

Are Climate Policies Marginal? A Welfare Evaluation of Environmental Reforms^{*}

Jean-François Fournel[†]

Harvard University

Sébastien Montpetit[‡]

University of Toronto

October 16, 2026

Latest version available [here](#).

Abstract

Which climate investments should policymakers prioritize, and how much should they spend? We develop a framework combining the marginal value of public funds (MVPF) with structural models of demand for green technologies to evaluate the welfare effects of environmental reforms. Our approach traces the MVPF across the entire policy spectrum using counterfactual simulations, thereby relaxing the constant elasticity and “small change” assumptions of sufficient-statistics methods. We apply our method to electric vehicles and public charging in Quebec, residential solar in Connecticut, and a fuel-economy label experiment in the United States. At the margin, public charging and the fuel-economy label deliver higher returns than purchase subsidies or taxes. The returns to subsidies fall with spending, whereas the returns to public charging and the fuel-economy label do not. For programs with increasing returns, the optimal reform is all-or-nothing, and the implementation depends on the return of the whole program rather than on the return of the marginal dollar spent.

Keywords: climate policy, sufficient statistics, structural models, non-marginal reforms, environmental externalities.

JEL codes: H23, Q58, D61.

^{*}We thank Benjamin Goldman, James K. Hammitt, Claire Leroy, David Paton, Jim Stock, and workshop participants at Harvard, Warwick, the QMUL Economics and Finance workshop, and the NIE Symposium for useful discussions and comments on this project. All remaining errors are our own.

[†]The Salata Institute for Climate and Sustainability, Harvard University. Email: jfournel@fas.harvard.edu.

[‡]Munk School of Global Affairs & Public Policy, University of Toronto. Email: sebastien.montpetit@utoronto.ca.

1 Introduction

It is well established that addressing climate change requires ambitious actions and massive investments to decarbonize the economy (Nordhaus, 2019). Such large-scale interventions have major implications for public finances. Given that public funds are limited, which investments should policymakers prioritize? How much should they spend on green initiatives? Gathering more systematic empirical evidence on the social returns of large environmental reforms is a crucial step in addressing the climate crisis.

Three methodological challenges hinder our ability to rigorously evaluate and compare environmental policies. First, existing evidence on the social returns of such policies lacks a unified metric that enables consistent comparisons across interventions (see Hahn et al., 2026, who address this issue for marginal reforms). Second, the empirical tools for ex-post welfare analysis of non-marginal reforms, such as large subsidies on green technologies, remain underdeveloped. Third, the optimal scale of an intervention and the optimal mix of climate policy instruments may lie far from the status quo, where causal estimates must be extrapolated, and may involve programs that were never implemented, for which no causal estimates exist.

In this paper, we develop a framework that combines the marginal value of public funds (henceforth MVPF) (Hendren, 2016; Hendren and Sprung-Keyser, 2020) with structural estimation to address these challenges. The MVPF is defined as the ratio of the willingness-to-pay for a small change in a policy, around the status quo, to the net cost to the government. We show how to map these two components to the deep parameters of any fully specified structural model.¹ In doing so, we provide a representation of the MVPF for both marginal and non-marginal reforms, as well as conditions for choosing an optimal mix of several policy instruments, given a fixed budget constraint.

We define the MVPF of a non-marginal reform as the ratio of the willingness-to-pay for the whole program to the total cost of the program. We show that this ratio is a weighted average of the instrument-specific MVPFs along the path of the reform. Under diminishing returns, the optimal policy equalizes the MVPF of the last dollar invested across policy instruments, and the government should expand spending until this common MVPF falls to the cost of raising one more dollar of public funds.

Our approach produces valid MVPF measurements that are fully consistent with, and comparable to, those obtained from causal estimates and the sufficient-statistics approach of

¹For the particular case of random coefficient discrete choice models (Berry et al., 1995), we derive a closed-form solution to this mapping, which can be easily implemented by practitioners.

Chetty (2009) and Hendren (2016). In this sense, we view our analysis of marginal reforms as complementary to those of Hendren and Sprung-Keyser (2020), Hahn et al. (2026), and other reduced-form studies (see Hendren et al., 2026, for an exhaustive list of MVPF estimates in the literature). We show that imposing structural assumptions on model primitives has several advantages, and broadens the set of policies which can be studied within the MVPF framework. First, our structural approach delivers the MVPF of each policy instrument separately, at any point of the policy space, including instruments that were never implemented. Second, the structural model allows us to evaluate changes in several policy instruments at once, and to identify the policy mix and the level of spending that maximize welfare. Most importantly, our structural approach extends the MVPF framework to non-marginal reforms without relying on the envelope theorem logic that is ubiquitous in the public finance literature.

We apply our methodology to three empirical studies of environmental policies in North America, which taken together cover price instruments (e.g., subsidies and taxes), public-good provision (e.g., public charging for electric vehicles), and informational nudges (e.g., green labels). The price instruments we study are subsidies to green products, including rebates on electric vehicles in Quebec (Fournel, 2026) and financial incentives on residential solar panels in Connecticut (Gillingham and Tsvetanov, 2019). In Quebec, local governments provide public charging stations to satisfy the demand for charging services from electric vehicle owners. Our analysis treats these public stations as a public good that complements the rebates on electric vehicles. The nudge we study is a fuel-economy label presented to potential car buyers in the United States, evaluated in the online experiment of Allcott et al. (2025). Consumers in the experiment underweight fuel costs, and the label induces some of them to switch to the fuel-efficient car, correcting their bias (internality) and reducing carbon emissions (externality). While the applications based on Fournel (2026) and Gillingham and Tsvetanov (2019) use their estimated structural models, our implementation of Allcott et al. (2025) estimates a new structural model of car choice, based on their theoretical framework.

In our applications, the returns to subsidies diminish with government spending. Intuitively, this occurs because each consumer that adopts the green product receives the full subsidy, so the fiscal cost paid to marginal consumers increases over the policy space, while the marginal reduction in environmental damages remains relatively constant. We show that this is sufficient to allow for an interior solution, although our empirical settings exhibit cases where one or more policy instruments are not used at the optimum. In contrast, the returns to public goods and nudges do not necessarily diminish with the size of the program.

The cost of an informational nudge per consumer is typically flat (e.g., printing a label), and a public good is shared by many consumers, so the cost per marginal consumer does not rise over the policy space as it does for subsidies. Without diminishing returns, the optimal policy is all-or-nothing (i.e., no interior solution exists). In our applications, this happens both for public charging and the fuel-economy label.²

Our structural approach allows us to trace the MVPF over the policy space, which is informative about the returns to policy instruments away from the status quo, including instruments that are not yet in place. For example, in our application to electric vehicles, the first dollar spent on battery electric vehicles returns nearly twice as much as the last dollar, and the return falls below one just above the observed rebate (\$13,000 per vehicle). Among the instruments that are not in place in 2020, a rebate on hybrid vehicles (available in Quebec until 2018) would return about 50% more per dollar spent than the rebates on electric vehicles, at the observed level.³ Conversely, a tax on gas guzzlers would raise its first dollars at a social cost slightly above one dollar per dollar raised, and the social cost increases with the level of the tax.

The marginal conditions of our framework give both the optimal mix of instruments and the optimal level of spending on a reform. For residential solar panels in Connecticut, the optimal policy is the Pigouvian subsidy: a state rebate that lowers the price paid by consumers by the environmental damage avoided per watt installed. The federal tax credit is not used, because the credit is proportional to the price of a system, while the damage avoided is proportional to the size of the system.⁴ The optimal level of spending lies between the 2014 and the 2021 budgets: the 2014 program spent too much, and the 2021 reform cut spending too far. For electric vehicles in Quebec, the rebates in place are close to their optimal levels. Most of the gain from the best policy comes from introducing a sizeable rebate on hybrid vehicles, funded by slightly lower rebates on electric vehicles, while the tax on gas guzzlers remains unused. In Quebec, the best policy raises total surplus by about 3% of the observed spending.

²In our application to public charging, increasing returns occur in some large counties partly because the demand model of [Fournel \(2026\)](#) does not allow for congestion effects. As a result, the policymaker finds it optimal to invest all government funds in a few counties, and divest everywhere else. Introducing congestion would likely restore diminishing returns everywhere.

³In South Korea, [Hong et al. \(2026\)](#) find larger price elasticities for hybrids than for battery electric vehicles. In Quebec, [Fournel \(2026\)](#), Table 6, estimates a smaller market-level elasticity for hybrids than for battery electric vehicles (-2.75 against -3.15), and the first dollar spent on battery electric vehicles returns more than the first dollar spent on hybrid rebates (1.83 against 1.48). The advantage of the hybrid rebate at the observed policy comes from its starting point at zero.

⁴The ranking of the two subsidies depends on the budget. The federal credit is the only instrument used for budgets below \$11.5 million, and the state rebate the only instrument used above \$19.2 million, see [Figure 11](#).

Most climate policies are non-marginal in nature because their intended goal is often precisely to induce large behavioral responses by consumers. The decision to implement such a policy rests on whether the non-marginal MVPF of the full program exceeds the cost of raising public funds for that program. In our empirical analysis of subsidies, this condition holds whenever the marginal conditions call for spending, since the first dollars of a subsidy return more than the last dollar.⁵ In our applications to public charging and the fuel-economy label, the non-marginal MVPF of the full program alone decides whether to implement the policy.

The fuel-economy label of [Allcott et al. \(2025\)](#) illustrates this case. The solution to the welfare problem is a trigger policy: the government should either implement the label at full coverage or abandon the label. In our application, the implementation cost of the label is not observed, and the label is fully implemented if it costs less than \$18.7 per consumer and vehicle-year. [Allcott et al. \(2025\)](#) propose a corrective tax on fuel-inefficient cars, which accounts for the carbon externality and for consumers' residual misperception of fuel costs after seeing the label. With this tax in place, the threshold falls to \$8.4, since each consumer who switches to the fuel-efficient car stops paying the tax. The decision to implement the label therefore depends on the other policies in place.

Related literature. This paper contributes to three strands of the economic literature. First, we add to the extensive literature assessing the benefits and costs of environmental reforms. Previous work has focused on estimating the marginal abatement cost or the environmental impacts of specific climate policies (see [Gillingham and Stock \(2018\)](#) and [Hoppe et al. \(2023\)](#) for reviews).⁶ Other studies compare the returns of several programs. For instance, [Gaarder et al. \(2026\)](#) study the misallocation of public funds across green investment programs in Norway and find that an optimal reallocation of funds across programs would nearly double emission reductions. We contribute by showing how the public finance toolkit can be combined with structural models to perform welfare analysis. Specifically, structural assumptions on preferences and technology allow us to extend the MVPF framework beyond the sufficient-statistics approach of [Hahn et al. \(2026\)](#), to non-marginal reforms and to policies not yet implemented.

⁵For electric vehicles, the best policy has a marginal MVPF of one and a non-marginal MVPF of 1.23, see [Table 3](#).

⁶See also [Timilsina \(2022\)](#) for a review of the literature on carbon taxes. On specific climate policies, studies estimating costs and benefits of government programs include [Allcott \(2013\)](#), [Jacobsen and Kotchen \(2013\)](#), [Holland et al. \(2016\)](#), [Allcott and Kessler \(2019\)](#), [Hahn and Metcalfe \(2021\)](#), [Jacobsen et al. \(2023\)](#), [List et al. \(2023\)](#), [Prest and Stock \(2023\)](#), [Allcott and Greenstone \(2024\)](#), [Allcott et al. \(2024\)](#), and [Kotchen \(2025\)](#).

The closest paper to ours is [Hahn et al. \(2026\)](#), who estimate the MVPF of 96 environmental reforms from existing studies, using a sufficient-statistics approach.⁷ Relative to their work, our main contribution is to measure the MVPF of non-marginal reforms, such as the introduction of a new program. Our empirical analysis shows that this distinction matters for policies with increasing returns, such as the fuel-economy label of [Allcott et al. \(2025\)](#). In these cases, the decision to implement the policy rests on the MVPF of the full program rather than on the MVPF of a marginal change around the status quo. Our framework also evaluates policies that were never implemented, such as a tax on gas guzzlers in Quebec, and solves for the combination of policy instruments and the level of spending that maximize welfare.

Beyond environmental policy, our work contributes to the public finance literature on empirical welfare analysis. Seminal papers show that, under standard assumptions, reduced-form causal estimates can be used as a sufficient statistic for beneficiaries’ willingness to pay for a small policy perturbation around the status quo ([Chetty, 2009](#); [Hendren, 2016](#)). Studying non-marginal reforms using this approach is typically limited to first-order approximations of welfare effects ([Finkelstein and Hendren, 2020](#); [Kleven, 2021](#)). [Kang and Vasserman \(2025\)](#) recently proposed welfare bounds for non-marginal changes in fiscal policy, as a robustness measure for the standard assumption of linearity of demand. Our paper, instead, identifies the MVPF of large changes in government policy by tracing the evolution of the policy elasticities (and the corresponding MVPF) across the full policy space. Our framework also provides a benchmark for these approximations: in all our applications, the trapezoidal approximation of [Kleven \(2021\)](#), with the local estimates extrapolated at a constant elasticity as in [Hahn et al. \(2026\)](#), deviates from the non-marginal MVPF by 2.5% on average. For marginal reforms, our estimates remain directly comparable to reduced-form estimates (e.g., [Hendren and Sprung-Keyser, 2020](#); [Hahn et al., 2026](#); [Boning et al., 2025](#); [Hendren et al., 2026](#)). Structural models necessarily impose additional assumptions on the economic environment, but making these assumptions explicit reveals some of the implicit assumptions behind sufficient-statistics methods.

Finally, our paper relates to the literature in empirical industrial organization that uses structural models of demand to evaluate environmental policies. Most of this work studies the car market (see [Van Biesebroeck and Verboven, 2026](#), for a review), from subsidies for

⁷[Hahn et al. \(2026\)](#) also incorporate learning-by-doing into the MVPF of subsidies for green technologies. Our framework accommodates these dynamic effects whenever the structural model includes them (e.g., through marginal costs that fall with cumulative production according to Wright’s rule ([Wright, 1936](#))), but the models of our applications do not, so our estimates for electric vehicles and solar panels exclude this channel.

electric vehicles and charging infrastructure (Li et al., 2017; Springel, 2021; Guo and Xiao, 2023; Dorsey et al., 2025; Remmy, 2026; Hong et al., 2026; Fournel, 2026), to taxes, feebates, and scrappage programs (Bento et al., 2009; Li et al., 2013; Durrmeyer, 2022; Gillingham et al., 2022), emission standards and mandates (Durrmeyer and Samano, 2018; Reynaert, 2021; Armitage and Pinter, 2026; Alé-Chilet et al., 2026), and urban driving restrictions and road tolls (Durrmeyer and Martínez, 2026), while a smaller body of work studies subsidies and electricity tariffs for residential solar panels (Gillingham and Tsvetanov, 2019; De Groote and Verboven, 2019; Feger et al., 2022). Several recent papers in this literature already estimate MVPFs of non-marginal reforms (Allcott and Greenstone, 2024; Allcott et al., 2024; Gerarden et al., 2025). We formalize the mapping between the MVPF and random coefficient discrete choice models of demand (Berry et al., 1995), for marginal and non-marginal reforms. Compared to this body of work, we contribute by proposing a method for comparative analysis across various markets using a common welfare criterion. In this way, our work creates another bridge between reduced-form (causal) and structural work beyond those identified in Todd and Wolpin (2023).⁸

The rest of the paper is structured as follows. Section 2 details our conceptual framework and our structural approach to welfare analysis. Section 3 presents an application of the theory to the special case of discrete choice models. Section 4 presents our empirical applications and detailed results. Finally, Section 5 concludes.

2 Conceptual framework

In this section, we describe the conceptual framework used for our welfare analysis. Our welfare criterion is the marginal value of public funds (MVPF), which is defined as the ratio between the policy’s benefit to its beneficiaries (measured as their willingness to pay for that policy) and the policy’s net cost to the government (Hendren, 2016). That is,

$$\text{MVPF} = \frac{\text{Beneficiaries' willingness to pay}}{\text{Net cost to government}}. \quad (1)$$

The net cost to the government is given by the sum of the direct government expenditure on the policy and fiscal externalities (i.e., indirect impacts on the government budget from

⁸Scholars such as Heckman (2010), Todd and Wolpin (2023), and Buera et al. (2023) discuss the merits of combining “the best of both worlds” in empirical research (i.e. randomized or quasi-natural experiments and structural models of behavior). On the one hand, structural models can help interpret the mechanisms through which a given policy change impacts relevant outcomes and allow for counterfactual policy experiments that deviate from the implemented policy. On the other hand, ex-post policy evaluation can be used to identify and discipline behavioral models.

changes in behavior).⁹ In what follows, we show how to adapt the sufficient-statistics toolkit to discrete choice models of demand for green technologies.

2.1 The model

We begin by deriving the MVPF in a general environment that imposes minimal structure on preferences, technology, or market conduct. The purpose of this section is twofold. First, it defines the welfare objects, namely, the willingness-to-pay and the net government cost, that any empirical implementation must measure. Second, it identifies precisely which features of the economic environment those objects depend on. The structural exercise in the next sections then consists of supplying these features from an estimated structural model rather than from (reduced-form) causal estimates. The exposition draws on the theoretical framework of [Hahn et al. \(2026\)](#).

A crucial distinction between our structural approach and the sufficient-statistics framework ([Hendren and Sprung-Keyser, 2020](#); [Hahn et al., 2026](#)) is that it is not limited to marginal changes in the economic environment. The following sections detail how the methodology can solve some of the limitations of the sufficient-statistics toolkit, at the cost of additional structure on the behavior of agents.

In what follows, we consider a government policy $\theta = (\tau, g, n)$ that includes a vector of price instruments τ (e.g., subsidies for electric vehicles, where negative values are taxes), the provision of public goods g (e.g., charging stations for electric vehicles), and nudges n (e.g., providing information about the potential cost savings of driving an electric vehicle). We assume throughout that price instruments affect consumers only through the expansion or contraction of the budget constraint, that public goods affect both the consumers' utility directly and their budget constraint (through the equilibrium price mechanism), and that nudges affect consumers' experienced utility (e.g., warm glow effect), consumers' budget constraints (e.g., equilibrium price response), and consumers beliefs about the targeted green technology (wedge between decision and experienced utility).

The first part of this section focuses on the price instruments τ and the provision of public goods g . We return to the role of nudges n at the end of this section, for expositional simplicity.¹⁰

⁹An important point to note here is that, by construction, we should not inflate costs to account for the deadweight loss of raising public expenditures in the MVPF framework (see [Hendren and Sprung-Keyser, 2022](#), for a discussion).

¹⁰This classification covers most instruments of environmental policy. For instance, mandates and performance standards can usually be converted to producer taxes and subsidies whenever compliance credits are tradeable, as in the US emission markets, and are therefore covered by the price instruments. [Jacobsen et al.](#)

Setup. Consider an economy populated by N individuals, indexed by $i = 1, \dots, N$. Each individual chooses an action vector x_i satisfying a budget set $B_i(\theta)$, where $\theta = (\tau, g)$ is a vector collecting policy instruments under the government's control. Without loss of generality, we assume throughout that x_i is a k -dimensional consumption bundle, τ is a k -dimensional vector of subsidies (negative values are taxes), g is a vector of government-provided public goods, and m_i is consumer i 's income.

Consumption goods are produced by F firms indexed by $f = 1, \dots, F$. Let $\Pi(\mathbf{x}) = \sum_f \pi_f(\mathbf{x})$ be total profits in the economy. We assume that consumer i owns a share s_i of firms' profits. Under corporate tax rate τ_c , her consumption bundle thus satisfies the following budget constraint,

$$B_i(\theta) \equiv p(\theta) \cdot x_i \leq m_i + (1 - \tau_c) s_i \Pi(\mathbf{x}),$$

where $\mathbf{x} = (x_1, \dots, x_N)$ contains the collective decisions of all consumers in the economy.

We assume throughout that consumer i 's utility is additively separable in the externality generated by the collective decisions of all agents, denoted by $d_i(\mathbf{x})$. That is, the utility of consumer i is

$$V_i(x_i, \theta) = u_i(x_i, g) - d_i(\mathbf{x}). \quad (2)$$

Notice that additive separability is useful in this context to distinguish decision utility (i.e., what consumers internalize when choosing x_i) from welfare utility (which includes the externalities generated by x_i). For instance, with additively separable preferences and atomistic consumers that take the aggregate action profile $\mathbf{x}$ as given, consumers choose an optimal consumption bundle satisfying,

$$x_i^*(\theta) = \arg \max_{x \in \mathbb{R}^k} u_i(x, g) + \lambda_i (m_i + (1 - \tau_c) s_i \Pi(\mathbf{x}) - p(\theta) \cdot x), \quad (3)$$

where the Lagrange multiplier λ_i is the marginal utility of income. Note that the optimal consumption bundle does not depend on environmental damages, $d_i(\mathbf{x})$.¹¹ In contrast,

(2023) provides a detailed discussion on the equivalence between vehicle emission standards and taxes. [Hahn et al. \(2026\)](#) also evaluate fuel-economy and renewable portfolio standards through their implicit tax and subsidy structure. When compliance credits cannot be traded, a standard is a restriction on the consumer's choice set and the welfare analysis proceeds with the same envelope logic applied to the constrained problem. Similarly, price regulations such as price caps affect consumers through equilibrium prices and firms through their profits, without any direct spending by the government.

¹¹With N atomistic agents, the effect of i 's own action on the externality she experiences, $\partial d_i / \partial x_i$, is of order $1/N$ and hence negligible in the choice problem for large N . However, the sum of these effects across all affected individuals is first-order for welfare. This microfound the assumption, standard in the literature, that agents do not internalize the environmental damage they generate (see [Hahn et al., 2026](#)).

welfare depends on both parts of the utility profile, as it depends on the choices of all individuals.

Social welfare and total cost. Let ψ_i denote a set of Pareto welfare weights. Social welfare is then the weighted sum of agents' utilities, that is,

$$\mathcal{W}(\theta) = \sum_i \psi_i V_i(x_i^*(\theta), \theta), \quad (4)$$

$$= \underbrace{\sum_i \psi_i u_i(x_i^*(\theta), g)}_{\text{Consumers' indirect utility}} - \underbrace{\sum_i \psi_i d_i(\mathbf{x}^*(\theta))}_{\text{Externality}}. \quad (5)$$

The total cost to the government for implementing policy $\theta = (\tau, g)$ includes the total fiscal cost of the policy, minus the taxable revenues collected from firms affected by those policies, which constitute one of the fiscal externalities defined above. That is,

$$\mathcal{C}(\theta) = \underbrace{\sum_i \tau \cdot x_i^*(\theta)}_{\text{Cost of policy}} + c_g g - \underbrace{\tau_c [\Pi(\mathbf{x}^*(\theta)) - \Pi(\mathbf{x}^*(0))]}_{\text{Tax revenues}}. \quad (6)$$

where c_g is the cost of providing one additional unit of the non-price policy g and $\mathbf{x}^*(0)$ is the collective actions of all agents absent government policies.

Willingness-to-pay and the marginal value of public funds. We consider a small policy change around a currently implemented policy θ_0 , in the direction θ_j . In what follows, we assume the social welfare weights to be equal to the inverse of the marginal utility of income, $\psi_i = 1/\lambda_i$, such that the total social welfare is the sum of agents' valuation for the policy change, in dollars.¹² In this context, the (aggregate) willingness-to-pay for an infinitesimal change to the policy along θ_j coincides with the change in aggregate social

¹²This normalization is standard in the MVPF literature. Formally, under arbitrary weights ψ_i , the welfare effect of the reform is $\bar{\eta}_j \times \sum_i WTP_{ji}$, where $\bar{\eta}_j = \sum_i \eta_i WTP_{ji} / \sum_i WTP_{ji}$ is the average social marginal utility of income ($\eta_i \equiv \psi_i \lambda_i$) among the policy's beneficiaries, weighted by their share of the aggregate willingness-to-pay. Setting $\psi_i = 1/\lambda_i$ implies $\bar{\eta}_j = 1$, as in [Hendren and Sprung-Keyser \(2020\)](#), so that our MVPF measurements are directly comparable to theirs. Unless noted otherwise, we maintain this assumption throughout the paper; Appendix C.1 shows how to recover welfare statements under arbitrary weights.

welfare, that is,

$$WTP_j = \underbrace{(1 - \tau_c) \sum_i s_i \sum_\ell \frac{\partial \Pi}{\partial x_\ell^*} \frac{dx_\ell^*}{d\theta_j}}_{\text{Shareholders' net benefits}} - \underbrace{\sum_i x_i^* \frac{\partial p}{\partial \theta_j}}_{\text{Savings by inframarginals}} + \underbrace{\sum_i \frac{1}{\lambda_i} \frac{\partial u_i}{\partial g} \frac{\partial g}{\partial \theta_j}}_{\text{Benefits from public goods}} - \underbrace{\sum_i \frac{1}{\lambda_i} \sum_\ell \frac{\partial d_i}{\partial x_\ell^*} \frac{dx_\ell^*}{d\theta_j}}_{\text{Externality}}, \quad (7)$$

where, by the envelope theorem, consumers' utility gains,

$$\sum_i \frac{1}{\lambda_i} \frac{\partial u_i}{\partial x_i^*} \frac{dx_i^*}{d\theta_j} = (1 - \tau_c) \sum_i s_i \sum_\ell \frac{\partial \Pi}{\partial x_\ell^*} \frac{dx_\ell^*}{d\theta_j} - \sum_i x_i^* \frac{\partial p}{\partial \theta_j}$$

boil down to the impact of the policy on inframarginal consumers' budget constraint (with no behavioral response from marginal consumers). The economic interpretation of the envelope conditions is that, while inframarginal consumers benefit from the small increase in subsidies through an income effect, marginal consumers (e.g., compliers) do not, since they are indifferent between purchasing the green good or not at θ_0 .

The impact of a small increase in the policy intensity on government finances is

$$\frac{d\mathcal{C}}{d\theta_j} = \underbrace{\sum_i x_i^* \frac{\partial \tau}{\partial \theta_j}}_{\text{Transfer to inframarginals}} + \underbrace{\sum_i \tau \frac{dx_i^*}{d\theta_j}}_{\text{Additional subsidy take-up}} + \underbrace{c_g \frac{\partial g}{\partial \theta_j}}_{\text{Spending on public goods}} - \underbrace{\tau_c \sum_\ell \frac{\partial \Pi}{\partial x_\ell^*} \frac{dx_\ell^*}{d\theta_j}}_{\text{Additional corporate tax revenues}}. \quad (8)$$

The first term represents the change in direct cost of the policy (additional spending on inframarginals). The second term is the additional take-up of subsidies, that is, a fiscal externality due to the behavioral change from compliers. The third term is the additional spending on non-price interventions and the last term is the change in revenues from taxation, a fiscal externality due to firms' profit gains.¹³

The MVPF of a given policy θ_0 along dimension θ_j is then the ratio of the willingness-to-pay for the incremental change in the policy to the additional costs (or the “bang-for-buck” of the policy), evaluated at the baseline policy θ_0 , that is,

$$MVPF_j(\theta_0) = \left. \frac{WTP_j}{d\mathcal{C}/d\theta_j} \right|_{\theta=\theta_0}. \quad (9)$$

¹³It is straightforward to expand the framework to allow for increasing or decreasing marginal cost of public good provision. In this case, $c_g = c_g(g)$ and equation (8) would contain an extra term, $c'_g g \neq 0$, capturing these increasing or decreasing returns.

Nudges and internalities. We return to the role of nudges in the consumer problem. We redefine the vector of policy instruments under the government’s control as $\theta = (\tau, g, n)$. The derivation above assumes that agents’ decision utility, denoted $\tilde{u}_i(x, g, n)$ in what follows, coincides with their experienced utility $u_i(x, g, n)$. Nudges operate precisely where this assumption may fail, if, for example, agents have incorrect perceptions about the environmental performance of a green product (List et al., 2023) or of their own actions (e.g. Semken, 2025). In this context, the experienced utility, which matters for welfare, is realized only after consumers purchase a given consumption bundle x_i under imperfect beliefs. In other words, the consumers solve

$$x_i^*(\theta) = \arg \max_{x \in \mathbb{R}^k} \tilde{u}_i(x, g, n) + \lambda_i (m_i + (1 - \tau_c) s_i \Pi(\mathbf{x}) - p(\theta) \cdot x), \quad (10)$$

so the first-order conditions set $\partial \tilde{u}_i / \partial x = \lambda_i p(\theta)$, while welfare utility remains $V_i = u_i + d_i$.

This implies that the willingness-to-pay includes a direct experience effect from the nudge and an internality correction effect, that is,

$$\begin{aligned} WTP_j = & \underbrace{(1 - \tau_c) \sum_i s_i \sum_\ell \frac{\partial \Pi}{\partial x_\ell^*} \frac{dx_\ell^*}{d\theta_j}}_{\text{Shareholders' net benefits}} - \underbrace{\sum_i x_i^* \frac{\partial p}{\partial \theta_j}}_{\text{Savings by inframarginals}} \\ & + \underbrace{\sum_i \frac{1}{\lambda_i} \frac{\partial u_i}{\partial g} \frac{\partial g}{\partial \theta_j}}_{\text{Benefits from public goods}} + \underbrace{\sum_i \frac{1}{\lambda_i} \frac{\partial u_i}{\partial n} \frac{\partial n}{\partial \theta_j}}_{\text{Direct experience effect}} \\ & + \underbrace{\sum_i \frac{1}{\lambda_i} \left[\frac{\partial u_i}{\partial x_i^*} - \frac{\partial \tilde{u}_i}{\partial x_i^*} \right] \frac{dx_i^*}{d\theta_j}}_{\text{Internality correction}} - \underbrace{\sum_i \frac{1}{\lambda_i} \sum_\ell \frac{\partial d_i}{\partial x_\ell^*} \frac{dx_\ell^*}{d\theta_j}}_{\text{Externality}}. \end{aligned} \quad (11)$$

The direct experience effect captures the utility consequences of receiving the nudge itself, such as warm glow, guilt, or a hassle cost,¹⁴ while the internality correction term captures biased beliefs about the value of the green technology to consumers.

Let c_n be the unit cost of the nudge intervention such as producing and mailing a report on an energy-efficiency program, or an advertisement campaign. The marginal impact of the

¹⁴Allcott and Kessler (2019) show that these effects are quantitatively important in the case of home-energy reports: some recipients value the reports while others would pay to stop receiving them.

policy on government finances is then

$$\frac{d\mathcal{C}}{d\theta_j} = \underbrace{\sum_i x_i^* \frac{\partial \tau}{\partial \theta_j}}_{\text{Transfer to inframarginals}} + \underbrace{\sum_i \tau \frac{dx_i^*}{d\theta_j}}_{\text{Additional subsidy take-up}} + \underbrace{c_g \frac{\partial g}{\partial \theta_j} + c_n \frac{\partial n}{\partial \theta_j}}_{\text{Spending on non-price instruments}} - \underbrace{\tau_c \sum_\ell \frac{\partial \Pi}{\partial x_\ell^*} \frac{dx_\ell^*}{d\theta_j}}_{\text{Additional corporate tax revenues}}. \quad (12)$$

2.2 Estimating MVPF in structural models

We now explain how a fully structural approach can deliver the same MVPF object as the sufficient-statistics framework (which instead uses causal effects) under different sets of assumptions. We first briefly present how the sufficient-statistics approach is used to estimate welfare effects of marginal changes in policy. Then, we discuss how the structural approach can similarly recover the willingness-to-pay and the net costs of marginal investments. Last, because causal effects are local to an observed policy, we detail how a structural model can expand the set of policies that can be studied under the MVPF framework to include (i) non-marginal programs, (ii) policy bundles, that is, simultaneous changes in more than one policy instrument, and (iii) policies that have not yet been implemented (i.e., ex ante policy evaluation).

Marginal interventions. The sufficient-statistics approach to welfare analysis shows how to map willingness-to-pay and costs to a relatively small set of causal effects (Chetty, 2009; Hendren, 2016; Finkelstein and Hendren, 2020). In our setup, equations (7) and (8) make explicit that a relatively small set of causal effects and externally calibrated parameters (e.g. tax rates, the social cost of pollution, and a firm markup) are sufficient to estimate the MVPF of a marginal change in subsidy τ without additional structure on utility. Indeed, the willingness-to-pay and costs are fully captured by the causal impact of the subsidy on market demand $dx/d\tau$ (or equivalently, by its elasticity to price, $\varepsilon = dx/d\tau \cdot p/x$).¹⁵

For public goods and nudges, a complication arises because there may be non-pecuniary benefits (e.g., from $\partial u/\partial g$ and $\partial u/\partial n$) that have first-order impacts on utility (Hendren, 2016). These can be substantial in some settings.¹⁶ To estimate these gains using causal effects, the researcher must estimate the dollar valuation of a unit of the public good (or

¹⁵To see this, setting $\partial g/\partial \theta_j = 0$ in equation (7) eliminates the public-goods term and leaves only the demand response $dx_\ell^*/d\theta_j$ and externally calibrated parameters (externalities and markups), confirming that the willingness-to-pay for a price instrument requires no additional structure on utility beyond the causal effect itself.

¹⁶For instance, Montpetit et al. (2026) study a universal childcare program using a structural model of childcare demand and document that non-pecuniary benefits for mothers can be even larger than earnings gains.

of a nudge) to beneficiaries. [Hendren \(2016\)](#) provides examples of the types of estimates required in the cases of food stamps, job training programs, and housing vouchers. [List et al. \(2023\)](#) similarly examine nudge policies in the cigarette, vaccine, and household energy markets.

Our thesis is that the same objects can be estimated in a structural model, provided we are ready to introduce additional assumptions on the fundamentals of the model. Discrete choice models as in [Berry et al. \(1995\)](#) are particularly well suited to estimate demand (hence demand elasticities) while incorporating “rich consumer heterogeneity to generate flexible substitution patterns between products [...] and account[ing] for imperfect competition.” ([Van Biesebroeck and Verboven, 2026](#)). We show the exact mapping to those demand models in Section 3. Intuitively, by imposing a particular functional form on utility, the full set of derivatives contained in equations (7) and (8) can be computed analytically or solved numerically (see [Small and Rosen, 1981](#); [Heckman and Vytlačil, 2007](#)).

Non-marginal policies. A crucial advantage of the sufficient-statistics framework is that it does not require additional assumptions to evaluate marginal investments. Therefore, the approach is particularly useful to identify where the government should invest the next dollar of public funds. However, as [Kleven \(2021\)](#) points out, “real-world reforms are never infinitesimal. If they were, we would not be very interested in them (p. 528)”. This is especially true for climate policy, for which massive public investments are needed ([Nordhaus, 2019](#)).

For a non-marginal reform, the sufficient-statistics framework delivers a first-order Taylor approximation of welfare effects ([Kleven, 2021](#)). Extrapolation of the local approximation implicitly assumes that demand is linear in the policy instrument. [Kang and Vasserman \(2025\)](#) show how the robustness of this approximation to deviations from linear demand varies across empirical applications. Flexible substitution patterns between products in structural models can generate such non-linearities, which is typically supported by the data.

Policy mix. A policymaker looking to adopt an optimal environmental program may consider using multiple policy instruments in tandem. [Hendren and Sprung-Keyser \(2020\)](#) recognized the empirical challenge this poses for estimating welfare impacts. In particular, when a policy change includes both a carrot and a stick (e.g., increasing a subsidy while imposing stricter eligibility requirements), whether the willingness-to-pay is positive or negative is unknown. Therefore, to study budget-neutral programs, the authors suggest comparing

the MVPF of an expenditure policy to that of a revenue-raising policy. Studying two programs in isolation and comparing their welfare impacts cannot capture potential interactions between two policies implemented simultaneously.¹⁷

Interactions between policy instruments might be substantial for policies implemented in markets with strong network effects, such as interventions in the automobile market. In our structural framework, two policies can be simulated jointly, allowing flexible interactions between policy instruments. Capturing these interactions is crucial for balancing policies with each other and determining the optimal mix of multiple policy instruments.

Policies not yet implemented. The ability to forecast counterfactual policies is a recognized advantage of structural models (Chetty, 2009; Todd and Wolpin, 2023). In the treatment-effects literature, Heckman and Vytlacil (2005, 2007) discuss how the marginal treatment effects framework can be used for counterfactual analysis. They conclude that “when the same policy questions addressed by the structural literature are asked of the treatment effect literature, the assumption sets used in the two literatures look very similar” (p. 673). While the marginal treatment effect has the advantage of relying on program evaluation methods, it requires an instrument with substantial support over the propensity score. Therefore, for programs where such an instrument is not available, such as subsidies for green technologies, structural models are a better-suited alternative.

In the rest of this Section, we show explicitly how the structural approach can be used to estimate welfare impacts of non-marginal policy changes and to determine the optimal policy mix when multiple policy instruments are implemented simultaneously. We conclude this Section with a brief discussion of inference.

2.3 Marginal vs. non-marginal policy changes

The framework developed in Section 2.1 considers small policy deviations, and measures the welfare return on the next dollar spent on a policy instrument θ_j , evaluated at the status quo policy θ_0 . Most real-life reforms, however, are not marginal: a government contemplating the introduction of a new subsidy program, for example, typically implies a large discrete change in the policy environment. In this section, we extend the concepts of willingness-to-pay, net

¹⁷List et al. (2023) perform a comparative welfare analysis using existing evaluations similar to the large MVPF studies of Hendren and Sprung-Keyser (2020) and Hahn et al. (2026). In particular, they also study joint nudge and taxation policies within the same setting. When analyzing joint programs, they have to parameterize the dependence structure between behavioral biases and the elasticity to the tax, so their approach (for joint programs) is in fact a structural one.

cost, and MVPF to non-marginal reforms, and we characterize the relationship between the marginal and non-marginal objects.

In what follows, we denote the willingness-to-pay and the MVPF of the non-marginal reform θ_1 as $\overline{WTP}(\theta_0 \rightarrow \theta_1)$ and $\overline{MVPF}(\theta_0 \rightarrow \theta_1)$, respectively.

Willingness-to-pay for a non-marginal reform. Consider a reform that moves the policy from θ_0 to θ_1 , and decompose it into a continuum of marginal reforms along a smooth path $\theta(t)$, $t \in [0, 1]$. Define the willingness-to-pay for each marginal reform spanned by t as the vector collecting the J instrument-specific willingness-to-pay, that is, $WTP(\theta(t)) = (WTP_1(\theta(t)), \dots, WTP_J(\theta(t)))$. Then, the willingness-to-pay for the non-marginal reform $\theta_1 = \theta(1)$ from the status quo $\theta_0 = \theta(0)$ is simply the sum of the small policy changes,

$$\begin{aligned} \overline{WTP}(\theta_0 \rightarrow \theta_1) &= \int_0^1 WTP(\theta(t)) \theta'(t) dt, \\ &= \sum_i \underbrace{\left[\frac{V_i(x_i^*(\theta_1), \theta_1) - V_i(x_i^*(\theta_0), \theta_0)}{\lambda_i} \right]}_{\text{Equivalent variation} = EV_i(\theta_1)}, \\ &= \mathcal{W}(\theta_1) - \mathcal{W}(\theta_0). \end{aligned} \tag{13}$$

The equivalent variation is equal to the change in social welfare if and only if the social weights are equal to the inverse of the marginal utility of income, $\psi_i = 1/\lambda_i$ (an assumption we maintain throughout the paper).¹⁸

It is important to note that the integral in (13) is path-independent: for any sequence of intermediate reforms the government implements to move θ_0 to θ_1 , the accumulated willingness-to-pay is the same and delivers exactly the transfer that consumers would accept in exchange for the reform.¹⁹ The same logic applies to government costs: the net cost function of the non-marginal reform integrates over the marginal changes in cost over the path $\theta(t)$, $t \in [0, 1]$,

¹⁸For non-marginal reforms, the equivalent variation is the income transfer at the baseline policy that would leave an agent indifferent between the status quo and the reform, that is, $V_i^*(\theta_0; m_i + EV_i) = V_i^*(\theta_1; m_i)$, where $V_i^*(\theta; m_i)$ denotes the indirect utility of consumer i at income m_i . The compensating variation evaluates instead the transfer at the new policy, that would make agent i indifferent with the status quo, that is, $V_i^*(\theta_0; m_i) = V_i^*(\theta_1; m_i - CV_i)$. With constant marginal utility of income, the two coincide; this would not hold if λ_i was a function of the policy, e.g., $\lambda_i = \lambda_i(\theta)$ (an assumption we relax in Appendix C.2). We work with the equivalent variation throughout, since they are the same object given our main specification.

¹⁹With income effects, that is, λ_i varying along the path, the marginal willingness-to-pay measures need not integrate to a money metric, and the evaluation of a non-marginal reform depends on the path of intermediate reforms.

and we have that $\Delta\mathcal{C}(\theta_0 \rightarrow \theta_1) = \mathcal{C}(\theta_1) - \mathcal{C}(\theta_0)$.

MVPF of a non-marginal reform. The MVPF of a non-marginal reform is the ratio of the willingness-to-pay and the net cost of the non-marginal reform, that is,

$$\overline{MVPF}(\theta_0 \rightarrow \theta_1) = \frac{\overline{WTP}(\theta_0 \rightarrow \theta_1)}{\Delta\mathcal{C}(\theta_0 \rightarrow \theta_1)} = \frac{\mathcal{W}(\theta_1) - \mathcal{W}(\theta_0)}{\mathcal{C}(\theta_1) - \mathcal{C}(\theta_0)} \quad (14)$$

We want to emphasize that we define the MVPF of a non-marginal reform as a ratio of integrals, and not the integral of the ratio that forms the MVPF of marginal reforms. To see the difference, parametrize a path $\theta(t)$, $t \in [0, 1]$, and let $MVPF(\theta(t))$ denote the MVPF of a marginal reform at $\theta(t)$ along the path traced by t . Then, we can show that

$$\overline{MVPF}(\theta_0 \rightarrow \theta_1) = \int_0^1 MVPF(\theta(t)) \omega(t) dt, \quad \omega(t) = \frac{d\mathcal{C}(\theta(t))/dt}{\mathcal{C}(\theta_1) - \mathcal{C}(\theta_0)}, \quad (15)$$

$$\neq \int_0^1 MVPF(\theta(t)) dt \quad (16)$$

that is, the non-marginal MVPF is the *cost-weighted* average of the marginal MVPFs along the path, where each intermediate reform is weighted by its share of the incremental cost. Provided that the cost increases along the path, the weights $\omega(t)$ are positive and sum to one.

Trapezoidal approximation. In practice, causal estimates deliver the marginal objects only at the observed policies, and extrapolating to non-marginal reforms typically requires assumption about the evolution of the sufficient-statistics over the policy space (e.g., constant elasticity, constant semi-elasticity). [Kleven \(2021\)](#) explains that a linear extrapolation of local elasticities akin to the [Harberger \(1964\)](#) triangular approximation of deadweight loss can be used to generate a trapezoidal approximation of welfare effects. The area below the linear projection over the relevant interval yields a first-order Taylor approximation of the WTP and cost of a non-marginal program.

When the marginal WTP and cost measures are available at both endpoints of the reform, the integral in (13) can be approximated by a similar trapezoid rule; the numerator and the denominator of (14) are each replaced by the average of their endpoint derivatives, multiplied by the size of the reform ([Kleven, 2021](#)). The approximation is exact when willingness-to-pay and costs are quadratic in the policy, and require no structural model beyond the two local estimates. Since our approach delivers exact integrals, we compare the accuracy of the

trapezoid approximation using each of the two methods (linear extrapolation and connecting both endpoints) in our empirical applications, see Section 4.

Normalization to the no-policy counterfactual. Unlike its marginal counterpart, the non-marginal MVPF is a two-point object: it depends on the baseline policy as much as on the reform itself, and comparisons across reforms are meaningful only from a common baseline. For this reason, we adopt the no-policy counterfactual as the reference point for the remainder of the paper, and write

$$\overline{MVPF}(\theta) \equiv \overline{MVPF}(0 \rightarrow \theta) = \frac{\mathcal{W}(\theta) - \mathcal{W}(0)}{\mathcal{C}(\theta)}, \quad (17)$$

where the denominator uses the normalization $\mathcal{C}(0) = 0$ (i.e., the cost function of Section 2.1 was normalized such that no reform implies no impact on government finances). The object $\overline{MVPF}(\theta)$ measures the average social return of the program as a whole, and it provides us an adequate benchmark to compare the returns of reforms across different empirical applications. We show in the next section how to use the non-marginal MVPF, evaluated from the no-policy baseline, to recover the optimal policy mix (e.g., the optimal share of the budget that should be allocated to each policy instrument θ_j) and the optimal scale of the policy (e.g., the total budget that should be spent on θ).

2.4 Determining the optimal policy

So far, we have shown how to derive the MVPF of marginal and non-marginal policies using a very general structural model. In practice, the MVPF framework allows for comparative analysis between policies and answers the question: *which policy instrument should the policymaker prioritize?*

Consider as an example the case of electric vehicles. A policymaker has access to two policy instruments to boost adoption and reduce pollution: subsidies to electric vehicles and the installation of public chargers. In the MVPF framework of [Hendren \(2016\)](#) and [Hendren and Sprung-Keyser \(2020\)](#), the policymaker should prioritize the policy with the highest MVPF.²⁰ That is, at the margin, the government should transfer one dollar away from subsidies to invest in charging infrastructures if the MVPF of providing public charging is greater than the MVPF of electric vehicle subsidies at the observed policy levels. Because the change is at the margin, they can abstract from the impact of that reallocation on the

²⁰We discuss the special case where the social MVPF is not equal to the monetary MVPF, which occurs whenever $\psi_i \neq 1/\lambda_i$, in Appendix C.1.

government's finances (i.e., they do not need to “close” the budget constraint). However, their approach is silent about what an optimal set of policies should look like, and what should be the optimal level of spending on the set of environmental policies. In this section, we leverage our structural framework to answer both questions.

Optimal reform given fixed government spending. We first determine the properties of an optimal set of policy instruments under a fixed government budget. Without loss of generality, we take as baseline the government's observed spending at the status quo policy θ_0 , denoted by $\bar{\mathcal{C}} = \mathcal{C}(\theta_0)$. The pairwise comparison in the electric vehicles example above suggests what an optimal policy vector must look like: since moving one dollar of spending from instrument k to instrument j changes welfare by $MVPF_j - MVPF_k$, some budget-neutral reallocation raises welfare as long as the MVPFs differ across policy instruments (Hendren and Sprung-Keyser, 2020). The reallocation opportunities are exhausted at a candidate policy θ^* only when each policy instrument delivers the same marginal return, that is,

$$MVPF_j(\theta^*) = \mu, \quad \forall j, \quad (18)$$

for some common value μ .²¹ Formally, we maintain the following assumptions throughout this section. The policy space Θ is convex and compact, and it contains the no-policy point $\theta = 0$. Social welfare $\mathcal{W}(\theta)$ and the cost function $\mathcal{C}(\theta)$ are continuously differentiable, with the normalization $\mathcal{C}(0) = 0$ of Section 2.3. The welfare weights are $\psi_i = 1/\lambda_i$, and the observed budget is positive, $\bar{\mathcal{C}} > 0$. Under these assumptions, equation (18) is the first-order condition of the following constrained welfare maximization problem, at an interior solution with a binding budget constraint and non-zero marginal costs,

$$\begin{aligned} \theta^* &= \arg \max_{\theta \in \Theta} \mathcal{W}(\theta) \quad \text{s.t.} \quad \mathcal{C}(\theta) \leq \bar{\mathcal{C}}, \\ &= \arg \max_{\theta \in \Theta} \mathcal{W}(\theta) + \mu \left(\bar{\mathcal{C}} - \mathcal{C}(\theta) \right), \end{aligned} \quad (19)$$

where the common value $\mu \geq 0$ is the shadow value of public funds.

Some of the results below also require that public spending exhibits diminishing returns, in the sense that the welfare gain from an optimally spent budget is a concave function of

²¹Equation (18) characterizes interior optima, but corner solutions are possible. The government may optimally leave some instruments unused, e.g., set a subsidy constrained to be non-negative. At such an optimum, the optimality condition is replaced by the complementary slackness condition $MVPF_j \leq \mu$, with equality for every instrument in use. That is, an instrument remains unused whenever its marginal social return falls short of the shadow value of public funds μ .

the budget. In our applications, welfare and costs are both convex in the policy. Welfare is convex because the expected utility of consumers, the log-sum of Section 3, is convex in the subsidies at given producer prices. Costs are convex because a subsidy is paid on a base that grows with the subsidy. Diminishing returns therefore require costs to be more convex than welfare, that is, the MVPF of the marginal reform must fall as spending rises. Concavity of welfare together with convexity of costs is sufficient for this, but it is too strong for our purposes. We state explicitly which results use the assumption, and we discuss in Appendix B.9 what happens when it fails, as it does for some of the policies that we study. In that case, the optimal mix is a corner solution, the equalized MVPF rises along the expansion path, and there is no interior optimal scale.

In the following proposition, we show that we can also solve for θ^* by maximizing the non-marginal MVPF subject to government spending being at least $\bar{C}$, that is,

$$\theta^* = \arg \max_{\theta \in \Theta} \overline{MVPF}(\theta) \quad \text{s.t.} \quad C(\theta) \geq \bar{C}, \quad (20)$$

where $\overline{MVPF}(\theta) = (\mathcal{W}(\theta) - \mathcal{W}(0))/C(\theta)$ is the non-marginal MVPF of the previous section, evaluated from the no-policy baseline.

Proposition 1. (*Equivalency of the welfare and MVPF problems*)

- (a) Let θ^* be an interior solution of (20) at which the spending constraint binds, $C(\theta^*) = \bar{C}$, and at which every instrument has a non-zero marginal cost, $\partial C(\theta^*)/\partial \theta_j \neq 0$. Suppose that welfare is not satiated below the budget, that is, some policy of cost exactly $\bar{C}$ solves the welfare problem (19). Then θ^* solves (19) and satisfies the first-order condition (18), with a shadow value of public funds

$$\mu = \overline{MVPF}(\theta^*) - \nu \bar{C} \geq 0, \quad (21)$$

where $\nu \geq 0$ is the Lagrange multiplier on the spending constraint of problem (20).

- (b) Conversely, let $\tilde{\theta}$ be a solution of (19) at which the budget constraint binds, and suppose that no larger budget, spent optimally, delivers a higher social return per dollar than $\bar{C}$. This is the case under diminishing returns to public spending. Then $\tilde{\theta}$ solves (20). If $\tilde{\theta}$ is in addition interior with non-zero marginal costs, it satisfies (18).

The intuition for the proof is as follows. With spending fixed, the denominator of the non-marginal MVPF is constant, so maximizing the social return per dollar is equivalent to maximizing the social return itself. The two conditions of part (a) cannot be dispensed

with. If welfare is satiated below the budget, the policy that maximizes the non-marginal MVPF still maximizes welfare on the budget surface, but a cheaper policy does better. The condition of part (b) is also necessary, since a larger budget with a higher return per dollar would be feasible in (20) and preferred. The full proof is relegated to Appendix B.3.

Proposition 1 has an interesting consequence. In the setting of part (a), $MVPF_j(\theta^*) = \overline{MVPF}(\theta^*) - \nu \bar{\mathcal{C}} \leq \overline{MVPF}(\theta^*)$ for all j , since $\nu \geq 0$. The inequality is strict when the returns to public spending are strictly diminishing, and an interior solution is not enough for it (Appendix B.4). The gap between the two objects measures the curvature of social returns to the policy, and we report both in our applications. The MVPF of the non-marginal reform is an upper bound to the MVPF of the marginal reform, therefore, it is generally not comparable to the MVPF estimated from causal estimates (for example, the estimates reported by Hahn et al. 2026).²²

The two problems in (19) and (20) deliver the same optimal policy even though the iso-MVPF curves and iso-welfare curves are distinct objects. Proposition 2 below formalizes why.

Proposition 2. *Let $\mathcal{C} = \{\theta \in \Theta : \mathcal{C}(\theta) = \bar{\mathcal{C}}\}$ denote the government’s budget surface. Then,*

- (i) *Restricted to $\mathcal{C}$, the non-marginal MVPF is an increasing affine transformation of social welfare, $\overline{MVPF}(\theta) = (\mathcal{W}(\theta) - \mathcal{W}(0))/\bar{\mathcal{C}}$. The two objectives therefore have the same maximizers on the budget surface.*
- (ii) *At any point of $\mathcal{C}$ where $\nabla \mathcal{C} \neq 0$, the iso-MVPF curve is tangent to the budget surface if and only if the iso-welfare curve passing through the same point is tangent to the budget surface.*²³
- (iii) *Away from $\mathcal{C}$, the iso-MVPF curves and the iso-welfare curves generally do not coincide. With at least two instruments, the non-marginal MVPF is a transformation of $\mathcal{W}$ on an open set of policies only if the marginal MVPFs are equalized across instruments throughout that set.*

The proof is relegated to Appendix B.5.

²²The inequality reverses when the return per dollar of public spending is still rising at the observed scale, e.g., for a public good whose benefits grow with an S-shaped adoption while its costs do not. A small program can then have a marginal MVPF above the non-marginal MVPF. The comparison of the two objects along a policy path is the empirical test of diminishing returns along that path, and we document reversals in our applications. Appendix B.9 discusses this case.

²³The level curve of a function is tangent to the budget surface at a point when the function is stationary along the surface at that point, i.e., when its gradient is proportional to $\nabla \mathcal{C}$.

Figure 1: Optimal policy mix under a fixed budget




Figure 1 illustrates these results in a two-dimensional policy space, where the government chooses, for example, the level of an electric-vehicle subsidy θ_1 and the number of public charging stations θ_2 . Panel (a) depicts the iso-welfare curves against the fixed budget constraint and panel (b) the iso-MVPF curves.

Notice that the two families of level curves are very different. Welfare increases outwards, while the non-marginal MVPF is highest close to the origin, where the first dollars of public spending generate the largest returns. Nevertheless, both families are tangent to the budget curve at the same point θ^* , as stated in Proposition 2.²⁴ The figure also shows a status-quo policy θ_0 with the same total cost as θ^* . Since θ_0 lies on a lower iso-MVPF curve (and a lower iso-welfare curve), the same budget generates a smaller social returns. In this case, a reallocation from the low-MVPF policy instrument to the high-MVPF policy instrument generates more social welfare for the same cost.²⁵

²⁴The two problems push against the budget curve from opposite sides; welfare maximization spends up to the budget, non-marginal MVPF maximization down to it, and both meet at the tangency. Taken together, Propositions 1 and 2 imply that we can solve for the optimum using either problem, and recover at the tangency point the policy where the marginal MVPFs are equalized.

²⁵Figure 1 represents the solution of a well-behaved problem that yields an interior solution. In practice, a plethora of solutions are possible. In our empirical application, we find that corner solutions are frequent (e.g., removing a tax or divesting a subsidy entirely). In the case of nudges and public goods, it is also possible that the MVPF is not decreasing along the policy path (over some range) since these interventions have first-order effects on utility that can surpass their implementation cost. When such a situation occurs, we would have typically that $\overline{\text{MVPF}}(\theta) < \text{MVPF}(\theta)$ in the direction of that policy instrument.

Direction-independence at the optimum. Recall from Section 2.1 that, when the policy space is multidimensional, the MVPF of a marginal reform depends on the direction of the reform. It is straightforward to show that the MVPF of a reform in any direction v with a nonzero marginal cost, $\nabla \mathcal{C} \cdot v \neq 0$, and with a non-zero marginal cost for every instrument that the reform moves, is a weighted average of the MVPFs along each dimension (Appendix B.6), that is,

$$\frac{\nabla \mathcal{W}(\theta) \cdot v}{\nabla \mathcal{C}(\theta) \cdot v} = \sum_j \omega_j MVPF_j(\theta), \quad \omega_j = \frac{(\partial \mathcal{C} / \partial \theta_j) v_j}{\sum_k (\partial \mathcal{C} / \partial \theta_k) v_k}, \quad (22)$$

where the weights are given by the contribution of each policy instrument to the net cost of the reform.²⁶ At the optimum, each $MVPF_j$ is equal to μ by the first-order condition in (18). In that case, the MVPF of any reform with a non-zero fiscal cost, including a proportional expansion of all instruments, is also equal to μ . In other words, the MVPF is direction-independent at the optimal mix. Away from the optimum, the dispersion of MVPFs across instruments measures the extent of the misallocation and indicates which reallocation increases welfare. For instance, if the MVPF of charging stations exceeds the MVPF of vehicle subsidies at the observed policy, welfare rises by shifting spending from subsidies toward the network. We quantify this diagnostic in our empirical applications in Section 4.

Optimal scale of the reform. We now turn to the second question raised at the beginning of this section: *what is the optimal level of government spending?* We describe in what follows how the mix and scale problems can be solved simultaneously. As the level of spending expands, the tangency points between the iso-welfare curve, the iso-MVPF curve, and the budget constraint trace a curve in the policy space. We call this curve the *budget expansion path*, that is, the locus of optimal policy mixes at every level of government spending. Interestingly, budget curves are not needed to construct this curve. Since both families of level curves are tangent to the budget constraint at every optimum on the path, they are tangent to each other. The expansion path is therefore contained in the locus of tangency

²⁶The weights are positive for reforms that expand spending on every instrument, e.g., a proportional expansion of the program. For reforms that cut spending on some instruments while raising it on others, some weights are negative, and the directional MVPF can lie outside the range of the $MVPF_j$'s. By construction, the weights ω_j sum to one (even if some of the ω_j are negative). Movements along the budget constraint, $\nabla \mathcal{C} \cdot v = 0$, are excluded because the marginal cost of the reform is zero. Away from the optimum, these are exactly the budget-neutral reallocations of Hendren and Sprung-Keyser (2020): whenever the MVPFs are not equalized, reallocating spending from the low- to the high-MVPF instrument generates a strictly positive willingness-to-pay at exactly zero net cost (i.e., the reform “pays for itself”), so the MVPF of the reform is $+\infty$ in the direction of the optimum and $-\infty$ in the opposite direction.

points between iso-welfare and iso-MVPF curves, as shown in panel (a) of [Figure 2](#), and it coincides with that locus in the well-behaved case. We formalize this result in the following corollary.

Corollary 1. *Define the expansion path as the locus of optimal policy mixes at every level of government spending, that is,*

$$\mathcal{E} = \{\theta^*(C) \in \Theta : C > 0\},$$

where $\theta^*(C)$ is an interior solution to the welfare problem (19) with budget C . Let $\tilde{\theta}$ be an interior policy such that $\partial \mathcal{W}(\tilde{\theta})/\partial \theta_j$ and $\partial \mathcal{C}(\tilde{\theta})/\partial \theta_j$ are non-zero and of the same sign for all j , and suppose that $\overline{MVPF}(\tilde{\theta}) \neq 0$ and $\mathcal{C}(\tilde{\theta}) \neq 0$.


- (a) *If $\tilde{\theta}$ lies on the expansion path, the iso-MVPF curve and the iso-welfare curve passing through $\tilde{\theta}$ are tangent at $\tilde{\theta}$.*
- (b) *Conversely, if the two curves are tangent at $\tilde{\theta}$, then $\tilde{\theta}$ satisfies the first-order condition (18) of the welfare problem at the budget $\mathcal{C}(\tilde{\theta})$, with a positive shadow value of public funds. Whenever this first-order condition is sufficient, e.g., if welfare is concave and costs are convex, $\tilde{\theta}$ lies on the expansion path.*

In particular, the expansion path is contained in the locus of tangency points between the two families of level curves, and coincides with it under the sufficiency condition of part (b). It is identified from the two families of level curves alone, without reference to the budget constraint $C = \mathcal{C}(\tilde{\theta})$.

The proof is relegated to [Appendix B.7](#). In general, the tangency points are the critical points of welfare on the budget surface, and the expansion path is the subset of those that are maxima of welfare at their own budget. At each point of the path, the equalized $MVPF_j$ of equation (18) is the shadow value of public funds at that level of spending, so the tangency points map out a function $\mu(C)$, displayed in panel (b). By the envelope theorem, $\mu(C)$ measures the welfare gain from one additional dollar of budget spent optimally, and it declines along the path under diminishing returns ([Appendix B.8](#)). The scale problem then reduces to identifying the point on the $\mu(C)$ schedule where the marginal welfare gain equals the marginal cost of public funds: the government should expand the program as long as an additional dollar spent optimally generates more welfare than raising that dollar costs, that is, until $\mu(C)$ falls to the MVPF of the marginal source of public funds.

When raising one dollar costs society exactly one dollar, the optimal scale solves $\mu(C^*) = 1$.

Figure 2: Optimal scale of the reform



Note: The shape of the μ function in Panel (b) assumes the policy has diminishing returns.

Notice that the policy recovered this way is a second-best solution. With a complete set of tax instruments, the optimal policy at the optimal scale would be the Pigouvian schedule, that is, a tax on each product equal to its marginal external damage.²⁷ In our empirical applications, the set of instruments is restricted (i.e., we are limited to the policies featured in each work, plus any policy we can reasonably implement from the estimated parameters) and product markets can be imperfectly competitive.²⁸ The expansion path therefore characterizes the best policy within the set of instruments actually available to us, and the welfare distance to the Pigouvian benchmark is measurable within each model.

In the illustration of Figure 2, the observed budget $C(\theta_0)$ exceeds the optimal scale C^* : the shadow value of public funds at the observed spending level lies below one, so implementing a second-best policy entails scaling the program back along the expansion path, even with the mix chosen optimally at every level of spending. Notice also that the expansion path is not a ray: as the budget grows, the composition of the optimal mix changes with the relative returns of the policy instruments.

²⁷As pointed out by [Hahn et al. \(2026\)](#), the MVPF of a Pigouvian tax is equal to exactly 1, by construction.

²⁸More generally, Pigouvian taxes are often impossible to implement for political-economy reasons.

2.5 Inference

A practical advantage of our structural approach is that inference on the MVPF is straightforward. Every object entering the willingness-to-pay and the net cost (e.g., demand levels, behavioral responses, passthrough, profits) is a known and differentiable function of the estimated structural parameter vector, denoted $\hat{\zeta}$. The MVPF of any reform θ , marginal or non-marginal, is therefore a statistic of the form $MVPF_j = \phi_j(\hat{\zeta}, \theta)$, and the sampling uncertainty of $\hat{\zeta}$ translates into uncertainty about the MVPF through standard methods (e.g., the delta method, bootstrap).²⁹

Because all our welfare measures derive from a single estimated model, their joint distribution is available, not only their marginal distributions. Consider our example from the previous section, where a government both subsidizes electric vehicles and provides public charging stations. Since the same structural model encompasses both policies, inference on the comparison of the two policy instruments becomes possible (e.g., a test for equality of their MVPFs). This contrasts with welfare comparisons assembled from separately published causal estimates, where the different components of the MVPF are estimated on different samples and the correlation across estimates is unknown. In this case, comparisons across policies require independence assumptions (see [Hendren and Sprung-Keyser \(2020\)](#); [Hahn et al. \(2026\)](#) for the treatment of this problem by simulation).

Additionally, we can conduct inference on any not-yet-implemented policy instrument, as long as the relevant primitives are included in the estimated structural model. For example, a model of car choice can be used to compare an electric vehicle subsidy to a tax on gas guzzlers, even if the latter policy was never implemented by the government.

Finally, some inputs of the welfare calculation are calibrated rather than estimated; the social cost of carbon is a notable example. We do not treat them as random. Instead, we report how the MVPF varies with the calibrated inputs, in the spirit of the specification analysis in [Hahn et al. \(2026\)](#).

²⁹In practice, we favor the bootstrap, which avoids computing analytical gradients and delivers the entire distribution of the MVPF in one pass. We draw B parameter vectors $\zeta^{(b)}$, $b = 1, \dots, B$, either from the asymptotic distribution or by resampling the data, and we recompute the market equilibrium and the MVPF at each draw and each point on the policy grid. The resulting pointwise confidence bands cover the entire MVPF curve, so the same simulations that trace the curve also deliver confidence intervals. The procedure can be easily parallelized, and no re-estimation of the model is required.

3 Application to discrete choice demand models

We now illustrate how to adapt our methodology to a fully specified econometric model. We follow the canonical random coefficient discrete choice model of demand of [Berry et al. \(1995\)](#), the workhorse model of empirical industrial organization.³⁰ We consider an economy with N agents choosing among J products, indexed by $j = 0, 1, \dots, J$, where $j = 0$ indicates the outside option of making no purchase. The set of products consists of a number of green and brown goods, and the government's policy instrument is a set of subsidies and taxes $\theta = (\tau_1, \dots, \tau_J)$, where $\tau_j > 0$ indicates a subsidy.³¹

Setup. Each product j is represented by a set of observed and unobserved product characteristics, denoted by x_j and ξ_j respectively, and is associated with some environmental damages e_j (e.g., present value of lifetime emissions). We define consumer i 's utility for product j as

$$u_{ij}(\zeta_i, \theta, \epsilon) = x_j \beta_i + \lambda_i (y_i - p_j + \tau_j) + \xi_j + \epsilon_{ij} \equiv \delta_{ij} + \lambda_i y_i + \epsilon_{ij}, \quad (23)$$

where $y_i = m_i + (1 - \tau_c) s_i \Pi(\mathbf{q})$ is the consumer's total income (including their share of firms' profits, which depends on the aggregate consumption profile $\mathbf{q}$), p_j is the price of product j , and $\zeta_i = (\beta_i, \lambda_i)$ are consumer-specific preference parameters (including the marginal utility of income). The utility of the outside good is normalized to $u_{i0} = \lambda_i y_i + \epsilon_{i0}$, i.e., $\delta_{i0} = 0$.

Assuming that the taste shocks ϵ_{ij} follow an i.i.d. Gumbel distribution, the probability that agent i purchases product j is

$$P_{ij}(\zeta_i, \theta) = \frac{\exp(\delta_{ij} + \lambda_i y_i)}{\sum_{k=0}^J \exp(\delta_{ik} + \lambda_i y_i)} = \frac{\exp(\delta_{ij})}{1 + \sum_{k=1}^J \exp(\delta_{ik})}, \quad (24)$$

and the aggregate consumption profile is

$$q_j(\theta) = \sum_i P_{ij}(\zeta_i, \theta). \quad (25)$$

Consumers experience a negative externality from the collective decisions of all agents. We

³⁰We do not discuss the estimation of the model presented in this section. We presume that the model was estimated in a previous step, and that we have obtained consistent estimates of all relevant parameters, including marginal costs. We refer the reader to [Conlon and Gortmaker \(2020\)](#) for the current best practice in demand estimation.

³¹We abstract from public goods and nudges; we reintroduce them in specific applications where they are relevant.

denote the collective decisions of all agents by the aggregate consumption in the economy, $\mathbf{q} = (q_1, \dots, q_J)$. The externality experienced by agent i is then $d_i(\mathbf{q}) = \sum_j q_j d_{ij}(e_j)$, which is a function of the environmental damage e_j associated to each product, aggregated over the collective consumption profile. The total utility of agent i upon purchasing product j is given by

$$V_{ij}(\zeta_i, \theta, \epsilon) = u_{ij}(\zeta_i, \theta, \epsilon) - d_i(\mathbf{q}). \quad (26)$$

We close the model by considering the firms producing the various products, indexed by $f = 1, \dots, F$. In general, we consider a multi-product oligopoly, with each firm offering a subset of the products $J_f \subset J$ (with a slight abuse of notation). Aggregate industry profits at a given consumption profile $\mathbf{q}$ are given by

$$\Pi(\mathbf{q}) = \sum_{j=1}^J (p_j - c_j) q_j, \quad (27)$$

where c_j is the marginal cost of product j . Marginal costs are generally not directly observed, but they can be recovered from the first-order conditions of multi-product Bertrand competition, using the estimated demand system (see [Berry et al. 1995](#); [Conlon and Gortmaker 2020](#)).³² The first-order conditions of the firms determine (jointly with demand) the equilibrium prices $p_j(\tau)$. We do not restrict subsidies and taxes to be fully passed through to consumers.

Social welfare and total cost. We now express the welfare utility formally at optimal consumer choices. Because the taste shocks are i.i.d. Gumbel, the expected maximized utility has the well-known closed form solution described by [Small and Rosen \(1981\)](#), up to

³²Notice that we are abstracting from learning-by-doing effects for tractability. Introducing dynamic effects requires formally modeling the marginal cost function, for example, using Wright's rule ([Wright, 1936](#)), and transforming the supply side into a dynamic model, since firms have an incentive to lower the price today to accelerate learning and increase profits in the future. Since learning-by-doing is not a central feature of our applications, we ignore its contribution to the MVPF. See [Hahn et al. \(2026\)](#), who derive a sufficient statistic for society's willingness to pay for these dynamic effects, combining the slope of the learning curve with the demand elasticity, and show how to incorporate it into the MVPF of production subsidies for green technologies.

an additive constant γ ,

$$\begin{aligned}
V_i^*(\zeta_i, \theta) &= E_\epsilon \left(\max_j V_{ij}(\zeta_i, \theta, \epsilon) \right), \\
&= \ln \left(\sum_{j=0}^J \exp(\delta_{ij} + \lambda_i y_i - d_i) \right) + \gamma, \\
&= \lambda_i (1 - \tau_c) s_i \Pi(\mathbf{q}) + \ln \left(1 + \sum_{j=1}^J \exp(\delta_{ij}) \right) - d_i(\mathbf{q}) + \gamma_i,
\end{aligned} \tag{28}$$

where the last equality factors the income term and the externality, common to all alternatives, out of the log-sum (the term $\lambda_i m_i$ is absorbed by the additive constant $\gamma_i = \gamma + \lambda_i m_i$). Notice that income and the externality, which drop out of the computation of the choice probabilities (hence are irrelevant for consumers' decisions) reappear in the welfare calculation through the expected maximum, mirroring the organizing principle of Section 2.1. We define the total social welfare as

$$\mathcal{W}(\theta) = \sum_i \psi_i V_i^*(\zeta_i, \theta) = \sum_i \frac{V_i^*(\zeta_i, \theta)}{\lambda_i}, \tag{29}$$

where the social weights are $\psi_i = 1/\lambda_i$ as previously. The total cost of implementing policy θ is

$$\mathcal{C}(\theta) = \sum_j q_j \tau_j - \tau_c [\Pi(\mathbf{q}(\theta)) - \Pi(\mathbf{q}(0))]. \tag{30}$$

Willingness-to-pay for a marginal reform. We now consider a reform that moves the status quo subsidy schedule θ_0 along a direction $\theta_j \equiv \tau_j$. Differentiating (29) with respect to τ_j yields the willingness-to-pay along direction τ_j , that is,

$$\begin{aligned}
WTP_j &= (1 - \tau_c) s \underbrace{\sum_k \left[q_k \frac{\partial p_k}{\partial \tau_j} + (p_k - c_k) \frac{\partial q_k}{\partial \tau_j} \right]}_{\text{Shareholders' net benefits}} + \underbrace{q_j - \sum_k q_k \frac{\partial p_k}{\partial \tau_j}}_{\text{Savings by inframarginals}} - \underbrace{\sum_k D_k(e_k) \frac{\partial q_k}{\partial \tau_j}}_{\text{Externality}}, \tag{31}
\end{aligned}$$

where $s = \sum_i s_i$ is the aggregate firm ownership share and $D_j = \sum_i d_{ij}/\lambda_i \equiv \kappa e_j$ is the monetary value of the environmental externality from product j , over all consumers (e.g., present value of lifetime emissions e_j times the social cost of carbon κ). The full derivation is relegated to Appendix B.1.

Equation (31) is the discrete choice counterpart of (7). Consumers' willingness to pay for a small change along τ_j depends on the additional profits generated from ownership in the firms, the mechanical benefit of the reform on inframarginal consumers, and the environmental externality. Notice that the envelope theorem applies here: taking the derivative of the log-sum term gets rid of behavioral responses and only the incremental gains to inframarginal consumers matter. Importantly, equation (31) shows that the willingness-to-pay depends on two key quantities: the derivative of the market demand, which has a simple closed form expression, and the passthrough, which typically does not. Nevertheless, it is straightforward to calculate $\partial p_k / \partial \tau_j$ using numerical derivatives. This means that, given structural parameters ζ_i , our approach allows for evaluating the willingness-to-pay at any alternative policy, say $\tilde{\theta}$, by recomputing the market equilibrium at $\tilde{\theta}$ and re-evaluating the derivative of the market demand and the passthrough at $\tilde{\theta}$.

The impact of the reform on government's finances is

$$\frac{\partial \mathcal{C}}{\partial \theta_j} = \underbrace{q_j}_{\text{Transfer to inframarginals}} + \underbrace{\sum_k \tau_k \frac{\partial q_k}{\partial \tau_j}}_{\text{Additional subsidy take-up}} - \underbrace{\tau_c \sum_k \left[q_k \frac{\partial p_k}{\partial \tau_j} + (p_k - c_k) \frac{\partial q_k}{\partial \tau_j} \right]}_{\text{Additional corporate tax revenues}}, \quad (32)$$

which includes the mechanical transfer to inframarginal consumers and two fiscal externalities: the additional subsidy payments induced by behavioral responses, and the change in corporate tax revenues collected on firms' profits.³³

Willingness-to-pay for a non-marginal reform. We now consider a non-marginal policy change from the status quo θ_0 to some counterfactual policy θ_1 , along the direction θ_j .³⁴ Since we are now considering a discrete change in the policy environment, the envelope theorem no longer applies: consumers who adjust their behavior in response to the reform may experience utility gains that are first-order. As mentioned in Section 2.3, an agent's willingness to pay for a non-marginal reform is measured by the equivalent variation, that

³³Notice that incomplete passthrough drives a wedge between the two sides of the ledger: the government spends the full transfer q_j on inframarginal consumers, while these consumers only save $q_j - \sum_k q_k \frac{\partial p_k}{\partial \tau_j}$. The difference, being the share of the subsidy captured by producers through higher equilibrium prices, reappears in firms' profits and is redistributed to consumers and to the government through ownership shares and the corporate tax.

³⁴We focus on a non-marginal policy change along θ_j for comparability with the marginal change described above. Our results on the MVPF of non-marginal policies hold more generally for any policy path between θ_0 and θ_1 .

is,

$$EV_i(\theta_1) = \frac{V_i^*(\zeta_i, \theta_1) - V_i^*(\zeta_i, \theta_0)}{\lambda_i}. \quad (33)$$

Replacing V_i^* by its closed form expression and summing over all consumers, the willingness-to-pay for the reform is

$$\begin{aligned} \overline{WTP}_j = & \underbrace{\sum_i \frac{1}{\lambda_i} \left[\ln \left(1 + \sum_{k=1}^J \exp(\delta_{ik}(\theta_1)) \right) - \ln \left(1 + \sum_{k=1}^J \exp(\delta_{ik}(\theta_0)) \right) \right]}_{\text{Consumers' willingness to pay}} \\ & + \underbrace{(1 - \tau_c) s \left[\Pi(\mathbf{q}(\theta_1)) - \Pi(\mathbf{q}(\theta_0)) \right]}_{\text{Shareholders' willingness to pay}} - \underbrace{\sum_{k=1}^J D_k \left[q_k(\theta_1) - q_k(\theta_0) \right]}_{\text{Society's willingness to pay}}. \end{aligned} \quad (34)$$

The willingness-to-pay for the reform therefore decomposes into the willingness-to-pay of three groups: consumers, whose money-metric utility gain is the change in the expected maximum utility converted into dollars; firm owners, who receive the after-tax share of the change in profits; and the individuals bearing the environmental damages. Because our utility specification is quasi-linear in income, the first two terms coincide with the changes in consumer and producer surplus familiar from welfare analysis in industrial organization; the log-sum difference is the expected compensating variation of [Small and Rosen \(1981\)](#). We adopt the public finance labels because welfare weights and the government's ledger attach to groups of individuals rather than to sides of the market, which is what allows the incidence of the reform to be traced consumer by consumer in our applications. The net cost of the policy is straightforward,

$$\Delta \mathcal{C}_j = \underbrace{\left[\tau_j(\theta_1) - \tau_j(\theta_0) \right] q_j(\theta_1)}_{\text{Transfer to inframarginals}} + \underbrace{\sum_k \tau_k(\theta_0) \left[q_k(\theta_1) - q_k(\theta_0) \right]}_{\text{Additional subsidy take-up}} - \underbrace{\tau_c \left[\Pi(\mathbf{q}(\theta_1)) - \Pi(\mathbf{q}(\theta_0)) \right]}_{\text{Additional corporate tax revenues}}, \quad (35)$$

which is the discrete counterpart of equation [\(32\)](#).

4 Applications to the empirical studies

We now apply the methodology developed in the previous sections to three empirical studies of environmental policies. We describe each work and its associated policy environment in

what follows. Our empirical analysis of environmental policies covers the market for electric vehicles in Quebec (Fournel, 2026), residential solar panels in Connecticut (Gillingham and Tsvetanov, 2019) and an online experiment on car choice in the United States (Allcott et al., 2025). These studies span the three types of policy instruments described in the previous sections: taxes and subsidies (e.g., rebates on electric vehicles, taxes on fuel-inefficient vehicles), public goods (e.g., public charging stations), and informational nudges (e.g., fuel-economy labels).

Table 1 describes the policies under consideration in each study. The policy space includes the set of observed policies that are in place in the data and a set of counterfactual policies that could have been under consideration by the policymakers at the time the observed policies were implemented. Unless stated otherwise, the MVPF of each policy is computed at the observed vector of policy instruments, with the counterfactual policies set to zero. Our structural approach allows us to evaluate the MVPF at arbitrary policy levels, which we use to determine optimal policy mixes, whenever it is relevant (see Section 4.6).

Appendix D describes each study’s structural model and the computation of the MVPFs. We implement our methodology directly on the models and data of Fournel (2026) and Gillingham and Tsvetanov (2019). Allcott et al. (2025) present a structural model, but their analysis relies solely on the experimental variation to identify welfare impacts. We chose to estimate their structural model by maximum likelihood on their experimental data. The full methodology and the estimation details are available in Appendix D.3.2.

We keep the fiscal environment of each study and only update the social cost of carbon. We use the recent estimates from the Environmental Protection Agency (2023), expressed in the year and currency of each study. Our MVPFs therefore remain “in context”, in the terminology of Hahn et al. (2026), with a valuation of carbon that reflects the most recent estimates.

4.1 Taxes and subsidies

Two of our studies evaluate price instruments, the electric vehicle subsidies in Quebec (Fournel, 2026) and the subsidies to residential solar panels in Connecticut (Gillingham and Tsvetanov, 2019). We discuss the tax on the low-MPG car of Allcott et al. (2025) together with the fuel-economy label in Section 4.3.

Electric vehicles. Fournel (2026) studies electric vehicle incentives in Quebec, Canada, and the supply of public charging infrastructures, using vehicle registration and charging

Table 1: Policies under study in our empirical applications

Policy	Instrument	Level
<i>Panel A. Electric vehicles, Fournel (2026), Quebec 2020</i>		
Observed policies:		
Battery electric rebate	Subsidy	\$8,000 prov. + \$5,000 fed.
Plug-in hybrid rebate	Subsidy	\$4,000 prov. + \$2,500 fed.
Charging stations (public provision)	Public good	440 added, \$19,000/station
Counterfactual policies:		
Hybrid rebate	Subsidy	Not yet implemented
Tax on gas guzzlers (>210 g/km)	Tax	Not yet implemented
<i>Panel B. Residential solar, Gillingham and Tsvetanov (2019), Connecticut 2014</i>		
Observed policies (at their 2014 levels):		
State rebate	Subsidy	\$0.935 per watt (\$7,300 per system)
Federal tax credit	Subsidy	30% ad valorem (\$7,000 per system)
Counterfactual policies (at their 2021 levels):		
State rebate	Subsidy	\$0.28 per watt (\$2,200 per system)
Federal tax credit	Subsidy	26% ad valorem (\$7,200 per system)
<i>Panel C. Fuel-economy labels, Allcott et al. (2025), United States 2019 experiment</i>		
Observed policies:		
Fuel-economy label	Nudge	Universal, assumed cost \$1, \$5 or \$10
Counterfactual policies:		
Fuel-economy label, with \$265 corrective tax	Nudge	As above

Notes: This table lists the observed policies evaluated in each application, in the currency and year of the source. The electric-vehicle rebates are the amounts for standard models, i.e., battery electric vehicles and short-range plug-in hybrids; long-range plug-in hybrids receive the battery electric amounts. The solar rebate is the average of a declining step schedule in 2014, and the federal tax credit applies to the price net of the rebate; per-system amounts are for the mean installation of 7.82 kW. The fuel-economy label was shown to every participant in the treatment arms of the experiment; its implementation cost is not observed, so we report results for three assumed values per exposed consumer and vehicle-year.

station data from 2012 to 2020. The provincial government in Quebec initiated its rebates program 2012, subsidizing heavily battery electric vehicles and plug-in hybrids by \$8,000 and \$4,000 respectively.³⁵ The federal government added its own rebate program in 2019, providing an additional \$5,000 subsidy to battery electric vehicles and \$2,500 to plug-in hybrids. We consider two additional policy instruments for our study of the Quebec car market. The first one is a subsidy on non-rechargeable hybrids (available in Quebec until 2018). The second one is a tax on fuel-inefficient vehicles, with tailpipe emissions above 210 grams per kilometer (this includes several but not all mass-market sedans and SUVs). To ensure our welfare analysis is not based on first adopters, we compute all MVPFs on year

³⁵A long-range plug-in hybrid, the Chevrolet Volt, was eligible for the same rebate as battery electric vehicles. For the purpose of this exercise, we will classify this model as a battery electric vehicle.

Table 2: MVPF of our empirical studies

Policy	Marginal MVPF	Non-marginal MVPF	Trapezoid approximation (structural)	Trapezoid approximation (extrapolated)
<i>Panel A. Electric vehicles, Fournel (2026), Quebec 2020</i>				
Subsidies and taxes:				
Battery electric rebate	1.018 (0.037)	1.269 (0.061)	1.176 (0.050)	1.154 (0.045)
Plug-in hybrid rebate	0.993 (0.046)	1.225 (0.067)	1.159 (0.060)	1.133 (0.054)
Hybrid rebate	1.480 (0.106)			
Tax on gas guzzlers	1.047 (0.092)			
Public charging:				
Public charging stations	1.957 (0.190)	1.953 (0.180)	1.953 (0.180)	1.880 (0.169)
<i>Panel B. Residential solar, Gillingham and Tsvetanov (2019), Connecticut 2014</i>				
Incentives at their 2014 levels:				
State rebate	0.912 (0.021)	0.970 (0.011)	0.963 (0.014)	0.961 (0.013)
Federal tax credit	0.897 (0.021)	0.998 (0.012)	0.975 (0.017)	0.972 (0.015)
Incentives at their 2021 levels:				
State rebate	1.035 (0.008)	1.055 (0.011)	1.054 (0.011)	1.054 (0.011)
Federal tax credit	1.030 (0.008)	1.142 (0.027)	1.118 (0.022)	1.115 (0.021)
<i>Panel C. Fuel-economy labels, Allcott et al. (2025), United States 2019 experiment</i>				
Fuel-economy label:				
Implementation cost of \$1	4.482 (0.492)	4.347 (0.443)	4.347 (0.444)	4.361 (0.466)
Implementation cost of \$5	2.525 (0.345)	2.475 (0.318)	2.473 (0.318)	2.495 (0.334)
Implementation cost of \$10	1.634 (0.247)	1.608 (0.231)	1.607 (0.231)	1.626 (0.242)
Fuel-economy label with corrective tax:				
Implementation cost of \$1	1.496 (0.084)	1.486 (0.082)	1.486 (0.082)	1.492 (0.083)
Implementation cost of \$5	1.192 (0.082)	1.177 (0.079)	1.177 (0.079)	1.191 (0.081)
Implementation cost of \$10	0.951 (0.078)	0.934 (0.076)	0.935 (0.076)	0.951 (0.078)


Notes: Marginal MVPF of a small change at the observed policy; non-marginal MVPF of the reform from no policy to the observed level; trapezoid approximation of Kleven (2021), using the marginal objects at both ends of the reform, either from the model (structural) or extrapolated from the observed policy at a constant price elasticity (extrapolated), as in Hahn et al. (2026). Standard errors from parametric bootstrap draws in parentheses.

2020, the last year of the data.

At the observed policy, an additional dollar of rebate returns \$1.02 of benefits for battery electric vehicles and \$0.99 for plug-in hybrids (see Panel A of Table 2). Essentially, no significant gains can be realized by reallocating government funds from plug-in hybrids to battery electric vehicles, that is, the policy is almost at its optimal mix on the budget surface. The policies generate on average \$1.27 and \$1.23 in benefits for each dollar spent, because the first dollars spent are the most effective (i.e., marginal consumers generate the same environmental benefits at a lower cost). The returns of these instruments diminish with the level of spending, consistent with the diminishing returns assumed in Section 2.4.

Figure 3 shows how the marginal MVPF falls over the support of the policy space. As stated above, each additional buyer receives the full rebate, so the cost of the marginal adoption rises with the rebate while the benefits do not, leading to diminishing returns. We can decompose these MVPFs into their sources of funding, e.g., by the effect of the provincial rebate and the federal rebate top up. Each dollar spent by the provincial government returned on average \$1.4 dollars in benefits, while the federal program, added afterwards, generated \$1.1 in benefits for each dollar spent beyond the provincial program (see Table A4).

Figure 3: MVPF of subsidies and taxes over policy space



Notes: This figure displays MVPF estimates over the policy space of the observed and counterfactual subsidy programs in Quebec, based on Fournel (2026). The shaded region represents the 90% confidence interval computed from 199 bootstrap samples.

Finally, we consider two counterfactual policies, a rebate to non-rechargeable hybrids and a tax on gas guzzlers, defined as any vehicle with an emission rating above 210 grams per kilometer. The results of Table 2 and Figure 3 suggest that a potential reallocation of funds from the observed policies to finance hybrid subsidies would be welfare improving (e.g., the MVPF of hybrid subsidies is 1.48 versus 1.02 and 0.99). However, raising revenues through a tax on gas guzzlers could be ineffective, unless the funds are used to finance subsidies to hybrids. The MVPF of the tax is above one, meaning that each dollar raised through the


tax reduces welfare by more than one dollars (e.g., consumers have a high willingness-to-pay to avoid the tax). The returns to the tax are also strictly increasing, consistent with the theory developed in Section 2.4.

Residential solar panels. Gillingham and Tsvetanov (2019) study the demand for residential solar photovoltaic systems in Connecticut, using the purchased systems installed in 1,534 Census block groups from 2008 to 2014. Two separate financial incentives are available to consumers. The first is a state incentive is a subsidy to the installer, on a per watt basis, provided through the Residential Solar Investment Program.³⁶ The second incentive is an *ad valorem* federal tax credit on the transaction price, net of the state rebate. A higher state rebate therefore lowers the federal credit paid on every system.

Panel B of Table 1 reports the policy levels in 2014, at the time of the study, and the policy levels in 2021. The state incentives were substantially reduced, from \$0.935 to \$0.28 per watt (\$7,300 and \$2,200 per system) and the federal tax credits was also reduced from 30% to 26% (\$7,000 and \$7,200 per system). We compute all MVPFs on the installations of 2014, the last year of the data, and we consider the impact of both the 2014 and the 2021 policies.

An additional dollar on either subsidy returns about \$0.90 in benefits at the 2014 levels, and about \$1.03 at the 2021 levels. Figure 4 traces the MVPF of each policy over their respective support, with the dashed lines representing the 2014 and the 2021 levels (the other policy is maintained at the 2014 baseline). In both cases, we observed strict diminishing returns.

Figure 4: MVPF of the solar subsidies over the policy space



Notes: This figure displays the marginal MVPF of each solar subsidy over the level of the subsidy, with the other subsidy at the 2014 level. The dashed vertical lines mark the 2014 and 2021 policy levels. The shaded region represents the 90% confidence interval computed from 199 bootstrap samples.

³⁶The installers keep 16% of the rebate, the rest accrues to the consumers.

4.2 Public goods

Public charging. Fournel (2026) studies the public provision of charging infrastructures by local government in Quebec. Local governments provide charging to respond to the demand for charging services from new and past electric vehicle owners, and the provincial and federal governments encourage electric vehicle adoption through subsidies. In what follows, we consider the case where the provincial government in Quebec invest in charging infrastructures directly, crowding out and replacing local planners.³⁷ Public charging then becomes a new policy lever that the government can use, jointly with subsidies, to boost electric vehicle adoption and reduce emissions.


We conduct our welfare analysis on year 2020, similarly to our evaluation of taxes and subsidies. There were 440 new station installations across all counties, at an assumed cost of \$19,000 per station. We emphasize that each dollar spent on public charging generates additional electric vehicle adoptions, which in turn generate fiscal externalities at the observed policy levels which increase costs beyond the mechanical dollars spent on public charging (e.g., the government needs to pay the electric vehicle subsidy to marginal consumers).

The results of the welfare analysis are available in Table 2, panel A. We compute a single MVPF for public charging that expand networks along the direction that preserves the county-level distribution of the 440 new chargers installed in 2020. We return to distributional effects at the end of the section. Each dollar spent on public charging yields almost two dollars in benefits, meaning that investments in public charging are about twice as effective as electric vehicles subsidies at the observed policy levels. The willingness-to-pay for charging infrastructures is high, about \$8.24 per mechanical dollar spent on charging stations, and each mechanical dollar spend on charging stations generates \$2.96 of additional costs in the form of subsidy take-up. Without subsidies, the MVPF of public charging would be significantly higher, around 3.94.

Figure 5 traces the MVPF of installing public charging stations over the policy space. Starting from the 2019 network, we add stations in every county in proportion to the 440 stations added in 2020. The key observation is that the MVPF does not fall with the policy level. This is also visible from Table 2, that is, the non-marginal MVPF is slightly lower than the marginal MVPF, indicating that the first dollars spent on the policy are slightly less efficient than future dollars. This is the reversal discussed in Section 2.4.

³⁷The crowding out of local planners is a result of the model rather than an assumption. In every counterfactual where we let the provincial government install chargers, local governments reduce supply one for one.

Figure 5: MVPF of expanding the charging network



Notes: Marginal MVPF of public charging stations in Quebec, adding stations in every county in proportion to the stations added in 2020. The dashed vertical line marks the 440 stations added in 2020. The shaded region represents the 90% confidence interval computed from 199 bootstrap samples.


In our application with public charging, the welfare function is convex, but costs increase almost linearly in the number of stations. For this reason, the budget surface is less convex than the welfare function and we observe non-decreasing MVPF over the policy space.³⁸ These types of welfare problems result in corner solutions, where the government either spends all its resources on the policy instrument in question, or fully divest from that policy instrument.

Public charging at the county level. Our analysis of public charging assumed that the policymaker installs chargers in proportion to the observed equilibrium outcome in the data. In practice, where we install public charging matters: users in cities may have a higher willingness-to-pay for public charging if home charging is not readily available. In what follows we break down the network decision into county-level policy instruments, all under the control of the central government, and we perform local welfare analyses.

Figure 6 represents the MVPF of adding charging capacity in each of the 96 counties in Quebec, evaluated at the observed policy levels. The counties are ranked by population. There is significant heterogeneity across country, the MVPF of adding one more station

³⁸We can generate convex budget surfaces by assuming that the cost of the public good, c_g is increasing in the level of public good provision. This could happen if, for example, opening new charging sites is increasingly costly due to the lack of installation teams, the scarcity of good charging locations, or technical challenges such as connecting a large number of infrastructures to the grid requiring local grid upgrades. We could not find evidence in the data and the procurement records we had access to that this assumption was reasonable.

Figure 6: MVPF of expanding network, by county



Notes: MVPF of expanding charging networks, by county (left), and difference between the MVPF of the marginal versus non-marginal policy (right). Counties are ranked by population, in decreasing order. The brackets represent the 90% confidence intervals, computed from 199 bootstrap samples. The non-marginal MVPF is calculated from each county's 2019 network to its 2020 network. Blue bars indicate diminishing returns and red bars indicate increasing returns.

varies from 0.27 to 2.36, and there is a strong relationship with population. Installing a station in a dense urban center benefits more consumers for the same cost, and this translates to stronger welfare effects.

The right panel of [Figure 6](#) shows the difference between the MVPFs of the marginal and the non-marginal policies, which indicates the counties with diminishing returns (blue) and increasing returns (red). For the overwhelming majority of counties, we observe diminishing returns. This could lead to an interior solution of a corner solution, depending on total spending and how fast the MVPF are falling in comparison to other policy instruments.

For a few large counties, we instead find increasing returns from more public charging. Since these counties include some of the most populous regions, the overall effect of adding public charging in all counties also exhibits increasing returns, as documented in [Figure 5](#). Because these counties also have among the highest willingness-to-pay per dollar spend, the optimal network is a corner solution where the government spend the full share of the budget devoted to public charging in the subset of large counties with increasing returns, and nothing elsewhere. This raises concerns about equity, which a policymaker could address by selecting a different set of welfare weights that favor disadvantaged counties.³⁹

³⁹See [Appendix C.1](#) for details on how to specify the welfare function using alternative welfare weights, similarly to [Hahn et al. \(2026\)](#).

4.3 Nudges

Fuel-economy label. Allcott et al. (2025) study the effect of a fuel-economy label using an online experiment on car choice in the United States. In the experiment, participants are asked to choose, for a three-year lease, between a fuel-inefficient sedan (rated at 23 miles per gallon) and a fuel-efficient sedan (rated at 30 miles per gallon) with otherwise similar characteristics. The online experiment is run on 1,267 participants, and 1,002 participants are exposed to an information label about the fuel efficiency of both options.⁴⁰ Allcott et al. (2025) find that consumers overvalue the fuel-inefficient car by \$135 per vehicle-year on average (e.g., because of limited attention to fuel costs) and that the label lowers the relative valuation of the fuel-inefficient car by \$59 on average (e.g., reducing their misperception about fuel costs).

The policy intervention we evaluate is a fuel-economy label shown to a fraction n of consumers.⁴¹ Our baseline estimates of the MVPF are all computed at the full exposure benchmark, e.g., $n = 1$ where all consumers are presented with the label. This is equivalent to a mandatory fuel-economy label displayed on each car in every showroom in the country. The cost of showing the label is not observed in the data, so we report the MVPFs for an implementation cost of \$1, \$5, and \$10 per exposed consumer and vehicle-year.

Our baseline estimates are reported in Table 2, panel C. The informational nudge generates between \$1.63 and \$4.48 of benefits per dollar spend on the policy, depending on the implementation cost. Among the chosen empirical studies, informational nudges are by far the most efficient from a MVPF point of view. The willingness-to-pay for the label is very high, around \$23 per consumer, mostly driven by bias correction (\$11.36) and avoided carbon damages (\$10.73). The cost to the government is split between the cost of the label (between \$1 and \$10), forfeited gasoline taxes (\$3.11), and lower corporate taxes (\$1.05).

Allcott et al. (2025) propose a corrective tax to account for the remaining consumer bias and the environmental externality. In our setting, this is a \$265 tax per vehicle-year on the fuel-inefficient vehicle.⁴² We perform the welfare analysis with this corrective tax vehicle in

⁴⁰Their design includes four different labels, assigned at random and equally among the 1,002 treated individuals. The first label reports the miles per gallon of each alternative. The second label reports the average yearly fuel cost of each car, based on the average mileage in the population. The third label is a personalized cost label, based on each respondents reported yearly mileage. Finally, the last group is exposed to a SmartWay label, an EPA certification for green vehicles (only the fuel-efficient vehicle has that certification).


⁴¹As explained previously, Allcott et al. (2025) presents a structural model, and calibrate it from moments in their experimental data to recover the welfare effects. We instead estimate their model on their data, using maximum likelihood. Computational details are relegated to Appendix D.3.2.

⁴²Allcott et al. (2025) use a social cost of carbon of \$51 per ton, which leads to a corrective tax of

place. The MVPFs are much smaller in this context, between 0.95 and 1.50. With the tax in place, the willingness-to-pay for the label barely changes, but each consumer who switches to the fuel-efficient car stops paying the tax. The lost tax revenue raises the net cost of the label, from \$5.16 to \$15.69 per mechanical dollar paid on the label (see Table A5). Since the tax already corrects the average distortion, the label mainly adds value by moving the most biased consumers more than the others.

To better understand the role of bias and the label design on the welfare analysis, we perform a series of additional counterfactual experiments. The results are presented in Figure 7. All experiments assume a \$1 implementation cost per consumer and vehicle-year. As a comparison, we report the MVPF of the label with and without the corrective tax (in grey) and the other experiments change elements of the model described in Appendix D.3.1 one by one. Implementation details regarding these additional counterfactual experiments are relegated to Appendix D.3.4.

Figure 7: MVPF of the fuel-economy label in alternative cases



Notes: Marginal MVPF of showing the fuel-economy label to every consumer, at an implementation cost of \$1 per exposed consumer and vehicle-year. The two baselines are the estimates of Table 2, without tax and with the \$265 corrective tax. Every other case changes one element of the model, without tax, as described in Appendix D.3.4. The intervals are the 90% confidence interval computed from 999 bootstrap samples.

We propose three experiments that change the design of the label. The debiasing label sets each consumer's bias to zero, after seeing the label. The fully targeted label reduces the bias on average by \$59, but levels out difference in biases across consumers (e.g., it has a larger effect the further away a consumer is from the average bias). The untargeted label reduces the bias by \$59 for all consumers, preserving differences across consumers. The associated

$135 + 50 - 59 = \$126$. We use a social cost of carbon at recent estimates, \$193 per ton, which gives a corrective tax of $135 + 189 - 59 = \$265$.

MVPFs are 7.02, 5.57, and 1.07, that is, improving targeting significantly increases the effectiveness of the information label.

In a second set of experiments, we study different assumptions about consumers' bias. We evaluate the effect of the label when the policymaker does not value the bias correction effect, and the willingness-to-pay for the label depends almost exclusively on the associated reduction in environmental externality. The MVPF is lower in this case, around 2.28. If instead consumers themselves are unbiased, then the label has only a small effect, statistically indistinguishable from zero.

In our final experiment, we impose a flat supply schedule, such that shifts in the relative demand of each car does not entail a change in their relative price. In our main specification, the relative price of the fuel-efficient car increases when consumer switch from the low-MPG to the high-MPG sedan, mitigating the effect of the information label. Shutting down the price channels therefore leads to a MVPF slightly higher, 6.01 versus 4.48, showing that equilibrium effects have non-negligible impacts on the effectiveness of the label.


4.4 Non-marginal reforms

We now focus on non-marginal reforms. We first evaluate an observed reform, the changes to the residential solar subsidies in Connecticut. The financial incentives were substantially reduced between 2014 and 2021, the state subsidy fell from \$0.935 to \$0.28 per watt and the federal tax credit from 30% to 26%. Second, we explore what happens when the MVPF of the marginal reform exceeds the MVPF of the non-marginal reform, as is the case for the fuel-economy label. In this case, the optimal policy is at either end of the policy space, and the choice between the two depends on the non-marginal MVPF of the whole program.

Reforming an existing program. The non-marginal MVPFs of [Table 2](#) evaluate each program against the no-policy counterfactual, but the non-marginal MVPF can also evaluate any reform between two existing programs. We consider the change in the residential solar program described in [Table 2](#). As previously, we perform the welfare analysis on year 2014, the final year of available data.

Both policies are represented in [Figure 8](#). The 2021 reform (red arrow) saves \$12.9 million of public spending and lowers welfare by \$12.5 million. The MVPF of that non-marginal reform is 0.966, that is, saving \$1 of government funds costs \$0.966 in welfare to society. Because the non-marginal MVPF is below 1, the 2021 program raises total surplus (i.e., welfare net of costs) by \$0.44 million compared with the 2014 program.

Figure 8: Non-marginal MVPF for residential solar policies



Notes: Solar subsidies in Connecticut, in the market of 2014, in the plane of the state rebate per watt and the federal tax credit rate. The solid curve gives the policies with the net public cost of the 2014 policy (\$27.9 million), and the dashed curve the policies with the net public cost of the 2021 policy (\$15.0 million). The red arrow goes from the 2014 policy to the 2021 policy, and the blue arrow from the 2014 policy to the policy that maximizes the willingness-to-pay on the 2021 budget. Each arrow reports the non-marginal MVPF of the move, $\overline{MVPF}$, i.e., the change in willingness-to-pay divided by the change in net public cost. Both moves cut spending, so a lower $\overline{MVPF}$ means less willingness-to-pay lost per dollar saved.

We compare these results with an alternative reform (blue arrow), set at the optimum on the 2021 budget surface. The MVPF of that non-marginal reform is 0.964, slightly lower than the observed 2021 reform. Since the reform reduces spending, a lower non-marginal MVPF indicates a better policy: welfare decreases by less for each dollar of savings. The total change in welfare is 0.47 million against 0.44, a modest gain over the observed reform. The policymakers in Connecticut kept the federal tax credit mostly intact while reducing the per watt subsidy significantly, effectively favoring one policy instrument over the other.


Qualitatively, our analysis of the non-marginal MVPFs reveals that government funds should have been directed instead towards the other policy instruments, a larger per watt subsidy with almost no *ad valorem* tax credit. In practice, the two policy instruments are set by two different levels of government. Coordination among policymakers in this case leads to a better outcome for the same total spending.

Implementing a policy with increasing returns. We focus on Allcott et al. (2025) and their fuel-economy experiment to study what happens when a policy exhibit increasing returns. Recall that the policy variable is the share of consumers that are exposed to the label, so the support of the policy goes from no coverage to full coverage. At all points on

the policy space, the MVPF of a marginal increase in coverage is greater than the MVPF of the non-marginal policy from the no coverage counterfactual. The solution of the welfare problem of Section 2.4 is then to expand the budget constraint to its maximum. In this case, this is a corner solution at the full coverage scenario, since every new dollar of spending generates more surplus than the last.

We show how we can refine on this equilibrium with a simple rule. We suggest that a policymaker may not want to implement a reform that, on average, generate surplus at a rate lower than one dollar of benefits for one dollar of additional cost. Figure 9 shows the non-marginal MVPF of the label at full coverage, with or without the corrective tax, over a range of potential implementation costs. In both cases, the non-marginal MVPF decreases with the implementation cost, as the information label generates roughly the same benefits over the increasing cost schedule. The figure spans three regions, defined by the points where the non-marginal MVPF curves drop below one.⁴³

Figure 9: Implementation thresholds of the fuel-economy label



Notes: Non-marginal MVPF of showing the fuel-economy label to every consumer, from $n = 0$ to $n = 1$, against the implementation cost per exposed consumer and vehicle-year, without tax (solid curve) and with a tax of \$265 per vehicle-year on the low-MPG car (dashed curve). The vertical lines mark the implementation costs at which each non-marginal MVPF equals 1, \$8.41 with the tax and \$18.69 without.

When the implementation cost is low enough (e.g., this is the first region of Figure 9), the policymaker would implement the label at full coverage. However, as the implementation cost increases, one or both policies can become inefficient: the non-marginal MVPF drops below one, and it is optimal for the policymaker to eliminate the reform. Policies with increasing returns behave fundamentally differently than policies with decreasing returns.

⁴³The choice of the threshold is arbitrary, some policymakers may want to set a different threshold, based for example on the MVPF of the best funding source for the policy.

Interior solutions are never optimal, and the policymaker alternates between corner solutions (e.g., full coverage or no coverage).

The same pattern is observed for public charging. Recall from [Figure 6](#) that some large municipalities exhibit both a high MVPF and increasing returns. The optimal program in the public charging application shares the same features as the fuel-economy program: it is optimal to increase public charging in counties with increasing returns, and eliminate public charging altogether in all other counties.⁴⁴

4.5 Trapezoidal approximations

As discussed in [Section 2.2](#), the sufficient statistics approach and our structural approach give comparable measures for the MVPF of marginal policies, evaluated at the observed policy environment. Computing the MVPF of non-marginal policies from causal estimates requires an extrapolation from the observed reform and additional assumptions on the policy elasticities.

We use our framework to assess the quality of the trapezoidal approximation proposed by [Kleven \(2021\)](#) and described in [Section 2.3](#). In the traditional MVPF framework, we would approximate the integrals below the willingness-to-pay and the cost curves over the policy space by extrapolating local policy elasticities linearly and calculating the area of the resulting trapezoids.

The results are presented in the last two columns of [Table 2](#). We provide two approximations. The first one uses the structural model to estimate the policy impact both ends of the reform, and compute the approximation described in [Section 2.3](#) exactly (see column 3). The second approximation extrapolates from the estimated policy impact at the observed outcome to the no-policy outcome (see column 4). We assume a constant policy elasticity, in line with what we would do if we had only access to causal estimates.⁴⁵

The trapezoidal approximation works remarkably well when it is based on structural estimates at both ends of a reform path. The mean absolute deviation is 1.4% (median: 0.07%)

⁴⁴The corner solution is a consequence of the absence of congestion effects in the demand model of [Fournel \(2026\)](#). One more station always raises the utility of electric vehicle owners, and does not depend on the installed base of electric vehicle owners. In this context, adding one charger in high-population, high-adoption counties raises the welfare of many consumers too much compared to the fixed cost of adding one charger, and leads to corner solutions where you invest in some counties without bounds, and divest everywhere else.

⁴⁵The policy elasticities at the observed outcome are the same in both cases, only the policy elasticities at the counterfactual outcome differs. We use the structural model to recover the policy elasticities at both ends of the reform. In the sufficient statistics approach, we would use instead an estimated causal effect at the policy and extrapolate from a constant elasticity.

and the maximum absolute deviation is 7.3% for battery electric vehicle subsidies. The extrapolation also provides a good approximation of the non-marginal MVPF in all applications. For the residential solar and the fuel-economy label cases, the constant elasticity assumption matches the specification of the structural model. The same is not true for the application to electric vehicles: the utility function is linear in prices and has the Box-Cox functional form in public charging stations, and thus, the associated elasticities are not constant over the policy space. Nevertheless, the approximation based on an extrapolation performs only slightly worse: the mean absolute deviation is 2.5% (median: 1.2%) and the maximum absolute deviation is 9.1% (for battery electric vehicle subsidies).

We test the quality of the trapezoidal approximation under different structural assumptions when extrapolating from the observed reform. When assuming a constant semi-elasticity rather than a constant elasticity (e.g., the utility is linear in prices rather than log-linear), we reduce the maximum absolute deviation from 9.1 to 7.5% of the non-marginal MVPF, which is similar to the approximation that uses the structural model at both ends of the reform (7.3%). Therefore, we highlight that extrapolating away from the observed outcomes using causal estimates, while not entirely assumption-free, may be quite robust to misspecification in our chosen empirical applications.


4.6 Optimal policy and the expansion path

We have established so far that, for policies with increasing returns, the optimal program is a corner solution at either end of the policy space (see Section 4.4). Among the empirical settings under study, this occurred for public goods (i.e., public charging) and informational nudges (i.e., fuel-economy label). For price instruments, we find instead evidence of diminishing returns, and an interior solution to the welfare problem may exist.

In what follows, we study an optimal program in the two empirical settings that include subsidies and taxes, that is, [Gillingham and Tsvetanov \(2019\)](#) on residential solar, and [Fournel \(2026\)](#) on electric vehicles. In both cases, we solve for two different types of optimum: the first preserves the observed government spending, so we are essentially solving for a more efficient policy mix, and the second finds both the optimal policy and the optimal government spending jointly.

Residential solar panels. We begin with the empirical application on residential solar subsidies, based on [Gillingham and Tsvetanov \(2019\)](#). There are two policy instruments: a state rebate per watt of installed capacity, and a federal tax credit applied to the net price of the system. We will assume throughout that a single policymaker decides on both

Figure 10: Optimal solar subsidies under a fixed budget




Notes: Solar subsidies in Connecticut, in the market of 2014. The grey curves are iso-MVPF curves, the blue curve is the budget of the observed policy, and the dashed red curve is the iso-MVPF curve through the best policy on that budget.

policies. While this is not an exact representation of the true policy state, it is still useful to understand what would be optimal under perfect coordination between the state and federal governments (and pooled government resources).

Since there are only two policy instruments, a clean graphical representation of the optimal policies is possible. Figure 10 presents the optimal policy mixes given the 2014 budget (panel (a)) and the 2021 budget (panel (b)). We represent the budget surface in blue, and the iso-MVPF curves in gray. The iso-MVPF curve of the optimum is the dashed red line. We find that the federal tax credit should be significantly lower on the 2021 budget, around 8%, and completely eliminated on the 2014 budget. The savings from the federal tax credit are reinvested in flat subsidies per watt (\$1.92 and \$1.03 per watt, respectively). Since both policies are price interventions on the same product, and only the delivery differs (e.g., flat versus *ad valorem*) the state program crowds out part of the federal program, and a slight advantage in targeting from either program (e.g., price versus size) can translate in one policy instrument dominating the other over a range of government spending, as is the case at the 2014 budget. Under the 2021 budget, the iso-MVPF curve at the optimum almost coincides with the budget surface, both policy are essentially equivalent, and we get an interior solution.

We now consider the problem of finding both the optimal policy mix and the optimal gov-

Figure 11: Optimal solar subsidies under a free budget



Notes: Solar subsidies in Connecticut, in the market of 2014. Panel (a) shows the expansion path $\theta(C)$, with the budgets of the 2014 and 2021 policies (dashed) and the budget at the optimal scale C^* (blue). Panel (b) shows the shadow value of public funds $\mu(C)$ along the path.

ernment spending, described in Section 2.4. We construct the expansion path, that is, the set of optimal reforms $\theta(C)$ over all possible total costs C , as the set tangency point between the iso-MVPF curves and the budget surfaces spanned by C . The result is presented in Figure 11, panel (a). We uncover an interesting pattern. Essentially, we can split the set of potential budgets in three sections, low budget (below \$11.5M), intermediate budget (\$11.5M to \$19.2M), and high budget (above \$19.2M). On the low budget segment, the *ad valorem* tax credit dominates and there is no per watt rebate, on the high budget segment, there is no *ad valorem* tax credit and the per watt rebate dominates, and we get an interior solution with both incentives in between. The best policy instrument changes as we increase the budget, and we move from one corner solution to the other.

Panel (b) of Figure 11 shows the shadow value of public funds, against the net public cost of the reform. This is the direct counterpart of the expansion path of panel (a). We depict the shadow value and the net public cost of the best policy achievable under the 2014 and the 2021 budgets, as well as the point where the shadow value is equal to one. We select this point to represent the best policy achievable under any budget constraint, e.g., the point at which one dollar spent generates one dollar of welfare. The associated policy sits at the kink where the *ad valorem* tax credit becomes inefficient. The total government spending reaches \$19.2M and the associated best policy is a \$1.52 per watt subsidy. Since installers keep about 16% of the rebate, this lowers the price paid by consumers by \$1.28 per watt, which

is the environmental damage avoided per watt installed. The optimal policy is therefore the Pigouvian subsidy discussed in Section 2.4. The federal credit is not used, because it is proportional to the price of each system, while the damage avoided per watt is the same for every system.

Electric vehicles. We repeat the same experiment for the empirical setting in Fournel (2026). We consider the two subsidies already in place, and two alternative policy instruments, a subsidy on non-rechargeable hybrids and a tax on gas guzzlers.⁴⁶ The results are presented in Table 3. The first column is the observed policy, the second column is the best policy on the same budget, and the third column is the best policy with the shadow value of public funds equal to one.

Table 3: Optimal policy mix and scale for electric vehicles

	Observed policy	Best policy, same budget	Best policy
<i>Panel A. Policy instruments</i>			
Battery electric rebate	13,000	12,221	13,566
Plug-in hybrid rebate	6,500	5,778	6,505
Hybrid rebate	0	4,440	5,280
Tax on gas guzzlers	0	14	0
<i>Panel B. MVPF and cost</i>			
MVPF	1.013	1.050	1.000
Non-marginal MVPF	1.257	1.281	1.234
Total cost (million)	264.5	264.5	323.1

Notes: Electric vehicles in Quebec in 2020, based on Fournel (2026). The best policy at the same budget keeps the total cost of the observed policy, and the best policy sets the MVPF of every instrument in use to one. Abated emissions are the lifetime emissions of the 2020 vehicles relative to no policy.

We want to highlight a few findings. Under the observed budget, the government should cut the rebates on battery electric vehicles and plug-in hybrids by about \$750 each and introduce a hybrid rebate of \$4,440 per vehicle. The reallocation equalizes the marginal MVPFs of the three rebates at 1.05 and raises total surplus by \$6.2 million. The tax on gas guzzlers is never implemented. Raising a dollar with the tax costs about \$1.05 of in benefits to , so the optimal tax under the observed budget is negligible, around \$14 per vehicle.⁴⁷

⁴⁶We keep public charging endogenous, following the network supply model described in Appendix D.1.1. With public charging as an additional instrument, the optimal policy would be trivial: remove all subsidies and invest on network without bounds.

⁴⁷Under lower government spending, it becomes efficient to tax gas guzzlers to raise additional funds and reinvest in subsidies.

We next evaluate the best policy jointly with the optimal government budget. The government should spend \$323 million, 22% more than in 2020, with all three rebates above their levels at the observed budget. The additional spending raises total surplus further by \$1.4 million, so most of the gain from the best policy comes from the hybrid rebate rather than from the level of spending.

5 Conclusion

In this paper, we propose a method to estimate the marginal value of public funds in structural models of demand. At the cost of additional assumptions on consumers' utility, our framework addresses several limitations of the sufficient-statistics approach. The structural model allows us to measure the MVPF at any point of the policy space, including policies that have never been implemented, to evaluate non-marginal reforms, and to find the optimal mix and scale of several policy instruments under a budget constraint. For marginal reforms, our estimates are consistent with, and comparable to, those of the sufficient-statistics approach, which remains the natural tool for small policy changes. Our method therefore extends the MVPF framework to any study based on a structural model, and we view this work as complementary to program-evaluation methods.

We apply the methodology to three studies of environmental policies: electric vehicles and public charging in Quebec, residential solar panels in Connecticut, and fuel-economy labels in the United States. For price instruments, such as subsidies to electric vehicles and solar panels, returns diminish with spending, and the welfare problem allows for an interior solution. For public goods and informational nudges, such as public charging and fuel-economy labels, returns do not diminish, and the optimal policy is a corner (e.g., a label with full coverage or no coverage). The decision to implement these policies then depends on the non-marginal MVPF of the whole program, rather than on the marginal MVPF.

Our study highlights the potential of structural demand models for studying a broader set of policies under the MVPF framework. However, this approach comes at the cost of structural assumptions on the economic problem, and conclusions about policies far from the status quo depend on these assumptions. Developing intermediate approaches that preserve the benefits of causal inference methods is an important direction for future research.

References

- Allcott, H. (2013). The welfare effects of misperceived product costs: Data and calibrations from the automobile market. *American Economic Journal: Economic Policy*, 5(3):30–66.
- Allcott, H., Cohen, D., Morrison, W., and Taubinsky, D. (2025). When do ”nudges” increase welfare? *American Economic Review*, 115(5):1555–96.
- Allcott, H. and Greenstone, M. (2024). Measuring the welfare effects of residential energy efficiency programs. Working Paper 23386, National Bureau of Economic Research.
- Allcott, H., Kane, R., Maydanchik, M. S., Shapiro, J. S., and Tintelnot, F. (2024). The effects of “buy american”: Electric vehicles and the inflation reduction act. Working Paper 33032, National Bureau of Economic Research.
- Allcott, H. and 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.
- Alé-Chilet, J., Chen, C., Li, J., and Reynaert, M. (2026). Colluding against environmental regulation. *The Review of Economic Studies*, 93(1):35–71.
- Armitage, S. and Pinter, F. (2026). Cross-state spillovers of regulation under national pricing strategies: Evidence from electric vehicles. Technical report, Technical report, 1-56. 6.
- Bento, A. M., Goulder, L. H., Jacobsen, M. R., and von Haefen, R. H. (2009). Distributional and efficiency impacts of increased us gasoline taxes. *American Economic Review*, 99(3):667–99.
- Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.
- Boning, W. C., Hendren, N., Sprung-Keyser, B., and Stuart, E. (2025). A welfare analysis of tax audits across the income distribution. *The Quarterly Journal of Economics*, 140(1):63–112.
- Buera, F. J., Kaboski, J. P., and Townsend, R. M. (2023). From micro to macro development. *Journal of Economic Literature*, 61(2):471–503.
- Chamberlain, G. (1980). Analysis of covariance with qualitative data. *Review of Economic Studies*, 47(1):225–238.
- Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1(1):451–488.
- Conlon, C. and Gortmaker, J. (2020). Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics*, 51(4):1108–1161.
- De Groote, O. and Verboven, F. (2019). Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems. *American Economic Review*, 109(6):2137–2172.

- Dorsey, J., Langer, A., and McRae, S. (2025). Fueling alternatives: Gas station choice and the implications for electric charging. *American Economic Journal: Economic Policy*, 17(1):362–400.
- Durrmeyer, I. (2022). Winners and losers: The distributional effects of the french feebate on the automobile market. *The Economic Journal*, 132(644):1414–1448.
- Durrmeyer, I. and Martínez, N. (2026). The welfare consequences of urban traffic regulations. Working paper, Toulouse School of Economics.
- Durrmeyer, I. and Samano, M. (2018). To rebate or not to rebate: Fuel economy standards versus feebates. *The Economic Journal*, 128(616):3076–3116.
- Environmental Protection Agency (2023). Epa report on the social cost of greenhouse gases: estimates incorporating recent scientific advances.
- Feger, F., Pavanini, N., and Radulescu, D. (2022). Welfare and redistribution in residential electricity markets with solar power. *The Review of Economic Studies*, 89(6):3267–3302.
- Finkelstein, A. and Hendren, N. (2020). Welfare analysis meets causal inference. *Journal of Economic Perspectives*, 34(4):146–67.
- Fournel, J.-F. (2026). Electric vehicle subsidies: Cost-effectiveness and emission reductions. Technical report, AEJ Microeconomics.
- Gaarder, I., Grindaker, M., Meling, T. G., and Mogstad, M. (2026). Green waste. Working Paper 34649, National Bureau of Economic Research.
- Gerarden, T. D., Bollinger, B., Gillingham, K., and Xu, D. Y. (2025). Strategic avoidance and the welfare impacts of U.S. solar panel tariffs. Working Paper 34401, National Bureau of Economic Research.
- Gillingham, K., Iskhakov, F., Munk-Nielsen, A., Rust, J., and Schjerning, B. (2022). Equilibrium trade in automobiles. *Journal of Political Economy*, 130(10):2534–2593.
- Gillingham, K. and Stock, J. H. (2018). The cost of reducing greenhouse gas emissions. *Journal of Economic Perspectives*, 32(4):53–72.
- Gillingham, K. and Tsvetanov, T. (2019). Hurdles and steps: Estimating demand for solar photovoltaics. *Quantitative Economics*, 10(1):275–310.
- Guo, X. and Xiao, J. (2023). Welfare analysis of the subsidies in the chinese electric vehicle industry. *The Journal of Industrial Economics*, 71(3):675–727.
- Hahn, R. W., Hendren, N., Metcalfe, R. D., and Sprung-Keyser, B. (2026). A welfare analysis of policies impacting climate change. *American Economic Review*, 116(7):2368–2421.
- Hahn, R. W. and Metcalfe, R. D. (2021). Efficiency and equity impacts of energy subsidies. *American Economic Review*, 111(5):1658–88.
- Harberger, A. C. (1964). The measurement of waste. *The American Economic Review*, 54(3):58–76.

- Heckman, J. J. (2010). Building bridges between structural and program evaluation approaches to evaluating policy. *Journal of Economic literature*, 48(2):356–398.
- Heckman, J. J. and Vytlačil, E. (2005). Structural equations, treatment effects, and econometric policy evaluation 1. *Econometrica*, 73(3):669–738.
- Heckman, J. J. and Vytlačil, E. J. (2007). Econometric evaluation of social programs, part i: Causal models, structural models and econometric policy evaluation. *Handbook of econometrics*, 6:4779–4874.
- Hendren, N. (2016). The policy elasticity. *Tax Policy and the Economy*, 30(1):51–89.
- Hendren, N. and Sprung-Keyser, B. (2020). A unified welfare analysis of government policies. *The Quarterly Journal of Economics*, 135(3):1209–1318.
- Hendren, N. and Sprung-Keyser, B. (2022). The case for using the mvpf in empirical welfare analysis. Working Paper 30029, National Bureau of Economic Research.
- Hendren, N., Sprung-Keyser, B., Stuart, E., et al. (2026). Policy impacts library. [online], Policy Impacts, Massachusetts Institute of Technology. Last accessed: 2026-10-14.
- Holland, S. P., Mansur, E. T., Muller, N. Z., and Yates, A. J. (2016). Are there environmental benefits from driving electric vehicles? the importance of local factors. *American Economic Review*, 106(12):3700–3729.
- Hong, Y., Kim, I. K., and Verboven, F. (2026). Which green technology to subsidize? evidence from electric vehicles in south korea. Technical report, arXiv preprint arXiv:2607.14446.
- Hoppe, J., Hinder, B., Rafaty, R., Patt, A., and Grubb, M. (2023). Three decades of climate mitigation policy: What has it delivered? *Annual Review of Environment and Resources*, 48(1):615–650.
- Interagency Working Group on Social Cost of Greenhouse Gases (2021). Technical support document: Social cost of carbon, methane, and nitrous oxide. interim estimates under executive order 13990. Technical report, United States Government, Washington, DC.
- Jacobsen, G. D. and Kotchen, M. J. (2013). Are building codes effective at saving energy? evidence from residential billing data in florida. *Review of Economics and Statistics*, 95(1):34–49.
- Jacobsen, M. R., Sallee, J. M., Shapiro, J. S., and van Benthem, A. A. (2023). Regulating untaxable externalities: Are vehicle air pollution standards effective and efficient?*. *The Quarterly Journal of Economics*, 138(3):1907–1976.
- Kang, Z. Y. and Vasserman, S. (2025). Robustness measures for welfare analysis. *American Economic Review*, 115(8):2449–2487.
- Kleven, H. J. (2021). Sufficient statistics revisited. *Annual Review of Economics*, 13(1):515–538.

- Kotchen, M. J. (2025). Taxing externalities: Revenue versus welfare gains with an application to us carbon taxes. *Review of Environmental Economics and Policy*, 19(1):25–47.
- Li, S., Linn, J., and Spiller, E. (2013). Evaluating “cash-for-clunkers”: Program effects on auto sales and the environment. *Journal of Environmental Economics and management*, 65(2):175–193.
- Li, S., Tong, L., Xing, J., and 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., and Sun, G. K. (2023). Judging nudging: Understanding the welfare effects of nudges versus taxes. Working Paper 31152, National Bureau of Economic Research.
- McFadden, D. (1999). Computing willingness-to-pay in random utility models. In Melvin, J. R., Moore, J. C., and Riezman, R., editors, *Trade, Theory and Econometrics: Essays in Honour of John S. Chipman*, pages 253–274. Routledge, London.
- Montpetit, S., Carrer, L., and Beaugard, P.-L. (2026). Beyond earnings: Availability and the non-pecuniary returns of universal childcare. Working paper.
- Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica*, 46(1):69–85.
- Nordhaus, W. (2019). Climate change: The ultimate challenge for economics. *American Economic Review*, 109(6):1991–2014.
- Prest, B. C. and Stock, J. H. (2023). Climate royalty surcharges. *Journal of Environmental Economics and Management*, 120:102844.
- Remmy, K. (2026). Adjustable product attributes, indirect network effects, and subsidy design: The case of electric vehicles. *American Economic Journal: Economic Policy*, 18(2):107–40.
- Reynaert, M. (2021). Abatement strategies and the cost of environmental regulation: Emission standards on the european car market. *The Review of Economic Studies*, 88(1):454–488.
- Semken, C. (2025). The marginal impact of emission reductions: Estimates, beliefs and behavior. Working paper, SIEPR.
- Small, K. A. and Rosen, H. S. (1981). Applied welfare economics with discrete choice models. *Econometrica: Journal of the Econometric Society*, pages 105–130.
- Springel, K. (2021). Network externality and subsidy structure in two-sided markets: Evidence from electric vehicle incentives. *American Economic Journal: Economic Policy*, 13(4):393–432.
- Timilsina, G. R. (2022). Carbon taxes. *Journal of Economic Literature*, 60(4):1456–1502.

- Todd, P. E. and Wolpin, K. I. (2023). The best of both worlds: Combining randomized controlled trials with structural modeling. *Journal of Economic Literature*, 61(1):41–85.
- Van Biesebroeck, J. and Verboven, F. (2026). Demand, competition, and public policy in the automobile industry. *Journal of Economic Literature*, 64(3):984–1033.
- Willig, R. D. (1976). Consumer’s surplus without apology. *American Economic Review*, 66(4):589–597.
- Wright, T. P. (1936). Factors affecting the cost of airplanes. *Journal of the Aeronautical Sciences*, 3(4):122–128.

A Appendix Tables and Figures

Table A1: Calibrated parameters, [Fournel \(2026\)](#)

Parameter	Value	Source
<i>PANEL A: Environmental externality</i>		
Social cost of carbon, in 2020	\$245/t	ECCC (2023)
Discount rate (damages and fiscal flows)	2%	ECCC (2023)
Expected vehicle lifetime	12.02 years	Fournel (2026)
Average yearly mileage	22,053 km	Fournel (2026)
Present-value lifetime mileage	238,230 km	Computed
Social cost of carbon, present value over vehicle's life	\$269/t	Computed
Tailpipe CO ₂ rating, sales-weighted mean		
Internal combustion vehicle	215 g/km	Fournel (2026)
Battery electric vehicles	0 g/km	Fournel (2026)
Plug-in hybrid vehicles	72 g/km	Fournel (2026)
Hybrid vehicles	134 g/km	Fournel (2026)
Lifetime damages per vehicle, sales-weighted		
Internal combustion vehicle	\$13,792	Computed
Battery electric vehicles	\$0	Computed
Plug-in hybrid vehicles	\$4,597	Computed
Hybrid vehicles	\$8,557	Computed
<i>PANEL B: Fiscal externalities</i>		
Corporate tax rate (τ_c)	27%	Fournel (2026)
Gasoline tax	\$0.192/litre	Fournel (2026)
Gasoline CO ₂ content	2,345 g/litre	Computed
Lifetime gas-tax revenue per vehicle, sales-weighted		
Internal combustion vehicle	\$4,197	Computed
Battery electric vehicles	\$0	Computed
Plug-in hybrid vehicles	\$1,399	Computed
Hybrid vehicles	\$2,605	Computed
Federal levy on fuel-inefficient vehicles	\$1,000 – \$4,000	Fournel (2026)
<i>PANEL C: Policy instruments, in 2020</i>		
Battery electric rebate	\$13,000	Fournel (2026)
Plug-in hybrid rebate	\$6,500	Fournel (2026)
Government spending	\$248.5M	Fournel (2026)
Charging stations installed	440	Fournel (2026)
Charging stations stock (end of 2019)	2,371	Fournel (2026)
Cost of a charging station	\$19,000	Assumption
<i>PANEL D: Estimated parameters</i>		
Marginal utility of income (λ_i), mean (std.dev.)	0.769 (0.148)	Fournel (2026)

Notes: Inputs of the MVPF calculations for the Quebec application, evaluated in 2020. All monetary values are reported in 2020 Canadian dollars. The social cost of carbon is the Environment and Climate Change Canada (ECCC) path, and damages and fiscal flows over the expected vehicle life are discounted at 2%. The externality covers tailpipe CO₂ only (i.e., excluding other pollutants and emissions from production), valued at the product level. The vehicle-type figures are sales-weighted averages over the 2020 registrations, and are reported as a reference.

Table A2: Calibrated parameters, [Gillingham and Tsvetanov \(2019\)](#)

Parameter	Value	Source
<i>PANEL A: Environmental externality</i>		
Social cost of carbon, in 2014	\$156/t	EPA (2023)
Social cost of carbon, present value over the system's lifetime	\$193/t	Computed
Discount rate (damages)	2%	EPA (2023)
System lifetime	25 years	Gillingham and Tsvetanov (2019)
Lifetime generation per watt installed	32.26 kWh	Gillingham and Tsvetanov (2019)
Marginal emission rate of displaced electricity	258 g CO ₂ /kWh	Gillingham and Tsvetanov (2019)
Lifetime avoided emissions per kilowatt installed	8.32 t	Computed
Present-value avoided damages per watt installed	\$1.282	Computed
<i>PANEL B: Fiscal externalities</i>		
Corporate tax rate, federal and state (τ_c)	39.9%	Statutory rates
Federal tax credit, share of the post-rebate price	30%	Gillingham and Tsvetanov (2019)
Lost electricity-tax revenue per watt installed	\$0.00	Assumption
<i>PANEL C: Policy instruments and market, in 2014</i>		
State rebate, average per watt	\$0.935	Gillingham and Tsvetanov (2019)
Post-rebate price per watt, installation-weighted mean	\$2.98	Gillingham and Tsvetanov (2019)
System size, mean	7.82 kW	Gillingham and Tsvetanov (2019)
Installations (model; observed 1,974)	2,012	Computed
Installed capacity (model; observed 15,429 kW)	15,733 kW	Computed
State rebate spending	\$14.7M	Computed
Federal tax credit spending	\$14.1M	Computed
<i>PANEL D: Estimated parameters</i>		
Passthrough of the rebate to households (ρ)	0.844 (0.045)	Gillingham and Tsvetanov (2019)
Price coefficient, any installation in the block group	-0.201 (0.109)	Estimated
Price coefficient, number of installations	-0.428 (0.133)	Estimated
Price elasticity of installed capacity	-0.95	Computed

Notes: Inputs of the MVPF calculations for the Connecticut application, evaluated in 2014, the last year of the estimation sample of [Gillingham and Tsvetanov \(2019\)](#). All monetary values are reported in 2014 U.S. dollars. The externality covers the CO₂ emissions avoided by displacing grid electricity over the 25-year life of a system, valued at the social cost of carbon used by [Gillingham and Tsvetanov \(2019\)](#) and discounted at 2%; the lifetime generation per watt and the marginal emission rate of displaced electricity are taken from their Section 7.3. The state rebate is the average rebate per watt paid in 2014 (their Table 9). The federal tax credit is 30% of the price net of the state rebate. Rebates are passed through to households at the rate estimated in their Table 8, and we assume that the federal credit accrues entirely to households. Installations, installed capacity, and spending on both instruments are model predictions at the observed 2014 policy, and are reported as a reference next to the observed values. The price coefficients are our estimates of the two parts of the hurdle demand model of [Gillingham and Tsvetanov \(2019\)](#), the probability that a block group has any installation and the number of installations, estimated with correlated random effects at the block-group level, with standard errors clustered at the town level in parentheses. The price elasticity of installed capacity is the corresponding model elasticity in 2014.

Table A3: Calibrated parameters, [Allcott et al. \(2025\)](#)

Parameter	Value	Source
<i>PANEL A: Environmental externality</i>		
Social cost of carbon	\$193/t	EPA (2023)
Gasoline CO ₂ content	8.87 kg/gallon	Allcott et al. (2025)
Fuel economy, low-MPG and high-MPG car	23 and 30 MPG	Allcott et al. (2025)
Average yearly mileage	10,803 miles	Allcott et al. (2025)
Additional gasoline of the low-MPG car, per year	110 gallons	Computed
Carbon damages of the low-MPG car, per vehicle-year	\$187.6	Computed
<i>PANEL B: Fiscal externalities</i>		
Corporate tax rate, federal and state (τ_c)	25.8%	Statutory rates
Federal gasoline tax	\$0.184/gallon	Statutory rate
State gasoline tax, sample mean	\$0.318/gallon	FHWA (2020)
Gasoline-tax revenue of the low-MPG car, per vehicle-year	\$54.6	Computed
<i>PANEL C: Policy instrument and market</i>		
Exposure to the label	0 to 100%	Policy experiment
Implementation cost per exposed consumer, per vehicle-year	\$0 – \$10	Assumption
Passthrough of a tax on the low-MPG car (ρ)	0.8	Allcott et al. (2025)
Markup	0	Allcott et al. (2025)
Market share of the low-MPG car without the label	20.9%	Computed
Change in the market share under the label	−5.4 p.p.	Computed
Participants (choice observations)	1,267 (2,534)	Allcott et al. (2025)
<i>PANEL D: Estimated parameters</i>		
Bias in the valuation of the low-MPG car (γ_i), mean	134.0 (16.9)	Estimated
Bias, standard deviation	381.6 (18.4)	Estimated
Share of consumers who respond to the label	0.549 (0.038)	Estimated
Effect of the label on responders (τ_i), mean	−59.9 (16.4)	Estimated
Effect of the label, slope on the bias	−0.382 (0.074)	Estimated
Effect of the label, slope on the participant taste (β_κ)	−0.447 (0.057)	Estimated
Effect of the label, residual standard deviation	149.9 (19.4)	Estimated
Slope of demand for the low-MPG car, per dollar	−0.00052	Computed

Notes: Inputs of the MVPF calculations for the fuel-economy label application, based on the cars experiment of [Allcott et al. \(2025\)](#). All monetary values are in US dollars per vehicle-year, and valuations are those of the low-MPG car relative to the high-MPG car, as elicited in the experiment. The externality covers tailpipe CO₂ only, valued at the social cost of carbon used by the authors; it is respondent-specific in the model (i.e., it depends on self-reported mileage), and the table reports sample means. The corporate tax rate combines the federal statutory rate of 21% with the average state rate. Gasoline taxes combine the federal rate with the rate of each respondent’s state of residence, taken from the Federal Highway Administration’s Highway Statistics (Table MF-205, rates as of December 31, 2019); these exclude general sales taxes that some states apply to gasoline. The policy is the share of the market exposed to the label, and the MVPF is evaluated at full exposure. The implementation cost is per consumer exposed to the label (i.e., not per consumer whose choice changes); since it is not observed, we report the MVPF over a range of values. The supply side follows the authors: a linear relative supply curve calibrated so that the passthrough of a tax on the low-MPG car is 0.8, with no markup. The estimated parameters come from our structural model of the experiment’s multiple price list responses, estimated by maximum likelihood, with standard errors clustered by participant in parentheses. The bias is the wedge between the decision and the experienced valuation of the low-MPG car, and the label shifts the decision valuation of the consumers who respond to it; the slope on the bias indicates how much of a respondent’s bias the label undoes. The slope of demand is the one implied by the model at the baseline price ([Allcott et al. \(2025\)](#) report −0.00060).

Table A4: MVPF decomposition (part 1)

Policy instrument	Willingness-to-pay					Net cost						MVPF
	Consumers	Internality correction	Shareholders	Environmental externality	Total	Transfer to inframarginals	Additional take-up	Public-good or nudge spending	Additional tax revenues	Other tax revenues	Total	
<i>Panel A. Electric vehicles, Fournel (2026), Quebec 2020</i>												
Marginal reforms:												
Battery electric rebate	1.085		0.284	0.670	2.040	1.000	0.887	0.014	0.105	-0.207	2.003	1.018
Plug-in hybrid rebate	1.071		0.274	0.299	1.643	1.000	0.648	0.014	0.101	-0.094	1.654	0.993
Hybrid rebate	1.003		0.274	-0.019	1.258	1.000	-0.045	-0.001	0.101	0.003	0.850	1.480
Tax on gas guzzlers	-1.039		-0.252	0.579	-0.712	-1.000	0.042	0.001	-0.093	-0.184	-0.680	1.047
Public charging stations	4.768		1.192	2.211	8.171	0.000	2.927	1.000	0.441	-0.688	4.174	1.957
Non-marginal reforms (individual reforms):												
Battery electric rebate	0.742		0.187	0.448	1.376	1.000	0.003	0.012	0.069	-0.138	1.085	1.269
Plug-in hybrid rebate	0.800		0.205	0.230	1.235	1.000	0.000	0.012	0.076	-0.072	1.008	1.225
Public charging stations	4.798		1.196	2.241	8.235	0.000	2.962	1.000	0.442	-0.697	4.216	1.953
Non-marginal reforms (full program):												
Provincial and federal rebates	0.753		0.191	0.394	1.338	1.000	0.000	0.013	0.071	-0.122	1.065	1.257
Provincial rebates alone	0.870		0.205	0.446	1.522	1.000	0.000	0.017	0.076	-0.138	1.079	1.410
Federal top-up given provincial rebates	0.885		0.252	0.481	1.618	1.000	0.460	0.012	0.093	-0.149	1.527	1.059
<i>Panel B. Residential solar, Gillingham and Tsvetanov (2019), Connecticut 2014</i>												
Marginal reforms:												
State rebate, 2014 level	0.591		0.117	0.332	1.039	1.000	0.470		0.077	0.253	1.139	0.912
Federal tax credit, 2014 level	1.000		0.038	0.552	1.590	1.000	0.799		0.025	0.000	1.773	0.897
State rebate, 2021 level	0.625		0.100	0.304	1.029	1.000	0.280		0.066	0.219	0.994	1.035
Federal tax credit, 2021 level	1.000		0.010	0.478	1.487	1.000	0.450		0.006	0.000	1.444	1.030
Non-marginal reforms (individual reforms):												
State rebate, 2014 level	0.527		0.094	0.277	0.898	1.000	0.214		0.062	0.226	0.925	0.970
Federal tax credit, 2014 level	0.839		0.028	0.412	1.280	1.000	0.301		0.019	0.000	1.282	0.998
State rebate, 2021 level	0.605		0.094	0.288	0.986	1.000	0.209		0.062	0.213	0.935	1.055
Federal tax credit, 2021 level	0.857		0.007	0.361	1.225	1.000	0.078		0.005	0.000	1.073	1.142
Non-marginal reforms (full program):												
2014 program (rebate and credit)	0.682		0.048	0.310	1.039	1.000	0.000		0.032	0.000	0.968	1.073
2021 program (rebate and credit)	0.799		0.022	0.329	1.149	1.000	0.000		0.014	0.000	0.986	1.166
2014 program to 2021 program	-0.738		-0.104	-0.387	-1.228	-1.000	-0.549		-0.069	-0.208	-1.272	0.966

Notes: Each entry is expressed in dollars of the mechanical term of the reform, that is, the transfer to inframarginal buyers for a price instrument and the spending for a public good or a nudge. Non-marginal reforms start from the no-policy counterfactual and are normalized by their full outlay. *Consumers* collects what accrues to households given their choices: the incidence of the transfer and of equilibrium price changes, the benefits from public goods (e.g., in Panel A, the value of the charging stations that local planners add in response to the rebates), and any direct experience effect of a nudge, which is set to zero in the label application. *Internality correction* is the wedge between decision and experienced utility times the behavioural response, and is therefore carried only by consumers whose choice changes. *Shareholders* are after-tax profits, and the corporate tax on profits appears under additional tax revenues. *Additional take-up* is the subsidy paid on induced purchases, and *public-good or nudge spending* includes, in Panel A, the stations induced by the rebates. *Other tax revenues* are gasoline taxes and, in Panel B, the interaction between the state rebate and the federal tax credit. A column is left empty when the application does not feature the channel. The MVPF is total willingness-to-pay over total net cost.

Table A5: MVPF decomposition (part 2)

Policy instrument	Willingness-to-pay					Net cost						MVPF
	Consumers	Internality correction	Shareholders	Environmental externality	Total	Transfer to inframarginals	Additional take-up	Public-good or nudge spending	Additional tax revenues	Other tax revenues	Total	
Panel C. Fuel-economy labels, <i>Allcott et al. (2025)</i> , United States 2019 experiment												
Marginal reforms:												
Fuel-economy label, cost of \$1	4.077	11.364	-3.025	10.726	23.142		0.000	1.000	-1.052	-3.112	5.164	4.482
Fuel-economy label, cost of \$5	0.815	2.273	-0.605	2.145	4.628		0.000	1.000	-0.210	-0.622	1.833	2.525
Fuel-economy label, cost of \$10	0.408	1.136	-0.303	1.073	2.314		0.000	1.000	-0.105	-0.311	1.416	1.634
Fuel-economy label with \$265 tax, cost of \$1	1.525	14.529	-1.131	8.547	23.470		11.810	1.000	-0.393	-2.483	15.687	1.496
Fuel-economy label with \$265 tax, cost of \$5	0.305	2.906	-0.226	1.709	4.694		2.362	1.000	-0.079	-0.497	3.937	1.192
Fuel-economy label with \$265 tax, cost of \$10	0.152	1.453	-0.113	0.855	2.347		1.181	1.000	-0.039	-0.248	2.469	0.951
Non-marginal reforms (individual reforms):												
Fuel-economy label, cost of \$1	4.731	11.193	-3.511	10.566	22.980		0.000	1.000	-1.221	-3.066	5.287	4.347
Fuel-economy label, cost of \$5	0.946	2.239	-0.702	2.113	4.596		0.000	1.000	-0.244	-0.613	1.857	2.475
Fuel-economy label, cost of \$10	0.473	1.119	-0.351	1.057	2.298		0.000	1.000	-0.122	-0.307	1.429	1.608
Fuel-economy label with \$265 tax, cost of \$1	1.909	13.942	-1.416	8.210	22.644		11.359	1.000	-0.492	-2.385	15.237	1.486
Fuel-economy label with \$265 tax, cost of \$5	0.382	2.788	-0.283	1.642	4.529		2.272	1.000	-0.098	-0.477	3.847	1.177
Fuel-economy label with \$265 tax, cost of \$10	0.191	1.394	-0.142	0.821	2.264		1.136	1.000	-0.049	-0.239	2.424	0.934

Notes: Each entry is expressed in dollars of the mechanical term of the reform, that is, the transfer to inframarginal buyers for a price instrument and the spending for a public good or a nudge. Non-marginal reforms start from the no-policy counterfactual and are normalized by their full outlay. *Consumers* collects what accrues to households given their choices: the incidence of the transfer and of equilibrium price changes, the benefits from public goods (e.g., in Panel A, the value of the charging stations that local planners add in response to the rebates), and any direct experience effect of a nudge, which is set to zero in the label application. *Internality correction* is the wedge between decision and experienced utility times the behavioural response, and is therefore carried only by consumers whose choice changes. *Shareholders* are after-tax profits, and the corporate tax on profits appears under additional tax revenues. *Additional take-up* is the subsidy paid on induced purchases, and *public-good or nudge spending* includes, in Panel A, the stations induced by the rebates. *Other tax revenues* are gasoline taxes and, in Panel B, the interaction between the state rebate and the federal tax credit. A column is left empty when the application does not feature the channel. The MVPF is total willingness-to-pay over total net cost.

B Proofs

The initial version of the proofs in Section B.2 to B.8 were obtained from a combination of the authors' work and prompts on Claude Code (model: Fable 5.1; effort level: Ultracode). The final version was verified and streamlined using Aristotle (version 2.1) and Claude Code (model: Opus 5.5; effort level: Max).

B.1 Derivation of equations (31) and (32)

In the discrete choice model of Section 3, the money-metric utility of consumer i is

$$\frac{V_i^*(\zeta_i, \theta)}{\lambda_i} = (1 - \tau_c) s_i \Pi(\mathbf{q}) + \frac{1}{\lambda_i} \ln \left(1 + \sum_{k=1}^J \exp(\delta_{ik}) \right) - \frac{d_i(\mathbf{q})}{\lambda_i} + \frac{\gamma_i}{\lambda_i},$$

where $\delta_{ik} = x_k \beta_i - \lambda_i(p_k - \tau_k) + \xi_k$. The policy moves this expression through the mean utilities δ_{ik} , through profits and through the externality. Prices $p(\theta)$ and demand $\mathbf{q}(\theta)$ are the equilibrium objects, which we assume locally unique and differentiable in θ , and $\partial/\partial\tau_j$ lets all prices adjust. The derivation holds for any form of conduct. Conduct matters only for the passthrough rates $\partial p_k/\partial\tau_j$.

Consumers. The mean utilities respond as $\partial\delta_{ik}/\partial\tau_j = \lambda_i(\mathbb{1}\{k=j\} - \partial p_k/\partial\tau_j)$: consumers gain the subsidy on product j and bear the change in every price. The derivative of the log-sum with respect to δ_{ik} is the choice probability P_{ik} (Small and Rosen, 1981), so

$$\frac{1}{\lambda_i} \frac{\partial}{\partial\tau_j} \ln \left(1 + \sum_{k=1}^J \exp(\delta_{ik}) \right) = P_{ij} - \sum_{k=1}^J P_{ik} \frac{\partial p_k}{\partial\tau_j}. \quad (36)$$

This is the envelope theorem of Section 2.1: the reallocation of consumers across products has no first-order effect on their utility.⁴⁸ Summing over consumers, with $q_k = \sum_i P_{ik}$, gives the savings of inframarginal consumers, $q_j - \sum_k q_k \partial p_k/\partial\tau_j$.

Shareholders and externality. With constant marginal costs,

$$\frac{\partial \Pi}{\partial\tau_j} = \sum_k \left[q_k \frac{\partial p_k}{\partial\tau_j} + (p_k - c_k) \frac{\partial q_k}{\partial\tau_j} \right],$$

⁴⁸The result does not need the Gumbel distribution. For any distribution of taste shocks with a density that does not depend on the policy, the derivative of the expected maximum utility with respect to δ_{ik} is P_{ik} (the Williams–Daly–Zachary theorem, see McFadden, 1981). Income, including profits, enters utility linearly with a marginal utility λ_i that does not depend on the policy, so there are no income effects and demand does not depend on Π .

and consumers receive the share $(1 - \tau_c)s$ of this change, $s = \sum_i s_i$. Since $d_i(\mathbf{q}) = \sum_k q_k d_{ik}(e_k)$, the money-metric change in the externality is $\sum_k D_k \partial q_k / \partial \tau_j$, with $D_k = \sum_i d_{ik} / \lambda_i$. The envelope theorem does not apply to these two terms, because consumers do not internalize the effect of their choices on profits or on the externality. Adding the three terms gives equation (31).

Demand responses. Since $\partial P_{ik} / \partial \delta_{i\ell} = P_{ik}(\mathbb{1}\{k = \ell\} - P_{i\ell})$,

$$\frac{\partial q_k}{\partial \tau_j} = \sum_i \lambda_i P_{ik} \left[(\mathbb{1}\{k = j\} - P_{ij}) - \sum_{\ell=1}^J (\mathbb{1}\{k = \ell\} - P_{i\ell}) \frac{\partial p_\ell}{\partial \tau_j} \right]. \quad (37)$$

The first term is the response at fixed prices and the second the response to the price changes. We compute the passthrough rates numerically, by solving the pricing equilibrium at $\tau_j \pm h$. This is the only step that uses the firms' conduct.

Net cost. Differentiating the cost function (30), in which $\Pi(\mathbf{q}(0))$ is a constant, gives equation (32): the mechanical transfer q_j , the subsidies paid on the demand responses, and the change in corporate tax revenue. ■

B.2 Preliminaries for the optimal policy results

Notation. Without loss of generality, we normalize $\mathcal{W}(0) = 0$ and $\mathcal{C}(0) = 0$ in what follows, to simplify the exposition. Under this normalization, the non-marginal MVPF is $\overline{MVPF}(\theta) = \mathcal{W}(\theta) / \mathcal{C}(\theta)$ and the marginal MVPF of instrument j is $MVPF_j = (\partial \mathcal{W} / \partial \theta_j) / (\partial \mathcal{C} / \partial \theta_j)$. The value of public spending is the welfare of an optimally spent budget,

$$\mathcal{V}(C) = \max_{\theta \in \Theta} \{ \mathcal{W}(\theta) : \mathcal{C}(\theta) \leq C \}, \quad 0 \leq C \leq C_{\max} = \max_{\theta \in \Theta} \mathcal{C}(\theta), \quad (38)$$

and we write $\theta^*(C)$ for a policy that attains the optimum at C . At the budget $\bar{\mathcal{C}}$, this is the welfare problem (19). Its constraint $\mathcal{C}(\theta) \leq \bar{\mathcal{C}}$ has a multiplier $\mu \geq 0$, and stationarity reads $\nabla \mathcal{W} = \mu \nabla \mathcal{C}$. The constraint $\mathcal{C}(\theta) \geq \bar{\mathcal{C}}$ of the MVPF problem (20) has a multiplier $\nu \geq 0$, and stationarity reads $\nabla \overline{MVPF} + \nu \nabla \mathcal{C} = 0$. We say that welfare is *satiated below* $\bar{\mathcal{C}}$ when no solution of the welfare problem costs exactly $\bar{\mathcal{C}}$.

Assumptions. We maintain the assumptions of Section 2.4: (1) Θ is convex and compact and contains $\theta = 0$, (2) $\mathcal{W}$ and $\mathcal{C}$ are continuously differentiable, (3) the welfare weights are $\psi_i = 1/\lambda_i$, and (4) $\bar{\mathcal{C}} > 0$. Some results also use the following assumption, and we say when they do.

Assumption D (diminishing returns). *The value of public spending $\mathcal{V}$ is concave on $[0, C_{\max}]$.*

Lemma 1 (Gradient of the non-marginal MVPF). *At any policy with $\mathcal{C}(\theta) \neq 0$,*

$$\nabla \overline{MVPF}(\theta) = \frac{\nabla \mathcal{W}(\theta) - \overline{MVPF}(\theta) \nabla \mathcal{C}(\theta)}{\mathcal{C}(\theta)}. \quad (39)$$

Proof. This is the quotient rule applied to $\mathcal{W}/\mathcal{C}$. ■

Component by component, $\partial \overline{MVPF} / \partial \theta_j = (\partial \mathcal{C} / \partial \theta_j) [\overline{MVPF}_j - \overline{MVPF}] / \mathcal{C}$. When costs are positive, expanding instrument j raises the non-marginal MVPF if and only if the marginal dollar returns more than the average dollar.

Lemma 2 (Value of public spending).

(i) $\mathcal{V}$ is non-decreasing, and $\mathcal{V}(0) \geq 0$.

(ii) (Envelope.) *Let $\theta^*(C)$ be an interior solution at which the constraint binds and $\nabla \mathcal{C} \neq 0$. Then $\nabla \mathcal{W} = \mu(C) \nabla \mathcal{C}$ at $\theta^*(C)$ for a unique $\mu(C) \geq 0$, so $\overline{MVPF}_j(\theta^*(C)) = \mu(C)$ for every instrument with a non-zero marginal cost. Moreover, $\mu(C)$ lies between the left and right derivatives of $\mathcal{V}$ at C ,*

$$\limsup_{h \rightarrow 0^-} \frac{\mathcal{V}(C+h) - \mathcal{V}(C)}{h} \leq \mu(C) \leq \liminf_{h \rightarrow 0^+} \frac{\mathcal{V}(C+h) - \mathcal{V}(C)}{h}. \quad (40)$$

If $\mathcal{V}$ is concave near C , as under Assumption D, then $\mathcal{V}$ is differentiable at C and $\mathcal{V}'(C) = \mu(C)$.

(iii) *If $\mathcal{V}$ is concave on $[0, \tilde{C}]$, the average return $\mathcal{V}(C)/C$ is non-increasing on $(0, \tilde{C}]$, and $\mathcal{V}'(C) \leq \mathcal{V}(C)/C$ wherever $\mathcal{V}$ is differentiable.*

(iv) (Curvature along a path.) *Let $\theta(t)$ be a twice differentiable path along which $\mathcal{W}' > 0$ and $\mathcal{C}' > 0$. The marginal MVPF along the path, $\mathcal{W}'/\mathcal{C}'$, is non-increasing at t if and only if $\mathcal{W}''/\mathcal{W}' \leq \mathcal{C}''/\mathcal{C}'$ at t . If this holds along the whole path, welfare is a concave function of cost along the path.*

Proof.

(i) A larger budget enlarges the feasible set, and $\theta = 0$ is feasible at $C = 0$.

(ii) Since $\theta^*(C)$ is interior and $\nabla \mathcal{C} \neq 0$, the Kuhn–Tucker conditions hold, with a multiplier that is unique because $\nabla \mathcal{C} \neq 0$. Dividing by $\partial \mathcal{C} / \partial \theta_j \neq 0$ gives $\overline{MVPF}_j = \mu(C)$. For (40),

move along the line $\theta(t) = \theta^*(C) + t \nabla \mathcal{C} / \|\nabla \mathcal{C}\|^2$. Its cost is $c(t) = C + t + o(t)$ and its welfare $\mathcal{V}(C) + \mu(C)t + o(t)$. Since $\theta(t)$ is feasible at the budget $c(t)$, its welfare is at most $\mathcal{V}(c(t))$. For $t > 0$, dividing by $c(t) - C > 0$ and letting $t \rightarrow 0$ gives the right-hand inequality. For $t < 0$, the division by $c(t) - C < 0$ reverses the inequality and gives the left-hand one. A concave function has a left derivative at least as large as its right derivative, so under concavity (40) forces both to equal $\mu(C)$.

(iii) For $0 < C_a < C_b \leq \tilde{C}$, concavity and $\mathcal{V}(0) \geq 0$ give $\mathcal{V}(C_a) \geq (C_a/C_b) \mathcal{V}(C_b) + (1 - C_a/C_b) \mathcal{V}(0) \geq (C_a/C_b) \mathcal{V}(C_b)$. Similarly, concavity gives $\mathcal{V}(0) \leq \mathcal{V}(C) - C \mathcal{V}'(C)$, so $\mathcal{V}'(C) \leq \mathcal{V}(C)/C$.

(iv) The derivative of $\mathcal{W}'/\mathcal{C}'$ is $(\mathcal{W}''\mathcal{C}' - \mathcal{W}'\mathcal{C}'')/\mathcal{C}'^2$, which is non-positive if and only if $\mathcal{W}''/\mathcal{W}' \leq \mathcal{C}''/\mathcal{C}'$. Since $\mathcal{C}$ is strictly increasing along the path, welfare along the path is a differentiable function of cost, with derivative $\mathcal{W}'/\mathcal{C}'$. If this derivative is non-increasing in t , it is non-increasing in cost, and the function is concave. ■

On Assumption D. In our applications, welfare and costs are both convex in the policy. The log-sum is convex in the subsidies at given producer prices, and a subsidy is paid on a base that grows with the subsidy. Concave welfare and convex costs are sufficient for Assumption D, but too strong for our purposes.⁴⁹ What matters is the curvature of welfare relative to costs along the expansion path, the path of $\theta^*(C)$ as C grows. Along that path, welfare as a function of cost is $\mathcal{V}$ itself. By part (iv), when the expansion path is smooth with rising costs, Assumption D holds on a range of budgets if and only if the marginal MVPF falls as spending rises along the path in that range. The same condition along another path, e.g., one instrument moving from zero to its observed level, describes the returns along that path only.

B.3 Proof of Proposition 1

The condition of part (b) is that no larger budget, spent optimally, has a higher average return,

$$\frac{\mathcal{V}(C)}{C} \leq \frac{\mathcal{V}(\bar{C})}{\bar{C}} \quad \text{for every } C \in [\bar{C}, C_{\max}]. \quad (41)$$

It holds under Assumption D, by part (iii) of Lemma 2. It also holds for an S-shaped $\mathcal{V}$ at any budget beyond the one at which the average return peaks.

⁴⁹Let θ_a and θ_b attain $\mathcal{V}(C_a)$ and $\mathcal{V}(C_b)$. For $\alpha \in [0, 1]$, convexity of $\mathcal{C}$ makes $\alpha \theta_a + (1 - \alpha) \theta_b$ feasible at the budget $\alpha C_a + (1 - \alpha) C_b$, and concavity of $\mathcal{W}$ gives it a welfare of at least $\alpha \mathcal{V}(C_a) + (1 - \alpha) \mathcal{V}(C_b)$.

Proposition 1 (restated).

- (a) Let θ^* be an interior solution of (20) at which the spending constraint binds, $\mathcal{C}(\theta^*) = \bar{\mathcal{C}}$, and at which every instrument has a non-zero marginal cost. Suppose that welfare is not satiated below $\bar{\mathcal{C}}$. Then θ^* solves the welfare problem (19) and satisfies (18), with

$$\mu = \overline{MVPF}(\theta^*) - \nu \bar{\mathcal{C}} \geq 0,$$

where $\nu \geq 0$ is the multiplier of the spending constraint of (20). Under Assumption D, $\mu = \mathcal{V}'(\bar{\mathcal{C}})$.

- (b) Conversely, let $\tilde{\theta}$ be a solution of (19) at which the budget constraint binds, and suppose that (41) holds. Then $\tilde{\theta}$ solves (20), and (41) is also necessary for this. If $\tilde{\theta}$ is in addition interior with non-zero marginal costs, it satisfies (18).

Part (a). The idea is that Lemma 1 turns the first-order condition of the MVPF problem into that of the welfare problem, and that the absence of satiation turns stationarity into optimality. Since $\nabla \mathcal{C}(\theta^*) \neq 0$, the first-order condition of the MVPF problem holds, $\nabla \overline{MVPF}(\theta^*) + \nu \nabla \mathcal{C}(\theta^*) = 0$. With Lemma 1 and $\mathcal{C}(\theta^*) = \bar{\mathcal{C}}$, it reads

$$\nabla \mathcal{W}(\theta^*) = [\overline{MVPF}(\theta^*) - \nu \bar{\mathcal{C}}] \nabla \mathcal{C}(\theta^*). \quad (42)$$

Dividing by $\partial \mathcal{C} / \partial \theta_j \neq 0$ gives $MVPF_j(\theta^*) = \mu$ for every j , with $\mu = \overline{MVPF}(\theta^*) - \nu \bar{\mathcal{C}}$. It remains to show that θ^* solves the welfare problem. Since welfare is not satiated below $\bar{\mathcal{C}}$, some policy $\tilde{\theta}$ attains $\mathcal{V}(\bar{\mathcal{C}})$ at a cost of exactly $\bar{\mathcal{C}}$. This policy is feasible in the MVPF problem, so $\mathcal{W}(\theta^*) / \bar{\mathcal{C}} \geq \mathcal{W}(\tilde{\theta}) / \bar{\mathcal{C}} = \mathcal{V}(\bar{\mathcal{C}}) / \bar{\mathcal{C}}$. Hence θ^* attains $\mathcal{V}(\bar{\mathcal{C}})$ and solves the welfare problem. Its Kuhn–Tucker multiplier is non-negative and, by (42), equal to μ , so $\mu \geq 0$. Under Assumption D, $\mu = \mathcal{V}'(\bar{\mathcal{C}})$ by part (ii) of Lemma 2.

Remark 1. The two hypotheses of part (a) are needed, and $\mu \geq 0$ is not a substitute for the absence of satiation.

- (1) Let $\mathcal{W}(\theta) = \sum_{j=1,2} (2\theta_j - \theta_j^2)$ on $[0, 2]^2$, $\mathcal{C}(\theta) = \theta_1 + \theta_2$ and $\bar{\mathcal{C}} = 3$. The MVPF problem is solved at $\theta^* = (1.5, 1.5)$, where $MVPF_j = -1$. The welfare problem is solved at $(1, 1)$, which costs 2 and gives a welfare of 2 against 1.5. Welfare is satiated below the budget, and θ^* only maximizes welfare on the budget surface (part (i) of Proposition 2).

(2) Let $\mathcal{C}(\theta) = \theta$ and $\mathcal{W}(\theta) = 6\theta - 4.5\theta^2 + \theta^3$ on $[0, 2.3]$, with $\bar{\mathcal{C}} = 2.1$. The MVPF problem is solved at $\theta^* = 2.1$, where $\mu = \mathcal{W}'(2.1) = 0.33 > 0$. The welfare problem is solved at $\theta = 1$, with a welfare of 2.5 against 2.016. Here $\mathcal{V}$ is concave, so Assumption D holds as well. What fails is the absence of satiation.

Remark 2 (Satiation and the spending constraint). Let an interior $\hat{\theta}$ maximize welfare on the budget surface, with $\nabla \mathcal{C}(\hat{\theta}) \neq 0$ and $\nabla \mathcal{W}(\hat{\theta}) = \mu \nabla \mathcal{C}(\hat{\theta})$. If $\mu < 0$, welfare is satiated below the budget: moving along $-\nabla \mathcal{C}$ lowers the cost and raises welfare. If $\mu \geq 0$, $\mathcal{W}$ is concave and $\mathcal{C}$ is convex, $\hat{\theta}$ solves the welfare problem (as in the proof of Corollary 1). Conversely, the spending constraint of the MVPF problem binds when $\mathcal{V}(\bar{\mathcal{C}})$ is attained at a cost of $\bar{\mathcal{C}}$ and the average return falls strictly beyond $\bar{\mathcal{C}}$, since a policy of cost $C > \bar{\mathcal{C}}$ has a non-marginal MVPF of at most $\mathcal{V}(C)/C < \mathcal{V}(\bar{\mathcal{C}})/\bar{\mathcal{C}}$. Under Assumption D and without satiation, it is slack only under constant returns from the first dollar, $\mathcal{V}(C) = kC$ on $[0, C_1]$ for some $C_1 > \bar{\mathcal{C}}$.⁵⁰ Then $\nu = 0$, and a solution solves the welfare problem at its own cost.

Part (b). Let $\tilde{\theta}$ solve the welfare problem at a cost of exactly $\bar{\mathcal{C}}$. It is feasible in the MVPF problem. For any feasible policy θ , with cost $C \geq \bar{\mathcal{C}}$,

$$\overline{MVPF}(\theta) = \frac{\mathcal{W}(\theta)}{C} \leq \frac{\mathcal{V}(C)}{C} \leq \frac{\mathcal{V}(\bar{\mathcal{C}})}{\bar{\mathcal{C}}} = \overline{MVPF}(\tilde{\theta}),$$

by the definition of $\mathcal{V}$ and by (41). Hence $\tilde{\theta}$ solves the MVPF problem, and part (ii) of Lemma 2 gives (18) when $\tilde{\theta}$ is interior with non-zero marginal costs. Condition (41) is also necessary. If it fails at some $C > \bar{\mathcal{C}}$, a policy θ_C that attains $\mathcal{V}(C)$ has a higher non-marginal MVPF than $\tilde{\theta}$. It is feasible in the MVPF problem, since a cost below $\bar{\mathcal{C}}$ would imply $\mathcal{V}(C) \leq \mathcal{V}(\bar{\mathcal{C}})$ and hence $\mathcal{V}(C)/C \leq \mathcal{V}(\bar{\mathcal{C}})/\bar{\mathcal{C}}$, contradicting the failure of (41). ■

B.4 Marginal and non-marginal MVPF at the optimum

In part (a) of Proposition 1, $MVPF_j(\theta^*) = \overline{MVPF}(\theta^*) - \nu \bar{\mathcal{C}} \leq \overline{MVPF}(\theta^*)$, since $\nu \geq 0$ and $\bar{\mathcal{C}} > 0$. The value of public spending gives the inequality its meaning. Under Assumption D, $\mu = \mathcal{V}'(\bar{\mathcal{C}})$ is the return on the marginal dollar and $\overline{MVPF}(\theta^*) = \mathcal{V}(\bar{\mathcal{C}})/\bar{\mathcal{C}}$ the return on the average dollar, and part (iii) of Lemma 2 states that the first is at most the second. The multiplier of the MVPF problem is the rate at which the average return falls with the budget, $\nu = -d[\mathcal{V}(C)/C]/dC$ at $C = \bar{\mathcal{C}}$.

⁵⁰If $\mathcal{V}(C)/C = k$ at $\bar{\mathcal{C}}$ and at $C_1 > \bar{\mathcal{C}}$, the average return, which is non-increasing, is constant on $[\bar{\mathcal{C}}, C_1]$, so $\mathcal{V}$ is linear there. A concave function lies below the extension of any of its linear pieces, so $\mathcal{V}(C) \leq kC$ on $(0, \bar{\mathcal{C}}]$, while $\mathcal{V}(C)/C \geq k$ there. Hence $\mathcal{V}(C) = kC$ on $(0, C_1]$, and $\mathcal{V}(0) = 0$.

The inequality is strict if $\mathcal{V}(0) > 0$, e.g., when a budget-neutral combination of taxes and subsidies raises welfare, or if $\mathcal{V}$ is strictly concave on $[0, \bar{\mathcal{C}}]$. An interior solution is not enough.⁵¹ ■

When the inequality reverses. Since the derivative of the average return is $[\mathcal{V}'(C) - \mathcal{V}(C)/C]/C$, the marginal MVPF exceeds the non-marginal MVPF at a budget C if and only if the average return is still rising at C . This requires $\mathcal{V}$ to be non-concave on $[0, C]$. For instance, with an S-shaped $\mathcal{V}$ and $\mathcal{V}(0) = 0$, the reversal holds from the first dollar up to the budget at which a ray from the origin is tangent to $\mathcal{V}$, beyond the inflection point. At an optimal mix, Lemma 1 gives $\nabla \overline{MVPF} \cdot v = (\nabla \mathcal{C} \cdot v) [\mu - \overline{MVPF}]/C$, so the non-marginal MVPF then rises along any direction v that raises costs.

B.5 Proof of Proposition 2

Let $\mathcal{C} = \{\theta \in \Theta : \mathcal{C}(\theta) = \bar{\mathcal{C}}\}$ be the budget surface. At a point of $\mathcal{C}$ where $\nabla \mathcal{C} \neq 0$, the level curve of a function f is tangent to $\mathcal{C}$ when f is stationary along $\mathcal{C}$, i.e., when ∇f is proportional to $\nabla \mathcal{C}$.

Proposition 2 (restated).

- (i) *On $\mathcal{C}$, the non-marginal MVPF is an increasing affine transformation of welfare, $\overline{MVPF}(\theta) = \mathcal{W}(\theta)/\bar{\mathcal{C}}$. The two objectives therefore have the same maximizers on $\mathcal{C}$.*
- (ii) *At any point of $\mathcal{C}$ where $\nabla \mathcal{C} \neq 0$, the iso-MVPF curve is tangent to $\mathcal{C}$ if and only if the iso-welfare curve through the same point is tangent to $\mathcal{C}$.*
- (iii) *With at least two instruments, the non-marginal MVPF is a transformation of $\mathcal{W}$ on an open set of policies with non-zero costs only if the marginal MVPFs are equalized at every point of that set where $\mathcal{W} \neq 0$ and $\nabla \mathcal{C} \neq 0$.*

Part (i). On the budget surface $\mathcal{C}$, $\overline{MVPF}(\theta) = \mathcal{W}(\theta)/\bar{\mathcal{C}}$, an increasing affine function of welfare, so the two have the same maximizers on $\mathcal{C}$.

Part (ii). At $\theta \in \mathcal{C}$ with $\nabla \mathcal{C}(\theta) \neq 0$, tangency of the level curve of f means $\nabla f(\theta) \cdot v = 0$ for every direction v with $\nabla \mathcal{C}(\theta) \cdot v = 0$. For such directions, Lemma 1 gives $\nabla \overline{MVPF}(\theta) \cdot v = \nabla \mathcal{W}(\theta) \cdot v / \bar{\mathcal{C}}$, so one vanishes for all of them if and only if the other does.

⁵¹With $\mathcal{W}(\theta) = 2\sqrt{\theta_1\theta_2}$ and $\mathcal{C}(\theta) = \theta_1 + \theta_2$, the optimum is interior at every budget, $\mathcal{V}(C) = C$, and $MVPF_j = \overline{MVPF} = 1$.

Part (iii). Let U be an open set of policies with non-zero costs, and let $\theta \in U$ with $\mathcal{W}(\theta) \neq 0$, $\nabla \mathcal{C}(\theta) \neq 0$, and $\nabla \mathcal{W}(\theta)$ not proportional to $\nabla \mathcal{C}(\theta)$. Then $\nabla \mathcal{W}(\theta) \neq 0$, and there is a direction v with $\nabla \mathcal{W}(\theta) \cdot v = 0$ and $\nabla \mathcal{C}(\theta) \cdot v \neq 0$. By the implicit function theorem, the iso-welfare surface through θ contains a curve with tangent v at θ , along which the cost strictly changes. Two nearby points on this curve have the same welfare, different costs and therefore different non-marginal MVPFs, so $\overline{MVPF}$ is not a function of $\mathcal{W}$ on U . Hence, if $\overline{MVPF}$ is a transformation of $\mathcal{W}$ on U , the gradients of $\mathcal{W}$ and $\mathcal{C}$ are proportional wherever $\mathcal{W} \neq 0$ and $\nabla \mathcal{C} \neq 0$, i.e., the marginal MVPFs are equalized. For instance, with $\mathcal{W}(\theta) = 2\sqrt{\theta_1} + \sqrt{\theta_2}$ and $\mathcal{C}(\theta) = \theta_1 + \theta_2$, the marginal MVPFs are $1/\sqrt{\theta_1}$ and $1/(2\sqrt{\theta_2})$. They are equal only on the ray $\theta_1 = 4\theta_2$, and the iso-welfare and iso-MVPF curves cross everywhere else. ■

B.6 Derivation of equation (22), and averages along a path

Marginal reforms. Let v be a direction with $\nabla \mathcal{C} \cdot v \neq 0$, and let $\partial \mathcal{C} / \partial \theta_j \neq 0$ for every instrument with $v_j \neq 0$. Writing $\partial \mathcal{W} / \partial \theta_j = MVPF_j \partial \mathcal{C} / \partial \theta_j$,

$$\frac{\nabla \mathcal{W} \cdot v}{\nabla \mathcal{C} \cdot v} = \frac{\sum_j MVPF_j (\partial \mathcal{C} / \partial \theta_j) v_j}{\sum_k (\partial \mathcal{C} / \partial \theta_k) v_k} = \sum_j \omega_j MVPF_j,$$

with the weights of (22), which sum to one. If $\nabla \mathcal{W} = \mu \nabla \mathcal{C}$, the MVPF of the reform is μ in every direction with a non-zero cost.

Non-marginal reforms. Let $\theta(t)$, $t \in [0, 1]$, be a continuously differentiable path with $\mathcal{C}'(t) \neq 0$ and $\Delta \mathcal{C} = \mathcal{C}(1) - \mathcal{C}(0) \neq 0$, and let $\omega(t) = \mathcal{C}'(t) / \Delta \mathcal{C}$, which integrates to one. Since $\omega(t) \mathcal{W}'(t) / \mathcal{C}'(t) = \mathcal{W}'(t) / \Delta \mathcal{C}$,

$$\frac{\mathcal{W}(1) - \mathcal{W}(0)}{\mathcal{C}(1) - \mathcal{C}(0)} = \int_0^1 \omega(t) \frac{\mathcal{W}'(t)}{\mathcal{C}'(t)} dt. \quad (43)$$

The non-marginal MVPF is the average of the marginal MVPFs along the path, weighted by the share of the cost incurred at each point. It does not depend on the path, because its numerator and denominator do not.

Trapezoid rule. If $f(t) = a + bt + ct^2$, then $f(1) - f(0) = b + c = [f'(0) + f'(1)]/2$. The trapezoid rule of Section 2.3 is therefore exact when welfare and costs are quadratic along the path. ■

B.7 Proof of Corollary 1

Corollary 1 (restated). *Let the expansion path $\mathcal{E}$ be the set of interior solutions $\theta^*(C)$ of the welfare problem (19) at budgets $C > 0$. Let $\tilde{\theta}$ be an interior policy at which $\partial\mathcal{W}/\partial\theta_j$ and $\partial\mathcal{C}/\partial\theta_j$ are non-zero and of the same sign for every j , with $\overline{MVPF}(\tilde{\theta}) \neq 0$ and $\mathcal{C}(\tilde{\theta}) \neq 0$.*

- (a) *If $\tilde{\theta}$ lies on the expansion path, the iso-MVPF curve and the iso-welfare curve through $\tilde{\theta}$ are tangent at $\tilde{\theta}$.*
- (b) *Conversely, if the two curves are tangent at $\tilde{\theta}$, then $\tilde{\theta}$ satisfies (18) at the budget $\mathcal{C}(\tilde{\theta})$, with a positive shadow value of public funds. Whenever this condition is sufficient, e.g., if $\mathcal{W}$ is concave and $\mathcal{C}$ is convex, $\tilde{\theta}$ lies on the expansion path.*

The expansion path is therefore contained in the locus of tangency points between the two families of level curves, and it coincides with that locus under the sufficiency condition of part (b).

Let $C = \mathcal{C}(\tilde{\theta})$. The two level curves through $\tilde{\theta}$ are tangent when $\nabla\overline{MVPF}(\tilde{\theta}) = \sigma \nabla\mathcal{W}(\tilde{\theta})$ for some scalar σ .

Part (a). On the expansion path, $\nabla\mathcal{W}(\tilde{\theta}) = \mu \nabla\mathcal{C}(\tilde{\theta})$, with $\mu \neq 0$ because marginal welfare is non-zero. By Lemma 1,

$$\nabla\overline{MVPF}(\tilde{\theta}) = \frac{\mu - \overline{MVPF}(\tilde{\theta})}{C} \nabla\mathcal{C}(\tilde{\theta}) = \frac{\mu - \overline{MVPF}(\tilde{\theta})}{\mu C} \nabla\mathcal{W}(\tilde{\theta}),$$

so the curves are tangent.⁵²

Part (b). If the curves are tangent, Lemma 1 gives $(1 - \sigma C) \nabla\mathcal{W}(\tilde{\theta}) = \overline{MVPF}(\tilde{\theta}) \nabla\mathcal{C}(\tilde{\theta})$. The right-hand side is non-zero, so $1 - \sigma C \neq 0$ and $\nabla\mathcal{W} = \mu \nabla\mathcal{C}$ with $\mu = \overline{MVPF}(\tilde{\theta})/(1 - \sigma C)$. Component by component, $MVPF_j(\tilde{\theta}) = \mu$, which is positive because marginal welfare and marginal cost have the same sign. This is condition (18) at the budget C . The condition is sufficient when $\mathcal{W}$ is concave and $\mathcal{C}$ is convex. The function $L = \mathcal{W} - \mu\mathcal{C}$ is then concave and stationary at $\tilde{\theta}$, so any θ with $\mathcal{C}(\theta) \leq C$ satisfies $\mathcal{W}(\theta) \leq L(\theta) + \mu C \leq L(\tilde{\theta}) + \mu C = \mathcal{W}(\tilde{\theta})$, and $\tilde{\theta}$ lies on the expansion path.

⁵²If $\mu = \overline{MVPF}(\tilde{\theta})$, the gradient of the non-marginal MVPF vanishes and tangency holds trivially. Under Assumption D, this requires constant returns from the first dollar. Otherwise, it occurs where the average return peaks (Section B.9).

Remark 3. *Without such a condition, tangency does not imply optimality. With $\mathcal{W}(\theta) = \theta_1^2 + \theta_2^2$, $\mathcal{C}(\theta) = \theta_1 + \theta_2$ and $\tilde{\theta} = (1, 1)$, every hypothesis of the corollary holds and $\nabla \overline{MVPF}(\tilde{\theta}) = (1/2, 1/2) = \nabla \mathcal{W}(\tilde{\theta})/4$. Yet $(2, 0)$ has the same cost and a welfare of 4 against 2. Welfare is convex, and $\tilde{\theta}$ is a minimum of welfare on the budget surface. In general, tangency points are the critical points of welfare on the budget surface, of which the expansion path is the set of maxima. The same-sign condition is needed too, since a tangency point with $\mu < 0$ lies beyond the satiation point of welfare.*

The tangency condition involves only the gradients of $\mathcal{W}$ and $\overline{MVPF}$ at $\tilde{\theta}$. The expansion path can therefore be traced without the budget surfaces, and the budget of each point is recovered as $C = \mathcal{C}(\tilde{\theta})$. ■

B.8 The shadow value of public funds along the expansion path

Under Assumption D, $\mu(C) = \mathcal{V}'(C)$ at every budget with an interior solution, a binding constraint and $\nabla \mathcal{C} \neq 0$ (Lemma 2, part (ii)). The derivative of a concave function is non-increasing, so $\mu(C)$ falls as the budget grows. When $\mathcal{V}$ is strictly concave, the optimal scale C^* with $\mu(C^*) = 1$ is unique, $\mu(C) \geq 1$ below it and $\mu(C) \leq 1$ above it, so the rule “expand as long as the marginal dollar returns more than a dollar” is valid. The non-marginal MVPF, $\mathcal{V}(C)/C$, also falls with the budget and stays above $\mu(C)$ (part (iii)), with a gap of νC . Neither statement holds without Assumption D. ■

B.9 Increasing returns to public spending

Assumption D can fail. The diagnostic compares the marginal and the non-marginal MVPF along a policy path. With $v(c)$ the welfare along the path as a function of its cost c , the marginal MVPF at the end of the path is $v'(c)$, the non-marginal MVPF is $v(c)/c$, and $d[v(c)/c]/dc = [v'(c) - v(c)/c]/c$. A marginal MVPF above the non-marginal MVPF therefore means that the return per dollar is still rising at the observed scale. This describes returns along that path, and it bears on Assumption D only along the expansion path. The two MVPFs must be computed along the same path. In our applications, the price instruments have diminishing returns at the observed policy, while the returns to charging stations and to the fuel-economy label are rising (Sections 4 and 4.4). Public goods are natural candidates for increasing returns. Their benefits accrue to every user, so the willingness-to-pay for a marginal unit grows with the number of users that earlier units attracted, while its cost does

not.⁵³

What survives. Lemma 1, parts (i), (ii) and (iv) of Lemma 2, part (a) of Proposition 1, Proposition 2, Section B.6 and part (a) of Corollary 1 do not use Assumption D. An interior solution of the welfare problem still equalizes the marginal MVPFs, at $\mu = \mathcal{V}'(C)$ wherever $\mathcal{V}$ is differentiable. The MVPF is still direction-independent there, and the iso-welfare and iso-MVPF curves are still tangent along the expansion path.

What fails when $\mathcal{V}$ is convex. Three results change. First, the marginal MVPF exceeds the non-marginal MVPF, and $\mu(C) = \mathcal{V}'(C)$ rises along the expansion path. Second, the optimal mix is a corner. Suppose that welfare is separable across J projects, e.g., the charging networks of J regions, with $\mathcal{W} = \sum_j \mathcal{V}_j(C_j)$ and each $\mathcal{V}_j$ convex in the spending C_j . A convex function reaches its maximum on the budget set $\{C_j \geq 0, \sum_j C_j = \bar{C}\}$ at a vertex, so some projects are funded fully and the others not at all. An interior allocation with equalized marginal MVPFs is a minimum of welfare on the budget surface. The complementary slackness condition of Section 2.4 still holds at the optimum but does not single it out, and the marginal MVPFs can point to the wrong corner. With $\mathcal{V}_1(C) = 2C^2$, $\mathcal{V}_2(C) = C^2$ and a budget of 4, the allocation (1, 3) has marginal MVPFs of 4 and 6. Following them leads to the corner (0, 4), with a welfare of 16, while the optimum is (4, 0), with a welfare of 32. Third, there is no interior optimal scale. The problem $\max_C \mathcal{V}(C) - C$ maximizes a convex function, so its solution is no policy or the largest budget, and the rule $\mu(C^*) = 1$ selects a minimum. The MVPF problem is then solved at the budget that maximizes $\mathcal{V}(C)/C$ over $[\bar{C}, C_{\max}]$. A strictly convex $\mathcal{V}$ with $\mathcal{V}(0) = 0$ violates (41), so the spending constraint is slack and $\nu = 0$.

The S-shaped case. If $\mathcal{V}$ is convex at a small scale and concave at a large scale, the rule $\mu(C^*) = 1$ has no solution or, generically, two, a local minimum on the convex part and a local maximum on the concave part. The global optimum is either no policy or the larger solution, and the choice between them is a non-marginal comparison: the policy should be implemented at the scale C^* if and only if $\mathcal{V}(C^*) - \mathcal{V}(0) \geq C^*$. When $\mathcal{V}(0) = 0$, this requires the non-marginal MVPF of the whole policy to exceed one. At that scale, the marginal MVPF is below the non-marginal MVPF, since $\mathcal{V}'(C^*) = 1 \leq \mathcal{V}(C^*)/C^*$. The average

⁵³An S-shaped adoption curve does not produce increasing returns for a subsidy, because the inframarginal transfers grow with the adoption that earlier dollars induced. With given producer prices, the marginal MVPF of a subsidy first falls and then rises back toward one as the market saturates, since the subsidy becomes a pure transfer. The reversal then occurs with both MVPFs below one, and it is no sign of a policy worth expanding.

return peaks where a ray from the origin is tangent to $\mathcal{V}$, beyond the inflection point, and (41) holds at every larger budget. Proposition 1 therefore applies beyond that point even though Assumption D fails.

A marginal MVPF computed at a small scale says little about this comparison. A policy with increasing returns can have a marginal MVPF below one at a small scale and be worth implementing at a large one. Conversely, the first dollars of a policy with sharply diminishing returns can return much more than the policy as a whole. When welfare is more convex than costs, a policy must be evaluated as a whole, with the non-marginal MVPF.

C Welfare weights and income effects

Throughout the paper, we make two important assumptions on the welfare weights and the marginal utility of income. First, we assume that the Pareto weights equal the inverse of the marginal utility of income, $\psi_i = 1/\lambda_i$, so that the change in social welfare equals the aggregate willingness-to-pay. Second, we assume that the marginal utility of income does not depend on the policy instruments θ . We discuss both assumptions in what follows.

C.1 Arbitrary welfare weights

As in Hahn et al. (2026), the weights do not enter the willingness-to-pay or the MVPF, and they are applied only when comparing policies.

Marginal reforms. Let WTP_{ji} denote the willingness-to-pay of individual i for a small change in the policy along θ_j , so that $WTP_j = \sum_i WTP_{ji}$. Under arbitrary weights ψ_i , the change in social welfare is

$$\frac{d\mathcal{W}}{d\theta_j} = \sum_i \psi_i \frac{dV_i}{d\theta_j} = \sum_i \eta_i WTP_{ji} = \bar{\eta}_j WTP_j,$$

where $\eta_i = \psi_i \lambda_i$ is the social value of one dollar given to individual i , and

$$\bar{\eta}_j = \frac{\sum_i \eta_i WTP_{ji}}{\sum_i WTP_{ji}}$$

is the average weight of the beneficiaries of the policy, weighted by their willingness-to-pay. With $\psi_i = 1/\lambda_i$, we have $\eta_i = 1$ for all i and $\bar{\eta}_j = 1$. Notice that the willingness-to-pay, the net cost and the MVPF are measured in dollars and do not depend on the weights. Under arbitrary weights, one dollar of net spending on policy j delivers $\bar{\eta}_j MVPF_j$ in social welfare

instead of $MVPF_j$.

Comparing policies. A budget-neutral reform that raises spending on policy j and finances it with policy k increases social welfare if and only if

$$\bar{\eta}_j MVPF_j > \bar{\eta}_k MVPF_k,$$

which is equation (2) of [Hahn et al. \(2026\)](#). The ranking of two policies by their MVPFs therefore holds under any weights that give their beneficiaries the same average weight. Similarly, the optimal policy of Section 2.4 sets $\bar{\eta}_j MVPF_j = \mu$ for every instrument in use, where μ is now the social value of a dollar of public funds. Proposition 1 holds as stated with the welfare function $\mathcal{W}$ defined by the weights ψ_i , that is, with $\bar{\eta}_j MVPF_j$ and $\bar{\eta} \overline{MVPF}$ in place of $MVPF_j$ and $\overline{MVPF}$.

Non-marginal reforms. With a constant marginal utility of income, the same results hold for non-marginal reforms, with the equivalent variation EV_i in place of WTP_{ji} . The change in social welfare is then $\sum_i \eta_i EV_i = \bar{\eta} \overline{WTP}(\theta_0 \rightarrow \theta_1)$, where $\bar{\eta}$ is the average weight of the beneficiaries weighted by their equivalent variation, and one dollar of net spending on the reform delivers $\bar{\eta} \overline{MVPF}(\theta_0 \rightarrow \theta_1)$ in social welfare.

Weights by group. In our applications, the willingness-to-pay splits into components that accrue to distinct groups, namely consumers (including the internality correction), shareholders and the individuals who bear the environmental externality (see [Table A4](#) and [Table A5](#)). If all individuals in group g share a common weight η_g , the average weight is

$$\bar{\eta}_j = \sum_g \eta_g \frac{WTP_{jg}}{WTP_j},$$

where WTP_{jg} is the willingness-to-pay of group g in the decomposition tables. The social welfare of each policy can thus be computed from our tables for any choice of group weights. For instance, a lower weight on shareholders than on consumers (e.g., because shareholders have higher incomes) lowers $\bar{\eta}_j$ for policies that raise profits, such as electric-vehicle rebates, and raises it for policies that lower profits, such as the fuel-economy label. Since some components are negative, $\bar{\eta}_j$ need not lie between the smallest and the largest weight. Weights on future generations or on residents of other countries apply to the externality column only. As in [Hahn et al. \(2026\)](#), the social cost of carbon we use is not equity-weighted, so this column measures a willingness-to-pay.

C.2 Marginal utility of income and income effects

The results for non-marginal reforms in Appendix C.1 assume that the marginal utility of income λ_i does not depend on the policy. The assumption plays no role for marginal reforms. The willingness-to-pay for a small change around θ_0 converts utility into dollars at the marginal utility of income $\lambda_i(\theta_0)$, so the marginal MVPF is the same whether or not λ_i varies with the policy. A non-marginal reform, however, can induce a change from $\lambda_i(\theta_0)$ to $\lambda_i(\theta_1)$, for example, when a large subsidy raises the income of the subsidized households. We discuss the implications in what follows.

Welfare change. In our four applications, the models are linear in income, so λ_i varies across individuals but not with the policy. In general, the change in social welfare along a path $\theta(t)$ from θ_0 to θ_1 is

$$\mathcal{W}(\theta_1) - \mathcal{W}(\theta_0) = \int_0^1 \sum_i \eta_i(\theta(t)) \sum_j WTP_{ji}(\theta(t)) \theta'_j(t) dt, \quad \eta_i(\theta) = \psi_i \lambda_i(\theta).$$

With a constant marginal utility of income, the weights do not move along the path, and the integral reduces to $\sum_i \eta_i EV_i$, where $EV_i = [V_i(\theta_1) - V_i(\theta_0)]/\lambda_i$ is also the compensating variation. The result is the same along any path, which is what the trapezoid approximations of Section 2.3 rely on. With a diminishing marginal utility of income, instead, the weight η_i falls along the reform for the individuals who gain and rises for the individuals who lose. Holding the weights constant then overstates the welfare gain of a large transfer and understates the welfare loss of a large tax. The bias grows with the transfer to each beneficiary, and matters most for the rebates on electric vehicles and the solar subsidies, which pay several thousand dollars to each buyer.

When income effects matter. A second issue concerns the measurement of the willingness-to-pay. With income effects as described above, the difference in log-sums is no longer equal to the equivalent variation or the compensating variation (see [McFadden, 1999](#)). [Willig \(1976\)](#) gives the size of the error. As a share of the gains of household i , the error is about $\varepsilon_i EV_i/(2y_i)$, where ε_i is the income elasticity of the demand for the good and y_i is the household's income. The error is small when the gain of each household is small relative to income, as for the fuel-economy label of [Allcott et al. \(2025\)](#) and the audit subsidies of [Allcott and Greenstone \(2024\)](#). With discrete choices, the error concerns only the marginal consumers who change their choice. For inframarginal consumers, the subsidy is a transfer of money. The error is therefore largest for the non-marginal reforms that move

many consumers, such as the rebates on electric vehicles.

D Computational details

We provide an overview of the structural model and the computational details behind our MVPF calculations for each of our empirical applications. We want to disclose that the description of each application was streamlined using Claude Code (model: Opus 5.5; effort level: Max), to ensure that they are consistent with each other and with the model of Section 2.1.

Our application based on [Allcott et al. \(2025\)](#) construct all MVPF estimates from the structural model described in Section 1 of their paper. Their analysis uses instead causal estimates, and the structural model is never estimated. Since our methodology requires deep preference parameters, we went ahead and estimated their model by maximum likelihood. We provide our parameter estimates and a comparative table to assess the fit against their welfare estimates.

The applications based on [Fournel \(2026\)](#) and [Gillingham and Tsvetanov \(2019\)](#) are direct implementation of our methodology, because a fully specified model and data were readily available or in the public domain.

D.1 Electric vehicles – Fournel (2026)

D.1.1 Model

[Fournel \(2026\)](#) estimates a structural model of the market for new cars in Quebec and of the supply of public charging stations, with data at the county-year level for the 96 counties of Quebec from 2012 to 2020. The provincial and the federal government give rebates on electric vehicles at the point of sale. County governments, the local planners, install the public charging stations. We follow the notation of [Fournel \(2026\)](#), except for the taste for the network, written α_{mt} here, since θ denotes the policy. We compute the MVPFs on the last year of the data, 2020.

Demand. In year t , consumer i in county m chooses one of the available products or the outside option. The utility of product j is

$$u_{ijmt} = x_{jt}\beta_i - \lambda_i(p_{jt} - \tau_{jt}) + f_j(G_{mt}) + \xi_{jmt} + \epsilon_{ijmt},$$

where p_{jt} is the price, τ_{jt} the rebate, ξ_{jmt} the unobserved quality and ϵ_{ijmt} an extreme value shock. The tastes β_i and the marginal utility of income λ_i vary with the demographics of the county and with random coefficients. Battery electric vehicles and plug-in hybrids give utility from the county's network of G_{mt} stations,

$$f_j(G_{mt}) = \alpha_{mt} \frac{(1 + G_{mt})^\zeta - 1}{\zeta} \quad \text{for electric vehicles,} \quad f_j(G_{mt}) = 0 \quad \text{otherwise,}$$

where the “1” is charging at home and the county's taste for the network α_{mt} depends on the demographics of the county. The estimated curvature is $\zeta = 0.21$, so each additional station adds less utility than the previous station.

Firms. In each year, the firms compete in prices as multiproduct Bertrand competitors, and each firm sets one price per product for the whole province. The marginal costs are recovered from the firms' first-order conditions at the observed prices.

Local planners. The planner of county m installs public stations to maximize the value of the network to the current and future owners of electric vehicles in the county, net of the cost of the stations. [Fournel \(2026\)](#) derives the network supply of these forward-looking planners. The network of county m in year t is the expected network given the fleet of electric vehicles, plus a residual,

$$G_{mt} = E(G_{mt} \mid Q_{mt}^{ev}, \gamma) + \epsilon_{mt},$$

with

$$E(G_{mt} \mid Q_{mt}^{ev}, \gamma) = \sum_{g=1}^{\infty} \Phi \left(\gamma_1 \ln \Delta f_{mt}(g) + \gamma_2 \ln \left(Q_{mt}^{ev}(g) + \frac{\rho}{1-\rho} \Delta q_{mt}^{ev}(g) \right) \right),$$

where Φ is the cumulative distribution function of the standard normal. The fleet of electric vehicles given a network of size g evolve according to the law of motion

$$Q_{mt}^{ev}(g) = (1 - d) Q_{m,t-1}^{ev} + q_{mt}^{ev}(g),$$

where d is the depreciation rate of the fleet and $q_{mt}^{ev}(g)$ are the sales of electric vehicles in county m and year t . The marginal sales due to station g are $\Delta q_{mt}^{ev}(g) = q_{mt}^{ev}(g) - q_{mt}^{ev}(g-1)$, so network supply depends on the future benefits to these marginal owners. The average value of station g , in dollars, to the owners of electric vehicles is $\Delta f_{mt}(g)$, and is constructed from the indirect utility of charging services $f(g_{mt})$. The parameters of the planner are γ_1 , γ_2 , and ρ is a composite discount factor equal to 0.871. Finally, the residual ϵ_{mt} is estimated

from the baseline, $\hat{\epsilon}_{mt} = G_{mt} - E(G_{mt} \mid Q_{mt}^{ev}, \hat{\gamma})$ and held fixed in counterfactuals.

Equilibrium. At a given policy, the prices of the firms and the networks of the planners solve a joint fixed point, with the fleet of the previous year held fixed. In our counterfactuals, the derivative of the expected network with respect to the fleet enters the firms' first-order conditions, so the firms set prices knowing that sales of electric vehicles bring stations. At the observed policy, the model reproduces the observed prices, sales and networks. When the network is the policy, the province replaces the planners, since public stations fully crowd out the planners' stations.

Externality. The externality of a vehicle is the damage from the tailpipe CO₂ that the vehicle emits over an expected life of 12.02 years, at 22,053 km per year. The emissions are valued at \$245 per ton in 2020, and discounted at 2 percent. The emission profile of vehicles only account for their lifetime tailpipe emissions.

D.1.2 Computation of the MVPFs

This section describes how we compute the MVPFs of the electric vehicle policies (see Panel A of [Table 2](#) and [Table A4](#)). All welfare objects are computed on the sales of the last year of the data, with the lifetime flows of the vehicles discounted at 2 percent. Amounts are in 2020 dollars.

Policies. The policy is the rebate τ_j on each product, the sum of the provincial and the federal rebates. The observed rebates vary across models. The standard rebate is \$13,000 for a battery electric vehicle (\$8,000 provincial and \$5,000 federal) and \$6,500 for a plug-in hybrid (\$4,000 and \$2,500). A reform of the rebates on battery electric vehicles moves the rebate on every battery electric vehicle in proportion to the observed rebate, and likewise for plug-in hybrids. The two counterfactual instruments are a rebate on every hybrid and a tax on every vehicle that emits more than 210 grams of CO₂ per km, both introduced at zero.

Willingness-to-pay and net cost. As in [Section 3](#),

$$\begin{aligned}\mathcal{W}(\theta) &= CS(\theta) + (1 - \tau_c) \Pi(\theta) - \sum_j D_j q_j(\theta), \\ \mathcal{C}(\theta) &= \sum_{j,m} \tau_j(\theta) q_{jm}(\theta) + c_g g(\theta) - \tau_c [\Pi(\theta) - \Pi(0)] - [R(\theta) - R(0)],\end{aligned}$$

where CS is the log-sum of Section 3 in dollars, summed over the households of the 96 counties, Π is the firms' variable profit, D_j is the externality of vehicle j of Appendix D.1.1, $g(\theta)$ is the number of stations installed in the year, summed over the counties, and $c_g = \$19,000$ is the one-time cost of a station. R is the lifetime revenue of the provincial gasoline tax (\$0.192 per litre) and of the federal levy on fuel-inefficient vehicles, and $\tau_c = 0.27$ is the combined federal and Quebec corporate tax rate. Consumer surplus includes the utility of the network.

Charging stations. The stations enter the MVPFs in two ways, depending on the policy instrument.

- *Rebates and taxes.* The local planners choose the stations, as in Appendix D.1.1. A rebate raises the sales of electric vehicles, and the local planners respond by installing more stations, e.g., these are positive network effects. The observed rebate schedule induces 171 of the 440 stations installed in 2020. Without the rebates, the local planners would install 269 stations. The value of these stations to the buyers is part of consumer surplus. The cost of these stations is part of the net cost.
- *Stations as the policy instrument.* Public stations installed by the province crowd out the stations of the planners one for one. When the network is the policy, we therefore remove the planners, and the network becomes one more instrument of the province. The province installs g stations, allocated across the counties as the observed installations, and $g = 440$ is the observed policy.

Decomposition. Table A4 splits the willingness-to-pay into consumers ΔCS , shareholders $(1 - \tau_c) \Delta \Pi$ and the externality $-\sum_j D_j \Delta q_j$. The table splits the net cost into the transfer, the take-up, the spending on stations $c_g \Delta g$, the additional tax revenues $\tau_c \Delta \Pi$ and the other tax revenues ΔR . For a rebate, the transfer is the rebate paid to the buyers who would have bought the vehicle anyway, and the take-up is the rebate paid to the additional buyers. Every term is divided by the mechanical term of the reform: the transfer for a rebate, the tax revenue at fixed sales for the tax, and the spending on stations for the stations.

Reforms. The marginal MVPF of each rebate is computed at the observed policy, and the marginal MVPF of the stations at $g = 440$. The hybrid rebate and the tax do not have a non-marginal MVPF as they are zero at baseline. The non-marginal MVPF of each rebate is the MVPF of the reform that raises the rebate from zero to the observed level, with the other rebate at the observed level. The non-marginal MVPF of the stations is the MVPF of the reform from $g = 0$ to $g = 440$. Table A4 adds three reforms of the rebate program:

both rebates from zero, the provincial rebates alone, and the federal rebates added to the provincial rebates (the federal rebate was added as a top-up in 2019). The marginal objects use the formulas of Section 3, with the responses of prices and stations computed by central differences of the equilibrium (analytical derivatives are possible for the charging networks). The non-marginal objects are differences in levels between the two equilibria.

Trapezoidal approximations. The structural trapezoid uses the marginal objects of the model at both ends of the reform. The extrapolated trapezoid replaces the objects at zero by values extrapolated from the observed policy, as if only the observed policy were known. For the rebates, the extrapolation holds constant the price elasticity of the vehicles that receive the rebate. For the stations, the extrapolation holds constant the responses to the network.

Standard errors. We compute standard errors by parametric bootstrap with 199 replications. In each replication, we draw the parameters from the estimated joint distribution of the parameters and recompute the MVPFs. The standard error is the standard deviation of the MVPFs across replications.

D.2 Residential solar – Gillingham and Tsvetanov (2019)

D.2.1 Model

Gillingham and Tsvetanov (2019) (GT hereafter) model the demand for residential solar photovoltaic (PV) systems in Connecticut. Two subsidies lower the price of a purchased system: a state rebate per watt, paid to the installer, and a federal tax credit of 30 percent of the price net of the rebate. We follow their notation. In particular, λ_{it} denotes the parameter of a Poisson distribution, not the marginal utility of income as in Section 2.1.

Demand. Let Y_{it} be the number of systems installed in block group i and year t , and p_{it} the price per watt after the rebate and before the credit. GT uses a hurdle model for Y_{it} . A logit model gives the probability of at least one installation, and, conditional on having at least one installation, a zero-truncated Poisson gives the number of installations :

$$\Pr(Y_{it} > 0) = \Lambda(\gamma_1 p_{it} + X_{it}\delta_1 + \alpha_i^L),$$

$$\Pr(Y_{it} = y \mid Y_{it} > 0) = \frac{\lambda_{it}^y}{y! (e^{\lambda_{it}} - 1)}, \quad \lambda_{it} = \exp(\gamma_2 p_{it} + X_{it}\delta_2 + \alpha_i^P),$$

where Λ is the logistic function, X_{it} contains demand shifters (a Solarize campaign, demographics and party registration) and year effects, and α_i^L and α_i^P are block-group effects. The two parts form one model of installations, with expected value

$$E[Y_{it}] = \Lambda(\gamma_1 p_{it} + X_{it}\delta_1 + \alpha_i^L) \frac{\lambda_{it}}{1 - e^{-\lambda_{it}}}.$$

The price affects both parts of the model, and the price elasticity of installations is the sum of the elasticities of the two parts.

Prices. The rebate is paid to the installer, who keeps 15.6 percent of it (GT, Table 8). The price thus falls by $\rho = 0.844$ per dollar of rebate. As in GT’s counterfactual simulations, we hold installers’ costs and the passthrough fixed across policies, and apply it to the 2014 average rebate, $R_0 = \$0.935$ per watt. We assume that the federal credit passes through fully to households. At rebate R and credit rate f , a household in block group i pays

$$(1 - f)p_i(R), \quad p_i(R) = p_i + \rho(R_0 - R),$$

per watt, where p_i is its price in 2014. Demand depends on the two instruments only through this net price.

Externality. A watt of capacity generates 32.26 kWh over a lifetime of 25 years, and displaced generation emits 258 grams of CO₂ per kWh (GT, Section 7.3), so a watt avoids 8.3 kilograms of CO₂. We count carbon dioxide only. We value the emissions of each year with the social cost of carbon of [Environmental Protection Agency \(2023\)](#) at a 2 percent discount rate, in 2014 dollars, from \$156 per ton in 2014 to \$237 per ton in 2038. Values before 2020 are extrapolated linearly, as in [Hahn et al. \(2026\)](#). With generation spread evenly over the 25 years and discounted at 2 percent, the avoided damage is $D = \$1.28$ per watt.

D.2.2 Estimation

GT estimate the model on a panel of block groups. Their estimators identify the price coefficients but not the level of demand at other prices, which are needed to identify willingness-to-pay and net costs. We therefore re-estimate the model on their data, with a slight improvement that allows for recovering the demand levels. This section describes our methodology, the data, and compares our estimated parameters to the published estimates of GT.

Data. The panel covers the purchased systems installed in 1,534 Census block groups from 2008 to 2014, or 10,738 block-group-years. Y_{it} is zero in 70 percent of them. The price is the average price per watt in the block group and year, in 2014 dollars. When a block group has no installation, it is the average price in its town, or in its county when the town has none. The controls are an indicator for a Solarize campaign, population density, median household income, median age, two education shares and the shares of registered Republicans and Democrats in the town.

GT’s estimator. GT estimates the model in a two-part procedure. First, they estimate the logit model by conditional maximum likelihood. Prices are endogenous, and GT addresses the endogeneity issue with a control function approach. The instruments include the state rebate and the county wage of roofing contractors. Second, they estimate the zero-truncated Poisson by generalized methods of moments (the instruments enter the moment conditions). Both parts include block group and year fixed effects.

Identifying demand levels. With (at most) seven years of data per block group, the block group fixed effects are not identified which would bias the price coefficients (the incidental parameters problem in nonlinear models). To circumvent this issue, GT use conditional estimators which difference out the fixed effects. These estimators identify the price coefficients but not the level of demand, because in a nonlinear model the response of installations to the price depends on the fixed effect. In the logit, for example, the effect of a price change on the probability of an installation is $\gamma_1 \Lambda(1 - \Lambda)$, which depends on α_i^L through Λ .

Estimating MVPF, however, requires demand levels. To identify them, we replace the fixed effects with correlated random effects, whose distribution we estimate. Demand at any price is then obtained by integrating over this distribution.⁵⁴

Our estimator. We keep GT’s data, controls, instruments, and first stage. Following [Mundlak \(1978\)](#) and [Chamberlain \(1980\)](#), the random effects depend on the block-group averages of the regressors,

$$\alpha_i^h = \bar{x}_i \omega_h + \sigma_h \epsilon_i^h, \quad \epsilon_i^h \sim N(0, 1), \quad h \in \{L, P\},$$

where $\bar{x}_i$ contains a constant and the averages of the price and the controls of block group i over 2008–2014. The random effects of the two parts are independent (the correlation

⁵⁴We integrate over the distribution of the random effects conditional on the observed installations of each block group. Integrating over the unconditional distribution changes the MVPFs by less than 0.001.

is within part), so we can estimate the logit and the truncated Poisson separately. We estimate both by maximum likelihood, using the first-stage residual as a control function, and integrate over the random effects with 24-node Gauss–Hermite quadrature. Standard errors are clustered by town and do not account for the estimation of the first stage. As a check, we also replicate GT’s estimates, and reproduce the price coefficients of all four columns of their Table 6 to within 0.002.

Estimates. Table D1 compares the estimates from our augmented model with GT’s main specification. We reproduce GT’s estimate reasonably well, with the exception of the price sensitivity parameter in the truncated Poisson model, which are about one standard error apart (-0.43 vs. -0.29).

Table D1: Demand for solar PV, estimates from GT and ours

	At least one installation (logit)		Number of installations (truncated Poisson)	
	GT (1)	Ours (2)	GT (3)	Ours (4)
Price (\$ per watt)	-0.201 (0.0999)	-0.201 (0.109)	-0.292 (0.132)	-0.428 (0.133)
Solarize campaign	0.863 (0.204)	0.902 (0.185)	0.877 (0.161)	0.800 (0.158)
Population density	-0.0002 (0.00009)	-0.00026 (0.00008)	0.0004 (0.0005)	0.00021 (0.00039)
Median household income	-0.0014 (0.00152)	-0.0013 (0.0014)	0.0013 (0.00267)	0.0026 (0.0025)
Median age	-0.0096 (0.0053)	-0.0104 (0.0048)	-0.008 (0.012)	-0.0109 (0.0113)
Share with some college	0.002 (0.004)	0.003 (0.004)	0.005 (0.007)	0.003 (0.006)
Share with a graduate degree	0.006 (0.005)	0.005 (0.004)	-0.005 (0.008)	-0.005 (0.007)
Share of registered Republicans	0.036 (0.041)	0.035 (0.033)	0.153 (0.059)	0.173 (0.059)
Share of registered Democrats	0.011 (0.0248)	0.013 (0.019)	0.102 (0.0384)	0.101 (0.034)
Control-function residual	–	0.867 (0.145)	–	0.325 (0.122)
SD of the block-group effect	–	0.171 (0.113)	–	0.685 (0.035)
Price endogeneity	Control function	Control function	GMM	Control function
Block-group effects	Fixed	Correlated random	Fixed	Correlated random
Observations	10,738	10,738	3,238	3,238

Notes: Columns (1) and (3) are columns (3) and (4) of Table 6 of GT, with block-group fixed effects. Columns (2) and (4) are our estimates with correlated random effects, which also include the block-group averages of the price and the controls. All columns include year effects. The dependent variable is the number of installations of purchased systems in a block group and year. The truncated Poisson uses the 3,238 observations with at least one installation. Standard errors are in parentheses. GT bootstrap them in column (1) and cluster them by town in column (3). Ours are clustered by town and treat the first-stage residual as known.

Fit. In 2014, the model predicts 2,012 installations and 15.7 MW of capacity, against 1,974 installations and 15.4 MW in the data. In 2008–2013, it underpredicts installations by 2 to 6 percent per year. The average price elasticity of installations over 2008–2014 is -0.95.⁵⁵

⁵⁵GT report a hurdle elasticity of -0.65, the sum of -0.53 for the logit and -0.12 for the truncated Poisson. Their Proposition 2 multiplies the price coefficient of the truncated Poisson by the average price. Applied to their estimates, it gives -0.48 for the truncated Poisson and -1.00 for the hurdle.

D.2.3 Computation of the MVPFs

This section describes how we compute the MVPFs of the solar subsidies (see Panel B of [Table 2](#) and [Table A4](#)).

Policies. We define the policy space as $\theta = (k, f)$, where k scales the 2014 average rebate, $R = kR_0$, and f is the credit rate. The observed policy is $\theta = (1, 0.30)$. All welfare objects are computed on the installations of 2014, the last year of GT's data. Amounts are in 2014 dollars. Installed capacity at policy θ is

$$Q(\theta) = \sum_i a_i E[Y_i](\theta),$$

where a_i is the average system size in block group i in 2014 and $E[Y_i](\theta)$ is its expected number of installations at the net price $(1 - f)p_i(R)$ of [Appendix D.2.1](#).

Consumer surplus. The hurdle model has no utility function, so we cannot compute consumer surplus from preferences as in [Section 2.1](#). Instead, we treat the expected installations in block group i as the demand of the households of block group i , without income effects. Let $x_i(\theta) = (1 - f)p_i(R)$ be the price paid per watt in block group i under policy θ , and $y_i(x)$ the expected number of installations in block group i at a price paid of x . The change in consumer surplus in block group i between policies θ_0 and θ_1 is the area under the demand for capacity between the prices paid $x_i(\theta_0)$ and $x_i(\theta_1)$:

$$\begin{aligned} \Delta CS &= - \sum_i a_i \int_{x_i(\theta_0)}^{x_i(\theta_1)} y_i(x) dx \\ &= - \sum_i a_i \int_{x_i(\theta_0)}^{x_i(\theta_1)} E_{\alpha^L} \left[\Lambda \left(\gamma_1 \frac{x}{1 - f_0} + X_i \delta_1 + \alpha^L \right) \right] \\ &\quad \times E_{\alpha^P} \left[\lambda_i \left(\frac{x}{1 - f_0} \right) \middle/ \left(1 - \exp \left\{ - \lambda_i \left(\frac{x}{1 - f_0} \right) \right\} \right) \right] dx, \end{aligned}$$

where a_i is the average system size in block group i (in kW), $\lambda_i(p) = \exp(\gamma_2 p + X_i \delta_2 + \alpha^P)$, and X_i holds the regressors of block group i , including the control function. The price in Λ and λ_i is the price before the federal credit, as in GT. For a price paid of x , this price is $x/(1 - f_0)$, the price before the credit that gives x at the credit rate of $f_0 = 0.30$ in GT's data. The expectations E_{α^L} and E_{α^P} are over the random effects of block group i . The integral has no closed form and is computed numerically. Without income effects, ΔCS is

the consumers' part of the equivalent variation $\sum_i EV_i(\theta_1)$ of Section 2.3.

Willingness-to-pay and net cost. As in Section 2.1,

$$\begin{aligned}\mathcal{W}(\theta) &= CS(\theta) + (1 - \tau_c) \Pi(\theta) + D Q(\theta), \\ \mathcal{C}(\theta) &= R Q(\theta) + f \mathcal{E}(\theta) - \tau_c \Pi(\theta),\end{aligned}$$

where $CS(\theta)$ is consumer surplus, $\Pi(\theta) = (1 - \rho) R Q(\theta)$ is the part of the rebate kept by installers, D is the avoided damage per watt, $\mathcal{E}(\theta) = \sum_i p_i(R) a_i E[Y_i](\theta)$ is the base of the credit, and $\tau_c = 39.9$ percent is the corporate tax rate (the federal rate of 35 percent and the Connecticut rate of 7.5 percent, which is deductible from the federal base). Installers earn no other margin in the model. The net cost includes the spending and revenues of both the state and the federal government, but does not include the electricity-tax revenue lost on the power produced by the households.

For a marginal reform, the change in consumer surplus equals the installed capacity times the fall in the price paid, because the households at the margin are indifferent between installing and not installing (Section 2.1). A small increase in the rebate scale raises the rebate by $R_0 dk$ per watt. The posted price falls by $\rho R_0 dk$ and the price paid falls by $(1 - f)\rho R_0 dk$, so consumer surplus rises by $(1 - f)\rho R_0 Q dk$. A small increase in the credit rate leaves the posted price unchanged and lowers the price paid in block group i by $p_i(R) df$, so consumer surplus rises by $\sum_i p_i(R) a_i E[Y_i](\theta) df = \mathcal{E}(\theta) df$.

Decomposition. Table A4 splits the willingness-to-pay into consumers (ΔCS), shareholders $((1 - \tau_c) \Delta \Pi)$ and the externality ($D \Delta Q$). The table splits the net cost into the transfer, the take-up, the additional corporate tax revenues ($\tau_c \Delta \Pi$) and the other tax revenues. For a marginal reform, the transfer is the outlay of the instrument at fixed installations, $R_0 Q$ per unit of k and $\mathcal{E}$ per unit of f . For a non-marginal reform, the transfer is the full outlay of the instrument at the new installations. The take-up is the outlay of the two subsidies on the induced capacity. The other tax revenues are the credit saved by the federal government when the rebate rises. Since the credit is a share of the price net of the rebate, each dollar of rebate lowers the credit paid on the systems that would be installed anyway by $f\rho = 0.25$. Every term is divided by the transfer.

Reforms. The marginal MVPF of each instrument is computed at the observed 2014 policy. The non-marginal MVPF of each instrument is the MVPF of the reform that raises the instrument from zero to the 2014 level, with the other instrument at the 2014 level.

The counterfactual rows set both instruments at the 2021 levels: a rebate of \$0.28 per watt ($k = 0.296$) and a credit rate of 26 percent ($f = 0.26$). For these rows, the marginal MVPF of each instrument is computed at the 2021 policy, and the non-marginal MVPF is the MVPF of the reform that raises the instrument from zero to the 2021 level, with the other instrument at the 2021 level. [Table A4](#) adds three reforms of the whole program, in which both instruments move together: from no policy to the 2014 policy, from no policy to the 2021 policy, and from the 2014 policy to the 2021 policy. We compute the change in consumer surplus with the trapezoid rule, on 41 points along the reform (401 points for the programs). The other terms of the willingness-to-pay and the net cost are differences in levels.

Trapezoidal approximations. The structural trapezoid uses the marginal objects of the model at both ends of the reform. The extrapolated trapezoid replaces the objects at zero by values extrapolated from the marginal objects at the chosen policy and a constant price elasticity of capacity.

Standard errors. We compute standard errors by parametric bootstrap with 199 replications. In each replication, we draw the demand parameters from their estimated joint distribution and the passthrough from a normal distribution with GT’s standard error. The standard error is the standard deviation of the MVPFs across replications.

D.3 Informational nudge – Allcott et al. (2025)

D.3.1 Model

[Allcott et al. \(2025\)](#) (ACMT hereafter) study fuel-economy labels in an online experiment. Participants choose, for a hypothetical three-year lease, between a sedan rated at 23 miles per gallon (MPG), the Ford Fusion, and a similar sedan rated at 30 MPG, the Honda Accord, Nissan Altima or Subaru Legacy. All valuations and prices are relative, i.e., those of the low-MPG car minus those of the high-MPG car, in dollars per vehicle-year. We follow their notation. In particular, τ denotes the effect of the label, not a price instrument as in [Section 2.1](#), and t denotes a tax.

Choices. Consumer i values the low-MPG car at v_i relative to the high-MPG car. This experienced valuation includes the true difference in fuel costs. When choosing, consumers overvalue the low-MPG car by the bias γ_i (e.g., because they pay little attention to fuel costs), and the label shifts the valuation of a consumer who sees it by τ_i . Let $L_i = 1$ if

consumer i sees the label. At relative price p , consumer i buys the low-MPG car if

$$v_i + \gamma_i + \tau_i L_i > p.$$

The bias and the label change choices but not experienced utility. A consumer who buys the low-MPG car obtains $v_i - p$. ACMT allow τ_i to covary with v_i and γ_i and leave the joint distribution of (v_i, γ_i, τ_i) unrestricted. We specify and estimate it in Appendix D.3.2. ACMT also allow the label to affect utility directly (e.g., the discomfort of seeing a warning). They set this effect to zero for their car experiment.

Label and tax. ACMT describe the label by an intensity σ that scales the effect τ_i of every consumer. We describe it by the share n of consumers who see it, as the nudge of Section 2.1. The two coincide without the label ($\sigma = n = 0$) and when every consumer sees it ($\sigma = n = 1$), which is the change that ACMT evaluate. The government can also tax the low-MPG car by t per vehicle-year, in which case producers receive $p - t$.

Externality. Driving the low-MPG car instead of the high-MPG car emits more carbon. As in ACMT, the damage is

$$\phi_i = m_i \left(\frac{1}{23} - \frac{1}{30} \right) \chi SCC,$$

where m_i is the consumer's yearly mileage, $\chi = 8.87$ kilograms of CO₂ per gallon, and $SCC = \$193$ per ton. This is the value of [Environmental Protection Agency \(2023\)](#) for 2020 emissions at a 2 percent discount rate, in 2020 dollars. The average damage is \$188 per vehicle-year.⁵⁶

Market. ACMT model supply as symmetric imperfect competition. Their welfare analysis needs only the passthrough of a tax to the relative price, $\rho = dp/dt = 0.8$, and the relative markup, $\mu = 0$. We close the model with a linear relative supply curve, with price equal to marginal cost and a slope that gives a passthrough of 0.8 at the equilibrium without label and tax. We normalize the relative price to zero in that equilibrium, which is the midpoint of the price lists of the experiment (both cars leased at \$2,000 per year). The relative price $p(n, t)$ clears the market.

⁵⁶ACMT uses a social cost of carbon of \$51 per ton, a discount rate of 3%, in 2020 dollars, see [Interagency Working Group on Social Cost of Greenhouse Gases \(2021\)](#). The associated average damage is \$50 per vehicle-year.

Welfare. Total surplus is experienced consumer surplus, plus producer surplus, plus tax revenue, minus the externality. Per consumer, experienced consumer surplus and the externality are

$$CS(n, t) = E[(v_i - p) \mathbf{1}\{v_i + \gamma_i + \tau_i L_i > p\}], \quad \Phi(n, t) = E[\phi_i \mathbf{1}\{v_i + \gamma_i + \tau_i L_i > p\}],$$

where $p = p(n, t)$, the expectation is over consumers and over who sees the label ($L_i = 1$ with probability n), and surplus is measured relative to leasing the high-MPG car. Let $Q(n, t)$ be the share of consumers who buy the low-MPG car, and $\Pi(n, t)$ producer surplus, the area above the supply curve up to the producer price $p - t$. Total surplus is $TS(n, t) = CS(n, t) + \Pi(n, t) + tQ(n, t) - \Phi(n, t)$. Appendix D.3.3 adds the revenue from gasoline taxes, the corporate tax and the cost of the label.

The welfare formula of ACMT. ACMT evaluate the label with a second-order approximation that uses a few moments of the joint distribution of $(\gamma_i + \phi_i, \tau_i)$. Without tax, the change in total surplus from $n = 0$ to $n = 1$ is

$$\Delta W(t = 0) \approx \frac{1}{2} \left\{ \rho \left[(E[\gamma + \phi] - \mu + E[\tau])^2 - (E[\gamma + \phi] - \mu)^2 \right] + 2 \text{cov}[\gamma + \phi, \tau] + \text{var}[\tau] \right\} D'_p,$$

where $D'_p < 0$ is the slope of demand and the moments are population moments, which ACMT use in place of moments among marginal consumers. The covariance of τ_i with v_i does not enter the formula, but it changes the moments among marginal consumers. The first term is the effect of the average change in choices. The last two terms are the change in the variance of the distortion, from $\gamma_i + \phi_i$ to $\gamma_i + \phi_i + \tau_i$. With the optimal tax, $t^* = E[\gamma + \phi + \tau] - \mu$, the first term is zero. Appendix D.3.2 compares the moments and welfare numbers of ACMT with those of our model.

D.3.2 Estimation

ACMT estimate only the moments that enter their welfare formula. We estimate the joint distribution of valuations, biases and label effects by maximum likelihood on their experimental data, so that the model gives welfare at any policy. This section describes the data, the specification, the estimator, the estimates and the fit.

Data. The sample has the 1,267 participants of ACMT after their exclusions: 1,002 in the four label arms, which we pool as ACMT do, and 265 in the control group. Each participant values two pairs of cars, $j = 1$ (Nissan Altima against Ford Fusion) and $j = 2$ (Honda Accord or Subaru Legacy against Ford Fusion), at baseline ($s = 1$) and at endline ($s = 2$). Only

the endline screens show a label, and only to treated participants ($T_i = 1$). The price lists fix the high-MPG car at \$2,000 per year and vary the price of the low-MPG car from \$500 to \$3,500, so they reveal the relative willingness to pay w_{ijs} within a \$100 interval between $-\$1,500$ and $\$1,500$. We treat the 58 answers outside this range (out of 5,068) as \$100 intervals around the reported value. We also observe the benchmark valuation of ACMT, $\hat{v}_{ij} = f_{ij} - g_{ij}$, the willingness-to-pay if gasoline were free minus the gasoline cost implied by the participant's mileage and gas price.

Specification. The experienced valuation, the elicited willingness to pay and the benchmark are

$$\begin{aligned} v_{ij} &= \delta_j + \kappa_i + \varsigma_{ij}, \\ w_{ijs} &= v_{ij} + \gamma_i + b_{ij} + \tau_i T_i \mathbf{1}\{s = 2\} + e_{ijs}, \\ \hat{v}_{ij} &= v_{ij} + u_{ij}, \end{aligned}$$

where δ_j is the mean valuation of pair j , $\kappa_i \sim N(0, \sigma_\kappa^2)$ and $\varsigma_{ij} \sim N(0, \sigma_\varsigma^2)$ are tastes of the participant and of the pair, $\gamma_i \sim N(\mu_\gamma, \sigma_\gamma^2)$ and $b_{ij} \sim N(0, \sigma_b^2)$ are biases of the participant and of the pair, and e_{ijs} and $u_{ij} \sim N(0, \sigma_u^2)$ are noise in the price lists and in the benchmark. The consumer of Appendix D.3.1 has valuation v_{ij} , bias $\gamma_i + b_{ij}$ and label effect τ_i . The label effect is

$$\tau_i = R_i [\mu_\tau + \beta_\gamma(\gamma_i - \mu_\gamma) + \beta_\kappa \kappa_i + \alpha_\phi(\phi_i - \bar{\phi}) + \nu_i], \quad \nu_i \sim N(0, \sigma_\nu^2),$$

where $R_i \sim \text{Bernoulli}(p_r)$ indicates the participants who respond to the label. Among responders, μ_τ is the mean effect, and β_γ , β_κ and α_ϕ are the slopes on the bias, the taste and the externality. The label effect is the same for both pairs, so it can depend on the taste and bias of the participant but not on those of the pair. The noise e_{ij1} is drawn from $N(0, \sigma_e^2)$. At endline, the participant repeats the baseline noise ($e_{ij2} = e_{ij1}$) with probability π_s and draws new noise otherwise. All components are independent of each other and of T_i . The implied moments of the label effect are $E[\tau] = p_r \mu_\tau$, $\text{cov}[\gamma, \tau] = p_r \beta_\gamma \sigma_\gamma^2$ and $\text{cov}[v, \tau] = p_r \beta_\kappa \sigma_\kappa^2$.

We estimate two specifications. The first sets $\beta_\kappa = 0$, so that the label effect does not depend on the valuation given the bias. The second leaves β_κ free, as ACMT leave $\text{cov}[v, \tau]$ unrestricted. The welfare estimates of the paper use the second.

Each feature matches a feature of the data. In the control group, 27 percent of endline answers equal the baseline answer for the same pair, far more than independent noise produces, and π_s captures these repetitions. In the treated group the share is 21 percent. A treated

participant who repeats the baseline noise changes the answer by exactly τ_i , so this share requires many participants with $\tau_i = 0$, and R_i captures them. The benchmark measures the taste without the bias, which separates β_κ from β_γ . Compared with ACMT, the specification imposes normal layers, a bias independent of the externality, and a label effect linear in the bias, the taste, and the externality.

Estimator. We estimate the parameters by maximum likelihood, 15 in the first specification and 16 in the second. The likelihood of a participant is the probability of the four observed intervals times the density of the two benchmarks. Given $(\kappa_i, \gamma_i, \nu_i)$ and the pair-level taste and bias, the interval probabilities are normal probabilities in closed form, with repeated or new noise, and for responders or non-responders. The mixture over repeated noise is taken pair by pair, and the mixture over responders at the level of the participant, since R_i and ν_i are shared by the two pairs. We integrate over κ_i , γ_i and the pair-level shock by Gauss–Hermite quadrature, with 13 nodes for κ_i and γ_i and 11 for the pair-level shock.

The integral over ν_i needs more care. A treated responder who repeats the baseline noise in pair j changes the answer by exactly τ_i . The probability of the two observed intervals is then zero unless τ_i lies in a window of \$200, between the lower bound of the endline interval minus the upper bound of the baseline interval and the upper bound of the endline interval minus the lower bound of the baseline interval, and it has two kinks inside the window. Gauss–Hermite nodes are too far apart to resolve such a window. With 13 nodes, the log-likelihood of some treated participants is off by as much as 0.38 compared with a simulation with four million draws. We therefore split the likelihood of a responder according to whether each pair repeats the baseline noise. The term in which both pairs draw new noise is smooth in ν_i and keeps 13 Gauss–Hermite nodes. The other terms are integrated over the window of the repeating pair, or the intersection of the two windows, with four Gauss–Legendre nodes on each piece between kinks, placed in probability space ($u = \Phi(\nu_i/\sigma_\nu)$). The largest error is 0.006.

We maximize with L-BFGS-B and exact gradients (automatic differentiation in JAX), from four starting values with 9 nodes in each dimension, then with the final nodes. Standard errors are clustered by participant.

Estimates. [Table D2](#) reports the estimates. The two specifications agree on the valuations, the bias and the noise. About 55 percent of treated participants respond to the label. Among responders, the label lowers the relative valuation of the low-MPG car by \$60 to

\$65 on average, and by more for participants with a larger bias ($\beta_\gamma \approx -0.4$). In the second specification, it also lowers it more for participants with a larger taste for the low-MPG car ($\beta_\kappa = -0.45$). Freeing β_κ raises the log-likelihood by 32.6, and a likelihood ratio test rejects $\beta_\kappa = 0$ ($p < 0.001$). The slope on the externality is imprecise in both specifications. The standard deviation of the pair-level bias converges to zero in both.

Table D2: Estimates of the structural model of the fuel-economy label

	(1) $\beta_\kappa = 0$	(2) β_κ free
<i>Panel A: Valuation and bias, \$ per vehicle-year</i>		
Mean valuation, pair 1, δ_1	-558.4 (16.3)	-557.3 (16.3)
Mean valuation, pair 2, δ_2	-608.2 (16.2)	-607.7 (16.1)
SD of the participant taste, σ_κ	353.5 (16.2)	377.1 (16.4)
SD of the pair taste, σ_ς	128.2 (17.7)	134.5 (16.5)
Mean bias, μ_γ	135.4 (16.9)	134.0 (16.9)
SD of the participant bias, σ_γ	370.5 (18.3)	381.6 (18.4)
SD of the pair bias, σ_b	0.0	0.0
<i>Panel B: Label effect</i>		
Share of responders, p_r	0.539 (0.041)	0.549 (0.038)
Mean effect among responders, μ_τ	-64.5 (18.2)	-59.9 (16.4)
Slope on the bias, β_γ	-0.413 (0.085)	-0.382 (0.074)
Slope on the participant taste, β_κ	0	-0.447 (0.057)
Slope on the externality, α_ϕ	0.158 (0.212)	-0.249 (0.196)
SD of the idiosyncratic effect, σ_ν	169.6 (25.8)	149.9 (19.4)
<i>Panel C: Price lists and benchmark</i>		
SD of the price-list noise, σ_e	406.7 (8.3)	405.3 (8.4)
Probability of repeating the baseline noise, π_s	0.277 (0.015)	0.279 (0.015)
SD of the benchmark noise, σ_u	439.4 (19.5)	420.6 (19.1)
Log-likelihood	-34,601.1	-34,568.5
Number of parameters	15	16
Participants	1,267	1,267

Notes: This table presents the maximum likelihood estimates on the data of ACMT. Column (1) sets $\beta_\kappa = 0$. Column (2) leaves β_κ free and gives the welfare estimates of the paper. Valuations and biases are those of the low-MPG car relative to the high-MPG car, in dollars per vehicle-year. Standard errors in parenthesis are clustered at the participant-level.

Fit. Table D3 compares the moments and the welfare effects of the label in our model with the estimates of ACMT. Both specifications match their average bias, and their average label effect within one of their standard errors. The variance of the label effect is close to theirs without β_κ (29,260 against 30,428) and higher with it (40,696), within one of their standard errors.

The main difference is the covariance of the bias and the label effect, about $-30,500$ in both

specifications against $-7,744$ for ACMT. The two estimates are about one standard error of ACMT apart. Ours is more precise because the likelihood models the noise in the price lists and in the benchmark. ACMT recover the covariance from the interaction of the treatment with a measure of the bias, $w_{ij1} - \hat{v}_{ij}$, and about 70 percent of the variance of this measure is noise in our estimates.

Panel B reports the change in total surplus as defined by ACMT, from no label to the label for everyone. It excludes the gasoline taxes, the corporate tax and the cost of the label. Our model gives a gain of \$14 per vehicle-year without tax and of \$9.5 to \$10.4 with the tax, against losses of \$0.07 and \$4.43 for ACMT. The covariance explains the difference. With the moments of the second specification, their formula gives a gain of \$7.7 without tax. With their covariance instead of ours, it gives a loss of \$4.1, while replacing any other moment by theirs leaves a gain. ACMT do not report standard errors for these changes. Propagating the standard errors of their moments through their formula gives about \$13 in both cases, so our estimates are within about one standard error of theirs.

Table D3: Fit of the structural model to the estimates of ACMT

	(1) Allcott et al. (2025)	(2) Our model, $\beta_\kappa = 0$	(3) Our model, β_κ free
<i>Panel A: Moments of bias and label effects</i>			
Average bias, $E[\gamma]$	135 (17.6)	135 (17)	134 (17)
Average label effect, $E[\tau]$	-59 (23)	-35 (10)	-33 (10)
Variance of the label effect, $\text{var}[\tau]$	30,428 (11,925)	29,260 (6,045)	40,696 (6,695)
Covariance of bias and label effect, $\text{cov}[\gamma, \tau]$	-7,744 (21,348)	-30,568 (6,769)	-30,512 (6,720)
Covariance of externality and label effect, $\text{cov}[\phi, \tau]$	-37 (514)	158 (218)	-253 (200)
Covariance of valuation and label effect, $\text{cov}[v, \tau]$	-	0	-34,873 (5,170)
Demand slope, D'_p	-0.00060	-0.00053	-0.00052
<i>Panel B: Change in total surplus from the label, \$ per vehicle-year</i>			
Without tax	-0.07	14.09 (3.76)	13.99 (3.07)
With the optimal tax of \$126	-4.43	10.39 (2.78)	9.50 (2.17)
<i>Panel C: Likelihood</i>			
Log-likelihood	-	-34,601.1	-34,568.5
Number of parameters	-	15	16

Notes: Column (1) reports the estimates of ACMT for cars, pooling the four labels (their Tables 1 and 2), with their standard errors in parentheses. Columns (2) and (3) report the moments implied by our estimates and the exact changes in total surplus in our model, with standard errors from 999 draws of the parameters in parentheses. Moments are in dollars per vehicle-year (squared for variances and covariances), and D'_p is the change in the share of low-MPG buyers per dollar. Total surplus follows ACMT: consumer surplus plus producer surplus plus the revenue of the tax on the low-MPG car, minus the externality, from no label to the label for everyone. To be comparable with ACMT, this table values carbon at their \$51 per ton and uses their optimal tax of \$126. The changes in total surplus of ACMT come from their second-order formula (Appendix D.3.1) and have no standard errors.

D.3.3 Computation of the MVPFs

This section describes how we compute the MVPFs of the fuel-economy label (see Panel C of [Table 2](#) and [Table A5](#)).

Market. Each participant and pair of cars is a consumer. The noise terms e_{ijs} and u_{ij} of [Appendix D.3.2](#) are errors of the experiment and do not enter choices. For a consumer who does not see the label or does not respond to it, the decision valuation $v_{ij} + \gamma_i + b_{ij}$ is normal, with mean $\delta_j + \mu_\gamma$, variance $\sigma_\kappa^2 + \sigma_\zeta^2 + \sigma_\gamma^2 + \sigma_b^2$ and covariance $\sigma_\kappa^2 + \sigma_\zeta^2$ with v_{ij} . For a responder who sees the label, the decision valuation $v_{ij} + \gamma_i + b_{ij} + \tau_i$ is normal, with mean $\delta_j + \mu_\gamma + \mu_\tau + \alpha_\phi(\phi_i - \bar{\phi})$, variance $(1 + \beta_\kappa)^2 \sigma_\kappa^2 + \sigma_\zeta^2 + (1 + \beta_\gamma)^2 \sigma_\gamma^2 + \sigma_b^2 + \sigma_\nu^2$ and covariance $(1 + \beta_\kappa) \sigma_\kappa^2 + \sigma_\zeta^2$ with v_{ij} . For a normal decision valuation W with mean m , standard deviation s and covariance c with v_{ij} , demand and experienced surplus at price p are

$$\Pr(W > p) = \Phi(z), \quad E[(v_{ij} - p) \mathbf{1}\{W > p\}] = (\delta_j - p) \Phi(z) + \frac{c}{s} \varphi(z),$$

where $z = (m - p)/s$. Market demand $D(p; n)$, surplus and the externality average these expressions over consumers, exposure, and responder types. Relative supply is $S(p) = Q_0 + s_S(p - t)$, where $Q_0 = D(0; 0)$, $s_S = \rho |D'_p|/(1 - \rho)$ and D'_p is the slope of $D(p; 0)$ at $p = 0$, so that a tax passes through at $\rho = 0.8$. The price solves $D(p; n) = S(p)$, and producer surplus is $\Pi = Q_0(p - t) + s_S(p - t)^2/2$.

Welfare and cost. As in [Section 2.1](#),

$$\begin{aligned} \mathcal{W}(n) &= CS(n) + (1 - \tau_c) \Pi(n) - \Phi(n), \\ \mathcal{C}(n) &= c_n n - [R(n) - R(0)] - \tau_c [\Pi(n) - \Pi(0)] - t[Q(n) - Q(0)], \end{aligned}$$

where c_n is the cost of the label per exposed consumer and vehicle-year, R the revenue from federal and state gasoline taxes on the extra gallons of the low-MPG car, $\tau_c = 25.8$ percent the corporate tax rate, and t the tax on the low-MPG car. The willingness-to-pay and the fiscal externalities do not depend on c_n . We report the MVPF at $c_n = \$1, \$5$ and $\$10$.

Decomposition. The change in experienced surplus splits into the price savings of inframarginal buyers, $-Q dp$ (consumers), and the surplus $v_{ij} - p$ of the consumers whose choice changes (internality correction). The second term is zero without a wedge between decision and experienced utility. The label pays no transfer. In [Table A5](#), every term is per dollar of spending on the label. The take-up column holds the lost revenue of the tax on the low-MPG

car, a negative subsidy in place. The corporate tax is in additional tax revenues and the gasoline taxes in other tax revenues.

Reforms. The marginal MVPF is evaluated at $n = 1$, with the price response from the implicit function theorem. The non-marginal reform raises n from 0 to 1, the change evaluated by ACMT. The structural trapezoid uses the marginal objects at $n = 0$ and $n = 1$. The extrapolated trapezoid replaces the objects at $n = 0$ by values extrapolated from $n = 1$ at constant responses. The terms proportional to sales of the low-MPG car (price savings, profits and the corporate tax) are scaled by $[Q(1) - Q'(1)]/Q(1)$, where $Q' = dQ/dn$, and the other terms are unchanged. The counterfactual rows keep a tax of \$265 on the low-MPG car at both ends of the reform. This is the optimal tax with the label implied by the estimates of ACMT, $E[\gamma] + E[\phi] + E[\tau] = 135 + 189 - 59$.

Specifications. Panel C of [Table 2](#) and [Table A5](#) uses the specification with β_κ free. [Table D4](#) compares the MVPFs of the two specifications of [Appendix D.3.2](#). Without tax, the label moves twice as many buyers away from the low-MPG car when β_κ is free (5.4 against 2.7 percentage points). With $\beta_\kappa < 0$, the label lowers the valuation most for the consumers who value the low-MPG car most, and these are the one in five who buy it. The avoided carbon damages and the lost revenue from gasoline taxes and the corporate tax both double with the number of buyers who switch. The MVPF is lower at a cost of \$1, 4.5 against 6.2, but slightly higher at a cost of \$10, 1.6 against 1.5. The gap narrows and then reverses as the cost of the label increases (and takes a larger share of the net cost).

Table D4: MVPF of the fuel-economy label in the two specifications

	$\beta_\kappa = 0$		β_κ free	
	Marginal (1)	Non-marginal (2)	Marginal (3)	Non-marginal (4)
<i>Panel A: No tax on the low-MPG car</i>				
MVPF, label cost of \$1	6.20 (0.81)	6.06 (0.76)	4.48 (0.49)	4.35 (0.44)
MVPF, label cost of \$5	2.65 (0.54)	2.60 (0.52)	2.53 (0.34)	2.47 (0.32)
MVPF, label cost of \$10	1.54 (0.36)	1.52 (0.34)	1.63 (0.25)	1.61 (0.23)
Willingness to pay, \$	18.46 (4.94)	18.21 (4.82)	23.14 (4.19)	22.98 (4.02)
Fiscal externality without the label cost, \$	1.98 (0.52)	2.00 (0.54)	4.16 (0.54)	4.29 (0.57)
Change in the low-MPG share, percentage points	-2.7 (0.7)	-2.7 (0.7)	-5.5 (0.7)	-5.4 (0.7)
<i>Panel B: Tax of \$265 on the low-MPG car</i>				
MVPF, label cost of \$1	1.93 (0.14)	1.92 (0.14)	1.50 (0.08)	1.49 (0.08)
MVPF, label cost of \$5	1.29 (0.18)	1.27 (0.17)	1.19 (0.08)	1.18 (0.08)
MVPF, label cost of \$10	0.91 (0.16)	0.90 (0.16)	0.95 (0.08)	0.93 (0.08)
Willingness to pay, \$	15.52 (4.13)	15.22 (3.99)	23.47 (3.48)	22.64 (3.27)
Fiscal externality without the label cost, \$	7.05 (1.86)	6.95 (1.81)	14.69 (1.92)	14.24 (1.82)
Change in the low-MPG share, percentage points	-2.2 (0.6)	-2.1 (0.5)	-4.5 (0.6)	-4.3 (0.5)

Notes: MVPF of the label at a cost c per exposed consumer and vehicle-year, the willingness-to-pay divided by c plus the fiscal externality. Columns (1) and (2) set $\beta_\kappa = 0$, and columns (3) and (4) leave it free (Table D2). The marginal MVPF is at $n = 1$, and the non-marginal MVPF is for the change from $n = 0$ to $n = 1$. The willingness-to-pay and the fiscal externality are in dollars per exposed consumer and vehicle-year. The fiscal externality is the lost revenue from gasoline taxes, the corporate tax and, in Panel B, the tax on the low-MPG car. Standard errors in parenthesis are computed from 999 parametric bootstrap replications.

D.3.4 Counterfactual cases

We now discuss the counterfactuals presented in Figure 7. The details of each case are available in Table D5. We focus on our main specification, with β_κ free and a cost of \$1 per consumer exposed per vehicle-year, and we change one element of the model at a time. In each case, we recalibrate the relative supply as in Appendix D.3.3, by solving the new equilibrium relative price, and we compute the MVPF at $n = 1$.

Label design. The three label design experiments change the effect of the label, τ_i . The untargeted label scenario sets $\beta_\gamma = 0$, so the label effect does not depend on the bias. The fully targeted label scenario sets $\beta_\gamma = -1$, so the label removes the deviation of each responder's bias from the average bias. The debiasing label scenario sets $\mu_\tau = -\mu_\gamma$, $\beta_\gamma = -1$ and $\beta_\kappa = \alpha_\phi = \sigma_\nu = 0$, such that $\tau_i = -\gamma_i$, and responders choose on their experienced valuation (without bias). In all cases, the label effect does not affect non-responders.

Bias. For the unbiased consumers scenario, we set $\mu_\gamma = \sigma_\gamma = 0$ to remove the bias, and we keeps the label effect at our estimates. Without bias, each consumer already leases the car that gives the highest experienced utility, so every consumer that switches due to the label is

worse off. The label induces lower carbon emissions, and the decreased externality more or less offsets the loss in consumer utility. The no bias correction removes the bias correction term from the willingness-to-pay.

No-repricing. Without repricing, relative supply is flat ($\rho = 1$), so the relative price between low- and high-MPG cars is fixed for every n . As a result, producer surplus does not change and has no effect on the MVPF. Compared with the baseline, more consumers change to the high-MPG car, leading to a higher MVPF, since the repricing reduces the relative price of low-MPG cars, offsetting the effect of the nudge.

Table D5: MVPF of the fuel-economy label in alternative cases

Description	Parametrization	Willingness to pay	Fiscal externality	Change in the low-MPG share	MVPF
<i>Baseline</i>					
Our estimates, no tax	Table D2, column (2)	23.14 (4.19)	4.16 (0.54)	−5.5 (0.7)	4.48 (0.49)
Our estimates, corrective tax	Same, $t = \$265$	23.47 (3.48)	14.69 (1.92)	−4.5 (0.6)	1.50 (0.08)
<i>Label design</i>					
Debiasing label	Same, $\tau_i = -\gamma_i$	39.40 (4.43)	4.61 (0.47)	−6.3 (0.7)	7.02 (0.22)
Fully targeted label	Same, $\beta_\gamma = -1$	37.54 (4.24)	5.74 (0.56)	−7.9 (0.8)	5.57 (0.37)
Untargeted label	Same, $\beta_\gamma = 0$	3.56 (2.46)	2.32 (0.37)	−2.8 (0.5)	1.07 (0.71)
<i>Bias</i>					
No bias correction	Same, no bias correction in the WTP	11.78 (1.56)	4.16 (0.54)	−5.5 (0.7)	2.28 (0.08)
Unbiased consumers	Same, $\mu_\gamma = \sigma_\gamma = 0$	0.64 (0.47)	1.92 (0.30)	−3.0 (0.5)	0.22 (0.17)
<i>Other</i>					
No repricing (flat supply)	Same, $\rho = 1$	28.50 (4.55)	3.74 (0.47)	−6.7 (0.8)	6.01 (0.53)

Notes: Marginal MVPF of the label at $n = 1$ and a cost of \$1 per exposed consumer and vehicle-year. The willingness-to-pay and the fiscal externality are in dollars per exposed consumer and vehicle-year, and the change in the low-MPG share is in percentage points. The willingness-to-pay includes the bias correction term unless indicated otherwise. The fiscal externality is the lost revenue from gasoline taxes and the corporate tax, and includes the revenues from the corrective tax whenever applicable. Standard errors in parenthesis are computed from 999 parametric bootstrap replications.