NBER WORKING PAPER SERIES

A WELFARE ANALYSIS OF POLICIES IMPACTING CLIMATE CHANGE

Robert W. Hahn
Nathaniel Hendren
Robert D. Metcalfe
Ben Sprung-Keyser

Working Paper 32728
http://www.nber.org/papers/w32728

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2024, Revised October 2025

We thank Sarah Aaronson, Mira Chaskes, Beatrice Chen, Aidan Creeron, Jamie Emery, Charlie Hutchinson, Jack Kelly, Navya Kumar, Isaac Maruyama, Victor Mylonas, Sethu Odayappan, and Lukas Puschnig for excellent research assistance. We thank Joseph Aldy, Lucie Gadenne, Ken Gillingham, Michael Greenstone, Christopher Knittel, Mark Jacobsen, Al McGartland, Nick Muller, Ishan Nath, Jacquelyn Pless, Brian Prest, Imran Rasul, Joe Shapiro, Alex Teytelboym, Rick van der Ploeg, and Catherine Wolfram for their excellent comments and thank James Stock and Juan Carlos Su´arez Serrato for both their excellent discussions and additional comments. We also thank seminar participants at the NBER Summer Institute Environmental & Energy 2024, Bank of England, Carnegie Mellon, European Commission, HEC Lausanne, Harvard University, Institute for Fiscal Studies, IMF and the World Bank, London School of Economics, Queen Mary, Paris School of Economics, Toulouse School of Economics, Tulane University, University College London, University of California San Diego, University of Chicago, and University of Oxford. We thank Doyne Farmer and Rupert Way for helpful comments and for sharing their data. This research was supported by Policy Impacts, which receives support from the Spiegel Family Fund. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

A Welfare Analysis of Policies Impacting Climate Change
Robert W. Hahn, Nathaniel Hendren, Robert D. Metcalfe, and Ben Sprung-Keyser
NBER Working Paper No. 32728
July 2024, Revised October 2025
JEL No. H0, Q5

ABSTRACT

We study the welfare impacts of 96 climate-related tax and spending policies. We extend and apply the marginal value of public funds (MVPF) framework, most notably providing a new method for incorporating learning-by-doing spillovers. We find subsidies for the production of clean energy (such as wind production tax credits) have higher MVPFs than all other subsidies in our sample, including EV subsidies. Conservation nudges have large MVPFs when targeting regions with dirty grids. Fuel taxes and cap-and-trade policies are highly efficient means of raising revenue. We also construct traditional cost-per-ton estimates and compare and contrast the lessons they provide.

Robert W. Hahn
University of Oxford
and Carnegie-Mellon University
robert.hahn@smithschool.ox.ac.uk

Nathaniel Hendren
Massachusetts Institute of Technology
Department of Economics
and NBER
nhendren@mit.edu

Robert D. Metcalfe
Columbia University
Department of Economics
and NBER
rdm2176@columbia.edu

Ben Sprung-Keyser
Wharton School
The Wharton School
bsprungk@wharton.upenn.edu

1 Introduction

What are the best ways to address climate change? There is a robust and growing literature that helps shed light on this issue by examining the causal effects of climate-related policy changes. These papers often assess the effectiveness of policies by measuring the cost per ton of carbon dioxide (CO_2) abated. Yet, comparisons of these costs per ton across studies face several challenges. First, the input assumptions in these calculations vary across papers. Second, there are at least three distinct (and often conflated) definitions of the cost per ton of CO_2 found in the literature: (1) resource costs expended per ton of CO_2 abated (Grubb et al. 1993, Enkvist et al. 2007, Mullainathan & Allcott 2010, Greenstone et al. 2022), (2) government costs per ton of CO_2 abated (Gillingham & Tsvetanov 2019, Knittel 2009), and (3) social costs per ton of CO_2 abated (Hughes & Podolefsky 2015, Fournel 2024). Third, each of these metrics focuses on the cost of achieving a given CO_2 reduction as opposed to maximizing social welfare.

It is with these considerations in mind that we extend and apply the marginal value of public funds (MVPF) framework to examine the welfare consequences of historical US spending and revenue-raising policies addressing climate change. The MVPF is a form of benefit-cost ratio in which all benefits to individuals are incorporated in the numerator (measured by the sum of their willingness to pay) and net government costs are incorporated in the denominator. All else equal, policies with higher MVPFs are generally "better" spending policies because they deliver greater welfare gains per dollar spent. Conversely, those with lower MVPFs are "better" methods of raising revenue (or reducing spending) because they impose a lower welfare cost per dollar of revenue raised.[1]

We conduct our analysis for a comprehensive set of climate-related tax and spending policy interventions in the US that affect greenhouse gas emissions. We focus on academic research that uses experimental or quasi-experimental methods to rigorously evaluate policy interventions in the past 25 years. This yields a sample of 96 policy changes in three primary categories—subsidies, nudges and marketing, and revenue raisers—along with a selected set of international aid policies. In addition to constructing an MVPF for each of these 96 policies, we also construct each of the three common cost-per-ton measures. This allows us to compare the way each metric ranks the policies in our sample.

Across our MVPF estimates, we use a consistent method to translate a policy's causal effect on behavior into a valuation of that change in behavior. We proceed in two steps. First, we use a harmonized method to translate changes in behavior (e.g., changes in car purchases or electricity usage) into changes in emissions and other damaging outcomes (e.g., car accidents). For example, in the case of changes in electricity production or electricity use, we use estimates

[1]The MVPF statistic does not depend on preferences for equity, but one can readily incorporate equity concerns when using the MVPF to make policy decisions: Given two policies, policy 1 and policy 2, a decision-maker prefers a budget neutral policy that spends more on policy 1 financed by raising revenue from policy 2 if and only if that decision-maker prefers giving $\$MVPF_1$ to policy 1 beneficiaries rather than $\$MVPF_2$ to policy 2 beneficiaries.

from the AVERT model from the US Environmental Protection Agency (EPA) to measure associated changes in emissions resulting from compositional changes in the grid (EPA 2024).

For changes to vehicle purchases (e.g., electric vehicles (EVs) versus internal combustion), we estimate the change in gallons of gasoline used relative to a counterfactual vehicle. We measure the total CO_2 emissions associated with the upstream production of gasoline and its combustion. We combine that with measures of local pollutants released such as particulate matter.

Second, we apply a consistent dollar value for each externality. For the social cost of carbon (SCC), we draw from recent work by the US EPA (EPA 2023) that places the SCC at $193 in 2020 (and rising in the years to follow). We also explore the robustness of our results to alternate measures of the SCC, ranging from $76 to $1367 in 2020. For local pollutants, we use estimates of the social cost of NH_3, HC, NO_X, $PM_{2.5}$ and SO_2 from the AP3 integrated assessment model (Tschofen et al. 2019), which monetizes health impacts from air pollution exposure using estimates on mortality and an associated value of a statistical life (VSL).

Our primary methodological contribution is the introduction of a new sufficient statistics approach to quantify the benefits of "learning-by-doing" effects. There is a large literature showing that the prices of new technologies such as solar cells, wind turbines, and batteries have declined with cumulative global production (Way et al. 2022). These patterns often serve as a proposed justification for subsidizing particular low-carbon technologies: Subsidizing specific technologies with relatively high abatement costs today may generate learning-by-doing spillovers that lower the future cost of these technologies and generate future environmental benefits (Romer 1986, van Benthem et al. 2008, Bollinger & Gillingham 2019, Bistline et al. 2023). Our approach builds most closely on the work of van Benthem et al. (2008), who develop a dynamic model of learning by doing and use it to simulate the desirability of solar subsidies in California. Our methodological contribution is to show that when the marginal cost of production is an isoelastic function of cumulative production and when demand is an isoelastic function of price, the time path of production follows a second-order ordinary differential equation that can be solved to estimate the benefits from learning-by-doing.

Learning by doing generates two types of benefits: first, reductions in the future cost of low-carbon technologies increase consumer welfare due to lower future prices, and second, these price reductions serve to increase future take-up and, consequently, reduce future emissions. We apply our framework to study the potential implications of learning by doing for policies that increase the current production of solar cells, wind turbines, and batteries. While we focus here on learning-by-doing in the context of climate change, our framework can be used in other industrial policy settings where there may be learning-by-doing externalities.

Results We have three primary empirical findings from our MVPF analysis. First, we find that subsidies for investments that directly displace the dirty production of electricity have higher MVPFs than all other subsidies in our sample. Policies providing production tax credits

(PTCs) for wind power and subsidies for residential solar have MVPFs that generally exceed 2. In contrast, subsidies to consumers that provide appliance rebates, home weatherization, vehicle retirement, or subsidies for hybrid vehicle purchases have MVPFs around 1, indicating they provide roughly $1 of benefits for every $1 of government spending. Electric vehicle subsidies have MVPFs around 1.4.

We examine the robustness of our conclusions to a wide range of input assumptions. We find that the relative ordering of policies remains consistent but that two key assumptions impact the levels of our MVPF estimates: i) the social cost of carbon and ii) the inclusion of learning by doing externalities. For example, the inclusion of learning-by-doing effects amplifies the MVPFs of wind, solar, and EV subsidies. In the case of wind PTCs, the MVPF rises from 3.85 to 5.87 with learning by doing. In the case of residential solar, the MVPF rises from a relatively low value of 1.45 to 3.86. Without learning-by-doing effects, the MVPF of EVs falls to around 1, in line with other consumer subsidies. Higher values of the SCC lead to larger MVPFs for all subsidies in our sample. Despite this heterogeneity, our main result that production tax credits have higher MVPFs than other consumer subsidies continues to hold across a wide range of SCC estimates and learning-by-doing assumptions, along with other assumptions discussed more below.

Second, we find that nudges to reduce electricity consumption can deliver large welfare gains (MVPFs exceeding 5) in regions with relatively dirty electric grids. In areas with cleaner grids like California and the Northeast, the MVPFs fall below 1.[2] This finding also suggests that the effectiveness of these nudges will decrease over time as more electricity comes from low- or zero-carbon sources.

Third, we find that taxes on polluting goods can serve as an efficient means of raising revenue, as they have low MVPFs. We analyze taxes on gasoline, diesel, and jet fuel, along with changes to the number of auctioned permits in cap-and-trade systems. We find that nearly all of these revenue-raising policies have MVPFs below 1, with most having MVPFs below 0.7. This means that taxes on polluting goods impose a welfare cost of only $0.70 on society for every $1 of revenue raised.[3] This finding reflects the logic of Pigouvian taxation, quantifying the efficiency of raising rates when current tax rates fall below the associated environmental externalities.

While our primary focus is on US environmental policy, we also consider the welfare conse-

[2]This echoes the conclusions in Borenstein & Bushnell (2022), who suggest the returns to reducing energy consumption are lowest in areas with clean grids. We find support for this conclusion in the context of energy conservation nudges, despite the fact that previous work has found treatment effects of nudges are larger in more environmentalist areas (Allcott 2015).

[3]We find even lower MVPFs that fall below zero when studying cap-and-trade policies that directly target carbon emissions. Results from the introduction of several cap-and-trade programs suggest reducing the number of auctioned permits (equivalent to raising a carbon tax) can raise revenue while providing positive net benefits for individuals in the economy (environmental gains outweigh the permit costs faced by emitters). We caution, however, that these policies may have benefited from the ability of low-hanging fruit, like removing coal power plants. We discuss these generalizability concerns in Section 6.

quences of US spending abroad on policies that address climate change. We find such subsidies have the potential to produce high MVPFs, even when only considering the impact on US beneficiaries and US taxpayers. For example, we consider the case of subsidies for the take-up of efficient charcoal cookstoves in Kenya (Berkouwer & Dean 2022). Ignoring any benefits of these stoves to local residents and ignoring any non-US benefits of CO_2 reductions, the US-specific gains from reduced CO_2 emissions are 37 times larger than the net cost of the subsidy, generating a higher MVPF than any domestic subsidy in our sample. (When considering the full set of global benefits, the MVPF rises from 37 to 323). That said, there is substantial uncertainty associated with these international subsidy estimates. The estimated impacts of these policies often vary quite extensively, even within categories of similar policies. As we discuss in Section 7, the magnitude of the US-specific MVPF depends heavily on the incidence of the social cost of carbon. In particular, it depends on the extent to which CO_2 damages have incidence on US residents and US government tax revenue.[4]

Having outlined our primary MVPF findings, we then estimate the cost per ton of CO_2 abated for each policy in our sample. As highlighted above, the notion of "cost" in cost-per-ton varies widely in the existing literature, with at least three conceptually distinct definitions being used. We refer to these as the i) resource cost per ton, ii) the government cost per ton, and the iii) social cost per ton.[5] The resource cost per ton measures the economic resources needed to abate each ton of CO_2 (Enkvist et al. 2007).[6] The government cost per ton measures the net government outlay per ton of CO_2 abated (Gillingham & Tsvetanov 2019, Knittel 2009). The social cost per ton seeks to incorporate a broader range of costs, capturing the total government cost net of any CO_2 benefits (transfers, local pollutants, etc).

The differences in these definitions of cost can yield substantively different rankings both within and across policies. For example, in the case of subsidies for energy-efficient appliances, the cost per ton values range from -$2 to $474 across the three definitions. From a resource cost perspective, the energy savings from more efficient appliances offset any higher upfront cost. This delivers a negative resource cost per ton of -$2. However, providing government subsidies for these appliances generates large transfers to those who would have purchased these appliances anyway. The cost of those transfers are omitted from the resource cost per ton, but are included in the government cost per ton, yielding a value of $474 per ton.[7]

[4]Many models that agree on the level of the social cost of carbon still differ in the geographic incidence of those damages and the split between market and non-market damages (e.g., productivity declines versus mortality impacts). The impact on US tax revenue is determined by the fraction of damages that reflects US-specific productivity changes, as the US Treasury has an equity stake in those changes.

[5]The distinction between resource, government, and social cost is not always made clear in the papers in previous literature. For example, Table 2 of Gillingham & Stock (2018) compiles a set of cost-per-ton estimates from the existing literature. The best policy listed is a behavioral nudge for reducing energy where the net resource cost of the policy is reported. By contrast, residential solar panels appear to be one of the highest cost policies in their sample, but the reported cost per ton measures the government cost of the policy.

[6]See Grubb et al. (1993) for an early summary of the resource cost-per-ton approach. See International Energy Agency (IEA 2020) and the Environmental Defense Fund (Environmental Defense Fund 2021) for more recent implementations.

[7]The social cost per ton lies in between as it nets out these inframarginal transfers as both a cost and benefit

Each of these cost-per-ton metrics can be motivated by a decision-maker seeking to minimize a particular notion of cost subject to a fixed reduction in CO_2. This contrasts with the conceptual experiment behind the MVPF, which envisions maximizing social welfare subject to a government budget constraint. We therefore examine the lessons offered by each of these cost-per-ton analyses and study how they differ from the conclusions of our MVPF analysis.

In the case of resource cost per ton, we find large variation across categories even among those with similar MVPFs. For example, appliance rebates and vehicle retirement subsidies have resource costs per ton of -\$2 and \$1008 (their MVPFs are 1.164 and 1.047). The resource costs diverge because switching to high-efficiency appliances saves resources via energy savings, while building a new car requires significant resources.[8] But in both cases, the subsidies to encourage these behavioral changes are primarily transfers to people buying those products anyway. The MVPF captures the costs and benefits of the inframarginal transfers induced by these subsidies, driving those values toward one. The resource cost per ton omits these costs and benefits, as it is better suited to evaluate abatement decisions of private actors (e.g. firms).

In the case of the government cost per ton, the relative ordering of policies is broadly consistent with the ordering generated by the MVPF. However, EV subsidies have some of the highest government cost per ton in our sample (\$1,356), while their MVPF lies above most other consumer subsidies. This is because the government cost per ton omits all non-CO_2 benefits. It is designed for a decision-maker focused exclusively on reducing emissions at lowest government cost, without regard for other welfare impacts. In the case of EV subsidies, we find that most of the benefits are non-environmental. These include both the transfer benefits to inframarginal recipients and the learning-by-doing effects that lower the cost of future vehicle purchases. They result in an MVPF value modestly above 1, despite the fact that the cost per ton exceeds the social cost of carbon used to construct the MVPF.

Finally, we consider the case of the social cost-per-ton (SCPT). Consistent with the lessons from the MVPF, the SCPT of wind PTCs, residential solar, and EVs are better than all other subsidy policies. However, the MVPF and SCPT deliver opposite orderings within these three policies (e.g., EVs have a SCPT of -\$415 in contrast to -\$32 for wind PTCs.) This reordering occurs because of a feature of the SCPT when values are negative, as such policies abate CO_2 while also providing positive non-CO_2 benefits. In such cases, the SCPT no longer retains its Lagrange multiplier interpretation, which prevents informative comparisons across policies. Increased non-CO_2 benefits make the cost per ton more negative while increased abatement makes it less negative.[9]

and adds non-CO_2 benefits like reductions in local pollutants.

[8]The resource cost of vehicle retirement policies does not include the utility value provided by the new cars. Analogously, gas taxes have negative resource costs because they discourage the use of resources. However, the welfare cost imposed on the taxed individuals is omitted from resource cost.

[9]The SCPT also generally omits the opportunity cost of raising funds. In Section 8 we discuss how this yields SCPT estimates that are independent of the behavioral response to a policy. We also discuss an alternative SCPT approach, implemented by Fournel (2024), which incorporates the shadow price of the government budget constraint using a "marginal cost of funds" adjustment. We discuss the pros and cons of this approach and

Our paper relates to a large literature in climate and environmental economics. It draws on a large set of causal effects of policy changes discussed throughout the text below and builds on a body of work conducting comparative analyses of climate policies (Gillingham & Stock 2018). In addition to the cost-per-ton analyses discussed above, our paper also relates to a large literature on benefit-cost analysis and its applications. The MVPF approach extends the traditional approach to benefit-cost analysis, which tends to compare the benefits of a spending policy to the distortionary cost of raising revenue through a change in a linear income tax rate (Stiglitz & Dasgupta 1971, Atkinson & Stern 1974). The MVPF approach allows researchers to choose from a menu of policies to close the budget constraint, instead of a linear income tax. For example, if one treats individuals paying the gas tax and wind PTC beneficiaries as having similar social welfare weights, the comparison of an MVPF of 5.87 for wind PTCs to an MVPF of 0.67 for gas taxes suggests that every $1 of government revenue raised from a gas tax and spent on wind PTCs generates $5.20 (=5.87-0.67) in net benefits to individuals in society.

The rest of this paper proceeds as follows. Section 2 discusses the MVPF framework and outlines how it can be used to examine the welfare effects of policies impacting climate change. We also discuss the definitions of cost per ton used in previous literature and how they compare to the MVPF. Section 3 discusses our sample of policies and methods for harmonizing the measurement of externalities and the valuation of those externalities. Sections 4, 5, and 6 discuss our results for subsidy policies, nudge and marketing policies, and revenue-raising policies, respectively. Section 7 discusses our findings for a limited set of international subsidies. Section 8 constructs the cost per ton measures for each policy in our sample and compares the conclusions with those provided by the MVPF approach. Section 9 concludes. The Online Appendix provides a detailed description of the MVPF and cost per ton construction for each policy in our sample.

2 Measuring the Welfare Impact of Policies Affecting Climate Change

The goal of our analysis is to translate estimates of the causal effects of policies into insights about their impacts on social welfare. We begin by extending the MVPF framework to consider environmental externalities. We then discuss several cost-per-ton definitions employed in previous literature. Finally, we present our key theoretical contribution, which is a sufficient statistics method for valuing learning-by-doing externalities that can be incorporated into both the MVPF and cost-per-ton frameworks.

illustrate how the magnitude of these estimates vary with the MVPFs of the accompanying revenue-raising policy.

2.1 MVPF Approach

For any policy change, the MVPF approach measures the benefits it provides to individuals in the economy relative to its net cost to the government:

$$MVPF = \frac{\text{Benefits to Individuals}}{\text{Net Cost to Govt}}. \tag{1}$$

One dollar of additional net spending yields $\$MVPF$ in benefits to individuals; conversely one less dollar of spending or revenue raised yields a cost of $\$MVPF$ to individuals. Benefits are measured using an individuals' private willingness to pay (WTP) and include both direct and indirect beneficiaries (e.g., future generations benefiting from lower CO_2). To translate private benefits into social welfare, we multiply the MVPF by the average social welfare weights of the beneficiaries, η, where giving \$1 to beneficiaries leads to an increase of $\$\eta$ in social welfare. In total, one dollar of spending on the policy delivers $\eta MVPF$ in social welfare benefits.

We use the MVPF by comparing across policies to form hypothetical budget-neutral policies. For any two policies, a budget neutral policy that increases spending on policy 1 financed by raising revenue from 2 increases social welfare if and only if

$$\eta_1 MVPF_1 > \eta_2 MVPF_2 \tag{2}$$

where $MVPF_j$ is the marginal value of public funds of policy $j = 1, 2$.

For example, if policy 1 has an MVPF of 1 and policy 2 has an MVPF of 2, then raising revenue from reductions in spending on policy 1 to finance increased spending on policy 2 will increase social welfare if and only if the government prefers \$2 going to policy 1 beneficiaries to \$1 going to policy 2 beneficiaries (i.e., $\eta_1 > 2\eta_2$). Reasonable people may disagree about the relative value of giving benefits to policy 1 versus policy 2 beneficiaries, but these disagreements do not affect the MVPFs. Rather, the MVPF characterizes the trade-offs faced by a decision-maker. Holding welfare weights fixed, policies with lower MVPFs are better ways of raising revenue while policies with higher MVPFs are better ways of spending government resources. In addition, when welfare weights are the same for policy 1 and policy 2 beneficiaries, the difference between $MVPF_1$ and $MVPF_2$ reveals the welfare gain to individuals in the economy per dollar spent on policy 1 using net revenue raised from policy 2.

2.2 Model

In this section, we develop a simple model to examine the welfare impact of environmental policy changes. We use the model to achieve two goals: (1) Illustrate the key ideas behind the MVPF and how it relates to cost-per-ton metrics, and (2) show how to incorporate learning-by-doing effects into our analysis. Appendix A extends this model to include various features that are important for our empirical implementation, such as multiple goods and imperfect competition.

Consider a good x that generates an environmental externality of V per unit of x consumed. For example, x may be an electric vehicle or a gallon of gasoline. Let p denote the price of x paid by consumers and let τ denote the current subsidy (or $-\tau$ is the tax) on good x such that producers receive $q = p + \tau$. The willingness to pay for a small increase in the subsidy, $d\tau$, is given by

$$WTP = xd\tau + Vdx \tag{3}$$

The first term is the monetary value of the subsidy (holding behavior fixed due to the envelope theorem), and the second term is the WTP from the change in the environmental externality. These two terms are sufficient for measuring WTP if we assume perfect competition and full pass-through (assumptions that we relax in our empirical implementation).

The term dx is the causal effect of the policy change. A subtle but important point is that this needs to include any "rebound" or general equilibrium effects of a policy that might be missed in a reduced-form analysis. For example, a subsidy for wind turbines may lead to a lower price of electricity. This lower price can cause an increase in energy consumption in the economy, which diminishes the environmental benefits of the subsidy. These types of rebound effects should be included to accurately value the externality generated by the subsidy. In Appendix D, we show how we incorporate these rebound effects using estimates of the market supply and demand curves and discuss how we apply this to account for the rebound created by upward-sloping local supply curves in the US electricity markets.

The net cost to the government of the subsidy has two terms:

$$Cost = xd\tau + \tau dx \tag{4}$$

where the first term holds x fixed and the second term captures the fiscal impact of the change in x. This cost, τdx, is paid by the government but is not valued by individuals due to the envelope theorem.

The ratio of WTP to government costs yields the MVPF for a change in τ:

$$MVPF = \frac{xd\tau + Vdx}{xd\tau + \tau dx} \tag{5}$$

$$= \frac{1 + \frac{V}{p}(-\epsilon)}{1 + \frac{\tau}{p}(-\epsilon)} \tag{6}$$

where $\epsilon = \frac{dx}{dp}\frac{p}{x}$ is the price elasticity of demand so that $-\epsilon$ is the percentage change in consumption of x in response to a 1% increase in the consumer price. The WTP for a dollar's worth of mechanical subsidy exceeds its mechanical cost by the change in consumption times the value of the environmental externality relative to the price of the good, $\frac{V}{p}$. Conversely, the net cost to the government is the sum of the mechanical cost and the fiscal externality, which is given by the change in consumption times the tax rate relative to the price of the good $\frac{\tau}{p}$. The formula shows that policies with high MVPFs tend to have higher magnitudes of the price elasticity, $-\epsilon$, higher environmental benefits per dollar of spending on the good, V/p, and lower preexisting subsidies, τ/p.[10]

A natural benchmark is the case where $\tau = V$. In this case, the government fully internalizes the externality with a Pigouvian tax (or subsidy), generating an MVPF of 1. More generally, the MVPF measures the extent to which status quo policy deviates from the Pigouvian policy and quantifies the willingness to pay per dollar of spending when moving toward that optimum.

2.3 Cost Per Ton Approaches

The MVPF approach provides guidance to a decision-maker seeking to maximize social welfare when facing a government budget constraint. An alternative approach is to consider a decision-maker facing the constraint of needing to abate a ton of CO_2 and seeking the lowest cost method of doing so. The key question with this method is what definition of "cost" to use. As noted in the introduction, several definitions have been employed, and at times conflated, in previous literature. In this section, we define three notions of "cost" that span the most common estimands in the literature: resource cost, government cost, and social cost per ton. We discuss each measure's relationship to the MVPF approach and construct these measures alongside the MVPF for each policy in our empirical analysis below.

Resource Cost per Ton The "resource cost per ton" has a long history (see Grubb et al. (1993)) and was popularized in the McKinsey cost curve (Enkvist et al. 2007). It measures the resources consumed to produce and use the product, divided by the tons of carbon it abates. For example, the resource cost of an electric vehicle (EV) is the difference in its production cost relative to a similar internal combustion engine (ICE) car minus their difference in operating costs. To capture this in the model above, let Δp denote the difference in cost of producing good x relative to its alternative (e.g., EV versus ICE vehicle), and let Δe denote the difference in operating costs (often negative due to energy savings). The resource cost per ton is then $(\Delta p - \Delta e)/T$, where T is the tons of carbon reduced from the good x relative to its alternative.

In the case where there are no alternative comparison goods and no operating cost differences (e.g. if x is a one-time purchase of a device that reduces emissions), the resource cost per ton is simply p/T.

The resource cost per ton (RCPT) may be a particularly appropriate measure if a company

[10] In the presence of firm markups (e.g., due to market power), there are additional terms in this expression. In the numerator, dx is multiplied by the firm markup net of taxes, and, in the denominator, dx is multiplied by the corporate tax revenue from firm profits.

is looking to determine the least costly way to achieve a given reduction in CO_2 emissions.[11] In contrast to the MVPF, the RCPT is independent of the causal effect of a subsidy for the product on take-up, ϵ. Generating one new purchase of an efficient appliance through a million dollars of subsidies to inframarginal beneficiaries will have the same RCPT as a subsidy that spends \$1 to generate that same new purchase. Additionally, it ignores any non-resource costs or benefits that flow to individuals in the economy. Environmental taxes have very low or negative resource costs because they discourage the use of resources. They do so by placing a burden on consumers, which is not captured in the resource cost per ton.[12]

Government Cost per Ton The "government cost per ton" of carbon abated measures the reduction in tons of CO_2 emitted per dollar of net government outlay (Knittel 2009, Gillingham & Tsvetanov 2019).[13] As shown in equation (4), the government cost is $\tau dx + xd\tau$, which generates a reduction in CO_2 of Tdx. Taking ratios and re-arranging terms, the government cost per ton is given by

$$GCPT = \frac{\frac{p}{-\epsilon} + \tau}{T} \tag{7}$$

The first term, $p/(-\epsilon)$, is the inframarginal transfer the government makes to induce an additional purchase of x; the second term, τ, is the cost to the government from the marginal change in purchases.

Government cost per ton (GCPT) is the right conceptual measure for a government to use if it is only concerned with reducing CO_2 emissions. As a result of this, it omits all non-CO_2 benefits. This means that the GCPT can actually exceed the social cost of carbon even if a policy provides meaningful welfare gains.[14] As is the case with the resource cost per ton, the GCPT also omits the welfare costs placed on individuals as a result of environmental taxes.

Social Cost per Ton The third cost-per-ton measure found in the literature subtracts all non-CO_2 benefits from government costs to form its net non-CO_2 social costs (Christensen et al. 2023, Hughes & Podolefsky 2015). We refer to this measure as the "social cost per ton," or SCPT.

[11]Technically, this alignment requires the absence of taxes or subsidies, which could generate a wedge between private costs and total resources.

[12]Similarly, individuals may optimize over non-resource considerations. Individuals might prefer to drive ICE vehicles to EVs even if the EVs cost less to operate. Inducing an individual to switch to an EV does not necessarily generate a private welfare gain for that individual even if it produces a resource saving.

[13]This measure is also sometimes referred to as the "program cost per ton" (Gillingham & Tsvetanov 2019, Davis et al. 2014).

[14]This particular criticism has been expressed in previous literature. For example, Davis (2023) provides a discussion of the cost effectiveness of heat pumps and notes "[i]t is tempting to compare the [cost per ton of CO_2 estimates] to estimates in the literature for the social cost of carbon. For example, the U.S. government currently uses a social cost of carbon of \$51 per ton (U.S. Interagency Working Group, 2021) and one recent study finds a preferred social cost of carbon of \$185 per ton (Rennert et al. 2022). However, this is not an apples-to-apples comparison. Subsidies are transfers, not economic costs, and many households value subsidies at close to \$1-for-\$1." A similar criticism can be found in Knittel (2009).

Conceptually, the goal of the social cost per ton is to capture all costs and benefits associated with abating a fixed quantity of CO_2. So, if a government's goal is to reduce a fixed quantity of CO_2, it provides a social-welfare-based metric by which to judge the efficacy of that tonnage reduction. To calculate the social cost per ton in our model, we can express the total externality, V, as the sum of carbon and non-carbon benefits: $V = SCC * T + V_{NonCO_2}$, where V_{NonCO_2} is the value of non-CO_2 externalities per unit of x. In addition, we let u_x denote the private benefits provided by a unit of x to individuals. The social cost of the policy (excluding CO_2 benefits) is given by the sum of private and social net costs:

$$\text{Social Cost} = xd\tau + (p - u_x)dx - xd\tau - V_{NonCO_2}dx \tag{8}$$
$$= \tau dx - V_{NonCO_2}dx \tag{9}$$

where the second line invokes private optimization so that $p - u_x = \tau$. Taking the ratio relative to the tons of carbon abated, the social cost per ton is given by:

$$SCPT = \frac{\tau - V_{NonCO_2}}{T} \tag{10}$$

The magnitude of the subsidy minus any non-CO_2 benefits measures the magnitude of the social cost of inducing additional consumption of x.[15]

Although this metric is designed to capture the social costs of each ton of CO_2 abated, the opportunity cost of raising funds is generally not considered. This is because inframarginal transfers, $xd\tau$, cancel out as both a cost and a benefit. As was the case with the RCPT, this means that the SCPT for a subsidy for x does not depend on the causal effect of the subsidy on the consumption of x. If two policies induce the same dx, they have the same SCPT, regardless of how many inframarginal beneficiaries receive the transfer, $xd\tau$.

It is worth noting that there is an alternative form of SCPT that is used in Fournel (2024) but has not yet seen broader adoption in the literature measuring costs per ton. This approach seeks to incorporate the opportunity costs of raising revenue by assigning a welfare cost of $\phi \geq 1$ to raising government revenue. (Here, ϕ corresponds to the MVPF of the policy used to raise the revenue to form the budget-neutral policy.) The SCPT becomes:

$$SCPT_\phi = \frac{\phi\tau + (\phi - 1)\frac{p}{-\epsilon} - V_{NonCO_2}}{T}. \tag{11}$$

The elasticity, ϵ, no longer drops out of the expression, as it is needed to measure the cost of the raising revenue. It is straightforward to show that if $SCC = SCPT_\phi$, then the MVPF of a subsidy for x is equal to ϕ (instead of 1 for the standard SCPT). In this sense, testing

[15]When τ is set at its Pigouvian level, V, the SCPT is equal to the SCC. In this sense, the SCPT is related to the MVPF: We have $SCPT = SCC$ if and only if the $MVPF = 1$.

for whether $SCPT_\phi > SCC$ is equivalent to testing whether $MVPF > \phi$ – i.e., whether the MVPF of the spending policy exceeds the MVPF of raising the revenue.

Determining the right ϕ, however, is no easy task. Although the most common value assumed is the cost of raising the income tax, a central lesson from public economics over the last 50 years is that there is no single ϕ even when considering a change in the income tax schedule (Mirrlees (1971); Mirrlees (1976); Saez (2001); Kleven & Kreiner (2006); Jacobs et al. (2017); Bourguignon & Spadaro (2012); Hendren (2020)). Raising revenue from the bottom of the income distribution can generate values of ϕ around 1; raising revenue from the top of the income distribution generally yields estimates of ϕ around 1.5-2 (Hendren 2020), and potentially even higher than 2 using recent elasticity estimates from Kleven et al. (2024). In Section 8, we show how the value of the SCPT is influenced by the assumed value of ϕ of the policy used to close the budget constraint.[16]

One final note on the SCPT approach is that one should be cautious when interpreting the ordering of policies with negative SCPTs. Holding net social costs fixed, policies that abate more CO_2 actually have higher (less negative) SCPT. This is related to the fact that the SCPT no longer retains its interpretation as a Lagrange multiplier on a tonnage constraint for a decision-maker seeking to minimize social costs. We return to this issue below in Section 8, in particular when comparing the SCPT of EV and Wind subsidies.

Summary In summary, the MVPF measures the benefits per dollar of government spending. This estimand aligns with our goal of finding policies to maximize social welfare subject to a government budget constraint. The cost-per-ton approach focuses on the cost of achieving a given level of CO_2 reduction. The key question when constructing cost per ton is what definition of "cost" to use, as each of these definitions correspond to different objectives (minimize resources, government expenditures, or social costs). While we focus our primary analysis on the MVPF, we also construct each of the cost-per-ton metrics to study the lessons they provide and how they differ from the MVPF approach.

2.4 Learning by Doing

A key rationale for many clean energy subsidies and other environmental policies is learning by doing: subsidies that increase demand today can lower the future marginal cost of new technologies (Acemoglu et al. 2012, Bistline et al. 2023). Industries, particularly those characterized by rapidly changing technologies, may learn as a result of their production experience. The model environment above considers only the static environmental externality, V, from the purchase of a good, x. Here, we use a dynamic version of this model featuring learning by doing

[16] One way to think about the tradeoff between the MVPF and the SCPT in this context is that the MVPF requires making an assumption about the social cost of carbon and the SCPT requires making an assumption about the cost of raising revenue.

to show how the associated externalities can be measured.[17]

Existing evidence suggests that learning-by-doing effects may be present in the production of solar cells, wind turbines, and batteries. Appendix Figure 1 reproduces evidence from Way et al. (2022) showing the relationship between the marginal cost per kW of wind and solar energy (and per kWh of battery storage) plotted against cumulative production. Their analysis shows that a 1% increase in cumulative solar production is associated with a 0.319% reduction in price. For wind and EV batteries, the associated price reductions are 0.194% and 0.421%, respectively. If one believes that these patterns reflect causal spillovers from learning by doing,[18] to what extent should that change their views about the welfare effects of subsidies for those goods?

The contribution of this section is to provide a new sufficient statistics result that measures the benefits from learning by doing.[19] We show that when the marginal cost of production is an isoelastic function of cumulative production and when demand is an isoelastic function of price, this leads to a second-order ordinary differential equation that can be solved to estimate society's willingness to pay for the learning-by-doing effects. Intuitively, subsidizing production today causes the economy to "move forward in time," lowering the cost of production. Thus, future consumers pay lower prices, and this price reduction generates additional environmental benefits from further purchases of the the clean good.

To illustrate this, we return to our example of a subsidy for a good, x. We bring the model into a continuous time environment, where time is indexed by $t \geq 0$. We imagine the subsidy of interest is a short-term subsidy enacted at time t^*. We wish to incorporate the welfare benefits accruing in future periods, $t > t^*$. Let $x(t)$ denote consumption of x at each time t and let $X(t) = \int_0^t x(s)ds + X(0)$ denote cumulative production through time t. Motivated by the historical evidence in Appendix Figure 1, which is also known as "Wright's Law" (Wright 1936), suppose that the marginal cost of production at each point in time is an isoelastic function of cumulative demand,

$$c(X(t)) = \kappa X(t)^{\theta} \tag{12}$$

where $\theta < 0$ is the elasticity of marginal cost with respect to cumulative production. Suppose also that the choice of $x(t)$ at each point in time depends on the price with a constant price

[17]For simplicity of exposition, we omit some features of the environment that we include in our empirical implementation, such as firm markups. Appendix B provides this more general model that we use in our empirical implementation.

[18]The extent to which the curve represents learning spillovers has been debated (Nemet 2006, Nordhaus 2014*b*, Rubin et al. 2015). See Lafond et al. (2022) for an estimate of the causal impact of learning by doing on military production. In the context of this paper, we take these learning-by-doing effects as given and then show the robustness of our results to the omission of learning-by-doing effects. There is quasi-experimental work that has found evidence of potential spillovers in solar production (Banares-Sanchez et al. 2023) and in wind installations in California (Gillingham & Stock 2018). We supplement this with additional analysis in Appendix Table 2. It shows the learning patterns largely hold even after controlling linearly for time and contemporaneous production.

[19]Our approach relates to work by van Benthem et al. (2008), who develop a dynamic model of learning by doing, and Bistline et al. (2023), who incorporate learning by doing into their assessment of taxes and subsidies.

elasticity of demand, $\epsilon < 0$,

$$x(t) = ap(t)^{\epsilon}. \tag{13}$$

This specification assumes the elasticity of demand is sufficient for understanding how future demand will evolve as prices fall from learning by doing. Finally, assume for expositional purposes that there is perfect static competition at all points in time and no future subsidies so that prices are set equal to marginal cost, $p(t) = c(X(t))$. Appendix B shows how we relax this assumption in practice and allow for firm markups.

Learning by doing generates two types of externalities: a price externality and an environmental externality. The price externality arises because an increase in production of $x(t)$ today (e.g., at time $t = t^*$) will generate consumer surplus via a reduction in prices faced by future customers (at time $t > t^*$). Let $dp(t)$ denote this impact on prices at each time t. The envelope theorem implies that the WTP for the price decline at each time t is given by $-dp(t)x(t)$, where $x(t)$ is the planned consumption at time t. In other words, the welfare gain is given by the price reduction times the counterfactual path of consumption in the absence of the subsidy.[20] The environmental externality arises because the price reduction caused by the subsidy will increase future consumption of the good, $dx(t)$, and, consequently, generate a positive environmental externality. This externality is given by $V_t dx(t)$, where we now introduce a t subscript to allow the environmental externality to vary over time. For example, this allows the SCC to increase or the cleanliness of the electrical grid to improve over time. The key to measuring our two externality terms is that we need to know how much prices decline, $dp(t)$, and how much consumption increases, $dx(t)$, in response to an increase in consumption of x today (e.g., at time t^*). With those terms in hand, we can then integrate over all the future price benefits, $-dp(t)x(t)$, and environmental benefits, $V_t dx(t)$, over time $t > t^*$.

How can we use this setup to measure the future price and quantity impacts of a policy that increases demand today? Our analysis relies on two key insights. First, we know that the impact of a subsidy $x(t)$ at some time, t^*, will affect future prices proportional to the amount that it increases cumulative production. While this effect can be mathematically complicated, the use of an autonomous supply and demand system allows us to re-frame the problem: we can think of the subsidy as moving us forward in time by some amount, dt. That shift in time is proportional to the size of the subsidy and the magnitude of the demand response when the subsidy is operating at time t^*.

[20] We assume learning by doing provides knowledge externalities to the entire market. It could be that learning by doing occurs within firms and is fully internalized. In that latter case, a subsidy might have no learning-by-doing price benefits for consumers. Moreover, learning-by-doing externalities are different from economies of scale, which are about reducing the fixed costs of production. As Borenstein (2012) notes, this difference might have important implications for public policy. In our modeling, we provide an optimistic interpretation of current subsidies lowering future costs through learning-by-doing externalities. In particular, we assume no internal capture of learning-by-doing benefits and no economies of scale, although this assumption has been questioned in the solar and wind industries (Nemet 2006, Söderholm & Sundqvist 2007). Such concerns would dampen the magnitude of the true learning-by-doing benefits we estimate using our approach, but as we discuss below, this would not affect our core empirical lessons.

Moving forward in time lowers marginal costs at each point in time (and thus prices) by $dp(t)$, given by

$$
\begin{aligned}
dp(t) &= c'(X(t))X'(t)dt && (14)\\
&= c'(X(t))x(t)dt && (15)\\
&= \kappa\theta X(t)^{\theta-1}x(t)dt && (16)
\end{aligned}
$$

Also, moving forward in time leads to a change in consumption of the good given by $dx(t) = X'(t)dt$.

Our second insight is that our demand and cost equations imply that the future time path of $x(t)$ is the solution to a second-order autonomous ordinary differential equation. To see this, note that $\log(x(t)) = \log(a) + \epsilon\log(p(t))$ and $\log(c(t)) = \log(\kappa) + \theta\log(X(t))$. Totally differentiating yields

$$
\begin{aligned}
d\log(x(t)) &= \epsilon d\log(p(t)) && (17)\\
&= \epsilon d\log(c(t)) && (18)\\
&= \epsilon\theta d\log(X(t)) && (19)\\
& && (20)
\end{aligned}
$$

Noting that $X'(t) = x(t)$ and the formula for the derivative of logs yields

$$\frac{X''(t)}{X'(t)} = \epsilon\theta\frac{X'(t)}{X(t)} \tag{21}$$

which is a second order autonomous ODE that we show has a closed-form solution. Combining these two insights leads to the core result in Theorem 1.

Theorem 1. *(Learning by Doing). Let the marginal cost be given by equation 12 and demand be given by equation 13. Suppose prices are set at marginal cost in all periods. Then, the willingness to pay for the future change in prices, DP, is given by*

$$DP = \frac{\theta\epsilon}{1-\theta\epsilon}(t^*)^{-\theta\frac{(1+\epsilon)}{1-\epsilon\theta}}\int_{t^*}^{\infty} e^{-\rho(t-t^*)}t^{-1+\theta\frac{1+\epsilon}{1-\epsilon\theta}}dt \tag{22}$$

where

$$t^* = \frac{X_{init}}{x_{init}(1-\epsilon\theta)} \tag{23}$$

is the normalized ratio of cumulative to flow production at the time the subsidy is enacted. Given the time path of the value of environmental externalities, V_t, the willingness to pay for *future environmental benefits, DE, is given by*

$$DE(\{V_t\}) = -\frac{\epsilon^2\theta}{(1-\epsilon\theta)c(X(t^*))} t^{*-\frac{\epsilon\theta}{1-\epsilon\theta}} \int_{t^*}^{\infty} e^{-\rho(t-t^*)} t^{\frac{2\epsilon\theta-1}{1-\epsilon\theta}} V_t dt \tag{24}$$

Combining, the MVPF of a subsidy at time t^* is given by

$$MVPF = \frac{1 + \frac{V}{p}(-\epsilon) + DP + DE}{1 + \frac{\tau}{p}(-\epsilon)}. \tag{25}$$

The present value of tons of carbon removed from the policy change is now $T+DE(\{SCC_t\})/SCC_0$. And the social cost per ton now subtracts the price and non-CO_2 benefits, $DP+DE(\{V_{NonCO_2}\})$, so that the social cost per ton now becomes

$$SCPT = \frac{\tau - V_{NonCO_2} - DP - DE(\{V_{NonCO_2}\})}{T + DE(\{SCC_t\})/SCC_0} \tag{26}$$

where the numerator contains the price and non-CO_2 benefits while the denominator adds the PDV of future carbon reduction.

Proof: See Appendix B.

This theorem illustrates how one can incorporate learning-by-doing externalities into a welfare analysis of subsidies or taxes. Calculating these new terms, DE and DP, requires four inputs: (1) the elasticity of demand with respect to price, ϵ, (2) the elasticity of marginal cost with respect to cumulative production, θ, (3) cumulative production at the time of the subsidy $X(t^*)$, and (4) product cost at the time the subsidy, $c(X(t^*))$. The first and fourth terms, ϵ and $c(X(t^*))$, are generally necessary for the construction of the static MVPF, indicating that only two new terms, θ and $X(t^*)$, are needed to construct these learning-by-doing welfare estimates. We use estimates of historical sales numbers to construct $X(t^*)$ and use estimates from Way et al. (2022) of the relationship between cumulative production and price to construct our cost curve parameter θ. The price elasticities, ϵ, come directly from each paper in our sample.

In our analysis below, we incorporate these learning-by-doing effects into our estimates for the MVPFs and cost-per-ton estimates of subsidies for wind, solar, and electric and hybrid vehicles (and the indirect effects of gasoline taxes on EVs).

3 Data and Sample

3.1 Sample

We analyze the welfare impact of 96 US spending and revenue-raising policies that affect greenhouse gas emissions and have been rigorously evaluated in the last 25 years using quasi-experimental or experimental methods. These policies span subsidies, nudges, and revenue raisers. As detailed in Appendix E, we form our sample from the full set of articles in 18 major journals in economics from January 1999 through December 2023, and supplement that with articles cited within these papers (a.k.a. a "snowball" sample). Within the category of subsidies, we analyze seven sub-categories: wind production tax credits, residential solar subsidies, electric vehicle subsidies, hybrid vehicle subsidies, vehicle buyback rebates, energy efficiency subsidies, and weatherization subsidies. Within the category of revenue raisers, we analyze four sub-categories: gasoline taxes, other fuel taxes (such as jet fuel and diesel taxes), other revenue raisers (including the California Alternative Rates for Energy), and cap-and-trade policies.[21] We also supplement this sample with a selected set of international policies that have been evaluated in the past ten years.

Appendix Table 1 presents a list of all of our policies. For each policy, we list the category, sub-category, year(s) of implementation, location of implementation, and the paper(s) estimating its causal effects. In certain cases, we observe some, but not all, of the relevant inputs necessary to construct an MVPF. In those instances, we provide an MVPF for the policy (under assumptions outlined in each policy's appendix) but only include it in our "extended" sample (denoted by "*" in Appendix Table 1). Extended sample policies are excluded from any category averages reported in the paper.

Publication Bias Our analysis is inevitably constrained by the set of studies available in the literature. This sample is potentially biased due to the fact that statistically significant studies are more likely to be published. In Appendix F, we present evidence of modest publication bias in the environmental economics literature: estimates are roughly two times more likely to be published if they cross a t-stat of around 2. To assess how this could impact our broad conclusions, we use the methods of Andrews & Kasy (2019) to correct for publication bias. This leaves our estimates largely unchanged and our conclusions unaffected. We present the unadjusted estimates throughout the remainder of the paper.

In-Context versus Baseline MVPFs For each policy change in our sample, we form two conceptually distinct welfare estimates. First, we consider the welfare impact in the context (year and location) in which the policy change occurred. For example, if we have estimates from an EV subsidy program in California in 2014, we use measures of the CA electric grid in 2014 to quantify the externalities due to reductions in gasoline usage offset by increased electricity use. We use the CA gasoline tax rate in 2014 to quantify the lost state government revenue from reduced gas purchases. These "in-context" measures capture the welfare impact of the policy as it was enacted.

[21]While regulatory policies are generally beyond the scope of our analysis, Appendix G includes analysis of CAFE standards and renewable portfolio standards. The Appendix outlines how to use the MVPF to compare a regulatory policy to a distributionally-equivalent tax and spending policy. A detailed assessment of this approach is left to future work, but this comparison provides a potentially unified treatment of tax/expenditure and regulatory policy within the MVPF framework.

Second, we construct the welfare impacts for each policy as if it was implemented nationally in the US in 2020. The key assumption required for this exercise is that the original elasticity estimated in each paper would also determine the behavioral response to federal policy in 2020. We then apply 2020 measures of the tax rates and values of externalities to measure the environmental and fiscal externalities from the policy. This approach harmonizes welfare comparisons across policies holding the contextual environment fixed. We refer to this as our "baseline" specification.

3.2 Valuing Environmental Externalities

We seek to apply a consistent and comprehensive method to value the range of externalities generated from each policy. We discuss these valuations briefly here and provide further details in Appendix C.

Greenhouse Gas Emissions CO_2 is a key greenhouse gas contributing to climate change. Our baseline estimates place a monetary cost on CO_2 emissions following the Environmental Protection Agency's 2023 guidance regarding the social cost of carbon at a 2% discount rate (EPA 2023). This model implies that the social cost of carbon is $193 per ton of emissions in 2020 and increasing over time. We also show the robustness of our results to other values of the SCC, including $76, $337, and $1367.[22]

We use the time path of the SCC to measure the environmental externality from each policy when calculating the MVPF. For example, a subsidy that leads to the installation of a wind turbine in 2020 will reduce emissions from 2020 through 2045. We use the year-specific SCC to value the associated externalities. For consistency, we apply the 2% discount rate to translate costs and benefits into 2020 present-value dollars. In addition to CO_2, we also incorporate costs from other greenhouse gases where available, including methane (CH_4), nitrous oxide (N_2O), carbon monoxide (CO), and hydrocarbons (HC). These valuations comprise a much smaller fraction of the greenhouse gas benefits we estimate and are discussed in detail in Appendix C.

There are three key things to note about our approach to quantifying the value of reducing greenhouse gas emissions. First, we require the SCC to be the sum of individuals' *private* willingnesses to pay for reduced CO_2 emissions. This is consistent with typical Integrated Assessment Models (IAMs). RICE and DICE focus on GDP or GDP-equivalent damages, which correspond to private measures of damages. Other IAMs, such as the GIVE model, infer an SCC from VSL estimates and use private VSLs that are not adjusted with welfare

[22] The 2% discount rate is the typical approach assumption in the environmental economics. literature (Nesje et al. 2023). The SCC of $193 in 2020 aligns closely with several other estimates from integrated assessment models (IAMs), such as the GIVE model in Rennert et al. (2022). The $76 (calculated with a 2.5% discount rate) SCC comes from Interagency Working Group (2021) and represents the largest SCC estimate for 2020 presented in earlier guidelines. The $337 (calculated with a 1.5% discount rate) represents the largest SCC for 2020 reported in the EPA's most recent guidelines (EPA 2023). The $1367 is the main point estimate from Bilal & Känzig (2024).

weights. These models also generate an SCC that corresponds to a private willingness to pay. By contrast, some have proposed equity-weighted social costs of carbon that adjust for welfare weights when forming the SCC (Prest et al. 2024). While the MVPF framework allows for equity weights, such weights are most appropriately excluded from the MVPF and instead applied ex-post when making policy comparisons, as in equation (2).

Second, the SCC embeds within it a real discount rate (2% in our baseline case) that captures the real cost to society of moving resources across periods. The application of this discount rate normalizes the willingness to pay in units of 2020 dollars for all comparisons, even across future generations. This discount rate does not, however, make any claims about the decision-maker's preferences across time. If a decision-maker places greater (or lower) weight on future generations, they will simply place a higher (lower) social welfare weight on those future beneficiaries. In the context of equation (2), this represents a modification of η to reflect weights on future generations.

Third, our calculations rely on estimates of the incidence of the social cost of carbon. In particular, the MVPF (and government cost per ton) requires isolating the portion of the social cost of carbon that is borne by the government. Correctly calculating these components requires identifying the incidence of the SCC. To account for this, in our baseline specification we assume a US incidence of 15%, following the US share of GDP in the global economy. This also corresponds to the assumption made in many models such as DICE (Nordhaus 1993).[23] Within this 15%, we assume in our baseline specification that 50% of this valuation is the result of changes in productivity that have direct effects on tax revenue (e.g., due to changes in agricultural productivity). We assume a tax rate of 25.54% as this is the 2020 tax-to-GDP ratio for the US (OECD 2021). (This captures both corporate and individual (labor income) taxes.) These numbers together imply that $13\%(= 15 - .2554 * .5 * 15)$ of the incidence from changes in carbon emissions falls directly on US residents while just under $2\%(= .2554 * .5 * 15)$ falls on the US government as changes in tax revenue. As it turns out, accounting for this fiscal externality has no bearing on any of our results for domestic subsidies, nudges, or revenue raisers. By contrast, the US-specific fiscal externality can get quite large for international policies. In that section, we analyze the robustness of our conclusions to those incidence assumptions.

Local Pollutants While greenhouse gases yield global externalities, a range of local pollutants produce negative health effects on individuals near the source of emissions. In order to value these externalities, we use the AP3 integrated assessment model (Tschofen et al. 2019), which measures the marginal health impacts of additional emission of NH_3, HC, NO_X, $PM_{2.5}$, and SO_2 in each county in the US. When measuring the local pollution externality from increased electricity usage, we take county-level damages estimated in AP3 and weight by fuel

[23]Other IAMs explicitly measure the distributional incidence of global damages. For example, Nordhaus (2014*a*, 2017) notes that the three models from the Interagency Working Group (Interagency Working Group 2021) on the social cost of carbon report US incidences of 10% for RICE2010 (Nordhaus 2010), 17% for FUND2013 (Anthoff & Tol 2010, 2013*b*,*a*), and 7% for PAGE2011 (Hope 2006, 2008).

consumed for electricity generation. When considering increased gasoline vehicle usage, we weight by county-level total vehicle miles traveled. Following (Tschofen et al. 2019), we monetize those health impacts using a VSL of $9.5 million (EPA 2010).

From Causal Effects to Externalities For each policy in our analysis, we translate its causal effect (e.g., purchases of EVs in response to subsidies) into the externalities it generates (e.g., the various pollutants discussed above) using a consistent approach across all policies. For example, consider policies that alter electricity usage. Some of these policies, such as residential solar subsidies, might generate new sources of electricity. Other policies, such as rebates for energy-efficient appliances, might reduce existing electricity usage. In order to identify the change in emissions from changes in electricity generation, we use estimates from EPA's Avoided Emissions and Generation Tool (AVERT) (EPA 2024). This provides year- and location-specific estimates of marginal emissions rates per kWh of electricity generated. We then use projections of the grid composition from Jenkins & Mayfield (2023) to measure future emissions. We also consider a class of policies that affect vehicle usage and gasoline consumption. In those cases, we estimate the change in gallons of gasoline used relative to a counterfactual vehicle. We measure the total CO_2 associated with the production and combustion of gasoline. We draw upon estimates from National Emissions Inventory, the Inventory of U.S. Greenhouse Gas Emissions and Sinks, as well as the EIA's reported CO_2 emissions coefficients. We describe these estimates in detail in Appendix C.

Appendix Figure 2 presents the environmental damages over time from driving and using electricity. Panel A presents the dollar value of the local and global externalities generated per gallon of gasoline used by the average light-duty, gasoline-powered vehicle. It shows that average non-CO_2 emissions have declined over the last several decades, and there has been a shift in the share of total pollution externalities driven by CO_2 emissions.[24] Panel B reports average emissions from the electricity grid over time. It shows a gradual reduction in emissions as more clean energy (and lower-carbon energy) has come online. This is supplemented by evidence in Panel C, which shows the geographic variation across the US in emission externalities, as measured in 2020. The Northeast and California have the cleanest grids (lowest environmental externality per mWh), and the Midwest and has the dirtiest electric grid. We discuss below how this leads to heterogeneity in the welfare impacts of policies that are targeted to different regions of the US.

[24]The graph also includes the impact of other vehicle externalities – congestion and accidents. For vehicle accidents, we use results from Jacobsen 2013, who estimates that a 1% reduction in vehicle miles traveled leads to 263 fewer fatalities in the US. We again apply a VSL of $9.5 million to yield a $0.08 per-mile externality. For congestion due to light-duty vehicles, we take an average of externality measures from Parry & Small (2005), Parry et al. (2014), and Couture et al. (2018) to yield an externality of $0.03 per mile.

4 Subsidies

The next four sections of the paper present the MVPF results for subsidies, marketing and nudges, revenue raisers, and international policies. We begin with a detailed description of the way in which we construct MVPF estimates for EV subsidies. We use this example because it utilizes nearly all of the machinery we develop to construct environmental MVPFs and cost-per-ton metrics.

Subsidies for Electric Vehicles We begin with a discussion of the MVPF for EV subsidies using estimates from the California Enhanced Fleet Modernization Program (EFMP) studied in Muehlegger & Rapson (2022). The authors use zip code variation in subsidy eligibility to estimate a price elasticity of demand of -2.1. They also find that 85% of the subsidy is passed to consumers while 15% is captured by dealers via higher prices. To evaluate the MVPF, we consider a $1 increase in EV subsidies. Figure 1 outlines each of the components of our MVPF estimate.

To begin, consumers who would have bought EVs anyway are willing to pay $0.85 for the $1 higher subsidy, while dealers are WTP $0.15. The incidence does not affect the overall WTP but can affect the average welfare weight, η, one might assign to the beneficiaries of the policy.

The $1 subsidy generates environmental externalities from the increase in EV purchases. We use estimates from Holland et al. (2016) to calculate the fuel economy of the counterfactual car that a marginal EV customer would have purchased. (This is generally cleaner than an average new car.) We combine this with an estimate of vehicle miles traveled for EV purchasers from Zhao et al. (2023) to measure the reduction in CO_2 from gasoline consumption. Valuing this using our SCC model of $193 in 2020 yields a welfare gain of $0.17.

The environmental gains from reduced gasoline consumption are partially offset by the emissions generated from increased electricity consumption over the life of the EV. We use estimates from AVERT (EPA 2024), combined with future grid forecasts from Princeton REPEAT Project (Jenkins & Mayfield 2023) to measure the current and future CO_2 emissions required to power the EVs. We also account for the fact that greater electricity use will raise prices – a so-called "rebound" effect. Using a demand elasticity of -0.19 and a supply elasticity of 0.78 from (DOI 2021) suggests that roughly 20% of the increased electricity demand from EVs displaces other uses of the electricity due to higher prices. On net, this suggests the increased electricity use generates $0.10 in global damages.

Finally, the production of the EV is generally more carbon intensive than an ICE vehicle due to the battery production. We incorporate estimates from Winjobi et al. (2022) that suggest that battery production releases 0.06 tons of CO_2 per kWh, which corresponds to a CO_2 externality of $838 per EV or -$0.03 per dollar of EV subsidy.

In addition to CO_2, we also incorporate the benefits and costs of local pollutants, including

NO_X, $PM_{2.5}$, HC, CO, SO_2, and NH_3. The reduction in gasoline consumption generates benefits of around \$0.02. This is nearly perfectly offset by the increase in emissions from electricity production, for a net gain of less than \$.01. Summing across all initial environmental externalities, we obtain a welfare gain of \$0.07.

This \$0.07 is the initial environmental benefit of the \$1 subsidy. The next bars in Figure 1 present the estimated externalities due to learning by doing in battery production. Way et al. (2022) estimate that a 1% increase in cumulative battery production leads to a reduction in battery costs of 0.42% ($\theta = -0.42$). Combining this with the demand elasticity of $\epsilon = -2.1$ suggests that the increased future demand for EVs yields environmental benefits of $DE =$ \$0.04 per dollar of the mechanical subsidy and benefits from lower future prices of $DP = 0.31$. We again note that the inclusion of these benefits relies on the assumptions that the reductions in battery costs are (a) the causal effect of cumulative production and are (b) generating spillovers across firms (i.e. the gains are not internalized within firms or in patent markets). Thus, we present results below with and without these learning-by-doing terms.

The last components of WTP in Figure 1 consider the willingness to pay from changes in firm profits. A change in the mix of goods consumed in the economy generates changes in profits to the extent to which markups vary across goods (Kaplow 2023). We assume EVs and ICE vehicles have the same markups, so that shifts towards EVs do not directly generate changes in producer profits. We do incorporate differential markups among gasoline and energy suppliers. Combining estimates from EIA (2024*c*,*b*,*a*), Favennec (2022), and EIA (2022), we calculate an average markup per gallon of gas of \$0.61 per gallon, or 27% of the 2020 retail price. Estimates from EIA (2023), Wiser et al. (2023), and Bolinger et al. (2021) suggest a markup on electricity of 20.9%. Both values exceed the average 8% markups found in De Loecker et al. (2020).[25] Applying these net markups to the change in electricity and gasoline consumption, we estimate a WTP of \$0.01 per \$1 of subsidy. Summing across all of the WTP components yields a WTP of \$1.38 per dollar of mechanical subsidy.

Next, we turn to the denominator of the MVPF. The government cost of the subsidy sums the \$1 mechanical cost and the fiscal externalities resulting from changes in EV purchases. The first such cost is the added cost from pre-existing EV subsidies. In 2020, the average federal credit for an EV purchase was \$43 (Tesla had exhausted its subsidies in 2020) while the average state subsidy was \$604. The increased EV purchases thus generates a fiscal externality of \$0.001 from the federal subsidy and \$0.02 from an average state subsidy. In addition, the reduced gasoline consumption generates a loss of \$0.04 in gas tax revenue and a loss of \$0.01 in corporate tax revenue per dollar of subsidy.

Finally, we incorporate a positive impact on the US government's budget due to reductions in climate damages. According to a wide class of IAMs, the SCC is driven by a combination of health and productivity effects. These productivity effects can have a direct effect on US

[25] These markups also generate changes in government cost due to corporate tax revenue and the fact that 28% of utilities are publicly owned. We incorporate these in our net cost estimates.

government revenue. In our baseline specification, we assume that half of the SCC is due to productivity effects and that 15% falls on the US economy (proportional to its share of global GDP). Applying a 25.5% tax rate to these productivity gains yields a fiscal externality equal to $0.003 for every $1 in subsidies. These "climate fiscal externality" effects are quite small for all domestic policies in our sample, but we return to them in Section 7 when we analyze the MVPFs of international policies.

Adding these costs together, we estimate a net cost of $1.07 for every $1 in mechanical subsidy costs. When we take the ratio of the willingness to pay and the net cost, we arrive at a baseline MVPF of 1.30. One dollar of net government expenditures on EV subsidies in 2020 generates $1.30 in benefits to individuals in society.

The bars in Figure 1 also shed light on the incidence of EV subsidies: 95% of the benefits of the government expenditure flow to individuals buying and selling EVs, while just 5% of the benefits flow to the rest of society through an improved environment. Most of these benefits to buyers and sellers are inframarginal transfers, resulting in an MVPF not much above 1. Inducing a new EV purchase costs the government roughly $30,000[26], much larger than the environmental and learning-by-doing benefits of the subsidy.

Table 1 presents the baseline MVPFs for the rest of the subsidy policies in our sample.[27] For EVs, we consider two other estimates from the literature – Clinton & Steinberg (2019) who find a price elasticity of -2.93 and Li et al. (2017) who find a price elasticity of -2.61.[28] These lead to MVPFs of 1.56 and 1.47, respectively. We then construct a category average MVPF by envisioning a policy that splits $1 in upfront spending equally across the three policies. This means we compare the average willingness to pay per dollar of upfront spending to the average net cost per dollar of upfront spending. This yields a category average MVPF of 1.45. We can also use the standard errors of each price elasticity in our analysis to construct bootstrapped confidence intervals. This yields a 95% confidence interval of [1.23, 1.94] for EV subsidies. Our confidence intervals for each policy are presented in Appendix Table 4.

Wind Production Tax Credits We find MVPFs for wind production tax credits (PTCs) that generally exceed 4, which is well above our estimates for EV subsidies. These subsidies pay producers a fixed payment per kilowatt hour of production of clean energy, typically for ten years after installation. We draw upon three papers estimating the elasticity of wind turbine investment with respect to these production tax credits in the US: Hitaj (2013), Metcalf (2010),

[26] EV prices in 2020 were approximately $54,000. The product of the price elasticity and pass-through rate from Muehlegger & Rapson (2022) is -1.78, implying a payment of approximately $30,000 per induced purchase

[27] Appendix Table 3 reports all MVPFs for our in-context specifications that use externalities and prices from the time and place the policy is implemented as opposed to our national 2020 harmonized baseline.

[28] We do not include any estimates of the impact of charging station subsidies (there is no quasi-experimental evidence on the effects of such policies in the US). However, work by Cole et al. (2023) uses a calibrated structural model to compare subsidies for EVs with subsidies for charging infrastructure. They find that $10B in subsidies generate a reduction in 884 million tons of CO_2 in subsequent decades. Applying our baseline SCC model and ignoring non-CO2 benefits suggests an MVPF in excess of 10, exceeding most MVPFs in our sample.

and Shrimali et al. (2015). We also supplement these results with six elasticity estimates from papers studying the impact of variation in feed-in-tariff rates in Europe.[29]

Figure 2 Panel A presents the components of WTP and net government cost using the elasticity from Hitaj (2013). The environmental benefits from the subsidy are much larger than those for EVs and sum to \$3.45 per dollar of mechanical subsidy. This is not because the price elasticity of wind turbines is larger (the price elasticity is -1.13 as opposed to -2.1 for EFMP above). Rather, it is because \$1 of induced spending on a wind turbine delivers more than \$3 of global environmental benefits from displacing the dirty production of electricity, while \$1 of induced spending on an EV generates less than \$0.04 in global environmental benefits.

The next bars capture learning by doing benefits. Way et al. (2022) estimate that a 1% increase in cumulative production leads to a reduction in wind turbine costs of 0.19% ($\theta = -0.19$). This leads to \$1 in future environmental benefits and \$0.46 in benefits from lower future prices of wind turbines. Combining together all our willingness to pay components produces a net WTP of \$5.90 per dollar of mechanical wind PTC.

For the net government cost, we begin with the \$1 mechanical cost of the policy and add the fiscal externality associated with the baseline PTC subsidy. In 2020 there was a PTC subsidy equal to 1.5 cents per kWh, which leads to a fiscal externality of \$0.35 per dollar of mechanical subsidy. Long-run climate benefits also generate a negative fiscal externality of \$0.08. Taken together we estimate a net cost of \$1.28. Dividing the WTP of \$5.90 by this net cost yields an MVPF of 4.63.

Figure 2 Panel B plots the MVPF estimates for wind subsidies and shows how our baseline MVPFs vary with the magnitude of the price elasticity. The other two studies we consider have elasticities of -1.3 (Metcalf 2010) and -1.75 (Shrimali et al. 2015), yielding MVPFs of 5.30 and 7.55, respectively. This yields a category average MVPF of 5.87 for wind PTCs, with a confidence interval of $[2.73, \infty]$, and we reject the hypothesis that EVs have a higher category average MVPF than wind PTCs with a p-value less than 0.001. In order to ensure that this high MVPF is not driven by the small sample of available quasi-experimental estimates, we also compare our estimates to those derived using elasticity estimates from “feed in tariff” policies in Europe that guarantee producers elevated prices for their clean energy generation. Figure 2 Panel B shows that these elasticities have a wide but roughly similar range to those in the US. Including European elasticities in our category average calculation would generate an MVPF of 5.93 as opposed to 5.85 using only US-based studies.

Residential Solar Subsidies The US federal government and many US states have enacted large subsidies to encourage residential solar installation. Figure 3 Panel A presents the components of the WTP and net cost of the MVPF using estimates from Pless & van Benthem (2019). We find an MVPF of 2.71, in between the MVPFs for EV subsidies and Wind PTCs.

[29] We do not provide in-context estimates for non-US studies, but instead focus on the implications of their price elasticity estimates for the US 2020 MVPF of wind subsidies.

As compared to wind PTCs, residential solar has a similar price elasticity but a smaller environmental benefit per dollar of induced spending. Residential solar also has a high learning rate from Way et al. (2022) ($\theta = -0.32$) leading to large learning-by-doing effects. Figure 3 Panel B presents broadly similar patterns using four causal effect estimates of solar subsidies from the literature. We find a category average MVPF of 3.86 (CI of [1.97,33.89]) that falls to 1.45 in the absence of learning by doing effects.

Even with learning-by-doing effects, residential solar subsidy MVPF estimates are substantially lower than our estimates for wind PTCs (3.86 versus 5.87). This difference may be driven by the distinction between utility-scale and residential energy production, rather than the distinction between wind and solar. With falling solar prices, the 2020 (levelized) cost of energy via utility-scale solar was roughly on par with the costs of utility-scale onshore wind. By contrast, the costs of residential solar remained more than two times higher than utility scale solar. While there are no quasi-experimental estimates of the impact of utility-scale solar, we can return to our wind PTC setting and imagine a similar subsidy for solar installations. Assuming the elasticity of solar installations is similar to historical wind PTC elasticities (-1.3), we can use the utility-scale solar costs per kWh to estimate an MVPF. Here, one motivation for assuming the -1.3 elasticity is similar for utility-scale wind and solar is that it captures a structural user cost elasticity that is plausibly constant across investment types. Under that assumption, we find the MVPF of utility-scale solar subsidies would be 10.97, well above our estimates for the wind PTC. Given this, a natural conclusion from our analysis is that subsidies to utilities for either wind or solar have higher MVPFs than residential solar subsidies.

Other subsidies Figure 4 presents the MVPFs for each subsidy in our sample along with the category average MVPF and its 95% confidence interval. In contrast to the higher MVPFs discussed above, we find smaller MVPFs for the majority of the remaining subsidies in our sample, including appliance rebates, hybrid vehicle subsidies, vehicle retirement programs, and weatherization subsidies. These policies have MVPFs near 1, with category averages whose confidence intervals fall below the confidence intervals for EVs, residential solar, and wind PTCs. While these other consumer subsidies can induce some changes in purchases, Table 1 shows that the magnitudes of the resulting environmental benefits are small relative to the size of the mechanical transfer offered by these policies. For every dollar of government spending, the average appliance rebate in our sample delivers twice as many inframarginal benefits (\$0.87) as environmental benefits (\$0.45). In other words, these policies are mostly transfers to those who would have purchased energy-efficient appliances anyway.

Summary The main lesson from this analysis is that subsidies for investments that directly displace the dirty production of electricity—namely, wind PTCs and residential solar subsidies—have the highest MVPFs. In particular, production tax credits for firms that produce wind energy have the highest MVPFs, generally exceeding 5. Subsidies to individuals who install residential solar panels also have high MVPFs exceeding 3.[30] By contrast, EV subsidies have MVPFs around 1.45. All other subsidies tend to have smaller MVPFs, with values around 1 ± 0.2. These results suggests the potential for meaningful welfare gains if climate spending is focused on policies that displace the production of dirty electricity. For example, every dollar of expanded spending on wind PTCs (with MVPFs above 5) financed by less spending on EV subsidies (with MVPFs around 1.5) would deliver \$3.50 in net benefits to society.

Robustness Appendix Tables 5–11 present the WTP and cost details for several alternative specifications. In particular, Appendix Tables 5 and 6 consider alternative models of the social cost of carbon corresponding to values of \$76 (with a 2.5% discount rate) and \$337 (with a 1.5% discount rate).[31] Higher values of the SCC accentuate the patterns of MVPFs we observe, but our core conclusions remain unchanged: wind PTCs have the highest MVPFs followed by solar, EVs, and then other subsidies.[32]

Panel A of Appendix Figure 3 visualizes seven alternative specifications for the construction of the MVPF, each reported separately for these three models of the SCC. First, the ex icons ($\times$) omit learning by doing benefits from the MVPF (see also Appendix Table 8). Without learning by doing, the values for EVs fall from 1.45 to 0.96, and the values for residential solar fall from 3.86 to 1.45. By contrast, even without learning by doing, subsidies for utility-scale wind produce relatively high MVPFs, with a category average of 3.85.

Next, we consider how our MVPFs change under alternative electricity grids. The diamond icons (♦) assume the US has the CA grid, which has the lowest CO_2 emissions. The plus icons (+) assume a grid corresponding to the Midwestern region of the US, which has the highest CO_2 emissions. Cleaner grids depress the MVPFs for wind and solar and slightly raise the MVPFs for EVs. Despite those changes, we continue to find an MVPF of 3.2 for wind PTCs under the CA grid. The MVPF for EVs increases to 1.53 in CA, but remains well below the MVPFs for wind PTCs and residential solar. In addition to US grid specifications, we have also considered robustness to the average European Union electricity grid (see Appendix Table 9). This delivers MVPFs similar to the CA grid, with category average MVPFs for wind, residential solar, and EVs of 3.84, 2.73, and 1.51, respectively.

[30] As a point of comparison, such estimates exceed estimates of the welfare cost of raising revenue through the income tax, which tend to range from 1 to 2 (Hendren 2020, Kleven & Kreiner 2006) and expansion of tax enforcement that has been estimated to be around 1.1-1.3 (Boning et al. 2023).

[31] Appendix Table 7 considers a much higher SCC of \$1367 from Bilal & Känzig (2024), which delivers similar orderings of policies but some policies with MVPFs near 1 now start to have MVPFs that exceed 2 (e.g. weatherization, vehicle retirements).

[32] Our baseline specification uses an SCC that captures benefits accruing to individuals around the world. Appendix Figure 4 shows, in blue bars, how the MVPF changes when only considering benefits to US residents and ignoring the benefits to the rest of the world. While the relative ordering again remains unchanged, the MVPF values decrease substantially. The wind and solar categories have MVPFs of 1.89 and 1.18 while other categories are often below 1. This is because only 13.1% of the global externality benefits are estimated to flow to US citizens and so the numerator of the MVPF falls in cases where the are meaningful global environmental benefits.

In Panel B of Appendix Figure 3, we consider how the MVPFs change if the electric grid sourced more energy from non-emissive sources. Eventually, the benefits of Wind PTCs and residential solar fall below EVs, but not until over 90% of the grid is sourced from renewable energy. Under a perfectly clean grid with no local or global pollution, the MVPF for wind is 1.03 and the MVPF for residential solar is 1.37. This illustrates that if one were to achieve a clean grid in the future, the value of subsidies to clean electricity production would diminish. In contrast, the MVPF for EVs rises to 1.61. This is because, as noted in (Gillingham et al. 2025), the environmental harms from electricity diminish while the environmental damage from gasoline, combined with learning by doing benefits of EV batteries, remains. The MVPF approach could be used to track the decreasing returns to such subsidies as the grids follow the transition to clean energy.

In addition to examining the effects of varying grid emissions, we also explore assumptions about the structure of the electricity market. Our baseline specification uses estimates of supply and demand curves for electricity consumption to estimate a rebound effect of 19.6%. The vee icons ($\vee$) double the magnitude of the electricity and natural gas rebound, while the hats ($\wedge$) assume no rebound in these markets. Larger rebound effects lower the value of wind PTCs slightly, as the MVPF falls from 5.87 to 4.99 when the rebound effect is doubled. But the broad patterns are similar. Our baseline specification also calculates markups in the electricity market and includes changes in profits as an additional WTP. The squares (■) show how assuming competition and omitting firm profits affects the MVPF (see also Appendix Table 10). These changes have minimal effects on the results.

Our baseline specification assumes that people are rational when purchasing energy-efficient goods so that the energy savings they provide are not included as an additional benefit beyond what is expressed via revealed preference. The triangles (▲) consider a specification that assumes that people purchasing energy-efficient goods were not aware of (or not internalizing) the energy savings they provide (see also Appendix Table 11). This generally increases the benefits of energy-efficient subsidies, but the MVPFs continue to fall below those of policies that directly displace dirty electricity production.

Lastly, our MVPF calculations for most goods envision a small (marginal) change in the policy relative to 2020 subsidy levels. We can also use the framework to explore non-marginal changes in subsidy levels. Appendix Figure 5 illustrates the MVPF of a non-marginal change in EV subsidies. It examines the effect of increasing the federal subsidy to $7500, the amount provided under the Inflation Reduction Act (IRA). The first dollar of the subsidy has an MVPF of 1.30. As the subsidy increases, the MVPF falls slightly. This is because the fiscal externalities are increasing in the size of the pre-existing subsidy. The MVPF on the 7500th dollar is 1.02. Integrating over all the marginal policy changes yields an average (non-marginal) MVPF of 1.15. Similarly, we can consider the case of non-marginal residential solar subsidies, where the IRA prevented the expiration of the subsidy and set the rate at 30%. Performing the same non-marginal analysis, we obtain an MVPF of 4.43, which is relatively close to but slightly above our marginal category average of 3.86. We can repeat the same exercise for the wind PTC, examining the effect of increasing the PTC from 0 to 2.6 cents per kWh, which is the change implemented under the IRA. That policy change results in an MVPF of 5.80 as compared to our baseline marginal MVPF estimate of 5.87. In sum, the MVPFs we observe for our marginal policy experiments yield similar conclusions to the non-marginal policy changes and broadly reinforce our primary conclusion about climate subsidies: those that directly displace the dirty production of electricity have the highest MVPFs.

5 Nudges and Marketing

Next, we consider policies that employ nudges or marketing strategies to reduce residential energy consumption. The Home Energy Report (HER) designed by Opower (now Oracle) is perhaps the most well-studied environmental nudge. The HER provides information on how to be more energy efficient in the home and often includes an element of social pressure (e.g., comparisons of a household's energy use with 100 similar neighbors). There have been over 200 rigorous RCTs showing the causal impact of such nudges on energy demand in the United States and around the world (Allcott 2011). Here, we show how to translate these estimates of causal impacts into MVPFs of nudges. First, we use estimates from Allcott (2011) of the national average treatment effect of HERs aimed at reducing electricity use. We then consider the effects of nudges to reduce electricity use in different regions using 166 treatment effect estimates obtained from Opower. Lastly, we discuss additional nudges designed to reduce natural gas usage and other marketing policies designed to increase uptake of clean technologies.

We begin with the WTP for the Opower nudge. HERs targeting electricity usage cause a reduction in consumption, which reduces both emissions and the profits of utility companies. Combining these treatment effects with the externality from electricity production in the US, we estimate that every \$1 invested in these nudges leads to \$3.87 in global environmental benefits and \$0.44 in local environmental benefits. These benefits are partially offset by rebound effects of \$0.76 and \$0.09 due to the increased demand that results from reduced energy prices. We also estimate that utility companies experience a decrease in profits of \$0.24 for each \$1 spent on the Home Energy Report (HER) nudge.

In our baseline specification, we assume that people were optimizing their energy use so that they do not place any additional value on private energy savings. We also assume that the nudge does not generate costs or benefits from any value of shame or pride or value of information from the nudges. We acknowledge, however, that these sources of WTP may be important and assess the robustness to including such estimates below (Allcott & Kessler 2019, Butera et al. 2022, List et al. 2023).[33]

[33] For example, Allcott & Kessler (2019) suggest that individuals would be willing to pay on average about half (49%) of the energy savings that they experience from the nudge. As a conservative approach, Appendix Table 11 presents the results when we add in 100% of the energy savings, and shows that our conclusions remain

On the government cost side, we assume that the government pays for the electricity HER and thus include administrative and logistical costs as a government cost.[34] Government revenue collected from utilities decreases by \$0.13, but the long-run climate fiscal externality saves the government \$0.06. Combining the willingness to pay and government costs, we obtain an MVPF of 3.01.

While this 3.01 estimate corresponds to an average electricity HER, the MVPF varies sharply across regions of the US. Figure 5 illustrates the MVPF for HER nudges across five US regions where field experiments have been conducted and evaluated. The Mid-Atlantic, Northwest, and Midwest have high MVPFs with average values of 5.68, 5.50, and 3.76, respectively. By contrast, in California and New England, the MVPFs are 0.52 and 0.24, respectively.[35] The key driver of these patterns is the relatively cleaner grid in New England and California, which decreases the environmental benefits. In these areas, any environmental benefits are roughly offset by the loss of profits to the utility companies.[36,37] We also note the value of nudges depends heavily on the global externalities from the grid, but the regional patterns we observe are robust to those SCC variables. At an SCC of \$76 rather than \$193, the category average MVPF falls from 3.01 to 1.34. In that case, regions with dirty grids have MVPFs in the 1.92 to 2.76 range while regions with cleaner grids have MVPFs near 0.

While we find large MVPFs for nudges to reduce electricity consumption, we find much smaller MVPFs for nudges to reduce natural gas consumption with an average value of 0.45. These lower MVPFs are driven by a combination of smaller treatment effects relative to electricity nudges (the average natural gas nudge reduces consumption by 0.14% while the average electricity nudge reduces consumption by 0.26%) and the fact that the environmental benefits of reduced natural gas consumption are smaller than the benefits of reducing electricity consumption in areas with dirty grids.

We also analyze nudges specifically targeted to manage demand during peak periods. These nudges can help avoid costly blackouts or expensive marginal generation caused by the intermit-

broadly similar.

[34]This appears to be a reasonable approximation of what happens in practice, but it is also true that energy companies may pay for nudges. This means that we measure the MVPF of the nudge *as if* the government were to enact the policy or pay utilities to enact the policy.

[35]It is possible that the effects of the nudge persist beyond the measured time periods in these studies. However, the MVPFs for CA and New England remain at 0.72 and 0.36 even if we assume that half of the treatment effects persist for two years after the nudge (Brandon et al. 2017, Allcott & Rogers 2014).

[36]Excluding the loss in firm profits, the MVPFs for CA and New England increase to 2.02 and 0.96, respectively. They continue, however, to be substantially smaller than the MVPFs in the three regions with dirtier grids: 5.81 (Mid-Atlantic), 5.50 (Northwest), 3.86 (Midwest). We note that this dependence of the welfare effects on firm profits is similar to the argument in Buchanan (1969), who considers welfare with corrective taxes under competition and monopoly.

[37]The Northwest is categorized as a dirty electric grid despite the substantial levels of hydroelectric power in the region. This is due to both (i) the high level of marginal emissions estimated in the AVERT model (as distinct from average emissions) and (ii) the nature of the regional aggregation used in the AVERT model of marginal emissions. The northwest region includes states with very high levels of grid emissions, such as Utah. Omitting the Northwest from our analysis does not change the broad trajectory of our findings regarding regional variation in nudge MVPFs.

tency of renewable energy sources. An example of such nudges is the peak energy report, which informs consumers of their energy consumption during peak periods. Brandon et al. (2019) find that these led to a 4% reduction in energy use during peak hours. Constructing the MVPF requires placing a social value on this reduction in peak energy use. Here, we focus on the extent to which the marginal cost of peak production exceeds the price. We consider marginal costs ranging from ranging from $500 per MWh to $1000 per MWh (CAISO 2021) and find associated MVPFs from 0.70 to 1.60. If the demand reduction succeeds in avoiding a blackout, rather than simply avoiding the costly generation of electricity, these MVPF estimates could rise as high as 5.30.[38]

In addition to energy reports, we study marketing strategies and information treatments designed to encourage the adoption of clean technologies and reduce electricity usage. For example, we examine the Solarize program that sought to increase residential solar installations by providing municipalities with a designated solar installer, group pricing, and an information campaign led by volunteer ambassadors. Translating estimates of the impact of this program from Gillingham & Bollinger (2021), we estimate an MVPF of 1.81. These exceeds the MVPFs of other marketing policies designed to encourage the adoption of weatherization technologies, which have MVPFs near 1.

Summary of MVPFs for Nudges and Marketing Nudges to reduce electricity consumption can yield high MVPFs — on average exceeding 1.5 in our 2020 baseline specification. However, the effectiveness of these policies depends heavily on the cleanliness of the electric grid. The MVPFs exceed 3 in areas with dirty grids and fall below 1 in areas with cleaner grids such as California and New England. Nudges aimed at reducing natural gas consumption have lower MVPFs than those targeting electricity consumption due to the smaller treatment effects and lower environmental damages relative to electricity production. Finally, marketing strategies can also yield modest MVPFs when targeting adoption of goods with particularly high environmental benefits like residential solar. But these marketing policies tend to have lower MVPFs than direct subsidies for the same goods.

6 Revenue Raisers

The classic solution to an environmental externality is to tax emissions or sources of emissions. Such policies can reduce emissions while also raising government revenue. For revenue-raising policies, the MVPF measures the welfare burden imposed on individuals per dollar of government revenue raised. This means that, all else equal, lower MVPFs correspond to better methods of raising revenue. We focus here on the MVPF for two types of policies: taxes and cap-and-trade.

[38]This calculation assumes a value of lost load (VOLL) of $4,300 per MWh (Brown & Muehlenbachs 2024). We recognize the VOLL may vary across settings (Borenstein et al. 2023).

6.1 Taxes

As a benchmark, lump-sum taxes that do not affect behavior have an MVPF of 1 because they impose \$1 in welfare cost per each dollar of revenue raised. More generally, the same equation (5) that we used for subsidies characterizes the MVPF of a tax. Instead of thinking of τ as a subsidy, suppose now that $\tau < 0$ is a tax, and imagine that the environmental externality is a damage from consumption of x, $V < 0$. From the perspective of a tax change, the numerator of the MVPF reflects two countervailing forces. On the one hand, each dollar of tax imposes a \$1 of burden on the taxed individuals. On the other hand, the behavioral response to the tax changes consumption of the taxed good, x, generating environmental gains that partially offset the burden of the tax, $(-\epsilon)\frac{V}{p}$. That change in consumption is also reflected in the denominator because changes in consumption affect tax revenue and diminish the net revenue raised from the tax, $-\epsilon\frac{\tau}{p}$. The MVPF is 1 if the tax is set at it's Pigouvian level.[39] If the tax is below (above) the Pigouvian level, the MVPF of the tax will fall below (above) 1. In our empirical implementation, we use an extended version of this formula that includes externalities from imperfect competition and learning-by-doing effects (e.g., gas taxes induce the adoption of EVs, generating learning by doing externalities).

We construct 12 MVPFs for gasoline taxes using estimates of the response of gasoline consumption to price and tax changes. These estimates imply price elasticities that range from -0.04 (Hughes et al. 2008) to -0.46 (Davis & Kilian 2011). Figure 6 illustrates the construction of the baseline MVPF using the elasticity estimate of -0.33 from Small & Van Dender (2007). A \$1 increase in the gas tax leads to a WTP of consumers of \$1 to avoid the tax increase, as the literature tends to suggest full pass through of taxes to consumers (Marion & Muehlegger 2011). The reduced driving due to the tax leads to global environmental benefits of \$0.27, local pollution benefits of \$0.03, and local benefits from reduced accidents and congestion of \$0.21.

Higher consumer prices for gasoline can cause an increase in EV adoption (Bushnell et al. 2022). Motivated by this, we use Slutsky symmetry to quantify the potential impact of this substitution on our MVPF estimates. We translate the own-price elasticity of EV purchases of -2.1 (Muehlegger & Rapson 2022) into a cross-price elasticity between the price of gasoline and EV demand of 0.22.[40] The resulting benefits from EV purchases are, however, quite small. We find that induced EV purchases generate \$0.0008 in initial global and local damages from electricity generation. They also generate learning-by-doing benefits of \$0.002 from reduced future EV prices and \$0.0002 from future environmental benefits.[41] Lastly, we estimate that

[39]With taxes or subsidies on other goods, the Pigouvian tax internalizes both the environmental and other fiscal externalities.

[40]Under Slutksy symmetry, in combination with the assumption of no change in overall car demand (just shifting between EVs and ICE vehicles), the cross-price elasticity is given by the own-price elasticity multiplied by the ratio of the present discounted value of operating costs of a gasoline powered car relative to the price of an EV. See Appendix E.10 for our derivation.

[41]We also account for utilities' WTP for increased electricity usage by EVs as well as accompanying fiscal externalities associated with EV adoption. These effects are negligible.

gasoline producers are willing to pay \$0.07 WTP to avoid the tax, due to reduction in gasoline use. Summing, this implies a total WTP to avoid the tax of \$0.56.

On the cost side, the reduction in demand also leads to lost corporate and gas tax revenue of \$0.09.[42] The US government also gains \$0.01 in future revenue via the climate externality. Combining our WTPs and cost implies an MVPF of 0.60. A dollar of government revenue raised leads to a welfare cost of \$0.60 on individuals.

Figure 7 shows MVPFs of 0.44 to 0.95 for the other gas tax policies in our sample, with a category average of 0.67. We find slightly higher MVPFs around 0.8 for taxes on diesel and jet fuel, with detailed calculations in Appendix E.11.[43] In each of these cases, the MVPF falls below 1 because the externalities avoided (environmental, congestion, or accidents) are larger than the fiscal externality induced by the policy.

On the whole, the results suggest that fuel taxes raise revenue at a relatively low welfare cost. The MVPFs of these revenue raisers are well below the MVPF of changes to the income tax, which range from 1 to 2 depending on the income level of the taxed individuals (Kleven & Kreiner 2006, Hendren 2020, Hendren & Sprung-Keyser 2020). The MVPFs of fuel taxes are even below 1, the MVPF of a non-distortionary lump sum tax. Returning to equation (2), we can use the MVPFs to make statements about the welfare effects of budget-neutral policy experiments. For example, we can directly compare an MVPF of .6 for gasoline taxes with an MVPF of 1.1 for income taxes on low-income earners. If society places equal weight on the individuals impacted by each policy, then every dollar of revenue shifted from income taxes to gasoline taxes generates 50 cents in additional welfare (with 10 cents of this corresponding to the "double-dividend" obtained by reducing distortions in the income tax schedule).[44] If, by contrast, a decision-maker would prefer the status quo, it implies they must place a higher welfare weight on drivers relative to an average low-income individual.

6.2 Cap and Trade

Another common mechanism used to reduce carbon emissions is through a cap and trade scheme. These systems impose quantity limits on emissions and let firms trade the rights to such emissions. Our model setup can be used to evaluate cap-and-trade policy in much the same way that it can evaluate a tax change. We let x denote the number of permits issued and make a modification to introduce the notion of "leakage." We assume that one fewer permit leads to $(1 - L)$ reductions in emissions, where L is the leakage of emissions into areas not

[42] Consistent with the findings in West & Williams (2007) that gasoline is a relative complement to leisure rather than labor, we exclude any fiscal externalities from changes in labor income.

[43] Diesel taxes have a higher MVPF than gas taxes because diesel demand is less elastic than gasoline demand. This increases the MVPF, despite the fact that diesel vehicles impose a larger per-gallon externality than gas-powered vehicles. The jet fuel tax has a higher MVPF than gas taxes due to fewer local externalities.

[44] Even ignoring environmental benefits and focusing solely on accidents and congestion, gas taxes have an MVPF of 0.95, which continues to be lower than the MVPFs identified for tax changes at any point across the income distribution (Hendren 2020).

captured by the cap-and-trade program. In the MVPF derivation, this changes the externality per unit of x from $-V$ to $-(1-L)V$. Following equation (5), and multiplying by $x\frac{dp}{dx}$, the MVPF of changing the number of auctioned permits is

$$MVPF = \frac{-x\frac{dp}{dx} + V(1-L)}{-x\frac{dp}{dx} - p}. \tag{27}$$

The first term is the firms' willingness to pay to avoid the increase in permit prices induced by a reduction in permit supply. This is offset by the environmental damages avoided, $V(1-L)$, due to a one-unit change in the number of permits auctioned. On the cost side, the government receives the mechanical revenue from the higher prices, $-x\frac{dp}{dx} > 0$, but also loses p in revenue from the forgone permit no longer auctioned. (p does not appear in the numerator due to the envelope theorem: the marginal firm holding a permit has a marginal abatement cost equal to the permit price.)

These derivations also illustrate that the MVPF in equation (27) is isomorphic to the MVPF of a carbon tax. This interpretation requires: (i) the change in permit price that results from a change in quantity of permits, dp/dx, must also reveal the (inverse) response that the economy would experience from a change in price due to a carbon tax, dx/dp, and (ii) both instruments generate similar spillovers across markets (i.e. similar leakage). Under these conditions, the cap-and-trade MVPFs we estimate below also reveal the MVPF of a hypothetical carbon tax in these settings.

There are two cap-and-trade policies that have been evaluated in the US: the Regional Greenhouse Gas Initiative (RGGI) in the Northeast and mid-Atlantic, and the California Cap-and-Trade Program. We also briefly discuss the European Emissions Trading System (ETS).

We begin with the in-context estimates of the effect of RGGI on greenhouse gas emissions using results from Chan & Morrow (2019). Between 2009 and 2016, there were 816.2 million permits auctioned (per short ton of CO_2) at an average clearing price of \$3.19 (in 2016 dollars). The authors estimate that RGGI reduced 22 million short tons of CO_2 during this period. This implies that a one unit reduction in the quantity of permits sold led to a $\$1.45 \times 10^{-7}$ dollar increase in the permit price, or $dp/dx = -1.45 \times 10^{-7}$. To incorporate this into the MVPF formula above, we assume that this abatement curve has a constant slope, so that this estimate of dp/dx applies to the marginal permit issued.[45] This suggests that if RGGI had auctioned one fewer permit between 2009 and 2016, it would have lost \$3.19 from the price of the permit but

[45] An alternative approach that does not necessarily require a linear abatement curve is to construct a non-marginal valuation of the introduction of cap and trade. The 816.2 million permits auctioned at \$3.19 generate revenue of \$2.6B. That is also a welfare cost to the inframarginal firms who buy permits. In addition, firms cut 22M short tons of CO_2 emissions, which impose a welfare cost on them somewhere between 0 and \$3.19*22M=\$70.18M, with the exact value depending on the shape of firm demand for emissions (half of this would be a welfare loss if the abatement curve were linear). The reduction in emissions generates CO_2 and non-CO_2 benefits of \$1462M and \$2578M, respectively. Regardless of the shape of the abatement curve, these gains exceed the cost to firms, suggesting that the introduction of cap and trade delivered net gains to society and revenue to taxpayers (i.e. $MVPF < 0$).

gained approximately $-dp/dx * x = 1.45 * 81.62 = \$118.48$ in additional revenue from higher permit prices.[46]

Higher prices impose a cost of \$118.48 on firms purchasing permits. Some firms choose to reduce their emissions instead of purchasing a permit, but such firms are indifferent on the margin to that change. Chan & Morrow (2019) also find evidence of significant leakage: permitting one less short ton of CO_2 emissions leads to $1 - L = 0.49$ fewer short tons of CO_2 actually emitted. Valuing this using the 2016 SCC implies that this reduction in CO_2 emissions has an environmental benefit of \$65.20. Adding the reduction in local pollutants SO_2 and NO_X yields an additional gain of \$117.21. On net, these environmental benefits outweigh the cost to firms, generating a net benefit of \$63.93. Raising revenue via a reduction in auctioned permits as part of RGGI led to a net win for individuals and taxpayers. The policy raises revenue without imposing a welfare cost, so the $MVPF < 0$.

While our in-context estimates suggest RGGI led to significant benefits to taxpayers and individuals in society, we caution that it is potentially difficult to extrapolate our in-context estimates to a hypothetical 2020 policy reform. This is because one needs to know the marginal abatement cost curve in 2020 to understand how the number of permits would affect its price. If one assumes that the abatement curve is linear and stable over time, we find that greater restrictions in auctioned permits would continue to increase government revenue (\$123.01) while also delivering a net gain to individuals in society. The WTP for environmental damages (\$210.33) outweighs each dollar firms pay in permits (\$127.78). However, the primary channel through which RGGI affected emissions was by inducing a switch from coal to natural gas. It is less clear whether the same set of low-cost substitution options continue to exist today, as many coal plants have been retired. Consequently, it may be that dp/dx is larger in 2020 than in the early 2010s, which leads to fewer environmental benefits per dollar of cost imposed on those buying permits.

In addition to our analysis of RGGI, we also consider the MVPF of the introduction of the California Cap-and-Trade Program using estimates from Hernandez-Cortes & Meng (2023). They track outcomes for a 5% sub-sample of firms that are subject to the cap-and-trade system. If we assume no response from the other 95% of firms, we find an MVPF of 0.941. However, if the other 95% of firms have a similar response to the 5% of firms included in their sample, the environmental gains outweigh the cost to firms of higher permit prices. This would suggest that, like RGGI, the California Cap-and-Trade auctions raise revenue while also generating net welfare gains to society, leading to a negative MVPF. We also find a similar result when using estimates from Colmer et al. (2024) and Bayer & Aklin (2020) to study the introduction of the European Union's Emissions Trading System (ETS). Despite evidence of significant leakage, estimates in both papers suggest that reductions in permits raise revenue while also providing

[46]We estimate a fiscal externality on the government budget from the impact of changes in CO_2 emissions to be \$1.27, which suggests a net government revenue of \$116.56 from issuing one fewer permit. Motivated by the evidence in Colmer et al. (2024) and Metcalf & Stock (2023), we assume that cap and trade induces no reduction in the productive capacity of firms, and so there is no additional corporate tax fiscal externality.

positive benefits to society. More speculatively, exploiting the isomorphism between carbon taxes and reduced permits in auction, this also suggests that carbon taxes can potentially deliver government revenue at low or even negative net costs to society.

Summary of Revenue-Raiser MVPFs The key lesson of this section is that taxes and other restrictions on pollution-emitting activities offer paths to raising revenues at low welfare costs. The MVPFs of these policies fall consistently below 1, suggesting that they impose less than $1 in burden for each dollar of revenue raised. This lies in contrast with other traditional revenue raisers, such as increases in income tax rates, which consistently have MVPFs above 1. Returning to equation (2), the results suggest that a decision-maker setting tax policy would need to have high implicit social welfare weights on individuals engaged in pollution-emitting activities in order to justify status quo policies as optimal. For cap and trade, the results show that at the time these policies were implemented, there were large quantities of emissions that could be reduced at relatively low cost. The presence of this low hanging fruit meant that small prices on carbon generated a win for taxpayers and a net win for individuals affected by the policy. More broadly, our results suggest that the presence of these large environmental externalities creates opportunities for raising revenue at a low welfare cost relative to typical methods of raising revenue.

7 International Policies

The impacts of greenhouse gas emissions are felt worldwide, regardless of the source of the emissions. This means that many of the beneficiaries of US policies addressing climate change reside outside of the US, and that US residents are the beneficiaries of climate policies enacted in other countries.

In this section, we draw upon an illustrative set of climate-focused policies implemented in developing countries, largely by NGOs. We consider: to what extent is it beneficial to US residents to pay for policies implemented in other countries as a form of international aid. We consider 14 policies in five categories: cookstoves, deforestation payments for ecosystem services, payments to prevent rice field burning, wind subsidy offsets, and appliance and weatherization rebates.

We begin with subsidies for improved cookstoves in Kenya. Berkouwer & Dean (2022) find that small subsidies for these cookstoves help overcome credit constraints and encourage the purchase of these appliances. When offered a $30.37 subsidy (in 2020 dollars), 54.5% of individuals take up the cookstove. Almost all of those beneficiaries are marginal, as only 0.6% would have taken up the cookstove in the absence of the policy. The paper also finds that each new cookstove reduces CO_2 by about 7 tons.[47] This translates into $43.16 in global

[47]These calculations assume that charcoal is derived entirely from non-renewable biomass. If we were to use

environmental benefits for each mechanical dollar of the subsidy. We combine those global externality benefits with the transfer benefits of the subsidy and the value of private energy savings. This yields a total willingness to pay of \$50.82 for each mechanical dollar of the subsidy.

Next, we consider the net cost of the policy. In contrast to the policies considered up to this point, the impact of this policy on carbon emissions is sufficiently large such that the climate benefits can meaningfully affect future US tax revenue. In our baseline specification, \$3.70 of the \$193 SCC falls onto the US government in the form of increased tax revenue from increased future GDP.[48] This means that the \$30 upfront subsidy only costs the US government $\approx$ \$4. In other words, a mechanical dollar of subsidy ultimately costs the US government just \$0.157. When combined with the WTP for the policy, this yields an MVPF of 37 when only considering benefits to US residents and an MVPF of 323 when considering benefits to individuals globally.

A key factor in this calculation is the extent to which reductions in global warming have a positive impact on future US tax revenue (due to higher future productivity). Models that report the same social cost of carbon can generate different MVPFs because they differ in their incidence on the US federal budget. For example, we could have assumed that the entirety of the SCC was driven by changes in market productivity. This approach is motivated by a literature estimating damages functions that relate carbon to GDP (Nath et al. 2024). In this case, we find that the subsidy pays for itself. The net cost of the policy is -\$11.31 for each dollar of mechanical subsidy (and the US-only MVPF is infinite). By contrast, other models suggest that the incidence of emissions damages on the US taxpayer could be quite small. For example, estimates from PAGE (Nordhaus 2017) suggest the US-incidence of carbon damages is just 7%. Similarly, estimates from the GIVE model (Rennert et al. 2022) suggest that changes in productivity are concentrated outside the US. If we drop the US-specific fiscal externality to zero, the US-only MVPF falls to 4.91 and the MVPF including global benefits falls to 49.97. This highlights the importance of articulating incidence when constructing measures of the social cost of carbon. Total damages estimates can be reported in GDP-equivalent terms, but the distinction between the sources of damages can meaningfully impact the welfare consequences of a policy.

Figure 8 presents the MVPFs for the other international policies in our baseline sample.[49] MVPFs using only US benefits are shown in blue, and those including global benefits are shown in orange. There is substantial variation in MVPFs both within and across program categories. For example, the evidence from Berkouwer & Dean (2022) differs from the findings of previous

a fraction non-rewewable biomass of 45% estimated by the United Nations (2023), the carbon reduction would be 1.67 tons.

[48]The precise value of that fiscal externality depends on the model underlying the social cost of carbon. In our baseline specification, we assume the US experiences 15% of the benefits of carbon abatement in proportion to its share of global GDP. Across SCC models these benefits are typically a mix of mortality reductions and productivity increases. We therefore assume 50% of the benefits are changes in productivity and therefore taxed by the US government at a rate of 25.5% (the US tax to GDP ratio in 2020).

[49]Table 1 discusses results for additional policies in our extended sample, which includes some policies that are not a natural fit when considering hypothetical US-based funding. This includes, for example, nudges for energy reduction in foreign countries.

work on cookstove subsidies. Hanna et al. (2016) found that recipients simply did not use the cookstoves, which translates to an MVPF near zero. Similarly, we find large variation in returns to policies designed to prevent deforestation. We find that payments to farmers in Sierra Leone to prevent deforestation yield an MVPF of 15.9 even when only considering the benefits to US residents. However, not all deforestation programs are as effective. We find a smaller MVPF for a program in Mexico evaluated in Izquierdo-Tort et al. (2024), with a global MVPF of 1.71 and a US-only MVPF of 0.1.

We also find large MVPFs for policies that use unique incentive contracts to discourage rice field burning. We find MVPFs between 10-15 when including global benefits and in the 1.3-1.8 range when only including US benefits. Additionally, we find potentially high returns to policies encouraging the adoption of wind turbines in India, with a global MVPF of 7.64 and a US-only MVPF of 0.9.[50] As is the case with our primary estimates, we find the lowest MVPFs for policies that use rebates to encourage the purchase of efficient appliances.

In sum, we find potentially high returns - even from a US-only perspective - from policies that invest in reducing greenhouse gas emissions in developing countries. In fact, subsidies for cookstoves and deforestation subsidies in Sierra Leone have higher MVPFs than any domestic subsidy in our sample, even when only considering the benefits accruing to US residents. That being said, we reiterate three notes of caution. First, our exact MVPF estimates depend on the incidence of the social costs of carbon and, in particular, whether the benefits accrue in the form of increased US productivity. Such productivity benefits have US tax revenue implications that meaningfully impact the net cost of the subsidies to the US government. Second, we find high variance in our international MVPFs estimates, even within policy categories. Even when spending within a promising category, high returns are certainly not guaranteed. Finally, our analysis assumes the US government could implement these policies with the same cost structure as the NGO conducting the evaluation. The US government may face different administrative costs when scaling these programs. Ultimately, the key lesson from our analysis mirrors the conclusions of Glennerster & Jayachandran (2023): International aid policies can be a valuable part of the toolkit for addressing climate change.

8 MVPF Versus Cost per Ton

The preceding analysis applies the MVPF framework to analyze the welfare consequences of US climate change policies. How do these lessons compare and contrast to what one would learn from measures of cost per ton?

[50] We draw upon estimates from Calel et al. (2025) examining the impact of a wind subsidy in India on greenhouse gas emissions. The authors argue that at least 52% of installations are inframarginal, suggesting that the carbon offsets are not fully offsetting carbon emissions. We take that implied inframarginal fraction as given, rather than a bound, and show that it results in an implied elasticity of -2.2 and an implied MVPF of 7.64. We note that the 52% inframarginal share is a lower bound so the ultimate MVPF could be lower if the leakage is higher.

In Section 2.3, we outlined three definitions of cost per ton, each of which conceptualizes "cost" in a different manner. We discussed the theoretical advantages and potential drawbacks of each approach. In this section, we examine each measure empirically. Table 2 reports all three measures of cost per ton discussed in Section 8 for each policy subcategory alongside the associated MVPF (see Appendix Table 12 for each individual policy in our sample).[51]

The results in Table 2 show that differences in the definition of "cost" matter in practice. For example, appliance subsidies have cost-per-ton values that range from -$2 to $474. From a resource cost perspective, energy-efficient appliances save enough energy to overcome the difference in upfront price as compared to counterfactual appliances. This leads to a net resource cost per ton of -$2. The government cost per ton, however, is $474, because many subsidies go to people who would have purchased those appliances even in the absence of the subsidy.[52]

The wide variation in cost per ton across definitions within a policy category highlights the need to be consistent when constructing a measure of cost per ton. For example, Gillingham & Stock (2018) provides a ranking of policies according to their cost per ton of carbon abated. The lowest cost per ton policy in their list is the nudges studied in Mullainathan & Allcott (2010), who use a resource cost-per-ton measure—a measure that tends to be lower because it includes energy savings and omits inframarginal costs.[53] By contrast, solar subsidies are reported to have higher costs per ton, but some of these measure the government cost per ton (e.g., Gillingham & Tsvetanov (2019)).

Using the columns of Table 2, we can compare policy categories while maintaining a fixed definition of cost. We see how each of these cost-per-ton definitions yield conclusions that differ from the MVPF approach. We begin with the resource cost per ton.

Resource Cost per Ton (RCPT) There is large variation in the RCPT across categories, even in cases where the MVPF values are quite similar. For example, energy-efficient appliances have a negative resource cost per ton on average (-$2), which is well below the resource cost per ton of vehicle retirement ($1,007) and hybrid vehicles ($577). From a resource cost perspective, buying an energy-efficient appliance uses far fewer economic resources per ton of CO_2 abated relative to the replacement of an old car with a new car. From a social welfare perspective, however, the MVPF of subsidies for these goods is close to 1 (1.16, 1.05 and 1.01). This difference is due to two factors. First, the resource cost per ton does not include non-CO_2 benefits. For example, resource cost per ton omits the consumption value of the new car when

[51] The estimates in Table 2 include learning-by-doing benefits; Appendix Table 13 shows the equivalent table if we exclude these effects.

[52] The social cost per ton lies in between. It includes these inframarginal transfers as both a cost and benefit and adds non-CO_2 benefits like reductions in local pollutants. It excludes the value of energy savings due to the envelope theorem.

[53] The paper describes its measure of costs as capturing the "long-run marginal cost of electricity minus the program cost to the utility."

considering vehicle retirement.[54] Second, resource cost per ton analyzes products as opposed to policies. As shown in Section 2.3, the resource cost is independent of the price elasticity of demand. In the case of appliance rebates, these policies are primarily transfers to individuals who would have purchased the appliances anyway. That drives the MVPF toward 1 but is not reflected in the RCPT.

Turning from category averages to individual policies in our sample, we see a similar pattern. Appendix Table 12 shows that energy-efficient refrigerators have some of the lowest resource costs per ton in our sample, ranging from -$298 to -$512. These values fall below the resource costs per ton associated with wind turbines that range from -$96 to -$113. This particular ordering between energy-efficient appliances and wind turbines echoes a key finding in the influential resource cost-per-ton calculations constructed by McKinsey & Company (Enkvist et al. 2007). The MVPF approach tells a different story. It shows that $1 of government spending on subsidies for efficient fridges delivers between $1.01 and $1.04 in benefits to individuals per dollar of government cost. This is much less than the $4.63-$7.55 in benefits per dollar spent on subsidies for wind turbines.

Government Cost per Ton (GCPT) The government cost-per-ton approach considers the fiscal cost of each ton of CO_2 abated. Using this approach, we find that wind PTCs have the lowest GCPT of any subsidy in our sample. We estimate that it costs the government $46 to abate a ton of CO_2. Residential solar subsidies have the second lowest GCPT of any subsidy in our sample at $90 per ton. The relative superiority of wind and solar mirrors the findings of the MVPF approach. That said, the GCPT approach omits the benefits of inframarginal transfers and other non-CO_2 benefits. This generates some differences relative to the MVPF. For example, EVs have a GCPT of $1,356, substantially higher than the $474 cost for appliance rebates, but a higher MVPF (1.45 versus 1.16). As noted above, 95% of the benefits of EV subsidies are non-environmental, flowing to individuals who are buying or selling EVs. The omission of those benefits increases the GCPT of EVs.

The omission of non-CO_2 benefits also influences the ability to compare the GCPT to the SCC. At first glance, one might think that an EV subsidy with a government cost of $1,356 per ton is not a worthwhile expenditure if the SCC is $193 per ton (because 1356 exceeds 193). However, the existence of non-CO_2 benefits means that we cannot readily compare their GCPT to the SCC and draw conclusions about the social welfare impact of such a subsidy.

The GCPT is also not well suited to examine revenue-raising policies. For example, the GCPT of fuel taxes is the lowest of any policy in our sample, at -$995. A decision-maker focused on reducing CO_2 at the lowest cost to the government without any consideration for the welfare impact of a policy would generally find taxes appealing. When evaluating the welfare

[54] The omission of non-CO_2 benefits and costs is most clearly seen in the negative value of 104 for gas taxes. Gas taxes reduce gasoline consumption and thus reduce resource usage. But, they impose a welfare cost on the taxed individuals that is excluded from the resource cost.

impact of taxes, however, the key consideration is how the burden of the taxes on individuals compares to the revenue raised (and the externalities avoided). The GCPT omits the burden of those taxes.

Social Cost per Ton (SCPT) The final column of Table 2 reports the SCPT of each policy category. As described in Section 2.3, the social cost is equal to the government cost minus all (non-CO_2) benefits of the policy.

We find negative SCPTs for residential solar (-\$67) and wind PTCs (-\$32), indicating these policies can both abate carbon and provide net non-CO_2 benefits to society. The superiority of these two policies aligns with the conclusions of the MVPF approach. However, we also find low SCPTs for EVs (-\$415) and HEVs (-\$38). This contrasts with the MVPF results of 1.45 for EVs and 1.01 for HEVs, which fall well below the MVPFs for wind and residential solar.

The reason for this difference in orderings stems from the difference in objective functions associated with the SCPT versus MVPF metrics. Recall that the MVPF provides guidance to a decision-maker seeking to maximize social welfare subject to a government budget constraint, while the SCPT approach aims to achieve a particular CO_2 reduction at the lowest social cost. This means that the SCPT holds fixed the tons of carbon abated when comparing policies. As seen in the GCPT estimates, using EV subsidies to abate a ton of carbon costs the government \$1,356. In contrast, abating a ton of carbon costs the government just \$46 through a wind PTC. So, the SCPT of -\$415 for EVs tells us that if the government spent \$1,356 on EV subsidies it could reduce 1 ton of carbon and generate \$415 in additional social benefits. This \$415 exceeds the \$32 in additional net benefits created from abating a ton of carbon through the wind policy. But, if the government instead spent an equal amount (\$1,356) on Wind PTCs, it would abate 29.5 tons of carbon ($1356/46 = 29.5$) and generate additional net benefits of \$943 ($32*1356/46 = 943$). Those benefits far exceed those generated from EVs. The MVPF captures these higher returns per dollar of government expenditure. The SCPT implicitly compares a policy that spends \$1,356 on EV subsidies to a policy that spends \$46 on wind PTCs.

This example also highlights another potential consideration when measuring social cost: the government may face an additional welfare cost of raising revenue. In its canonical formulation, the SCPT statistic does not account for the welfare cost of raising revenue. As we noted in Section 2.3, a potential way to account for this within the SCPT approach is to include the welfare cost associated with raising government revenue, most commonly conceptualized as an increase in the linear income tax rate. Appendix Table 14 reports the SCPTs associated with assuming that the policy is financed by an income tax increase that has an MVPF of 1.1, 1.3, and 1.5, respectively.

The key takeaway from this analysis is that the SCPT can be quite sensitive to the welfare cost of raising revenue, especially for policies with large inframarginal transfers. The SCPT of EV policies ranges from -\$259 to \$260 as the cost of raising revenue changes. If raising revenue imposes a welfare burden of \$1.10 per dollar, then using EV subsidies to remove 1 ton of carbon would lead to a $259 benefit to society. By contrast, if raising revenue imposes a welfare burden of $1.50 per dollar, then abating a ton of carbon through an EV subsidy costs society $260. These adjustments also lead to similarly large changes in the SCPT for appliance rebates. Without any welfare cost of raising revenue, the SCPT is $111, but this rises to $349 when $\phi = 1.5$.

When calculating the welfare cost of raising revenue, previous literature has argued that there is no one-size-fits-all value of ϕ. Even focusing exclusively on changes in the income tax schedule does yield a single marginal cost of raising revenue, as noted in Section 8. Rather, the marginal cost depends on who you raise the revenue from: there are generally low welfare costs (e.g. $\phi < 1$) when raising money from low-income individuals and higher welfare costs (e.g. ϕ exceeding 2 or 3) when raising money from high-income individuals, even if the elasticity of taxable income is constant across the income distribution (Kleven & Kreiner 2006, Jacobs et al. 2017, Hendren 2020).

A key feature of the MVPF is that it does not embed an assumption about how the budget constraint is closed. The MVPF of a spending policy measures the welfare effects of that spending and the MVPF of a revenue raiser reflects the welfare cost of that policy. A decision-maker can therefore use the MVPF to choose from a menu of spending and revenue-raising policies to construct their preferred budget-neutral policy. For example, one can evaluate the welfare effects of a wind subsidy financed by an increase in the gas tax. If the beneficiaries of these policies have similar social welfare weights, the comparison of the 5.87 for wind PTCs to the 0.67 for gas taxes suggests that every $1 of government revenue raised from a gas tax and spent on wind PTCs generates $5.20 (=5.87-0.67) in benefits to individuals in society. Such a calculation requires an estimate of the SCC but avoids assumptions about the welfare cost of changes to the income tax schedule.

9 Conclusion

We conduct a comprehensive assessment of tax and expenditure policies that impact CO_2 emissions and have been rigorously evaluated using experimental and quasi-experimental methods. We draw three main lessons: First, subsidies for investments that directly displace the dirty production of electricity, such as production tax credits for wind power and subsidies for residential solar panels, have higher MVPFs (generally exceeding 3), than all other subsidies in our sample (with MVPFs generally around 1). Second, nudges to reduce energy consumption have large MVPFs, often exceeding 5, when targeted to regions of the US with a dirty electric grid. By contrast, nudges targeted toward areas with cleaner grids such as California and the Northeast have substantially smaller MVPFs (often below 1). Third, fuel taxes and cap-and-trade policies are highly efficient means of raising revenue (with MVPFs below 0.7) due to the presence of large environmental externalities. We also note that some of the highest MVPFs in our sample are international subsidies. These policies can produce high returns, even when only considering benefits to US residents and the incidence on US taxpayers.[55]

Methodologically, our approach integrates learning-by-doing externalities directly into our welfare analysis. Learning-by-doing effects have only modest impacts on the desirability of wind subsidies; however, the desirability of residential solar policies (and to some extent EV subsidies) depends heavily on the potential for learning-by-doing spillovers. We hope our framework is useful in future work analyzing subsidies for newer technologies such as carbon capture.

Alongside our MVPF calculations, we also compare and contrast these results with those obtained from more common cost per ton of CO_2 measures used in the literature. We first highlight the importance of being clear about the definition of 'cost' when measuring cost per ton, as the values differ depending on whether one talks about a resource cost, government cost, or social cost per ton. We also show that the key lessons from the MVPF analysis – especially that wind PTCs have higher welfare benefits per dollar of government spending compared to EV subsidies – would have been difficult to glean from cost-per-ton metrics alone.

The MVPF approach also facilitates comparisons across policy domains. The high MVPF values we find for spending on renewable energy generation exceeds the MVPFs found for many areas of spending on US adults documented in (Hendren & Sprung-Keyser 2020) and the Policy Impacts Library. The values rival, but are slightly less than, the MVPFs for spending on health and education for low-income children. By comparison, the MVPFs of climate-focused revenue raisers are far below the MVPFs of nearly all other policies studied in Hendren & Sprung-Keyser (2020) and common revenue raisers such as increasing tax rates or increasing tax enforcement (Boning et al. 2023). This suggests that climate policies present a unique opportunity to raise revenue comparatively a lower welfare cost.

We can also use the MVPF framework to examine whether historical environmental policy in the US has prioritized spending in areas with high returns. For example, we can compare the allocation of funds under the American Recovery and Reinvestment Act (ARRA) of 2009 with the allocation of funds under the Inflation Reduction Act (IRA) of 2022. The ARRA spent 3 times more on clean energy than on energy efficiency. By contrast, the IRA spent 9.4 times more on clean energy than energy efficiency. This represents a substantial relative reallocation, with far greater focus on spending in categories with higher MVPFs.[56] It is important to note, however, we also see a reallocation over time toward greater relative spending on EV subsidies, an area with comparatively lower returns. IRA funding on EVs exceeded its direct funding for clean energy while the ARRA spending on EVs was less than half its spending on clean energy.

In the end, we believe that the MVPF framework and the valuation methods used herein can serve as a useful tool for the analysis of climate policy. All of our code is available on GitHub. We hope this serves as an aid to researchers constructing their own future policy analyses.

[55] We note that such policies appear to have highly variable returns, and the incidence on climate damages on the US government remains uncertain. Nonetheless, the math suggests these types of policies have the potential to unlock large welfare gains.

[56] Details of this calculation can be found in Appendix J.

References

Acemoglu, D., Aghion, P., Bursztyn, L. & Hemous, D. (2012), 'The environment and directed technical change', *American Economic Review* **102**(1), 131–166.

Allcott, H. (2011), 'Social norms and energy conservation', *Journal of Public Economics* **95**(9-10), 1082–1095.

Allcott, H. (2015), 'Site selection bias in program evaluation', *The Quarterly Journal of Economics* **130**(3), 1117–1165.

Allcott, H. & Kessler, J. B. (2019), 'The welfare effects of nudges: A case study of energy use social comparisons', *American Economic Journal: Applied Economics* **11**(1), 236–276.

Allcott, H. & Rogers, T. (2014), 'The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation', *American Economic Review* **104**(10), 3003–3037.

Andrews, I. & Kasy, M. (2019), 'Identification of and correction for publication bias', *American Economic Review* **109**(8), 2766–2794.

Anthoff, D. & Tol, R. S. (2010), 'On international equity weights and national decision making on climate change', *Journal of Environmental Economics and Management* **60**, 14–20.

Anthoff, D. & Tol, R. S. (2013*a*), 'Erratum to: The uncertainty about the social cost of carbon: A decomposition analysis using fund', *Climatic Change* **121**, 413.

Anthoff, D. & Tol, R. S. (2013*b*), 'The uncertainty about the social cost of carbon: A decomposition analysis using fund', *Climatic Change* **117**, 515–530.

Atkinson, A. B. & Stern, N. H. (1974), 'Pigou, taxation and public goods', *The Review of Economic Studies* **41**(1), 119–128.

Banares-Sanchez, I., Burgess, R., Laszlo, D., Simpson, P., Van Reenen, J. & Wang, Y. (2023), 'Ray of hope? China and the rise of solar energy'.

Bayer, P. & Aklin, M. (2020), 'The European Union Emissions Trading System reduced CO_2 emissions despite low prices', *Proceedings of the National Academy of Sciences* **116**(16), 8804–8812.

Berkouwer, S. B. & Dean, J. T. (2022), 'Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households', *American Economic Review* **112**(10), 3291–3330.

Bilal, A. & Känzig, D. R. (2024), The macroeconomic impact of climate change: Global vs. local temperature, Working Paper 32450, National Bureau of Economic Research.

Bistline, J., Mehrotra, N. & Wolfram, C. (2023), Economic implications of the climate provisions of the Inflation Reduction Act, Technical report, National Bureau of Economic Research.

Bolinger, M., Seel, J., Warner, C. & Robson, D. (2021), *Utility-Scale Solar, 2021 Edition: Empirical Trends in Deployment, Technology, Cost, Performance, PPA Pricing, and Value in the United States*, Department of Energy.

Bollinger, B. & Gillingham, K. (2019), 'Learning-by-doing in solar photovoltaic installations', *Available at SSRN 2342406* .

Boning, W. C., Hendren, N., Sprung-Keyser, B. & Stuart, E. (2023), A welfare analysis of tax audits across the income distribution, Working Paper 31376, National Bureau of Economic Research.

Borenstein, S. (2012), 'The private and public economics of renewable electricity generation', *Journal of Economic Perspectives* **26**(1), 67–92.

Borenstein, S. & Bushnell, J. B. (2022), 'Do two electricity pricing wrongs make a right? cost recovery, externalities, and efficiency', *American Economic Journal: Economic Policy* **14**(4), 80–110.

Borenstein, S., Bushnell, J. & Mansur, E. (2023), 'The economics of electricity reliability', *Journal of Economic Perspectives* **37**(4), 181–206.

Bourguignon, F. & Spadaro, A. (2012), 'Tax–benefit revealed social preferences', *The Journal of Economic Inequality* **10**, 75–108.

Brandon, A., Ferraro, P. J., List, J. A., Metcalfe, R. D., Price, M. K. & Rundhammer, F. (2017), Do the effects of nudges persist? Theory and evidence from 38 natural field experiments, Working Paper 23277, National Bureau of Economic Research.

Brandon, A., List, J. A., Metcalfe, R. D., Price, M. K. & Rundhammer, F. (2019), 'Testing for crowd out in social nudges: Evidence from a natural field experiment in the market for electricity', *Proceedings of the National Academy of Sciences* **116**(12), 5293–5298.

Brown, D. P. & Muehlenbachs, L. (2024), 'The value of electricity reliability: Evidence from battery adoption', *Journal of Public Economics* **239**, 105216.

Buchanan, J. M. (1969), 'External diseconomies, corrective taxes, and market structure', *American Economic Review* **59**(1), 174–177.

Bushnell, J. B., Muehlegger, E. & Rapson, D. S. (2022), Energy prices and electric vehicle adoption, Technical report, National Bureau of Economic Research.

Butera, L., Metcalfe, R., Morrison, W. & Taubinsky, D. (2022), 'Measuring the welfare effects of shame and pride', *American Economic Review* **112**(1), 122–168.

CAISO (2021), 2020 Annual Report on Market Issues and Performance, Technical report, Department of Market Monitoring – California ISO.

Calel, R., Colmer, J., Dechezleprêtre, A. & Glachant, M. (2025), 'Do carbon offsets offset carbon?', *American Economic Journal: Applied Economics* **17**(1), 1–40.

Chan, N. W. & Morrow, J. W. (2019), 'Unintended consequences of cap-and-trade? Evidence from the Regional Greenhouse Gas Initiative', *Energy Economics* **80**, 411–422.

Christensen, P., Francisco, P., Myers, E. & Souza, M. (2023), 'Decomposing the wedge between projected and realized returns in energy efficiency programs', *Review of Economics and Statistics* **105**(4), 798–817.

Clinton, B. C. & Steinberg, D. C. (2019), 'Providing the spark: Impact of financial incentives of battery electric vehicle adoption', *Journal of Environmental Economics and Management* **98**.

Cole, C., Droste, M., Knittel, C., Li, S. & Stock, J. H. (2023), Policies for electrifying the light-duty vehicle fleet in the united states, *in* 'AEA Papers and Proceedings', Vol. 113, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 316–322.

Colmer, J., Martin, R., Muûls, M. & Wagner, U. J. (2024), 'Does pricing carbon mitigate climate Change? Firm-Level evidence from the European Union Emissions Trading System', *The Review of Economic Studies* .

Couture, V., Duranton, G. & Turner, M. A. (2018), 'Speed', *Review of Economics and Statistics* **100**(4), 725–739.

Davis, L. W. (2023), The economic determinants of heat pump adoption, Working Paper 31344, National Bureau of Economic Research.

Davis, L. W., Fuchs, A. & Gertler, P. (2014), 'Cash for coolers: Evaluating a large-scale appliance replacement program in Mexico', *American Economic Journal: Economic Policy* **6**(4), 207–238.

Davis, L. W. & Kilian, L. (2011), 'Estimating the effect of a gasoline tax on carbon emissions', *Journal of Applied Econometrics* **26**(7), 1187–1214.

De Loecker, J., Eeckhout, J. & Unger, G. (2020), 'The rise of market power and the macroeconomic implications', *Quarterly Journal of Economics* **135**, 561–644.

DOI (2021), Consumer Surplus and Energy Substitutes for OCS Oil and Gas Production: The 2021 Revised Market Simulation Model (MarketSim) , Technical report, US Department of the Interior.

EIA (2022), U.S. Total Gasoline DTW Sales Price by Refiners, Technical report, US Energy Information Administration.

EIA (2023), Annual Energy Outlook (2007–2023 Editions), Technical report, US Energy Information Administration.

EIA (2024*a*), Factors affecting gasoline prices, Technical report, US Energy Information Administration.

EIA (2024*b*), Refiner Acquisition Cost of Crude Oil, Data series, US Energy Information Administration.

EIA (2024*c*), U.S. Landed Costs of Crude Oil, Data series, US Energy Information Administration.

Enkvist, P., Nauclér, T. & Rosander, J. (2007), 'A cost curve for greenhouse gas reduction', *McKinsey Quarterly* **1**, 34.

Environmental Defense Fund (2021), A revamped cost curve for reaching net-zero emissions, Technical report.

EPA (2010), Guidelines for preparing economic analyses, Technical report, US Environmental Protection Agency, National Center for Environmental Economics, Office of Policy.

EPA (2023), Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances, Technical report, US Environmental Protection Agency.

EPA (2024), AVoided Emissions and geneRation Tool (AVERT), Technical report, US Environmental Protection Agency.

Favennec, J.-P. (2022), Economics of oil refining, *in* M. Hafner & G. Luciani, eds, 'The Palgrave Handbook of International Energy Economics', Palgrave Macmillan, Cham, pp. 59–74.

Fournel, J.-F. (2024), Electric vehicle subsidies: Cost-effectiveness and emission reductions, Working paper, Toulouse School of Economics.

Gillingham, K. & Stock, J. H. (2018), 'The cost of reducing greenhouse gas emissions', *Journal of Economic Perspectives* **32**(4), 53–72.

Gillingham, K. T. & Bollinger, B. (2021), 'Social learning and solar photovoltaic adoption', *Management Science* **67**(11), 7091–7112.

Gillingham, K. T., Ovaere, M., Weber, S. M. & Yates, A. J. (2025), 'Carbon policy and the emissions implications of electric vehicles', *Journal of the Association of Environmental and Resource Economists* **12**(2).

Gillingham, K. & Tsvetanov, T. (2019), 'Hurdles and steps: Estimating demand for solar photovoltaics', *Quantitative Economics* **10**, 275–310.

Glennerster, R. & Jayachandran, S. (2023), 'Think globally, act globally: Opportunities to mitigate greenhouse gas emissions in low-and middle-income countries', *Journal of Economic Perspectives* **37**(3), 111–135.

Greenstone, M., Mukhametkaliev, B., Stolove, J., Larsen, J., King, B., Kolus, H. & Herndon, W. (2022), 'Assessing the costs and benefits of clean electricity tax credits', *EPIC and Rhodium Group* .

Grubb, M., Edmonds, J., Ten Brink, P. & Morrison, M. (1993), 'The costs of limiting fossil-fuel CO_2 emissions: A survey and analysis', *Annual Review of Energy and the environment* **18**(1), 397–478.

Hanna, R., Duflo, E. & Greenstone, M. (2016), 'Up in smoke: The influence of household behavior on the long-run impact of improved cooking stoves', *American Economic Journal: Economic Policy* **8**(1), 80–114.

Hendren, N. (2020), 'Measuring economic efficiency using inverse-optimum weights', *Journal of Public Economics* **187**, 104198.

Hendren, N. & Sprung-Keyser, B. (2020), 'A unified welfare analysis of government policies', *The Quarterly Journal of Economics* **135**(3), 1209–1318.

Hernandez-Cortes, D. & Meng, K. C. (2023), 'Do environmental markets cause environmental injustice? Evidence from California's carbon market', *Journal of Public Economics* **217**.

Hitaj, C. (2013), 'Wind power development in the United States', *Journal of Environmental Economics and Management* **65**(3), 394–410.

Holland, S. P., Mansur, E. T., Muller, N. Z. & Yates, A. J. (2016), 'Are there environmental benefits from driving electric vehicles? The importance of local factors', *American Economic Review* **106**(12), 3700–3729.

Hope, C. (2006), 'The marginal impact of CO_2 from PAGE2002: An integrated assessment model incorporating the IPCC's five reasons for concern', *The Integrated Assessment Journal* **6**, 19–56.

Hope, C. (2008), 'Discount rates, equity weights and the social cost of carbon', *Energy Economics* **30**, 1011–1019.

Hughes, J. E., Knittel, C. R. & Sperling, D. (2008), 'Evidence of a shift in the short-run price elasticity of gasoline demand', *The Energy Journal* **29**(1), 113–134.

Hughes, J. E. & Podolefsky, M. (2015), 'Getting green with solar subsidies: Evidence from the California Solar Initiative', *Journal of the Association of Environmental and Resource Economists* **2**(2), 235—-275.

IEA (2020), GHG abatement costs for selected measures of the Sustainable Recovery Plan, Technical report, International Energy Agency.

Interagency Working Group (2021), Technical support document: Social cost of carbon, methane, and nitrous oxide, Technical report, Interagency Working Group on Social Cost of Greenhouse Gases, United States Government.

Izquierdo-Tort, S., Jayachandran, S. & Saavedra, S. (2024), Redesigning payments for ecosystem services to increase cost-effectiveness, Working Paper 32689, National Bureau of Economic Research.

Jacobs, B., Jongen, E. & Zoutman, F. (2017), 'Revealed social preferences of dutch political parties', *Journal of Public Economics* **156**, 81–100.

Jacobsen, M. R. (2013), 'Fuel economy and safety: The influences of vehicle class and driver behavior', *American Economic Journal: Applied Economics* **5**(3), 1–26.

Jenkins, J. D. & Mayfield, E. (2023), Rapid energy policy evaluation and analysis toolkit, Technical report, Princeton University Zero Lab.

Kaplow, L. (2023), 'Competition policy in a simple general equilibrium model', *Journal of Political Economy Microeconomics* **1**(1), 80–114.

Kleven, H. J. & Kreiner, C. T. (2006), 'The marginal cost of public funds: Hours of work versus labor force participation', *Journal of Public Economics* **90**(10-11), 1955–1973.

Kleven, H., Kreiner, C., Larsen, K. & Søgaard, J. (2024), 'Micro vs macro labor supply elasticities: The role of dynamic returns to effort', *American Economic Review (Forthcoming)* .

Knittel, C. R. (2009), 'The implied cost of carbon dioxide under the Cash for Clunkers program'.

Lafond, F., Greenwald, D. & Farmer, J. D. (2022), 'Can stimulating demand drive costs down? world war ii as a natural experiment', *The Journal of Economic History* **82**(3), 727–764.

Li, S., Tong, L., Xing, J. & Zhou, Y. (2017), 'The market for electric vehicles: Indirect network effects and policy design', *Journal of the Association of Environmental and Resource Economists* **4**(1), 89–133.

List, J. A., Rodemeier, M., Roy, S. & Sun, G. K. (2023), Judging nudging: Understanding the welfare effects of nudges versus taxes, Technical report, National Bureau of Economic Research.

Marion, J. & Muehlegger, E. (2011), 'Fuel tax incidence and supply conditions', *Journal of Public Economics* **95**, 1202–1212.

Metcalf, G. E. (2010), 'Investment in energy infastructure and the tax code', *Tax Policy and the Economy* **24**(1), 1–34.

Metcalf, G. E. & Stock, J. H. (2023), 'The macroeconomic impact of Europe's carbon taxes', *American Economic Journal: Macroeconomics* **15**(3), 265–86.

Mirrlees, J. A. (1971), 'An exploration in the theory of optimum income taxation', *The Review of Economic Studies* **38**(2), 175–208.

Mirrlees, J. A. (1976), 'Optimal tax theory: A synthesis', *Journal of Public Economics* **6**(4), 327–358.

Muehlegger, E. & Rapson, D. S. (2022), 'Subsidizing low-and middle-income adoption of electric vehicles: Quasi-experimental evidence from California', *Journal of Public Economics* **216**.

Mullainathan, S. & Allcott, H. (2010), 'Behavior and energy policy', *Science* **327**, 1204–1205.

Nath, I. B., Ramey, V. A. & Klenow, P. J. (2024), How much will global warming cool global growth?, Working paper.

Nemet, G. F. (2006), 'Beyond the learning curve: Factors influencing cost reductions in photovoltaics', *Energy Policy* **34**(17), 3218–3232.

Nesje, F., Drupp, M. A., Freeman, M. C. & Groom, B. (2023), 'Philosophers and economists agree on climate policy paths but for different reasons', *Nature Climate Change* **13**(6), 515–522.

Nordhaus, W. D. (1993), 'Optimal greenhouse-gas reductions and tax policy in the "DICE" model', *American Economic Review* **83**(2), 313–317.

Nordhaus, W. D. (2010), 'Economic aspects of global warming in a post-Copenhagen environment', *Proceedings of the National Academy of Sciences* **106**(26), 11721–11726.

Nordhaus, W. D. (2014*a*), 'Estimates of the social cost of carbon: Concepts and results from the DICE-2013R model and alternative approaches', *Journal of the Association of Environmental and Resource Economists* **1**(1/2), 273–312.

Nordhaus, W. D. (2014*b*), 'The perils of the learning model for modeling endogenous technological change', *The Energy Journal* **35**(1), 1–14.

Nordhaus, W. D. (2017), 'Revisiting the social cost of carbon', *Proceedings of the National Academy of Sciences* **114**(7), 1518–1523.

OECD (2021), *Revenue Statistics 2021*, OECD.

Parry, I. W. H. & Small, K. A. (2005), 'Does Britain or the United States have the right gasoline tax?', *American Economic Review* **95**(4), 1276–1289.

Parry, I. W., Heine, M. D., Lis, E. & Li, S. (2014), *Getting energy prices right: From principle to practice*, International Monetary Fund.

Pless, J. & van Benthem, A. A. (2019), 'Pass-through as a test for market power: An application to solar subsidies', *American Economic Journal: Applied Economics* **11**(4), 367–401.

Prest, B. C., Rennels, L., Errickson, F. & Anthoff, D. (2024), 'Equity weighting increases the social cost of carbon', *Science* **385**(6710), 715–717.
URL: *https://www.science.org/doi/abs/10.1126/science.adn1488*

Rennert, K., Errickson, F., Prest, B. C., Rennels, L., Newell, R. G., Pizer, W., Kingdon, C., Wingenroth, J., Cooke, R., Parthum, B., Smith, D., Cromar, K., Diaz, D., Moore, F. C., Müller, U. K., Plevin, R. J., Raftery, A. E., Ševčíková, H., Sheets, H., Stock, J. H., Tan, T., Watson, M., Wong, T. E. & Anthoff, D. (2022), 'Comprehensive evidence implies a higher social cost of co2', *Nature* **610**(7933), 687–692.

Romer, P. M. (1986), 'Increasing returns and long-run growth', *Journal of Political Economy* **94**(5), 1002–1037.

Rubin, E. S., Azevedo, I. M., Jaramillo, P. & Yeh, S. (2015), 'A review of learning rates for electricity supply technologies', *Energy Policy* **86**, 198–218.

Saez, E. (2001), 'Using elasticities to derive optimal income tax rates', *Review of Economic Studies* **68**, 205–229.

Shrimali, G., Lynes, M. & Indvik, J. (2015), 'Wind energy deployment in the U.S.: An empirical analysis of the role of federal and state policies', *Renewable and Sustainable Energy Reviews* **43**, 796–806.

Small, K. A. & Van Dender, K. (2007), 'Fuel efficiency and motor vehicle travel: The declining rebound effect', *The Energy Journal* **28**(1), 21–51.

Söderholm, P. & Sundqvist, T. (2007), 'Empirical challenges in the use of learning curves for assessing the economic prospects of renewable energy technologies', *Renewable Energy* **32**(15), 2559–2578.

Stiglitz, J. E. & Dasgupta, P. (1971), 'Differential taxation, public goods, and economic efficiency', *The Review of Economic Studies* **38**(2), 151–174.

Tschofen, P., Azevedo, I. L. & Muller, N. Z. (2019), 'Fine particulate matter damages and value added in the US economy', *Proceedings of the National Academy of Sciences* **116**(40), 19857–19862.

United Nations (2023), Development of default values for fraction of non-renewable biomass, Information Note CDM-MP92-A07, United Nations.

van Benthem, A., Gillingham, K. & Sweeney, J. (2008), 'Learning-by-doing and the optimal solar policy in California', *The Energy Journal* **29**(3), 131–152.

Way, R., Ives, M. C., Mealy, P. & Farmer, J. D. (2022), 'Empirically grounded technology forecasts and the energy transition', *Joule* **6**(9), 2057–2082.

West, S. E. & Williams, R. C. (2007), 'Optimal taxation and cross-price effects on labor supply: Estimates of the optimal gas tax', *Journal of Public Economics* **91**(3), 593–617.

Winjobi, O., Kelly, J. C. & Dai, Q. (2022), 'Life-cycle analysis, by global region, of automotive lithium-ion nickel manganese cobalt batteries of varying nickel content', *Sustainable Materials and Technologies* **32**.

Wiser, R., Bolinger, M., Hoen, B., Millstein, D., Rand, J., Barbose, G., Darghouth, N., Gorman, W., Jeong, S., O'Shaughnessy, E. & Paulos, B. (2023), *Land-Based Wind Market Report: 2023 Edition*, DOE.

Wright, T. P. (1936), 'Factors affecting the cost of airplanes', *Journal of the aeronautical sciences* **3**(4), 122–128.

Zhao, L., Ottinger, E. R., Yip, A. H. C. & Helveston, J. P. (2023), 'Quantifying electric vehicle mileage in the United States', *Joule* **7**(11), 2537–2551.

FIGURE 1: Electric Vehicle Subsidy
Baseline Estimates from Muehlegger and Rapson (2022)

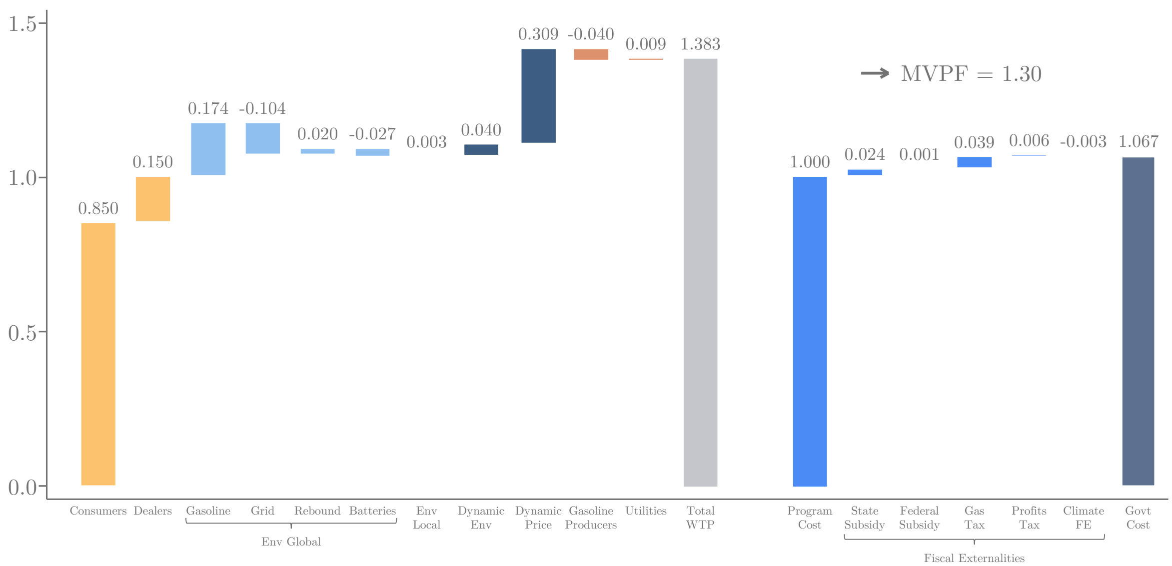

Notes: This figure presents the components of willingness to pay and net government cost for the EV subsidies in the California Enhanced Modernization Program (CEFMP) using the -2.1 price elasticity estimated in Muehlegger & Rapson (2022). We present estimates for our baseline specification that envisions a change to the federal 2020 subsidy. Each component is normalized relative to $1 of mechanical cost of the policy change. The first two bars show how this transfer is passed through to consumers and car dealers. The next three bars report the environmental externalities, including the global (GHG) externalities, local (e.g. $PM_{2.5}$) externalities, and rebound effects from higher prices in the electricity market. The next two bars report learning-by-doing externalities from both future environmental benefits (DE) and lower prices (DP) using the approach in Theorem 1 and Appendix B. The last two columns report impacts on producer profits due to markups in the oil/gasoline and utility sectors. The Cost components start with the mechanical cost of the $1 subsidy, then add the impact of the behavioral response on the cost of state and federal subsidies using national average subsidies in 2020, followed by the impact on changes in revenue from the gas tax and corporate profits taxes on oil/gasoline producers and utilities. Lastly, the climate FE term captures future tax revenue due to the impact of lower emissions today on future productivity. All numbers are calculated using our baseline path for the social cost of carbon ($193 in 2020) and a 2% discount rate.

FIGURE 2: Utility-Scale Wind Subsidies & Production Tax Credits

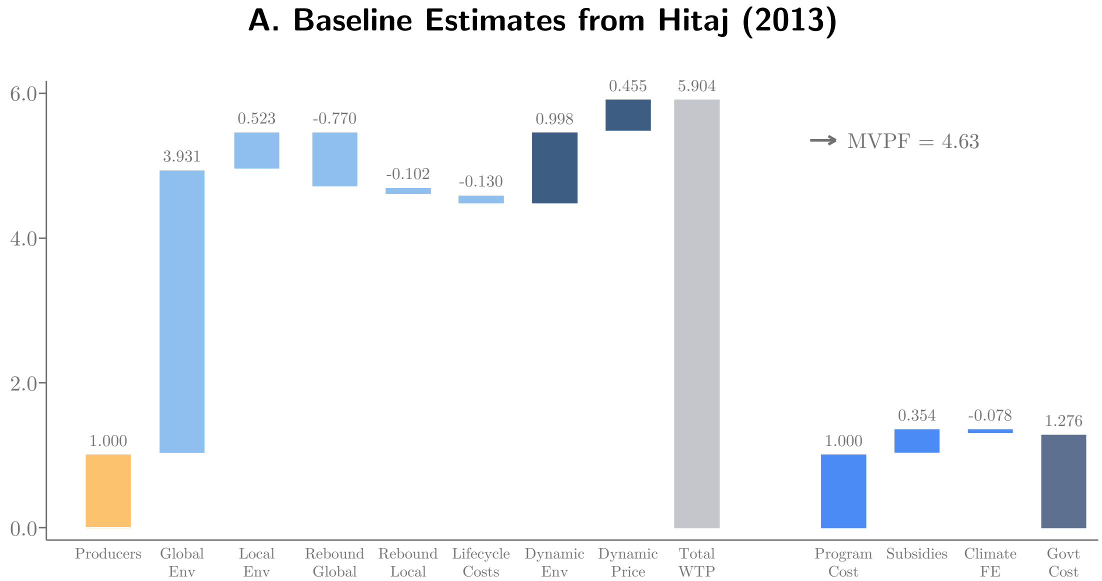

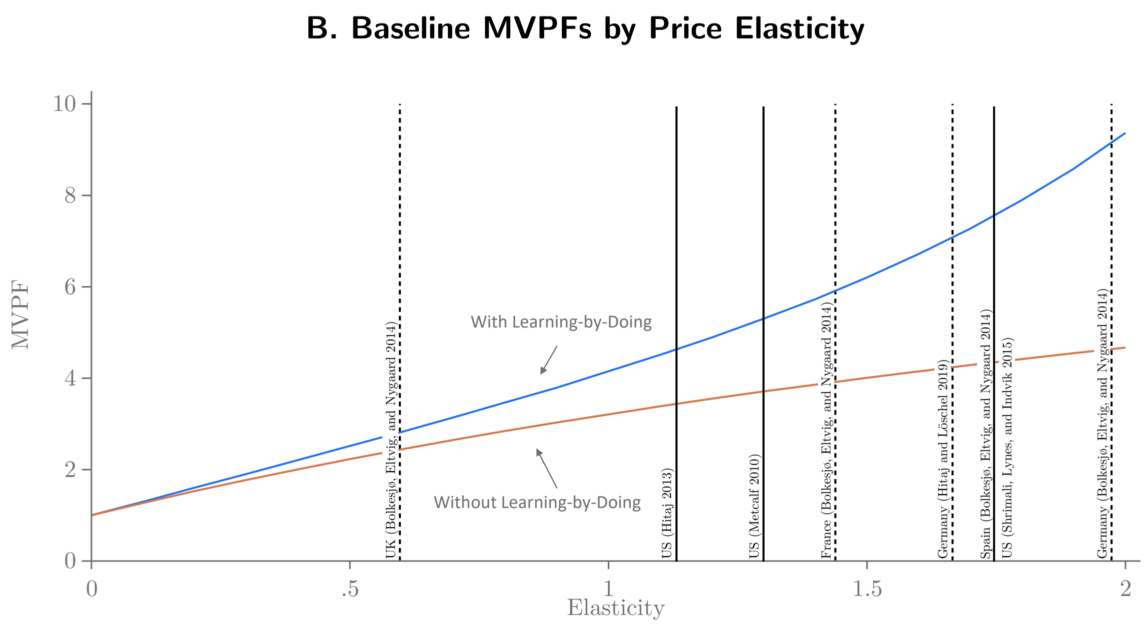

Notes: This figure illustrates the MVPF measurement for wind subsidies. Panel A shows the WTP and Cost components for the baseline specification for the wind production tax credit using a supply elasticity of 1.4 estimated in Hitaj (2013). The WTP components consist of the transfer (yellow), environmental externality (light blue), and learning by doing effects (dark blue). The subsidy cost is calculated using the wind PTC in 2020 of \$0.015 per KWh. Panel B shows how the MVPF varies with the elasticity of wind turbine installation with respect to the price paid to suppliers for wind energy. We place solid vertical lines at the US estimates of the elasticities in our main sample and dotted vertical lines for international estimates in our extended sample. All numbers are calculated using our baseline path for the social cost of carbon (\$193 in 2020) and a 2% discount rate.

FIGURE 3: Residential Solar Subsidies

A. Baseline Estimates from Pless and Van Bentham (2019)

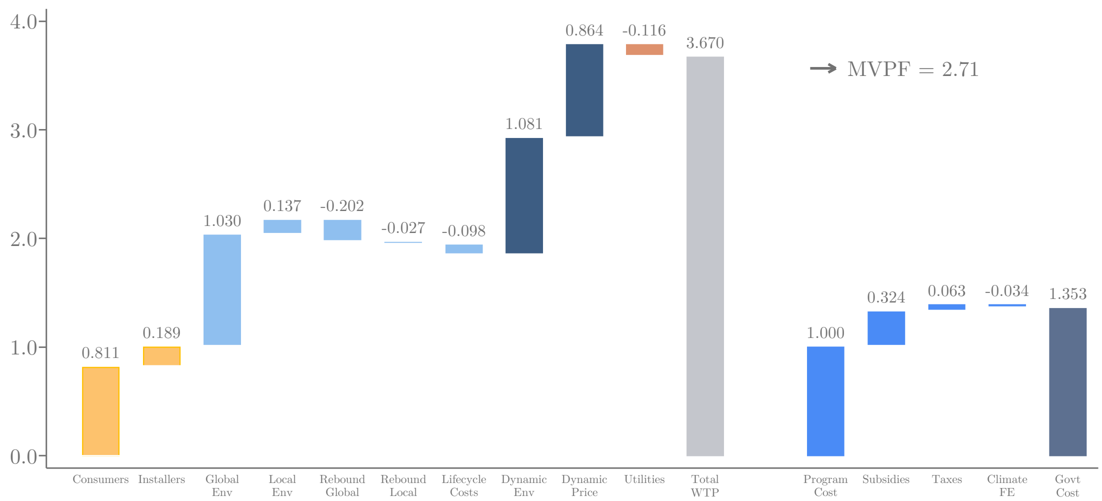

B. Baseline MVPFs by Price Elasticity

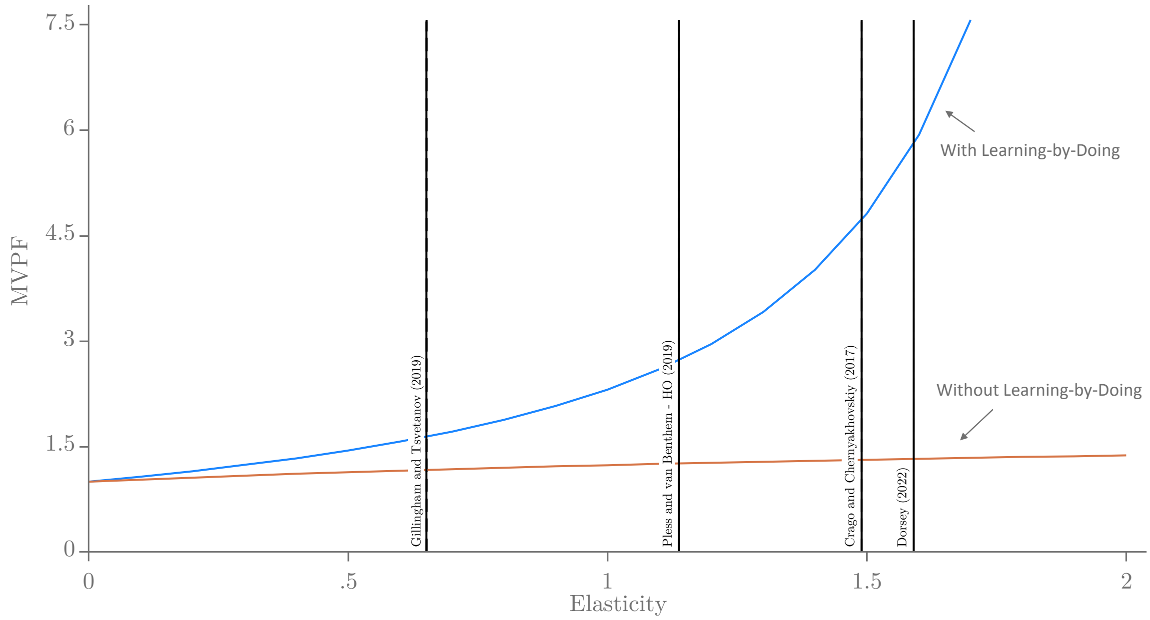

Notes: This figure illustrates the MVPF measurement for residential solar subsidies. Panel A shows the WTP and Cost components for our baseline specification for the California Solar Initiative using a demand elasticity of -1.14 estimated in Pless & van Benthem (2019). The WTP components consists of the transfer (yellow), environmental externality (light blue), learning by doing effects (dark blue), and utility profit loss (orange). The subsidy cost is calculated using the 26% investment tax credit for residential solar installations. Panel B shows how the MVPF varies with the elasticity of demand for residential solar panel capacity with respect to the price of residential solar panels. The MVPF with learning by doing is not shown above 7.5 for illustrative purposes. The solid lines represent the estimates of the elasticity in our sample. All numbers are calculated using our baseline path for the social cost of carbon ($193 in 2020) and a 2% discount rate.

FIGURE 4: Baseline MVPFs for Subsidies

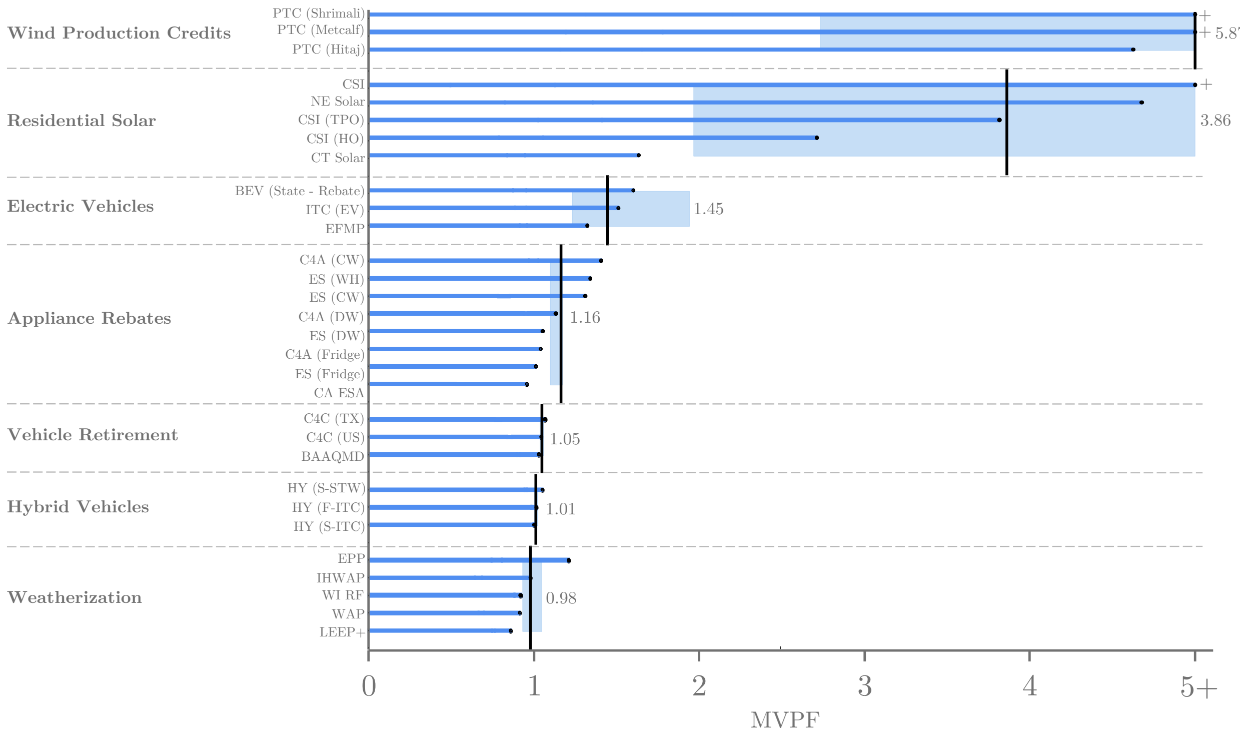

Notes: This figure shows the 2020 baseline MVPF estimates for all categorized subsidy policies in our main sample. We cap estimates at 5 with + signs indicating MVPFs above 5. The category average (shown by the black vertical lines) reports the MVPF associated with a conceptual experiment where $1 in initial program cost is split equally across each policy in the category, so that we take the average willingness to pay relative to the average net government cost within each category. The blue shading presents bootstrapped 95% confidence intervals for each category average MVPF, restricting to underlying estimates for which we have sampling uncertainty. See Appendix Table 1 for the full policy name and source to each label reference. See Appendix Table 3 for comparisons of the category averages on this subsample. All numbers are calculated using our baseline path for the social cost of carbon ($193 in 2020) and a 2% discount rate.

FIGURE 5: Baseline MVPF of Home Energy Reports

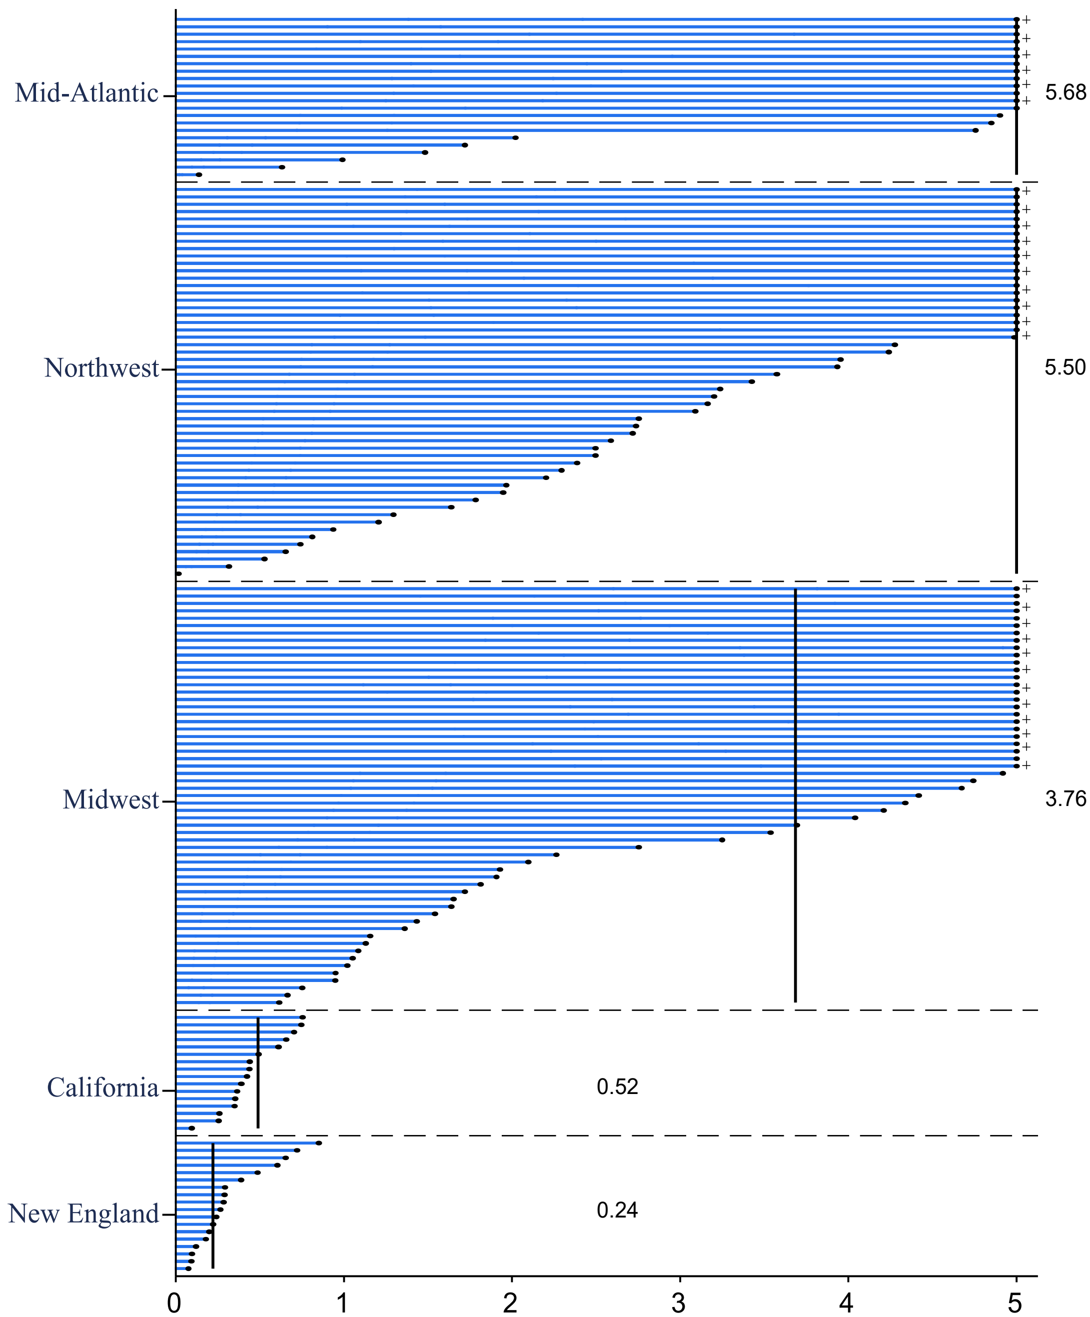

Notes: This figure illustrates the MVPF estimates for Opower Home Energy Reports split across the 5 AVERT model's electricity regions for which the experiments have been conducted. The benefits per dollar of government cost equal the environmental benefits minus the loss in utility profits. MVPFs above five are censored and the category averages are written to the right of each category. All numbers are calculated using our baseline path for the social cost of carbon ($193 in 2020) and a 2% discount rate.

FIGURE 6: MVPF of a Gasoline Tax
Baseline Estimates from Small & Van Dender (2007)

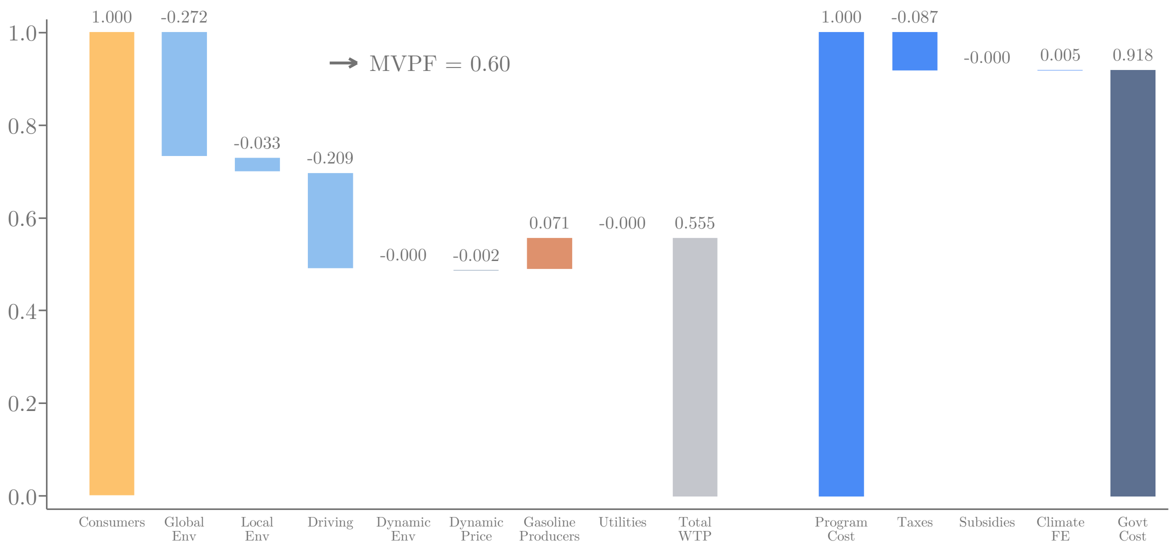

Notes: This figure presents the components of the baseline MVPF for the gasoline tax using a gasoline price elasticity of -0.334 from Small & Van Dender (2007). The WTP components include the transfer cost (yellow), global greenhouse gas benefits and local environmental externalities arising from accidents, congestion, and local pollutants (light blue), learning by doing benefits from increased EV purchases (bars not visible), and gasoline/electricity producer profits (orange). The tax cost arises from the impact of the response to the tax on gas tax revenue using the 2020 tax of $0.46 per gallon. All numbers are calculated using our baseline path for the social cost of carbon ($193 in 2020) and a 2% discount rate.

FIGURE 7: Baseline MVPFs of Revenue Raisers

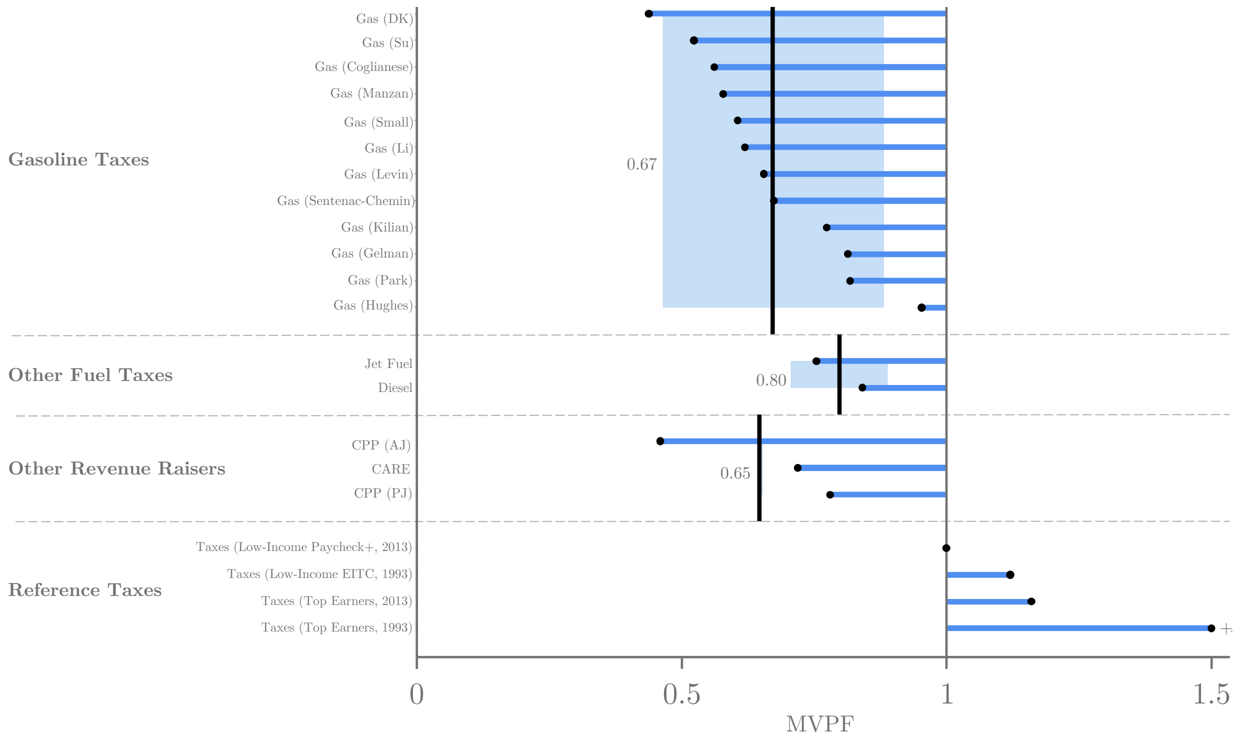

Notes: This figure illustrates the MVPF for revenue raisers in our sample. Note that the MVPF measures the welfare cost per dollar of revenue raised (or, equivalently, the welfare gain per dollar of net expenditures on tax cuts). We illustrate each MVPF relative to the MVPF of a non-distortionary lump sum tax of 1. The black lines are the category averages and the blue regions indicate the 95% confidence intervals computed via bootstrap. See Appendix Table 1 for the full policy name and source to each label reference. All numbers are calculated using our baseline path for the social cost of carbon ($193 in 2020) and a 2% discount rate.

FIGURE 8: Baseline MVPFs of International Policies

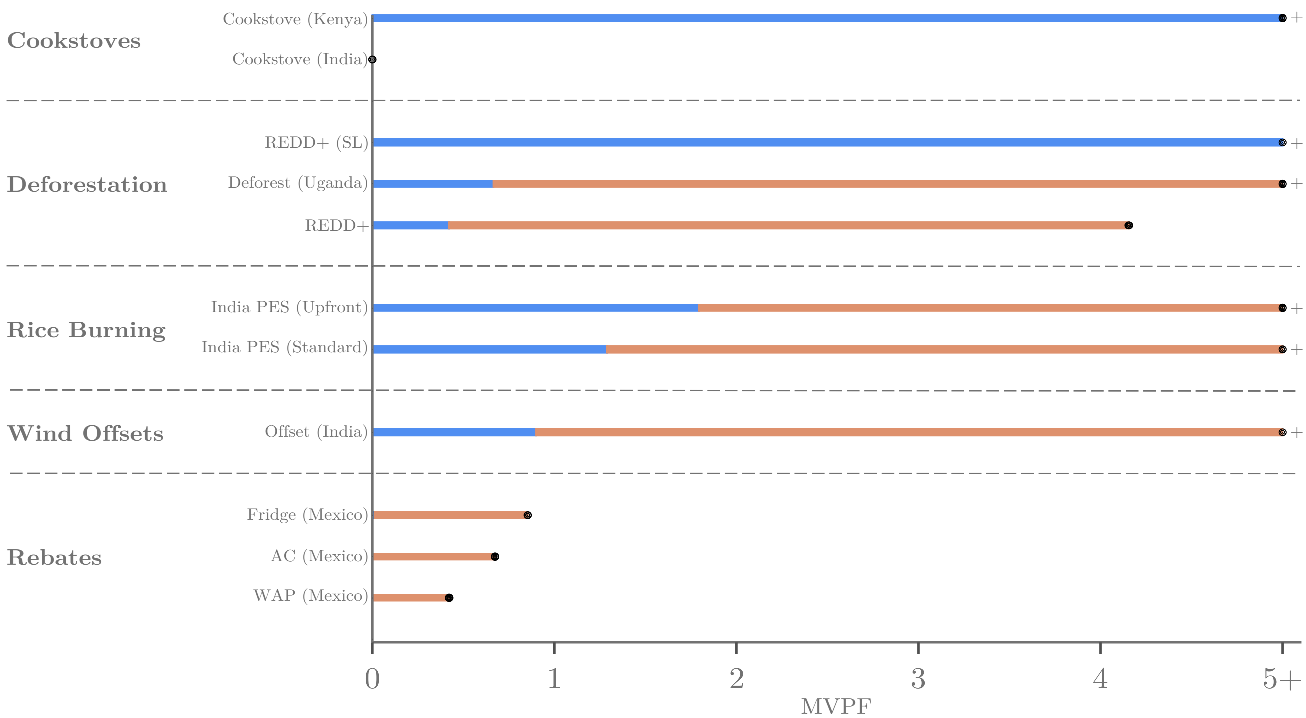

Notes: This figure illustrates the 2020 baseline MVPF estimates for US spending on international policies. The denominator is net cost to the US government and the numerator is the sum of US and non-US WTP for the subsidy. We cap estimates at 5 with + signs indicating MVPFs above 5. The blue bars represent the MVPF only including US beneficiaries and the orange bars illustrate how the MVPF increases if one includes benefits to non-US residents. See Appendix Table 1 for the full policy name and source to each label reference. All numbers are calculated using our baseline path for the social cost of carbon ($193 in 2020) and a 2% discount rate.

Table 1: Baseline MVPF Components

Full Version

	Willingness to Pay								Cost				
		Environmental Benefits			Learning by Doing					Fiscal Externalities			
Panel A. Subsidies	**Transfer**	**Global**	**Local**	**Rebound**	**Env.**	**Price**	**Profits**	**WTP**	**Program**	**Initial**	**Climate**	**Total**	**MVPF**
Wind Production Credits	**1.000**	**4.678**	**0.643**	**-1.074**	**1.900**	**0.645**		**7.793**	**1.000**	**0.435**	**-0.108**	**1.328**	**5.870**
PTC (Shrimali)	1.000	5.865	0.806	-1.346	3.277	0.920		10.522	1.000	0.546	-0.152	1.394	7.547
PTC (Metcalf)	1.000	4.368	0.601	-1.002	1.427	0.560		6.953	1.000	0.407	-0.094	1.312	5.298
PTC (Hitaj)	1.000	3.801	0.523	-0.872	0.998	0.455		5.904	1.000	0.354	-0.078	1.276	4.626
FIT (Germany - BEN) *	1.000	6.629	0.911	-1.521	4.841	1.170		13.030	1.000	0.617	-0.193	1.424	9.148
FIT (Spain) *	1.000	5.866	0.806	-1.346	3.277	0.920		10.522	1.000	0.546	-0.152	1.394	7.547
FIT (Germany - HL) *	1.000	5.596	0.769	-1.284	2.844	0.844		9.768	1.000	0.521	-0.140	1.381	7.072
FIT (France) *	1.000	4.837	0.665	-1.110	1.877	0.658		7.926	1.000	0.450	-0.110	1.340	5.913
FIT (UK) *	1.000	2.006	0.276	-0.460	0.223	0.199		3.243	1.000	0.187	-0.035	1.151	2.817
FIT (EU) *	1.000	0.546	0.075	-0.125	0.016	0.050		1.561	1.000	0.051	-0.009	1.042	1.498
Residential Solar	**1.106**	**1.718**	**0.252**	**-0.421**	**2.280**	**1.636**	**-0.214**	**6.356**	**1.000**	**0.714**	**-0.068**	**1.646**	**3.862**
CSI	1.000	4.299	0.631	-1.054	4.988	3.987	-0.535	13.316	1.000	1.787	-0.157	2.630	5.063
NE Solar	1.000	1.220	0.179	-0.299	3.132	1.610	-0.152	6.690	1.000	0.507	-0.076	1.431	4.676
CSI (TPO)	1.528	1.604	0.235	-0.393	1.982	1.371	-0.200	6.128	1.000	0.667	-0.061	1.606	3.815
CSI (HO)	1.000	0.932	0.137	-0.228	1.081	0.864	-0.116	3.670	1.000	0.387	-0.034	1.353	2.712
CT Solar	1.000	0.533	0.078	-0.131	0.216	0.346	-0.066	1.976	1.000	0.222	-0.012	1.209	1.634
ITC *	1.000	1.152	0.169	-0.282	3.825	1.944	-0.143	7.664	1.000	0.531	-0.088	1.443	5.312
Electric Vehicles	**1.000**	**0.057**	**0.000**	**0.032**	**0.073**	**0.452**	**-0.043**	**1.571**	**1.000**	**0.092**	**-0.004**	**1.087**	**1.445**
BEV (State - Rebate)	1.000	0.068	0.000	0.038	0.103	0.564	-0.051	1.722	1.000	0.108	-0.006	1.103	1.561
ITC (EV)	1.000	0.061	0.000	0.034	0.078	0.482	-0.046	1.609	1.000	0.097	-0.005	1.092	1.474
EFMP	1.000	0.042	0.000	0.023	0.040	0.309	-0.031	1.383	1.000	0.070	-0.003	1.067	1.296
BEV (State - ITC) *	1.000	-0.048	0.000	-0.027	0.000	0.000	0.036	0.961	1.000	-0.076	0.003	0.927	1.037
Appliance Rebates	**0.867**	**0.497**	**0.043**	**-0.089**			**-0.103**	**1.215**	**1.000**	**0.052**	**-0.009**	**1.044**	**1.164**
C4A (CW)	0.952	0.550	0.083	-0.124			-0.039	1.423	1.000	0.021	-0.009	1.012	1.405
ES (WH)	0.598	1.707	0.000	-0.201			-0.659	1.445	1.000	0.112	-0.033	1.078	1.340
ES (CW)	1.000	0.861	0.126	-0.193			-0.072	1.722	1.000	0.328	-0.014	1.315	1.310
C4A (DW)	0.929	0.243	0.037	-0.055			-0.017	1.138	1.000	0.009	-0.004	1.005	1.132
ES (DW)	1.000	-0.223	-0.033	0.050			0.019	0.813	1.000	-0.231	0.003	0.772	1.053
C4A (Fridge)	0.960	0.099	0.015	-0.022			-0.007	1.044	1.000	0.004	-0.002	1.002	1.042
ES (Fridge)	1.000	0.199	0.029	-0.045			-0.017	1.167	1.000	0.157	-0.003	1.154	1.011
CA ESA	0.500	0.541	0.083	-0.122			-0.034	0.968	1.000	0.018	-0.008	1.010	0.958
Vehicle Retirement	**0.910**	**0.280**	**0.102**	**-0.137**			**-0.049**	**1.106**	**1.000**	**0.060**	**-0.004**	**1.056**	**1.047**
C4C (TX)	1.000	0.410	0.030	-0.208			-0.074	1.157	1.000	0.091	-0.006	1.084	1.067
C4C (US)	1.000	0.271	0.020	-0.140			-0.049	1.102	1.000	0.060	-0.004	1.055	1.044
BAAQMD	0.730	0.161	0.255	-0.062			-0.025	1.059	1.000	0.031	-0.003	1.028	1.030

Hybrid Vehicles	**1.000**	**0.031**	**0.003**	**-0.026**	**0.000**	**0.014**	**-0.006**	**1.016**	**1.000**	**0.004**	**-0.001**	**1.004**	**1.012**
HY (S-STW)	1.000	0.069	0.007	-0.058	0.001	0.031	-0.014	1.036	1.000	0.010	-0.002	1.008	1.027
HY (F-ITC)	1.000	0.020	0.002	-0.017	0.000	0.009	-0.004	1.010	1.000	0.003	0.000	1.002	1.008
HY (S-ITC)	1.000	0.004	0.000	-0.004	0.000	0.002	-0.001	1.002	1.000	0.001	0.000	1.001	1.002
Weatherization	**0.774**	**0.297**	**0.029**	**-0.057**			**-0.054**	**0.989**	**1.000**	**0.017**	**-0.005**	**1.012**	**0.978**
EPP	0.750	0.593	0.083	-0.133			-0.057	1.237	1.000	0.031	-0.009	1.022	1.210
IHWAP	0.750	0.404	0.019	-0.064			-0.111	0.999	1.000	0.025	-0.007	1.019	0.980
WI RF	0.870	0.052	0.011	-0.012			-0.001	0.920	1.000	0.001	-0.001	1.000	0.920
WAP	0.750	0.297	0.013	-0.045			-0.088	0.927	1.000	0.018	-0.005	1.013	0.915
LEEP+	0.750	0.138	0.019	-0.031			-0.013	0.864	1.000	0.007	-0.002	1.005	0.859
Other Subsidies	**0.887**	**1.504**	**0.424**	**-0.234**			**-0.065**	**2.517**	**1.000**	**0.036**	**-0.025**	**1.010**	**2.492**
CA 20/20	0.882	2.090	0.297	-0.468			-0.131	2.671	1.000	0.071	-0.033	1.038	2.572
CRP	0.893	0.919	0.552	0.000			0.000	2.363	1.000	0.000	-0.018	0.982	2.407
Panel B. Nudges and Marketing													
Home Energy Reports													
HER (17 RCTs)	0.000	3.872	0.439	-0.844			-0.244	3.222	1.000	0.133	-0.061	1.072	3.006
Opower Elec. (166 RCTs)	0.000	3.246	0.368	-0.708			-0.205	2.701	1.000	0.111	-0.051	1.060	2.548
PER	0.000	0.230	0.064	0.000			0.695	0.989	1.000	-0.378	-0.004	0.618	1.600
Opower Nat. Gas (52 RCTs)	0.000	0.950	0.000	-0.112			-0.367	0.472	1.000	0.062	-0.016	1.046	0.451
Other Nudges	**0.507**	**4.799**	**0.613**	**-1.061**			**-0.659**	**4.199**	**1.000**	**2.243**	**-0.076**	**3.167**	**1.326**
Audit Nudge	0.000	8.678	1.333	-1.961			-0.542	7.507	1.000	2.683	-0.136	3.547	2.117
Solarize	1.145	15.001	2.200	-3.678			-1.844	12.824	1.000	6.320	-0.230	7.091	1.809
ES (WH) + Nudge	0.416	1.630	0.000	-0.192			-0.629	1.225	1.000	0.107	-0.032	1.075	1.140
IHWAP + Nudge (H)	0.739	0.517	0.019	-0.085			-0.105	1.085	1.000	0.023	-0.008	1.015	1.069
IHWAP + Nudge (L)	0.743	0.500	0.018	-0.082			-0.101	1.078	1.000	0.022	-0.008	1.014	1.062
WAP + Nudge	0.000	2.467	0.107	-0.371			-0.732	1.471	1.000	4.300	-0.041	5.259	0.280
Food Labels *	0.000	6.170	0.000	0.000			0.000	6.170	1.000	0.000	-0.120	0.880	7.015
Panel C. Revenue Raisers													
Gasoline Taxes	**1.000**	**-0.229**	**-0.204**		**0.000**	**-0.002**	**0.060**	**0.625**	**1.000**	**-0.074**	**0.004**	**0.931**	**0.671**
Gas (DK)	1.000	-0.374	-0.333		0.000	-0.002	0.098	0.388	1.000	-0.120	0.007	0.887	0.437
Gas (Su)	1.000	-0.323	-0.288		0.000	-0.002	0.084	0.472	1.000	-0.104	0.006	0.903	0.523
Gas (Coglianese)	1.000	-0.299	-0.267		0.000	-0.002	0.078	0.510	1.000	-0.096	0.006	0.910	0.561
Gas (Manzan)	1.000	-0.289	-0.257		0.000	-0.002	0.075	0.527	1.000	-0.093	0.006	0.913	0.578
Gas (Small)	1.000	-0.272	-0.242		0.000	-0.002	0.071	0.555	1.000	-0.087	0.005	0.918	0.605
Gas (Li)	1.000	-0.263	-0.234		0.000	-0.002	0.069	0.570	1.000	-0.084	0.005	0.921	0.619
Gas (Levin)	1.000	-0.240	-0.214		0.000	-0.002	0.063	0.607	1.000	-0.077	0.005	0.928	0.654
Gas (Sentenac-Chemin)	1.000	-0.228	-0.203		0.000	-0.002	0.060	0.627	1.000	-0.073	0.004	0.931	0.673
Gas (Kilian)	1.000	-0.161	-0.143		0.000	-0.002	0.042	0.736	1.000	-0.052	0.003	0.951	0.773
Gas (Gelman)	1.000	-0.133	-0.119		0.000	-0.002	0.035	0.781	1.000	-0.043	0.003	0.960	0.814
Gas (Park)	1.000	-0.130	-0.116		0.000	-0.002	0.034	0.786	1.000	-0.042	0.003	0.961	0.818
Gas (Hughes)	1.000	-0.034	-0.030		0.000	-0.002	0.009	0.943	1.000	-0.011	0.001	0.990	0.953

Gas (West) *	1.000	-0.373	-0.332		0.000	-0.002	0.097	0.391	1.000	-0.120	0.007	0.888	0.440
Gas (Tiezzi) *	1.000	-0.354	-0.315		0.000	-0.002	0.093	0.421	1.000	-0.114	0.007	0.893	0.472
Gas (Bento) *	1.000	-0.285	-0.254		0.000	-0.002	0.074	0.534	1.000	-0.091	0.006	0.914	0.584
Gas (Hughes - Ext) *	1.000	-0.272	-0.243		0.000	-0.002	0.071	0.554	1.000	-0.088	0.005	0.918	0.604
Gas (Kilian - Ext) *	1.000	-0.255	-0.227		0.000	-0.002	0.067	0.582	1.000	-0.082	0.005	0.923	0.630
Gas (Small - Ext) *	1.000	-0.054	-0.048		0.000	-0.002	0.014	0.910	1.000	-0.018	0.001	0.984	0.925
Other Fuel Taxes	**1.000**	**-0.185**	**-0.066**				**0.025**	**0.775**	**1.000**	**-0.033**	**0.004**	**0.970**	**0.798**
Jet Fuel	1.000	-0.310	-0.003				0.036	0.722	1.000	-0.048	0.006	0.958	0.754
Diesel	1.000	-0.059	-0.129				0.015	0.827	1.000	-0.019	0.001	0.982	0.842
Heavy Fuel *	1.000	-0.075	-0.001				0.007	0.931	1.000	-0.002	0.001	1.000	0.931
Crude (WPT) *	1.000	0.000	0.000				0.000	1.000	1.000	-0.002	0.000	0.998	1.002
Crude (State) *	1.000	-0.075	0.000				0.000	0.925	1.000	-0.364	0.001	0.637	1.451
E85 *	1.000	0.562	0.009				0.411	1.982	1.000	-0.361	0.011	0.650	3.051
Other Revenue Raisers	**0.979**	**-0.150**	**-0.014**	**0.012**			**-0.108**	**0.719**	**1.000**	**0.109**	**0.003**	**1.112**	**0.647**
CPP (AJ)	1.000	-0.107	-0.030	0.000			-0.323	0.540	1.000	0.176	0.002	1.178	0.459
CARE	0.936	-0.303	0.000	0.036			0.117	0.785	1.000	0.086	0.006	1.092	0.719
CPP (PJ)	1.000	-0.039	-0.011	0.000			-0.119	0.831	1.000	0.065	0.001	1.065	0.780
Cap and Trade													
RGGI	1.000	-0.657	-0.989					-0.646	1.000	-0.050	0.013	0.963	-0.671
CA CT	1.000	-0.061	-0.002					0.937	1.000	-0.006	0.001	0.996	0.941
ETS (BA) *	1.000	-9.192	0.000					-8.192	1.000	-0.900	0.180	0.280	-29.287
ETS (CMMW) *	1.000	-1.279	0.000					-0.279	1.000	-0.125	0.025	0.900	-0.310
Panel D. International													
Cookstoves													
Cookstove (Kenya)	7.656	43.161	0.000					50.817	1.000	0.000	-0.843	0.157	323.453
Cookstove (India)	0.545	-2.956	0.000					-2.410	1.000	0.000	0.058	1.058	-2.279
Deforestation													
REDD+ (SL)	0.000	35.840	0.000					35.840	1.000	0.000	-0.700	0.300	119.438
Deforest (Uganda)	0.421	4.538	0.000					4.959	1.000	0.000	-0.089	0.911	5.441
REDD+	0.965	2.951	0.000					3.916	1.000	0.000	-0.058	0.942	4.156
Deforest (Mexico) *	0.944	0.740	0.000					1.684	1.000	0.000	-0.014	0.986	1.709
Rice Burning													
India PES (Upfront)	0.972	10.642	0.000					11.614	1.000	0.000	-0.208	0.792	14.661
India PES (Standard)	0.915	8.128	0.000					9.043	1.000	0.000	-0.159	0.841	10.749
Wind Offset													
Offset (India)	1.000	9.355	0.000	-1.861				8.495	1.000	0.258	-0.146	1.112	7.641
International Rebates													
Fridge (Mexico)	0.750	0.125	0.000	-0.024				0.850	1.000	0.000	-0.002	0.998	0.852
AC (Mexico)	0.750	-0.094	0.000	0.018				0.675	1.000	0.000	0.001	1.001	0.674
WAP (Mexico)	0.500	-0.096	0.000	0.019				0.422	1.000	0.000	0.002	1.002	0.422

International Nudges													
Nudge (Qatar) *	0.000	7.201	0.000	-1.410				5.791	1.000	0.000	-0.113	0.887	6.529
Nudge (Germany) *	0.000	0.401	0.000	-0.079				0.323	1.000	0.000	-0.006	0.994	0.325

Notes: This table presents the WTP and cost components for each policy in our sample using the baseline specification. Each component is normalized per dollar of mechanical spending on the policy. The first column reports the size of the transfer. The next three columns report the environmental externality including local externalities, global greenhouse gas externalities, and rebound effects (both global and local). The next two columns report learning by doing components for both the environmental benefits and future price reductions. The next column reports impact on profits of oil/gas and utility sectors. The cost components report the mechanical cost, followed by the fiscal externalities (state and federal tax and subsidy impacts), and the climate fiscal externality from the impact of changes in climate on future GDP and thus future tax revenue. We report estimates for each policy in our sample along with category averages for each type of policy. We denote policies excluded from our primary sample by "*", and these policies are not included in our category average measures. All numbers are calculated using our baseline path for the social cost of carbon ($193 in 2020) and a 2% discount rate.

Table 2: MVPF Versus Cost Per Ton

		Cost Per Ton		
Panel A. With Learning by Doing	**MVPF**	**Resource**	**Government**	**Social**
Subsidies				
Wind Production Credits	5.870	-103	46	-32
Residential Solar	3.862	-77	90	-67
Electric Vehicles	1.445	-458	1,356	-415
Appliance Rebates	1.164	-2	474	111
Vehicle Retirement	1.047	1,007	876	148
Hybrid Vehicles	1.012	577	5,940	-38
Weatherization	0.978	194	779	207
Nudges and Marketing				
Opower Elec. (166 RCTs)	2.548	-41	77	70
Revenue Raisers				
Gasoline Taxes	0.671	-104	-770	-64
Panel B. Without Learning by Doing				
Subsidies				
Wind Production Credits	3.851	-42	69	-8
Residential Solar	1.446	4	237	83
Electric Vehicles	0.961	963	2,422	283
Appliance Rebates	1.164	-2	474	111
Vehicle Retirement	1.047	1,007	876	148
Hybrid Vehicles	0.998	659	6,087	43
Weatherization	0.978	194	779	207
Nudges and Marketing				
Opower Elec. (166 RCTs)	2.548	-41	77	70
Revenue Raisers				
Gasoline Taxes	0.673	-104	-768	-62

Notes: This table presents estimates of the MVPF and cost-per-ton measures using our three definitions: resource cost per ton, government cost per ton and social cost per ton. See text for precise definitions of each measure. We present estimates here for each policy category average; the Appendix provides estimates for each policy. All numbers are calculated using our baseline path for the social cost of carbon ($193 in 2020) and a 2% discount rate.