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 adjust our estimates for rebound effects. Appendix E (still in progress) 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).

A Model Appendix: Setup

The MVPF framework requires measuring the willingness-to-pay for each group in society along with the net costs 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 (Acemoglu et al. 2012, 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.[153]

[153]We 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

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 $\mathbf{x}$.[154] 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 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 $u_i(\mathbf{x;e})$ subject to a budget constraint, $\mathbf{p*x} \leq m_i$, 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:

$$WTP_i = \mathbf{x}_i * \mathbf{dp} + \mathbf{v}_i * \mathbf{de} + (1 - \tau_c))s_i d\Pi. \tag{35}$$

There are three reasons that individuals are willing to pay for a policy change: (a) it makes the goods they consume cheaper, $\mathbf{x}_i * \mathbf{dp}$, where $\mathbf{dp}$ is the causal effect of the policy on prices; (b) it changes the value of the externalities they experience, $\mathbf{v}_i * \mathbf{de}$, where $\mathbf{de}$ 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 $\Pi(\mathbf{x}) = \mathbf{x}\boldsymbol{*}(\mathbf{q} - \mathbf{c}(\mathbf{x}))$. The policy impact on firm profits is:

$$d\Pi = \mathbf{dx}\boldsymbol{*}(\mathbf{q}\text{-}\mathbf{c}(\mathbf{x})) + \mathbf{x}\boldsymbol{*}(\mathbf{dq} - \nabla\mathbf{c}(\mathbf{x}) \cdot \mathbf{dx}) \tag{36}$$

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, $\mathbf{dq}$, this increases (decreases) firm profits. If the policy increases firm costs, this reduces firm

costs).

[154]We 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

profits proportional to the Hadamard product of the derivatives of the cost function, $\triangledown \mathbf{c}$, and the change in production of each good, $\mathbf{dx}$. 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 $\triangledown \mathbf{E}$ 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:

$$\mathbf{v}_i * \mathbf{de} = \mathbf{v}_i * \triangledown \mathbf{E} * \mathbf{dx} \tag{37}$$

where $\mathbf{v}_i * \triangledown E$ 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 37 is summing across all the possible externalities experienced by individual i. This means we allow for individuals to experience externalities very differently.[155] In implementation, we will 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 suggest 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,[156] this suggests government tax revenue declines by \$.045SCC 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

$$Cost = \mathbf{x} * \mathbf{dt} + \mathbf{t} * \mathbf{dx} + \tau^c(\mathbf{dx}\text{*}(\mathbf{q}\text{-}\mathbf{c}(\mathbf{x})) + \mathbf{x}\text{*}(\mathbf{dq} - \triangledown \mathbf{c}(\mathbf{x}) \cdot \mathbf{dx})) \tag{38}$$

This is equivalent to the sum of the mechanical cost of any change to the subsidies or taxes ($\mathbf{x} * \mathbf{dt}$), the impact of the behavioral response on the cost of subsidies ($\mathbf{t} * \mathbf{dx}$), and the impact

[155]For 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.

[156]This further assumes government and private discount rates are equivalent

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 35 and 38 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 $\mathbf{dx}$. 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 "$\mathbf{dx}$" 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" 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. The contribution of this section is to provide a new sufficient statistics result (Theorem 1, introduced in the main text and stated precisely in a generalized form below) that translates cost 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, we care not only about the direct effects (e.g., the $\nabla\mathbf{c}(\mathbf{x})$ above), but also the 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.'[157] 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 $dx(t)$.[158] Formally, $dx(t)$ is a (Fréchet) differential of the time path of consumption of $x(t)$ with respect to the subsidy change, t^*, and η.[159] 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 $dp = -\gamma d\tau$. In our baseline specification, $\gamma = 1$, 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.[160]

In addition to environmental externalities, the subsidy also can affect firm profits in a non-competitive environment. However, introducing strategic behavior on the part of firms introduces substantial complexity in a dynamic environment with learning by doing. As a parsimonious alternative, 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

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

[158] We 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.

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

[160] In 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, $\nu_i(t) = \mathbf{v_i} \triangledown \mathbf{Edx(t)}$

implementations where relevant. [161]

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

$$MVPF = \frac{1+SE}{1+FE} \tag{39}$$

where

$$SE = \frac{\nu(0)}{p(0)}\epsilon \tag{40}$$

is the static externality benefit from additional consumption of x and

$$FE = \frac{\tau(0)}{p(0)}\epsilon j \tag{41}$$

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 - this is the $\triangledown \mathbf{c}(\mathbf{x})$ term in equation 13. 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

$$DP_i = \int_0^\infty -x_i(t)\frac{d[p(X(t))]}{x(0)\eta d\tau}e^{-\rho(t)}dt \tag{42}$$

[161] Adding 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 "DP" stands for the dynamic price reduction generated from the response. [162] Meanwhile, firms have a willingness-to-pay of

$$D\pi = \int_0^\infty \frac{d[\pi(X(t), x(t))]}{x(0)\eta d\tau} e^{-\rho(t)} dt \tag{43}$$

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:

$$DE_i = \int_0^\infty [\frac{dx(t)}{x(0)\eta d\tau} \nu_i(t)] e^{-\rho(t)} dt \tag{44}$$

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 $\nu_g(t)$.[163] The PDV impact on government costs is then

$$DFE = \int_0^\infty \frac{dx(t)}{x(0)\eta d\tau} \nu_g(t) e^{-\rho(t)} dt \tag{45}$$

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:

$$MVPF = \frac{1 + SE + DP + DE}{1 + FE + DFE} \tag{46}$$

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

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

[163] This term is given by the impact of the policy today on future consumption of goods in the economy, 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 $t > 0$.

how the production paths change with a shock to the initial conditions.[164]

What remains then is simply 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

$$dX(0) \approx -\gamma\eta\epsilon\frac{d\tau}{p(0)}x\left(0\right) \tag{47}$$

where.[165] 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 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 $\eta \to 0$ 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

$$dt = -\gamma\eta\epsilon\frac{d\tau}{p(0)} \tag{48}$$

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 $c\left(X(t)\right)$, 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 $\frac{d}{dt}c\left(X\left(t\right)\right)dt = c_X(X(t))x(t)dt$.

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

$$\frac{d[p(X(t)]}{\eta d\tau} = (\mu+1)\frac{d[c(X(t)]}{\eta d\tau} = -(\mu+1)\gamma\frac{\epsilon}{p(0)}c_X(X(t))x(t) \tag{49}$$

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 49 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:

[164] This 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 $X(0) = X_0$.

[165] To 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 $\epsilon\frac{x(t)}{p(t)} = \frac{dx(t)}{dp(t)}$

$$DP = (\mu+1)\int_0^\infty \frac{x(t)}{x(0)}\lim_{\eta\to 0}\left[\frac{-d[c(X(t))]}{\eta d\tau}\right]e^{-\rho t}dt \tag{50}$$

$$= (\mu+1)\gamma\int_0^\infty \frac{x(t)}{x(0)}\left[\epsilon\frac{x(t)}{p(0)}c_X(X(t))\right]e^{-\rho t}dt \tag{51}$$

$$= (\mu+1)\gamma\int_0^\infty e^{-\rho t}\epsilon\left(\frac{c_X\left(X(t)\right)}{c\left(X(t)\right)}X(t)\right)\frac{c(X(t))}{c(X(0))}\frac{x(t)}{X(t)}\frac{x(t)}{x(0)}dt \tag{52}$$

The first line passes the limit to the variables that depend on η ($dp(t)$ and η), the second line plugs in equation 49, 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 $d\pi = \mu\left(dc(t)x(t) + dx(t)c(t)\right)$, therefore

$$D\pi = \mu\left(\int_0^\infty \frac{x(t)}{x(0)}\lim_{\eta\to 0}\left[\frac{dc(t)}{\eta d\tau}\right]e^{-\rho t}dt + \int_0^\infty \frac{c(t)}{x(0)}\lim_{\eta\to 0}\left[\frac{dx(t)}{\eta d\tau}\right]e^{-\rho t}dt\right) \tag{53}$$

$$= -\mu\gamma\left(\int_0^\infty \frac{x(t)}{x(0)}\frac{c_X(X(t))x(t)}{p(0)}\epsilon e^{-\rho t}dt + \int_0^\infty \frac{c(t)}{x(0)}\frac{\epsilon x'(t)}{p(0)}e^{-\rho t}dt\right) \tag{54}$$

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:

$$dx(t) = -\gamma\eta x'(t)\epsilon\frac{d\tau}{p(0)} \tag{55}$$

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

$$DE = \int_0^\infty \nu(t)\lim_{\eta\to 0}\left[\frac{dx(t)}{x(0)\eta d\tau}\right]e^{-\rho t}dt \tag{56}$$

$$= \int_0^\infty \frac{\epsilon\nu(t)x'(t)}{x(0)c(X\left(0\right))}e^{-\rho t}dt \tag{57}$$

where the last line both substitutes equation 55 and uses the assumption that subsidies go away at $t = 0$ so that $c\left(X\left(0\right)\right) = p\left(0\right)$. 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*

$$MVPF = \frac{1 + SE + DP + DE}{1 + FE + DFE} \tag{58}$$

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 isoleastic specification:

$$x\left(p(t)\right) = ap(t)^{\epsilon} \tag{59}$$

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:

$$c\left(X\right) = \kappa X^{\theta} \tag{60}$$

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

$$\frac{d}{dt}\log(x) = \epsilon\frac{d}{dt}\log(p) = \frac{d}{dt}\frac{d\log\left(c\left(X\right)\right)}{d\log\left(X\right)}\frac{d\log\left(X\right)}{dt} = \epsilon\theta\frac{x(t)}{X(t)} \tag{61}$$

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):

$$\frac{X''(t)}{X'(t)} = \epsilon\theta\frac{X'(t)}{X(t)} \tag{62}$$

Equation 62 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. [166] This yields a general closed-form solution for $X(t)$ given by

$$X(t) = C_1(t + C_2)^{\frac{1}{1-\epsilon\theta}} \tag{63}$$

[166] Given 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.

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

$$x(t) = \frac{C_1}{1-\epsilon\theta}(t+C_2)^{\frac{\epsilon\theta}{1-\epsilon\theta}} \tag{64}$$

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 59 and the marginal cost is given by equation 60. Then,*

$$DP = \gamma(\mu+1)\frac{\theta\epsilon}{1-\theta\epsilon}C_2^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}}\int_0^\infty e^{-\rho t}(t+C_2)^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}}dt \tag{65}$$

where

$$C_2 = \frac{X(0)}{x(0)(1-\epsilon\theta)} \tag{66}$$

while

$$DE = -\frac{\gamma\theta\epsilon^2}{(1-\epsilon\theta)c(X(0))}C_2^{-\frac{\epsilon\theta}{1-\epsilon\theta}}\int_0^\infty e^{-\rho t}(t+C_2)^{\frac{2\epsilon\theta-1}{1-\epsilon\theta}}\nu(t)dt \tag{67}$$

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-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, $\frac{X(0)}{x(0)(1-\epsilon\theta)}$. Moreover, we can also do some comparative statics to establish some intuitions of how these dynamic components change with these parameters.

Theorem 2 *(Comparative statics)*
In our isoelastic setting, we have (a) $\lim_{C_2\to\infty} DP = 0$, and there exists $\bar{C}$ such that DP is strictly decreasing for C_2 greater than $\overline{C}$.
(b) For C_2 greater than some cutoff , $DP = 0$ is strictly decreasing in ϵ and θ.
(c) DP is always strictly increasing in γ (the pass through rate of the subsidy).

Part (a) We prove decreasingness first. For reference, we reproduce the expression for

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

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

$\frac{d}{dC_2}\left(C_2^{-\theta\frac{1+\epsilon}{1-\theta\epsilon}}\int_0^\infty e^{-\rho t}(t+C_2)^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}}dt\right)<0.$ [167]

The derivative of interest is

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

which has the same sign as

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

The first term, A, is clearly negative. For the second term, B, note that $\int_0^\infty e^{-\rho t}(t+C_2)^{-2+\theta\frac{1+\epsilon}{1-\epsilon\theta}}dt \le C_2^{-1}\int_0^\infty e^{-\rho t}(t+C_2)^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}}dt$. 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 negligable relative to A, implying there exists a cutoff $\overline{C}$ above which the statement holds. To see this, consider

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

where the RHS follows after canceling out common terms $e^{\rho C_2}$. By the preceeding 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 Liebniz' rule for differentiation under the integral sign yields

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

This again yields an indeterminate form of $\frac{0}{0}$, but, noting that $\int_{C_2}^\infty e^{-\rho t}t^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}}dt \le C_2^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}}\int_{C_2}^\infty e^{-\rho t}dt =$

[167] Note 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.

$\frac{e^{-\rho C_2}}{\rho} C_2^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}}$ (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

$$\lim_{C_2\to\infty} \frac{\frac{e^{-\rho C_2}}{\rho} C_2^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}}}{C_2^2 e^{-\rho C_2} C_2^{-2+\theta\frac{1+\epsilon}{1-\epsilon\theta}}}$$

$$= \lim_{C_2\to\infty} \frac{1}{\rho C_2} = 0$$

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

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

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

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

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 $\int_{C_2}^\infty e^{-\rho t} t^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}} \le C_2^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}} \frac{1}{\rho} e^{-\rho C_2}$ 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: $\int_x^\infty e^{-z} z^{a-1} dv \le a e^{-x} x^{a-1}, \forall a \in \mathbb{N}\backslash\{0\}$ and $x > a$. This is proven via induction on a. In the base case, start with $a = 1$. We have $\int_x^\infty e^{-z} z^{a-1} dv = \int_x^\infty e^{-z} dz = e^{-x} = a e^{-x} x^{a-1}$.

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

$$\int_x^\infty e^{-z} z^{a-1} = -e^{-z} z^{a-1}|_x^\infty + \int_x^\infty (a-1) e^{-z} z^{a-2} dz.$$

$$= e^{-x} x^{a-1} + (a-1)\int_x^\infty e^{-z} z^{a-2} dz$$

$$\le e^{-x} x^{a-1} + (a-1)(a-1) e^{-x} x^{a-2}$$

by the inductive hypothesis.

For$x > a-1$, this obeys

$$\le e^{-x} x^{a-1} + (a-1) x e^{-x} x^{a-2}$$

$$= e^{-x} x^{a-1} + (a-1) e^{-x} x^{a-1}$$

$$= ae^{-x}x^{a-1},$$

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 of part a.

Part (b)

Again we reproduce

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

We start with the comparative statics with respect to θ.

$$\begin{aligned}
\frac{\partial DP}{\partial \theta} &= \frac{\theta\epsilon}{1-\epsilon\theta}\gamma(\mu+1)\frac{1+\epsilon}{(1-\epsilon\theta)^2} C_2^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}} \int_{C_2}^{\infty} e^{-\rho(t-C_2)} t^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} \ln(t) dt \\
&- \frac{\theta\epsilon}{1-\epsilon\theta}\gamma(\mu+1)\frac{1+\epsilon}{(1-\epsilon\theta)^2} \ln(C_2)(1+\epsilon) C_2^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}} \int_{C_2}^{\infty} e^{-\rho(t-C_2)} t^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} dt \\
&+ \frac{\epsilon}{(1-\epsilon\theta)^2}\gamma(\mu+1) C_2^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}} \int_{C_2}^{\infty} e^{-\rho(t-C_2)} t^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} dt
\end{aligned}$$

which has the same sign as

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

which equals

$$-\theta\frac{1+\epsilon}{1-\epsilon\theta} C_2^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}} \int_{C_2}^{\infty} e^{-\rho(t-C_2)} t^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} \ln(\frac{t}{C_2}) - C_2^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}} \int_{C_2}^{\infty} e^{-\rho(t-C_2)} t^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} dt$$

Note that if $\theta\frac{1+\epsilon}{1-\epsilon\theta} > 0$, both terms are negative, concluding the proof.

If $\theta\frac{1+\epsilon}{1-\epsilon\theta} < 0$, the first term is positive while the second term is negative. Again, the proof amounts to showing that the first term is asymptotically negligable relative to the second.

To see, note that

$$\lim_{C_2\to\infty} \frac{C_2^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}} \int_{C_2}^{\infty} e^{-\rho(t-C_2)} t^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} \ln(\frac{t}{C_2}) dt}{C_2^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}} \int_{C_2}^{\infty} e^{-\rho(t-C_2)} t^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} dt}$$

$$= \lim_{C_2\to\infty} \frac{\int_{C_2}^{\infty} e^{-\rho t} t^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} \ln(\frac{t}{C_2}) dt}{\int_{C_2}^{\infty} e^{-\rho t} t^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} dt}$$

$$= \lim_{C_2\to\infty} \frac{-e^{-\rho C_2} C_2^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} \ln(C_2) + e^{-\rho C_2} C_2^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} \ln(C_2) - \frac{1}{C_2}\int_{C_2}^{\infty} e^{-\rho t} t^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} dt}{-e^{-\rho C_2} C_2^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})}}$$

$$= \lim_{C_2\to\infty} \frac{\int_{C_2}^{\infty} e^{-\rho t} t^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} dt}{e^{-\rho C_2} C_2^{\theta\frac{1+\epsilon}{1-\epsilon\theta}}}$$

$$\leq \lim_{C_2\to\infty} \frac{C_2^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} \int_{C_2}^{\infty} e^{-\rho t} dt}{e^{-\rho C_2} C_2^{\theta\frac{1+\epsilon}{1-\epsilon\theta}}}$$

$$= \lim_{C_2\to\infty} \frac{C_2^{(-1+\theta\frac{1+\epsilon}{1-\epsilon\theta})} e^{-\rho C_2}}{\rho e^{-\rho C_2} C_2^{\theta\frac{1+\epsilon}{1-\epsilon\theta}}}$$

$$= \lim_{C_2\to\infty} \frac{C_2^{-1}}{\rho} = 0.$$

We use a similar approach to sign the derivative with respect to ϵ. We get

$$\begin{aligned}
\frac{\partial DP}{\partial \epsilon} = & \frac{\epsilon\theta}{1-\epsilon\theta}\theta\frac{1+\theta}{(1-\epsilon\theta)^2}\gamma(\mu+1) C_2^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}} \int_{C_2}^{\infty} e^{-\rho(t-C_2)} t^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}} \ln(t) dt \\
& - \frac{\epsilon\theta}{1-\epsilon\theta}\theta\frac{1+\theta}{(1-\epsilon\theta)^2}\gamma(\mu+1) \left(C_2\right)^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}} \ln(C_2) \int_{C_2}^{\infty} e^{-\rho(t-C_2)} t^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}} dt \\
& + \frac{\theta}{(1-\epsilon\theta)^2}\gamma(\mu+1) \left(C_2\right)^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}} \int_{C_2}^{\infty} e^{-\rho(t-C_2)} t^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}} dt
\end{aligned}$$

which has the same sign as

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

The same argument as above establishes that the first two terms become negligible relative to the third (which is negative) as C_2 grows large.

Part c

Simply differentiating the WTP expression with respect to the pass through rate γ yields

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

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. .

C.1 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 environmental externalities (both GHGs and local pollutants) 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.

C.1.1 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 2024*b*). [168] 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 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 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 renewables may still have a high marginal emissions rate if natural gas is the marginal generation source.

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.[169] This is mostly trivial; a state is generally

[168] Consistent with vehicle externalities, NO_x, $PM_{2.5}$, and SO_2 are local pollutants and CO_2 is a global pollutant

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

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 for solar, wind, and energy efficiency programs is \$0.149, \$0.145, and \$0.159, respectively.[170]

C.1.2 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) 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 CO2e output emissions rates from EPA's eGRID, coal produces 2181 pounds of CO2e 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.

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

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.1.3 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. n.d., 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 2023*a*). 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. [171] 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 2023). [172] For the price of electricity, we use annual data on the retail price of electricity by state from the the BLS (BLS 2024). [173]

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:

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

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 cost through profit tax revenue from utilities. 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:

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

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 $\big(\overline{LCOE} + c_{td}\big)/p$ is constant over time and use retail prices from the BLS.

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

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

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

C.1.4 Rebound

For policies that cause an exogenous shock to electricity demand or supply, it is natural to expect that the shock's impact on the supply or demand curve will lead to a change in prices, and therefore quantity of electricity consumed. This so-called "rebound effect" is generally not captured by the reduced form elasticities estimated in the papers we consider. We therefore use a canonical supply and demand framework combined with estimates of the supply and demand elasticities for electricity to think about and include a correction for this rebound effect in our estimates.

Let $Q(p)$ and $S(p)$ denote the quantity demanded and supplied at price p. We now imagine there is a shock that expands the supply of electricity by C (e.g., a new wind turbine). Unless electricity demand is perfectly inelastic or electricity supply is perfect elastic, we would expect that an expansion in the supply of electricity as a result of a subsidy would not lead to a 1-1 reduction in dirty energy. We can model a positive electricity shock (C) to supply S(p) from policy such as a wind PTC or solar ITC as being given by:

$$Q(p) = S(p) - C$$

Total differentiation gives us a change in price:

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

which means the total change in the production of electricity is given by

$$S'dp = -dC * R$$

where R is defined as

$$R = \frac{1}{1 - \frac{Q'}{S'}}$$

Transforming this to elasticities, we have

$$\begin{aligned} \frac{Q'}{S'} &= \frac{Q'}{Q}\frac{S}{S'}\frac{Q}{S} \\ &= \frac{\epsilon_Q^{demand}}{\epsilon_S^{supply}} \end{aligned}$$

The rebound effect, defined as the percent of the supply shock that is offset by an increase in demand, is 1 - R. A higher rebound effect means that electricity price reductions from wind PTCs will lead to significant increases in the quantity of electricity demanded, offsetting the environmental impact of the initial PTC.

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 2023*c*). 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 9 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 estimates a residential price elasticity of -0.27. The EIA's 2020 Annual Energy Outlook 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%.

C.2 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.2.1 Environmental Externalities

We assume 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.

Following the approach in Appendix Section 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 Auffhammer & Rubin (2018). These elasticities lead to a natural gas rebound effect of 11.76%.

C.2.2 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 (2023*e*) 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 2021*a*). 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 2020*b*). Therefore, the producer profit and fiscal externality per MMBtu is \$4.40 and \$0.75, respectively.

C.3 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.[174] In this section, we discuss our approach to measuring the dollar value of the externalities generated by these vehicles. We provide a table (Appendix Table 12) containing the values of specific externalities in 2020, as well as a figure (Appendix Figure 2) that 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.

We consider two types of externalities: pollution and driving externalities. The driving externalities we consider are accidents, congestion, and pollution from tire and brake wear.[175] Section C.3.1 presents our calculations for pollution externalities. Section C.3.2 presents our calculations for accidents and congestion. Section C.3.3 consolidates these externalities into the average externality for a gallon of gas and discusses how our results vary as a function of time. Unless otherwise noted, all dollar values are in 2020 dollars.

[174] The EPA's definition of light-duty vehicles includes two regulatory classes: passenger cars and light trucks (EPA 2023*b*). 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.

[175] In our review of the literature, other driving externalities that have been measured seem to be negligible for light-duty vehicles. Noise pollution, for example, is much larger on a per-mile basis for buses and heavy-duty trucks. The same holds for road damage, with heavy-duty trucks appearing to cause the most damage per mile (Davis 2017, FHWA 1997).

C.3.1 Per-Gallon Externalities

The use of gas-powered vehicles results in pollution from a range of sources. We consider two categories of emissions in developing monetized estimates of pollution externalities from light-duty vehicles. The first includes emissions that result from producing a gallon of gasoline. We call these "upstream" emissions. The second includes emissions released during on-road driving. We ignore emissions from idling.

Upstream Pollution A vehicle's 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.[176] 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 for a gallon of gasoline 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

$$Upstream_{y,p,s} = \frac{Pollution_{y,p,s}}{RefineryYield_y} \tag{68}$$

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 (EIA 2024*f*,*e*).[177] In 2020, one barrel of crude oil produced on average 44.3 gallons of petroleum product.[178] The national refinery yield has remained roughly constant since the EIA began tracking refinery production data. The following paragraphs explain how we obtain values for the pollution emitted from a barrel of crude, $Pollution_{y,p,s}$.

[176]Since we consider petroleum extracted outside of the US, 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.

[177]The 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. 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.

[178]Output data come from the EIA's "U.S. Refinery and Blender Net Production of Crude Oil and Petroleum Products (Thousand Barrels)" series. 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 experience a processing gain and produce more than 42 gallons of product from one barrel of crude (EIA 2024*g*).

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.[179] 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.[180] We then divide the share of total CO_2e attributable to CH_4 and N_2O by the global warming potential (GWP) factors used by the authors to convert grams of non-CO_2 pollutant, p, 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.[181]

We then consider the 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). 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). 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 2024*e*). We then divide pollution released per barrel of crude oil by the refinery yield.[182] We again apply each pollutant's corresponding social cost to value emissions in dollars. 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, 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 in year y by the share of fuel ethanol in finished motor gasoline. We calculate this share using the approach outlined by the EIA (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

[179] This 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 2023*d*). For policies that target crude oil production in specific countries, we rely on the authors' country-specific carbon intensity measurements.

[180] We 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 as the authors we leave this pollutant in terms of CO_2e.

[181] Since the social cost of non-CO_2 pollutant, p, is roughly equal to the social cost of carbon scaled by the GWP factor of pollutant, p, 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.

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

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. The quantity of fuel ethanol supplied comes from the EIA's Monthly Energy Review (Table 10.3). For 2020, we multiply all upstream emissions by 0.95 to account for the 4.9% of ethanol in gasoline. We describe below how we adjust on-road emissions for the share of ethanol in gasoline.

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). We allow the carbon intensity of ethanol production to vary over time.[183] 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 carbon dioxide emissions increase from \$0.18 to \$0.22 per gallon.

On-Road Pollution Next, we consider emissions released while a vehicle is in use. For each pollutant, 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 across model years to measure the average 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.[184] 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 requires new vehicles to undergo emissions testing to ensure vehicles meet regulatory standards at the time of production (EPA 2024*a*). For some pollutants, we need to account for the fact that a vehicle's emission control system may become less effective over time.[185]. We split our analysis of on-road pollution into emissions that increase with vehicle age and those less affected by vehicle usage.

We begin with emissions that change as a vehicle ages, which consist of carbon monoxide (CO), hydrocarbons, (HC), and oxides of nitrogen (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

[183] This 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 MJ of energy (AFDC 2024*c*).

[184] For 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.3.3 below.

[185] Catalytic converters, for example, deteriorate over a vehicle's lifetime (Baronick et al. 2000)

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., the emissions control systems of newer vehicle models do not decay at different rates). Combining these parameters 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

$$\underbrace{EmissionRate_{y,m,p}}_{\text{Grams per Gallon}} = \underbrace{EmissionRate_{m,p}}_{\text{Grams per Mile}}(1 + AgeFactor_p)^{y-m} \times \underbrace{FuelEconomy_m}_{\text{Miles per Gallon}} \tag{69}$$

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 ($EmissionRate_{m,p}$) 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.[186] We assume no vehicles from model years earlier than 1957 remain in use. The EPA reports average fuel economy by model year ($FuelEconomy_m$). Fuel economy data for model years 1957–1975 come from EPA (1973), and data for model years 1975 onward come from the Automotive Trends Report (EPA 2023*b*).[187] 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 from vehicles using fuel containing different amounts of ethanol.[188] Using the emission rates reported in the authors' Tables S1, S2, and S3, we find that vehicles running on fuel with 9.8% ethanol emit 13.2% less NO_X, 6.8% more CO, and 13% more HC 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

[186]The 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.

[187]The 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.

[188]We do not adjust the emission rates for CH_4 or N_2O because 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 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).

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 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.

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 carbon dioxide (CO_2), sulfur dioxide (SO_2), particulate matter ($PM_{2.5}$), methane (CH_4), and nitrous oxide (N_2O). 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) reported by the EIA (2023*b*). As described above, we adjust our externalities to account for the share of ethanol in finished motor gasoline. For CO_2, we assume the share of gasoline that is ethanol 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 2023*b*, AFDC 2024*b*). 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).[189]

Because catalytic converters do not affect $PM_{2.5}$, CH_4, and N_2O emissions, we assume these three 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 (**MO**tor **V**ehicle **E**missions **S**imulator), a tool designed by the EPA to quantify pollution from mobile sources (EPA 2024*d*).[190] We use sales weights from the EPA (2023*b*) to average across vehicle classes included in MOVES' definition of light-duty vehicles (passenger cars and trucks and light-duty commercial trucks).[191] 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.[192]

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 annual gasoline consumption by vehicle age. To measure this, we use data on the age distribution and miles traveled for light-duty vehicles from the National Household Transportation Survey (FHWA 2017).[193] This survey provides a snapshot of the

[189] The EIA reports the average annual sulfur content of gasoline for 1997–2016. For years before 1997, we assume the sulfur content of gasoline equals the sulfur content 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).

[190] We use emission rates derived from MOVES but reported by Argonne National Laboratory's Greenhouse Gases, Regulated Emissions, and Energy Use in Transportation (GREET) model (Argonne National Laboratory 2013)

[191] MOVES includes these 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.

[192] As 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.

[193] When using data from the NHTS, we exclude recreational vehicles and motorcycles, as these are not included

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 assume the age distribution of the fleet and the distribution of VMT by vehicle age have remained 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. For example, in 2020, a vehicle from model year 2020 traveled approximately 13,962 miles in its first year. With a fuel economy of 25.38 MPG, this vehicle would use 550 gallons of gas in its first year. Multiplying by 0.065 (since 6.5% of vehicles in the NHTS were one year old) then gives us the weight we assign to vehicles of model year 2020. 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 its corresponding social cost to monetize damages.

C.3.2 Per-Mile 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.[194] However, some vehicle externalities arise on a per-mile basis. These externalities 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 $PM_{2.5}$ from tire and brake wear emission rate comes from MOVES (Argonne National Laboratory 2013), the same source from where we obtain per-mile emission rates for exhaust $PM_{2.5}$.[195] To value accidents, we use the annual fatalities avoided from a 1% reduction in VMT (263 fatalities avoided) estimated in Jacobsen (2013*b*), 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 2024*a*).[196] This calculation yields an average accident externality of $0.08 per mile. To value for congestion, we average per-mile externality estimates from three papers—Couture et al. (2018) ($0.02), Parry & Small (2005) ($0.05), and Parry et al. (2014) ($0.03)—for

in the Automotive Trends Report.

[194] If we assumed on-road local air pollution (except SO_2, which is a function of fuel composition) arose per-mile traveled instead of per-gallon consumed, the total externality reported in Appendix Table 12 would only fall by around $0.10 in 2020, assuming 52% of the price elasticity of gasoline arises from changes in VMT. This small change follows from the fact that local on-road pollution is a relatively small share of the total externality from a gallon of gasoline in 2020.

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

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

an average congestion externality of \$0.03 per mile.[197] 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. Weighting by the product of age shares and VMT, we obtain an average fuel economy of 23.1 MPG. Multiplying per-mile emission rates by this VMT- and age-weighted fuel economy yields per-gallon estimates for our three per-mile externalities.

Changes in gasoline consumption do not arise entirely from changes in VMT. As a result, we must know how much of a change in gasoline consumption us 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.5 an alternative approach where we apply β to the price elasticity of gasoline; each approach yields identical conclusions.

C.3.3 Summary of Fleet-wide Gasoline Externalities

Appendix Table 12 provides per-gallon estimates of each vehicle externality, the total externalities from pollution and driving, and the total vehicle externality for 2020. CO and HC include both local and global damages. On-road $PM_{2.5}$ includes emissions from vehicle exhaust and brake and tire wear. In 2020, $PM_{2.5}$ from exhaust contributed \$0.06 per gallon to on-road $PM_{2.5}$ emissions. $PM_{2.5}$ from brake and tire wear made up the remaining \$0.02 per gallon. Accidents, congestion, and $PM_{2.5}$ have been scaled by our preferred value of β (0.52), as described above. We do not observe on-road NH_3. We note that Appendix Table 12 only applies when considering a change in gasoline usage by the average vehicle in the fleet in a given year. Below, we describe how to calculate externalities for policies that target new vehicles or vehicles with specific characteristics (e.g., higher-than-average fuel economies).

In 2020, approximately 60% of the total externality produced by the average light-duty, gas-powered vehicle came from (upstream and on-road) pollution. Upstream emissions contributed 12%, or \$0.26, of total pollution damages. CO_2 made up 86% of total pollution damages and 52% of the total vehicle externality. Local air pollution made up 12% of total pollution damages. Accidents and congestion represented 40% of the total externality in 2020, or \$1.40 combined. Our results are similar for 2020 and 2022, with rising social costs driving the \$0.09 difference.

Appendix Figure 2 illustrates how the externalities from an additional gallon of gas consumed by light-duty, gasoline-powered vehicles have evolved (1990–2022). On-road CO_2 pol-

[197] For 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.

lution has increased both in magnitude and as a share of the total externality from the rising SCC. Total CO_2 damages increased from \$0.68 per gallon in 1990 to \$1.61 per gallon in 2020. Accidents and congestion have increased over time as the average annual fuel economy of the vehicle fleet has improved, as vehicles can drive more on a single gallon of gas. In dollar-per-gallon terms, the combined cost of accidents and congestion has grown from \$1.14 in 1990 to \$1.40 per gallon in 2020. On-road non-CO_2 pollution has experienced the most significant decline since 1990 due to improvements in the emission rates of new vehicle models, as documented by Jacobsen et al. (2023). In 1990, non-CO_2 pollution imposed an external cost of \$1.57 per gallon or 44% of the total externality in 1990. By contrast, total on-road non-CO_2 pollution contributed \$0.25 in 2020 or only 7% of the total externality. Between 1990 and 2020, NO_X fell from \$0.85 to \$0.08 per gallon, HC from \$0.23 to \$0.04 per gallon, CO from \$0.21 to \$0.05, and total $PM_{2.5}$ from \$0.20 to \$0.09 per gallon.

C.3.4 Lifetime Vehicle Externalities

In addition to valuing the externality imposed by consuming one gallon of gasoline, we 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 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 car being purchased. For example, if a subsidy induces the purchase of a vehicle in 2020, we consider the emission rates of a vehicle purchased in 2020, rather than a fleet-wide average emission rate. We still 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 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).[198] 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 pollutants with global damages. However, 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 decay. 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

[198]To calculate the lifetime of an average new light-duty vehicle, we take the 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 (2023*b*). This yields an average lifetime that rounds to 19 years.

the rebound in VMT due to improved vehicle fuel economy (which we include in our hybrid and vehicle retirement MVPFs), we do account for accidents, congestion, and $PM_{2.5}$ from tires and brakes, as vehicles generate these externalities when they travel more miles. In these instances, the per-mile externality does not differ between vehicle types even though per-gallon pollution externalities do vary as a function of fuel economy. 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.[199] 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 over the vehicle's lifetime.

D Rebound

When a policy causes people to consume more or less energy, this can affect the price of that energy, 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 section, 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

$$\begin{aligned} S'(p)dp &= dE\frac{S'(p)}{S'(p) - Q'(p)} \quad (70)\\ &= dE\frac{1}{1 - Q'(p)/S'(p)} \quad (71)\\ &= dE\left(1 - \frac{-\epsilon^D/\epsilon^S}{1-\epsilon^D/\epsilon^S}\right) \quad (72)\end{aligned}$$

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

[199] Phrased 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 damages from gasoline consumption are a small component of the local externality (especially when looking at 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.

$\epsilon^S = \infty$, 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.

In our empirical implementation, we account for rebound effects in the electricity generation market. We use estimates of the supply and demand curves elasticities of 0.78 and -0.19, respectively. This yields a rebound estimate of approximately 20%.

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. [200]

[200] While we do not incorporate rebound in the gasoline market, we note that many policies such as gas taxes can affect the cost of driving. We discuss in each policy context how we incorporate these effects. But, we note 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 this sense, estimates of the causal effect of the gas tax on gasoline consumption already incorporate this rebound effect.

E Policy Appendices

E.1 Battery Electric Vehicles

E.1.1 State-level Battery Electric Vehicle Financial Incentives

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'll 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 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 0.078, 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 that are full-sized and capable of 60 mph up until 2019. 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 model of BEV 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 0.326 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. is due to selection in the types of drivers that purchase BEVs independent of vehicle type. [201] 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.

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.1. 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 \$0.060 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

[201] It 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.

avoiding carbon emissions, which is 1.9% of the total amount (see Section 4). Written out this is $0.060 \cdot -2.931 \cdot (1 - 0.15 \cdot 0.2554 \cdot 0.5)$. Thus, our final values for the damages from using a BEV are -0.171 (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 1.982.

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. 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.3. Again, we take this amount and normalize it by the net MSRP of a BEV to get 0.099 for 2020 and 0.130 in-context, and multiply it by the elasticity to get 0.291 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 0.285 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 0.016 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 0.045 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 0.069 (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.1 and C.3. We calculate -0.009 in 2020 and -0.019 in-context in local damages from BEVs and 0.009 in 2020 and 0.011 in-context from the counterfactual ICE, which differenced out gives 0.000 and -0.008, respectively. After multiplying by the elasticity, we have 0.000 and -0.023.

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 C.1.4 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 0.038 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 inputs of the model. 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'll 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, marginal, and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries as well as the cost per kWh come from Ziegler and 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 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 0.064 (-0.233) 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 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 0.066/0.142. 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 0.015/0.044.

With the various components, we can now calculate the total WTP of 1.684 for 2020 and 1.221 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 that 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 $7107.55 in-context. Normalizing that value by the net MSRP gives us $0.001 in 2020 and $0.196 in context. Finally, multiplying by the elasticity gives us the federal fiscal externality of

0.002 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 0.033.

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 the text of Clinton and 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 Section E.5. 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 0.063/0.082.

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.1.3. The FE is calculated in the same way as the gas tax FE to get 0.008/0.024.

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*a*) 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 -0.005/0.000.

Thus, our final cost is 1.104/1.857, which gives us our MVPFs of 1.525 and 0.657.

E.1.2 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'll 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 $647.25.

For the in-context specification, the sales data by model comes from the Transportation Research Center at Argonne National Laboratory and covers all models that are full-sized and capable of 60 mph up until 2019. 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 text of 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 model of BEV 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 0.331 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 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)'s analysis finding that BEVs accumulate fewer annual miles than ICE vehicles: 7,165 versus 11,642. [202] 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.

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 each year 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.1. 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 0.060 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 $0.060 \cdot -2.611 \cdot (1 - 0.15 \cdot 0.3)$. Thus, our final values for the damages from using a BEV are -0.153 (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

[202] It 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.

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 1.700.

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.3. Again, we take this amount and normalize it by the net MSRP of a BEV to get 0.099 for 2020 and 0.118 in-context, and multiply it by the elasticity to get 0.259 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 $0.099 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 0.254 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 0.016 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 0.040 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 0.061 (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.1 and C.3. We calculate local damages from increased grid usage from BEVs of -0.009 in 2020 and -0.029 in context. We calculate savings from reduced gasoline consumption of 0.009 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 0.000 and -0.018, respectively. After multiplying by the price elasticity, it suggests \$1 of mechanical spending on the subsidy delivers local environmental benefits of \$0.000 in 2020 and \$-0.034in 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 C.1.4 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 0.034 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 inputs of the model. 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'll 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, marginal, and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries as well as the cost per kWh come from Ziegler and 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 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 0.049 (0.035) 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 2,857.1/3,225.4 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 0.059/0.128. 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 0.013/0.018.

With the various components, we can now calculate the total WTP of 1.580 for 2020 and 1.368 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 that 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 \$42.9789 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 \$0.001 in 2020 and \$0.168 in context. Finally, multiplying by the elasticity gives us the federal fiscal externality of 0.002 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 \$42.9789. This value normalized by the net MSRP and multiplied by the elasticity gives us the state fiscal externality of 0.002.

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 Section E.5. 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 2,857.1/3,225.4 and normalized by the net MSRP to get 0.056/0.074.

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.1.3. The FE is calculated in the same way as the gas tax FE to get 0.007/0.010.

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*a*) 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 -0.004/-0.004.

Thus, our final cost is 1.093/1.638, which gives us our MVPFs of 1.447 and 0.836.

E.1.3 Enhanced Fleet Modernization Program

Muehlegger and 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 $647.25.

For the in-context specification, the sales data by model comes from the Transportation Research Center at Argonne National Laboratory and covers all models that are full-sized and capable of 60 mph up until 2019. 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 and 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 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 model of BEV 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)'s analysis finding that BEVs accumulate fewer annual miles than ICE vehicles: 7,165 versus 11,642. [203] 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.

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 each year 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 description of the grid externalities calculations in Appendix C.1. 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 0.060 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 $0.060 \cdot -2.1 \cdot 0.85 \cdot (1 - 0.15 \cdot 0.2554 \cdot 0.5)$. Thus, our final values for the damages from using a BEV are -0.104 (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.

[203] It 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.

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 1.450.

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.3. Again, we take this amount and normalize it by the net MSRP of a BEV to get 0.099 for 2020 and 0.109 in-context, and multiply it by the elasticity and the pass-through rate to get 0.177 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 \$0.099 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 0.174 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 0.016 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 0.027 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 0.042 (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.3. We calculate local damages from increased grid usage from BEVs of \$-0.009 in 2020 and \$-0.003 in context. We calculate savings from reduced gasoline consumption of \$0.009 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 \$0.000 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 \$0.000 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 C.1.4 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 0.023 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 inputs of the model. 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'll 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 and 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 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 0.025 (0.105) 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 2,857.1/2,993.9 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 0.040/0.045. 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 0.009/0.045.

With the various components, we can now calculate the total WTP of 1.368 for 2020 and 1.477 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 that 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 $7021.44 in-context. Normalizing that value by the net MSRP gives us $0.001 in 2020 and $0.154 in-context. Finally, multiplying by the elasticity gives us the federal fiscal externality of 0.001 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 0.024.

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 the text of Muehlegger and 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.5. 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 2,857.1/2,993.9 and normalized by the net MSRP to get 0.039/0.041.

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.1.3. The FE is calculated in the same way as the gas tax FE to get 0.005/0.025.

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 (2023*a*) 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 -0.003/-0.005.

Thus, our final cost is 1.067/1.712, which gives us our MVPFs of 1.282 and 0.863.

E.2 Hybrid Electric Vehicles

E.2.1 HEV USA - Income Tax Credit

Gallagher and 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'll 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 that are full-sized and capable of 60 mph up until 2019. 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 model of HEV. 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. [204] 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.3. Since the VMT changes for each year of the car's lifetime, the damages each year 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 \$600.963 and \$445.390 respectively, of damages from gasoline. See a more detailed description of the gasoline externalities calculations in Appendix C.3. We then estimate the damages for each year of the vehicle's lifetime, which totals 8342.157 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 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 and Rapson (2022). 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

[204] It 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.

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 1.023.

In the first year of the counterfactual ICE vehicle's life, then, 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.3. Again, 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 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 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.3.2 and C.1. We calculate 0.028 in 2020 and -887.598 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 -887.553, 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 and 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 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 any necessary preliminary calculations as well as the data sources for the inputs of the model. 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'll 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 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, marginal, and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries as well as the cost per kWh come from Ziegler and 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 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 Lastly for WTP, 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 E.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 can now calculate the total WTP of 1.002 for 2020 and 1.011 for in-context.

Cost The cost of $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 and 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 the text of Gallagher and Muehlegger (2011). When normalized and multiplied by the elasticity, this gives us 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.5. 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%, so multiplying the gasoline producers' WTP by 0.21, we have -0.000 for the 2020 specification and 0.000 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 (2023*a*) 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.

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

E.2.2 Federal Income Tax Credit for Hybrid Vehicles

Beresteanu and 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'll 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 that are full-sized and capable of 60 mph up until 2019. This gives us an average MSRP of $25,758, which 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 model of HEV. 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". 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. [205] 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.3. Since the VMT changes for each year of the car's lifetime, the damages each year 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 600.963 and 600.963 respectively, of damages from gasoline. See a more detailed description of the gasoline externalities calculations in Appendix C.3. We then estimate the damages for each year of the vehicle's lifetime, which totals 8342.157 in damages in 2020 and 5510.528 in context. Normalized by the net MSRP, we have 0.249 for 2020 and 0.254 for in-context.

Finally, 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

[205] It 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.

externality in our baseline specification (see Section 4). Writing this out is $0.249 \cdot -1.985 \cdot (1 - 0.15 \cdot 0.25540.5)$. 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 and Rapson (2022). 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 40.62 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 1.580.

In the first year of the counterfactual ICE vehicle's life, then, 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.3. Again, 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.1 and C.3. We calculate 0.028 in 2020 and 0.038 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 and 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 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 any necessary preliminary calculations as well as the data sources for the inputs of the model. 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'll 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, marginal, and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries as well as the cost per kWh come from Ziegler and 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 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 Lastly for WTP, 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 E.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 cost of $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 and 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 0.198 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 the text of Gallagher and Muehlegger (2011). When normalized and multiplied by the elasticity, this gives us a state fiscal externality of 0.198.

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.5. 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 -0.001 for the 2020 specification and 0.002 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 (2023*a*) 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.2.3 HEV USA - Sales Tax Waiver

Gallagher and 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'll 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 that are full-sized and capable of 60 mph up until 2019. 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 model of HEV. 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. [206] 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.3. Since the VMT changes for each year of the car's lifetime, the damages each year 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 600.963 and 445.390 respectively, of damages from gasoline. See a more detailed description of the gasoline externalities calculations in Appendix C.3. We then estimate the damages for each year of the vehicle's lifetime, which totals 8342.157 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.

[206] It 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.

Finally, we take this amount and multiply it by the elasticity of -6.916 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 $0.249 \cdot -6.916 \cdot (1 - 0.15 \cdot 0.2554 \cdot 0.5)$. 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 and Rapson (2022). 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 1.231.

In the first year of the counterfactual ICE vehicle's life, then, 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.3. 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.1 and

C.3. We calculate 0.028 in 2020 and -887.598 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 -887.555, 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 and 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 any necessary preliminary calculations as well as the data sources for the inputs of the model. 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'll 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, marginal, and cumulative sales to describe. Marginal and cumulative sales of MWh of batteries as well as the cost per kWh come from Ziegler and 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 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 Lastly for WTP, 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 E.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 can now calculate the total WTP of 1.036 for 2020 and 1.188 for in-context.

Cost The cost of \$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 and 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 the text of Gallagher and Muehlegger (2011). When normalized and multiplied by the elasticity, this gives us a state fiscal externality of 0.388.

Gas Tax FE We also calculate a gasoline tax fiscal externality using the average state and federal gas tax rates as described in Appendix Section E.5. 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 -0.004 for the 2020 specification and 0.008 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 (2023*a*) 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 $0.15 \cdot 0.25540.5 = 1.9\%$ 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.3 In-Context Muehlegger and Rapson (2022) Discussion

Figure 1 Panel A presents the components of our WTP and net cost estimates used in the construction of the MVPF. All components are normalized by the mechanical cost of the subsidy change (i.e., the cost if individuals did not change their behavior). By construction, individuals are willing to pay $1 per $1 in mechanical subsidy cost. The fractional pass-through of subsidies means that this $1 is split between $0.85 for those purchasing the cars and $0.15 for the owners of CA dealerships that sell EVs.

The next bars in Figure 1 measure the environmental externalities associated with marginal EV purchases. We use estimates from Holland et al. (2016), who report the fuel economy of the counterfactual car that an EV customer would have purchased (they find that EVs displace a cleaner-than-average new light-duty car). We then combine this counterfactual fuel economy (39.4 MPG) with an estimate of the per-mile externality of a new car as well as data on the profile of VMT over a vehicle's lifetime.[207] Here, we use estimates of the average EVs' VMT from Zhao et al. (2023) that is roughly 61% of VMT for the average gas-powered car.[208] This provides the lifetime environmental benefits from *not* driving the counterfactual gas-powered vehicle.[209] This leads to a WTP of $0.0155 from local pollutants and $0.195 from global pollutants. This sums to a total benefit of $0.21 from the reduced gasoline consumption induced by the subsidy.

[207]To calculate the per-mile externality we use estimates from new cars in the year of policy implementation.

[208]This estimate is similar to those presented in Davis (2019) and Burlig et al. (2021).

[209]We assume the average EV and the average counterfactual car both have 17-year lifetimes (NHTS 2006). We also assume that EVs have the same total lifetime VMT and distribution of VMT over their lifetimes. Since we assume the same VMT between gas-powered vehicles and EVs, we do not need to account for damages from accidents, congestion, or $PM_{2.5}$ from tires and brakes, which arise per mile traveled. Appendix C describes this approach in further detail.

While the decrease in gasoline consumption leads to environmental benefits, the additional use of electricity leads to partially offsetting environmental damages. As described in detail in Appendix C, we incorporate the emissions from additional electricity usage with estimates from the EPA's Avoided Emissions and Generation Tool, AVERT (EPA 2024*b*).[210] The environmental externality per kWh in California between 2015 and 2018 is roughly six times lower than the national average. Combining the change in emissions with our valuations of those externalities discussed above for the electric grid, we find that an average EV leads to an increase in \$0.087 in welfare cost due to pollution from the grid (\$0.016 of this accrues to US residents). On net, society is willing to pay \$0.113 for these global benefits and an additional \$0.01 for the local benefits from substituting toward an EV.

Some of these increases in electricity usage from EVs are potentially moderated through a rebound effect: the increased electricity demand can push up prices leading to lower electricity consumption. To account for this, we use estimates of the demand and supply elasticities of electricity: We use a demand elasticity from the EPA of -0.19 (EIA 2021*b*) and a supply elasticity of 0.78. Combining, we show this implies that roughly 20% of the electricity demand is offset by reduced demand due to higher electricity prices.[211] This suggests that roughly \$0.017 of the environmental harms from increased electricity consumption are offset by the rebound effect.

In addition to environmental externalities from driving the car, we also account for the fact that the upstream production of EVs is more carbon-intensive than the production of ICE vehicles due to the nature of the battery production. We incorporate estimates from Winjobi et al. (2022) that suggest that batteries in particular release 0.06 tons of CO_2 per kWh. This suggests the average EV imposes a global externality from battery production of \$0.026 per EV, leading to an externality of \$0.026 per dollar of EV subsidy.

We now turn to the effects of learning-by-doing externalities, which potentially arise in the EV context in the production of batteries. We use estimates of the recent slope of the learning-by-doing curve for batteries of $\theta = -0.421$ from Way et al. (2022), indicating that a 1% increase in battery production leads to a reduction in battery costs of 0.42%. We combine this with an estimated demand elasticity of $\epsilon = -2.1$ and a future discount rate of 2%. As we discuss in Appendix C, we account for forecasted changes in the cleanliness of the electric grid over time (i.e., cleaner electric grids in the future) as well as improvements to the fuel economy of new light-duty cars. This yields an environmental benefit from increased future consumption of EVs due to lower future prices that is \$0.144 per dollar of mechanical cost of the subsidy, with \$0.03 flowing directly to US residents.[212]

[210] In the years after 2023, we use grid projections rather than direct emissions estimates from the AVERT model. In particular, we use the Princeton REPEAT Project's mid-range forecast from 2023-2050 to estimate the average combustion share of the grid (Jenkins & Mayfield 2023). We then use estimates from the AVERT model to map this combustion share of the grid to the environmental externality per kWh. We compute these estimates at both the state and national level. For projections beyond 2050, we repeat the same process but assume the combustion share remains constant.

[211] We do not incorporate a rebound effect for gasoline. This is because we treat gasoline as a global market where the price does not meaningfully change in response to the demand shock induced by EV purchases. This differs from the local electricity market.

[212] While the subsidy encourages the purchase of a new electric vehicle, we only incorporate learning-by-doing effects associated with the production of batteries. As a result, we find that the learning-by-doing effects of EV subsidies fall rapidly over time. Intuitively, as the battery costs decline, there is a limit to the extent to which lower battery prices can lower future EV costs. We show in Appendix B how we account for this dynamic in learning by doing.

We also estimate that learning by doing lowers the price of EVs for future purchasers. This generates a willingness to pay of $0.261 (in examining the incidence of these gains, we allocate 10% percent of this to the US, according to the US share of global EV purchases (IEA 2023).) Taken together, the learning-by-doing effects increase the value of the subsidy by $0.404 per dollar of EV subsidy. We note that including these benefits requires one to believe that the price declines as a function of cumulative production reflect spillovers across firms that are not internalized through the patent system or other means. In this sense, the inclusion of these benefits is potentially an upper bound on the MVPF for EVs.

The last benefit we consider is the impact of the policy change on the profits of gasoline and electricity producers. Our estimates suggest a marginal EV purchase between 2015 and 2018 reduced gasoline consumption by 2,994 gallons over the lifetime of the vehicle. We account for producer profits using an average markup per gallon of gas of $0.68 per gallon, or 27% of the retail price between 2015 and 2018.[213] This lies above the economy-wide average markup of 8% (De Loecker et al. 2020), leading to a relative profit loss by producers from the shift away from gasoline towards other goods.[214] Applying a 21% effective corporate tax rate, we calculate post-tax lost producer profits are equal to $0.045 per dollar of the subsidy.[215] By contrast, electricity suppliers benefit from increased electricity consumption. Electric utilities are a regulated industry with natural monopolies that sell electricity at a markup, which we estimate to be 12.9% on top of the 8% economy-wide markup.[216,217] We find that the subsidy generates an increase in electricity company profits of $0.045 per $1 of subsidy. The total willingness-to-pay for the policy change in California between 2015 and 2018 was $1.518 per $1 of subsidy.

The next several columns in Figure 1 consider the net cost of the subsidy to the government. During this time frame, the federal government offered $7,021 in tax credits for EV purchases.[218] This means that the estimated increase in EV purchases costs the federal govern-

[213]We calculate the markup on a gallon of gasoline by summing the profits at each step in the gasoline supply chain. First, we estimate crude oil producer profits by subtracting the landed cost of crude oil (EIA 2024*d*) from the US refiner crude oil acquisition cost (EIA 2024*b*). We divide by the average US refinery yield (EIA 2024*g*) to convert from barrels of crude oil to gallons of petroleum product. Next, we calculate refiner profits using the average cost of refining a barrel of crude (Favennec 2022) and the share of the gasoline price that comes from refiner costs and profits (EIA 2024*a*). We determine retailer profits by subtracting the average retail gasoline price (EIA 2024*c*) from the price retailers pay for a gallon of gasoline (EIA 2022), assuming no long-run costs to retailers.

[214]For ease of explanation, we use the term markup to refer to the producer profit rate. In De Loecker et al. (2020), this term is referred to as the average profit rate while markup is used to refer to the levels of prices relative to marginal costs (excluding capital expenditures, fixed costs, and overhead costs).

[215]We draw the corporate tax rate from Watson (2022). We also use that foregone tax rate estimate to adjust the net cost of the policy. This tax rate does not vary over time. In context, the pre-tax markup on gasoline was $0.27 per dollar spent on gas, or $0.21 per dollar spent on gas after adjusting for corporate taxes.

[216]Appendix C explains our approach in further detail. We start by using the EIA's levelized cost of electricity (LCOE) (EIA 2015). We construct the total cost per MWh at the state and national level by taking an average of the LCOEs weighted by the share of the grid made up of each generation source. We also add transmission and distribution costs from the EIA to the LCOE (EIA 2023*a*). The EIA does not report the state-specific LCOE for each generation source. It only reports the minimum, average, and maximum. We create 50 discrete equally spaced buckets from the minimum to maximum LCOE for each generation source and assign states to each bucket using the BLS's power generation industry wage index. We then compare these costs to the retail price of electricity by year and state from the BLS (BLS 2024).

[217]As in the gasoline market case, we split this markup into two components: after-tax profits and government revenue. In our baseline specification, we assume that 28% of utilities are publicly owned, that the effective corporate tax rate on private utilities is 10% (drawing from (DOT 2016)), and 100% on public utilities.

[218]The maximum federal subsidy was $7,500 during this period, but it was subject to manufacturer-specific

ment an additional \$0.275 for every \$1 in subsidy. (The source of this estimate can be seen in equation 10 above. We get \$0.275 by multiplying the elasticity by the size of the pre-existing subsidy as a fraction of the total price of the vehicle.) In addition to federal incentives, CA also provided a subsidy of \$9,000, which we estimate leads to a \$0.414 additional cost. The reduced gasoline consumption leads to a loss in gas tax revenue for the government of \$0.041 for every \$1 in subsidy.[219] This is offset by \$0.013 in additional government revenue collected on producer profits.[220] Finally, we incorporate a fiscal impact on the US government's budget due to the productivity reduction caused by climate change. This is equal to \$0.006 for every \$1 in subsidies. Adding these costs together, we estimate a net cost of \$1.711 for every \$1 in mechanical subsidy costs. When we take the ratio of the willingness-to-pay and these net costs, we arrive at an in-context MVPF of 0.887.

The MVPF of 0.887 means that a \$1 increase in net government subsidy spending in CA in 2015–2018 would have led to \$0.887 in benefits for members of society.

E.4 Weatherization

Our category average MVPF for weatherization programs is 0.98 with a 95% confidence interval of [0.93,1.04]. 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. We do, however, harmonize the externalities in the baseline specification. 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

caps that caused the subsidy to decline as more vehicles were produced. We use a purchase-weighted average for our main MVPF estimates, but we also consider a scenario with \$7,500 subsidies.

[219] Total gas taxes in California were \$0.54 per gallon. This is due to the federal gas tax of \$0.184 per gallon and a state gas tax of \$0.36 per gallon.

[220] In practice, utilities make profits, some of which flow to the government while gasoline producers generate losses. The effect on utilities is larger than the effect on gasoline producers.

the share of marginal beneficiaries. Our baseline MVPF assumes that 50% of households are marginal to the subsidy, and we present robustness to a 100% marginal assumption. For the marginal households, some of them 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.

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 Section C.2 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.4.1 Weatherization Assistance Program in Michigan

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 Section C.1 (electricity) and Appendix Section C.2 (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 Appendix Sections C.1 and C.2, 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 Section C.1 (electricity) and Appendix Section C.2 (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.

E.4.2 Home Weatherization Assistance Program in Illinois

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.[221]

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.

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 Section C.1 (electricity) and Appendix Section C.2 (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.

[221] Christensen, 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.

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 Sections C.1 and C.2, 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 Section C.1 (electricity) and Appendix Section C.2 (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.

E.4.3 Low-income Energy Efficiency Program in Florida

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 Section C.1. 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 Section C.1. 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 Section 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 Section C.1. 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.

E.4.4 Energy Retrofits in Phoenix

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 Section C.1. 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 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 Section C.1. 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 Section 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 Section C.1. 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.

E.4.5 Weatherization Assistance Program in Wisconsin

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 takeup 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.

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 takeup to measure consumer surplus from the subsidy.

In their paper, Alcott 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 takeup and a 2% change in retrofit investment takeup. 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 Section 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 Appendix Sections C.1 and C.2. 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 Section C.1. 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.5 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 increase 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 burning one gallon of gasoline imposed $2.12 in damages from pollution; $1.89 of these damages came from global pollutants, while the remaining $0.23 came from local pollution. Appendix C explains how we estimate these externalities. These damages include emissions released upstream during the oil extraction and refining processes. 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.[222] 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 $0.312 for a $1 increase in the gas tax rate, with $0.2777 for greenhouse gases that contribute to global warming and $0.0339 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 × 0.2554 × 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

[222] To 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 and weight by monthly data on the quantity of gasoline supplied from the EIA's "U.S. Product Supplied of Finished Motor Gasoline" series, 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, 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.

the remaining 1.92% of global benefits into the MVPF's denominator as increased long-run revenue. We multiply the \$0.2777 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.3.2 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 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 toward 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.[223] 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:

$$\frac{dQ_{EV}}{dP_{Gas}} \times \frac{P_{Gas}}{Q_{EV}} = \eta_{EV,\,Gas} \tag{73}$$

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 in the opposite direction to the change in own-good consumption: the increase in EV consumption is equal in magnitude to the decrease in gas-powered vehicle consumption. We can express this relationship as

$$\frac{dQ_{EV}}{dP_{Gas\,Car}} = \frac{dQ_{Gas\,Car}}{dP_{EV}} \tag{74}$$

$$\frac{dQ_{Gas\,Car}}{dP_{EV}} = \frac{-dQ_{EV}}{dP_{EV}} \tag{75}$$

where the price of owning a gas-powered vehicle ($P_{Gas\,Car}$) is the present discounted value of

[223] 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 (namely, gasoline producer profits, environmental benefits of reduced gas consumption, and gas tax revenue) need not be adjusted.

gas consumed over the vehicle's lifetime.[224] Changes in the price of gasoline enter linearly into the price of owning a gas-powered vehicle, such that

$$\frac{dQ_{EV}}{dP_{Gas\,Car}} = \frac{dQ_{EV}}{dP_{Gas}} \tag{76}$$

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:

$$\frac{dQ_{EV}}{dP_{Gas}} = \frac{dQ_{EV}}{dP_{Gas\,Car}} = \frac{-dQ_{EV}}{dP_{EV}} \tag{77}$$

Eq. 73 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.

$$\frac{-dQ_{EV}}{dP_{EV}} \times \frac{P_{Gas}}{Q_{EV}} = \eta_{EV,\,Gas} \tag{78}$$

$$(\frac{-dQ_{EV}}{dP_{EV}} \times \frac{P_{Gas}}{Q_{EV}}) \times \frac{P_{EV}}{P_{EV}} = \eta_{EV,\,Gas} \tag{79}$$

$$\underbrace{(\frac{-dQ_{EV}}{dP_{EV}} \times \frac{P_{EV}}{Q_{EV}})}_{\text{Price Elasticity of EVs } (\epsilon_{\text{EV}})} \times \frac{P_{Gas}}{P_{EV}} = \eta_{EV,\,Gas} \tag{80}$$

To calculate the cross-price elasticity implied by Eq. 80, we use the own price elasticity (-2.1) estimated by Muehlegger & Rapson (2022).[225] 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 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

[224] We 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 vehicles both remain in use for 20 years. We assume EVs and gas-powered cars travel the same number of miles over their lifetimes.

[225] We 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.

$$\eta_{EV,\,Gas}\frac{V_{EV}}{P_{Gas}}\frac{Q_{EV}}{Q_{Gas}} \tag{81}$$

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.[226] In 2020, US consumers purchased 238,540 EVs, generating a total of \$810.63 million and \$87.31 million in global and local damages, respectively. Dividing monetized damages from EV consumption by the product of the price of gasoline (\$2.27) and the number of gallons of gasoline consumed 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.[227] 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

$$\begin{aligned} WTP_{Society} = & \underbrace{\left(\epsilon_{Gas}\frac{(1-\sigma_{US\,Govt})\,V_{Gas,\,Global}}{P_{Gas}} + \eta_{EV,\,Gas}\frac{(1-\sigma_{US\,Govt})\,V_{EV,\,Global}Q_{EV}}{P_{Gas}Q_{Gas}}\right)}_{\text{Global Env (-\$0.272)}} \\ & + \underbrace{\left(\epsilon_{Gas}\frac{V_{Gas,\,Local}}{P_{Gas}} + \eta_{EV,\,Gas}\frac{V_{EV,\,Local}Q_{EV}}{P_{Gas}Q_{Gas}}\right)}_{\text{Local Env (-\$0.034)}} + \underbrace{\beta\epsilon_{Gas}\frac{V_{\text{Driving}}}{P_{Gas}}}_{\text{Driving (-\$0.209)}} \end{aligned} \tag{82}$$

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 -\$0.0338 (with -\$0.03389 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 \$0.2427 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.

[226]As 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.

[227]We 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). To those annual values, we add the total annual quantity of aviation gasoline supplied 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.

As described above, 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.

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 today.[228] 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 care not about a change in the subsidy amount but rather 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 (81) 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,

$$\eta_{EV,\,Gas}\left(V_{Dynamic}\frac{P_{EV}}{\epsilon_{EV}}\frac{Q_{EV}}{Q_{Gas}}\right) \tag{83}$$

Put differently, equation (83) 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.[229]

[228] We do not consider substitution toward EVs for in-context MVPFs calculated for years before 2011.

[229] For the price benefits, we take the \$0.368 of learning-by-doing benefits calculated using a 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 (\$0.042 in global benefits and \$0.006 in local benefits), which yields a learning-by-doing environmental benefit of \$0.00025 per dollar of spending on gasoline.

$$WTP_{LBD} = \underbrace{\eta_{EV,\,Gas}\left(((1-\sigma_{US\,Govt})\,V_{Dynamic,\;Env.\,Global} + V_{Dynamic,\;Env.\,Local})\,\frac{P_{EV}}{\epsilon_{EV}}\frac{Q_{EV}}{P_{Gas}Q_{Gas}}\right)}_{\text{Dynamic Env (-\$0.00025)}}$$
$$+ \underbrace{\eta_{EV,\,Gas}\left(V_{Dynamic,\;Price}\frac{P_{EV}}{\epsilon_{EV}}\frac{Q_{EV}}{P_{Gas}Q_{Gas}}\right)}_{\text{Dynamic Price (-\$0.0019)}} \tag{84}$$

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.

Producers' WTP 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. We also account for utilities' WTP for increased EV charging.

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*d*,*b*). 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.[230] In 2020, the average markup imposed by crude producers equaled $0.06 per gallon or 2.5% 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*a*). 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.[231]

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. 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.[232]

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

[231] Neither approach results in a negative markup in any period.

[232] This approach generates a negative markup for one month in our data (October 2019). We set markups to

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. 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).[233] 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.[234] 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

$$WTP_{Producers} = \underbrace{\epsilon_{Gas} \frac{(1 - \tau_{Corp,\, Gas})\, \mu_{Gas}}{P_{Gas}}}_{\text{Gasoline Producers (\$0.071)}} + \underbrace{\eta_{EV,\, Gas} \frac{(1 - \tau_{Corp,\, Utilities})\, \mu_{EV} Q_{EV}}{P_{Gas} Q_{Gas}}}_{\text{Utilities (-\$0.00005)}} \tag{85}$$

where μ is the pre-tax profit the producer earns 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. In 2020, gasoline producers were WTP $0.071 for the policy change, and utilities -$0.00005.

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 (-$0.515) 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. As described below, removing a tax also allows us to treat both taxes and subsidies as having a $1 mechanical cost.

$0 if this approach yields a negative markup.

[233] We do not vary across time the effective corporate tax rate gasoline producers face.

[234] We hold the price of electricity in 2020 constant over the vehicle's lifetime and discount using our preferred discount rate of 2%.

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 gasoline spending. In 2020, the federal gas tax rate was $0.184 per gallon (FHWA 2022*b*) while the average state tax on gasoline (weighted by gross gallons of gasoline taxed) was $0.281 per gallon (FHWA 2022*a*, 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.[235] 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, society's WTP in 2020 for *all* global benefits was $0.2777, implying the government generated $0.005 ($0.2777 $\times$ 0.0192) in revenue by abating carbon emissions today and promoting economic output tomorrow.[236] 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.[237] Collecting the mechanical revenue raised and the fiscal externalities, we can express the denominator of the MVPF as

[235] We 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.

[236] This total WTP for global damages includes global benefits and damages from EV substitution and learning-by-doing.

[237] To 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.

$$\text{Cost} = \underbrace{\epsilon_{Gas}\frac{\tau_{Gas,\,Federal} + \tau_{Gas,\,State}}{P_{Gas}} + \epsilon_{Gas}\frac{\tau_{Corp,\,Gas}\mu_{Gas}}{P_{Gas}} + \eta_{EV,\,Gas}\frac{\tau_{Corp,\,Utilities}\mu_{EV}Q_{EV}}{P_{Gas}Q_{Gas}}}_{\text{Taxes (-\$0.087)}}$$
$$+ \underbrace{\eta_{EV,\,Gas}\frac{\tau_{EV}Q_{EV}}{P_{Gas}Q_{Gas}}}_{\text{Subsidies (-\$0.0001)}}$$
$$+ \underbrace{\left(\epsilon_{Gas}\frac{\sigma_{US\,Govt}V_{Gas,\,Global}}{P_{Gas}} + \eta_{EV,\,Gas}\frac{\sigma_{US\,Govt}V_{EV,\,Global}Q_{EV}}{P_{Gas}Q_{Gas}} + \eta_{EV,\,Gas}\frac{\sigma_{US\,Govt}V_{Dynamic,\,Env.\,Global}P_{EV}Q_{EV}}{\epsilon_{EV}P_{Gas}Q_{Gas}}\right)}_{\text{Climate FE (\$0.005)}} \tag{86}$$

Each component is labeled using the corresponding label from Figure 7. 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 0.604 in 2020.

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.

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 -0.3340 (s.e. 0.0451). 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 (\$0.034), 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 of less 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 0.604 in 2020.

In 2001, the nominal gas price was \$1.47. Society was WTP -\$0.187 for reduced greenhouse gases, -\$0.127 for reduced local air pollution, and -\$0.194 for reduced driving externalities. We do 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.

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.

F.1 Estimating Publication Bias

We first form a dataset of the t-statistics for the studies underlying our estimates.[238] 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 93 distinct estimates with t-statistics.

Appendix Figure 10 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.[239] Panel B of Appendix Figure 10 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 2.200 (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 insignificant studies.[240]

[238] Most 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.

[239] Here, 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 10 shows approximate symmetry around 0, suggesting that ignoring the signs of the estimates sacrifices little information.

[240] While our baseline binning of .98 is relatively large, we obtain similar results for smaller bins, e.g., .49 and .28.

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.[241]

Appendix Figure 11 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 12 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 correction with higher degrees of estimated bias (e.g., using the approach in Andrews & Kasy (2019)) similarly preserves our main conclusions.

[241]Their 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.

G Regulation

Our primary results focus on the welfare benefits and costs of taxes and subsidies that affect greenhouse gas emissions. The MVPF approach is particularly well suited to analyze policies of this sort – policies with direct statutory impacts on the government's budget. Alternatively, one could seek to reduce emissions via changes in regulations. While regulations can have an impact on the government budget (e.g., a reduction in gasoline usage can reduce gas taxes collected), the impact on the government is often small. Instead, the key welfare effect regulation policy is the tradeoff it induces across multiple groups of beneficiaries. For example, improvements in fuel standards may reduce emissions but car owners may need to pay more for cars under stricter fuel economy standards.[242]

While it is not particularly informative to construct an MVPF corresponding to a change in regulation policy, the MVPF framework can still be used to examine the welfare consequences of these policies. This approach requires a different conceptual experiment than the one outlined in equation (6). Here, we recall 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.

The key question here is whether the environmental benefits obtained through regulation could be obtained more efficiency 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?

In this Appendix, we present a detailed description of the results from this exercise comparing gas taxes and income taxes to CAFE standards using estimates 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.

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

[242]In 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.

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).[243] 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.[244]

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 (26.38 MPG) fuel economy. This approach compares the average light-duty vehicle purchased in 2020 to a vehicle that achieved an additional mile per gallon burned, or a 3.9% more fuel-efficient vehicle.[245] We account for the rebound in miles traveled by the more fuel-efficient vehicle by multiplying the percent-change in the cost of driving one mile by the elasticity of VMT with respect to the cost of driving from Small & Van Dender (2007) (-0.2221).[246] This approach to calculating the benefits of CAFE standards assumes the size of the vehicle fleet remains constant while the fleet's composition changes.

Recall that Appendix Figure 6, Panel A illustrates the costs and benefits of increased CAFE standards. 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)?[247]

[243] As 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.

[244] Leard & 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.

[245] For example, while a new light-duty vehicle manufactured and purchased in 2020 generated $16,858.56 in global damages over its lifetime, a new light-duty vehicle with a 1 MPG higher fuel economy generated $15,923.97 in global damages ($15,923.97 divided by 1.039), before accounting for the rebound in VMT.

[246] We calculate the cost of driving one mile by dividing the cost of a gallon of gas ($2.27 in 2020) by the vehicle's fuel economy. We then calculate the percent change in the cost of driving relative to the average new light-duty vehicle released in 2020. Multiplying the percent-change in the cost of driving by the elasticity of VMT with respect to the cost of driving (-0.2221) yields a rebound in VMT of 0.84%. The positive rebound implies that VMT increases as vehicles become more fuel-efficient and the cost of driving one mile falls.

[247] This test of "efficiency" dates back to the classic definition of Kaldor (1939) and Hicks (1940) with the

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 \$2.25. The tax also generates \$2.09 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, \$1.93. The MVPF of raising revenue from the average consumer is around 1.2, suggesting that this costs the government \$1.61 (=\$1.93/1.2), 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 $-\$0.60 = -2.09 - 0.11 + 1.61$. On net, Appendix Figure 6, Panel B shows that the government can replicate the distributional incidence of CAFE using taxes and still run a \$0.60 surplus, in contrast to the \$0.39 deficit that CAFE generates. In other words, it is \$0.99 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 \$0.99 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 (2013*a*) (Panel B). We present the benefits and costs of the regulation in blue and the tax in orange.[248] 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 distortions that are partially orthogonal to the efforts to simply reduce gasoline consumption. Second, gas taxes reduce vehicle miles traveled which leads to reductions in accidents and congestion – benefits that are not achieved through CAFE regulation. A deep analysis of the theoretical mechanisms driving the results is beyond our scope; rather, we simply note that the empirical results suggest

modification that we use actual tax and transfer policies instead of lump-sum redistribution to neutralize distributional incidence.

[248]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*b*). For Jacobsen (2013*a*), we take the per-household estimates (reported in the author's Table 6) of the effects of a 1 MPG increment in CAFE standards on producer welfare, consumer welfare, gallons of gasoline used, and vehicle miles traveled; multiply by the number of households in the author's sample (20,429); and value the change in the total change in gasoline and the total change in VMT using the average per-gallon (\$2.12 per gallon) and per-mile (\$0.118 per mile) externalities we calculate in Appendix C. Although we typically assume that CAFE standards *increase* VMT, excluding the author's estimated reduction in VMT would only reinforce our conclusion that gas taxes are more efficient than CAFE standards.

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 the Renewable Portfolio Standards (RPS) regulation. 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 (2020), 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 $58-$298. We use the lower bound ($58) of the authors' reported cost per ton estimates, as our ultimate conclusion will be robust to assuming RPS imposes larger welfare costs on consumers. In addition, because the $58 estimate from Greenstone & Nath (2020) does not include learning-by-doing benefits or local pollution benefits, we harmonize our estimates to theirs by excluding these components when considering the wind PTC impacts. For global environmental benefits, we use our baseline $193 SCC, but since these are the only environmental benefits the results would be unchanged with other values of the SCC.

The results suggest that every $1 of environmental benefits provided by RPS imposes a cost on consumers of $0.31 and a $0.02 savings to the government due to the climate fiscal externality, which are displayed in Appendix Figure 7, Panel C.[249] 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. Income taxes that tax consumers enough to impose the same $0.31 cost that RPS imposes on them would generate $0.48 in revenue. This means one could construct a combined wind PTC and income tax regime that delivers $1 of environmental benefits and $0.31 in costs to consumers, but that generates $0.12 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.

[249]We follow Greenstone & Nath (2020) 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. In addition to the MVPF, the two most common alternative metrics in cost-benefit analysis are (a) the net benefits of a policy and (b) the benefit-cost ratio of a policy. We briefly discuss these metrics and explain the advantage of our focus on constructing the MVPF.

Net benefits of a policy change equal the difference between the total benefits provided and the total costs. This contrasts with our MVPF approach which focuses on the ratio of net benefits to individuals divided by the net government cost. 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.

The ratio used to construct an MVPF is related to a large literature discussing the pros and cons of benefit-cost ratios. 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 (6). 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).

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.

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.1). This adds up to $2,216. The 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 for a battery electric vehicle is $6,634.

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.3 and C.1, 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 and dividing it by the tons of carbon, we have a resource cost per ton of $900.06.

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 is simply the difference between these two LCOEs, which is -\$0.0167.

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 is -\$42.24.

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/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 is \$0.00291.

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 is \$4.43 per ton

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 and 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 State-level ENERGY STAR Rebate - Clothes Washers

To estimate the energy savings we use the kWh difference from Houde and Aldy (2017) of 201 kWh since this value is estimated closer to 2020 than the one reported by Datta and 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.4.3 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 $208.

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 and 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 and 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 and Aldy (2017) of 65 kWh per year since this value is estimated closer to 2020 than the one reported in Datta and 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 and Aldy (2017) of 34 kWh per year since this value is estimated closer to 2020 than the one reported in Datta and 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 (2019) 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 and 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, Linn, and Spiller 2013)

Li, Linn, and Spiller (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 two months of acceleration. We use a 3% interest rate and get a leasing cost of $153.21.

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 0.24. Thus the cost is $23.52.

Lastly, using the retail gasoline price net of the gasoline tax and markups as in the BEV calculations, the 0.24 mpg difference, and the lifetime of the vehicle, the present discounted value of the gas savings between the new car and its counterfactual is $83.15. Thus, the resource cost is $93.58.

The carbon number is the emissions saved from that difference in fuel economy over the lifetime of the car (see Appendix C.3 for details on the estimation of driving emissions), which is 0.92 tons. The resource cost per ton is $101.56.

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 $89.79. Thus, the resource cost is $3,310.

The carbon number is the emissions saved from that difference in fuel economy over the 3.8 years (see Appendix C.3 for details on the estimation of driving emissions), which is 1.26 tons. The resource cost per ton is $2,625.

I.5.3 Cash for Clunkers (Hoekstra, Puller, and West 2017)

Hoekstra, Puller, and West (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 3.57. Thus the cost is $350.

Lastly, using the retail gasoline price net of the gasoline tax and markups as in the BEV calculations, the 3.57 mpg difference, and the lifetime of the vehicle, the present discounted value of the gas savings between the new car and its counterfactual is $1,082. Thus, the resource cost is -$119.65.

The carbon number is the emissions saved from that difference in fuel economy over the lifetime of the car (see Appendix C.3 for details on the estimation of driving emissions), which is 11.99 tons. The resource cost per ton is -$9.98.

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

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 form 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 for a hybrid electric vehicle is $1,200.

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.3. 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 of $659.

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 0.009779 (details can be found in Appendix C.3). Thus, the resource cost per ton is -\$103.80.

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 \$0.49 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 -\$45.55.

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 Accounting for Learning by Doing in Resource Cost per Ton

All of the calculations above have assumed that there are no learning-by-doing effects. Here we discuss the accounting of those for wind, solar, BEVs, and HEVs within a resource cost-per-ton framework.

Recall that our model of learning-by-doing looks at the effect of $1 in mechanical subsidy on future prices of the object of interest. For example, the CEFMP policy leads to a dynamic price component of 0.309. This can be interpreted as a 31-cent reduction in future BEV prices for every dollar of subsidy. To use this effect in the resource cost per ton calculations, we convert to a per vehicle (or per other object for other programs) unit by dividing the component by the semi-elasticity and pass-through rate if relevant. For CEFMP, the semi-elasticity is 0.00003924 and the pass-through rate is 85%, so the per vehicle component is $9,264. Since this is a future benefit, we subtract it from the existing resource cost estimate (which for BEVs was $6,634) to get -$2,630.

Nothing changes with the carbon estimate when we add in learning-by-doing, so the new resource cost per ton number is -$378.72. The other learning-by-doing policies follow the same steps to account for the dynamic price effect.

I.13 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.14 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.[250] 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).[251]

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

[250] We 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

[251] While the CEA estimates this category as $19.9 billion, we subtract $2 billion, one for SEP and $1 billion for section 25D tax credits.

subsidies.

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. Sending 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 he 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			MVPF (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	1.42
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.)