, we provide more details on our procedure to estimate and correct for publication bias in the environmental economics literature. Our approach follows that outlined in Andrews & Kasy (2019), with some modifications that relax the assumptions required for identifying the degree of publication bias in our data. Broadly, we find modest evidence of publication bias. Studies are about twice as likely to be published if they have t-statistics above 1.96. While nonzero, this is less bias than found in many other literatures, and correcting for this bias has virtually no observable effect on our results. We also note below that the off-the-shelf procedure from Andrews & Kasy (2019) for identifying publication bias does not do a very good job of fitting the distribution of our estimates, as it imposes normality on the underlying effect sizes. We relax this normality assumption below, although we note that had we not made this adjustment our publication bias-corrected estimates would still deliver very similar results.

F.1 Estimating Publication Bias

We first form a dataset of the t-statistics for the studies underlying our estimates.²⁴⁵ We restrict attention to our baseline sample and drop all observations for which there are no reported measures of sampling uncertainty. Our focus is on the literature measuring elasticities and semi-elasticities of climate-relevant outcomes with respect to various policies. To that end, we drop estimates of pass-throughs and markups, which we view as ancillary to the main objects of interest. This yields a final sample of 103 distinct estimates with t-statistics.

Appendix Figure 9 provides heuristic evidence of the presence of publication bias in our sample. Here, we show a scatterplot of the standard errors for the studies in our sample against the corresponding point estimates. The dashed gray lines indicate slopes of $-1/1.96$ and $1/1.96$; data points above these lines are insignificant (assuming a conventional 5% cutoff), while those below are significant. This “funnelplot” shows substantial excess mass below the dashed lines. Assuming, as in Andrews & Kasy (2019), that standard errors and point estimates would be uncorrelated in the absence of publication bias, this offers suggestive evidence that conventionally significant estimates are more likely to be published.

While this offers evidence of the presence of publication bias, to correct for the distortions such biases induce, we require an estimate of the *degree* of publication bias. We do so via a regression-discontinuity-like design, comparing publication probabilities below and above the 1.96 cutoff.²⁴⁶ Panel B of Appendix Figure 9 visualizes our procedure. We form bins of t-stats of width .98 and count the number of published studies in each bin; our estimate of publication bias is given by the ratio of the number of studies in the bin $[1.96, 2.94)$ relative to the number in the bin $[-.98, 1.96)$. This yields a ratio of 2.18 (p-value for the null hypothesis that the ratio is 1: $< .0001$). Assuming that the underlying distribution of t-statistics is smooth in the neighborhood of the cutoff, this corresponds to the odds ratio of publication for significant vs

²⁴⁵Most papers in our sample report point estimates and standard errors or t-statistics directly. Some papers report p-values only; for these, we invert the p-value assuming 95% two-sided normal hypothesis tests to yield the corresponding t-statistics.

²⁴⁶Here, we conduct our analysis using the absolute value of the t-statistics. In addition to increasing statistical power, the signs in our baseline sample are often arbitrary (e.g., some demand elasticities are reported as negative; others are reported as positive and are implicitly understood to be absolute values). Moreover, Panel A of Appendix Figure 9 shows approximate symmetry around 0, suggesting that ignoring the signs of the estimates sacrifices little information.

insignificant studies.²⁴⁷

F.1.1 Comparison with Andrews & Kasy (2019) Method

Our approach to estimating publication bias differs slightly from the methodology proposed in Andrews & Kasy (2019). While their paper offers non-parametric identification results, in practice they estimate publication bias by specifying (1) a parametric hyperdistribution of the true effect sizes and (2) regions of different publication probabilities, corresponding to t-stats above or below conventional significance levels. By assuming a functional form for the hyperdistribution (e.g., Gaussian, T-), this approach imposes more assumptions, whereas our method imposes only the nonparametric requirement that the distribution of true effect sizes is continuous at the threshold between regions.²⁴⁸

Appendix Figure 10 presents the implied CDFs from our method and from that of Andrews & Kasy (2019), both compared to the empirical CDF of the (absolute value of the) t-stats in our data. Two important patterns stand out. First, except for the kink around a t-statistic of 2, the empirical CDF is relatively smooth in the regions below and above this cutoff. This suggests that focusing on the region local to the 1.96 cutoff allows us to capture the main source of publication bias in our sample without imposing parametric restrictions on the true effects. Second, the parametric approach from Andrews & Kasy (2019) appears to significantly overestimate the jump in publication probabilities around the cutoff, leading to a much steeper kink in the CDF than is observed in the data. Intuition for this result lies in the fact that the minimal t-statistic in our sample is roughly .5, implying a large degree of missing mass in the region $[0, .5)$. Since the approach in Andrews & Kasy (2019) imposes full support priors (e.g., Gaussian), a very high degree of publication bias is necessary to rationalize this “empty” region. In contrast, our local estimator around the t stat of 1.96 threshold does not infer publication bias from the extent of mass in the region between 0 and 0.5.

F.2 Correcting for Publication Bias

Armed with estimates of the degree of publication bias, we can use the approach in Andrews & Kasy (2019) to correct our data for the distortions such bias induces. They consider a setup in which a researcher observes a draw from a distribution centered at the true effect size but with some noise given by the study’s standard error. The draw is then published with possibly different probabilities depending on whether or not it is significant. In this setting, they show the studies’ standard errors and the estimates of publication bias allow for median-unbiased estimation of the true effects. In other words, we can compute the true effect size such that the published study is at the 50th percentile of the implied distribution of published effects.

Appendix Figure 11 shows the results of applying this bias-correction procedure to our raw estimates and re-creating the MVPFs in our baseline sample. It shows that our core conclusions remain unaffected by correcting our estimates: wind policies continue to dominate solar policies, which outperform the other subsidies in the sample. While these patterns emerge with the relatively modest degree of publication bias we find in our approach, applying the bias

²⁴⁷While our baseline binning of 0.98 is relatively large, we obtain similar results for smaller bins e.g., .49 and .28.

²⁴⁸Their approach allows for estimating publication bias over the entire range of the data, whereas ours forces us to focus on the region local to the cutoff.

correction with higher degrees of estimated bias (e.g., using the approach in Andrews & Kasy (2019)) similarly preserves our main conclusions.

G Regulation

Our primary results focus on the welfare benefits and costs of taxes and subsidies that affect greenhouse gas emissions. But, two of our core conclusions focus on the desirability of Wind PTC subsidies and gasoline taxes. In both cases, there exist comparable regulatory policies - namely renewable portfolio standards (RPS) for utility companies to purchase clean energy and fuel efficiency (CAFE) standards for automakers to increase gas mileage of their vehicles. In this section, we show how to relate our MVPF estimates for these policies to ask whether these regulations or taxes/subsidies are better methods to achieve the welfare gains identified from these policies in the MVPF framework.

While tax and spending policies involve direct tradeoffs between a group in the economy and the government budget, regulatory policies primarily involve tradeoffs between two groups of people in the economy. For example, improvements in fuel standards may reduce emissions but car owners may need to pay more for cars under stricter fuel economy standards.²⁴⁹ While regulations can have an impact on the government budget (e.g., a reduction in gasoline usage can reduce gas taxes collected), the fact that the government budget impacts tend to be smaller than the core tradeoffs between groups in the economy means that it is less informative to construct an MVPF corresponding to a change in regulation policy. However, we show that the MVPF estimates of tax and spending policies can be combined with MVPFs of tax and transfer policies to study the relative desirability of regulation versus tax+spending policies.

Our approach takes a slightly different conceptual experiment than the one outlined in equation (5) and instead follows in the tradition of excess burden calculations. The key question we ask is whether the environmental benefits obtained through regulation could be obtained more efficiently through taxes and subsidies. For example, can a combination of tax and spend policies be used to replicate the distributional incidence of regulation across all groups of beneficiaries? This question is motivated by the original ideas of Kaldor (1939) and Hicks (1940) who suggested we can use combinations of policies to neutralize distributional incidence when making policy comparisons. Kaldor and Hicks envisioned individual-specific lump-sum transfers to create these policy combinations. The MVPF framework, as outlined in Hendren (2020), allows us to extend this idea to consider feasible policy tools to neutralize this incidence.

In this Appendix, we present a detailed description of the results from this exercise where we use the MVPFs of gasoline tax and income tax policies to compare gas taxes and income taxes to CAFE standards using estimates of the impact of CAFE from Leard & McConnell (2017), Anderson & Sallee (2011), and Jacobsen (2013*a*). We then present a comparison of Wind PTCs to Renewable Portfolio Standards (RPS), which require utility companies to source a certain fraction of their energy from clean sources.

²⁴⁹In practice, one can still construct the MVPF of this policy. The net cost of the policy is determined by the change in gas tax revenue. The WTP is determined by the sum of the effects on consumers and global beneficiaries of emissions reductions.

G.1 Corporate Average Fuel Economy Standards (CAFE)

Corporate Average Fuel Economy (CAFE) standards have been an important method for regulating vehicle emissions in the US. These standards require automakers selling light-duty vehicles of a given model year in the US to meet specified fleet-wide average fuel economy ratings (typically stated in terms of miles per gallon). We show how to relate our results on the MVPF of the gas tax to results from papers estimating the costs and benefits of changes in the stringency of CAFE standards.

We begin by combining estimates of the costs of CAFE standards from Leard & McConnell (2017) with our calculations of the lifetime damages generated by the average new light-duty vehicle sold in 2020. Since 2012, vehicle manufacturers who over-comply with CAFE standards receive credits. Over-compliant firms can use these credits to cover under-compliant vehicles they manufacture over the next five years (or to retroactively cover vehicles from the previous three years that fell short of standards) or to offset under-compliant models (so that the firm’s vehicles average out to the CAFE standard).²⁵⁰ Additionally, over-compliant firms can sell credits to under-compliant firms; in a competitive market, the price at which credits are traded reveals the marginal cost of compliance with CAFE standards. While firms are not required to disclose credit prices, Leard & McConnell (2017) infer prices using SEC filings from Tesla, finding an average credit price between \$70 and \$119 (in 2014 dollars). We use an adjusted average credit price of \$99.22 (in 2020 dollars) to calculate the marginal cost of compliance and assume the entire cost is passed onto consumers through higher vehicle prices.²⁵¹

For benefits and costs proportional to fuel use, we calculate the difference in costs/benefits generated by the average new light-duty vehicle released in 2020 (25.38 MPG) and a new vehicle with a 1 MPG higher fuel economy.²⁵² This approach compares the average light-duty vehicle purchased in 2020 to a vehicle that achieved an additional mile per gallon.²⁵³ We account for

²⁵⁰As Leard & McConnell (2017) explain, credits lower costs by allowing firms to over-comply when manufacturing vehicles with lower marginal costs (such as cars) and under-comply with higher marginal cost vehicles (such as light-duty trucks). Credits also allow firms to smooth costs over time.

²⁵¹Leard & McConnell (2017) first calculate an implied credit price in terms of dollars per ton of CO_2 (author’s Table 4), which the authors convert to dollars per mile-per-gallon since carbon emissions are proportional to fuel use. We apply three transformations to harmonize our analysis of CAFE standards with similar policies. First, we take a simple average of the three permit prices (\$70, \$119, and \$80) inferred by the authors. Although the marginal cost of compliance is likely to rise as CAFE standards tighten further, we do not have enough information to estimate how credit prices change as compliance becomes more difficult. Second, we re-scale by the lifetime VMT of our estimated counterfactual vehicle (197,592 miles/author’s reported 195,264 miles) to harmonize with the parameters used to calculate lifetime damages. Lastly, we inflation adjust to 2020 dollars, yielding a credit price of \$99.22 in 2020 dollars.

²⁵²We note that this calculation is identical to how we calculate the environmental benefits from vehicle retirement programs when purchases are not accelerated. Using the baseline vehicle externalities reported in our analysis of either “Cash for Clunker” programs but instead using a 1 MPG fuel economy improvement provides the externality values we input into our analysis of Leard & McConnell (2017) and Anderson & Sallee (2011). Using our calculated externalities and accounting for the rebound in VMT, we find that tightening CAFE standards by 1 MPG generates (in 2020 dollars) \$457.63 in global benefits, -\$149.95 in net local benefits (including local pollution and driving externalities), -\$104.75 in post-tax benefits for producers, \$128.47 in lost corporate and gas taxes for the government, and -\$8.94 in savings from the climate fiscal externality. Net local benefits are negative since the increase in damages from driving externalities (-\$183.75) more than offsets the decrease in local pollution damages (\$33.79).

²⁵³For example, a new light-duty vehicle manufactured and purchased in 2020 received 25.38 MPG and generated \$15,654.55 in global damages over its lifetime. A 1 MPG improvement translates to a 3.94% improvement in fuel economy (1/25.38). Dividing the baseline externality by 1.0394 gives us an adjusted externality of \$15,061.13 for the more fuel-efficient vehicle, before accounting for the rebound in VMT.

the rebound in miles traveled by the more fuel-efficient vehicle using the method outlined in Appendix D.²⁵⁴ This approach to calculating the benefits of CAFE standards assumes the size of the vehicle fleet remains constant while the fleet's composition changes.²⁵⁵

Recall that Appendix Figure 6, Panel A illustrates the costs and benefits of increased CAFE standards, normalized per dollar of environmental benefits (including benefits from accidents and congestion). Every \$1 of environmental benefits leads to a cost on producers of \$0.34 and a cost on consumers of \$0.32. Additionally, lost gas and corporate tax revenue generates a cost to the government of \$0.39. This implies that more stringent CAFE regulation that creates \$1 of environmental benefits delivers an unweighted sum of net benefits to society of \$0.34. The question we now ask is rooted in the classic efficiency tests of Kaldor (1939) and Hicks (1940): Can we do better than this for the affected groups by finding a combination of gasoline taxes and income taxes that generate at least (a) \$1 of environmental benefits, (b) -\$0.34 in producer benefits, and (c) -\$0.32 in consumer benefits at a cost to the government that is less than \$0.39? In other words, can taxes replicate the distributional incidence of the CAFE standards at a lower cost to the government (so that the excess revenue could be redistributed to make everyone better off)?²⁵⁶

To assess this, the orange bars in Appendix Figure 6, Panel A show that every \$1 of environmental benefits provided by the gas tax generates a cost to producers of \$0.14 and a cost to consumers of \$2.31. The tax also generates \$2.15 in government revenue. Next, we combine the gas tax with a tax on producers of \$0.20 to equalize their willingness to pay under the tax regime as in the CAFE expansion (-\$0.34). We assume the MVPF of taxes on producers is 1.8, consistent with estimates of the MVPF of taxes on top earners from Hendren & Sprung-Keyser (2020). This suggests imposing the \$0.20 cost on producers raises \$0.11 ($= .20/1.8$) in revenue for the government. We present this in the second column of Appendix Figure 6, Panel B. Next, we compensate the consumers for the difference between their losses under CAFE versus the gas tax, \$1.99. The MVPF of raising revenue from the average consumer is around 1.2, suggesting that this costs the government \$1.65 ($= \$1.99/1.2$), which we present in the third column of Appendix Figure 6, Panel B. Therefore, the net cost to the government of the gas taxes plus income taxes that replicate CAFE is -\$0.61 ($= -2.15 - 0.11 + 1.65$). On net, Appendix Figure 6, Panel B shows that the government can replicate the distributional incidence of CAFE using taxes and still run a \$0.61 surplus, in contrast to the \$0.39 deficit that CAFE generates. In other words, it is \$1.00 ($\$0.39 - \text{\$}0.61$) cheaper for the government to generate the \$1 of environmental benefits through taxes instead of CAFE. In this sense, although CAFE generates a positive net surplus, our estimates would suggest that the gas tax is more efficient than CAFE at delivering those environmental benefits because one can redistribute back the \$1 in a way that would make each group better off.

Appendix Figure 7 present results for two other analyses of the costs and benefits of CAFE: Anderson & Sallee (2011) (Panel A) and Jacobsen (2013a) (Panel B). We present the benefits and costs of the regulation in blue and the tax in orange.

²⁵⁴The values used in Appendix D are identical to those used in this calculation.

²⁵⁵Jacobsen & van Benthem (2015) argue that tightening CAFE standards decreases vehicle scrappage as drivers respond to changing prices by holding onto older vehicles for longer. They estimate that this effect offsets 13-16% of the expected benefits from tightening CAFE standards. Accounting for this effect would only strengthen our finding that gas taxes are a more efficient means of abating vehicle emissions than fuel economy standards.

²⁵⁶This test of "efficiency" dates back to the classic definition of Kaldor (1939) and Hicks (1940) with the modification that we use actual tax and transfer policies instead of lump-sum redistribution to neutralize distributional incidence.

To evaluate Anderson & Sallee (2011), we repeat the exercise described above but substitute the credit price from Leard & McConnell (2017) with the marginal cost of compliance estimated by the authors. We take the midpoint of the 6 ranges reported in Table 8 of Anderson & Sallee (2011), take simple averages for cars and trucks, and calculate a single weighted cost of compliance using the 2020 car and truck production shares (0.44 and 0.56, respectively) reported in the Automotive Trends Report (EPA 2023*d*). This provides us with an inflation adjusted figures of \$17.75 in 2020 dollars, which we assume is entirely passed onto consumers.

Jacobsen (2013*a*) calculates the welfare effects of a 1 MPG increase in CAFE standards. The paper estimates that the equivalent variation (EV) per ton of CO₂ avoided is \$222 in 2001 dollars (\$324.58 in 2020 dollars). We use this as a measure of the change in consumer and producer welfare per ton of CO₂ abated as a result of tightening CAFE standards. To convert to a total change in surplus, we use the paper’s provided estimates to calculate total tons of CO₂ abated. The paper finds a 3.37% reduction in gasoline usage per household. With a baseline gasoline consumption of 828.89 gallons per household (author’s Table 6) and 20,429 households in the sample, the paper finds tightening CAFE standards abated 570,655.37 gallons of gasoline. Using the paper’s supplied carbon content (0.008887 tons/gallon), we calculate a total reduction of 5,071.41 tons of carbon. Multiplying by the welfare cost of \$324.58 per ton of CO₂ yields a total change in market surplus of \$1,646,065.90 (= 5,071.41 × \$324.58) in 2020 dollars. We use the paper’s estimates for the aggregate change (author’s Table 6) in consumer and producer surplus (-\$24.1 billion and -\$5.52 billion in 2001 dollars, respectively) to calculate the share of welfare losses that flow to consumers and producers (81.4% and 18.6%, respectively). The total change in consumer welfare is -\$1,339,304.10 (in 2020 dollars). Producer welfare decreases by \$306,761.77. Producer welfare here refers to automobile firms from the paper’s model. For consistency with other policies that affect gasoline consumption, we account for changes to gasoline producer welfare from CAFE by multiplying the per-gallon, post-tax markup on gasoline in 2020 (\$0.484 per gallon, as described in Appendix C.4.5) by the total change in gasoline consumption, which results in a total welfare effect on gasoline producers -\$276,156.71. Combining this with the welfare effect on vehicle producers, we find the total change in producer surplus is -\$582,918.49 (in 2020 dollars).

We calculate environmental benefits using the paper’s reported total change in gasoline consumption and vehicle miles traveled. Using the method described above to calculate the total change in gas consumption (570,655.37 gallons), we find a total change in VMT of -9.151 million miles.²⁵⁷ To calculate pollution benefits, we multiply the total change in gas consumption by the average gasoline externality from pollution in 2020 (\$2.12 per gallon). We multiply the change in VMT by the average driving externality in 2020 (\$0.12 per mile).²⁵⁸ This approach implies CAFE generated \$1,058,221.2 in global pollution benefit (adjusted for the share of benefits that flows to the US government), \$128,799.44 in local pollution benefits, and \$1,082,863.50 in abated accidents, congestion, and tire and brake PM_{2.5}.²⁵⁹

²⁵⁷Although we typically assume that CAFE standards *increase* VMT by decreasing the cost of driving, excluding the author’s estimated reduction in VMT would only reinforce our conclusion that gas taxes are more efficient than CAFE standards.

²⁵⁸Since we observe the change in VMT, we need not assume some share of the change in total gasoline consumption arises from changes in VMT (see Appendix C. Rather, we can simply apply our average per-mile externality to the observed change in VMT. The per-gallon pollution externality noted here excludes driving externalities (accidents, congestion, and PM_{2.5} from tires and brakes).

²⁵⁹As noted above, we normalize per dollar of environmental benefits, including benefits from changes in accidents and congestion. To calculate the bars in Appendix Figure 7, Panel C, one should sum these three components and divide the cost imposed on either consumers, producers, or the government’s budget by this

To find the effect on the government, we account for changes in gas and corporate tax revenue as well as revenue from abating carbon emissions. The change in corporate profits is calculated by multiplying the pre-tax, per-gallon markup (\$0.613 per gallon, as described in Appendix C.4.5) on gasoline by the tax rate (21%, from Watson (2022)) and the total change in gas consumption, resulting in a loss of \$73,408.75 in 2020. The same method is used to calculate the loss in gas tax revenue: with an average per-gallon tax rate in 2020 of \$0.465 per gallon, we calculate a loss in gas tax revenue of \$265,280.56. Additionally, we find the long term effect by using overall change in gasoline consumption and converting this to a change in global damages using the per gallon global externalities, for a fiscal externality of -\$20,270.23 (this fiscal externality raises revenue for the government in the long-run).

We display this decomposition in Appendix Figure 7. The blue bar on the far right replicates the analysis above by constructing the net cost to the government of replicating CAFE using taxes and transfers. Using each of these estimates of the welfare impact of CAFE, our estimates imply that taxes can replicate the CAFE benefits at a surplus to the government, in contrast to CAFE, which imposes a net cost to the government. In this sense, our MVPF results imply gasoline taxes are more efficient than CAFE in delivering environmental benefits through reduced gasoline consumption. In theory, two potential mechanisms could be driving this result. First, CAFE imposes implicit taxes on fleet characteristics beyond purely a tax on gasoline emissions (see Ito & Sallee (2018)). These additional taxes impose extra distortions which are not needed if the aim is to simply reduce gasoline consumption. Second, gas taxes reduce vehicle miles traveled, which leads to reductions in accidents and congestion – benefits that are (typically) not achieved through CAFE standards, since fuel economy standards lower the cost of driving and encourage drivers to travel more miles.^{260,261} A deep analysis of the theoretical mechanisms driving the results is beyond our scope; rather, we simply note that the empirical results suggest a superiority of gas taxes over CAFE standards.

G.2 Renewable Portfolio Standards (RPS)

Next, we consider the relative efficiency of wind subsidies compared with Renewable Portfolio Standards (RPS). These regulations, generally passed by states, require power companies to source a certain percentage of their energy from clean sources like wind and solar. Here, we use estimates from Greenstone & Nath (2024), who study the causal effect of these state-level standards. They find that every ton of carbon removed from the atmosphere leads to a reduction of consumer surplus between \$80-\$210 in 2022 dollars. We use the median estimate of \$145 in 2022 dollars (\$128.26 in 2020 dollars). Because the cost per ton measure does not include learning-by-doing benefits or local pollution benefits, we harmonize our estimates to theirs by excluding these components when considering the benefits from a wind PTC.²⁶²

The results suggest that every \$1 of environmental benefits provided by RPS imposes a cost on consumers of \$0.68 and a \$0.02 savings to the government due to the climate fiscal

sum.

²⁶⁰As noted above, however, Jacobsen (2013a) estimates that CAFE will reduce VMT as well, generating additional benefits. Our comparison using Jacobsen’s estimates still implies that gas taxes are relatively more efficient than CAFE standards.

²⁶¹CAFE standards can also encourage drivers to adopt lighter vehicles, which pose higher risks to drivers should they be involved in an accident (see Jacobsen (2013b)).

²⁶²We account for changes in global damages from the rebound effect and from lifecycle costs when comparing the benefits of wind PTCs to RPS.

externality, which are displayed in Appendix Figure 7, Panel C.²⁶³ In contrast, delivering \$1 of environmental benefits through wind PTC subsidies delivers \$0.27 in benefits to consumers and costs the government \$0.37. Producers have no willingness to pay for either policy in our analysis. Income taxes that tax consumers enough to impose the same \$0.68 cost that RPS imposes would generate \$0.79 ($=(\$0.27 - \text{\$}0.68)/1.2$) in revenue. This means one could construct a combined wind PTC and income tax regime that delivers \$1 of environmental benefits and \$0.68 in costs to consumers but generates \$0.42 in government revenue (in contrast to the \$0.02 from RPS). In this sense, the estimates suggest that wind subsidies are more efficient than RPS regulation.

In summary, these examples illustrate how the library of MVPFs we provide can readily be incorporated into welfare analyses of regulations to help assess the relative efficiency of regulation versus combinations of taxes and subsidies. For the estimates of the effects of CAFE and RPS, our results suggest tax and transfer policies are more efficient than regulation. That is, there is the potential to make all affected groups better off with tax and subsidy policies than with the specific regulatory alternative being assessed.

²⁶³We follow Greenstone & Nath (2024) and assume all costs associated with RPS are passed on to consumers. We also assume the wind PTC is passed on to consumers as lower electricity prices.

H Comparison to Net Benefits and Benefit-Cost Ratios

In critiquing cost-effectiveness ratios, we follow a large literature discussing the advantages of benefit-cost analysis over cost-effectiveness analysis because of its more comprehensive nature. Indeed, the MVPF is a particular form of benefit-cost ratio. Benefit-cost ratios are often criticized because it is not conceptually clear what constitutes a cost in the denominator versus a negative benefit in the numerator (Boardman et al. 2018). The MVPF solves this conundrum by being clear about the incidence of the policy: the government incidence is in the denominator; the beneficiaries of the policy are in the numerator. By making the ratio correspond to the incidence of well-defined groups, we remove any indeterminacy around measurement. Moreover, from the perspective of a policymaker seeking to maximize social welfare subject to a government budget constraint, the MVPF has a Lagrange multiplier interpretation: it helps characterize the extent to which social welfare can be increased per dollar of net government spending on a policy. In this sense, the MVPF is a key statistic for attempting to optimize policy choices.

In contrast to the MVPF, a more traditional benefit-cost ratio might place the net benefits to the government in the numerator relative to upfront government costs in the denominator (Heckman et al. 2010). Because the fiscal externalities are broadly quite small relative to programmatic cost, our conclusions would be similar if one were to use such a benefit-cost ratio for the analysis. While the results are similar, the clear conceptual advantage of the MVPF approach is that it does not require making assumptions about how the budget constraint is closed. As a result, the welfare conclusions do not depend on (often opaque) assumptions about the deadweight loss of income taxation or the “marginal cost of public funds”. This enables researchers to compare the desirability of wind PTC subsidies to spending on education, without worrying about the MCPF assumptions embedded in welfare analyses of the PTC and education studies. It also allows researchers to consider raising revenue from a gas tax instead of an income tax - indeed, rarely does one talk about the benefit-cost ratio of a gas tax. Instead, the MVPF provides a unified way of thinking about tax and spending policies. In doing so, it also provides a transparent method of incorporating preferences for equity in equation (5). An MVPF of 6 for wind PTCs vs. 1.5 for an income tax means we prefer the wind PTC if we want to give \$6 to its beneficiaries (roughly \$4 flows overseas but 2 goes to US residents).

Other work has also focused on constructing measures of the net benefits of policy changes, which can account for the fact that some policies might deliver greater benefits because of their differential scale. The MVPF of a policy provides a first order approximation to the net benefits of a policy by simply multiplying by the total spending of the policy. As we discuss in our section on EVs and the non-marginal analysis of the IRA, one can also integrate over the MVPF to get the average MVPF of a non-marginal change in spending.

A key advantage of our MVPF approach relative to traditional implementations of work measuring net benefits is that we do not impose ad hoc assumptions about how the budget constraint is closed. Researchers can compare MVPFs to decide how to close the budget constraint. The key idea behind the MVPF framework is that we can construct budget-neutral policy experiments for the decision-maker by comparing any two MVPFs. For example, if one policy costs the government one dollar and another raises revenue of one dollar, the two policies can be combined using their respective MVPFs to yield an expression for net benefits (it would equal the sum of the two MVPFs in this case). By constructing such budget-neutral policy experiments, MVPFs can be used to construct a benefit-cost analysis representing the sum of

willingness to pay across all individuals.

I Resource Cost per Ton

This appendix describes our approach to calculating the cost per ton for each policy in our sample. Because government cost per ton and net social cost per ton use the same inputs as the MVPFs, we defer readers to those appendices for details on the inputs for those calculations (and we provide a brief discussion at the end of this section on how we construct the government and net social cost per ton).

For most policies, the formula for resource cost follows: Difference in Sticker Prices + Difference in Use Costs where sticker prices are the upfront costs paid for vehicles or appliances for example and use costs are often the payments for fuels needed to power the item. In the following subsections, we detail the calculations for each specific policy or a general policy category when possible.

In cases where policies have potential learning-by-doing effects, we provide two measures of the resource cost per ton: one that includes LBD and one that does not. The presence of learning-by-doing presents an interesting conceptual question for how (if at all) to include them in the resource cost per ton calculation. When we include learning by doing effects, we follow the traditional resource cost approach that focuses on products not policies. Hence, our counterfactual experiment underlying the resource cost and number of abated tons involves a change in quantity of the product purchased today holding fixed all other current or future purchases. As we discuss below, this means that we include benefits from reductions in future costs of producing goods as a result of learning-by-doing, but we do not include subsequent changes in purchases of the good that these lower costs may induce (intuitively, these purchases would have their own resource cost and tons of carbon emitted).

I.1 Electric Vehicles

The resource cost is calculated as the difference in buying and fueling a battery electric vehicle (BEV) versus buying and fueling an internal combustion engine vehicle (ICEV). The difference in the price of a BEV in 2020 versus an ICEV comes from Vincentric's 2024 Electric Vehicle Cost of Ownership Analysis and is reported to be \$8,166 for 2023 models. Adjusting to 2020 dollars, we have \$6,937.08.

The cost of fueling a BEV is calculated as the present discounted value of the VMT in each year multiplied by the 2020 kWh per mile (0.29) multiplied by the average levelized cost of electricity (LCOE) (\$0.074/kWh) (details for calculating the LCOE can be found in Appendix C.2). This adds up to \$2,216. The resource cost of fueling an ICEV is, similarly, the PDV of the VMT in each year multiplied by the counterfactual MPG (41.23) in 2020 multiplied by the retail gasoline price (\$2.27) minus the gasoline tax (\$0.46) and markups (\$0.79). In total, this implies a lifetime gasoline cost of \$2,519. Overall, the resource cost without any learning-by-doing effects for a battery electric vehicle is \$6,634.

In the presence of learning-by-doing, the purchase of an EV today lowers future EV production prices and causes future EV purchases. Because resource cost measures typically focus on the resources associated with a given product, there is some ambiguity in how these costs

should be incorporated into resource costs and how the reduction in future tons of carbon emitted from induced new EV purchases should factor into the tons of carbon abated. Our interpretation of the resource cost approach is that it conceptualizes a counterfactual between a world with and without an additional purchase of the EV, holding all other purchases fixed. This means that the fact that purchasing an EV today lowers the cost of producing future EVs constitutes a reduction in real resources used in the economy to produce EVs. So, this dynamic cost of production benefit from learning-by-doing is included as a reduction in real resource costs. However, because we consider the counterfactual of one additional EV purchased today, we do not include the resource costs or tons of carbon abated from the induced EV purchases that occur as a result of learning-by-doing.

In practice, this means that to account for learning-by-doing, we take the resource cost and subtract off the reduction in future BEV production costs due to learning-by-doing. Since each BEV policy has a different elasticity, these price impacts vary across policies. For Clinton & Steinberg (2019), the dynamic price component is 0.564, which can be interpreted as a 56-cent reduction in future BEV prices for every dollar of mechanical subsidy. To use this effect in the resource cost per ton calculations, we convert to a per vehicle unit by dividing the component by the semi-elasticity. For Clinton & Steinberg (2019), the semi-elasticity is 0.0000549, so the per vehicle component is \$10,268. Since this is a future benefit, we subtract it from the existing resource cost estimate to get -\$3,634.

For Li et al. (2017), the dynamic price component is 0.482 and the semi-elasticity is 0.0000489, so the per vehicle component is \$9,854. Thus, the resource cost with learning-by-doing is -\$3,220.

Lastly, for Muehlegger & Rapson (2022), the dynamic price component is 0.309 and the semi-elasticity is 0.0000393. There is also an estimated pass-through of 85% in this paper, so the per vehicle component is \$9,245. Thus, the resource cost with learning-by-doing is -\$2,611.

To calculate the tons of carbon abated by purchasing a BEV, we take the carbon emissions from the ICEV lifetime gas consumption and subtract the carbon emissions from the BEV lifetime electricity consumption as well as the emissions from the production of BEV batteries. Details on the calculation of emissions from gasoline and the electricity grid can be found in Appendices C.4 and C.2, respectively. For ICEVs, we have 28.38 tons and for BEVs, we have 16.66 tons. The emissions from battery production are 59.5 kg per kWh of battery capacity. The average 2020 BEV battery capacity is 73 kWh. Thus, we have 4.34 tons of carbon from batteries. Overall, the tons of carbon abated from purchasing a BEV is 6.89 tons.

Taking the resource cost without learning-by-doing and dividing it by the tons of carbon, we have a resource cost per ton of \$962.70.

With learning-by-doing, the resource cost per ton for Clinton & Steinberg (2019) is -\$527.43, for Li et al. (2017) is -\$467.34, and for Muehlegger & Rapson (2022) is -\$378.96, with the differences arising because the strength of learning by doing effects depends on the magnitude of the price elasticity estimated in each paper.

I.2 Wind

For wind, we use utility-scale natural gas as the counterfactual since in 2021 it was the main source of new capacity to the grid coming from fossil fuels. The natural gas LCOE in 2020 is \$0.05/kWh and the wind one is \$0.033. The resource cost without any learning-by-doing effects

is simply the difference between these two LCOEs, which is $-\$0.0167$.

Just as with EVs, to account for learning-by-doing, we take the resource cost and subtract off the reduction in future wind production costs due to learning-by-doing. Since each wind policy has a different elasticity, these price impacts vary across policies. For Hitaj (2013), the dynamic price component is 0.455, which can be interpreted as a 46-cents reduction in future wind prices for every dollar of mechanical subsidy. To use this effect in the resource cost per ton calculations, we convert to a per kWh unit by dividing the component by the semi-elasticity for a \$1 change in price. For Hitaj (2013), the semi-elasticity is 21.27, so the per kWh component is $\$0.0214$. Since this is a future benefit, we subtract it from the existing resource cost estimate to get $-\$0.0047$.

For Metcalf (2010), the dynamic price component is 0.560 and the semi-elasticity is 24.45, so the per kWh component is $\$0.0229$. Thus, the resource cost with learning-by-doing is $-\$0.0062$.

Lastly, for Shrimali et al. (2015), the dynamic price component is 0.920 and the semi-elasticity is 32.82, so the per kWh component is $\$0.0280$. Thus, the resource cost with learning-by-doing is $-\$0.0113$.

The carbon amount is emissions from one kWh of natural gas minus the emissions from one kWh of wind energy. For natural gas, this is 0.0004074 tons and for wind, this is 0.000011. Thus, we have 0.0003964 tons of carbon abated per kWh of wind energy. Our final resource cost per ton number without learning-by-doing is $-\$42.24$.

With learning-by-doing, the resource cost per ton for Hitaj (2013) is $-\$11.86$, for Metcalf (2010) is $-\$15.64$, and for Shrimali et al. (2015) is $-\$28.51$.

I.3 Solar

Since all of the solar policies we analyze regard residential solar, we use the average energy mix from the grid as our counterfactual, meaning we use the $\$0.074/\text{kWh}$ average LCOE as in the BEV calculations. For the cost of a kWh of residential solar, we use the average cost per watt number from the National Renewable Energy Laboratory which is $\$2.77$ after adjusting the value to 2020\$. To convert this to a per kWh value, we divide it by the average lifetime of a solar system (25 years) and the average annual output from one watt (1.44 kWh) This gives us a per-kWh cost of $\$0.0769$. Thus, our resource cost without any learning-by-doing effects is $\$0.00291$.

To account for learning-by-doing, we take the resource cost and subtract off the reduction in future solar costs due to learning-by-doing. Since each solar policy has a different elasticity, these price impacts vary across policies. For Hughes & Podolefsky (2015), the dynamic price component is 3.99, which can be interpreted as a $\$3.99$ reduction in future solar prices for every dollar of mechanical subsidy. To use this effect in the resource cost per ton calculations, we convert to a per kWh unit by dividing the component by the semi-elasticity for a \$1 change in price. For Hughes & Podolefsky (2015), the semi-elasticity is 74.82, so the per kWh component is $\$0.0533$. Since this is a future benefit, we subtract it from the existing resource cost estimate to get $-\$0.0504$.

For Crago & Chernyakhovskiy (2017), the dynamic price component is 1.61 and the semi-elasticity is 21.23, so the per kWh component is $\$0.0758$. Thus, the resource cost with learning-by-doing is $-\$0.0729$.

For Pless & van Benthem (2019) with third-party owners, the dynamic price component is 1.37 and the semi-elasticity is 27.92, so the per kWh component is \$0.0491. Thus, the resource cost with learning-by-doing is -\$0.0462.

For Pless & van Benthem (2019) with host owners, the dynamic price component is 0.86 and the semi-elasticity is 16.22, so the per kWh component is \$0.0377. Thus, the resource cost with learning-by-doing is -\$0.0504.

Lastly, for Gillingham & Tsvetanov (2019), the dynamic price component is 0.35 and the semi-elasticity is 9.28, so the per kWh component is \$0.0373. Thus, the resource cost with learning-by-doing is -\$0.0343.

The carbon amount is emissions from one kWh of electricity from the grid using AVERT's model of the makeup of the grid that solar replaces minus the emissions from 1 kWh of solar electricity. The grid emissions are 0.0006968 and the solar emissions are 0.00004, which leaves us with 0.0006568 tons of carbon abated per kWh of solar electricity. Thus, our resource cost per ton number without learning-by-doing is \$4.43 per ton.

With learning-by-doing, the resource cost per ton for Hughes & Podolefsky (2015) is -\$76.72, for Crago & Chernyakhovskiy (2017) is -\$110.99, for Pless & van Benthem (2019) with third-party owners is -\$70.34, for Pless & van Benthem (2019) with host owners is -\$76.74 and for Gillingham & Tsvetanov (2019) is -\$52.29.

I.4 Appliance Rebates

For appliance rebates, the papers in our sample find varying reductions in energy usage when consumers move from non-Energy Star to Energy Star (ES) appliances. Thus, we calculate the resource cost per ton separately for each policy. In general, we calculate the resource cost as the sticker price minus the energy savings.

I.4.1 Cash for Appliances - Clothes Washers

To estimate the energy savings from purchasing an ES-rated clothes washer, we use the authors' reported difference between an ES and non-ES-rated clothes washer in 2010 as 201 kWh per year. We use this number for the kWh reduction in years 6-15 of the clothes washer's lifetime. For years 1-5, we compare the 2010 ES-rated clothes washer with a 2001 non-ES-rated clothes washer. The paper does not directly report this number. It does report the ratio of the rebate amount to the total lifetime reduction. Using this ratio, we calculate the kWh difference for the first five years of the ES-rated appliance to be 668 kWh per year. Taking the present discounted value of this energy consumption multiplied by the average LCOE (\$0.074/kWh), we get a lifetime energy cost savings of \$432.25.

The sticker price comes from Table 3 of Houde & Aldy (2017), which reports an ES manufacturer's suggested retail price of \$1,033 and a non-ES price of \$643. Taking the difference and inflation adjusting to 2020\$ (we assume the values are in 2011\$), we get \$448.82. Thus, the resource cost is \$16.57.

Using the same kWh numbers from above, we estimate the carbon abated from the 5,350 kWh saved over the clothes washer's lifetime using AVERT's reported marginal emissions coefficients and get 3.903 tons. Thus the resource cost per ton is \$4.24.

I.4.2 ENERGY STAR Rebate - Water Heaters

To estimate the energy savings we take the EIA's estimate for an average natural gas water heater in a four-person household of 22.7 MMBtu of natural gas (EIA 2018). An Energy Star water heater uses 8% less energy than a standard model (ENERGY STAR 2023). Therefore, we estimate that an ES-rated water heater saves 1.816 MMBtu per year. Consistent with the other appliance rebate MVPFs in our sample, we assume a lifetime of 15 years. Using the average Citygate price for natural gas in 2020 of \$3.56 per MMBtu, the lifetime energy savings is \$84.74.

The sticker price difference is calculated using the values in Table 1 and computing a weighted average across the four models within the standard and Energy Star categories. The average non-ES price is \$520.10 and the average ES price is \$862.75. Taking the difference and converting from 2012\$ to 2020\$, we have \$386.32, giving us a resource cost of \$301.58.

For the carbon abated, we have the 27.24 MMBtu of natural gas saved multiplied by the emissions from one MMBtu of 0.0531 (from the EPA) to get 1.45 tons of carbon. Thus, the resource cost per ton is \$209.

I.4.3 State-level ENERGY STAR Rebate - Clothes Washers

To estimate the energy savings we use the kWh difference from Houde & Aldy (2017) of 201 kWh (we prefer this value because is estimated closer to 2020 than the one reported by Datta & Gulati (2014)). Using the same 15-year lifespan, we have an energy savings of \$194.94. We use the same sticker price of \$448.82 from above, so the resource cost is \$253.88

For carbon, we have 3,015 kWh of electricity saved over the lifetime, which produces 1.49 tons of carbon. Thus, the resource cost per ton is \$169.92.

I.4.4 Cash for Appliances - Dishwashers

To estimate the energy savings from purchasing an ES-rated dishwasher, we use the authors' reported difference between an ES and a non-ES-rated dish washer in 2010 as 34 kWh per year. We use this number for the kWh reduction in years 6-15 of the dishwasher's lifetime. For years 1-5, we compare the 2010 ES-rated clothes washer with a 2001 non-ES-rated clothes washer. The paper does not directly report this number. It does report the ratio of the rebate amount to the total lifetime reduction. Using this ratio, we calculate the kWh difference for the first five years of the ES-rated appliance to be 234.5 kWh per year. Taking the present discounted value of this energy consumption multiplied by the average LCOE (\$0.074/kWh), we get a lifetime energy cost savings of \$98.03.

The sticker price comes from Table 3 of Houde & Aldy (2017), which reports an ES manufacturer's suggested retail price of \$764 and a non-ES price of \$624. Taking the difference and inflation adjusting to 2020\$ (we assume the values are in 2011\$), we get \$161.12. Thus, the resource cost is \$63.08.

Using the same kWh numbers from above, we estimate the carbon abated from the 1,512.5 kWh saved over the dishwasher's lifetime using AVERT's reported marginal emissions coefficients and get 0.91 tons. Thus, the resource cost per ton is \$69.08.

I.4.5 Cash for Appliances - Refrigerators

To estimate the energy savings from purchasing an ES-rated refrigerator, we use the authors' reported difference between an ES and a non-ES-rated refrigerator in 2010 as 65 kWh per year. We use this number for the kWh reduction in years 6-15 of the refrigerator's lifetime. For years 1-5, we compare the 2010 ES-rated refrigerator with a 2001 non-ES-rated refrigerator. The paper does not directly report this number. It does report the ratio of the rebate amount to the total lifetime reduction. Using this ratio, we calculate the kWh difference for the first five years of the ES-rated appliance to be 207.6 kWh per year. Taking the present discounted value of this energy consumption multiplied by the average LCOE (\$0.074/kWh), we get a lifetime energy cost savings of \$113.78.

The sticker price comes from Table 3 of Houde & Aldy (2017), which reports an ES manufacturer's suggested retail price of \$1,778 and a non-ES price of \$1,938. Taking the difference and inflation adjusting to 2020\$ (we assume the values are in 2011\$), we get -\$184.13. Thus, the resource cost is -\$297.92.

Using the same kWh numbers from above, we estimate the carbon abated from the 1,688 kWh saved over the fridge's lifetime using AVERT's reported marginal emissions coefficients and get 0.998 tons. Thus, the resource cost per ton is -\$298.42.

I.4.6 State-level ENERGY STAR Rebate - Refrigerators

To estimate the energy savings we use the kWh difference from Houde & Aldy (2017) of 65 kWh per year since this value is estimated closer to 2020 than the one reported in Datta & Gulati (2014). Using the same 15-year lifespan, we have an energy savings of \$63.04. We use the same sticker price of -\$184.13 from above, so the resource cost is -\$247.18.

For carbon, we have 975 kWh of electricity saved over the lifetime, which produces 0.48 tons of carbon. Thus, the resource cost per ton is -\$511.56.

I.4.7 State-level ENERGY STAR Rebate - Dishwashers

To estimate the energy savings we use the kWh difference from Houde & Aldy (2017) of 34 kWh per year since this value is estimated closer to 2020 than the one reported in Datta & Gulati (2014). Using the same 15-year lifespan, we have an energy savings of \$32.98. We use the same sticker price of \$161.12 from above, so the resource cost is \$128.14.

For carbon, we have 510 kWh of electricity saved over the lifetime, which produces 0.25 tons of carbon. Thus, the resource cost per ton is \$507.02.

I.4.8 California Energy Savings Assistance Program - Refrigerators

Blonz (2023) finds that 3,715 replacements were for qualified refrigerators compared to 1,261 for unqualified refrigerators. Therefore, about 75% of the replacements were for qualified fridges. The paper also finds that the people who qualified accelerated their replacement decisions by five years and those who should not have qualified accelerated their replacement decisions by six years. During this window, the paper estimates that the qualified refrigerators saved 73.45 kWh per month and the unqualified refrigerators saved 38.02 kWh per month. Since the paper

estimates the average change in purchase timing across all the beneficiaries, we assume that everyone is marginal to the policy and changes their decision by either 5 or 6 years depending on whether they should have qualified for the replacement. Consistent with the other appliance rebate policies, we assume that these appliances have a 15-year lifetime. Taking the present discounted value of this energy consumption multiplied by the average LCOE (\$0.074/kWh), we get a lifetime energy cost savings of \$290.59.

The sticker price comes from Table 3 of Houde & Aldy (2017), which reports an ES manufacturer's suggested retail price of \$1,778 and a non-ES price of \$1,938. Taking the difference and inflation adjusting to 2020\$ (we assume the values are in 2011\$), we get -\$184.13. Thus, the resource cost is -\$474.72.

Using the same kWh numbers from above, we estimate the carbon abated from the 3,984 kWh saved over the fridge's lifetime using AVERT's reported marginal emissions coefficients and get 2.94 tons. Thus, the resource cost per ton is -\$161.69.

I.5 Vehicle Retirement

Similarly to appliance rebates, we estimate the resource cost per ton separately for each vehicle retirement policy.

I.5.1 Cash for Clunkers (Li et al. 2013)

Li et al. (2013) find that the "Cash-for-Clunkers" policy had two effects: accelerating the purchase of a new car and shifting the new car to have a higher fuel economy than the consumer would have otherwise purchased. This creates three sources of resource cost: a leasing cost to quantify the acceleration of the purchase, an accounting of the cost of the increased MPG using the marginal cost of compliance for CAFE standards, and gas savings over the lifetime of the new car due to its higher MPG.

We calculate a weighted average of new and used vehicle prices in 2020 to use in the estimate of the leasing cost. According to CarGurus, the average used car price was \$27,409 and according to KBB, the average new car price was \$39,592. Using sales numbers from Statista, this gives a weighted average of \$30,643. The leasing cost is the interest over the seven months of acceleration. We use a 3% interest rate and get a leasing cost of \$536.25.

The cost of the increased MPG is the marginal cost of compliance, \$89.67 per MPG, multiplied by the difference in the new car's MPG and its counterfactual MPG, which is 2.2. Thus the cost is \$215.67.

Lastly, using the retail gasoline price net of the gasoline tax and markups as in the BEV calculations, the 2.2 MPG difference, and the lifetime of the vehicle, the present discounted value of the gas savings between the new car and its counterfactual is \$647.38. Thus, the resource cost is \$104.53.

The carbon number is the emissions saved from that difference in fuel economy over the lifetime of the car (see Appendix C.4 for details on the estimation of driving emissions), which is 7.43 tons. The resource cost per ton is \$14.07.

I.5.2 BAAQMD Vehicle Buyback Program

Sandler (2012) finds that the vehicle buyback program accelerated the purchase of a new car. This creates two sources of resource cost: a leasing cost to quantify the acceleration of the purchase and gas savings over the acceleration period due to the higher MPG of the new car compared to the retired car.

We calculate a weighted average of new and used vehicle prices in 2020 to use in the estimate of the leasing cost. According to CarGurus, the average used car price was \$27,409 and according to KBB, the average new car price was \$39,592. Using sales numbers from Statista, this gives a weighted average of \$30,643. The leasing cost is the interest over the 3.8 years of acceleration. We use a 3% interest rate and get a leasing cost of \$3,511.

Using the retail gasoline price net of the gasoline tax and markups as in the BEV calculations, the 2.68 MPG difference, and the 3.8-year acceleration, the present discounted value of the gas savings between the new car and the retired car is \$87.75. Thus, the resource cost is \$3,424.

The carbon number is the emissions saved from that difference in fuel economy over the 3.8 years (see Appendix C.4 for details on the estimation of driving emissions), which is 1.14 tons. The resource cost per ton is \$3,010.

I.5.3 Cash for Clunkers (Hoekstra et al. 2017)

Hoekstra et al. (2017) find that the “Cash-for-Clunkers” policy had two effects: accelerating the purchase of a new car and shifting the new car to have a higher fuel economy than the consumer would have otherwise purchased. This creates three sources of resource cost: a leasing cost to quantify the acceleration of the purchase, an accounting of the cost of the increased MPG using the marginal cost of compliance for CAFE standards, and gas savings over the lifetime of the new car due to its higher MPG.

We calculate a weighted average of new and used vehicle prices in 2020 to use in the estimate of the leasing cost. According to CarGurus, the average used car price was \$27,409 and according to KBB, the average new car price was \$39,592. Using sales numbers from Statista, this gives a weighted average of \$30,643. The leasing cost is the interest over the eight months of acceleration. We use a 3% interest rate and get a leasing cost of \$612.85.

The cost of the increased MPG is the marginal cost of compliance, \$89.67 per MPG (inflation-adjusted from 2014\$ to 2020\$), multiplied by the difference in the new car’s MPG and its counterfactual’s, which is 3.54. Thus the cost is \$347.

Lastly, using the retail gasoline price net of the gasoline tax and markups as in the BEV calculations, the 3.54 MPG difference, and the lifetime of the vehicle, the present discounted value of the gas savings between the new car and its counterfactual is \$976.27. Thus, the resource cost is -\$16.58.

The carbon number is the emissions saved from that difference in fuel economy over the lifetime of the car (see Appendix C.4 for details on the estimation of driving emissions), which is 11.23 tons. The resource cost per ton is -\$1.48.

I.6 Weatherization

For each weatherization policy, the resource cost per ton is the cost of the retrofits minus the energy savings all divided by the tons of carbon abated.

I.6.1 Energize Phoenix Program - Residential Buildings

Liang et al. (2018) found that the program reduces electricity consumption by 26%. The average baseline annual electricity usage for the 24 households before the energy upgrades was 14,350 kWh. This results in an annual reduction of approximately 3,740 kWh. Using the average LCOE of \$0.074 and assuming a weatherization lifetime of 20 years, the energy savings from the program are \$110,793.

The sticker price is the total spending from the program inflation-adjusted to 2020\$, \$192,590, so the resource cost is \$81,797.

Using the same kWh numbers from above, we estimate the carbon abated from the 1,795,200 kWh saved by the 24 households over the weatherization's lifetime using AVERT's reported marginal emissions coefficients and get 735 tons. Thus, the resource cost per ton is \$111.34.

I.6.2 Michigan Weatherization Assistance Program (WAP)

The average household in the paper's sample uses 79.44 MMBtu of natural gas and 7,543.65 kWh of electricity annually. The paper's main specification estimates that weatherization reduces natural gas consumption by 18.9% and electricity consumption by 9.5%. This translates into an annual 713 kWh and 14.5 MMBtu reduction. Given a 20-year lifetime, the average LCOE, and the Citygate natural gas price, the lifetime energy savings are \$1,742.

The sticker price is the total spending from the program inflation-adjusted to 2020\$, \$5,928, so the resource cost is \$4,184.

Using the same kWh and MMBtu numbers from above, we estimate the carbon abated from the 14,260 kWh and 290 MMBtu saved over the weatherization's lifetime using AVERT's reported marginal emissions coefficients for the grid and the EPA's emissions estimates and get 21.3 tons. Thus, the resource cost per ton is \$196.84.

I.6.3 Illinois Home Weatherization Assistance Program

The paper estimates the average treatment effect of IHWAP on the monthly change in electricity and natural gas consumption. Converting these estimates to annual changes, the average household in their sample reduces annual electricity consumption by 1,656 kWh and annual natural gas consumption by 19.48 MMBtu. Given a 34-year lifetime, the average LCOE, and the Citygate natural gas price, the lifetime energy savings are \$4,796.

The sticker price is the total spending from the program inflation-adjusted to 2020\$, \$10,196, so the resource cost is \$5,400.

Using the same kWh and MMBtu numbers from above, we estimate the carbon abated from the 56,304 kWh and 662.3 MMBtu saved over the weatherization's lifetime using AVERT's reported marginal emissions coefficients for the grid and the EPA's emissions estimates for

natural gas and get 53.5 tons. Thus, the resource cost per ton is \$100.89.

I.6.4 Gainesville Regional Utility LEEP Plus Program

Using household and time fixed effects, the paper finds that treated households reduce electricity consumption relative to control households by 7.1% following the weatherization. The average electricity usage of the households in their sample was 9,965.5 kWh per year, implying a reduction of 706.9 kWh. Using a 20-year lifetime and the average LCOE, the lifetime energy savings are \$872.

The sticker price is the total spending from the program inflation-adjusted to 2020\$, \$3,900, so the resource cost is \$3,028.

Using the same kWh number from above, we estimate the carbon abated from the 14,138 kWh saved over the weatherization's lifetime using AVERT's reported marginal emissions coefficients for the grid and get 5.79 tons. Thus, the resource cost per ton is \$523.

I.6.5 Wisconsin Energy Efficiency Retrofit Program

Using a randomized experiment and a structural model to evaluate two home energy retrofit programs, the paper finds that treated households reduced electricity consumption relative to control households by 1.142 kWh per day and reduced natural gas consumption by 0.396 MMBtu following the weatherization. Using a 20-year lifetime, the average LCOE, and the Citygate natural gas price, the lifetime energy savings are \$1,373.

The sticker price is the total spending from the program inflation-adjusted to 2020\$, \$2,096, so the resource cost is \$723.

Using the same kWh and MMBtu numbers from above, we estimate the carbon abated from the 8,336.6 kWh and 2,890.8 MMBtu saved over the weatherization's lifetime using AVERT's reported marginal emissions coefficients for the grid and the EPA's emissions estimates for natural gas and get 18.76 tons. Thus, the resource cost per ton is \$38.53.

I.7 Hybrid Vehicles

The resource cost is calculated as the difference in buying and fueling a hybrid electric vehicle (BEV) versus buying and fueling an internal combustion engine vehicle (ICEV). The prices of an HEV and an ICE in 2020 according to KBB are \$28,359 and \$27,012, respectively, so the difference is \$1,347.

The cost of fueling an HEV is calculated as the present discounted value of the VMT in each year multiplied by the 2020 average HEV fuel economy (42.52) multiplied by the retail gasoline price (\$2.27) minus the gasoline tax (\$0.46) and markups (\$0.79). This adds up to \$4,008. The cost of fueling an ICEV is, similarly, the PDV of the VMT in each year multiplied by the counterfactual MPG (40.62) in 2020 multiplied by the same gasoline cost. In total, this implies a lifetime gasoline cost of \$4,154. Overall, the resource cost without any learning-by-doing effects for a hybrid electric vehicle is \$1,200.

To account for learning-by-doing, we take the resource cost and subtract off the reduction in future battery costs due to learning-by-doing. Since each hybrid policy has a different elasticity,

these price impacts vary across policies. For Gallagher & Muehlegger (2011)'s sales tax waiver estimate, the dynamic price component is 0.031, which can be interpreted as a \$0.03 reduction in future hybrid prices for every dollar of mechanical subsidy. To use this effect in the resource cost per ton calculations, we convert to a per vehicle unit by dividing the component by the semi-elasticity for a \$1 change in price. For Gallagher & Muehlegger (2011), the semi-elasticity is 0.000207, so the per vehicle component is \$151.57. Since this is a future benefit, we subtract it from the existing resource cost estimate to get \$1,048.43.

For Beresteanu & Li (2011), the dynamic price component is 0.0089 and the semi-elasticity is 0.0000593, so the per vehicle component is \$149.13. Thus, the resource cost with learning-by-doing is \$1,050.87.

Lastly, for Gallagher & Muehlegger (2011)'s income tax credit estimate, the dynamic price component is 0.0019 and the semi-elasticity is 0.0000129, so the per vehicle component is \$148.38. Thus, the resource cost with learning-by-doing is \$1,051.62.

To calculate the tons of carbon abated by purchasing a HEV, we take the carbon emissions from the ICEV lifetime gas consumption and subtract the carbon emissions from the HEV lifetime gasoline consumption as well as the emissions from the production of HEV batteries. Details on the calculation of emissions from gasoline can be found in Appendix C.4. For ICEVs, we have 46 tons and for HEVs, we have 44 tons. The emissions from battery production are 234 kg per battery. Thus, we have 0.234 tons of carbon from batteries. Overall, the tons of carbon abated from purchasing an HEV is 1.82 tons.

Taking the resource cost and dividing it by the tons of carbon, we have a resource cost per ton without any learning-by-doing effects of \$659.

With learning-by-doing, the resource cost per ton for Gallagher & Muehlegger (2011)'s sales tax waiver estimate is \$576.06, for Beresteanu & Li (2011) is \$577.40, and for Gallagher & Muehlegger (2011)'s income tax credit estimate is \$577.81.

I.8 Home Energy Reports

For home energy reports, the papers in our sample find varying reductions in energy usage when consumers receive a report. Thus, we calculate the resource cost per ton separately for each policy. In general, we calculate the resource cost as the sticker price minus the energy savings.

I.8.1 Home Energy Reports (17 RCTs)

Across the 17 RCTs in this sample, the weighted average energy reduction is 243.26 kWh per household. Using the average LCOE, the energy savings are \$18. The Home Energy Report program costs \$8.83 per household on average, so the resource cost is -\$9.17.

Taking the 243.26 kWh and the marginal emissions coefficients from AVERT, the carbon abated per household is 0.1806 tons. Thus, the resource cost per ton is -\$50.76.

I.8.2 Opower Natural Gas Program Evaluations (52 RCTs)

Across the 52 RCTs in this sample, the weighted average natural gas reduction is 0.9416 MMBtu per household. Using the Citygate natural gas price, the energy savings are \$3.35. The Home Energy Report program costs \$9.96 per household on average, so the resource cost is \$6.61.

Taking the 0.9416 MMBtu and the emissions rate for natural gas from the EPA, the carbon abated per household is 0.05 tons. Thus, the resource cost per ton is \$132.

I.8.3 Peak Energy Reports

In this experiment, the average electricity reduction from receiving a PER is 0.1235 kWh. Assuming the LCOE at peak energy usage is \$1 per kWh, the energy savings are \$0.12. Each PER costs \$0.10, so the resource cost is -\$0.02.

Taking the 0.1235 kWh and assuming any energy reduction at peak usage is saving coal from being burned, the carbon abated per household is 0.0001213 tons. Thus, the resource cost per ton is -\$193.71.

I.8.4 Opower Electricity Program Evaluations (166 RCTs)

Across the 166 RCTs in this sample, the weighted average electricity reduction is 161 kWh per household. Using the average LCOE, the energy savings are \$11.89. The Home Energy Report program costs \$6.96 per household on average, so the resource cost is -\$4.93.

Taking the 161 kWh and the marginal emissions coefficients from AVERT, the carbon abated per household is 0.1194 tons. Thus, the resource cost per ton is -\$41.33.

I.9 Gasoline Taxes

For gasoline taxes, the resource cost is simply the retail gas price net of markups and taxes, which is \$1.02 per gallon. This is a savings though, so it is negative for our calculations. The carbon emissions from one gallon of gasoline are 0.009781 (details can be found in Appendix C.4). Thus, the resource cost per ton is -\$103.77.

I.10 Other Fuel Taxes

I.10.1 Tax on Jet Fuel

For a jet fuel tax, the resource cost is simply the retail jet fuel price net of markups and taxes, which is \$0.46 per gallon. This is a savings, though, so it is negative for our calculations. The carbon emissions from one gallon of jet fuel are 0.01085. Thus, the resource cost per ton is -\$42.27.

I.10.2 Tax on Diesel Fuel

For a diesel tax, the resource cost is simply the retail diesel price net of markups and taxes, which is \$1.12 per gallon. This is a savings, though, so it is negative for our calculations. The carbon emissions from one gallon of diesel are 0.01133. Thus, the resource cost per ton is -\$98.54.

I.11 Other Revenue Raisers

I.11.1 Critical Peak Pricing - Passive Joiners

At peak energy demand, we assume the LCOE is \$1 per kWh, so we take that as our resource cost. We also assume that at peak energy demand, the marginal kWh of electricity is coming from coal. One kWh of electricity produced solely with coal emits 0.0009823 tons of carbon. Thus, the resource cost per ton is -\$1,018.

I.11.2 California Alternate Rates for Energy

The resource cost is simply the citygate price for one MMBtu of natural gas, which is \$3.56. This is a savings though, so it is negative for our calculations. The carbon from one MMBtu is 0.0531 tons. Thus, the resource cost per ton is -\$67.06.

I.11.3 Critical Peak Pricing - Active Joiners

At peak energy demand we assume the LCOE is \$1 per kWh, so we take that as our resource cost. We also assume that at peak energy demand, the marginal kWh of electricity is coming from coal. One kWh of electricity produced solely with coal emits 0.0009823 tons of carbon. Thus, the resource cost per ton is -\$1,018.

I.12 Government Cost per Ton

As discussed in the main text, the government cost per ton measures the reduction in tons of CO_2 emitted per each dollar of net government outlay. The construction of the government cost per ton uses all of the same inputs as the MVPF, so we defer readers to the detailed appendix for the MVPF construction of each policy for information on how the numbers are constructed. Relative to the MVPF, it uses the denominator of the MVPF in its numerator (the net government cost of the policy) and compares this to the tons of carbon abated from the policy. To calculate the government cost per ton we take the Total Cost (see Table 2) of a policy and divide it by the sum of Global Environmental Benefits and the global portion of the Rebound Effect (including any portion captured by the climate FE) and divided by the social cost of carbon. While this doesn't account for the discount rate or the rising social cost of carbon, it is approximately equal to the tons of carbon from the policy.

If we are including the effect of learning by doing, then the denominator will be calculated by also including the global portion of the Learning by Doing Environmental Benefit.

I.13 Net Social Cost per Ton

Net social cost per ton is calculated as the ratio of the net government cost minus all of the non- CO_2 -related benefits of the policy and the abated tons of CO_2 . The construction of the net social cost per ton uses all of the same inputs as the MVPF, so we defer readers to the detailed appendix for the MVPF construction of each policy for information on how the numbers are constructed. The abated tons of carbon are calculated in the same way as for government cost per ton. For the numerator though, we take the Total Cost (again see Table 2) and subtract the Transfer, Profits, Local Environmental Benefits, and the local portion of the Rebound Effect.

If we are including the effect of learning by doing, then the numerator is calculated by also subtracting the Learning by Doing Price benefit. Again, the denominator is calculated in the exact same way as for government cost per ton.

J Federal Energy Policy over the last 15 years

There have been two main pieces of federal legislation over the last 15 years than have guided US energy policy: The American Recovery and Reinvestment Act (ARRA) enacted in 2009 and the Inflation Reduction Act (IRA), enacted in 2022. Here, we compare the relative spending in each Act for renewables, energy efficiency, and EVs.

J.1 ARRA

The aim of ARRA was to stimulate the economy following the Great Recession, which major objective being to create jobs, promote investment in infrastructure, and foster consumer spending. The energy component of ARRA aimed to modernize the energy sector, enhance energy efficiency, and promote renewable energy sources. Here, we break down the allocation of funding as part of the ARRA.

We draw the breakdown of funds in the ARRA from Table 1 from CEA (2016). We report values in 2009 prices, unless otherwise stated. The CEA reports that anticipated ARRA spending was \$90 Billion and that total spending was \$105 billion. We conservatively estimate that the ARRA spent \$49.8 billion on renewable technologies. This includes the \$26.6 billion that the CEA designated as renewable generation. That figure includes wind and solar production tax credits (PTCs) and investment tax credits (ITCs), as well as the 1603 Cash Grant program for renewables. To that \$26.6 billion we add \$3.5 billion for the Green Innovation & Job Training, \$3.4 billion for Carbon capture and Sequestration and \$2 billion for the State Energy Plan.²⁶⁴ The CEA (2016) also stated that total spending exceeded projected spending by \$15 billion. They cite the 1603 Cash Grant program and the clean energy manufacturing tax credit as sources of this cost overrun. In order to be conservative in our calculations, we assume that the full \$15 billion was allocated toward clean energy, although this is certainly an over-estimate. We also allocate a portion of Section 25 spending to the clean energy category. The program was dominated by Section 25C, which was focused on household energy efficiency, but we use estimates from the JCT to estimate the relative spending on Section 25C versus Section 25D (renewable generation) (Brown & Sherlock 2011). Assuming that 10% of total Section 25 spending went to clean energy, increases the total spending on clean energy by another \$1 billion.

We estimate that ARRA spending on energy efficiency spending was \$16.9 billion, of which was made up by weatherization, energy efficiency and conservation block grants, the energy efficiency tax credits of 25C, and state energy plan (CEA 2016, Goldman 2011).²⁶⁵

The remaining portion of ARRA spending is as follows: Transit had the next largest amount of investment, with \$18.1 billion. This was focused more on infrastructure, such as high-speed rail, but not on EVs. Next was grid modernization at \$10.5 billion, which focused on making the grid more efficient, with a great deal of spending on smart meters and technology (not renewables). Spending on advanced vehicles was \$6.1 billion, which focused on EV and battery subsidies.

²⁶⁴We omit the Clean Energy Equipment Manufacturing \$1.6 bn line item from renewable generation. This is consistent with our omission of advanced manufacturing spending for this calculation in the IRA

²⁶⁵While the CEA estimates this category as \$19.9 billion, we subtract \$2 billion, one for SEP and \$1 billion for section 25D tax credits.

Given these numbers, we calculate that subsidies (both grants and tax credits) for clean renewable energy were about 3 times those for energy efficiency. Subsidies for clean energy and energy efficiency were 8.2 times and 2.8 times larger the spending on EVs, respectively.

For the purposes of comparison to the IRA in the table below we also inflation our estimates of spending levels. Spending on clean energy was \$67.9 billion in 2022 dollars. Spending on energy efficiency was \$23 billion. Spending on EVs was \$8.3 billion.

J.2 IRA

The IRA aimed at addressing various economic and environmental issue in the US, such as reducing inflation, lowering healthcare costs, and investing in clean energy and climate change mitigation. Here, we focus on two major sources estimating realized IRA spending: reports by the Penn Wharton Budget Model (PWBM 2023) and Goldman Sachs (Della Vigna et al. 2023). We use the estimates from Goldman Sachs as our default comparison, but also report the robustness of our results to the estimates from the Penn Wharton Budget Model.

Estimates from (PWBM 2023) suggest that, by 2032, the IRA will lead to the following amounts of spending. Estimated subsidies that will be spent by are \$263 billion for clean renewable energy, \$393 billion for EVs, and \$28 billion for energy efficiency. Based on these numbers, subsidies for clean energy (excluding advanced manufacturing) are roughly 9.4 times those for energy efficiency. However, subsidies for EVs are 1.5 and 14 times the spending on clean energy and energy efficiency, respectively. These estimates are relatively similar to the figures from Goldman Sachs, who suggest that spending on clean energy versus energy efficiency is \$274 billion versus \$44 billion, a ratio of more than 6:1. (They estimate \$393 billion for EVs spending, the same figure as above.)

	ARRA Spending (2022 Prices)	IRA Spending			MVPPF (Our Estimates)
		CBO Estimate	Goldman Sachs	Penn Wharton Model	
Clean Energy	\$67.9bn	\$192bn	\$274bn	\$263bn	Wind – 5.87 Solar – 3.86
EVs	\$8.3bn	\$14bn	\$393bn	\$393bn	1.45
Energy Efficiency	\$23.0bn	\$2bn	\$44bn	\$28bn	~1

Note: ARRA numbers are in 2022 prices.

Interestingly, these same basic patterns can also be seen when comparing expected spending rather than realized spending. If we eliminate the \$15 billion cost overrun from the ARRA, we find that spending on clean energy relative to energy efficiency is 2:1. If we use the original CBO estimates of the IRA, we see that \$192 billion for clean renewable energy, \$14 billion for EVs, and \$2 billion for energy efficiency. That suggests a ratio an order of magnitude higher. Interestingly, this also suggests that expected EV spending relative to clean energy spending was lower under the IRA than in the ARRA. It is only the realized spending figures that reversed that pattern.

We have also excluded credits for advanced manufacturing from these calculations. They were expected to be \$37 billion under the IRA and are now projected to be \$193 billion according

to Goldman Sachs. If these values were included in our estimates of IRA spending on clean energy, it would only increase the relative spending on clean energy as ARRA spending on advanced manufacturing subsidies was far smaller by comparison (and is already included in part in the \$15 billion in cost overruns currently allocated to clean energy production.)