Who Can Afford to Go Green? Optimal Environmental Policy under Financing Constraints

Simon Margolin* UC Santa Barbara

Alicia Zhang† Princeton

October 28, 2026

Latest draft can be found here.

Abstract

Environmental policy often relies on pricing 'brown' externality-producing goods to internalize damages. This principle follows a classic result in the Pigouvian taxation literature which argues that directly taxing the externality is optimal. This paper reexamines optimal Pigouvian taxation in light of income dynamics, financing constraints, and a fixed cost of switching between a brown good and a substitutable 'green' good. We show that subsidizing the substitutable green good may be optimal in conjunction with, or instead of, taxing the brown good. A brown tax can only induce switching to the green good for households that can already finance the fixed cost, while a green subsidy expands the *extensive margin* by making adoption feasible. We test these predictions using evidence from electric vehicle subsidies in California. ZIP codes with lower credit scores exhibit larger responses to the incentive payment, consistent with our theory that borrowing constraints can affect adoption decisions. Quantitatively, we find that a subsidy can increase the share of individuals who can afford to switch to the green technology, leading to higher welfare while achieving a given adoption target.

JEL codes: H21, H23, Q58, G51

*Simon Margolin: margolin@ucsb.edu

†Alicia Zhang: aliciazhang@princeton.edu.

We thank Philippe Aghion, Zach Bethune, John Grigsby, Kyle Meng, George Nikolakoudis and Richard Rogerson for their comments.

1 Introduction

The classic "targeting principle", introduced by Sandmo (1975) and generalized by Kopczuk (2003), is one of the pillars of externality correction in public finance. It states that when the consumption of a good generates negative externalities, the optimal policy is to impose a Pigouvian tax on this good rather than subsidize its externality-free substitutes. Despite this strong theoretical result, national and local governments have increasingly relied on green subsidies rather than environmental taxation in recent decades, as illustrated in Figure 1. While this observed switch is often attributed to political constraints, we show in this paper that the use of green subsidies may instead be justified by the presence of liquidity constraints that prevent households from adopting costly new technologies.

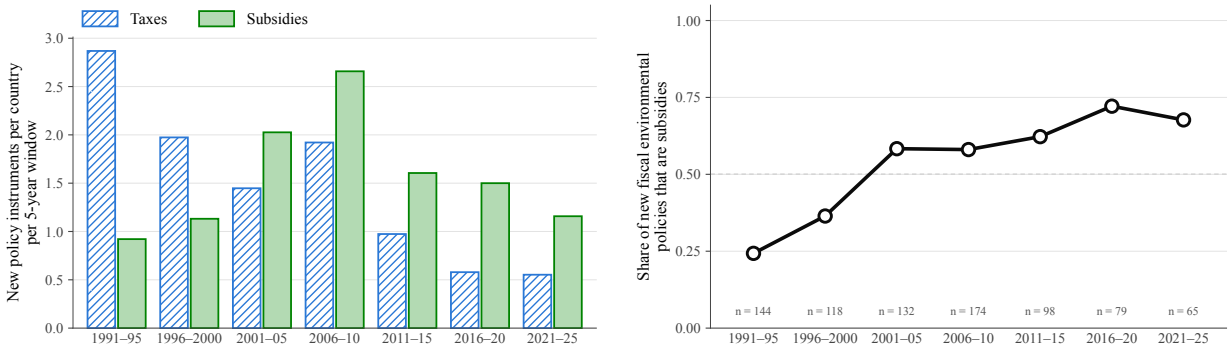
This paper studies the optimal design of taxes and subsidies in an economy with two empirically relevant features: replacing a polluting good with a green substitute requires paying a fixed adoption cost, and some households are unable to finance this cost because of low income or binding borrowing constraints. In this environment, a Pigouvian tax on the brown technology can induce switching for unconstrained households, but it does not generate adoption among liquidity constrained households who cannot afford the alternative option. In contrast, a green subsidy can operate through an additional extensive margin by relaxing the feasibility constraint.

We first characterize optimal taxes and subsidies in a two-period model with income inequality and borrowing constraints. Households choose between a "brown" externality-generating technology and a "green" technology, and adopting the green technology requires paying an upfront fixed cost. This setup speaks to many durable consumption decisions − buying an electric vehicle, investing in energy-efficient household appliances, installing solar panels, and more. In this model, a borrowing constraint generates a technology-specific feasibility region: some households that can afford to stick with the brown technology cannot pay the adoption cost to switch to the green technology.

In this environment, we derive sufficient statistics that highlight two channels governing the adoption response to policy: an *intensive margin* channel with a welfare effect that depends on the distribution of households that can always finance a switch, and an *extensive margin* channel driven by constrained households that are unable to afford switching absent policy intervention. These channels are absent in environments without income heterogeneity and financing constraints, and the extensive margin channel generates a deviation from the classic targeting principle. In particular, the extensive margin channel justifies including a green subsidy in the optimal policy mix even when the brown technology can be directly taxed.

For empirical validation, we take these predictions to the data by using California's EV incentive program as a case study. We extend the quasi-experimental design in Muehlegger and Rapson (2022) to test heterogeneity in the subsidy effect with respect to local credit conditions. We use ZIP code-level average credit scores measured in the pre-period as a proxy for local borrowing constraints. Consistent with the literature, we find that the subsidy increases EV adoption. Addi-

Figure 1: Environmental policy instruments in OECD countries



(a) New instruments

(b) Subsidy share

Notes: OECD PINE. New national climate and pollution instruments in five-year windows.

tionally, this effect is larger in ZIP codes with lower credit scores, even after controlling for baseline income, time, and ZIP fixed effects. This result indicates that the the extensive margin liquidity channel is an empirically relevant mechanism that affects the transmission and effectiveness of Pigouvian policy.

We extend our framework to an infinite horizon setting to quantitatively study transition dynamics. Borrowing constraints slow the diffusion of the green technology by preventing low-liquidity households from adopting optimally. In a calibrated model that matches baseline EV adoption and the cross-sectional gradient in subsidy effects by credit score, we compare tax-only, subsidy-only, and tax-plus-subsidy policies that implement a green adoption target. The subsidy-only policy reduces the welfare cost of achieving the adoption target by a multiple of three. Similarly, a planner that simply maximizes welfare will prefer a subsidy-only policy over using a tax. This exercise highlights that subsidies can increase the fraction of individuals who are able to adopt the green good and lead to higher welfare along the transition path.

Literature. This paper contributes to an extensive literature on optimal Pigouvian taxation. Sandmo (1975) provides a general framework for Pigouvian policy with multiple goods, and finds that in a second-best world where the government has an exogenous spending requirement, only the externality-generating commodity faces an additional Pigouvian tax while the other commodities only face revenue-generating taxes. This is similar to the 'targeting principle' in Kopczuk (2003). Compared to these papers, we introduce a fixed cost and borrowing constraint into the individual's problem and argue that these constraints are empirically relevant. These additional features generate a discontinuity in the planner's problem, which violates the targeting principle.

Recent papers study Pigouvian taxation in conjunction with different policy instruments. Smith (1992) surveys the literature on the 'double dividend' hypothesis, which posits that revenue raised from environmental taxes can reduce other distortionary taxes. We address this hypothesis from a different direction: If environmental taxes are regressive, does the double dividend still hold?

Cremer et al. (1998) studies the combination of income and commodities taxes when consumption may be unobservable. Farhi and Gabaix (2020) add a behavioral wedge into standard models of optimal taxation, and find in the Pigouvian case that the covariance between attention and the externality cost mediates the size of the optimal tax. Jaqua and Schaffa (2022) find that subsidizing a harmful activity may be optimal if the policymaker cannot tax a more harmful activity.

We also contribute to a literature on the transition to a green economy. Acemoglu et al. (2012) solve for the mix of taxes and subsidies to direct innovation towards clean inputs, and Aghion et al. (2016) study the interaction between fuel prices and directed technical change in the auto industry. While we do not incorporate innovation in this paper, we explore optimal Pigouvian policy to achieve a green transition from an individual perspective, as *consumption* of green goods is integral to the diffusion and adoption of green technologies after innovation. In a similar vein, Kremer and Willis (2016) consider the government's commitment problem to a path of anticipated subsidies over time. In future work, we hope to incorporate this dynamic problem into our quantification.

Finally, we draw from the empirical literature on household decision-making regarding green technologies to justify and measure the empirical importance of our liquidity mechanism. Fried et al. (2025) survey households on their reasoning for purchasing energy-saving equipment, and find that households have preferences for cutting emissions as well as saving on energy costs. Similar to Fried et al. (2025), we solve for the optimal subsidy for green adoption, and emphasize the importance of subsidies for encouraging adoption among lower-income households. In a developing country context, Berkouwer and Dean (2022) estimate the effect of a loan in inducing credit constrained households in Nairobi to switch to a more energy-efficient cookstove. Allcott et al. (2015) argue that subsidies are often poorly targeted−that is, they are largely taken up by wealthy, less credit-constrained, environmentalist homeowners. In our empirical section, we study a similar question: Is the treatment effect of a subsidy larger or smaller for credit-constrained consumers? To do this, we follow the approach of Muehlegger and Rapson (2022) in estimating the demand response of electric vehicle purchases to an EV subsidy, and extend their findings by studying the interaction between the subsidy effect and credit score.

Roadmap. Section 2 develops a two-period model with borrowing constraints and shows that the targeting principle fails when their is an active extensive margin of adoption that only subsidies can move. Section 3 extends this model to an infinite horizon with income and asset dynamics that can be mapped to the data. Section 4 describes the data and empirical strategy for investigating the relationship between consumer credit and subsidy effectiveness, and use our reduced-form estimates to discipline the calibration of the quantitative model. Section 5 describes the quantitative results and studies the transition path under optimal taxes and subsidies. Section 6 concludes.

2 A model of technology adoption under financing constraints

2.1 Environment

Our first model features two periods indexed by $t \in \{1, 2\}$, and a unit mass of households indexed by $i \in [0, 1]$. Each household draws a pair of income y, received in both periods, and credit limit κ, from the joint distribution $H(y, \kappa) \in [\underline{y}, \overline{y}] \times \mathbb{R}_+$. Households can save and borrow at rate R up to their idiosyncratic credit limit, and such that consumption $c_{i,t}$ satisfies the subsistence floor $c_{i,t} \geq \underline{c} > 0$ for all i and t. Households discount future utility with the discount factor $\beta = 1/R$. This assumption, together with the flat income profile, shuts down any borrowing motive other than the financing of technological transition.

Two technologies coexist, indexed by j: a brown technology ($j = b$) and a green technology ($j = g$). All households start with the brown technology, but at the beginning of period 1, they have an opportunity to switch permanently to the green one for an upfront fixed cost of $\phi \geq 0$. Throughout, we use electric-vehicle adoption as a running example. The fixed adoption cost would therefore represent the downpayment required to buy an electric vehicle (EV), the cost of home-charging installation, etc.

Aggregate pollution is a function of the share of the population that uses the brown technology, $K(1 - P)$, where $P \in [0, 1]$ is the fraction of green users. The pollution function $K(\cdot)$ is differentiable and strictly increasing over $[0, 1]$. Households take aggregate pollution as given and do not internalize the effect of their adoption choice on aggregate emissions. Their preferences are separable in consumption, pollution, and idiosyncratic tastes for technologies, and are represented by the utility function

$$U_{ij} = u(c_{i1}) + \beta u(c_{i2}) - (1 + \beta)z(K) + \varepsilon_{ij} \tag{1}$$

where $u(c)$ is strictly increasing and concave over $[\overline{c}, +\infty)$, $z(K)$ is differentiable and strictly increasing in K, and ε_{ij} is an idiosyncratic taste shock over technologies $j \in \{b, g\}$ that is i.i.d. and drawn from a type-1 extreme value distribution with scale parameter σ.

The government must finance exogenous expenditures $G \geq 0$, and controls three fiscal instruments: a tax on brown usage τ levied in each period on households who have not adopted the green technology, a green subsidy s that is paid at the time of adoption, and a proportional consumption tax $\tau_c \in [0, 1]$ levied in both periods. We assume that fiscal policies satisfy the following admissibility restrictions.

Assumption 1 (Admissibility) *The policy vector* (τ, s, τ_c) *satifies* $0 \leq \tau \leq \underline{y} - (1 + \tau_c)\underline{c}$ *and* $0 \leq s \leq \phi - \underline{y} + (1 + \tau_c)\underline{c}$.

The first restriction rules out pollution subsidies and ensures that every household has enough liquid resources to pay the brown tax. The second restriction rules out taxes on adoption of the green technology and that the green technology is weakly unaffordable for some households.

We further assume that credit limits are tighter than the natural borrowing limit of the poorest household, so that no one would lend an amount that the poorest houseold could not afford to

repay without going under its subsistence level of consumption. This bound on credit constraints prevents the possibility of default and is formally stated in the following assumption.

Assumption 2 (Natural borrowing limit) *Credit limits κ satisfy $\kappa \leq (\underline{y} - (1+\tau_c)\underline{c})/R$.*

Household i chooses consumption, savings, and technology to solve

$$\begin{aligned} V_i(y_i, \kappa_i; K) \equiv \max_{\substack{j\in\{b,g\},\, a_i,\\ c_{i1}, c_{i2} \geq \underline{c}}} \quad & u(c_{i1}) + \beta u(c_{i2}) - (1+\beta) z(K) + \varepsilon_{ij} \\ \text{s.t.} \quad & (1+\tau_c) c_{i1} + a_i = y_i - \tau\, \mathbb{I}\{j = b\} - (\phi - s)\, \mathbb{I}\{j = g\}, \\ & (1+\tau_c) c_{i2} = y_i + R a_i - \tau\, \mathbb{I}\{j = b\}, \\ & a_i' \geq -\kappa_i \end{aligned} \tag{2}$$

2.2 Policy margins

We want to compare the benefits and costs of taxes τ and subsidies s in correcting the environmental externality that arises from technology choices. These policies affect households through two distinct margins. Both taxes and subsidies change the relative value of technologies for households who are financially able to choose between the two technologies. The subsidy, in addition, changes the set of households who can choose at all.

To see this asymetry, let us first characterize households' consumption-savings choices. Given the flat income process and the assumption that $\beta R = 1$, we show in the following Lemma that households who stay brown and do not adopt the green technology never borrow.

Lemma 3 *Under assumption 1, a household that stays brown consumes*

$$c_b = \frac{y - \tau}{1 + \tau_c} \geq \underline{c}$$

in both periods, and its lifetime utility from consumption is $V_b(y) = (1+\beta) u(c_b)$. This choice is independent of credit limit κ.

On the contrary, households who switch to the green technology want to borrow, as they pay an additional $\phi - s \geq 0$ in period 1 and want to spread this cost over the two periods. As $\beta R = 1$, an unconstrained household's Euler equation implies a flat consumption plan, such that

$$c^u_{1,g} = c^u_{2,g} = c^u_g = \frac{y - \frac{R}{1+R}(\phi - s)}{1 + \tau_c}, \qquad a^u = -\frac{\phi - s}{1+R} < 0 \tag{3}$$

where c^u_g denotes the unconstrained level of consumption of a green adopter. This consumption-savings plan is feasible only if $\kappa \geq \kappa^u = (\phi - s)/(1+R)$. A second category of households can adopt the green technology but cannot borrow enough to reach the optimal unconstrained consumption plan. For such households, indexed by c (for "constrained"), the credit limit binds at $a = \kappa$, and

consumption is increasing over time:

$$c^c_{1,g} = \frac{y + s - \phi + \kappa}{1 + \tau_c} < c^c_{2,g} = \frac{y - R\kappa}{1 + \tau_c} \tag{4}$$

Finally, a third caqtegory of households face a constraint that is too tight for them to afford the switch to the green technology. Even at maximal borrowing, consumption in period 1 would fall below the subsistence threshold if they were to adopt the green technology. Adoption requires income y and maximum credit κ to be sufficient to pay for the cost of switching and the subsistence floor of consumption, that is

$$y \geq y^*(\kappa) = \phi - s + (1 + \tau_c)\underline{c} - \kappa \tag{5}$$

where $y^*(\kappa)$ represents the level of income needed for the green option to be affordable, conditional on credit limit κ. It traces the boundary of the population that has access to the green technology. This boundary expands as the credit constraint relaxes.

Comparative statics of this boundary of affordability with respect to policy shows the key asymetry between brown taxes and green subsidies:

$$\frac{\partial y^*}{\partial s} = -1, \quad \frac{\partial y^*}{\partial \tau} = 0 \tag{6}$$

A higher subsidy expands the set of households who have access to the green option, while the brown tax has not direct effect on this set. Both affect the intensive margin of switchers by making the green technology relatively more attractive, but only the subsidy makes this technology available for a larger fraction of households. The household types experiencing these effects consists of those who are potentially constrained by the credit margin of adoption, defined as

$$K^* = \{\kappa : y^*(\kappa) \in (\underline{y}, \overline{y})\}$$

2.3 Individual and aggregate adoption

Let $\Delta(y, \kappa) = V_g(y, \kappa) - V_b(y)$ denote the systematic utility gain from adopting, i.e. before accounting for taste shocks. With T1EV shocks with scale parameter σ, a household that can afford the green technology adopts with probability $\Lambda(\Delta/\sigma)$, where $\Lambda(x) = (1 + e^{-x})^{-1}$. The probability that a household of type (y, κ) adopts the green technology is

$$p(y, \kappa) = \mathbb{I}\{y \geq y^*(\kappa)\}\ \Lambda\big(\Delta(y, \kappa)/\sigma\big), \tag{7}$$

This probability shows the two margins that affect adoption: access, equivalent to the condition $y \geq y^*(\kappa)$, and the value of switching conditional on access, determined by the systematic utility gain $\Delta(y, \kappa)$.

Aggregate adoption is obtained by integrating switching probabilities across types:

$$P = \iint p(y,\kappa)\, dH(y,\kappa). \tag{8}$$

Let G denote the marginal distribution of credit limit κ and $F(\cdot \mid \kappa)$ the distribution of income conditional on κ. The following proposition derives the aggregate adoption response to subsidies and taxes.

Proposition 1 (Aggregate adoption responses) *Suppose that for every $\kappa \in K^*$, $F(\cdot \mid \kappa)$ admits a density $f(\cdot \mid \kappa)$ that is continuous at $y^*(\kappa)$. Then*

$$\frac{\partial P}{\partial s} = \underbrace{\iint \mathbb{I}\{y \geq y^*(\kappa)\} \frac{1}{\sigma} \Lambda'\left(\frac{\Delta(y,\kappa)}{\sigma}\right) \frac{u'(c_1^g(y,\kappa))}{1+\tau_c}\, dH(y,\kappa)}_{\textit{intensive margin}} + \underbrace{\int_{K^*} \Lambda\left(\frac{\Delta(y^*(\kappa),\kappa)}{\sigma}\right) f\big(y^*(\kappa) \mid \kappa\big)\, dG(\kappa)}_{\textit{extensive margin}}, \tag{9}$$

$$\frac{\partial P}{\partial \tau} = \underbrace{\iint \mathbb{I}\{y \geq y^*(\kappa)\} \frac{1}{\sigma} \Lambda'\left(\frac{\Delta(y,\kappa)}{\sigma}\right) \frac{(1+\beta)u'(c_b(y))}{1+\tau_c}\, dH(y,\kappa)}_{\textit{intensive margin}}. \tag{10}$$

Proposition 1 presents a core asymmetry between taxes and subsidies in the presence of income heterogeneity and liquidity constraints. Both instruments affect green adoption for households who already have access to both technologies. However, the subsidy creates an additional effect at the extensive margin by expanding the set of households who can effectively choose between the two options.

2.4 Planner's problem

In this section, we study the problem of a Ramsey planner who chooses the optimal mix of fiscal instruments to maximize social welfare.

Welfare The expected indirect utility of a household of type (y,κ), before idiosyncratic tastes are realized, is

$$V(y,\kappa) = \mathbb{E}_\varepsilon[\max\{V_b(y) + \varepsilon_b,\, V_g(y,\kappa) + \varepsilon_g\}] = V_b(y) + \mathbb{I}\{y \geq y^*(\kappa)\}\Omega(y,\kappa), \tag{11}$$

where $\Omega(y,\kappa) = \sigma \log(1 + e^{\Delta/\sigma}) > 0$ is the option value of getting access to the green technology. We can use this expression to define social welfare, net of environmental damages, as

$$W = \iint \big[V_b(y) + \mathbb{I}\{y \geq y^*(\kappa)\}\, \Omega(y,\kappa)\big]\, dH \;-\; (1+\beta)\, z\big(K(1-P)\big), \tag{12}$$

Government budget constraint Let $\bar{c}_t = \int\int [p(y,\kappa)c_t^g + (1-p(y,\kappa))c_b]dH(y,\kappa)$ denote aggregate consumption in period t. The government gets revenues from the consumption tax and the brown tax, while its expenditures consist of an exogenous expenditures flow G and the green subsidy paid out to the P switchers. The intertemporal government budget constraint (IGBC) is

$$(\tau_c, \tau, s) = \tau_c \left(\bar{c}_1 + \frac{\bar{c}_2}{R}\right) + \left(1 + \frac{1}{R}\right)\tau(1-P) - sP = \left(1 + \frac{1}{R}\right) G. \tag{13}$$

Budget-balanced policy margins We are interested in the welfare effects of changes in τ and s that satisfy the government budget constraint. Let

$$\mathcal{W}_x = W_x - \frac{\mathcal{R}_x}{\mathcal{R}_{\tau_c}} W_{\tau_c}, \qquad x \in \{\tau, s\}$$

denote the marginal welfare gain from a change in brown taxes or green subsidies net of the effect of the consumption tax adjustment needed to satisfy the government budget constraint. We assume that the consumption tax is on the upward-sloping part of its Laffer curve, that is $\mathcal{R}_{\tau_c} > 0$. This assumption rules out cases in which increasing green subsidies or reducing brown taxes would trigger a reduction in consumption taxes to balance the government's budget.

2.5 Targeting principle

2.5.1 Linear benchmark

We first derive a failure of the targeting principle in the special case of linear utlity, $u(c) = c$. Together with $\beta R = 1$, linear utility implies that consumption-savings choices are not necessarily unique. Any path that produces the same present value of consumption delivers the same utility. We further assume in this special case that damages are linear, such that $(1+\beta)z(K(1-P)) = z_0 - \mathcal{E}P$, where $\mathcal{E} \in \mathbb{R}^+$.

Conditional on adoption being feasible, the borrowing constraint does not affect utility. The only channel through which the borrowing limit bites is the affordability constraint: the new technology may still be unaffordable because of the inability of pulling future income forward. The feasibilty cutoff $y^*(\kappa)$ remains unchanged and is defined by (5). Adoption is still heterogenous at the extensive margin: some households have the option to switch while others do not. But the intensive margin is common across income and credit groups, as the systematic change in utility from adoption is no longer a function of group-specific characteristics:

$$\Delta = \frac{s - \phi + (1+\beta)\tau}{1+\tau_c} \tag{14}$$

With linear utility social welfare reduces to

$$W = P\left[\frac{\Omega(\Delta)}{p(\Delta)} - \Delta + \mathcal{E} - \phi\right] + (1+\beta)\mathbb{E}[y] - \left(1 + \frac{1}{R}\right) G \tag{15}$$

where $\Omega(\Delta)/p(\Delta) - \Delta$ represents the value of being able to choose according to one's idiosyncratic tastes. It is the private value of having access to the new technology, net of the systematic surplus from adoption. The second component inside brackets, $(\mathcal{E} - \phi)$, represents the net social benefit from adoption, which consists of its environmental benefits net of real adoption costs. Policy affects welfare only through the pair (Δ, y^*). Utility differential Δ governs the intensive margin of adoption, conditional on access, while y^* draws the boundary of access.

The targeting principle states that the optimal policy to correct an externality is to tax the externality-generating activity rather than subsidize its substitutes or use a mix of both instruments. To show that this principle breaks in this setting, we consider an initial interior policy $(s, \tau, \tau_c) \in [0, \overline{s}) \times (0, \overline{\tau}) \times (0, 1)$. We introduce a perturbation $(ds, d\tau, d\tau_c)$ such that

$$ds > 0, \qquad dP = 0, \qquad d\mathcal{R} = 0 \tag{16}$$

This perturbation introduces a strictly positive subsidy, while leaving the externality reduction P and government revenues $\mathcal{R}$ unchanged. The welfare effects from such a policy change can only spring from the access margin and fiscal incidence. In order to characterize this policy perturbation, we define the aggregate present value of consumption if every household adopts the new technology as $\widetilde{C} = ([1 + \beta]\mathbb{E}[y] + s - \phi)/(1 + \tau_c)$. The following Lemma describes how taxes (τ, τ_c) and the affordability threshold y^* respond to the subsidy increase.

Lemma 4 *If $\widetilde{C} > \underline{c}$, the budget-balanced path that marginally raises s and keeps P constant is unique, with*

$$\frac{d\tau_c}{ds} = \frac{(1-P)P_s^{\text{ext}} + P_s^{\text{int}}}{(1-P)\underline{c}P_s^{\text{ext}} + \widetilde{C}P_s^{\text{int}}} > 0, \qquad \frac{d\tau}{ds} = -\frac{(C - P\underline{c})P_s^{\text{ext}} + (C - P\Delta)P_s^{\text{int}}}{(1+\beta)\left[(1-P)\underline{c}P_s^{\text{ext}} + \widetilde{C}, P_s^{\text{int}}\right]} < 0,$$

and it lowers the affordability threshold by

$$-\frac{dy^*}{ds} = 1 - \underline{c}\frac{d\tau_c}{ds} = \frac{(\widetilde{C} - \underline{c})P_s^{\text{int}}}{(1-P)\underline{c}P_s^{\text{ext}} + \widetilde{C}P_s^{\text{int}}} \in (0, 1).$$

Lemma 4 shows that a subsidy increase that preserves budget balance and aggregate adoption leads to a unique policy path with a decrease in the brown tax τ and a compensating increase in the consumption tax τ_c. The former is necessary to keep adoption unchanged, while the latter is the only way to balance the budget after substituting a subsidy for a tax. This substitution expands the extensive margin of access to the green technology, as demonstrated by the reduction of the affordability threshold y^*. The following proposition characterizes the marginal welfare effect of a movement along this path.

Proposition 2 *At any admissible policy $(s, \tau, \tau_c) \in [0, \overline{s}) \times (0, \overline{\tau}) \times (0, 1)$ with $\widetilde{C} > \underline{c}$, the first-order welfare change from a marginal increase in the green subsidy $ds > 0$ that leaves adoption P and*

government revenue $\mathcal{R}$ unchanged is

$$\frac{dW}{ds} = m_0\,\Omega(\Delta)\Big(-\frac{dy^*}{ds}\Big)$$

which is strictly positive if $m_0 > 0$.

Proposition 2 shows that, as long as a strictly positive mass of household is priced out of green adoption at the margin, a higher subsidy is welfare-improving even at constant aggregate adoption. This policy change keeps the number of adopters constant, but changes their composition. It reduces adoption at intensive margin and expands it at the extensive margin. The swap is welfare-improving because it trades marginal adopters at the intensive margins, who are close to indifferent between the two technologies, for priced-out households who have a high preference for the green technology. This result breaks the targeting principle by showing that there is no interior optimum with a strictly positive brown tax and a zero subsidy. Any such point would benefit from a decrease in the brown tax and an increase in the subsidy.

2.5.2 General case

Let us now turn to the general case with concave utility and non-linear environmental damages. Let $\mu^j(y,\kappa), j \in \{b,g\}$ denote the Lagrange multiplier on the borrowing constraint of the household problem (2), conditional on the household income and credit type (y,κ) and the technology choice j. As shown in Lemma 3, $\beta R = 1$ implies that households who stay brown never borrow, so $\mu^b(y,\kappa) = 0$ for all (y,κ). The first-order condition of the household problem implies that, for households who adopt the green technology, this multiplier is equal to

$$\mu^g(y,\kappa) = \frac{u'(c_1^g(y,\kappa)) - u'(c_2^g(y,\kappa))}{1+\tau_c} \geq 0 \tag{17}$$

As in the linear case, we consider the a policy perturbation that increases the green subsidy by $ds > 0$ while keeping the aggregate share of switchers P and government revenues $\mathcal{R}$ unchanged. The welfare effect of such a reform is characterized in the following proposition.

Proposition 3 *Suppose the economy has an admissible policy $(s,\tau,\tau_c) \in [0,\overline{s}) \times (0,\overline{\tau}) \times (0,1)$ and the matrix $\left(\begin{smallmatrix} P_\tau & P_{\tau_c} \\ \tau & \tau_c \end{smallmatrix}\right)$ is nonsingular. Then the welfare effect of this policy change is The welfare effect of the substitution is*

$$\frac{dW}{ds} = \underbrace{\iint p(y,\kappa)\,\mu^g(y,\kappa)\left(1 - c_1^g(y,\kappa)\,\frac{d\tau_c}{ds}\right)dH(y,\kappa)}_{L\ \textit{(liquidity)}} + \underbrace{\left(1 - \underline{c}\,\frac{d\tau_c}{ds}\right)\mathcal{X}_\Omega}_{\mathcal{A}\ \textit{(access)}} + \underbrace{\iint \omega\, g\, dH(y,\kappa)}_{\Xi\ \textit{(redistribution)}},$$

where

$$\mathcal{X}_\Omega = \int_{K^*} \Omega(y^*(\kappa),\kappa) f(y^*(\kappa)|\kappa) dG(\kappa)$$

is the option value of marginal entrants, weighted by their mass at the access cutoff, $\omega = u'(c_b)/(1+\tau_c)$ *for a stayer,* $\omega = u'(c_2^g)/(1+\tau_c)$ *for an adopter, and* g *is the net dollar gain from the reform, that is* $g = 1 - (c_1^g + \beta c_2^g)\, d\tau_c/ds$ *for an adopter and* $g = -(1+\beta)\, d\tau/ds - (c_1^b + \beta c_2^b)\, d\tau_c/ds$ *for a stayer.*

Proposition 3 identifies three welfare effects of higher subsidies. First, the liquidity term L represents the benefits from the relaxation of the borrowing constraint for households that could not smooth consumption optimally. It reverts to a loss when the subsidy requires such a high consumption tax increase that it actually tightens liquidity constraints. The second term on the right-hand side of the welfare change, $\mathcal{A}$, values changes in the set of households that can afford adoption. Subsidies extend access to new households that may have strong preferences for green vehicles, which represents the extensive margin gain from higher subsidies. This welfare channel turns negative if consumption taxes rise so much in response to the subsidy that they actually tighten liquidity constraints and reduce access to the green technology. Finally, the last term, Ξ, captures redistribution. It is positive when gains from the reform accrue disproportionately to households with high marginal utility.

3 Quantitative model

We now extend the model to an infinite horizon framework with a gradual transition from brown to green technology. Households face persistent idiosyncratic income risk and make consumption-savings decisions subject to a borrowing constraint. Brown households periodically receive opportunities to replace their polluting cars with electric vehicles. The purchase of this new vehicle is financed through a loan with a down-payment requirement. The government uses the same instruments as in the two-period model: a tax on brown usage, a switching subsidy, and a consumption tax to balance its budget. As in the two-period model, this setting introduces income heterogeneity and liquidity requirements that limit access to the new technology, allowing subsidies to affect adoption at the extensive margin.

3.1 Environment

Time is discrete and periods are quarters. There is a unit mass of households indexed by $i \in [0,1]$. Individual log income follows the AR(1) process $y_{i,t} = \rho_y y_{i,t-1} + \varepsilon_{i,t}^y$ with $\varepsilon_{i,t}^y \sim \mathcal{N}(0, \sigma_y)$. Mean income is normalized to one. Households can save and borrow at an interest rate $r > 0$, subject to the asset constraint $a > \underline{a}$. They are also characterized by the same technology state $j \in \{b, g\}$ as in the two-period model, and by an idiosyncratic credit type d that determines the down payment required when purchasing a new car. A household's idiosyncratic state is summarized by (a, y, j, d).

Households have constant relative risk aversion (CRRA) preferences over consumption, repre-

sented by the utility function

$$u(c) = \begin{cases} c^{1-\gamma}/(1-\gamma) & \text{if } \gamma \neq 1 \\ \log(c) & \text{if } \gamma = 1 \end{cases}.$$

Households discount future utility flows by a factor β, with $\beta(1+r) < 1$. At each replacement opportunity, households draw idiosyncratic taste shocks for the two technologies from a T1EV distribution with scale parameter σ. Households using the green technology also receive a common flow utility benefit A_t. This time-varying benefit increases over time to capture improvements in the quality of the green technology. Environmental damages generate the same disutility δK_t to every household, where K_t is the amount of emissions in quarter t. Emissions are a linear function of the aggregate fraction of brown users $1 - P_t$, such that $K_t = e(1 - P_t)$.

Households receive replacement opportunities with probability $\omega \in [0, 1]$ per quarter. The total cost of switching from brown to green is $\phi > 0$, while we assume for simplicity that replacing a vehicle with a new one of the same type is costless. Once a household adopts the green technology, it remains green permanently. When purchasing a green vehicle, a brown household can borrow up to $B_t(d)$ against the vehicle:

$$B_t(d) = \min\{\phi - d, \phi - s_t\},$$

where d is a down-payment requirement, which characterizes households' idiosyncratic credit type, and s_t is the green adoption subsidy. The amount due at purchase is therefore $\max\{d - s_t, 0\}$. The subsidy lowers the cash requirement until it reaches the down payment amount d.

Households face a consumption floor $\underline{c}(y)$ that increases with income y. This floor captures spending commitments, such as mortgage payments or childcare expenses, that cannot easily be reduced to finance a car purchase.

The government finances an exogenous spending flow G and a subsidy s_t paid to households when they first purchase a green car. Revenue comes from a quarterly tax τ_t on brown usage and a proportional consumption tax at rate τ_t^c.

3.2 Household problem

Define a household's cash on hand as $x = (1+r)a + y$. A green household chooses its consumption and savings by solving the Bellman equation

$$V_t^g(a, y, d) = \max_{a'} \; u\Big(\tfrac{x - a'}{1+\tau_t^c}\Big) + A_t + \beta\, \mathbb{E}\big[V_{t+1}^g(a', y', d) \mid y\big], \quad a' \geq \min\{\underline{a} - \bar{B}(d), a\},\; c \geq \underline{c}(y). \tag{18}$$

Where $B_t(d)$ is the maximum loan on the vehicle switching cost that a household can take. A brown household's lifetime utility if it keeps its polluting vehicle is

$$v_t^b(a, y, d) = \max_{a' \geq \underline{a}} \; u\Big(\tfrac{x - \tau_t - a'}{1+\tau_t^c}\Big) + \beta\, \mathbb{E}\big[V_{t+1}^b(a', y', d) \mid y\big], \quad c \geq \underline{c}(y). \tag{19}$$

This value applies both to households that receive no replacement opportunity and to those that receive one but retain their polluting vehicles. Note that for the latter, the value $v_t^b(a, y, d)$ does not include the taste shock associated with this choice, which will be added below. A brown household that switches to the green technology solves a similar Bellman equation

$$v_t^s(a, y, d) = \max_{a' \geq \underline{a} - B_t(d)} u\Big(\tfrac{x-(\phi-s_t)-a'}{1+\tau_t^c}\Big) + A_t + \beta\, \mathbb{E}\big[V_{t+1}^g(a', y', d) \mid y\big], \quad c \geq \underline{c}(y). \tag{20}$$

Brown households considering a switch to the green technology face an access threshold similar to that described in Section 2, but expressed in terms of cash on hand, x. Switching is feasible if and only if

$$x \;\geq\; \bar{x}_t(y, d) \;\equiv\; \underline{a} + \max\{d - s_t, 0\} + (1 + \tau_t^c)\, \underline{c}(y). \tag{21}$$

As in the two-period model, the subsidy s_t lowers the access threshold, while the brown tax τ_t has no direct effect on access. Financing the subsidy through consumption taxes tightens the access constraint because a higher consumption tax increases the liquidity needed to meet the consumption floor. The main difference from the two-period model lies in the variable subject to the access constraint, rather than in the threshold itself. Because access now depends on cash on hand rather than income alone, households can save to reach the threshold and adopt the green technology.

The systematic surplus from the switch is $\Delta_t(a, y, d) = v_t^s - v_t^b$. This surplus is increasing in the quality of the green technology, A_t. Conditional on a replacement opportunity and financial access, a brown household adopts the green technology with probability $\Lambda(\Delta_t/\sigma)$. As in the two-period model, this conditional adoption probability captures the intensive margin of adoption. The unconditional probability of adoption is $\omega \mathbb{I}\{x \geq \bar{x}_t\}\Lambda(\Delta_t/\sigma)$. The value function for a brown household before replacement opportunities and taste shocks are realized is:

$$V_t^b(a, y, d) = v_t^b(a, y, d) + \omega\, \mathbb{I}\{x \geq \bar{x}_t(y, d)\}\, \Omega_t(a, y, d), \qquad \Omega_t = \sigma \log\big(1 + e^{\Delta_t/\sigma}\big), \tag{22}$$

where Ω_t is the option value of access to the green technology.

3.3 Adoption and aggregate transition

Let μ_t^b denote the distribution of brown households over states (a, y, d) at the beginning of period t, and let μ_t^g denote the corresponding distribution of green households. These distributions satisfy $\int d\mu_t^g = P_t$ and $\int d\mu_t^b = 1 - P_t$. We define

$$p_t(a, y, d) = \mathbb{I}\{x \geq \bar{x}_t(y, d)\}\Lambda\left(\frac{\Delta_t(a, y, d)}{\sigma}\right),$$

so that $\omega p_t(a, y, d)$ is the unconditional probability of adoption in period t. The mass of new adopters in period t, conditional on an initial distribution μ_0^b and an initial share of green households

P_0, is

$$N_t = \omega \int p_t(a, y, d)\, d\mu_t^b, \qquad P_{t+1} = P_t + N_t,$$

The green share is weakly increasing over time. The speed of diffusion depends on the arrival rate of switching opportunities, the distribution of financial resources among brown households, and the policy-induced switching probabilities.

Let a_t^b, a_t^s, a_t^g denote the savings policies of brown stayers, switchers, and incumbent green households. For each $j \in \{b, s, g\}$, define

$$Q_t^j(S \mid a, y, d) = \int \mathbb{I}\Big\{(a_t^j(a, y, d), y', d) \in S\Big\}\, f(y'|y) d(y'), \tag{23}$$

where S is a measurable subset of the state space. Thus, $Q_t^j(S \mid a, y, d)$ is the probability being in S next period under policy a_t^j. This formulation allows us to define the laws of motion of μ_t^b and μ_t^g as

$$\mu_{t+1}^b(S) = \int [1 - \omega p_t(a, y, d)]\, Q_t^b(S \mid a, y, d)\, d\mu_t^b(a, y, d) \tag{24}$$

and

$$\mu_{t+1}^g(S) = \int Q_t^g(S \mid a, y, d)\, d\mu_t^g(a, y, d) + \int \omega p_t(a, y, d)\, Q_t^s(S \mid a, y, d)\, d\mu_t^b(a, y, d). \tag{25}$$

3.4 Government

Aggregate consumption can be expressed in terms of households' choice-specific policy functions. Let $c_t^g(a, y, d)$ denote the consumption of a green household, $c_t^b(a, y, d)$ the consumption of a brown household that remains brown, and $c_t^s(a, y, d)$ consumption of a brown household that adopts the green technology. Then

$$\bar{c}_t = \int c_t^g(a, y, d)\, d\mu_t^g(a, y, d) + \int \Big[(1 - \omega p_t(a, y, d))\, c_t^b(a, y, d) + \omega p_t(a, y, d) c_t^s(a, y, d)\Big]\, d\mu_t^b(a, y, d).$$

Since the brown tax is paid only by households that remain brown after the switching decision, the brown tax base is $1 - P_{t+1}$. In period t, the subsidy is paid only to the N_t new adopters. The government also collects the proportional consumption tax on aggregate consumption, $\bar{c}_t$. Government revenue in period t is therefore $\mathcal{R}_t = \tau_t^c \bar{c}_t + \tau_t(1 - P_{t+1}) - s_t N_t$. The government must finance an exogenous spending flow G_t, and we assume that it balances its budget each period, so the period budget constraint is

$$\tau_t^c \bar{c}_t + \tau_t(1 - P_{t+1}) - s_t N_t = G_t \tag{26}$$

3.5 Planner problem

Taking the initial distributions of households across states $(a, y, d)\mu_0^b$ and μ_0^g and across technologies P_0, utilitarian social welfare, net of environmental damages, is

$$\mathcal{W} = \int V_0^b(a, y, d)\, d\mu_0^b(a, y, d) + \int V_0^g(a, y, d)\, d\mu_0^g(a, y, d) - \sum_{t=0}^{\infty} \beta^t\, \delta e\, (1 - P_t) \tag{27}$$

In the quantitative section, we will consider the problem of a planner that chooses constant paths of brown taxes τ and green subsidies s to maximize this welfare function, subject to value functions and distributions $\{\mu_t^b, \mu_t^g\}$ such that households solve the problem defined in 3.2, laws of motion (24) and (25), the government budget constraint in (26), which determines endogenous the consumption-tax path, and finally an environmental target P_T at horizon T. This exercise allows us to compare the welfare effects of taxes and subsidies in delivering the same environmental goal, and to compute the optimal mix of instruments that can reach that objective.

4 A calibrated quantitative model of technology switching

In this section, we augment our theoretical model to quantitatively study optimal policy in the presence of liquidity constraints. First, we present empirical evidence on liquidity constraints amplifying the treatment effect of electric vehicle subsidies in California. We use the estimated treatment effect of subsidies on EV purchases as a calibration target for the model, and show that the model matches untargeted moments that indicate a larger treatment effect for lower-credit areas. Finally, we construct optimal policy paths targeting a prescribed green transition and show that the usage of subsidies substantially increases welfare, partly through pushing households across the extensive financing margin and partly through distributional benefits.

4.1 Empirical evidence: Liquidity constraints amplify the effect of subsidies

To test the empirical relevance of liquidity constraints on policy efficacy, we use the natural laboratory of electric vehicle (EV) subsidies. Specifically, we follow Muehlegger and Rapson (2022) in studying the effect of the Enhanced Fleet Modernization Program (EFMP), an EV subsidy program in California implemented in 2015, on EV sales. The program provides cash incentives to eligible households to retire older, high-emitting vehicles and purchase a cleaner replacement vehicle. Importantly, incentives are delivered at the point of sale, directly lowering the upfront purchase cost and allowing this case study to speak directly to our model about fixed costs.

The model from Section 2 provides a mechanism where, holding all else fixed, the marginal effect of subsidy dollars on adoption may be larger when households are more likely to be borrowing constrained. To test this mechanism, we use credit scores as our empirical proxy for liquidity constraints. We interact pre-period credit scores with the subsidy treatment variable. We find that a decrease in the ZIP-level credit score leads to a larger treatment effect of the EV subsidy

on sales, consistent with predictions from the theory: Individuals with lower credit scores are more responsive to a subsidy. Then, we use these empirical results to calibrate a quantitative model of technology switching with different credit types in the remainder of this section.

4.1.1 Data and institutional background

Launched in 2015, the Scrap-and-Replace program in the Enhanced Fleet Modernization Program subsidizes low- and middle- income households to buy cleaner vehicles in exchange for scrapping older, dirtier vehicles. The 2015 version of this program provided up to $4,500 to lower-income drivers to buy a cleaner vehicle, with an additional $5,000 top-up incentive through the Clean Cars 4 All program for participants living near a disadvantaged community, classified by California's Environmental Protection Agency. The initial EFMP was available only for residents in the San Joaquin Valley Air Pollution Control District and the South Coast Air Quality Management District (AQMD); these regions expanded to the Bay Area AQMD in 2019, the Sacramento Metro AQMD in 2020, and San Diego County AQMD in 2024. This staggered rollout generates variation in program availability over time.

We use a variety of data sources on EV sales, rebates, and household financial variables. Information on EV sales numbers and shares comes from the California Energy Commission (2026), and rebates through Clean Cars 4 All are reported by the California Air Resources Board (2026). Data on credit scores are obtained through Bricker and Li (2024), who calculate average credit scores by ZIP code using the Federal Reserve Bank of New York Consumer Credit Panel/Equifax data. Finally, income distributions are taken from the American Community Survey 5-year estimates (U.S. Census Bureau (2012)). Income and credit variables are taken from 2012, which is during the pre-period before the policy was implemented in 2015. We choose 2012 as the pre-period year because it is far enough from treatment for relevant variables to be predetermined, and avoids conditioning on ZIP code characteristics measured during the same period as baseline outcomes. Our full sample of ZIP-quarter observations is from 2014-2025, and we perform robustness checks to restricting the sample to 2014-2018 as in Muehlegger and Rapson (2022).

Figure A1 illustrates the geographic distribution of credit scores, income, EV sales, and EFMP rebates. Credit scores and income tend to follow a similar pattern, with higher income areas also generally having higher average credit scores. EV sales are clustered both in high income/credit score locations as well as locations with available EFMP rebates, which would imply a positive effect of rebates on EV sales in areas with lower average income.

4.1.2 Methodology

We follow Muehlegger and Rapson (2022) by regressing EV sales on a measure of EFMP treatment. Our contribution is to also include interactions with the financial variables of income and credit score:

$$Y_{zt} = \beta_T T_{zt}^{\omega} + \beta_{credit}\left(T_{zt}^{\omega} \times credit\ score\right) + \beta_{income}\left(T_{zt}^{\omega} \times income\right) + FEs + \epsilon_{zt} \tag{28}$$

Where Y_{zt} is the inverse hyperbolic sine of EV quantity in ZIP z at time t, and T^{ω}_{zt} is the treatment variable that is described further below. Fixed effects specifications include ZIP, quarter, and ZIP $\times$ quarter. Additionally, we include specifications that interact treatment variable T^{ω}_{zt} with the financial variables of average credit score and/or income in an ZIP code, in order to measure the differential effect of the subsidy on households with different financial situations. The interaction of the treatment variable with the credit score β_{credit} is our main coefficient of interest, which is the effect of a borrowing constraint − proxied by a credit score − on the treatment. The theory predicts $\beta_{credit} < 0$: subsidies should have higher effect in lower credit areas. We also control for the interaction between the treatment variable and income to ensure that we are measuring specific credit score effects, rather than conflating the borrowing constraint with income dynamics.

Ideally, the treatment is the DiD estimate $\hat{\beta} = (Q_{11} - Q_{10}) - (Q_{01} - Q_{00})$, where Q_{ap} is the quantity of EVs, $a = 1$ in treated regions ($a = 0$ is untreated), and $p = 1$ in the post-period ($p = 0$ in pre-period). Because not all individuals are eligible for the means-tested treatment, we want to obtain the ToT by scaling the ITT estimate by the fraction of treatment-eligible households in the population.

We consider multiple treatment variables, two to replicate the Muehlegger and Rapson (2022) realized treatment (T^{h}_{zt}, T^{l}_{zt} as defined in the Muehlegger and Rapson (2022) paper) and one based on predicted treatment exposure. We cannot observe the fraction of eligible households because in addition to being income-tested, individuals must also have a car to scrap to be eligible. Because both treatment variables used in Muehlegger and Rapson (2022) rely on *realized* treatment quantities, which may be correlated with the income and credit score characteristics we study here, we construct a new treatment variable that depends on *predicted* treatment from the statutory policy. This treatment variable largely replicates the main results in Muehlegger and Rapson (2022), as shown in Table A1.

Predicted treatment T^{pred}_{zt}. We construct predicted subsidy exposure using statutory award amounts and pre-period household income distributions. Let r denote household income expressed as a percentage of the federal poverty level (FPL) applicable to the household's size and income-reference year. Using pre-period ACS data, we separate households in ZIP code z into three income brackets based on how far they are from the FPL:

$$\pi_{z,225} = \Pr_{z}(r \leq 225), \tag{29}$$

$$\pi_{z,300} = \Pr_{z}(225 < r \leq 300), \tag{30}$$

$$\pi_{z,400} = \Pr_{z}(300 < r \leq 400). \tag{31}$$

These groups are held fixed throughout the sample period. Because the brackets are mutually exclusive, $\pi_{z,225} + \pi_{z,300}$ is the share of households with income at or below 300% of FPL, while the sum of all three shares gives the share at or below 400% of FPL and are thus eligible for the policy.

Next, we consider the delineations of disadvantaged community status in the program rules and

incorporate a revision to the award and eligibility schedule that began in 2023. Let D_z indicate disadvantaged community status under the program's geographic eligibility rules, and let A_{zt} indicate whether the program is available in ZIP code z in quarter t. In 2023Q1, subsidy awards increased to $10,000 outside disadvantaged communities and $12,000 within disadvantaged communities, and the income eligibility ceiling fell from 400% to 300% of FPL (California Air Resources Board (2022b)).

Our predicted exposure treatment variable, measured in thousands of dollars per household, is

$$\begin{aligned} T_{zt}^{\text{pred}} &= S_{zt}^{\text{pred}} \\ &= A_{zt} \begin{cases} (1 - D_z)\left(4.5\pi_{z,225} + 3.5\pi_{z,300} + 2.5\pi_{z,400}\right) & \\ \quad + D_z\left(9.5\pi_{z,225} + 7.5\pi_{z,300} + 5.5\pi_{z,400}\right), & t < \text{2023Q1}, \\ \left[10(1 - D_z) + 12D_z\right]\left(\pi_{z,225} + \pi_{z,300}\right), & t \geq \text{2023Q1}. \end{cases} \end{aligned} \tag{32}$$

This treatment variable combines the statutory program award schedules with predetermined income shares, so it does not depend on realized participation or payments. The policy exposure of this variable depends on income and ZIP-code designations used in the policy design.

4.1.3 Main results

First, we replicate the results from Muehlegger and Rapson (2022) to validate our sample and treatment variables construction. Table A1 shows the replication results. We are generally able to match the results in Muehlegger and Rapson (2022), shown in the last two columns labeled (MR), although our data differs slightly as we use public ZIP-aggregated data that is available starting in 2014 while Muehlegger and Rapson (2022) use private transaction data starting in 2012. As in Muehlegger and Rapson (2022), the lower bound of the subsidy treatment effect is about a 3% increase in the probability of purchasing an EV for ZIP codes where all buyers are eligible for the program, while the upper bound estimate shows a 12-15% treatment effect per $1,000 subsidy. The effect of our predicted treatment variable falls squarely in the center of the high- and low- weighted treatment variables, which implies that this treatment variable is capturing a similar effect as the others. We will use T^{pred} as our main treatment variable for the remaining analysis, because T^{pred} is the treatment effect of *predicted* policy exposure rather than the treatment effect on *realized* treatment. This is particularly important for our use case because using realized treatments may conflate the effect of credit scores and income on the treatment effect. Since predicted exposure is statutory, we prefer this treatment variable because credit scores and income are not mechanically tied to subsidy take-up.

Table 1 reports the main DiD estimates using the predicted treatment variable and interactions with credit scores and income. Average ZIP-level credit scores and incomes are standardized. All specifications include ZIP and time fixed effects, while Columns (5) and (6) add geographic $\times$ time fixed effects. Column (1) recalls the baseline estimate from Column (3) in the replication A1. Column (2) shows the additional interaction between the predicted treatment variable and the

ZIP-level standardized credit score, with a negative coefficient on the interaction variable: A higher ZIP-level credit score is associated with a *lower* subsidy treatment effect. Column (3) includes the interaction between the treatment variable and income, indicating that higher income ZIP codes are associated with larger treatment effects.

Our main result is in Column (4), which interacts the treatment effect with both incomes and credit scores. Including both income and credit score interactions ensures that the coefficient on the credit score does not incorporate with income effects. A one standard deviation *increase* in a ZIP's credit score leads to a 8.6% *decrease* in the treatment effect of a subsidy. This is slightly *larger* than the baseline subsidy effect of around 6%, which implies that the credit score is an important component of the subsidy response. This result corroborates one of the main mechanisms in our theory: Subsidies are particularly effective for those with binding borrowing constraints, proxied empirically by lower credit scores.

Columns (5) and (6) confirm this result by adding geography by time fixed effects: Column (5) includes a county × year fixed effect, while Column (6) includes a AQMD × year fixed effect. Both recover a similar coefficient on the interaction between the average credit score and the predicted treatment variable, although the pooled treatment variable becomes less precise. These specifications show that the credit gradient survives finer geographic fixed effects, and that the credit gradient is not driven solely by local trends. All specifications in Table 1 recover a positive baseline treatment effect, a negative interaction coefficient between the treatment and credit score, and a positive interaction coefficient between the treatment and income.

Table 1: Differential treatment effects by credit score and income

	(1)	(2)	(3)	(4)	(5)	(6)
T^{pred}	0.0648***	0.0339***	0.0928***	0.0675***	0.0025	0.0077
	(0.0062)	(0.0076)	(0.0071)	(0.0074)	(0.0095)	(0.0098)
T^{pred} × credit score		-0.0377***		-0.0864***	-0.0751***	-0.0764***
		(0.0058)		(0.0079)	(0.0072)	(0.0076)
T^{pred} × income			0.0393***	0.1032***	0.0692***	0.0765***
			(0.0068)	(0.0087)	(0.0085)	(0.0087)
ZIP FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Quarter	Quarter	Quarter	Quarter	Quarter	Quarter
Geographic time FE	No	No	No	No	County×year	AQMD×year
Observations	7.8e+04	7.8e+04	7.8e+04	7.8e+04	7.8e+04	7.8e+04
Adjusted R-squared	0.9236	0.9242	0.9240	0.9260	0.9380	0.9361

Standard errors clustered by ZIP.

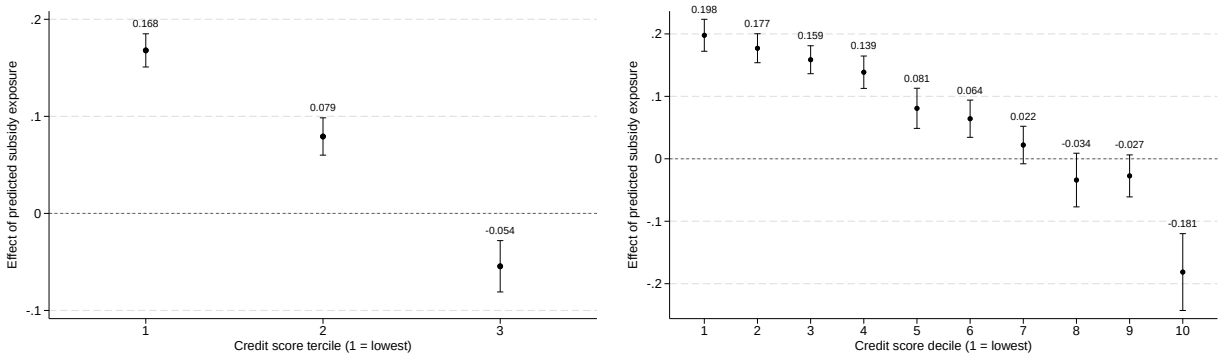
Next, we run a similar regression but include interactions between the treatment variable and credit score quantiles to capture potential nonlinearities. This specification also provides the treatment effects that we will bring to the calibration. We sort ZIP codes into credit score quantiles,

where the lowest quantile has the lowest credit score, and estimate the regression:

$$Y_{zt} = \vec{\beta}_{groups}\big(T_{zt}^{pred} \times \mathbb{I}[credit\ score\ groups]\big) + \beta_{income}\big(T_{zt}^{pred} \times income\big) + \nu_z + \gamma_t + \epsilon_{zt} \quad (33)$$

We omit the baseline uninteracted treatment in this specification, so each coefficient is interpreted as the treatment effect for a particular credit score group netting out the effect on income. Figure 2 visualizes heterogeneity by estimating separate treatment interactions for credit score terciles and deciles. The estimated treatment effect is largest in the lowest credit score group and declines roughly monotonically with credit score. The lowest credit score tercile experiences a treatment effect of 16.8%, compared with a treatment effect of -5.44% for the highest credit score tercile. The treatment effect differences across groups are statistically significant. Similarly, the lowest credit score decile experiences a treatment effect of 19.8%, while the highest credit score decile experiences a treatment effect of -18.1%. Estimating these treatment effects by credit score quantile also confirms that the credit score treatment effect is roughly linear, so the linear functional form used in Equation 28 is a good fit for the data.

Figure 2: Treatment effects by credit score quantile



(a) By credit score tercile

(b) By credit score decile

The main identification concern in these estimates is that low- and high-credit ZIP codes may have followed different EV adoption trends in the absence of the policy. Figure 3 plots the event study treatment effect estimates around the first quarter of program availability. The pre-period estimates do not show a persistent divergence by credit group. After the program rollout, the EV sales probability increases more in low credit ZIP codes than in high credit ZIP codes, as in the reduced form estimates. This timing is consistent with the treatment effect of the subsidy generating the change in EV sales by credit group, rather than an exogeneous trend.

4.1.4 Empirical moments for the quantitative model

The quantitative model is disciplined by three moments from the empirical analysis: EV adoption in never-treated ZIP codes over the sample period, the treatment effect of the subsidy on EV adoption from Table 1, and the total number of subsidies granted. The flow value of the green technology is

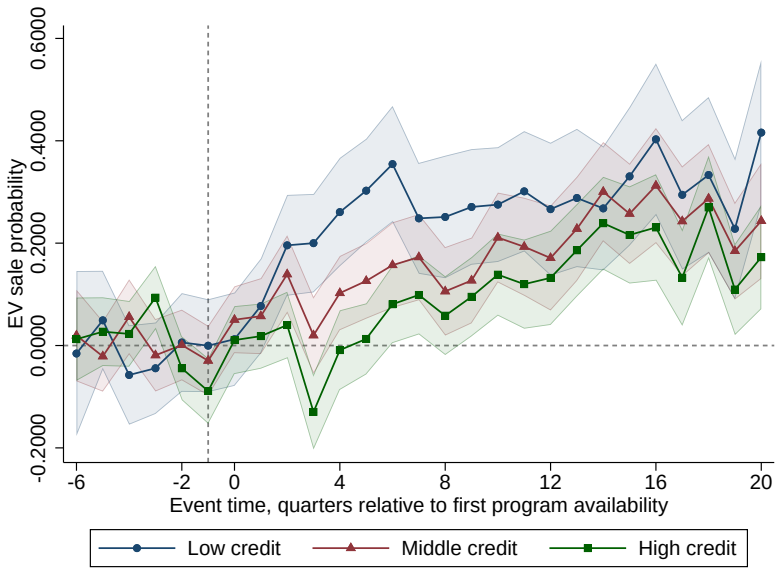


Figure 3: EV adoption event study, by credit group

calibrated to match total green adoption; the T1EV preference parameter that governs switching is calibrated to match the pooled treatment effect of the subsidy; and the subsidy take-up parameter is calibrated to match the number of subsidies granted in the data.

We use the Column (4) coefficient on the pooled treatment effect from Table 1 as a targeted moment for indirect inference that disciplines the scale of the subsidy adoption response. The quantitative model produces ZIP-by-quarter EV registration counts on simulated household data calibrated to the income and credit distribution in the 2014-2025 panel, and we run the exact same regression used in Column (4) to replicate this coefficient in the model. We do not target the credit score nad income interaction coefficients in Column (4), but report them as untargeted moments in the model. Table 2 shows the three key empirical moments that we use in the quantitative calibration.

Table 2: Empirical moments used in the quantitative model

Empirical moment	Data	Quantitative parameter
EV sales share, never-treated counties	2.2% $\rightarrow$ 18.1%	Green flow value A_t
Treatment response (Table 1, Column (4))	0.0675 (0.0074)	Preference scale σ
Total subsidies granted	27,779	Program uptake s_e

4.2 Parametrization and calibration

Next, we extend the model in Section 2 to a dynamic structure and add a parsimonious credit and income structure. Households decide whether to switch from a brown to a green vehicle when a replacement opportunity arrives, and they can finance the fixed switching cost with a loan subject

to a downpayment requirement. We calibrate this model to the subsidy treatment effect observed in the data as in Table 2, and study transition dynamics and optimal policy in this environment.

4.2.1 Timing, household parameters, and states

The model is calibrated at a quarterly frequency. The household state variables are their asset holdings, income, credit type, current brown/green technology, and downpayment requirement: (a, ℓ, k, χ, d), where $\chi \in \{brown,\ green\}$. All model units are as a share of average quarterly household income.

Households cannot use unsecured debt, so their usual borrowing constraint is $\underline{a} = 0$. However, households can borrow against the value of their green vehicle by taking a loan at the time of purchase, as described in the following section. The initial distribution of household assets is calibrated to the 2013 Survey of Consumer Finances. Period utility is logarithmic.

Vehicle purchase, or switching, opportunities arrive at rate λ, which is calibrated to annual household vehicle sales in California. Because there is a secular rise in green vehicle purchases during this time period, we also add a flow value of the green technology A_t which captures the *non-policy* change in green adoption (e.g. preferences such as in Fried et al. (2025), energy prices such as in Bushnell et al. (Forthcoming)). We assume a linear functional form $A_t = c_0 + c_1(t/T)$ and fit (c_0, c_1) to the EV sales share in never-treated counties.

4.2.2 ZIP-specific income and credit score distributions

Income follows an AR(1) process and is discretized on the grid $\ell_l \in [\underline{\ell}, \bar{\ell}]$ where $l = 1, \ldots, L$ are individuals in the population of size L. The stationary income weights are π_l^ℓ, where $\sum_l \pi_l^\ell = 1$. Average income is normalized to one: $\sum_l \pi_l^\ell \ell = 1$.

To apply the empirical results on the interaction between credit scores and the treatment effect to the model, we assign individuals to three credit types $k \in \{1, 2, 3\}$, each of equal mass $\pi_k = \frac{1}{3}$. These map to the credit terciles in the data. ZIP codes are calibrated to match their empirical average income and credit scores as used in Section 4.1. The baseline distribution between incomes and credit types is drawn from a Gaussian copula:

$$(Z_1, Z_2) \sim \mathcal{N}\left(\begin{bmatrix}0\\0\end{bmatrix}, \begin{bmatrix}1 & \rho\\ \rho & 1\end{bmatrix}\right), \qquad U = \Phi(Z_1), \quad V = \Phi(Z_2).$$

The correlation between income and credit scores is set to 0.29 following Beer et al. (2018). Define $U = \Phi(Z_1)$, $V = \Phi(Z_2)$. Income and credit draws are then mapped into discrete income and credit type bins based on the original weights for each distribution, which gives the discrete joint

distribution :

$$f_{lk}^0 = P(U \in I_i, V \in C_k)$$
$$where \;\; I_i = \left(\sum_{j<i} \pi_j^\ell, \sum_{j \leq i} \pi_j^\ell \right], \; C_k = \left(\sum_{m<k} \pi_m^c, \sum_{m \leq k} \pi_m^c \right]$$

Next, each ZIP code z is exponentially tilted from the baseline distribution to match the standardized mean income and standardized mean credit score from the data. Let $\tilde{y}_z^{data}$ and $\tilde{c}_z^{data}$ be standardized average income and credit scores from the data, respectively. Similarly, let $\tilde{y}_i$ and $\tilde{c}_k$ be standardized log income and credit scores in the the model. The ZIP-specific joint distribution is:

$$f_z(i,k) = \frac{f_0(i,k) \exp(\lambda_{1z}\tilde{y}_i + \lambda_{2z}\tilde{c}_k)}{\sum_{i',k'} f_0(i',k') \exp(\lambda_{1z}\tilde{y}_{i'} + \lambda_{2z}\tilde{c}_{k'})}$$

Where $(\lambda_{1z}, \lambda_{2z})$ are chosen so that the average income and credit scores in the model match the data:

$$\sum_{i,k} f_z(i,k)\, \tilde{y}_i = \tilde{y}_z^{data},$$
$$\sum_{i,k} f_z(i,k)\, \tilde{c}_k = \tilde{c}_z^{data}$$

4.2.3 Vehicle financing

In Section 2, households must pay the full switching cost out of cash-on-hand subject to the borrowing constraint. In the data and accordingly in the quantitative model, households can finance their vehicle purchase with a loan secured against the value of the vehicle. We calibrate the minimum downpayment share that households must pay at time of purchase to posted car loan downpayment schedules (e.g. Sidney Federal Credit Union (2026), Westlake Financial Services (2025), Santander Consumer USA (2025)). Car loan downpayment requirements are generally segmented by credit score, so we assign households to different minimum downpayment requirements based on their credit scores.

The model's credit types are terciles of the credit score distribution averaged at the ZIP code level. In order to get downpayment schedules at the household level, we assume that a household's score is normally distributed around the ZIP code mean, with a standard deviation of 102 as computed from the FICO score distribution in Black (2024). We base the downpayment schedule specifically on Sidney Federal Credit Union (2026), which explicitly states financing terms for credit score groups. Our three financing types j have downpayment requirements described in Table 3.

That is, the highest credit score group has no downpayment requirement; the second highest has a requirement of 10%, and the third highest has a requirement of 20% of the vehicle price ϕ. Conditional on credit type, financing type is fixed for a household and drawn independently from

Table 3: Financing-type weights π_{jk} by credit tercile

	Cash requirement d_j			
Credit tercile k	0	$0.10\,\phi$	$0.20\,\phi$	Mean cash share
Lowest	0.65	0.10	0.25	0.060
Middle	0.78	0.08	0.15	0.037
Highest	0.86	0.06	0.09	0.023

Notes: Each row gives the probability that a household in credit tercile k faces cash requirement d_j, computed from the lender schedule under a normal within-ZIP score distribution with standard deviation 102 centered on the tercile's household-weighted mean score (670, 708, and 739). Rows sum to one. The last column is the mean of d_j/ϕ across types.

all other variables. Thus, π_{jk} is the probability of falling into financing group j conditional on being in credit group k. For a given financing group j, a household that switches at date t with subsidy s_t can borrow up to

$$B_t(j) = \min\left\{\max(\phi - d_j, 0), \max(\phi - s_t, 0)\right\}$$

The cash that the household must contribute at purchase is $\max(d_j - s_t, 0)$, so the subsidy *relaxes* the downpayment requirement. Switching to the green technology is feasible only if the household has enough cash-on-hand given the financing and minimum consumption constraints:

$$x + s_t - \phi \geq \underline{a} - B_t(j) + (1 + \tau_t^c)\underline{c}(\ell)$$

After financing their vehicle, the additional borrowing amount $B_t(j)$ is subject to the same interest rate r as the standard household asset a.

4.2.4 The policy schedule

Within each ZIP code, households are allocated to their federal poverty level award tiers to match the share of households in each tier based in the ACS. The calibration simulates the full program window used in the empirics, 2014-2025. Each household's value function includes their available subsidy, which is their statutory award path based on their FPL tier, DAC status, and date. Households learn about the subsidy when it is rolled out to their district, but they do not anticipate future changes in the policy. The predicted exposure T_{zt}^{pred} from Section 4.1 is used in the treatment regression to best mimic the empirical result.

4.2.5 Calibration objective

The calibration searches over four structural parameters: The parameters governing the value of the green technology, the T1EV scaling parameter, and the subsidy take-up level:

$$\theta = (c_0, c_1, \sigma, s_e).$$

The switching cost ϕ, the discount factor $\bar{\beta}$, the borrowing limit $\underline{a}$, the financing schedules (π_{jk}, d_j), the minimum consumption schedule $\underline{c}(\ell)$, the income-credit correlation parameter ρ, the annual turnover rate, and the initial green share are all fixed outside the search.

We calibrate these parameters sequentially. First, we calibrate the green flow value $A_t = c_0 + c_1(t/T)$ to match EV adoption in never-treated counties. This stage can be solved independently as it does not depend on the subsidy policy. Second, we jointly calibrate the subsidy take-up level s_e to match the flow of subsidies in the data and the T1EV preference parameter σ to the Table 1, Column (4) coefficient on the main treatment effect of the subsidy. These two parameters are solved iteratively to a joint fixed point, as subsidy take-up is affected by the preference parameter. We do *not* target the interaction coefficients β_{credit} and β_{income}, but we are able to largely match them in the calibrated model as untargeted moments.

4.2.6 Matching empirical treatment effects

We follow the empirical strategy in Section 4.1 to match the treatment effect and subsidy take-up in the model. For each income-credit group (ℓ, k) in ZIP code z, let $S_{z\ell kt}(\theta)$ be the share of households that switch in period t, and $P^g_{z\ell kt}(\theta)$ be the share of already-green adopters at the start of time t. The total number of EV purchases in ZIP z, quarter t, within each (ℓ, k) group is the sum of new switchers and already-green adopters who replace their vehicles at the replacement arrival rate λ:

$$Q^{EV}_{z\ell kt}(\theta) = S_{z\ell kt}(\theta) + \lambda P^g_{z\ell kt}(\theta)$$

Summing across groups within a ZIP code, the share of switchers at time t and ZIP z is:

$$q^{EV}_{zt}(\theta) = \sum_{\ell,k} f_z(\ell, k) Q^{EV}_{z\ell kt}(\theta)$$

This aggregation maps the household decision problem into a ZIP-level sales share comparable to the data. We then turn these rates into switching quantities transformed by an inverse hyperbolic sine to match the y-variable used in the empirical section. To do so, we first apply a Poisson count draw to the switching rates in order to transform them into switching counts, then transform using an inverse hyperbolic sine:

$$Y^{model}_{zt} = \mathbb{E}[asinh(N_{zt})] \quad where \quad N_{zt} \sim Poisson(n_z q^{EV}_{zt}(\theta))$$

Where n_z is the number of households in ZIP code z. This transformation is particularly important in small zip codes where noise may be a larger concern.

We then feed in the empirical treatment variable T^{pred}_{zt} and run an analogous two-way fixed effect regression to Equation (28):

$$Y^{model}_{zt} = \beta_T T^{pred}_{zt} + \beta_{credit}\big(T^{pred}_{zt} \times c_z\big) + \beta_{income}\big(T^{pred}_{zt} \times inc_z\big) + \nu_z + \gamma_t + \epsilon_{zt}, \tag{34}$$

where c_z denotes standardized credit score, inc_z denotes standardized income, ν_z are ZIP-code fixed effects, and γ_t are time fixed effects. We jointly search over $\theta = (\sigma, s_e)$ to match the administrative award count, which determines s_e, and the pooled response β_T, which determines σ. The model-implied interactions β_{credit} and β_{income} are *untargeted*, but the fit for these interaction treatment effects is reasonably close in the calibrated model.

4.2.7 Calibrated parameters and moments

We calibrate the model using standard parameters from the literature, as well as targeting the treatment effect coefficient and EV diffusion trends described above. Table 4 shows our calibrated and target parameters. The model is initialized at an EV share of 0.004, which matches the share of EVs in California at the start of the sample window in 2014Q1. The remaining calibrated parameters are fairly standard at the quarterly frequency. We calibrate the consumption floor $\underline{c}(\ell)$ to the share of income in committed consumption in the CEX following Chetty and Szeidl (2016). We compute this consumption share within vehicle owners in the 2015-2019 CEX. We estimate this share within deciles of the income distribution and interpolate so the consumption floor is a function of income.

Externality damages from using the brown technology are linear in the share of households that own brown vehicles $1 - P_t$, or $D(P_t) = \delta(1 - P_t)$. We use the social cost of carbon to calibrate δ, or the quarterly externality imposed by one brown vehicle in utility units. A gasoline vehicle emits about 4.6 tons of CO_2 per year (U.S. Environmental Protection Agency (2023b)), or 1.15 tons per quarter. We take this estimate and multiply it by the social cost of carbon and scale by average marginal utility, $\bar{u}'(c_0) = \mathbb{E}\left[\frac{1}{c_0}\right]$:

$$\delta = \frac{SCC \times 1.15}{\bar{c}_0}$$

In our main specification, we use $\delta = 0.01$, which is implied by the EPA's central estimate of \$190 per ton of carbon in 2020 dollars (U.S. Environmental Protection Agency (2023a)). In robustness tests, we use values of the social cost of carbon that range up to \$1,200 as in Bilal and Känzig (2026) (BK SCC). We deflate the social cost of carbon to 2012 dollars using the CPI.

Table 5 shows the fit between the model moments, the empirical targets, and untargeted empirical moments. Panel A shows the targeted calibration fit. The linear green flow value A_t is calibrated to hit the quarter-by-quarter EV takeup in never-treated counties. The model-equivalent moment is the adoption path of never-treated ZIP codes. The correlation between the data and model for the never-treated switching share is 0.97, and the model approximately matches the empirical increase in EV takeup from the start to the end of the sample. The average sales share in never-treated counties is matched almost exactly by the model.

The preference parameter σ is targeted to hit the scale of the subsidy response from the main treatment coefficient β_T in Table 1, Column (4). The model moment from running the exact same regression on simulated data is 0.0684, which is very close to the measured 0.0675 coefficient from

Table 4: Calibration

Parameter	Value	Description
β	0.985	Discount factor (quarterly)
r	0.005	Interest rate (quarterly)
γ	1.0	CRRA coefficient
$\underline{c}(\ell)$	0.61–0.28	Subsistence floor by income decile, CEX
$\underline{a}$	0	Borrowing constraint
ρ_ℓ	0.9873	Income: AR(1) persistence (annual 0.95)
σ_ε	0.1019	Income: innovation st.d. (annual 0.20)
G	0.18	Gov't spending (share of GDP)
ϕ	1.468	Gross EV price (CC4A median)
ω	0.128	Annual vehicle turnover, CEC LDV sales per household
d_j	$(0, 0.1, 0.2)\phi$	Cash requirement by type (lender LTV 100%, 90%, 80%)
$\Pr(d_j > 0 \mid k)$	$(0.35,\ 0.22,\ 0.14)$	Share facing a cash floor, by credit tercile (low to high)
ρ	0.29	Income-credit copula correlation (Beer et al. (2018))
P_0	0.004	Initial EV stock share (2014Q1)
δ	0.01	Brown externality cost per unit, using SCC $\approx$ \$190
Calibrated parameters		
(c_0, c_1)	$(-0.077,\ 0.069)$	Green flow value A_t
σ	0.627	Logit preference scale
s_e	0.053	Subsidy uptake

the data and well within one standard error. Finally, the subsidy uptake in the model matches exactly the number in the data.

Panel B compares untargeted moments from the data and the model. The model naturally reproduces the signs of the credit and income interactions in the data: The credit interaction is negative, while the income interaction is positive. The model understates the magnitude of the credit score gradient but is much closer to the magnitude of the income gradient. The qualitative fit of these untargeted moments indicates that the income, credit, and financing structure are able to generate the motivating fact on lower credit scores being correlated with greater subsidy treatment effects, which reflects the key insight of the theory that loosening credit constraints can improve financing ability.

5 Optimal policy and welfare

In this section, we solve the planner's problem from Section 3 in the calibrated economy. Starting from the initial period in 2014Q1, the planner chooses the path of policy instruments that persist through the end of the sample in 2025Q2. The government budget constraint is closed in each period through the consumption tax τ^c. In our baseline specification, the planner uses a utilitarian social welfare function. We compare three regimes that differ in the policy instruments available to the planner: (1) A brown tax only $(\tau, 0)$; (2) A green subsidy only $(0, s)$, and (3) The brown tax

Table 5: Targeted and untargeted moments

Moment	Data	Model
Panel A. Targeted moments		
Never-treated EV sales share:		
Average	0.093	0.092
Start → end	0.022 → 0.181	0.014 → 0.218
Data and model correlation	0.97	–
Table 1, Column (4) coefficient β_T	0.0675 (0.0074)	0.0684
Subsidies granted (thousands)	27.8	27.8
Panel B. Untargeted moments (Table 1, Column (4))		
Credit interaction β_{credit}	−0.0864 (0.0079)	−0.0459
Income interaction β_{income}	0.1032 (0.0087)	0.0879

jointly with the green subsidy (τ, s). In the baseline specification, the policy rate is chosen before the sample starts and remains constant throughout the sample.

We conduct two policy exercises in this economy. First, we mandate that the planner targets a fixed terminal adoption share, and we compare the welfare cost of achieving this adoption share under each policy regime. This policy exercise is in line with explicitly stated policy goals of reaching a certain adoption of EVs through the California EV rebates (e.g. Cart (2017)) as well as similar national policy such as the Inflation Reduction Act (e.g. Minott and Nguyen (2023)). In this specification, we choose the terminal adoption share to be 1 percentage point greater than last-period adoption in the model with no policy, and we show robustness to different adoption targets.

Second, we allow the planner to choose the policy parameters with the goal of maximizing welfare without an adoption target. We apply a variety of social welfare weights to study how optimal policy varies based on the government's welfare priorities. For simplicity when mkaing welfare calculations, we assume that the amount of green adoption in the last period continues beyond the policy window, which allows us to fully compute discounted social welfare without making modeling adoption decisions after the sample period. We find that under both the fixed adoption target and no adoption target regimes, the subsidy is able to reduce the number of households that cannot finance the green purchase. This leads to higher welfare under regimes that include the subsidy.

5.1 Welfare with an adoption target

The first exercise compares the welfare cost of meeting a prespecified green adoption share by the end of the sample period. We solve the planner's problem conditional on a target share of adoption by allowing the planner to impose a constant policy regime at the start of the sample. The terminal adoption share is calibrated to the no-policy terminal adoption share plus 1%, which

is a total adoption share of 10.58%. California's green vehicle policy takes exactly the target adoption approach; recent vehicle regulations and policies are written to target a certain share of zero emission vehicles (e.g. California Air Resources Board (2022a)). We compare three regimes: A brown tax only $(\tau, 0)$; a green subsidy only $(0, s)$, and a brown tax with a green subsidy (τ, s). The government closes its budget constraint with a consumption tax.

For the tax- and subsidy-only policies, we find the constant instrument that induces the adoption target by using a bracketed scalar root, as the outcome variable is one dimensional. For the joint policy, we search for the optimal subsidy given different starting values for taxes on a five-point grid, then run refinements around the best tax point to select the optimal tax instrument conditional on the selected subsidy. We solve for the joint policy after the single instrument policies and ensure that the joint policy can do at least as well as the single instrument policy.

Table 6 reports the constant policy rates and welfare results for the fixed adoption comparison. All policies result in net welfare costs; the reason is because the 1% policy adoption target is slightly higher than the welfare-optimizing adoption rate given the calibrated social cost of carbon. However, the subsidy-only policy leads to a 3x smaller welfare cost compared to the tax-only policy. The differences across regimes largely comes from changes in household lifetime values instead of avoided environmental damages, as the last-period target adoption is the same under both regimes which leads to largely similar externality costs. Under the subsidy-only policy, about 4% more households are able to finance switching compared to the tax-only policy. The tax and subsidy regime obtains a slight welfare increase compared to the single instrument policies by setting a low tax rate and a high subsidy to increase the number of people who can afford to switch while funding the subsidy partially with the brown tax.

Table 6: Welfare across policy regimes with an adoption target

Regime	Welfare cost of target, CE (%)	Welfare change vs. no policy			Terminal adoption	Constant policy		Unable to finance
		Total	Private	Damages		τ	s	
No policy	–	–	–	–	0.0958	0	0	0.0694
Tax only	0.0241	-0.01169	-0.01430	+0.00261	0.1058	0.00197	0.00000	0.0645
Subsidy only	0.0088	-0.00426	-0.00675	+0.00249	0.1058	0.00000	0.13133	0.0263
Tax + subsidy	0.0063	-0.00307	-0.00559	+0.00252	0.1058	0.00049	0.08142	0.0381

Notes: The social welfare function is utilitarian. The welfare cost is the percent reduction in each household's consumption over the full period, relative to the no-policy economy. The welfare change from private values indicates changes in households' lifetime value functions excluding the externality, and the welfare change from damages stem from the externality.

Figure 4 further shows how the financing relief offered by the subsidy contributes to the welfare advantage of the subsidy over the tax. We split each policy's welfare gain into an intensive margin that includes the set of households who can always afford to switch, as well as an extensive margin composed of households that are able to afford switching due to the liquidity benefits of the subsidy. The extensive margin channel accounts for just under half of the increase in welfare of the subsidy over the tax, while the intensive margin cost accounts for the remaining difference. The tax plus subsidy regime is able to maintain a large positive effect of the extensive margin and a relatively

smaller intensive margin cost, leading to a smaller overall welfare cost.

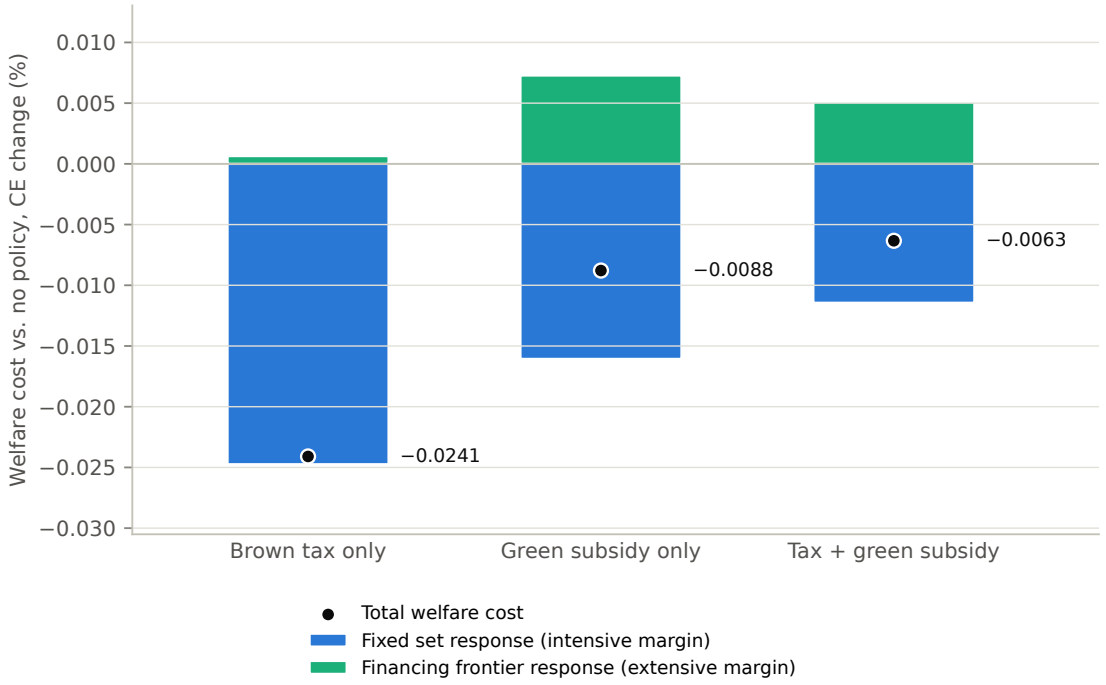


Figure 4: Welfare decomposition across policy regimes

Much of the welfare advantage of the subsidy comes from increasing the number of households that are able to finance a switch. Figure 5 plots the share of brown households that cannot afford switching across the subsidy horizon. In general, the tax-only and no policy regimes have about double the number of households that cannot afford to switch compared to the subsidy-only and tax plus subsidy regimes. This *extensive margin* response from the subsidy contributes substantially to the welfare benefit of the subsidy, as households with lower incomes that could not otherwise afford switching are able to change their vehicle type while avoiding the brown tax after switching.

Finally, we test the sensitivity of these welfare results to a collection of different parameters. In Panel A, we vary the slope of climate damages δ from 0.005, which is smaller than our baseline value, to 0.06, which is motivated by the social cost of carbon in Bilal and Känzig (2026). This experiment holds fixed the 1 percent target in increased adoption, which means that the target taxes and subsidies remain the same to reach this adoption rate. The welfare cost of hitting 1 percent higher adoption decreases as the externality increases, with the subsidy-only and tax plus subsidy policies reaching *positive* welfare results starting from $\delta = 0.05$. Under the externality calibrated to the social cost of carbon in Bilal and Känzig (2026), all policies lead to positive welfare gains. Under each value of the externality cost, the subsidy results in higher welfare than the tax, while the tax plus subsidy policy leads to the highest welfare.

In Panel B, we vary the welfare weights on the social welfare function. The planner's social

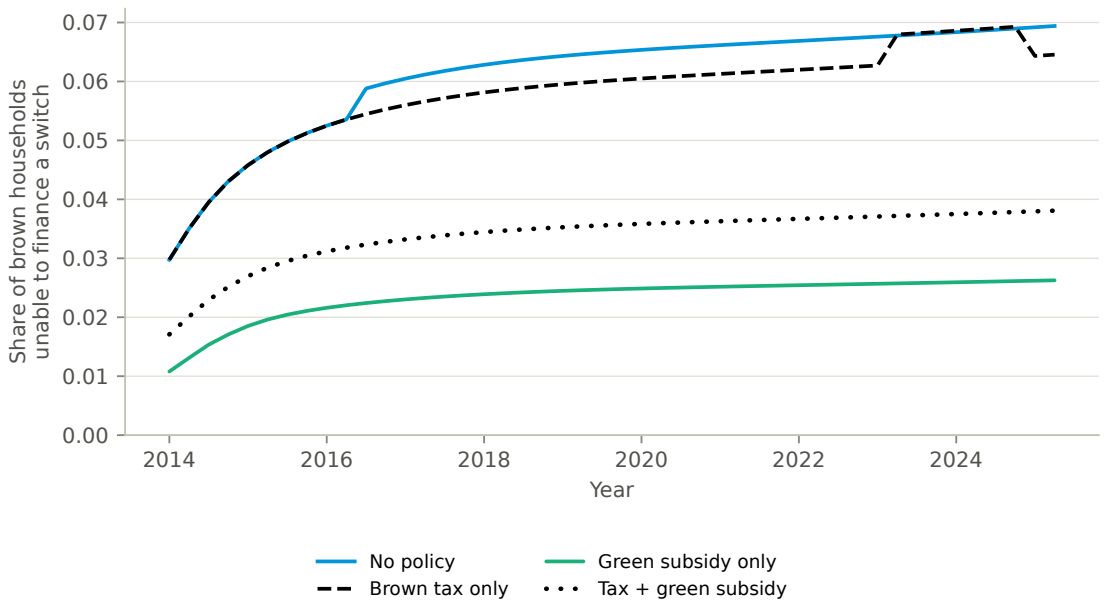


Figure 5: Share of brown households unable to finance a switch

welfare function with weights w_ℓ on income groups is:

$$SWF(\tau, s) = \sum_\ell \pi_\ell w_\ell V_\ell(\tau, s) - D(\tau, s), \quad \sum_\ell \pi_\ell w_\ell = 1$$

In the utilitarian benchmark, $w_\ell = 1$ for all ℓ. Our sensitivity tests add three additional welfare weights: inverse income weights ($\omega_\ell \propto \ell^{-1}$), square root income weights ($\omega_\ell \propto \sqrt{\ell}$), and Negishi weights $\omega_\ell \propto \ell_0$. These weights depend on time 0 income. Inverse income weights emphasize the distributional incidence of policy, leading to lower welfare costs for policies that include the subsidy and a higher welfare cost for the tax-only policy. Square root income moves in the other direction and gives higher weights to higher income households. This reduces the welfare cost of policies that include the tax, but increases the welfare cost of the subsidy. Finally, Negishi weights, which are also used in popular environmental frameworks like the Kyoto Protocol, give households weights equal to their time-0 income. This substantially reduces the welfare cost of the tax but increases the cost of the subsidy. In general, these different social welfare weights emphasize the regressivity of the tax and the progressivity of the subsidy.

In Panel C, we vary the government's adoption target. We start with the no-policy plus 1 percent target as in the baseline, and increase this target up to 4 percent. Increasing the target monotonically increases the welfare cost across all policies as it pushes adoption away from its optimal value, but the joint tax plus subsidy policy experiences a smaller increase in the welfare cost than the single instrument policies as this regime is able to incorporate the intensive margin effects from the tax as well as the extensive margin effects from the subsidy.

Table 7: Welfare and policy sensitivity with an adoption target

	Tax only		Subsidy only		Tax + subsidy		
Scenario	CE (%)	τ	CE (%)	s	CE (%)	τ	s
A. Externality cost ($\Delta P_T = 1$ pp; utilitarian)							
$\delta = 0.005$	-0.0268	0.00197	-0.0114	0.13133	-0.0089	0.00049	0.08142
$\delta = 0.01$	-0.0241	=	-0.0088	=	-0.0063	=	0.08142
$\delta = 0.02$	-0.0187	=	-0.0037	=	-0.0011	=	0.08142
$\delta = 0.05$	-0.0026	=	+0.0117	=	+0.0144	=	0.08142
BK SCC ($\delta \approx 0.06449$)	+0.0052	=	+0.0192	=	+0.0219	=	0.08142
B. Welfare weights ($\delta = 0.01$; $\Delta P_T = 1$ pp)							
Utilitarian	-0.0241	0.00197	-0.0088	0.13133	-0.0063	0.00049	0.08142
$w_\ell \propto \ell_0^{-1}$	-0.0465	=	-0.0058	=	-0.0058	0.00000	0.13133
$w_\ell \propto \sqrt{\ell_0}$	-0.0141	=	-0.0102	=	-0.0048	0.00049	0.08134
$w_\ell \propto \ell_0$ (Negishi)	-0.0056	=	-0.0114	=	-0.0036	0.00050	0.08108
C. Adoption increment ($\delta = 0.01$; utilitarian)							
$\Delta P_T = 1$ pp	-0.0241	0.00197	-0.0088	0.13133	-0.0063	0.00049	0.08142
$\Delta P_T = 2$ pp	-0.0470	0.00378	-0.0270	0.26483	-0.0270	0.00000	0.26483
$\Delta P_T = 3$ pp	-0.0688	0.00548	-0.0782	0.45576	-0.0481	0.00137	0.27933
$\Delta P_T = 4$ pp	-0.0897	0.00707	-0.1585	0.68502	-0.0646	0.00412	0.17087

*N*otes: = denotes the same policy parameter as in the panel's first row. All rows report the same log utility consumption-equivalent percentage, and each change is relative to welfare in the no-policy economy.

5.2 Welfare maximization without an adoption target

Next, we solve the full planner's problem as in Section 3 by allowing the planner to optimize welfare freely without being bound to an adoption target. As before, we have tax-only, subsidy-only, and tax plus subsidy regimes. Our baseline calibration includes utilitarian welfare weights and externality damages of $\delta = 0.01$. For a given policy vector (τ, s) we solve for household decisions and the budget-balancing consumption tax. Then, we search over different values for (τ, s) using a bounded quasi-Newton method to converge at the policy with the highest welfare.

Table 8 reports the welfare results from each optimized policy regime. The optimal tax-only candidate is no taxation at all with no change in welfare, as increasing the tax leads to reduced welfare from the private fiscal incidence of the brown tax. The optimal subsidy-only policy is setting the subsidy at 0.0075, which is a lower rate than the adoption target subsidy of 0.13. The joint policy maximization lands at the corner solution with subsidies only and achieves the same welfare gain, 0.00109, as in the subsidy-only policy. Although this welfare maximization problem landed at the no-tax corner solution, sensitivity results in Table 9 show that increasing the externality damages can lead to positive taxes in the tax-only and joint regimes.

As with the constant adoption regimes, we also study the sensitivity of the welfare ordering and policy structure under different damages δ and welfare weights ω. Increasing damages from 0.005 to 0.065 substantially increases the welfare benefits of policy, which is able to mitigate a larger externality. Additionally, $\delta \geq 0.05$ leads to a positive value of the tax in the tax only and

Table 8: Welfare-maximizing policy regimes

Regime	Welfare change, CE (%)	Welfare change vs. no policy: Total	Private	Damages	Terminal adoption	Constant policy: τ	s	Unable to finance
No policy	–	–	–	–	0.0958	0	0	0.0694
Tax only	+0.0000	+0.00000	+0.00000	+0.00000	0.0958	0.00000	0.00000	0.0694
Subsidy only	+0.0022	+0.00109	+0.00077	+0.00032	0.0971	0.00000	0.00750	0.0548
Tax + subsidy	+0.0022	+0.00109	+0.00077	+0.00032	0.0971	0.00000	0.00750	0.0548

Notes: The welfare change is relative to the no policy economy.

tax plus subsidy regimes. In particular, the tax plus subsidy regimes for large δ shifts from being subsidy-only to tax-only, as the tax becomes high enough to induce most of the socially optimal switching given the large externality.

Table 9: Sensitivity of welfare-maximizing constant policies without an adoption target

Scenario	Tax only: CE (%)	τ	P_T (%)	Subsidy only: CE (%)	s	P_T (%)	Tax + subsidy: CE (%)	τ	s	P_T (%)
A. Externality cost (utilitarian)										
$\delta = 0.005$	+0.0000	0.00000	9.576	+0.0019	0.00606	9.693	+0.0019	0.00000	0.00679	9.700
$\delta = 0.01$	+0.0000	0.00000	9.576	+0.0022	0.00750	9.707	+0.0022	0.00000	0.00750	9.707
$\delta = 0.02$	+0.0000	0.00000	9.576	+0.0039	0.08060	10.311	+0.0039	0.00000	0.08060	10.311
$\delta = 0.05$	+0.0895	0.03415	38.673	+0.0216	0.30278	12.020	+0.0895	0.03415	0.00000	38.673
BK SCC ($\delta \approx 0.06449$)	+0.4225	0.04838	51.973	+0.0400	0.30380	12.031	+0.4225	0.04838	0.00000	51.973
B. Welfare weights ($\delta = 0.01$)										
Utilitarian	+0.0000	0.00000	9.576	+0.0022	0.00750	9.707	+0.0022	0.00000	0.00750	9.707
$w_i \propto y_{i0}^{-1}$	+0.0000	0.00000	9.576	+0.0029	0.00856	9.717	+0.0029	0.00000	0.00856	9.717
$w_i \propto \sqrt{y_{i0}}$	+0.0000	0.00000	9.576	+0.0020	0.00750	9.707	+0.0020	0.00000	0.00750	9.707
$w_i \propto y_{i0}$ (Negishi)	+0.0000	0.00000	9.576	+0.0017	0.00728	9.705	+0.0017	0.00000	0.00728	9.705

Notes: P_T under no policy is 9.5760%. The welfare change is relative to the no policy economy.

Figure 6 shows the amount of green adoption in each of the policy regimes. On the left panel, the adoption target regimes all hit the same adoption share of 0.1058 by the end of the sample period, compared to 0.0958 in the no policy economy. In the baseline calibration with $\delta = 0.01$, the welfare maximizing policies target only a slightly higher amount of adoption than the no policy case, and see small welfare gains that mostly stem from relaxing the liquidity constraint. However, in the calibration using the much higher Bilal and Känzig (2026) social cost of carbon, green adoption is much higher in all of the policy regimes. The brown tax is able to encourage more adoption than the subsidy and earn substantially higher welfare benefits as a result.

In general, comparing adoption and welfare across different externality costs reveals a pattern: At lower externality costs (e.g. $\delta < 0.05$), the subsidy is more effective at increasing welfare by relaxing the borrowing constraint and encouraging a small amount of adoption. However, at larger externality costs, the tax becomes a larger priority for the planner to encouraging large-scale adoption within the intensive margin of households.

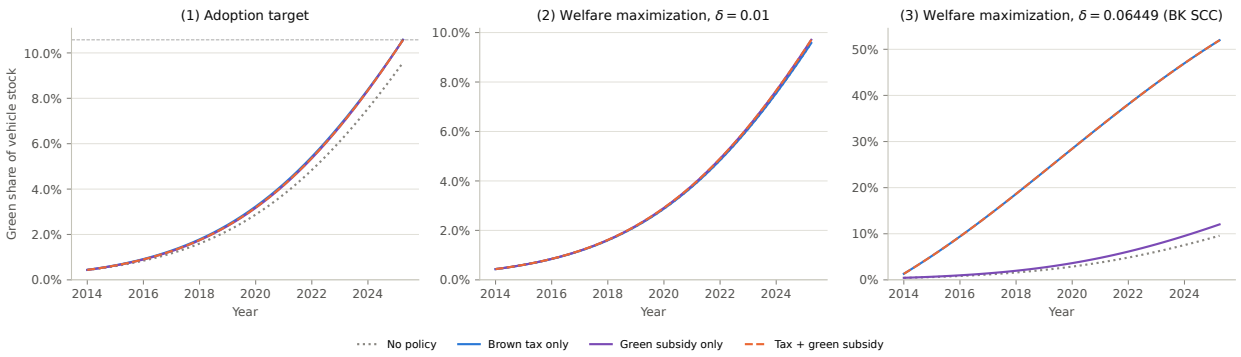


Figure 6: Green adoption across policy regimes

6 Conclusion

This paper revisits optimal Pigouvian policy under individual heterogeneity in income and financing ability. When adopting a green good requires paying an upfront fixed costs, borrowing constraints may prevent households from financing adoption. In this environment, the usual intuition that optimal policy constitutes a Pigouvian tax on the brown good may not hold. While a tax on the brown technology induces switching through the intensive margin, a subsidy on the green technology operates through an additional extensive margin by relaxing households' feasibility constraints. We derive theoretical results that identify conditions under which this extensive margin can lead to optimally adopting a welfare-improving green subsidy, even when the brown good can be taxed directly.

Consistent with this mechanism, we study California's electric vehicle incentive program and find that ZIP codes with lower average credit scores are *more* responsive to an EV subsidy, suggesting an empirically relevant extensive margin. We then calibrate a quantitative model to study adoption and welfare under policy combinations of taxes and subsidies. In our baseline calibration which matches the empirical treatment effect and credit score gradient of the subsidy, we find that subsidies reduce the welfare cost of achieving a green adoption target by relaxing the liquidity constraint. Additionally, a welfare-maximizing planner can increase total welfare by prioritizing subsidies to relax households' borrowing constraints. Our results imply that Pigouvian policy should take household financing conditions seriously when planning for a green transition, and future work on alternative policy instruments in these environments may be fruitful.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous, "The Environment and Directed Technical Change," *Am. Econ. Rev.*, February 2012, *102* (1), 131–166. 4

Aghion, Philippe, Antoine Dechezleprêtre, David Hémous, Ralf Martin, and John Van Reenen, "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry," *J. Polit. Econ.*, February 2016, *124* (1), 1–51. 4

Allcott, Hunt, Christopher Knittel, and Dmitry Taubinsky, "Tagging and targeting of energy efficiency subsidies," *Am. Econ. Rev.*, May 2015, *105* (5), 187–191. 4

Beer, Rachael, Felicia Ionescu, and Geng Li, "Are Income and Credit Scores Highly Correlated?," *FEDS Notes*, August 2018. 23, 28

Berkouwer, Susanna B and Joshua T Dean, "Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households," *Am. Econ. Rev.*, October 2022, *112* (10), 3291–3330. 4

Bilal, Adrien and Diego R. Känzig, "The Macroeconomic Impact of Climate Change: Global Versus Local Temperature," *The Quarterly Journal of Economics*, 2026, *141* (2), 889–944. 27, 31, 34

Black, Michelle L., "Average FICO Credit Score in America," https://www.self.inc/info/average-credit-score-in-america/ April 2024. Self Financial. Distribution of U.S. consumers across FICO score bands as of April 2024, tabulated from FICO data. Accessed: 2026-09-28. 24

Bricker, Jesse and Geng Li, "Your friends, your credit: Social capital measures derived from social media and the credit market," *Econ. Lett.*, September 2024, *242* (111881), 111881. 17

Bushnell, James, Erich Muehlegger, and David Rapson, "Energey Prices and Electric Vehicle Adoption," *J. Polit. Econ. Microecon.*, Forthcoming. 23

California Air Resources Board, "Advanced Clean Cars II Regulations," https://ww2.arb.ca.gov/our-work/programs/drive-forward-light-duty-vehicle-program/advanced-clean-cars 2022. Adopted August 25, 2022; 100 percent of new passenger vehicles zero-emission by model year 2035. Accessed: 2026-09-28. 30

_ , "Appendix I: Annual Performance Goals and Evaluation for the Clean Cars 4 All and Enhanced Fleet Modernization Programs—Fiscal Years 2021/2022 and 2022/2023," Technical Report, California Air Resources Board 2022. Appendix to the Proposed Fiscal Year 2022–23 Funding Plan for Clean Transportation Incentives. 19

_ , "Outcomes and Results for Clean Cars 4 All," https://ww2.arb.ca.gov/our-work/programs/clean-cars-4-all/outcomes-and-results-clean-cars-4-all 2026. Accessed: 2026-02-10. 17

California Energy Commission, "ZEV and Infrastructure Stats Data," https://www.energy.ca.gov/files/zev-and-infrastructure-stats-data 2026. Accessed: 2026-02-10. 17

Cart, Julie, "California Tries to Goose Electric Car Sales with Rebates—but Is It Enough?," September 2017. Updated June 23, 2020; accessed September 9, 2026. 29

Chetty, Raj and Adam Szeidl, "Consumption commitments and habit formation," *Econometrica*, March 2016, *84* (2), 855–890. 27

Cremer, Helmuth, Firouz Gahvari, and Norbert Ladoux, "Externalities and optimal taxation," *J. Public Econ.*, December 1998, *70* (3), 343–364. 4

Farhi, Emmanuel and Xavier Gabaix, "Optimal taxation with behavioral agents," *Am. Econ. Rev.*, January 2020, *110* (1), 298–336. 4

Fried, Stephie, David Lagakos, and Hannah Rhodenhiser, "Homework in Climate Economics: Household Production, Environmental Preferences, and Climate Policy," *Working Paper*, February 2025. 4, 23

Jaqua, Daniel and Daniel Schaffa, "The case for subsidizing harm: constrained and costly Pigouvian taxation with multiple externalities," *Int. Tax Publ. Fin.*, April 2022, *29* (2), 408–442. 4

Kopczuk, Wojciech, "A note on optimal taxation in the presence of externalities," *Econ. Lett.*, July 2003, *80* (1), 81–86. 2, 3

Kremer, Michael and Jack Willis, "Guns, latrines, and land reform: Dynamic Pigouvian taxation," *Am. Econ. Rev.*, May 2016, *106* (5), 83–88. 4

Minott, Owen and Helen Nguyen, "IRA EV Tax Credits: Requirements for Domestic Manufacturing," February 2023. Accessed September 9, 2026. 29

Muehlegger, Erich and David S Rapson, "Subsidizing low- and middle-income adoption of electric vehicles: Quasi-experimental evidence from California," *J. Public Econ.*, December 2022, *216* (104752), 104752. 2, 4, 16, 17, 18, 19, 39

Sandmo, Agnar, "Optimal taxation in the presence of externalities," *Swed. J. Econ.*, March 1975, *77* (1), 86. 2, 3

Santander Consumer USA, "How much should a car down payment be?," March 2025. Published March 27, 2025. Accessed September 28, 2026. 24

Sidney Federal Credit Union, "Auto Loans: Rates and Maximum Loan-to-Value by Credit Score," https://sfcuonline.org/loans/vehicle/auto-loans-2/ 2026. Posted auto-loan schedule, rates effective September 25, 2026. Maximum loan-to-value of MSRP or retail value including tax, tags, GAP, and warranty by credit score: 740+ 125%, 670–700 120%, 630–669 100%, no score 100%, 600–629 90%, 599 and below 80%. Accessed: 2026-09-25. 24

Smith, Stephen, "Taxation and the Environment: A Survey," *Fisc. Stud.*, 1992, *13* (4), 21–57. 3

U.S. Census Bureau, "American Community Survey 5-Year Estimates: Median Household Income," https://data.census.gov/ 2012. Accessed: 2026-02-10. 17

U.S. Environmental Protection Agency, "EPA Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances," https://www.epa.gov/system/files/documents/2023-12/epa_scghg_2023_report_final.pdf 2023. Docket ID No. EPA-HQ-OAR-2021-0317, November 2023. Accessed: 2026-09-29. 27

_ , "Greenhouse Gas Emissions from a Typical Passenger Vehicle," https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle 2023. EPA-420-F-23-014, June 2023. Accessed: 2026-09-29. 27

Westlake Financial Services, "Independent Dealer Rate Sheet, v02.25," https://blob.westlakepromos.com/media/2025/03/0225-Independent-Rate-Sheet.pdf February 2025. Minimum down payment as a share of price by program: Standard (scores 0–599) 10%, Gold (600–699) 7.5%, Platinum (700–749) 5%, Titanium (750+) 5%; maximum loan-to-value including backend 140–150%. Accessed: 2026-09-25. 24

Appendix

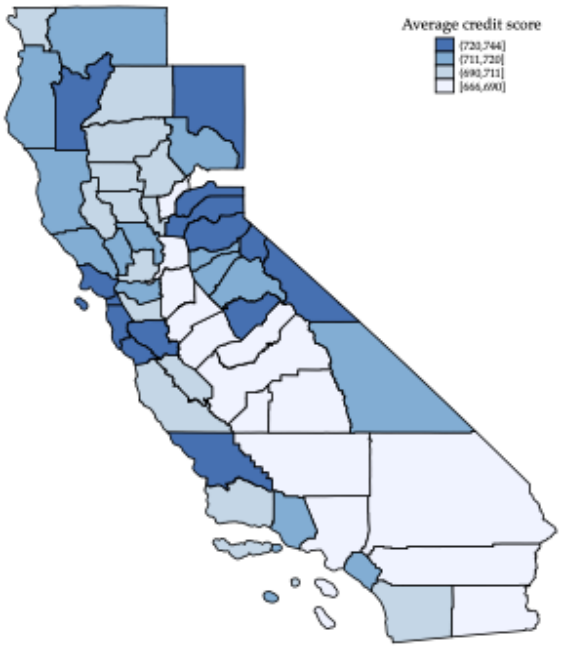
A Additional tables and figures

Table A1: Replication of Muehlegger and Rapson (2022)

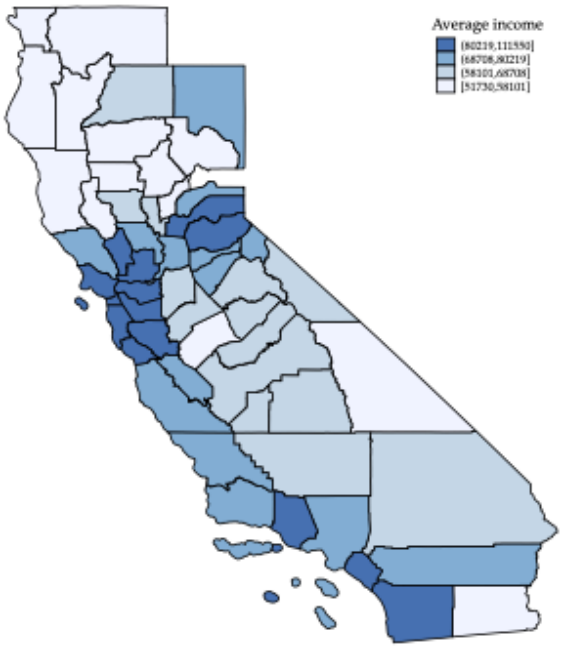
	This paper						MR (2022)	
	2014Q1–2025Q2			2014Q1–2017Q4			Table 4 estimates	
	(1) T^h	(2) T^l (IV)	(3) T^{pred}	(4) T^h	(5) T^l (IV)	(6) T^{pred}	MR (1) T^h	MR (4) T^l (IV)
T^h	0.0367***			0.0311***			0.0280***	
	(0.0028)			(0.0034)			(0.0057)	
T^l		0.1545***			0.0900***			0.1200***
(instrumented)		(0.0211)			(0.0128)			(0.0110)
T^{pred}			0.0648***			0.0338***		
			(0.0062)			(0.0035)		
ZIP FE	Yes	Yes	Yes	Yes	Yes	YES	Yes	Yes
Time FE	Quarter	Quarter	Quarter	Quarter	Quarter	Quarter	Quarter	Quarter
Observations	7.80e+04	7.80e+04	7.80e+04	2.71e+04	2.71e+04	2.71e+04	15,801	15,621
Adj. R^2	0.9231	–	0.9236	0.9238	–	0.924	0.8700	–
First stage F-stat	–	175.2439	–	–	314.4410	–	–	170.2

Standard errors clustered by ZIP; credit scores and incomes are standardized. All specifications include ZIP and time fixed effects. Columns (1)-(2) and (4)-(5) use the same regression specifications as in Muehlegger and Rapson (2022), shown in the last two columns (MR). Columns (3) and (6) use the OLS regression and the predicted treatment variable, which estimates a treatment effect in between T^h and T^l. We include specifications using our full sample of 2014-2025 in Columns (1)-(3), while Columns (4)-(6) use 2014Q1-2017Q4 only to better match Muehlegger and Rapson (2022) and demonstrate that both sample lengths generate similar treatment effects.

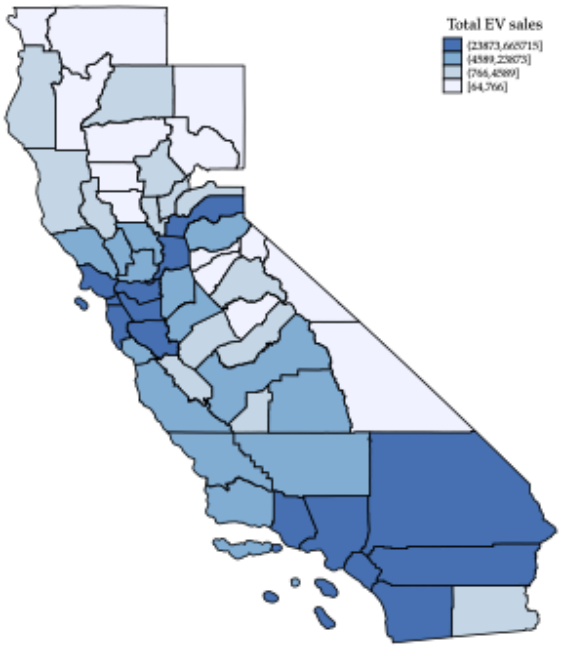
Figure A1: Geographic distribution of main variables in 2018



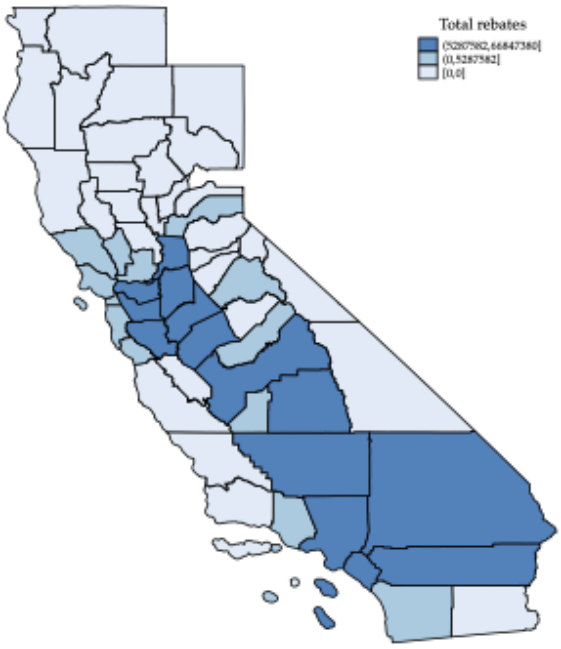
(a) Map: Credit scores (avg)



(b) Map: Income (avg)



(c) Map: EV sales (total)



(d) Map: Rebates (total)

Table A2: Difference-in-Differences Regressions: Effect of Subsidies on EV Sales (2015-2018)

	(1)	(2)	(3)	(4)	(5)
T^h	.02023***		.02658***	.03012***	.02826***
	(.003035)		(.003828)	(.003382)	(.003789)
T^l		.1094***			
		(.01086)			
Credit score $\times$ T^h			.006808*		-.006108
			(.003777)		(.006486)
Income $\times$ T^h				.01467***	.02038***
				(.004293)	(.007482)
Time $\times$ ZIP FEs	X	X	X	X	X
N	26668	26668	26668	26668	26668
Adjusted R^2	.8959	-.1341	.896	.896	.896
First-stage F-stat		252.903			
Hansen Test p-value		—			

Notes: Standard errors clustered at the ZIP level in parentheses. Credit score and income are standardized. T^h uses the "high" eligibility weight (share of households below 400% FPL). T^l uses the "low" eligibility weight (EFMP share of EV purchases), instrumented using predicted subsidy exposure based on pre-period income brackets.