

Green subsidies with demand distortions

Susanna B. Berkouwer[§]

Joshua T. Dean[†]

April 10, 2026

Latest version available [here](#).

Abstract

Standard Pigouvian theory predicts that externalities should be corrected at the margin. However, demand distortions such as credit constraints or behavioral biases create a wedge between marginal benefit and marginal cost. While these distortions can lower aggregate abatement, they can *increase* the efficiency of green subsidy spending. In theory, this happens through two channels: by shifting the marginal adopter toward higher private and social benefits and by increasing demand elasticity. We test these predictions by cross-randomizing fixed cost subsidies, marginal cost subsidies, and loan access for an induction stove among 2,134 charcoal users in Kenya. Marginal cost subsidies that lower electricity costs by up to 75% have a precise zero effect on both adoption and usage. Fixed cost subsidies abate at just US\$13 per ton of CO₂e, and demand distortions are responsible for making this cost low: reducing credit constraints raises abatement costs to US\$22 per tCO₂e. These efficiency gains operate through the two hypothesized channels: demand distortions increase the marginal positive externality by 19% and lower the subsidy cost per marginal abatement by 30%. We estimate the model to generate counterfactual simulations and find that, without any distortions, abatement costs would reach US\$122 per tCO₂e. The social welfare gain would be US\$3.1 per subsidy dollar; demand distortions increase this to US\$20. These results suggest that contexts with larger demand distortions, including many low- and middle-income economies, could generate some of the lowest-cost opportunities on the abatement cost curve.

JEL: O12, Q56

[§]The Wharton School, University of Pennsylvania and NBER; sberkou@wharton.upenn.edu. [†]The Booth School of Business, University of Chicago; joshua.dean@chicagobooth.edu. For generous financial support we thank the International Growth Center, the Weiss Fund for Research in Development Economics, the Penn Global Holman Africa Fund, the Mack Institute, the Wharton Behavioral Lab, and the University of Chicago Development Economics Research Fund. We thank numerous seminar participants for helpful comments. Busara (in particular Suleiman Amanela and Debra Opiyo) superbly implemented field activities. We thank Marta Westerstahl and Rosemary Zhang for excellent research assistance. This study has research approval in Kenya ([MIRERC 069/2025](#) and Nakuru County Unified Trade Permit No.2025/H23H7831401) and the US (University of Pennsylvania IRB Protocol #855529). A [Pre-Analysis Plan](#) was filed with the AEA RCT Registry ([ID: 16832](#)). A disclosure statement is [available here](#).

1 Introduction

The foundational insight of Pigouvian economics is that externalities should be corrected at the margin: a first-best policy sets a tax or subsidy equal to the marginal external cost or benefit (Pigou, 1920). This logic underlies the design of many green subsidy programs worldwide: electric utilities and electric car manufacturers offer subsidized electricity tariffs for electric vehicle owners (PG&E, 2025; Volvo, 2025), and utilities in Kenya and Uganda have introduced e-cooking tariffs to promote electric stoves (Kenya Ministry of Energy, 2024; UEDCL, 2025). In theory, marginal cost subsidies offer key advantages over fixed cost subsidies of equivalent net present value. They can improve selection on the extensive margin by encouraging adoption among higher-usage agents, who stand to benefit more from lower operating costs, and induce treatment effects on the intensive margin by incentivizing greater usage of the clean technology conditional on adoption.

However, many markets face additional distortions like behavioral biases and financial market failures, which cannot be corrected through first best policy in the short run.¹ This creates a wedge between marginal benefit and marginal cost—or between observed demand and a hypothetical frictionless benchmark—at the market equilibrium (Atkin et al., 2025; Bergquist, Lashkari, and Verhoogen, 2026; Ghatak and Mookherjee, 2025; Indarte et al., 2026) and calls for second best policies (Lipsey and Lancaster, 1956).

Demand distortions can lower global abatement: for example, weak institutions in many low- and middle-income countries have exacerbated global capital misallocation, undermining the green transition as cleaner technologies often require more upfront capital (Acemoglu, Johnson, and Robinson, 2005; Nunn, 2007; Shapiro, 2025; Pande et al., 2025). Demand distortions can also dampen responsiveness to price signals, reducing the impact of marginal cost subsidies on both adoption and usage (Gillingham and Palmer, 2014). However, unlike this dampening effect on abatement and on marginal cost subsidies, this paper argues that demand distortions can *increase* the efficiency of *fixed cost* subsidies. We offer theory and evidence to demonstrate that distortions increase green subsidy efficacy not just relative to marginal cost subsidies, but in absolute terms.

This paper’s first contribution is to document in theory how two key channels can cause demand distortions to increase the welfare gains from green subsidies. The first channel is that distortions can push high-benefit agents below the purchase threshold. A fixed cost subsidy then reaches marginal adopters who generate larger consumer surplus. When private and social benefits are positively correlated, as when fuel savings reduce emissions and fuel expenditures, this also increases marginal environmental benefits. The second channel is that distortions can increase demand elasticity, increasing the percentage of additional adopters and reducing the share of expenditure on inframarginal buyers—a widespread concern for environmental subsidy programs (Sallee, 2025; Boomhower and Davis, 2014; Segerson et al., 2024; Gaarder et al., 2026).

Testing the theory requires a setting where we can (i) identify the private and social benefits

¹Solving credit market frictions requires institutional intermediation between borrowers and lenders that overcome classic information asymmetries at the root of moral hazard and adverse selection problems (Bernanke, 2023). Weak institutions pervasive in low- and middle-income countries undermine this (Acemoglu, Johnson, and Robinson, 2005).

of a green technology, (ii) experimentally vary both fixed and marginal cost subsidies, and (iii) experimentally manipulate demand distortions. We do so in the context of induction cookstoves among 2,134 charcoal-cooking households in Nakuru County, Kenya, where capital market distortions are severe: even with sophisticated credit technologies, paying in instalments incurs 75% in interest over 6 months, corresponding to an annualized rate 215% above the central bank lending rate. We cross-randomize three treatments: fixed cost subsidies off the US\$82 price (of 10% or 75%), marginal cost subsidies that lower the electricity cost of operating the stove (by 0%, 25%, or 75%), and access to a loan. The marginal cost subsidy applies only to electricity consumed by the induction stove. Willingness-to-pay is elicited using an incentive-compatible Becker, Degroot, and Marschak (1964) mechanism. Randomly assigned prices generate exogenous variation in adoption, allowing us to causally estimate private and social benefits. Temperature sensors measure charcoal stove usage continuously and electric stove usage is transmitted in real time.

The paper’s second main finding is that induction stoves generate large private and social benefits. Adoption decreases daily charcoal cooking by 31 minutes and reduces monthly charcoal expenditures by US\$10 (66%). After accounting for increased electricity spending of US\$2.3 per month, net monthly energy savings are US\$8.9 (36%), implying a 112% annual return on investment over the stove’s 1.9 year lifetime. Credit market failures heavily distort demand in this context: average willingness-to-pay is only 10% of the net present value of future discounted fuel savings, but when given access to an interest-free loan this increases to 18%. Assuming a 38% fraction of non-renewable biomass (Ghilardi and Bailis, 2024), we estimate that adoption abates 2.5 tCO_{2e} per year across all energy sources. Modest reductions in LPG usage offset a small increase in electricity emissions.

Our third finding is that marginal cost subsidies have no impact. Despite providing detailed information and drawing attention to the reduction in cooking costs before eliciting willingness-to-pay, each US\$1 in electricity subsidy has a US\$0.01 impact on WTP (we can rule out more than US\$0.16). On the intensive margin, despite transferring subsidies weekly via Kenya Power’s standard token system, we estimate a demand elasticity of -0.06 and can rule out that a 1% price decrease increases electric stove usage by more than 0.11%. Both null results hold with and without access to a loan, suggesting the presence of additional demand distortions beyond credit constraints alone, such as inattention to marginal electricity costs or myopia.

Our fourth finding is that fixed cost subsidies increase adoption dramatically. Each US\$1 of subsidy spending generates US\$20 in aggregate welfare gains when valuing reductions in climate change damages at a social cost of carbon (SCC) of US\$230 per ton of CO_{2e} avoided, including US\$2.8 in private welfare gains for adopters. In terms of abatement costs, the subsidies incur US\$13 in government expenditure per ton of CO_{2e} avoided. This is significantly lower than most abatement technologies available today. The resource cost of abatement is *negative*: fuel savings exceed the stove’s production cost, so each ton of CO_{2e} abated saves US\$24 in economic resources.

Combining these results yields the paper’s main, overarching finding: demand distortions in this context decrease abatement cost from US\$22 to US\$13 per ton of CO_{2e} (a reduction of 41%) and increase the welfare gain per US\$1 of government spending by 69%, from US\$12 to US\$20.

Distortions operate through both channels identified in the theory. First, marginal adopters with larger demand distortions generate 30% more in private fuel savings than those whose distortion was partly relaxed (from US\$154 to US\$200 over the lifetime of the stove) and abate 19% more CO₂e (from 2.2 tCO₂e to 2.6 tCO₂e per year), increasing the social benefit from each marginal adopter. Second, credit constraints increase demand elasticity from -0.9 to -2.3, reducing inframarginal expenditure from 38% down to 6.5%. This reduces the subsidy expenditure per marginal sale by 30%, from US\$93 to US\$65, which results from a combination of a greater percentage point increase in marginal adopters and a reduction in the number of inframarginal adopters (who receive the subsidy transfer but generate no emissions reductions).

The experiment only partially relaxes distortions: loan access reduces the estimated demand wedge from 0.11 to 0.19, but willingness-to-pay remains well below discounted private benefits in both conditions. A natural question is then: how efficient would fixed cost subsidies be in a setting without distortions, where willingness-to-pay equals the private benefit (termed ‘frictionless’ demand in Indarte et al. (2026)? To answer this, we estimate the theoretical model using two-step Generalized Method of Moments and simulate counterfactuals across a wider range of distortions. The model is identified using the randomized variation from the experiment. Fully eliminating distortions such that WTP equals the discounted stream of fuel savings would raise abatement costs to US\$122 per tCO₂e (13 times the cost under the credit control condition) and decrease the MVPF down to US\$3.1. We use the model to show how this result depends on two key features: the correlation between private and social benefits, and the heterogeneity of distortions across agents. The positive correlation that characterizes fuel-saving green technologies is what makes distortions beneficial for subsidy efficacy in this class of problems: higher users save more money and generate larger externalities.

Annual carbon mitigation financing exceeds US\$1 trillion (Buchner, 2024). Efficient abatement requires allocating each marginal dollar towards the lowest cost abatement technology. Our results suggest that many of the lowest cost portions of the global abatement supply curve could be found in these contexts. For example, the U.S. has a sophisticated financial market for consumer goods, with around 60% of new cars sold with financing and an additional 20% leased (Experian, 2023). The mechanisms identified in this paper could partly explain why U.S. electric vehicle subsidies cost over US\$1,300 in government expenditure per tCO₂e abated (Hahn et al., 2026)—more than 100 times more costly than the US\$13 per ton abatement cost in our setting. Even holding the underlying technology fixed, our results indicate that green subsidy expenditure could be significantly more effective in contexts with more constrained credit markets.

This paper contributes to several literatures. First, we contribute to a longstanding literature on optimal environmental policy instruments and second best environmental policies (Jacobsen et al., 2020; Knittel and Sandler, 2018; Holland et al., 2016) and optimal environmental policy in the presence of two or more market failures (Fowlie, Reguant, and Ryan, 2016), including those related to behavioral biases (Allcott and Taubinsky, 2015) and those motivating fixed cost subsidies (Carlton and Loury, 1980). We show that demand distortions such as behavioral biases and credit

constraints can *increase* the efficiency of some green subsidies.

Second, we contribute to a literature on selection and mechanism design for subsidy allocation, and in particular using graphical intuition to examine selection (Einav, Finkelstein, and Cullen, 2010). In the absence of demand distortions, a negative correlation between private WTP and social benefits can generate adverse selection into abatement (Aspelund and Russo, 2026). A positive correlation can cause advantageous selection into abatement, but this can then cause adverse selection into subsidy programs (as in Ryan, 2018). Demand distortions can mitigate these effects in contexts with a positive correlation between private and social benefits: this characterizes a broad class of abatement opportunities, including green technologies that reduce fuel usage and thus lower both private energy expenditures and emissions.

Third, we contribute to a literature studying how demand distortions (or ‘wedges’) affect market outcomes (Atkin et al., 2025; Bergquist, Lashkari, and Verhoogen, 2026; Ghatak and Mookherjee, 2025). This includes a large literature in development economics on credit constraints and technology adoption (Banerjee and Duflo, 2014; Duflo, Kremer, and Robinson, 2008; Berkouwer and Dean, 2022; Fowlie and Meeks, 2021). This literature has documented that credit constraints depress adoption of privately beneficial technologies. We show that credit constraints also change the efficiency properties of subsidy instruments: not just whether people adopt, but how cost-effective it is to induce them to do so.

A large empirical literature estimates extensive margin subsidy impacts (Houde and Aldy, 2017; Boomhower and Davis, 2014; Hahn et al., 2026; Borenstein and Davis, 2025) and intensive margin elasticities, largely in high-income countries.² The optimal policy is actively debated (Campbell, 2025), including in the context of marginal versus average cost misperception (Ito and Zhang, 2025; Ito, 2014; Kahn and Wolak, 2013). The precise null on marginal cost subsidies contributes to the literature on energy cost inattention from a very different context. We provide the first experimental comparison of both subsidy instruments in a setting with severe demand distortions. This paper speaks directly to the policy debate around e-cooking tariffs in East Africa and subsidized electricity rates for electric vehicles in the United States (Borenstein, 2025). Research has studied the impacts of innovative pricing schemes on electric vehicle charging, but due to the technological constraints subsidies often inadvertently incentivize the use of electricity for other purposes (Borenstein, 2025).

Finally, we contribute to a growing literature on electric cooking. Using experimental variation, Desbureaux et al. (2025) and Mahadevan et al. (2025) find that fixed cost subsidies increase adoption of electric cooking technologies in the DRC and in Nepal, respectively, which increases electricity usage and generates important social benefits. Gould et al. (2023) find that the nationwide transition from gas towards electric cooking reduced Ecuador’s CO_{2e} emissions due to its predominantly hydroelectric powered grid. Other recent descriptive research and laboratory experiments have

²Inattention (as described in DellaVigna (2009)) to energy costs has been documented in the context of cars (Grigolon, Reynaert, and Verboven, 2018; Sallee, West, and Fan, 2016; Sallee, 2014; Allcott and Wozny, 2014; Busse, Knittel, and Zettelmeyer, 2013; Allcott, 2011; Gillingham, Houde, and Van Benthem, 2021; Kahn, 1986; Klier and Linn, 2010; Bushnell, Muehlegger, and Rapson, 2025), housing (Myers, 2019; Myers, 2020; Palmer and Walls, 2015; Myers, Puller, and West, 2022), and other household technologies (Hausman, 1979; Allcott and Taubinsky, 2015; De Groot and Verboven, 2019; Sejas-Portillo, Moro, and Stowasser, 2025; Gerster and Kramm, 2024).

shed light on perceptions, usage patterns, energy efficiency, and potential for scale of both induction stoves and pressure cookers (Aemro, Moura, and Almeida, 2021; Leary et al., 2021; Saha, Razzak, and Khan, 2021; Clements et al., 2020; Alem, Hassen, and Köhlin, 2014; Banerjee et al., 2016; Rubinstein et al., 2022).

2 Motivation

Electrification is a key component of the world’s decarbonization strategy. Electrification involves replacing technologies that directly combust fuel—such as gasoline vehicles, gas or charcoal stoves, gas water heaters, or gas furnaces—with electric substitutes for those same technologies: electric vehicles, induction stoves, electric water heaters, or heat pumps. It is one of the largest and most fundamental overhauls of our economy that the world has ever embarked on. In combination with rapid decarbonization of grid electricity, electrification has the potential to enable deep and immediate decarbonization, and to mitigate climate change.

Pigouvian taxes are theoretically first-best (Pigou, 1920). However, carbon taxes are often politically intractable: in Canada, a residential carbon tax efficiently designed with lump sum transfers was recently repealed due to political pressures (Department of Finance Canada, 2025). Energy expenditures tend to be inelastic, such that low-income households spend a disproportionate share of their incomes on energy costs. Rising electricity prices in Kenya have seen significant political backlash (CNN, 2023), preventing meaningful political action on this front. Taxation may also be logistically infeasible: Kenya’s charcoal sector is largely informal and recent government intervention trying to ban or fine logging have seen very poor enforcement (Sola and Cerutti, 2021; Daghar, 2021). Technological requirements can also make accurately pricing the marginal externality cost prohibitive (Jacobsen et al., 2020).

Many environmental principals (including for example multilateral donors, environmental agencies, the UNFCCC, voluntary carbon offsets markets, philanthropists, and other investors) furthermore lack the legal authority to levy a tax. Instead, their focus is on subsidizing climate change mitigation. Doing so efficiently requires identifying technologies and contexts with the lowest marginal abatement costs, starting at the lowest end of the abatement supply curve and then moving up over time (Gillingham and Stock, 2018; Kolstad, 2011). Demand distortions may reduce the *level* of abatement but generate the largest return to *marginal* abatement expenditure. Abatement spending would then be most effective in contexts with more severe demand distortions: for example, low- and middle income economies with significant financial market failures.

Well-functioning credit markets require institutional intermediation between borrowers and lenders to overcome classic information asymmetries (Bernanke, 2023). However, institutional failures are pervasive in low- and middle-income countries (Acemoglu, Johnson, and Robinson, 2005), weakening credit markets. Particular institutional shortcomings include lack of a centralized credit bureau to intermediate borrower information; incomplete property rights undermining collateral usage; widespread informality (1.7 billion people lack a bank or mobile money account); relational

economics, which can make threats of punishment non-credible; a lack of deposit insurance to stimulate saving; and constrained legal systems impeding debt recovery. These exacerbate the classic information asymmetries, adverse selection, and moral hazard problems facing financial markets.

As a result, credit market failures are pervasive in many low- and middle-income countries, creating pecuniary or non-pecuniary borrowing costs (such as high interest rates, short repayment periods, short grace periods, and appeals to social networks for repayment) that depress adoption of even technologies with large private benefits by firms (Banerjee and Duflo, 2014; Mel, McKenzie, and Woodruff, 2008; Fafchamps et al., 2014; Kremer et al., 2013; McKenzie and Woodruff, 2008; McKenzie and Woodruff, 2006; Banerjee and Duflo, 2005; Quinn and Woodruff, 2019) and households (Blattman, Fiala, and Martinez, 2014; Duflo, Kremer, and Robinson, 2008; Anagol and Udry, 2006), and in the adoption of green technology in particular (Adhvaryu, Kala, and Nyshadham, 2020; Berkouwer and Dean, 2022; Rom, Günther, and Pomeranz, 2023; Fowlie and Meeks, 2021). Agents may also face psychological borrowing frictions that serve to prevent costly default (Martínez-Marquina and Shi, 2024; Prelec and Loewenstein, 1998; Sussman and Shafir, 2011).

Demand distortions vary in the degree to which a policy maker can correct them directly. Plenty of demand distortions could be addressed directly through a first-best policy intervention: as an example, widespread undervaluation of potential energy savings could be corrected through an information campaign. In this case, the policy maker might consider the costs of targeting the distortion directly against the fixed cost subsidy to determine the optimal policy. This paper focuses specifically on distortions that cannot be targeted directly, such as credit market failures resulting from underlying institutional capacity constraints as described above. For example, these constraints can also hamper the availability of appropriate insurance instruments, exacerbating risk aversion and distorting demand. The model is largely agnostic as to the source of the demand distortion.

Low- and middle-income countries are also home to billions of people who lack modern cooking technologies (World Bank, 2020). Across Africa 73% of households (more than 1 billion people) use charcoal or firewood as their primary cooking fuel. Given steady population growth across Africa, projections forecast approximately one billion people across Africa to still be cooking with biomass as their primary cooking fuel in 2050 (IEA, 2025). At the same time, in the past two decades more than 2 billion people have gained access to electricity (World Bank, 2024). The rapid rise of household electricity access combined with the persistent use of biomass fuel for cooking has generated widespread policy interest in ‘leapfrogging’ populations towards electric cookstoves. Per the Organization of the Petroleum Exporting Countries (OPEC) fund for international development, “most specialists agree that electric cooking from renewable energy sources remains the ideal solution over the longer term” (OPEC, 2024).

3 Theory

An agent consumes energy services (e.g. transportation or cooking) using an incumbent dirty technology (e.g. a gasoline vehicle or a charcoal cookstove). The agent may adopt a cleaner substitute (e.g. an electric vehicle or an electric cookstove) at a one-time cost p equal to the marginal cost of producing the technology, and then choose how much of their energy consumption to shift to the clean technology. This framework builds on Allcott, Mullainathan, and Taubinsky (2012; 2014) and the theory of the second best developed in Lipsey and Lancaster (1956).

Agents are heterogeneous in their demand for energy services θ_i and in a (dis)taste parameter α_i for the clean technology, distributed $g(\alpha)$ with an increasing hazard rate.³ Each agent receives income Y and derives linear utility from a numéraire good. We assume that total demand for energy services is perfectly inelastic: θ_i is fixed and does not respond to prices. The agent’s only choice is how to allocate consumption between the dirty and clean technologies. We also assume no intertemporal discounting; the one-time adoption cost and the stream of per-period fuel savings are directly comparable.

The dirty technology has per-unit fuel cost f_H and emits ϕ_H per unit of energy service. The clean technology has per-unit fuel cost $f_L < f_H$ and emits $\phi_L < \phi_H$. Define, respectively, the per-unit cost saving and the per-unit externality reduction from using the clean technology:

$$\pi \equiv f_H - f_L > 0 \quad \phi \equiv \phi_H - \phi_L > 0$$

Policy instruments The principal has access to three instruments: (1) a fixed cost subsidy s_1 that reduces the one-time adoption price to $p - s_1$; (2) a marginal cost subsidy s_2 on the clean fuel, reducing its per-unit cost to $f_L - s_2$; (3) a marginal cost tax τ on the dirty fuel, raising its per-unit cost to $f_H + \tau$. Under policies s_2 and τ , the effective relative fuel price becomes:

$$\pi' = (f_H + \tau) - (f_L - s_2) = \pi + \tau + s_2$$

A tax τ and a marginal subsidy of $s_2 = \tau$ differ in the direction of the fiscal transfer but have identical effects on the relative fuel price signal facing the agent.

The agent’s problem At $t = 1$, agent i makes an adoption decision $d_{1i} \in \{0, 1\}$: whether to buy the clean technology at price $p - s_1$. If they adopt, at $t = 2$ they choose a usage fraction $d_{2i} \in [0, 1]$: the share of energy services θ_i supplied with the clean technology. Each unit shifted to the clean technology saves π' in fuel costs but generates a non-pecuniary disutility $\mu'(d_{2i}\theta_i - \alpha_i)$, capturing switching costs such as learning or habit change. Higher- α_i agents have less distaste for the clean technology. $d_{2i}\theta_i$ denotes the total quantity of clean energy services consumed, and $d_{2i}\theta_i - \alpha_i$ denotes the total quantity of clean energy services consumed in excess of their taste parameter, α_i .

³The economic content of this distributional assumption is that we don’t have too many “super users”, which ensures that the demand elasticity is monotonically increasing in price. This is satisfied for many common distributions including uniform, normal, and exponential. It is violated for heavy-tailed distributions like pareto or log-normal.

We assume $\mu(0) = 0$, $\mu'(0) = 0$, $\mu'' < 0$, and $\lim_{x \rightarrow \infty} \mu'(x) = -\infty$, so that $\mu(x) \leq 0$ for all $x \geq 0$. Total utility from usage conditional on adoption is:

$$f(d_{2i}, \pi', \theta_i, \alpha_i) = d_{2i}\theta_i\pi' + \mu(d_{2i}\theta_i - \alpha_i)$$

Allowing non-pecuniary impacts from adoption ($\mu(0) \neq 0$) does not change the results meaningfully. We discuss this in [Section 3.5](#) below with formal proofs in [Section III.4](#).

The agent chooses d_{2i}^* to maximize $f(\cdot)$. The agent shifts clean-energy consumption until the marginal disutility equals the per-unit fuel saving: $\mu'(d_{2i}^*\theta_i - \alpha_i) = -\pi'$. Define x^* as the unique solution to $\mu'(x^*) = -\pi'$. An agent who adopts receives:

$$u_1 = Y - (p - s_1) + f(d_{2i}^*, \pi', \theta_i, \alpha_i)$$

The agent adopts iff $u_1 > Y$, i.e. iff their willingness-to-pay $wtp_i \equiv f(d_{2i}^*, \pi', \theta_i, \alpha_i)$ exceeds $p - s_1$. Market-level demand at price p is given by $Q(p) = \Pr(wtp_i > p)$.

It will be useful to define the constant a as capturing the common component of WTP shared by all agents, which are the fuel savings and switching costs evaluated at the common usage in excess of their taste parameter x^* :

$$wtp_i = (x^* + \alpha_i)\pi' + \mu(x^*) = \underbrace{x^*\pi' + \mu(x^*)}_{\equiv a} + \alpha_i\pi'$$

Given that a and π' are common across all agents, heterogeneity in WTP depends only on α_i . WTP and abatement both increase in α_i : agents with a stronger taste for the clean technology use it more, abate more, and save more on fuel, and are willing to pay more.

Lemma 1. *Willingness-to-pay is increasing in α_i and does not depend on θ_i , and heterogeneity in usage and willingness-to-pay depend only on α_i .* Proof in [Section III.1](#)

Intertemporal distortions (γ). Intertemporal demand distortions reduce the perceived value of wtp_i at the time by $\gamma \in (0, 1]$ of the adoption decision but do not affect usage conditional on adoption. Examples include credit constraints (where $\gamma = \frac{1}{1+r}$ reflects an inflated borrowing rate), myopia, or present bias:

$$\widetilde{wtp}_i = \gamma wtp_i = \gamma[(x^* + \alpha_i)\pi' + \mu(x^*)]$$

Marginal distortions (δ). Marginal distortions reduce the agent's perceived relative fuel price to $\delta\pi'$ with $\delta \in (0, 1]$. This reduces usage $\tilde{x}^* < x^*$ which in turn decreases savings and non-pecuniary costs, both of which decrease WTP. Examples include inattention to fuel costs or incorrect beliefs about savings. The agent discounts the savings portion of the technology but experiences and anticipates switching costs in full. Define \tilde{a} to be the distorted counterpart of a and define $k = \tilde{a} - \delta a$.

WTP is then lowered by δ and by the reduction in expected usage ($k \leq 0$):

$$\widetilde{wtp}_i = \delta wtp_i + k \leq \delta wtp_i$$

Lemma 2. *Marginal distortions reduce WTP by at least as much as proportional scaling by δ .*

Proof in [Section III.1](#).

This is distinct from intertemporal distortions, where the agent neglects both the future savings and the future switching costs by the same fraction. Both classes of distortion dampen the efficacy of instruments that operate through the relative fuel price by reducing their impact on adoption. Marginal distortions further reduce their efficacy by also dampening their impact on usage.

Lemma 3. *Demand distortions reduce the efficacy of marginal cost instruments, but fixed cost subsidies bypass the relative fuel price channel.*

Proof in [Section III.1](#).

3.1 Why distortions increase the efficacy of fixed cost subsidies

Demand distortions increase fixed cost subsidy efficiency through two channels. We first present intuition for both channels. [Section 3.2](#), [Section 3.3](#), [Section 3.4](#) then define and derive results for a social planner, a social principal, and an environmental principal.

Channel 1: The marginal adopter generates a larger externality.

Lemma 4. *Demand distortions shift the marginal adopter to have higher private benefit, and if private and social benefits are positively correlated then this shifts the marginal adopter and thus the average adopter to generate larger social benefits.*

$$\frac{\partial \underline{\alpha}}{\partial \gamma} < 0 \quad \frac{\partial \underline{\alpha}}{\partial \delta} < 0$$


Proof in [Section III.1](#).

Define the marginal adopter as $\underline{\alpha} : \gamma wtp_i(\underline{\alpha}) = p - s_1$. For any fixed subsidized price $p - s_1$, any decrease in γ increases the marginal agent's wtp_i . Since wtp_i is positively correlated with α_i and thus with ϕ_i (since $\phi_i = (\tilde{x}^* + \alpha_i)\phi$), this increases the level of abatement for the marginal adopter.⁴ Similarly, any increase in the marginal distortion will reduce $wtp_i(\alpha_i)$ for all α_i and increase the α_i required to be marginal, increasing the level of abatement for the marginal adopter.

This mechanism does not depend on the model's specific structure. It operates whenever a distortion reduces WTP and there is a positive correlation between WTP and the externality reduction. Demand distortions push high-value agents to just below the adoption threshold. Those agents generate the largest social benefits of all non-adopters: a fixed cost subsidy that brings these agents back into the adopter pool therefore captures more abatement per marginal adopter than it would in a market without distortions. This holds under both intertemporal and marginal distortions.

⁴We discuss the case when private and social benefits are not positively correlated below.

Figure 1: How demand distortions and the shape of the externality determine the role of selection



Notes: Demand distortions increase the marginal positive externality if willingness-to-pay and social benefits are positively correlated ($A1 > A2$) but decrease it when they are negatively correlated ($B1 < B2$). A positive correlation generates advantageous selection that lowers the marginal externality of a subsidy in regular markets ($A2 < B2$) but increases it when markets face distorted demand ($A1 > B1$).

This positive correlation characterizes a broad class of green technology adoption problems where higher users generate larger private benefit and abate more emissions—for example because higher users save more money or derive more personal utility from switching, and because abatement is proportional to usage. Panel A of Figure 1 shows intuitively the average social benefit exceeds the externality at the margin using the framework of Einav, Finkelstein, and Cullen (2010). In line with Lemma 4 above, a demand distortion increases the marginal externality, such that $A1 > A2$, and therefore the average externality.

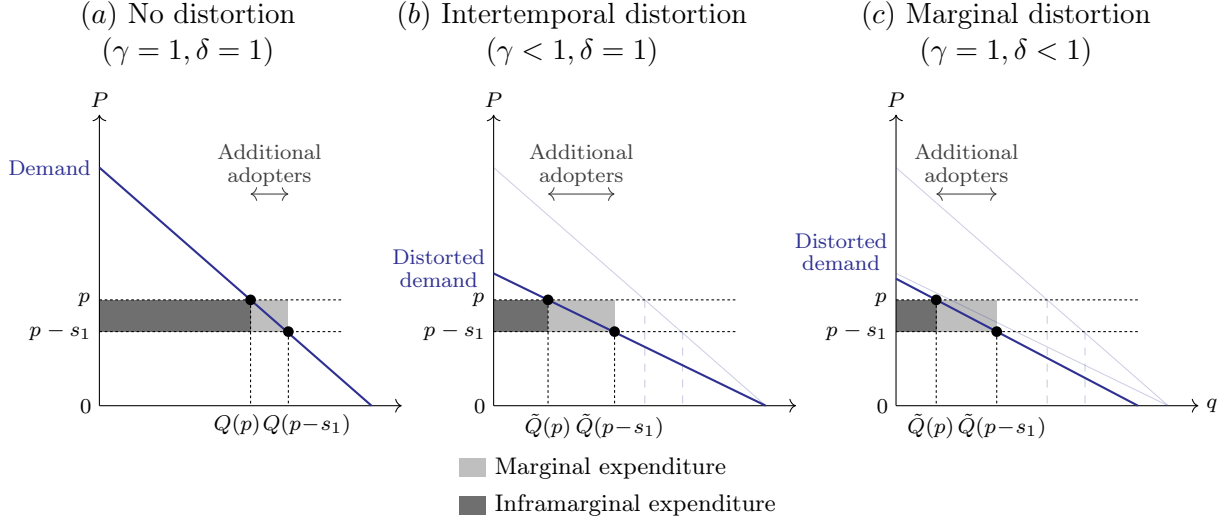
In undistorted markets, this positive correlation can generate advantageous selection into the market resulting in adverse selection into who remains marginal to the subsidy ($A2 < B2$), lowering the benefit of subsidies (as in Chen, Ryan, and Xu, 2025). The presence of a demand distortion reverses this effect, with the positive correlation now increasing the marginal externality: $A1 > B1$. When private and social benefits are negatively correlated (for example as in Aspelund and Russo, 2026), the distortion lowers the marginal externality: $B1 < B2$. Section 3.5 discusses this further.

Channel 2: Demand distortions increase demand elasticity. Demand distortions can increase demand elasticity, which decreases the share of subsidy recipients that are inframarginal, reducing the cost per marginal adopter. Figure 2 illustrates this intuition graphically: the distortion shrinks the inframarginal region (dark) relative to the marginal region (light).

Lemma 5. *Demand distortions increase demand elasticity $|\mathcal{E}_D|$, increasing the number of marginal adoptions per dollar of subsidy spending and decreasing the marginal cost per additional adoption.*

The distorted elasticity at price p equals the undistorted elasticity at the higher effective price p/γ . Since $|\mathcal{E}|$ is increasing in price for most standard demand curves (with the exception of isoelastic

Figure 2: Marginal and inframarginal subsidy expenditures with and without distortions



Notes: All panels show subsidy expenditures from reducing the price from p to $p - s_1$. Dark regions represent spending on inframarginal buyers; light regions represent spending on marginal buyers. Panel (b): an intertemporal distortion (γ) compresses the demand curve toward the origin by rotating around the quantity intercept, since at zero price all agents still adopt. Panel (c): a marginal distortion (δ) both rotates and shifts the demand curve leftward. In both cases, the number of marginal adopters increases and spending on inframarginal adopters decreases.

demand, where it is constant), intertemporal distortions increase demand elasticity.

$$\tilde{\mathcal{E}}(p) = \frac{p}{\tilde{Q}(p)} \tilde{Q}'(p) = \frac{p}{Q(p/\gamma)} \cdot \frac{Q'(p/\gamma)}{\gamma} = \mathcal{E}(p/\gamma)$$

Marginal distortions also increase demand elasticity (derivation in [Section III.6](#)). This requires a stronger condition (our assumed increasing hazard rate) than the γ case (which only needed $|\mathcal{E}|$ increasing in p), but holds for a broad class of empirically relevant distributions of α_i .

Define the principal's cost $C(s_1) = s_1 Q(p - s_1)$. Marginal cost per adoption can be expressed as a function of demand elasticity, with an increase in subsidy equaling an increase in expenditures to inframarginal consumers plus additional expenditures to marginal consumers:

$$\frac{dC(s_1)}{dQ(p - s_1)} = \frac{\partial C(s_1)}{\partial s_1} \cdot \frac{\partial s_1}{\partial Q(p - s_1)} = s_1 - \frac{Q(p - s_1)}{Q'(p - s_1)} = s_1 - \frac{p - s_1}{\mathcal{E}_D}$$


Aggregate impact of distortions on green subsidy efficacy

Proposition 1. *Demand distortions increase total abatement per dollar of fixed subsidy expenditure.*

Letting A be either the sum of private benefits and emissions reductions, or emissions reductions alone, distortions then impact the cost-efficiency of abatement as follows:

$$\frac{\partial}{\partial \gamma} \left(\frac{\partial A}{\partial C} \right) = \frac{\frac{\partial^2 A}{\partial Q, \partial \gamma} \cdot \frac{\partial C}{\partial Q} - \frac{\partial A}{\partial Q} \cdot \frac{\partial^2 C}{\partial Q, \partial \gamma}}{\left(\frac{\partial C}{\partial Q} \right)^2}$$

Figure 3: Aggregate effect of demand distortions on private benefit, externality, and subsidy cost



Notes: The dotted green area represents the total positive externality generated, the striped blue area represents the consumer surplus generated, and the gray area represents the subsidy expenditure incurred when a subsidy lowers the price to $p - s_1$. Subsidy efficiency is the ratio of the sum of externality benefit and consumer surplus to total subsidy expenditures. Through the two channels identified above, demand distortions can increase this ratio.

From Lemma 5 we have $\frac{\partial^2 C}{\partial Q \partial \gamma} > 0$: a larger demand distortion ($\gamma < 1$) reduces average subsidy cost. From Lemma 4 we have $\frac{\partial^2 A}{\partial Q \partial \gamma} < 0$: a larger demand distortion increases average abatement. Given somewhat trivially that $\frac{\partial C}{\partial Q}, \frac{\partial A}{\partial Q} > 0$ we can sign the LHS to be negative: distortions (lower γ) increase abatement per dollar of subsidy expenditure. Figure 3 shows these results intuitively. Under an increasing hazard rate, the marginal distortions case is symmetric.

3.2 The social planner's problem

The planner chooses (s_1, s_2, τ) to maximize total welfare, which includes the externality reduction plus the private value of adoption (fuel savings net of switching costs) minus the technology cost for all adopters. Define the *social value* of agent i adopting:

$$SV_i(\alpha_i) \equiv \underbrace{(\tilde{x}^* + \alpha_i)\phi}_{\text{externality reduction}} + \underbrace{(\tilde{x}^* + \alpha_i)\pi + \mu(\tilde{x}^*)}_{\text{private value (true price, actual usage)}} - \underbrace{p}_{\text{technology cost}}$$

The planner's instruments affect both who adopts and how much adopters use. Without distortions, the first-best policy taxes each fuel at its marginal external cost ($\tau^* = \phi$), with no fixed cost subsidy ($s_1^* = 0$). At this tax $\pi' = \pi + \phi$, so the agent's WTP internalizes the externality. The adoption rule $wtp_i \geq p$ is therefore equivalent to $SV_i \geq 0$, which is the social optimum. This is the standard Pigouvian result (Pigou, 1920).

Proposition 2. *Intertemporal demand distortions ($\gamma < 1$) increase the social planner's optimal fixed subsidy while marginal demand distortions ($\delta < 1$) increase the optimal marginal tax:*

$$s_1^* = p(1 - \gamma) \quad \tau^* = \frac{\phi + \pi(1 - \delta)}{\delta}$$

Intuitively, the δ distortion can be fully corrected by the optimal tax, which exceeds the Pigouvian benchmark of $\tau = \phi$. The fixed cost subsidies can then compensate for γ .

3.3 The social principal's problem

As discussed in [Section 2](#), this paper focuses on contexts where a planner faces a budget constraint and a tax on the polluting good is infeasible. With only subsidies available, the principal faces a genuine tradeoff between fixed and marginal cost subsidies. The principal chooses (s_1, s_2) to maximize aggregate social welfare subject to a fixed budget B :

$$(s_1^*, s_2^*) = \arg \max \int_{\underline{\alpha}}^{\infty} SV_i(\alpha_i) g(\alpha) d\alpha \quad \text{s.t.} \quad s_1 Q + s_2 U \leq B$$

where $\underline{\alpha}$ corresponds to the α_i of the marginal adopter such that $wtp_i(\underline{\alpha}) = p - s_1$, $Q \equiv 1 - G(\underline{\alpha})$ (number of adopters), and $U \equiv \int_{\underline{\alpha}}^{\infty} (\tilde{x}^* + \alpha_i) g(\alpha) d\alpha$ (total clean energy usage by adopters).

The first term is the fixed cost subsidy paid to every adopter; the second is the marginal cost subsidy paid on every unit of clean energy used. Both q and U depend on the instruments through $\underline{\alpha}$ (who adopts) and \tilde{x}^* (how much each adopter uses): s_1 lowers the adoption threshold directly, while s_2 raises the effective fuel price $\pi' = \pi + s_2$, which increases both usage and WTP. The agent responds to π' but the principal values fuel savings at the economic cost difference of π .

The FOC for the fixed cost subsidy s_1 is as follows:

$$\underbrace{(\tilde{x}^* + \underline{\alpha})(\pi + \phi) + \mu(\tilde{x}^*)}_{\text{fuel savings, abatement, and switching costs from one new adopter}} - \underbrace{p}_{\text{technology cost from one new adopter}} = \lambda \left[\underbrace{s_1 + s_2(\tilde{x}^* + \underline{\alpha})}_{\text{subsidy to new adopter}} + \underbrace{\frac{p - s_1}{|\tilde{\mathcal{E}}(p - s_1)|}}_{\text{inframarginal cost}} \right] \quad (1)$$

Proof in [Section III.2.](#)

Unlike s_1 , the marginal cost subsidy s_2 generates abatement on both the intensive margin (existing adopters use more) and the extensive margin (new adopters enter). The FOC for s_2 is given by:

$$\underbrace{\left[(\pi + \phi) + \mu'(\tilde{x}^*) \right] \frac{d\tilde{x}^*}{ds_2} Q}_{\text{intensive: social value of additional usage}} + \underbrace{\left[(\tilde{x}^* + \underline{\alpha})(\pi + \phi) + \mu(\tilde{x}^*) - p \right] g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_2} \right)}_{\text{extensive: social value of new adopters}} \\ = \lambda \left[\underbrace{U}_{\text{cost of existing usage}} + \underbrace{s_2 \frac{d\tilde{x}^*}{ds_2} Q}_{\text{cost of additional usage}} + \underbrace{[s_1 + s_2(\tilde{x}^* + \underline{\alpha})] g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_2} \right)}_{\text{subsidy package to new adopters}} \right] \quad (2)$$

Proof in [Section III.2.](#)

Combining the two FOCs characterizes the optimal allocation between s_1 and s_2 . Write $\bar{\alpha} \equiv$

$E[\alpha_i \mid \alpha_i \geq \underline{\alpha}]$ for the mean taste among adopters, $U/Q = \tilde{x}^* + \bar{\alpha}$, and define $MSV \equiv (\pi + \phi) + \mu'(\tilde{x}^*)$ (the marginal social value of one additional unit of clean energy usage). Define $\kappa \equiv (p - s_1)/|\tilde{\mathcal{E}}(p - s_1)|$ (the inframarginal cost per new adopter) and $\Delta \equiv MSV \cdot (\tilde{x}^* + \underline{\alpha}) - SV(\underline{\alpha}) = \mu'(\tilde{x}^*)(\tilde{x}^* + \underline{\alpha}) - \mu(\tilde{x}^*) + p$ (how much the value of an additional unit of usage by the marginal adopter differs from the total social value generated by their adoption). The optimal combination of s_1, s_2 is then pinned down by the budget constraint and by:

$$\frac{MSV \cdot (s_1 + \kappa) + s_2 \cdot \Delta}{SV(\underline{\alpha})} = \frac{\overbrace{(\tilde{x}^* + \bar{\alpha}) - \gamma\delta(\tilde{x}^* + \underline{\alpha})}^{\text{extensive margin disadvantage of } s_2}}{\underbrace{\frac{d\tilde{x}^*}{ds_2}}_{\text{intensive margin advantage of } s_2}} \quad (3)$$

Proof in [Section III.2](#).

When $\gamma\delta < 1$, each dollar of s_2 raises the marginal adopter's WTP by only $\gamma\delta(\tilde{x}^* + \underline{\alpha})$ rather than $(\tilde{x}^* + \underline{\alpha})$, widening s_2 's extensive margin disadvantage.

How does an intertemporal distortion γ affect the optimal fixed cost subsidy? We set $\delta = 1$ throughout to isolate γ , so that $\tilde{x}^* = x^*$ and usage is not distorted.

$$\frac{ds_1^*}{d\gamma} = \underbrace{-wtp(\underline{\alpha}^*)}_{\text{direct effect}} - \underbrace{\gamma\pi' \cdot \frac{d\underline{\alpha}^*}{d\gamma}}_{\text{re-optimization}} \quad (4)$$

Proof in [Section III.2](#).

The first term is the direct effect, which is always negative: when γ rises (distortions weaken), agents value the technology more, so the principal can reduce s_1^* while holding the adoption threshold fixed. The second term captures the principal's re-optimization of the adoption threshold. The appendix shows that under IHR, $ds_1^*/d\gamma < 0$ for all demand elasticities.

Proposition 3. *Intertemporal distortions $\gamma < 1$ increase the social principal's optimal fixed cost subsidy s_1^* .*

Proof in [Section III.2](#).

3.4 The environmental principal's problem

We now consider an environmental principal whose only objective is to reduce carbon emissions. The budget constraint, adoption condition, and cost structure are identical to the social principal's problem. The only difference is the objective: the principal maximizes total abatement $A = \phi U$ rather than social welfare $\int SV_i g(\alpha) d\alpha$. Deriving and combining the two FOCs then characterizes the optimal policy (s_1, s_2) :

$$\underbrace{\frac{s_1 + \kappa}{\tilde{x}^* + \underline{\alpha}}}_{\text{cost of } s_1 \text{ per unit abated}} = \frac{\overbrace{(\tilde{x}^* + \bar{\alpha}) - \gamma\delta(\tilde{x}^* + \underline{\alpha})}^{\text{extensive margin disadvantage of } s_2}}{\underbrace{\frac{d\tilde{x}^*}{ds_2}}_{\text{intensive margin advantage of } s_2}} \quad (5)$$

Since the budget constraint and $ds_1^*/d\gamma$ decomposition are identical to the social principal's, and the environmental principal's instrument mix equation is a specialization of the social principal's (with $MSV = \phi$, $SV(\alpha) = \phi \underline{y}$, and $\Delta = 0$), Proposition 3 extends to the environmental principal under IHR; [Section III.3](#) shows why.

3.5 Model discussion and possible extensions

Non-pecuniary benefits Allowing non-pecuniary benefits to clean technology adoption ($\mu(0) = c$) does not change the results above meaningfully, because the mechanisms on which they depend do not involve the level of $\mu(0)$. Channel 1 persists because WTP and abatement both remain increasing in α_i such that distortions exclude the highest-externality agents. WTP and abatement are both affine in α_i with the same-sign coefficient when $\pi' > 0$. Marginal distortions continue to reduce the WTP response to fuel price instruments ($\partial \widetilde{wtp}_i / \partial \tau < \partial wtp_i / \partial \tau$). The envelope theorem and the inequality $\widetilde{x}^* < x^*$ both hold whenever $\pi' > 0$, regardless of $\mu(0)$. The social planner's result remains unchanged because $SV_i + p = \widetilde{wtp}_i$ at the optimal tax is independent of $\mu(0)$. The structure of equation (3) is also unchanged; it depends on the adoption and usage conditions and the Lagrangian, none of which involve $\mu(0)$. These results are derived more formally in [Section III.4](#).

Negative correlation between private and social benefits The correlation between private and social benefits determines whether demand distortions increase or decrease the externality of the marginal adopter. When WTP and abatement are negatively correlated, distortions push low-externality agents below the threshold, reversing the mechanism of channel 1. For example, suppose the dirty fuel is cheaper than the clean fuel: then WTP is *decreasing* in α_i (high-taste agents incur large fuel premiums) while abatement remains increasing in α_i for interior agents. The agents who value adoption most (low α_i , who enjoy the non-pecuniary benefit without heavy usage) generate the *least* abatement, while the agents who abate the most (high α_i , heavy users) have the lowest WTP. In the EV example, status buyers who rarely drive might adopt first, but the committed daily drivers who would displace the most gasoline are loath to adopt. The new marginal adopter under the distortion has lower α_i and generates less abatement. Channel 1 now works *against* fixed cost subsidies: a subsidy that brings excluded agents back into the adopter pool reaches marginal adopters who generate less abatement per dollar than they would absent the distortion. This is the mirror image of the baseline result.

Even with negatively correlated WTP and abatement, channel 2 (increased demand elasticity) still operates. Which channel dominates is then an empirical question—the counterfactual estimation in [Section 8](#) examines this in our context. [Section III.4](#) discusses these mechanisms more formally and offers some additional implications.

Fiscal considerations Carbon taxes and marginal cost subsidies are symmetric in their price signal but differ fiscally. The revenue of a marginal cost tax could fund fixed cost subsidies, com-

binning the fiscal benefits of taxation with the behavioral effectiveness of upfront price reductions. A marginal cost subsidy on the other hand is pure expenditure. When γ is small, agents do not respond to the price signal, so most expenditure flows to inframarginal usage with little behavioral change. Under severe demand distortions, the optimal policy may be a carbon tax (for revenue) paired with a fixed cost subsidy (for behavioral change), rather than a marginal cost subsidy.

Standards and bans Standards and bans are common alternatives to the direct pricing of externalities (Fowlie, Reguant, and Ryan, 2020; Goulder and Parry, 2008), and can be more efficient than taxes when customers misperceive energy costs (Houde and Myers, 2025). These can be effective not only in high-income countries (where most evidence is generated) but also in low- and middle-income countries: as an example, Kenya has promoted the adoption of energy-efficient lighting through CFL distribution programs and efficiency standards (Kovacova, 2026; Wambui, 2010), and is now one of Africa’s largest LED markets (IMARC Group, 2025).

In the context of cooking however, these policy tools face both political and enforcement hurdles. Much of Kenya’s charcoal-related deforestation is informal. The Kenyan government has in the past passed charcoal logging bans, but these have been largely ineffective due to the gaps in enforcement (Sola and Cerutti, 2021; Daghar, 2021). Traditional cookstoves are often produced locally and thus suffer from similar enforcement constraints caused by widespread informality. Given widespread informality in both the charcoal production sector and in stove manufacturing, policy instruments like bans, standards, or taxes are inefficiently costly to enforce.

Open questions Future research could explore a range of model extensions. First, political constraints might require a policy maker earmark an environmental tax for an accompanying environmental subsidy. The model above could accommodate a budget neutral policy such as this.

Second, heterogeneity in γ across agents could shift the optimal instrument mix towards marginal pricing as this increases targeting on the externality—the counterfactual estimation in Section 8 assumes a functional form on the distribution of γ across agents.

Third, loss aversion could play a role if agents weight losses (in this context, taxes) more than gains (subsidies). Additional work could explore whether this restores efficacy of carbon taxes under distortions, and estimate how large this asymmetry would need to be to generate the marginal distortions observed in many papers studying marginal subsidies (including ours).

Finally, solving explicitly for the optimal instrument mix (s_1^*, s_2^*, τ^*) as a function of γ could offer a threshold at which the distortions get sufficiently large at which point marginal pricing alone is no longer optimal.

Additional functional form assumptions Some of the results above simplify greatly when making additional form assumptions. As an example, Section III.5 presents results assuming quadratic switching costs $\mu(x) = c - \frac{x^2}{2}$.

4 Experimental design and measurement

4.1 Setting

East Africa is at the forefront of the global push towards electric cooking. 65% of Kenyan households use wood or charcoal as their primary cooking fuel (Table A7). The most common biomass stove in Kenya is a *jiko*, shown on the left in Figure 4. As a back of the envelope calculation, at baseline respondents use on average 3.1 kilograms of charcoal per day, the production and burning of which emits approximately 4.9 tCO_{2e} per year (Table 1 presents these and additional household statistics; Section 4.6 discusses the conversion from kilograms of charcoal to tons of CO_{2e} emitted). This is roughly the same as the annual emissions of residential gasoline vehicles: the average U.S. household drives a gasoline vehicle 11,500 miles per year at 22.2 miles per gallon of gasoline, emitting 8.9 kilograms of CO₂ per gallon or 4.6 tons of CO₂ per year (EPA, 2025).

To mitigate the negative externalities associated with charcoal production and to stimulate revenue for the financially distressed electric utility, in 2023 Kenya’s national electricity distribution launched its “Pika na Power”, or “cook with electricity” program. It aims to increase uptake of electric cooking “to 5% (500,000 customers) in the short term and to 10% in the medium term” (Kenya Power, 2023). The Kenyan government’s “National Electric Cooking Strategy Action Plan” (Kenya Power, 2024) announced several programs to incentivize electric cooking, including VAT waivers for new electric appliances, behavior change campaigns, expanded support for consumer financing, investments in grid capacity and reliability. This also included a lower ‘e-cooking’ electricity tariff for customers consuming between 31-100kWh per month designed to make electric cooking more attractive relative to LPG and biomass stoves. Electric cooking is a particularly effective climate mitigation strategy given the low average emissions of many African grids. In 2023, 48% of Kenya’s annual on-grid electricity generation was geothermal, 21% was hydroelectric, 18% was wind power, and 3% was solar power (Kenya Power, 2023).

Other East African countries are on similar paths. Uganda’s 2022 Nationally Determined Contribution (NDC) was for “electricity to reach 50% of cooking fuel share by 2025”, including “decreasing the share of biomass energy used for cooking from 88% at baseline to 40% in 2030” (Republic of Uganda, 2022). In 2021, Tanzania adopted a National Environmental Policy stating that “the government shall promote affordable, accessible and reliable alternative energy to charcoal and firewood so as to reduce wood-biomass energy dependency” (Republic of Tanzania, 2021).

4.2 The ECOA induction stove

We study the induction stove manufactured by ECOA (formerly Burn Manufacturing) shown on the right in Figure 4. Stoves are manufactured at their factory in Ruiru outside Nairobi, Kenya. Purchase of an induction stove also includes an induction pot, allowing the buyer to continue using their old pot on their old charcoal stove if they wish (‘stacking’) — for example during a power outage or because some foods taste better when cooked on charcoal. The induction stove has a power rating of 200W–2000W depending on the desired functionality (from low simmer to high

Figure 4: Charcoal stove and induction stove



Notes: On the left is a charcoal stove instrumented with a high-frequency temperature sensor. On the right is the induction stove we study with the three pots that are included in the purchase bundle.

boil) and a voltage rating of 240–250VAC and 50–60Hz.

At the time of our study launch in August 2025, more than 50,000 induction stoves had been sold, primarily through women’s group and pilot programs implemented through their sales agent programs in Tanzania and Kenya. That said, as the stoves were not yet commercially available, our study population was generally not familiar with the product.

We estimate the expected lifetime of the stove by applying a Cox proportional hazard model to the 23,444 stoves that had been sold on or before December 31, 2024.⁵ Using this method, the median induction stove sold by Burn has a lifespan of 1.9 years (Figure C1 shows the survival curve).

The marginal cost of using the stove Domestic residential customers in Kenya face an increasing block tariff. At the start of our study, customers consuming on average less than 30 kWh per month paid US\$0.166 per kWh, customers consuming on average between 30–100 kWh per month paid US\$0.20 per kWh, and customers consuming on average more than 100 kWh per month paid US\$0.22 per kWh.⁶ In our sample, 86% consume less than 30 kWh per month and 99% consume less than 100 per month; since bands can change, to minimize error we use the lowest tariff band in our calculations.

The fixed cost of buying the stove At the time of the study, the stove with one pot was priced at US\$72 in Kenya, with additional pots available for purchase separately. During the study, we offer participants a bundle consisting of the stove plus three pots, which had a marginal cost of US\$82.

⁵The time variable is the number of days between first usage and last usage. A breakage event is the last observed usage if that usage happened more than 3 months before the end of our dataset.

⁶Kenya’s electricity tariff has a fixed component as well as a floating component that is pegged to international fuel prices and changes monthly. There were no changes in the fixed component over the study duration. The total price per kWh for customers consuming on average more than 100 kWh per month fluctuated between US\$0.2129 and US\$0.2219 over the eight month study duration.

Buying the stove on credit Subject to a small number of eligibility criteria (such as presentation of a national identification card and registration of a next of kin), anyone who wanted to buy the stove in installments was eligible to do so. The company uses many of the sophisticated credit technologies that have been developed in recent years to help address widespread credit market failures. Customers can choose whether to pay using monthly, weekly or even daily deadlines depending on what would best facilitate their repayment. In the lead-up to each payment deadline, participants who have not yet paid receive SMS reminders. Delinquent accounts are manually checked and followed-up with by a dedicated agent team (BURN Manufacturing, 2023). If the customer does not make the required payment within seven days, the device is remotely turned off and only turned back on once the owner makes the payment (a ‘digital collateral’ practice that is standard in East Africa; see Gertler, Green, and Wolfram, 2024). If after 14 days the owner has still not made the required payment, the agent visits the respondent in the field to either restructure the loan (if non-payment is due to unforeseen circumstances) or, if needed, to repossess the device. To cover the cost of these sophisticated procedures, buyers paying in installments must pay back 175% of the amount they borrowed. For those paying on a monthly plan, an initial down payment of US\$15 would be followed by six monthly payments of US\$19 each for a total price of US\$132. Respondents could repay their loans early if they wished, but this would not reduce the total payment amount. The payment structure corresponds to an effective annualized percentage rate (APR) of 224%, which is 215% above Kenya’s central bank rate of 9.75% during the experiment (Central Bank of Kenya, 2026).⁷

4.3 Implementation

Study enumerators recruited 2,511 people into the study residing in the urban and peri-urban areas around the cities of Naivasha, Gilgil, and Nakuru in Nakuru county in Kenya. Respondents were enrolled by enumerators walking around the study areas. To be eligible for study participation, participants had to spend at least half of their cooking fuel expenditures on charcoal for a standard Kenyan stove, shown on the left in Figure 4. They also had to be at least 18 years of age, have a pre-paid Kenya Power electricity meter in their homes, have a household size of at least 3 individuals, have basic literacy skills, spend at least US\$4 per week on charcoal, and not currently own an electric or energy efficient charcoal stove.

Respondents were visited up to four times. During the baseline survey enumerators completed a short survey consisting of demographic and socioeconomic questions. The main survey was conducted around 50 days after the baseline visit. During the main survey the enumerators implemented the random treatments, provided (randomized) information about the induction stove, and elicited willingness-to-pay (WTP). All respondents were given the opportunity to buy the induction stove shown in Figure 4 at this time, which also included three pieces of induction cookware (two pots and one frying pan). An endline survey was conducted 40 days after the main survey visit (at-

⁷For comparison, in 2025 U.S. auto loans had an average interest rate of 6.8% for new vehicles and 11.5% for used vehicles (Experian, 2025).

Table 1: Summary statistics from respondent surveys at baseline

	Mean	SD	25 th	50 th	75 th	N
Number of residents	5.0	1.7	4	5	6	2134
Age	35.8	10.1	28	35	42	2134
Male (=1)	0.1	0.3	0	0	0	2134
Charcoal cooking time (minutes/day)	186.0	71.1	140	180	220	2134
Number of lpg stoves	0.6	0.6	0	1	1	2126
Number of wood stoves	0.1	0.4	0	0	0	2134
Charcoal spending (USD/month)	19.4	6.2	14	16	23	2134
LPG/ethanol spending (USD/month)	5.2	10.2	0	7	8	2128
Daily charcoal usage (KG per day)	3.1	1.0	2	3	4	2134
Annual emissions from charcoal usage (tCO _{2e})	4.9	1.5	3	4	6	2134
Respondent income (USD/month)	101.5	115.3	46	69	116	2127
Household income (USD/month)	246.9	204.8	116	196	300	1815
Monthly rent paid by household (30d; USD)	24.4	16.6	15	23	31	1797
Max willing to pay if phone broke (USD)	32.9	38.9	8	12	46	2134
Hours of power outage (30d)	13.9	27.5	1	4	12	2134
Hours of voltage fluctuation (30d)	2.0	10.9	0	0	1	2134

Notes: Mean, standard deviation, and 25th, 50th, 75th for key socio-economic variables collected during the baseline survey.

trition is balanced by random treatment assignment; stove adopters are 2.2 percentage points less likely to attritl see [Table C2](#)). Finally, enumerators visited respondents around 180 days after the initial sensor deployment to download temperature data. [Figure A1](#) presents a more detailed study timeline.


The sections below describe the randomized treatments and the data collection in more detail. [Table 1](#) presents summary statistics collected during the baseline survey for participants who also subsequently participated in the main study.

4.4 Randomized treatments

Between the baseline and main surveys, each respondent is assigned to several randomized treatments, as shown in [Figure 5](#). All treatments are fully cross-randomized and stratified on each other. All treatment assignments are furthermore stratified by baseline LPG ownership, baseline electricity expenditure, and baseline charcoal expenditure. Treatment assignment is balanced on observable baseline characteristics ([Table C1](#)).

Fixed cost subsidy Participants were cross-randomized into different subsidy amounts for the fixed cost of buying the induction stove bundle (including the three pots, as discussed above), which had a marginal production cost of US\$82 at the time of the experiment. 49% of participants received

Figure 5: Randomized treatments



Notes: Random assignment for the 2,134 study participants. [Section 4.4](#) discusses the randomized treatments in more detail.

a discount of 10% (for a price of US\$73) and 46% received a discount of 75% (for a price of US\$20).⁸

Marginal cost subsidy Participants in the usage subsidy treatment groups receive a subsidy on the cost of the electricity that is used to operate the induction stove. During the main visit, prior to the WTP elicitation, respondents in the treatment groups are told that they would receive this subsidy, and the level they were assigned. They are informed that only electricity used by the electric stove is subsidized; they do not receive any subsidies on electricity used for other purposes. One-third of participants do not receive a marginal cost subsidy, one-third receive a subsidy of 25%, and one-third receive a subsidy of 75%. For a sense of magnitude, given ex post realized stove usage rates of on average 12 kWh per month at a tariff of US\$0.166, total marginal cost subsidy payouts averaged US\$3 and US\$10 for respondents in the 25% and 75% subsidy groups, respectively.

All respondents have pre-paid electricity meters, which is the norm in Kenya and across many African utilities.⁹ Consuming electricity runs down the meter: a customer must periodically ‘top up’ their meter to avoid it reaching zero (at which point electricity gets shut off).¹⁰ They can do so by purchasing a unique 20-digit token code either at a kiosk or via the Kenya Power mobile application, and then manually entering this code on their meter. We implement electricity transfers through this same channel: the research team collects respondents’ account numbers to purchase electricity tokens on their behalf, and then sends customers the 20-digit token code via SMS, which they then enter manually into their meter. The code is unique to the meter and cannot be transferred, sold, or returned. We process electric stove usage data, calculate subsidy amounts, and send respondents token codes on a weekly basis. For comparison, the median respondent in our sample purchases

⁸The remainder of participants had discounts of 46% (0.2% of participants), 50% (2.5%), or 69% (1.8%). All prices had positive probability of being selected from the feasible range of US\$20 (since a down payment of \$12 was required by the stove firm) and US\$73 (for ethical reasons all study participants received a discount of at least 10%).

⁹87% of all Kenyan residential electricity customers have pre-paid meters, and 95% of all residential customers who were connected since 2015—which includes most low-income connected households—do.

¹⁰For 59% of our sample this had not happened in the past month. 31% it had happened once or twice in the past month.

electricity four times per month, or once every 7.5 days, suggesting our implementation matches respondents’ own purchasing habits.

On the one hand, the manual entry of tokens could add an additional hurdle that lowers the value of marginal cost subsidies. On the other hand, it could make the transfer more salient than an automatic bill rebate. Rather than trying to bound these opposing forces, we only note that the utility recently used this same method of token transfers to implement a subsidy pilot study.¹¹ Our implementation methodology thus reflects the realistic implementation of such a program by the electric utility and the empirical estimates can be interpreted as policy relevant.

Payment structure Participants in the ‘pay up front’ group had to pay the full (randomly assigned) price during the main visit. Participants in the ‘loan’ group had to pay a down payment of US\$12 during the main visit.¹² One half of participants were randomly assigned to pay up front and one half to pay with a loan.

For those paying with a loan, starting one month after the purchase date ($t = 31$) the amount due would increase by the daily installment amount for the 90 subsequent days. Respondents could pay early but this would not reduce the total amount owed. If a respondent did not meet the amount due by a particular date, the stove would shut off until they make another daily payment, at which the new amount due stream is calibrated to the new payment date.

Information treatment All respondents receive an information sheet that shows what meals can be cooked with a charcoal stove and an information sheet that shows what meals can be cooked with an electric stove (Figure A2). Participants in the information treatment group receive two information sheets about the cost of the electricity used to cook various meals on an induction stove, as well as the cost of the charcoal used to cook those same meals on a charcoal stove. For participants in the usage subsidy groups, the information sheets show prices both with and without the subsidy amounts. 14% of participants are assigned to the information control group and 86% are assigned to the information treatment group.

4.5 Data collection

During the baseline visit and the endline visit, enumerators conducted a socio-economic survey aimed primarily at measuring fuel expenditures. The survey specifically asked about the ownership of cookstoves with different fuel types, including charcoal, wood, electricity, and liquefied petroleum gas (LPG) or liquid bioethanol fuel. During the main visit, they measured beliefs and WTP.

During the baseline visit, enumerators also installed charcoal stove monitors that continuously collected temperature data throughout the study period. The continuous collection of induction

¹¹Kenya Power’s e-cooking tariff development study states that “Rebates are provided as electricity tokens sent via SMS to pre-paid customers” (2024).

¹²On the first two days of surveying, September 23 and 24, a down payment of \$20 was used due to a miscommunication with the study partner. We omit surveys collected on these days (which comprise 5% of the sample) from analyses concerning the impact of credit.

stove usage and loan repayment data began after the main visit. [Figure A1](#) presents a more detailed timeline of data collection activities.

Charcoal stove usage During the baseline visit, enumerators install temperature sensors to monitor usage of households’ charcoal stoves. The sensors record a temperature measurement once every two minutes and store this on a local memory card that we download once per month. For households who owned two or more charcoal stoves at baseline, we installed monitoring devices on all stoves that were less than three years old (older stoves than this often caused the device to malfunction). The temperature sensors are drilled into the stove and securely fastened without damaging the functionality of the stove. [Figure 4](#) shows a charcoal stove instrumented with a temperature sensor. The installation occurs during the baseline visit such that we are able to estimate a within-individual treatment effect. Prior research has found minimal Hawthorne effects from electronic stove monitoring;¹³ still, any Hawthorne effect resulting from the installation of stove sensors would be symmetric across treatment groups.

Willingness-to-pay We measure Willingness-to-pay (WTP) using the Becker, Degroot, and Marschak (1964) mechanism, building on the implementations developed in Berry, Fischer, and Guiteras (2020), Dean (2024), and Berkouwer and Dean (2022). Neither the respondent nor the enumerator know the respondent’s randomly assigned price. The enumerator conducts a binary search over the range of Ksh 1,500 to Ksh 18,000 (US\$12 to US\$139), first asking “*If the price of the ECOA is Ksh 9,792 [US\$76] would you want to buy it?*”, then proceeding to a higher or lower price based on the respondent’s answer, and repeating this process until arriving to the nearest Ksh 10 (US\$0.08). After arriving at a final WTP, the envelope is opened and the respondent buys the stove for the price in the envelope if and only if it is less than or equal to their WTP. Since the price is fixed ex ante, the mechanism is incentive compatible. Since upfront prices are randomly assigned, purchasing is random conditional on WTP within the support of random prices stove. Practice randomized BDM and take-it-or-leave-it (TIOLI) elicitation result in similar estimates.

Crucially, respondents with a WTP between the minimum and maximum of the support of random prices are the adopters for whom the instrumental variables (IV) approach estimates the local average treatment effect (LATE) of purchasing a stove. These same respondents are defined as marginal (or ‘additional’) adopters in the subsidy efficacy calculations, as they would buy the stove if and only if there was a subsidy. The fact that these groups are naturally equivalent allows us to conduct several relevant policy calculations.

Induction stove usage The electric stoves record power usage in kilowatt-hours (kWh) on a 15-minute interval basis. This is transmitted to remote servers in real-time via a continuous wireless connection. This enables the seller to disconnect the stove in case of non-payment, and—crucially

¹³For example, in a randomized trial of improved cookstoves in Rwanda, Thomas et al. (2016) randomly assigned respondents to have a monitor installed in either a visible or hidden part of the cookstove. The visible monitor increases usage by 6% in the first week but causes no difference in cooking behavior from week two onwards.

for our study—allows us to subsidize only electricity used for the induction stove. The household surveys also record self-reported measures of electricity spending, which we convert to usage using the marginal tariff relevant for most customers, which is US\$0.166 per kWh.

Market price survey We collect a panel dataset of charcoal prices by conducting market surveys over the course of the experiment across all three cities. This allows us to convert household charcoal expenditure data to kilograms of charcoal purchases.

4.6 Converting energy usage to CO₂e emissions

Charcoal We assume a fraction of non-renewable biomass (fNRB) of 38% reflecting the most recent Modeling Fuelwood Savings Scenarios (MoFuSS) estimates for Kenya (Ghilardi and Bailis, 2024; UNFCCC, 2023). Still, our conclusions and results are relatively stable even when assuming an fNRB of 10% or 100% (Section 9.1).

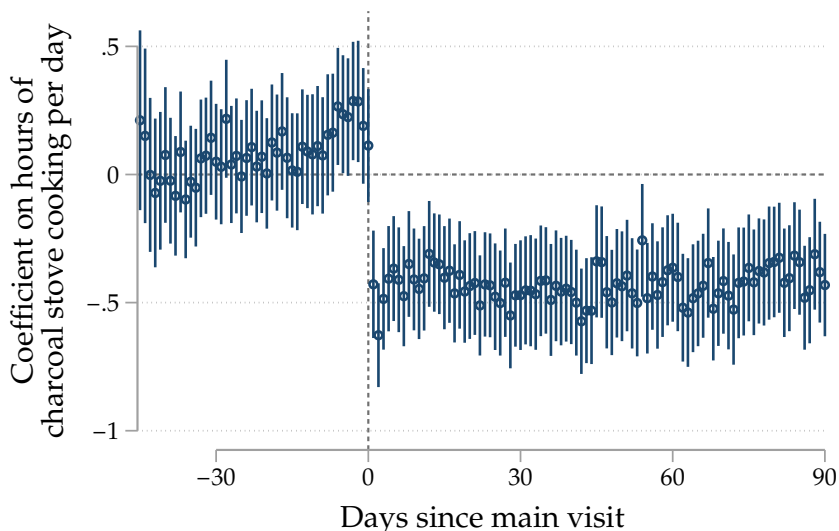
We use estimates of combustion emissions and non-combustion (production and processing) emissions of CO₂, CH₄ and N₂O from Floess et al. (2023). We convert this to total CO₂-equivalent emissions assuming a 50-year time horizon. This results in an emissions factor of 161 tCO₂e per terajoule, consisting of 36 tons of combustion CO₂e and 124 tons of non-combustion CO₂e (Gill-Wiehl, Kammen, and Haya, 2024; Whitman and Lehmann, 2011). With a calorific value of 0.027 terajoules per tonne of charcoal, this yields an estimated 4.3 kilograms of CO₂e per kilogram of charcoal.¹⁴ Sample average household usage of 3.1 kilograms of charcoal per day thus emits 4.9 tCO₂e per year.

Electricity Renewable sources of energy (primarily geothermal, hydropower, and wind) generate 90% of Kenya’s annual grid electricity (Kenya Power, 2023). While Kenya Power meets short-term demand increases with fossil fuels, almost all capacity expansions in recent years have been renewable. When considering the electrification of the cooking sector as a whole, Kenya’s average grid emissions are therefore a reasonably representative margin. To convert the remaining 10% of generation (primarily relatively dirty coal) to CO₂e emissions we conservatively assume an emissions factor of 1 tCO₂e per MWh of electricity, resulting in an average emissions intensity of 154 gCO₂e per kWh. We also report outcomes if Kenya’s grid had the average emissions intensity of the U.S. grid, which is approximately 366 grams of CO₂e per kWh (EIA, 2024).

Liquefied Petroleum Gas We convert LPG expenditures to cylinders purchased using local prices, and convert cylinders to CO₂e emissions using standard conversion ratios.

¹⁴Intuitively, the mass of oxygen is approximately 16 g/mol and the mass of carbon (which charcoal largely consists of) is approximately 12 g/mol. Each carbon particle combines with two oxygen particles to generate CO₂ such that 12 grams of charcoal generates 44 grams of CO₂, or 3.6 kilograms CO₂ per kilogram of charcoal. However, this does not factor in production, processing, and transportation; the purity of charcoal in East Africa (it does not consist exclusively of carbon); nor reforestation rates, and is thus a more naive approximation.

Figure 6: Charcoal stove usage from Stove Use Monitors



Notes: Coefficients from a fixed effects regression. Estimates from regressions with data collapsed to pre- and post-period differences in [Table A1](#). The outcome variable is hours per day that the respondents cooked with the charcoal stove, which is identified using the raw 2-minute temperature data.

Firewood 17% of households report ever using firewood to cook. Because we do not track the quantity of wood collected nor firewood prices, we refrain from estimating the CO₂e abatement resulting from a reduction in the use of firewood for cooking. As a result, the abatement and cost estimates are if anything conservative.

5 Private and social impacts of adoption

5.1 Charcoal stove usage

We convert the temperature data collected by the Stove Use Monitors (SUMs) to an hourly panel dataset of cooking time by respondent, spanning from 45 minutes before and 90 days after the main visit. We then estimate how induction stove adoption affects this measure of cooking time.

[Figure 6](#) estimates separate coefficients of the daily difference between adopters and non-adopters from 45 days before the main visit until 90 days after the main visit. Standard errors are clustered by respondent. The corresponding OLS regression (Column 1 of [Table A1](#)) shows a decrease of 26 minutes per day and instrumental variables (IV) estimate (Column 5 of [Table A1](#), which uses price, credit, and their interaction as instruments for adoption) shows a reduction of 31 minutes per day, or a 34% reduction in daily cooking time caused by adoption of the electric stove.

5.2 Energy expenditures

[Table 2](#) presents estimates of the causal effect of stove adoption on energy outcomes. The regressions use an instrumental variables approach using the randomly assigned stove subsidy and the

Table 2: Impact of induction stove purchase (instrumental variables)

(a) Energy usage

	Hours per day cooking with charcoal	Monthly expenditures (USD)					Total
		Charcoal	LPG	Wood	Electricity		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Bought induction stove (=1)	-0.51*** (0.16)	-10.03*** (0.82)	-0.83 (0.56)	-0.33 (0.21)	2.29*** (0.38)	2.17*** (0.66)	-8.90*** (1.06)
Observations	5296128	2078	2078	2078	2078	703	2078
Control Mean	1.51	15.11	5.51	0.75	3.51	0.33	24.88
F-Stat		241.22	241.22	241.22	241.22	45.32	241.22

(b) Emissions (tons of CO_{2e} per year)

	Charcoal	LPG	Electricity	Total
	(1)	(2)	(3)	(4)
Bought induction stove (=1)	-2.53*** (0.20)	-0.03 (0.02)	0.02*** (0.00)	-2.54*** (0.20)
Observations	2078	2078	2078	2078
Control Mean	3.77	0.18	0.04	3.99
F-Stat	121.98	121.98	121.98	121.98

Notes: Instrumental variables estimates use the randomized treatment price, the randomized credit condition, and their interaction as exogenous instruments. LPG is liquefied petroleum gas. Panel A shows the impact of induction stove purchasing on daily cooking and monthly energy expenditures. Column 1 uses hourly SUMS data (standard errors clustered by respondent and by date). Column 5 adds subsidies to electricity expenditure, Column 6 estimates among subsidy control group. Panel B shows the impact on emissions. [Section 4.6](#) describes the methodology for converting energy expenditures to emissions. ‘Control mean’ evaluated among non-adopters in the complier sample.

randomized credit treatment as instruments for adoption, which consistently yields an F -stat above 100.¹⁵

Panel A of [Table 2](#) shows a significant reduction in charcoal usage. Daily time spent cooking using a charcoal stove, as measured by the Stove Use Monitors (SUMs), decreases by 31 minutes per day (34%). Self-reported monthly charcoal expenditures, shown in Column 2, decrease by US\$10 (66%). Using the methodology described in [Section 4.6](#), this reduction in charcoal usage corresponds to emissions abatement of 2.5 tCO_{2e} per year.

It is worth emphasizing that this estimate is an average measure across all buyers that are compliers in the instrumental variables approach. Thus, it factors in behaviors like ‘stacking’ (where households continue to use their old stoves) as well as households choosing not to use their stoves. Without those behaviors, the estimated CO_{2e} reduction would likely have been higher.

Columns 5 and 6 show that self-reported monthly electricity expenditures increase by US\$2.3

¹⁵ Access to a loan increases willingness-to-pay (WTP) by US\$16 over a mean WTP of US\$20 among those paying up front ([Table A2](#)). Together with the randomized prices, this generates a strong first stage. Among respondents with access to a loan and a high subsidy 56% bought the induction stove, whereas among those paying up front and receiving a low subsidy fewer than 1% did.

(65%) per month.¹⁶ This is slightly lower than the induction stoves’ average consumption of 18.3 kWh per month, which equates to US\$3.0 per month at the Kenya Power tariff of US\$0.166 per kWh. The under-estimate could result from a range of sources, including for example that other household members may be responsible for purchasing household electricity tokens such that the household’s primary cook may not have an accurate sense. There are additional, smaller decreases in wood and LPG expenditure.

Aggregate monthly energy expenditures decrease by US\$8.9 (36%). Assuming a 10% annual discount rate, this corresponds to US\$184 over the 1.9 year lifetime of the stove. For buyers paying for the stove up front, the US\$82 price of the stove then generates a 112% annual return on investment (ROI) after 1.9 years. At a price of US\$132 (the total amount paid using the installment plan described in [Section 4.2](#)), the annual ROI is 64%.

It is worth noting that the commercially available payment plan has an APR of 224%: for most households, the monthly energy savings are thus not sufficient to cover the monthly repayment requirements, rendering adoption with credit prohibitively expensive despite the 112% rate of annual return on investment.

5.3 Emissions

We convert energy expenditures to CO_{2e} emissions using the methodology described in [Section 4.6](#). Panel B of [Table 2](#) shows an annual emissions reduction of 2.5 tCO_{2e}. This effectively entirely reflects the reduction in charcoal-related emissions, since the modest increase in electricity-related emissions is offset by a modest decrease in LPG-related emissions (we do not quantify fuelwood reductions).¹⁷ While electricity expenditures constitute approximately 36% of total household energy expenditure among those who bought the stove, they constitute less than 3% of their total household emissions.¹⁸ The increase in electricity usage thus generates a negligible increase in grid emissions of 0.02 tCO_{2e} per year. Induction stoves are highly efficient in converting energy to heat: cooking with an induction stove uses significantly less energy than cooking with charcoal.

Kenya’s grid is furthermore 90% renewable. Even when assuming U.S. average grid emissions, the emissions increase from electricity usage would only increase to 0.046 tCO_{2e} per year, still well below the decrease in charcoal-related emissions.¹⁹

¹⁶To address the confounding effect of the electricity subsidies, Column 5 defines total electricity spending as self-reported electricity spending plus subsidy receipts. Column 6 estimates self-reported electricity spending among the subsidy control group.

¹⁷Due to uncertainty in its origin, we refrain from converting the reduction in firewood usage to CO_{2e} emissions: the aggregate emissions reduction is thus a lower bound.

¹⁸These numbers exclude transportation-related emissions or other emissions generated outside the home.

¹⁹As discussed in [Section 4.6](#), average U.S. grid emissions are 366 gCO_{2e}/kWh as opposed to Kenya’s 154 gCO_{2e}/kWh. As a back of the envelope calculation: if people consume around 0.5 kWh per day, or 183 kWh per year, then at 366 gCO_{2e}/kWh this would constitute 67 kgCO_{2e} or 0.07 tCO_{2e}.

5.4 Other impacts

We see a 10% decrease in the fraction of people who cook ugali, mukimo, or matoke (three common starch staples) and an 11% increase in the fraction of people who cook rice, with no other meaningful changes in which foods people cook, in line with the fact that rice is relatively easier to cook on an electric stove than other starches. We see statistically significant reductions in self-reported health outcomes among respondents and their children, but these are not correlated with WTP and do not translate into significant economic impacts such as hours worked or school attendance. While respondents almost never use a charcoal stove exclusively for space heating, this is a co-benefit of charcoal cooking and could thus reduce the welfare gain from electric stove adoption (Berkouwer and Dean, 2022).

These outcomes are important and valuable from the perspective of poverty reduction and welfare improvement. However, they should not affect the cost of subsidies per ton of CO₂e abated other than by moving WTP and usage directly (in which case our measurements of WTP and usage already incorporate them). They matter for a social planner whose decisions incorporate welfare gains such as these, however even then they remain difficult to quantify. We therefore refrain from trying to quantify them. Berkouwer and Dean (2026b) presents additional detail on the drivers and impacts of electric stove adoption that relate to health, socioeconomic status, and the electrical grid.

6 Green subsidy efficacy

Without additional access to a loan, marginal buyers of induction stoves in our context abate 2.6 tCO₂e per year, or 5 tCO₂e over the 1.9 year lifetime of the stove (Table 5).²⁰ This was causally estimated using the IV compliers, who have a willingness-to-pay (WTP) between the low price $p_l = \text{US}\$20$ and the high price $p_h = \text{US}\$73$ and thus correspond exactly to adopters who are marginal to a subsidy that lowers the price from p_h to p_l . To estimate the subsidy cost per additional adoption we leverage the demand curve elicited through the BDM exercise (panel A of Figure 8). The cost per ton of CO₂e abated can then be calculated by combining these numbers.

We offer four different costs per ton, all of which might be relevant to a company, policy-maker, or organization depending on their goals: subsidy cost per ton of CO₂e, resource cost per ton of CO₂e, marginal value of public funds, and abatement per stove sold.

We first evaluate the subsidy cost of incentivizing a marginal adoption. At baseline, the high price is the competitive market price $p_h = p_{mc}$. A subsidy s_1 lowers this price to $p_l = p_{mc} - s_1$. Denote ρ_l and ρ_h to be the fraction of individuals who would buy the stove at the low or high price, respectively. The subsidy cost required to incentivize one additional adoption can then be denoted

²⁰Section 9.1 quantifies subsidy costs under alternative assumptions, such as a 1-year or 3-year average lifetime.

as follows:

$$\mathbb{C}_s = \frac{\text{Additional subsidy cost}}{\text{Additional adoptions}} = \frac{\rho_l s_l - \rho_h s_h}{\rho_l - \rho_h} \quad (6)$$

In the context of this study, 40% of respondents would buy the stove at the low price whereas 2.6% of respondents would buy the stove at the high price. Applying Equation 6, it costs the principal US\$65 in subsidy expenditures to incentivize one additional adoption.

Subsidy cost per ton of CO₂e Consider an environmental principal who wishes to maximize CO₂e abatement while minimizing subsidy expenditure (referred to as the ‘government cost per ton’ in Hahn et al., 2026). This could be a budget-constrained government agency or a cost-minimizing company or organization that wishes to meet an abatement target.

Given the estimated abatement per additional adoption $\hat{\phi}$, the additional subsidy expenditure required to abate one additional ton of CO₂e is then given by:

$$\text{Subsidy cost per tCO}_2e = \frac{\text{Subsidy cost per additional adoption}}{\text{Abatement in tCO}_2e \text{ per adoption}} = \frac{\mathbb{C}_s}{\hat{\phi}} \quad (7)$$

Given that each additional adoption abates 5 tCO₂e, subsidies for induction stoves in this context abate greenhouse gas emissions at a cost of US\$13 per tCO₂e. This is a relatively low abatement cost compared to other technologies (Section 9 discusses the policy implications of this finding in more detail).

Marginal Value of Public Funds We evaluate the efficacy of green subsidies from the perspective of a social principal using the Hendren and Sprung-Keyser (2020) Marginal Value of Public Funds (MVPF) approach. The private benefit includes both the marginal fuel savings as well as the direct cash transfer, which is valued at US\$1. Defining the fraction of inframarginal adopters $\mathbb{I} = \frac{\rho_h}{\rho_h + \rho_l}$, the MVPF is then given by:

$$\begin{aligned} \text{MVPF} &= \frac{\text{Benefit of inframarginal adopters} + \text{Benefit of marginal adopters}}{\text{Subsidy cost per additional adoption}} \\ &= \frac{\mathbb{I}\mathbb{C}_s + (1 - \mathbb{I})(\hat{b} - p_{mc} + \mathbb{C}_s + \phi)}{\mathbb{C}_s} \\ &= 1 + \frac{1 - \mathbb{I}}{\mathbb{C}_s}(\hat{b} + \phi - p_{mc}) \end{aligned} \quad (8)$$

Each marginal adoption costs the principal \mathbb{C}_s in subsidy expenditure and generates net energy savings of US\$200, abates 5 tCO₂e, and incurs US\$82 in stove cost. This yields an MVPF of US\$2.8 when excluding environmental benefits and an MVPF of either US\$12, US\$20, or US\$29 per US\$1 of spending when using a Social Cost of Carbon (SCC) of US\$120, US\$230, or US\$340, respectively.²¹

²¹These SCC numbers reflect Environmental Protection Agency estimates using a 2.5%, 2%, or 1.5% discount rate, respectively (EPA, 2023).

Resource cost per ton of CO₂e A social planner values the economic cost of producing each stove p_{mc} and fuel savings $\hat{b} = \pi'\hat{\theta}\hat{d}_2$ but disregards inframarginal transfers. The resource cost of abatement can then be defined as:

$$\text{Resource cost per tCO}_2e = \frac{\text{Resource cost per adoption}}{\text{Abatement in tCO}_2e \text{ per adoption}} = \frac{p_{mc} - \hat{b}}{\phi} \quad (9)$$

Column 8 of [Table 5](#) (multiplied by 23 months to reflect the stove’s 1.9 year expected lifespan, discounted 10% per year) shows US\$200 in total energy savings. Since this exceeds the US\$82 marginal cost of the stove, the resource cost per ton is *negative*: each cookstove will *save* on average US\$128 in economic resources in this context. Abating one ton of CO₂e thus *reduces* total resource costs by US\$24.

Abatement per stove sold Finally, we consider a principal who wishes to contract on the total number of stoves sold. Under the high subsidy, 6.5% of stoves that are sold are not additional and would have been sold even with the smaller subsidy. Thus, while adoption of an additional stove causes abatement of 2.6 tons of CO₂e per year, only 94% of sales are additional to a subsidy. Any single stove sale therefore abates 2.5 tCO₂e per year in expectation, or 4.6 tCO₂e if we assume a 1.9 year stove lifetime. This would be the relevant number to use for example when the principal contracts carbon credit issuances on aggregate sales.

These abatement costs are significantly lower than many other green technologies. Panel A of [Table 3](#) summarizes the abatement cost estimated in this paper as well as in other papers evaluating energy efficient technologies in contexts where adopters likely face significant demand distortions. Panel B lists estimates of green household technologies receiving subsidies in the U.S. By and large, subsidies for green technologies in low- and middle-income contexts generate larger social benefits per subsidy dollar spent than similar subsidies for green household technologies in high-income contexts.

These differences are partly driven by technological differences. Each US\$82 stove abates on average 2.5 tCO₂e per year: for comparison, a US\$40,000 electric vehicle abates between 0.1–1.5 tCO₂e per year relative to the counterfactual (Muehlegger and Rapson, 2023). However, even conditioning on technology, the channels identified in this paper—differences in the marginal adopter and the demand elasticity—drive increases in subsidy efficiency. [Section 7](#) first disentangles how the two channels contribute to energy efficiency. [Section 8](#) then estimates the model presented in [Section 3](#) to disentangle how much demand distortions per se reduce abatement costs relative to estimate from contexts with fewer demand distortions.

7 Effect of demand distortions on subsidy efficacy

One of this paper’s central findings is that the relatively low abatement costs discussed above are caused at least in part by demand distortions. [Section 5](#) showed that the stove’s average discounted

Table 3: Efficiency of selected green subsidies

Technology subsidized	MVPF (US\$)	Subsidy spend per tCO _{2e} abated (US\$)	Source
<i>(a) Low- and middle-income contexts</i>			
Electric stoves	20	13	<i>This paper</i>
Solar lanterns	13	10	Rom, Günther, and Pomeranz (2023)
Energy efficient motors	5.1	^a	Chaurey et al. (2025)
Improved charcoal stoves	323	6	Berkouwer and Dean (2026a), Hahn et al. (2026)
<i>(b) High-income contexts</i>			
Electric vehicles	1.4	1,356	Hahn et al. (2026)
Electric vehicles	1.3	1,202	Allcott et al. (2026)
Wind production credits	5.9	46	Hahn et al. (2026)
Residential solar	3.9	90	Hahn et al. (2026)
Appliance rebates	1.2	474	Hahn et al. (2026)
Vehicle retirement	1.0	876	Hahn et al. (2026)
Hybrid vehicles	1.0	5,940	Hahn et al. (2026)
Weatherization	1.0	779	Hahn et al. (2026)

Notes: The Marginal Value of Public Funds (MVPF) is the total private and social value generated per US\$1 subsidy (Hendren and Sprung-Keyser, 2020; Hahn et al., 2026). ^aChaurey et al. (2025) estimate the impact of disseminating energy efficient motors for free to non-adopters.

fuel savings are US\$184 after 1.9 years. Respondents paying up front have an average WTP of US\$20. Defining willingness to pay as the distortion times the private benefit ($wtp = \gamma_c \cdot b$), this suggests a demand distortion of $\gamma_0 = 0.11$. This also corresponds to an implicit discount factor of 0.887 per week, or 0.002 per year. In other words, to rationalize WTP with discounting alone, respondents would need to be indifferent between receiving US\$1 today or receiving US\$499 one year from today.

Access to credit in this context increases WTP by 81% (Figure 9 shows the distributions of WTP by credit treatment status). Respondents given access to a loan have an average WTP of US\$35 suggesting a demand distortion of $\gamma_1 = 0.19$. How does this difference in demand distortion affect subsidy efficacy?

Table 4 presents abatement cost and MVPF estimates for respondents whose demand was more distorted and less distorted. Among respondents given access to a loan to buy the induction stove, who face a lower demand distortion, fixed cost subsidies decrease emissions at a cost of US\$22 per tCO_{2e}, which is 70% higher than the US\$13 per tCO_{2e} abatement cost among those with a higher demand distortion. Similarly, the MVPF is 69% higher, at US\$20 when using an SCC of US\$230 per tCO_{2e}.

Resource cost differences are exacerbated by the fact that estimated private energy savings are

Table 4: Impact of demand distortions on subsidy efficacy through the two channels

	Less distorted ($\bar{\gamma} = 0.19$)	More distorted ($\bar{\gamma} = 0.11$)	Difference
<i>Channel 1: Marginal adopter</i>			
Fuel savings per marginal adoption (\hat{b})	US\$154	US\$200	+30%
tCO ₂ e abated per marginal adoption ($\hat{\phi}$)	4.2 tCO ₂ e	5 tCO ₂ e	+19%
<i>Channel 2: Subsidy cost</i>			
Extensive margin elasticity to fixed subsidy	$\mathcal{E}_d = -0.9$	$\mathcal{E}_d = -2.3$	
Fraction of subsidy expenditures inframarginal	38%	6.5%	
Subsidy expenditure per marginal adoption	US\$93	US\$65	-30%
<i>Aggregate difference</i>			
Welfare gain per \$1 spending (MVPF)	US\$12	US\$20	+69%
Subsidy spending per tCO ₂ e abated	US\$22	US\$13	-41%

Notes: How demand distortions affect subsidy efficacy. ‘Less distorted’ are respondents with access to a loan; ‘more distorted’ are respondents who pay the full price up front. [Section 3.1](#) describes the theoretical underpinnings of channels 1 and 2. [Section 7.1](#) presents empirical estimates for channel 1. [Section 7.2](#) presents empirical estimates for channel 2.

US\$154 for adopters with access to credit, which is 23% lower than average energy savings among the credit control group. In aggregate, CO₂e abatement saves 34% more in economic resources among those with the largest distortions.

WTP among respondents who are given access to a loan is still well below what we might call an efficient level. It is reasonable to think that our short intervention did not fully alleviate all demand distortions experienced by households in this context. In addition, given that average monthly savings were US\$8.9, a risk neutral agent facing complete savings constraints might only be able to pay US\$74 over a six month period (when also including the US\$12 down payment), which is still well below the total discounted savings of US\$184 after 1.9 years. The model and counterfactuals in [Section 8](#) aim to quantify outcomes in alternative scenarios where the demand distortion is alleviated further.

Next, we document how the two channels hypothesized in theory contribute to this aggregate empirical result. [Section 7.1](#) quantifies in more detail how demand distortions affect the externality generated by the marginal adopter (channel 1). [Section 7.2](#) quantifies how they affect demand elasticity and subsidy efficacy (channel 2). [Section 7.3](#) compares the results with marginal cost subsidies.

7.1 Channel 1: Fuel reduction of marginal adopters

We first test the impact of demand distortions on the average externality of the marginal adopter. Throughout, we define a marginal adopter as having a WTP between the low price and the high price, so that they would adopt the stove if and only if there was a subsidy lowering the price from

Table 5: Abatement among marginal adopters by credit condition

	Charcoal				Electricity		Total			
	(1) USD	(2) USD	(3) CO _{2e}	(4) CO _{2e}	(5) USD	(6) USD	(7) USD	(8) USD	(9) CO _{2e}	(10) CO _{2e}
Credit (=1)	-1.0 (1.4)	-1.0 (1.4)	-0.2 (0.4)	-0.2 (0.4)	0.1 (0.8)	0.1 (0.8)	-0.7 (2.0)	-0.7 (2.0)	-0.2 (0.4)	-0.2 (0.4)
Purchased stove (=1)	-10.5*** (0.6)		-2.6*** (0.2)		2.3*** (0.3)		-9.7*** (0.9)		-2.6*** (0.2)	
Purchased X No credit		-10.5*** (0.6)		-2.6*** (0.2)		2.3*** (0.3)		-9.7*** (0.9)		-2.6*** (0.2)
Purchased X Credit	1.6* (0.9)	-8.9*** (0.6)	0.4* (0.2)	-2.2*** (0.2)	-0.1 (0.5)	2.2*** (0.3)	2.3* (1.2)	-7.4*** (0.8)	0.4* (0.2)	-2.2*** (0.2)
Observations	753	753	753	753	753	753	753	753	753	753
Control Mean	15.7	15.7	3.9	3.9	3.6	3.6	25.5	25.5	4.1	4.1

Notes: Instrumental variables regressions using the randomly assigned price as an instrument for adoption to estimate its impact on spending and emissions among marginal adopters (i.e. $WTP \in (20, 73)$ US\$). IV compliers correspond to marginal adopters. Odd columns test whether the marginal adopters' average externality differs by credit condition. Even columns estimate the positive externality by credit condition.

the high price down to the low price (such that $p_h - s_1 < wtp < p_h$).

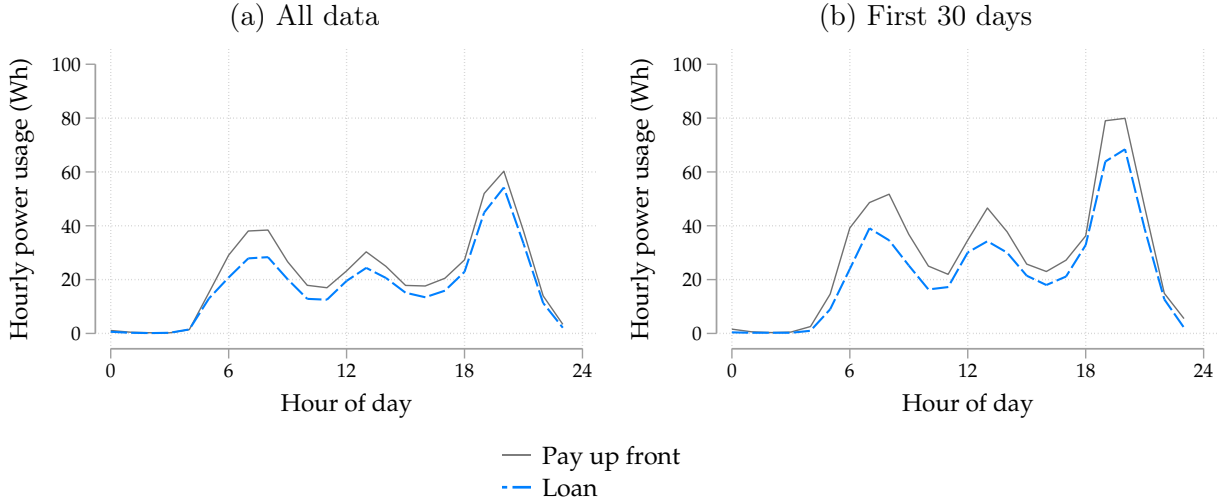
In the context of the model, when agents face larger demand distortions, agents who are marginal to a subsidy should unambiguously have higher average private benefits. Columns 7 and 8 of [Table 5](#) offer a test of this. Marginal adopters with a larger demand distortion generate on average US\$200 in discounted private savings over the course of 1.9 years of ownership, whereas adopters whose demand distortion was reduced through access to credit generate US\$154, a statistically significant difference.

Since social and private benefits are positively correlated in this context, this translates into a similar difference in positive externalities. Columns 9 and 10 of [Table 5](#) show that this is the case. Marginal adopters in the credit control group abate 5 tCO_{2e} after 1.9 years whereas marginal adopters in the credit treatment group abate 4.2 tCO_{2e}. In other words, the credit constraints in this context increase the abatement caused by marginal adopters by 19%.

[Figure 7](#) plots induction stove usage over the course of the day among adopters who were marginal to the subsidy, by credit treatment group. Marginal adopters in the credit control group have an average usage of 599 Wh per day, which is 34% higher than marginal adopters in the credit treatment group, who have an average daily usage of 445 Wh per day ([Table A3](#) shows the estimates corresponding to panel A).

This difference could reflect a treatment effect rather than selection. For example, having to pay back a loan leaves less money available for electricity expenditures. However, panel B shows that these effects begin within the first month, when adopters paying up front would have just faced a large direct liquidity shock from adoption. If anything, this would bias the estimates toward finding less usage among adopters paying up front. The difference in usage across the two groups is therefore likely to be a result purely of selection rather than a treatment effect. Either way, the

Figure 7: Electric stove usage among marginal adopters by credit condition



Notes: Hourly usage in Watt-hours by randomized credit condition among marginal adopters, defined as having a $p_{mc} - s < wtp_i < p_{mc}$. Table A3 presents point estimates, which are significant with $p < 0.01$. The IV F -stat is 3,151 for all regressions.

interpretation of this wedge does not affect the subsidy efficacy calculations.

7.2 Channel 2: Subsidy expenditure per marginal adoption

The second channel through which demand distortions increase subsidy efficacy is by increasing demand elasticity, which both increases the number of agents who are marginal to the subsidy, and also reduces the number of inframarginal payments. Figure 8 presents intuition on how these effects operate in the context of the demand curves elicited through the BDM. Area A as a fraction of areas $A \cup B \cup C$ denotes the fraction of additional subsidy expenditures disbursed to inframarginal buyers. The width of area $B \cup C$ represents the number of buyers that are marginal to a subsidy increase. Comparing panels A and B, demand distortions in this context both decrease payouts to inframarginal payments and increase the number of marginal buyers.

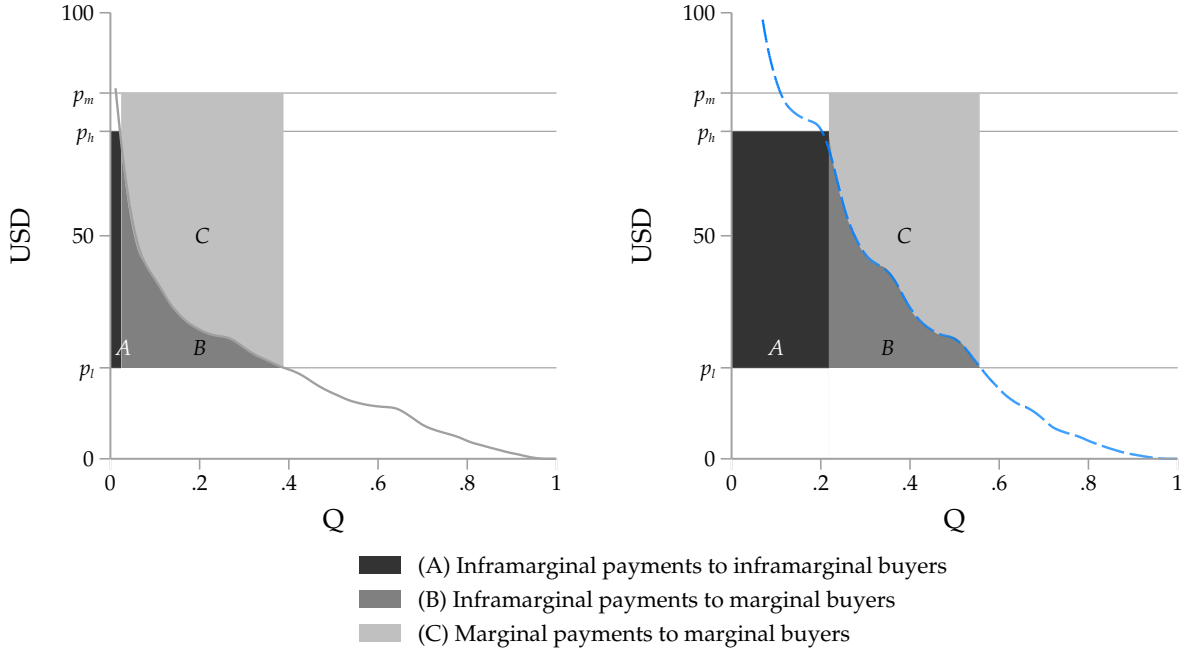
Intuitively, the cost increase resulting from a subsidy increase can consist of subsidy expenditures towards both inframarginal and marginal buyers. In other words, the numerator in Equation 6 can be disaggregated as follows:

$$\rho_l s_l - \rho_h s_h = \underbrace{\rho_h (s_l - s_h)}_{\text{Increased payouts to inframarginal buyers}} + \underbrace{s_l (\rho_l - \rho_h)}_{\text{Full payouts to marginal buyers}} \quad (10)$$

In terms of inframarginal adopters, 2.6% of people in the credit control group would buy the stove at the high price p_h whereas 21% of people in the credit treatment group would. As a result, among the credit treatment group 38% of subsidy expenditures flow to inframarginal buyers whereas in the credit control group only 6.5% of expenditures do.

The subsidy also increases sales by 35 percentage points among respondents in the credit treat-

Figure 8: Marginal and inframarginal expenditures by credit condition
(a) More distorted demand (b) With access to a loan



Notes: Demand curves constructed using WTP values elicited through the BDM mechanism. Demand distortions increase the proportion of marginal buyers and decreases the proportion of inframarginal buyers that result from a subsidy decreasing the price from p_h to p_l .

ment group, but by 37 percentage points among the credit control group. The same subsidy thus has a slightly larger increase in adoption rates among more credit constrained subsidy recipients.

Taken together, the environmental principal incurs a total subsidy cost of US\$93 per additional sale among respondents with access to a loan, whereas respondents in the credit control group require additional subsidy spending of US\$65 per additional sale. In other words, relaxing demand distortions increases the subsidy cost of achieving one additional sale by 43%.

These impacts can be summarized by the change in the demand elasticity, which as a percentage term reflects both marginal adopters (in the numerator) and inframarginal adopters (in the denominator). Table 6 estimates the price-quantity coefficients. Column 5 shows a demand elasticity of -2.3 among participants required to pay the entire price up front and Column 6 shows -0.9 among people paying with a loan. An alternative, policy relevant measure of demand elasticity would be to take p_h as the policy baseline and then to estimate the percentage quantity and price change from this starting point. This results in a demand elasticity among the loan group of $\mathcal{E}_D = -2.5$ and among the upfront payment group of $\mathcal{E}_D = -30.8$.

Columns 2 and 3 indicate that the increased elasticity includes an increase in the numerator: there are more marginal adopters per US\$1 subsidy. Given that there are also fewer inframarginal adopters (the denominator in the elasticity calculation), Column 5 shows respondents with a larger demand distortion have an even larger demand elasticity.

Table 6: Extensive margin demand elasticities by credit treatment group

	Quantity (%)			Log(Quantity)		
	(1)	(2)	(3)	(4)	(5)	(6)
Price (USD)	-0.85*** (0.01)	-0.96*** (0.02)	-0.80*** (0.01)			
Log(Price)				-1.23*** (0.01)	-2.33*** (0.02)	-0.88*** (0.01)
Observations	766	388	378	766	388	378
Sample		C=0	C=1		C=0	C=1

Notes: We estimate elasticities over the policy relevant sample of compliers, i.e. anyone with willingness to pay between p_l and p_h . Column 1 show that a US\$1 increase in price causes a 0.8 percentage point decrease in the quantity demanded. Columns 2 and 3 show that this decrease is steeper for respondents in the control group than in the treatment group. The lower rate of inframarginal adoption exacerbates the difference in elasticities in Columns 5 and 6, as the small denominator among the credit control group causes a larger demand elasticity.

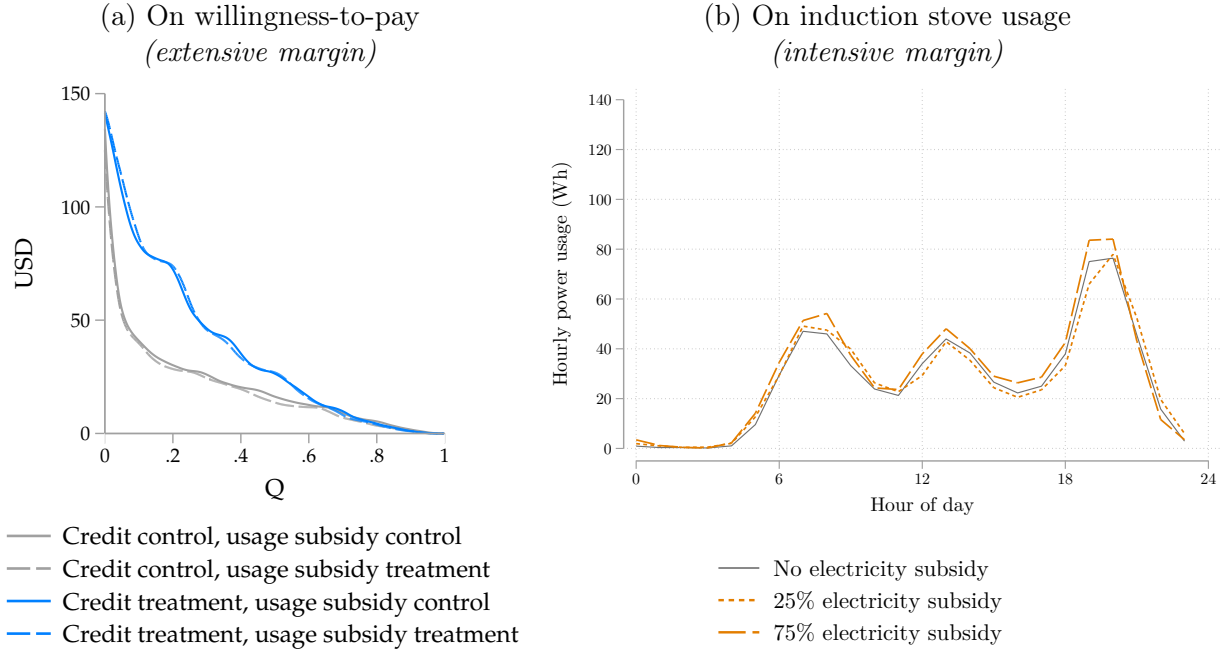
7.3 Marginal cost subsidies

The absolute efficacy of fixed cost subsidies does not inform their relative efficacy compared with marginal cost subsidies, which could offer two advantages over a fixed cost subsidy of equivalent net present value. First, marginal cost subsidies could encourage selection into adoption by higher users. Second, they could incentivize switching conditional on adoption. Both channels could increase average usage levels conditional on adoption, increasing the efficiency of marginal cost subsidies in terms of CO₂e abated per dollar of subsidy expenditure relative to an equivalent fixed cost subsidy.

On the other hand, if respondents are myopic or imperfectly attentive to electricity expenditures, they might undervalue electricity subsidies that accrue in the future. This could undermine the efficacy of marginal cost subsidies. The relative efficacy of these two subsidy instruments is thus an empirical question.

We evaluate the impact of the randomly assigned electricity subsidies, which lowered the cost of cooking by 25% (one-third of the sample) or 75% (one-third of the sample). Since we announced the electricity subsidy prior to eliciting WTP, the announcement of the usage subsidy could have affected usage both through a direct effect as well as a selection channel. [Figure 9](#) therefore shows the impact of usage subsidies on both the extensive margin (purchasing the stove) and the intensive margin (using the stove). Panel A shows WTP by whether the respondent was given access to a loan and whether they received a marginal cost subsidies. The electricity subsidy has no statistically detectable effect on average WTP. We can rule out that each US\$1 in electricity subsidy has more than a US\$0.16 impact on WTP, with a point estimate of 0.01 (Column 3 of [Table A4](#)). Pooled regressions and separate regressions by subsidy level all show no statistically significant impact on WTP ([Table A2](#)). The interaction term between credit and subsidies is positive, but this effect is not significant and very noisy so we do not report it: electricity subsidies may have a positive impact on WTP when respondent’s credit constraints are relaxed, or credit causes selection by people who

Figure 9: Impacts of marginal cost subsidies



Notes: Panel A shows the impact of access to credit and electricity subsidies on demand, with demand curves constructed using WTP values elicited through the Becker-DeGroot-Marschak mechanism (point estimates in Table A2). In Panel B, electric stove usage in Watt-hours recorded by the induction stoves, separated by whether the respondent was in the electricity subsidy control group, the 25% electricity subsidy group, or the 75% electricity subsidy group (point estimates in Table A4).

are slightly more price elastic, but we are unable to cleanly evaluate whether this is the case in this context.

Panel B of Figure 9 presents average hourly induction stoves usage by electricity subsidy treatment group. We estimate a demand elasticity of -0.06, and we can rule out that a 1% increase in price causes more than a 0.11% decrease in usage (Column 6 of Table A4). The point estimate for the 75% subsidy is twice as large as the point estimate for the 25% subsidy, but the point estimate on even the 75% subsidy increased stove usage is 10% and this is not statistically different from zero.

What causes the comprehensive unresponsiveness to electricity subsidies in this context? We explore several hypotheses. First, households may perceive a hassle cost in topping up a prepaid meter: as discussed in Section 4.4, recipients had to manually enter a unique code on their electricity meter to activate the subsidy each week. However, this is Kenya Power’s standard method for buying electricity, which households in our survey sample do on average once per week anyway—the same frequency at which we transfer electricity subsidies. Alternatively, cooking habits may simply be very inelastic to price.

8 Model estimation and counterfactuals

The goal of the model is to guide an understanding of how two parameters of particular interest—the demand distortion and the private-social benefit correlation—affect subsidy efficacy. We are agnostic about what drives this heterogeneity. For example, in another context a clean technology might generate a reduction in pollution exposure that increases WTP even in the absence of any financial savings: this would still cause a positive correlation between private and social benefits, and the model’s insights would still apply. Since the model is parameterized with only four parameters, other moments of the data fit less well, and we refrain from more sophisticated welfare analyses using these counterfactuals. That said, given the sensible fit of the seven targeted moments with only four parameters, we think the model can be informative in understanding how demand distortions affect subsidy efficacy.

We estimate the stylized model specified in [Section 3](#). We assume agents i are distributed inverse Fréchet by their inelastic demand for energy services $\theta_i \sim \text{Weibull}(\alpha_1, \alpha_2)$. Private benefits are defined $b_i = \hat{b} \cdot \theta_i$. To allow variation in the correlation between social benefit and private willingness-to-pay we define $\phi_i = \hat{\phi} [\alpha_3 \theta_i + (1 - \alpha_3) \epsilon_i]$, with $\hat{\phi} = 4.8 \text{ tCO}_2e$ (from panel B, Column 4 of [Table 2](#)), $\alpha_3 \in [0, 1]$, and ϵ_i an independent draw from the same Weibull distribution.

We assume a heterogeneous demand distortion distributed $\gamma_i \sim \text{Beta}(\alpha_4, \beta_c)$, where α_4 is a model parameter to be estimated and β_c is pinned down by $E[\gamma_i] = \hat{\gamma}_c$. For each credit condition $c \in \{0, 1\}$ the mean demand distortion γ_c is defined by $wtp_c = \gamma_c \cdot \hat{b}$, with discounted private benefits after 1.9 years $\hat{b} = \text{US\$184}$ (from panel A, Column 7 of [Table 2](#)) and observed mean willingness-to-pay $wtp_0 = \text{US\$20}$ and $wtp_1 = \text{US\$35}$ (from [Table A2](#)), which results in $\hat{\gamma}_0 = 0.11$ and $\hat{\gamma}_1 = 0.19$. Individual willingness-to-pay is then given by $wtp_i = \gamma_i \cdot b_i$

The distributions of wtp_i and ϕ_i can be used to calculate demand elasticity ϵ^D and the average externality of adopters who are marginal to the subsidy, which we define:

$$\bar{\phi} = \mathbb{E}[\phi_i | p_l < wtp_i < p_h]$$

Where $p_l = \text{US\$20}$ and $p_h = \text{US\$73}$ as in the experiment.

We use method of simulated moments to target seven empirical moments: extensive margin elasticity ϵ_c^D , average marginal externality $\bar{\phi}_c$, and cost per marginal adoption—all of which vary by credit groups $c = \{0, 1\}$ —as well as a common $\rho = \text{corr}(b_i, \phi_i)$. For each moment n we define the simulated moment error $e_n(\alpha)$ as a normalized function of the empirical moment $m_n(x)$ and the simulated moment $\hat{m}_n(\tilde{x})$:

$$e_n(\alpha) = \frac{\hat{m}_n(\tilde{x}) - m_n(x)}{m_n(x)}$$

We estimate the model parameters $\alpha = \{\alpha_1, \alpha_2, \alpha_3, \alpha_4\}$ using 2-step Generalized Method of Moments (2GMM), defining the estimator:

$$\hat{\alpha}_{2GMM} = \alpha : \min_{\alpha} e(\alpha)^T e(\alpha)$$

Table 7: Simulated estimates from observed and counterfactual parameters

(a) Estimation

Moment	RCT: Pay up front ($\gamma = 0.11$)			RCT: Pay with loan ($\gamma = 0.19$)		
	Empirical	Simulated	Deviation	Empirical	Simulated	Deviation
Correlation (ρ)	0.8	0.8	-3%	0.8	0.8	-3%
Marginal externality (ϕ)	5	4.8	-3%	4.2	4.4	5%
Extensive margin elasticity (\mathcal{E}_D)	-2.3	-2.2	7%	-0.9	-1.4	-55%
USD per marginal adoption	65.4	65.4	0%	93.4	73.5	-21%

(b) Counterfactual simulations: Varying the distortion γ

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Assumed demand distortion (γ)	0.05	0.11	0.19	0.25	0.5	0.75	1
Marginal externality (ϕ)	5.5	4.8	4.4	4.2	3.6	3.2	2.5
Extensive margin elasticity (\mathcal{E}_D)	-3.7	-2.2	-1.4	-1.1	-0.5	-0.3	-0.2
USD per marginal adoption	62.2	65.6	73.7	80	117.1	178.6	306.3
USD per tCO ₂ e	11.4	13.7	16.9	19.2	32.4	56.6	121.5
MVPF per US\$1	25.4	21	17	15.1	9.2	5.6	3.1

Notes: Panel A shows reasonable congruency between empirical moments and the simulated moments under the estimated model parameters. Panel B uses the model parameters to estimate key moments, as well as aggregate abatement cost, for counterfactual scenarios with higher or lower demand distortions, holding constant the correlation between private benefits and social benefits to match the empirical moment. [Figure 10](#) and [Table A5](#) also vary the correlation parameter ρ .

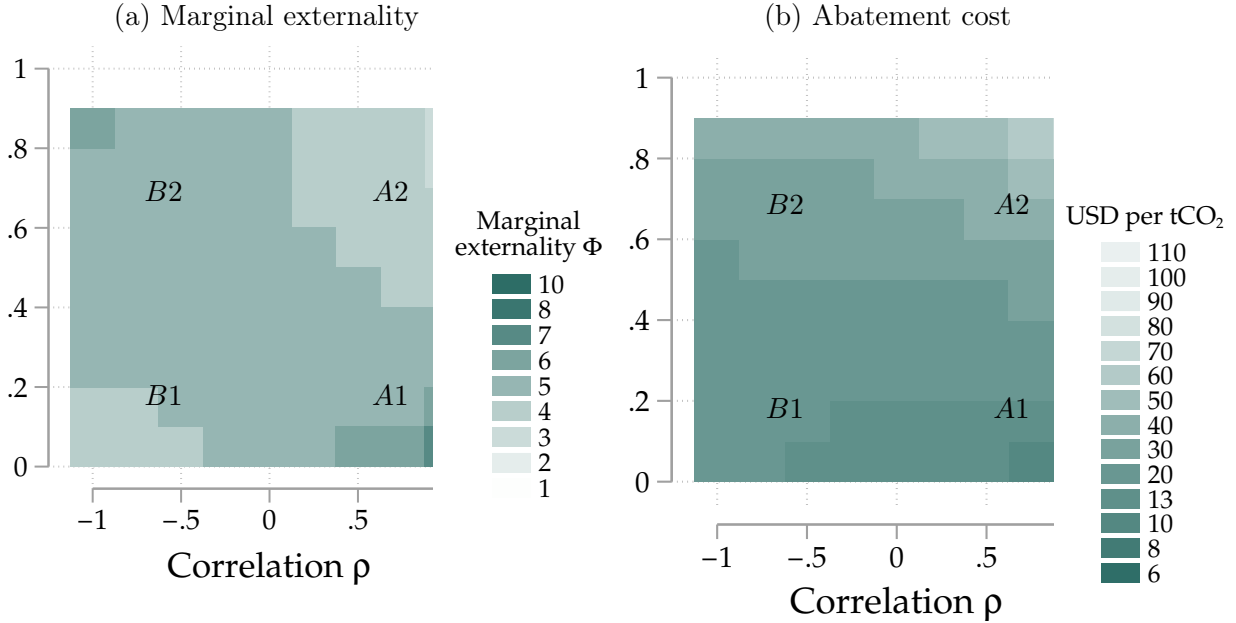
The estimated coefficients are $\hat{\alpha}_1 = 1.8$, $\hat{\alpha}_2 = 0.9$, $\hat{\alpha}_3 = 0.6$, and $\hat{\alpha}_4 = 1.2$. Panel A of [Table 7](#) shows the congruency between the moments simulated with the estimated parameters and the empirical moments calculated using the data.

Using the simulated moments to estimate abatement costs yields US\$13 per tCO₂e and US\$16 per tCO₂e for the credit control and treatment groups, respectively (shown in panel B). Despite these not being targeted moments, these aggregate efficiency numbers reasonably approximate the empirically estimated abatement costs of US\$13 and US\$22, respectively.

8.1 Counterfactual simulations

We run counterfactuals by simulating $\theta_i \sim \text{Weibull}(\hat{\alpha}_1, \hat{\alpha}_2)$ and $\epsilon_i \sim \text{Weibull}(\hat{\alpha}_1, \hat{\alpha}_2)$. The first set of counterfactuals varies only α_4 (which pins down γ). The second set of counterfactuals varies both α_3 and α_4 to understand how a negative or positive correlation can cause a demand wedge to generate adverse selection or beneficial selection, respectively. The third set of counterfactuals varies the heterogeneity of γ across agents.

Figure 10: Counterfactual estimates



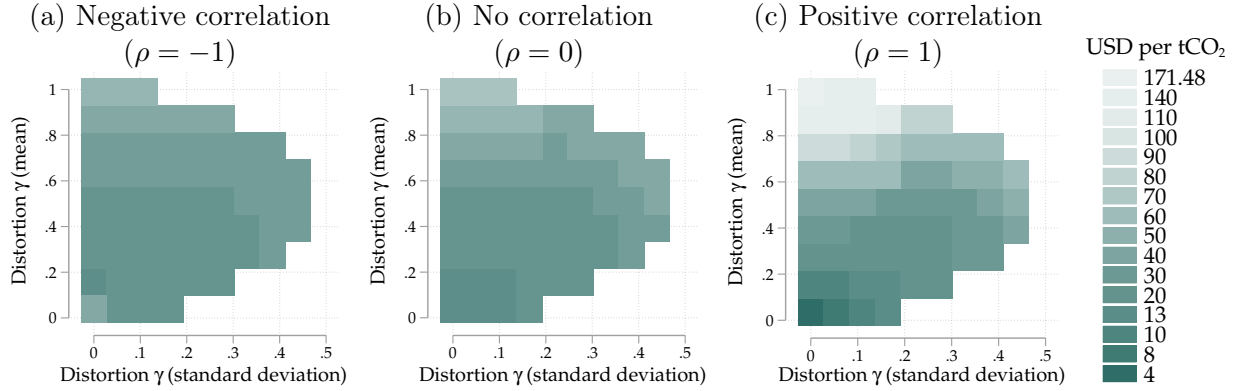
Notes: Both figures show how demand distortions as well as the correlation between willingness to pay and social benefits affect subsidy efficacy. Panel A shows the marginal externality. The labels A1, A2, B1, and B2 correspond to the labels in panels A and B of Figure 1, which presents the theoretical underpinnings for channels 1 and 2. Panel B shows the combined effect of the marginal externality and the elasticity on abatement cost. The demand distortion decreases abatement costs everywhere in this context, regardless of correlation. Table A5 presents a subset of these estimates in the format of Table 7.

Counterfactual I: Demand distortion Panel B of Table 7 shows the results implied by the model for different assumptions of γ . The estimated average marginal externalities follow the intuition offered by Figure 1: demand distortions that dampen demand increase the positive externality generated by the marginal adopter. The estimated extensive margin elasticities follow the intuition offered by Figure 2: a decrease in γ increases demand elasticity. Through these two channels, a demand distortion that dampens willingness-to-pay decreases the abatement cost in terms of US\$ per tCO₂e. In the neoclassical case where private benefits equal willingness-to-pay—corresponding to no demand distortion—the abatement cost increases to US\$122 per tCO₂e.

Counterfactual II: The correlation between willingness-to-pay and social benefits We extend the counterfactuals to include alternative correlations between willingness-to-pay and social benefits. Figure 10 shows the results when varying not only the demand distortion parameter γ but also the correlation parameter α_3 . To help build intuition, we use labels A1, A2, B1, and B2 to demarcate the approximate externalities labeled in panels A and B of Figure 1.

In the neoclassical framework, where willingness-to-pay reflects private benefits and there is no distortion ($\gamma = 1$, such that $wtp_i = b_i$), a positive correlation lowers the marginal externality and induces adverse selection, decreases the marginal positive externality ($A2 < B2$). On the other hand, under a demand distortion that dampens willingness-to-pay by a factor $\gamma < 1$, a positive correlation

Figure 11: Counterfactual estimates with heterogeneous demand distortions



Notes: This graph explores how different demand distortions and correlations affect abatement costs when demand distortions are allowed to be orthogonally heterogeneous across agents. Since we constrain $\gamma \in [0, 1]$, not all combinations of mean and standard deviation are feasible.

means compliers generate larger positive externalities, inducing beneficial selection ($A1 > B1$). The effect of a demand distortion thus hinges on the correlation between willingness-to-pay and social benefits: when this correlation is negative, a demand distortion decreases the average externality of compliers ($B1 < B2$). It is only when it is positive that the distortion increases the average externality of compliers ($A1 > A2$).

Panel B shows how these dynamics subsequently affect the subsidy expenditure cost per ton of CO₂e abated, factoring in the marginal externality shown in panel A as well as any changes in the demand elasticity caused by the demand distortion. Even when the correlation is negative and the demand distortion lowers the marginal externality, the demand distortion also simultaneously increases the extensive margin elasticity, regardless of the correlation between willingness-to-pay and social benefits. In this context, this effect outweighs the opposing effect on the average externality of compliers among negative correlation coefficients (such that even though $B1 < B2$ in panel A, we have $B1 > B2$ in panel B). As a result, the demand distortion unambiguously decreases abatement costs in this context, regardless of the correlation coefficient ($B1 > B2$ and also $A1 > B2$).

Counterfactual III: Heterogeneous demand distortions across agents Finally, we explore how heterogeneity in the demand distortion affects subsidy efficacy. For simplicity we assume the γ_i are uncorrelated with any of the other parameters of interest (willingness-to-pay and social benefits). We conduct the same counterfactual exercises as above, holding fixed the correlation to be either perfectly negative, zero, or perfectly positive. We draw random distributions of $\gamma_i \sim \text{Beta}(\beta_1, \beta_2)$, using β_1, β_2 to vary the mean and standard deviation of the γ_i while restricting $\gamma_i \in [0, 1]$.

Figure 11 shows the results. When the standard deviation is 0, the graphs follow the patterns shown in Panel B of Figure 10, with the mean gamma mattering more in the case of a positive correlation than under a negative correlation. All three panels show that this pattern weakens as the standard deviation increases.

9 Discussion

9.1 Results under alternative assumptions

The benchmark scenario evaluated in this paper assumes 1.9 years estimated stove durability and uses Kenya’s average grid emissions. Alternatively, assume that the stove’s lifespan is on average 1 year or 3 years. This changes the subsidy cost of abatement among respondents who were not also given access to a loan to US\$25 and US\$8, respectively (scenarios 1 and 2 of [Table A6](#)).

There is significant uncertainty around the fraction of non-renewable biomass (fNRB), which is the fraction of deforestation that is permanently removed (as opposed to forest that regrows naturally). Our primary estimates use the 38% rate that Ghilardi and Bailis ([2024](#)) estimate using their Modeling Fuelwood Savings Scenarios (MoFuSS) model, per the recommendations of the United Nations Article 6.4. To measure the sensitivity of our results to this parameter, we also calculate costs under 10% and 100% rates. This changes the subsidy cost per ton of CO_{2e} abated to US\$16 or US\$10, respectively (scenario 3 and 4 of [Table A6](#)). These differences are relatively minor because approximately half of charcoal-related emissions are incurred not during combustion but during production and processing, which are not affected by the assumed fNRB (Floess et al., [2023](#)).

Finally, Kenya’s electricity generation mix is significantly cleaner than most other countries’, at 90% renewable and fewer than half the emissions per kWh than the U.S. grid. If Kenya’s grid emissions were equivalent to the U.S. grid, the abatement cost would remain essentially unchanged (scenario 5 of [Table A6](#)). The difference with the benchmark scenario is minor because cooking with charcoal is significantly dirtier than cooking with electricity, such that the reduction in charcoal cooking emissions is only very slightly offset by the increase in electricity cooking emissions (as discussed in [Section 5.3](#)).

9.2 Scale

Approximately 1.4 million Kenyan households use charcoal on a daily basis. Our results suggest that a price subsidy could incentivize adoption by 37% of households.²² The induction stove usage data indicates that the 90th percentile of peak usage is on average 88 Wh per customer. Thus, scaling adoption in line with demand from our experiment would require an additional 45 MW in grid capacity. For context, Kenya’s current peak demand is 2,177 MW with an installed capacity of 3,243 MW (Kenya Power, [2024](#)). While the intermittency of renewable resources means the grid generally does not generate at installed capacity, Kenya’s electric utility appears capable of meeting any demand increase associated with widespread electric stove adoption.

Approximately 54 million households across the world today use charcoal as their primary cooking fuel, including approximately 31 million households that live in Africa ([Table A7](#)). At annual

²²This assumes universal electricity access among charcoal users, which is not completely unreasonable. Charcoal is used primarily in urban and peri-urban areas with access to markets. Urban electricity access increased from 50% in 2000 to 98% in 2022 (World Bank Data, [2025](#)). The rural access rate lags but now exceeds 66% and continues to increase rapidly in large part due to Kenya’s nationwide electrification programs.

abatement of 2.6 tCO_{2e} per year and with 37% adoption at subsidized prices, this would abate approximately 30 million tons of CO_{2e} per year.

In addition, more than 300 million households across the world still cook primarily with firewood. Households currently cooking primarily with firewood could generate additional abatement. And, the IEA (2025) estimates that the number of households cooking with wood or charcoal in Africa in 2050 will if anything be higher than the numbers today. A fruitful area of future research would be to better understand global cooking transition pathways given the heterogeneity of demand distortions across countries.

10 Conclusion

Demand distortions such as financial market frictions and behavioral biases are pervasive across many low- and middle-income countries, often operating by dampening willingness-to-pay (WTP) for profitable technologies. This paper shows that this has important implications for the efficiency of Pigouvian price instruments. In theory, fixed cost subsidies for green technologies generate a larger gain in consumer surplus as well as a larger positive externality when agents face demand distortions that dampen demand. This operates through two channels: increasing the WTP of the marginal adopter, and increasing demand elasticity.

We test this theory in a field experiment with more than 2,000 households in Nakuru County, Kenya, where we randomize marginal cost subsidies, fixed cost subsidies, and demand distortions for the purchase of an energy efficient technology: induction stoves. We first use the random variation in prices to estimate the causal impacts of stove adoption. We find that the stove abates 2.6 tCO_{2e} per year and reduces (discounted) household energy expenditures by US\$184 over the 1.9 year lifetime of the stove. WTP among the control group is US\$20: access to a loan increases WTP to US\$35. Even factoring in continued use of charcoal stoves and imperfect additionality of the subsidies, green subsidies in this context offer a very low-cost abatement opportunity, at only US\$13 per ton of CO_{2e} abated. This compares favorably to many other abatement technologies available today. Factoring in the private benefit to adopters, each dollar of subsidy generates US\$20 in total social welfare gain.

As predicted by the theory, demand distortions in this context contribute to these favorable numbers. When we relax demand distortions by randomly offering access to credit, the cost of abatement increases to US\$22 per ton of CO_{2e} and the MVPF decreases to US\$12 per subsidy dollar. This operates through the two channels identified in the model: demand distortions increase the marginal adopter’s positive externality by 19% and increase demand elasticity from -0.9 to -2.3.

The loan treatment increases WTP in this context but still does not close the adoption gap. To estimate abatement costs in the absence of any demand distortions, we estimate the model using 2-step Generalized Method of Moments (2GMM), using the random variation from the experiment to identify key parameters. Counterfactuals suggest that if WTP equaled the discounted stream of fuel savings the abatement cost would have increased more than tenfold, to US\$122 per tCO_{2e}, and

decreased the MVPF to US\$3.1 per subsidy dollar.

The mechanisms identified in this paper suggest that many of the most effective abatement opportunities may be found in contexts with larger demand distortions. As an example, the U.S. has relatively sophisticated financial markets for consumer durables, with the vast majority of new cars sold with financing. Our paper suggests that this could contribute to the high cost of abatement through electric vehicle subsidies, at more than US\$1,000 per ton.

Decarbonization through rapid electrification to slow climate change is motivating more than US\$1 trillion in investments towards climate change mitigation each year. Maximizing the impact of these expenditures requires allocating this funding towards abatement opportunities at the lowest portion of the abatement cost curve. Given the results of this study, many of these opportunities may be found in low- and middle-income countries, where demand distortions are pervasive. Additional research is needed to identify specific mitigation opportunities.

References

- Acemoglu, D., S. Johnson, and J. A. Robinson (2005). “Chapter 6 Institutions as a Fundamental Cause of Long-Run Growth”. In: ed. by P. Aghion and S. N. Durlauf. Vol. 1. *Handbook of Economic Growth*. Elsevier, pp. 385–472.
- Adhvaryu, A., N. Kala, and A. Nyshadham (2020). “The Light and the Heat: Productivity Co-Benefits of Energy-Saving Technology”. *The Review of Economics and Statistics* 102.4, pp. 779–792.
- Aemro, Y. B., P. Moura, and A. T. de Almeida (2021). “Experimental evaluation of electric clean cooking options for rural areas of developing countries”. *Sustainable Energy Technologies and Assessments* 43, p. 100954.
- Alem, Y., S. Hassen, and G. Köhlin (2014). “Adoption and disadoption of electric cookstoves in urban Ethiopia: Evidence from panel data”. *Resource and Energy Economics* 38, pp. 110–124.
- Allcott, H. (2011). “Consumers’ Perceptions and Misperceptions of Energy Costs”. *American Economic Review* 101.3, pp. 98–104.
- Allcott, H., R. Kane, M. S. Maydanchik, J. S. Shapiro, and F. Tintelnot (2026). “The Effects of “Buy American”: Electric Vehicles and the Inflation Reduction Act”. *Working Paper*. March 4, 2026.
- Allcott, H., S. Mullainathan, and D. Taubinsky (2012). “Energy Policy with Externalities and Internalities”. *National Bureau of Economic Research*. Working Paper Series. Working Paper #17977.
- (2014). “Energy Policy with Externalities and Internalities”. *Journal of Public Economics* 112, pp. 72–88.
- Allcott, H. and D. Taubinsky (2015). “Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market”. *American Economic Review* 105.8, pp. 2501–2538.
- Allcott, H. and N. Wozny (2014). “Gasoline Prices, Fuel Economy, and the Energy Paradox”. *The Review of Economics and Statistics* 96.5, pp. 779–795.
- Anagol, S. and C. Udry (2006). “The Return to Capital in Ghana”. *American Economic Review* 96.2, pp. 388–393.
- Aspelund, K. M. and A. Russo (2026). *Additionality and Asymmetric Information in Environmental Markets: Evidence from Conservation Auctions*. Tech. rep.
- Atkin, D., B. Bernadac, D. Donaldson, T. Garg, and F. Huneeus (2025). *The Incidence of Distortions*. Documento de Trabajo No. 1026. Working Paper Series. Banco Central de Chile.
- Bagnoli, M. and T. Bergstrom (2005). “Log-Concave Probability and Its Applications”. *Economic Theory* 26.2, pp. 445–469.
- Banerjee, A. V. and E. Duflo (2005). “Growth Theory through the Lens of Development Economics”. In: *Handbook of Economic Growth*. Ed. by P. Aghion and S. Durlauf. Vol. 1. Chap. 7, pp. 473–552.
- (2014). “Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program”. *The Review of Economic Studies* 81.2, pp. 572–607.
- Banerjee, M., R. Prasad, I. H. Rehman, and B. Gill (2016). “Induction stoves as an option for clean cooking in rural India”. *Energy Policy* 88.C, pp. 159–167.
- Bangladesh Bureau of Statistics (2023a). *Report on Socio-Economic and Demographic Survey*. Tech. rep.
- (2023b). *Statistical Yearbook*. Tech. rep.
- Bank, W. (2025). *Access to Electricity, Kenya*. Accessed: March 31, 2026.
- Becker, G. M., M. H. Degroot, and J. Marschak (1964). “Measuring utility by a single-response sequential method”. *Behavioral Science* 9.3, pp. 226–232.

- Bergquist, L. F., D. Lashkari, and E. Verhoogen (2026). *Wedges: A Microeconomic Perspective on Misallocation*. Working Paper 34756. National Bureau of Economic Research.
- Berkouwer, S. and J. Dean (2026a). “The cost-effectiveness of clean cookstove carbon mitigation after adjusting for additionality and impact”. *Working Paper*. Unpublished manuscript.
- Berkouwer, S. B. and J. T. Dean (2022). “Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households”. *American Economic Review* 112.10.
- (2026b). “Watts for Supper? The socioeconomic drivers and impacts of electric stove adoption in Kenya”. Working paper.
- Bernanke, B. S. (2023). “Nobel Lecture: Banking, Credit, and Economic Fluctuations”. *American Economic Review* 113.5, pp. 1143–69.
- Berry, J., G. Fischer, and R. P. Guiteras (2020). “Eliciting and Utilizing Willingness to Pay: Evidence from Field Trials in Northern Ghana”. *Journal of Political Economy* 128.4, pp. 1436–1473.
- Blattman, C., N. Fiala, and S. Martinez (2014). “Generating skilled self-employment in developing countries: Experimental evidence from Uganda”. *Quarterly Journal of Economics* 129.2, pp. 697–752.
- Boomhower, J. and L. W. Davis (2014). “A credible approach for measuring inframarginal participation in energy efficiency programs”. *Journal of Public Economics* 113, pp. 67–79.
- Borenstein, S. (2025). *Should the Price of Electricity Depend On What You Use It For?* URL: <https://energyathaas.wordpress.com/2025/03/10/should-the-price-of-electricity-depend-on-what-you-use-it-for/> (visited on 03/10/2025).
- Borenstein, S. and L. W. Davis (2025). “The Distributional Effects of US Tax Credits for Heat Pumps, Solar Panels, and Electric Vehicles”. *National Tax Journal* 78.1, pp. 263–288.
- Buchner, B. (2024). “COP29’s climate investment imperative”. *Science* 386.6722, pp. 601–601.
- Burkina Faso National Institute of Statistics and Demography and the DHS Program (2021). *Demographic and Health Survey*. Tech. rep.
- BURN Manufacturing (2023). *Burn Credit Policy*. Unpublished internal policy document. Version 1. Dated 23 May 2023.
- Bushnell, J., E. Muehlegger, and D. S. Rapson (2025). “Energy Prices and Electric Vehicle Adoption”. *Journal of Political Economy Microeconomics*.
- Busse, M. R., C. R. Knittel, and F. Zettelmeyer (2013). “Are Consumers Myopic? Evidence from New and Used Car Purchases”. *American Economic Review* 103.1, pp. 220–256.
- Campbell, A. G. (2025). *Choosing Between Increasing Subsidies or Lowering Rates for Electrification*. <https://energyathaas.wordpress.com/2025/11/17/choosing-between-increasing-subsidies-or-lowering-rates-for-electrification/>. Energy Institute Blog, accessed 2026-01-06.
- Carlton, D. W. and G. C. Loury (1980). “The Limitations of Pigouvian Taxes as a Long-Run Remedy for Externalities*”. *The Quarterly Journal of Economics* 95.3, pp. 559–566.
- Central Bank of Kenya (2026). *Central Bank Rate*. <https://www.centralbank.go.ke/rates/central-bank-rate/>. Accessed: 2026-01-06.
- Chaurey, R., G. Nayyar, S. Sharma, and E. Verhoogen (2025). “Social Learning among Urban Manufacturing Firms: Energy-Efficient Motors in Bangladesh”. *Working Paper*. September 18, 2025.
- Chen, Q., N. Ryan, and D. Xu (2025). *Firm Selection and Growth in Carbon Offset Markets: Evidence from the Clean Development Mechanism*. Working Paper 33636. National Bureau of Economic Research.
- China National Bureau of Statistics (2020). *China Population Census Yearbook 2020*. Tech. rep.
- Clements, W., K. Silwal, S. Pandit, J. Leary, B. Gautam, et al. (2020). “Unlocking Electric Cooking on Nepali Micro-Hydropower Mini-Grids”. *Energy for Sustainable Development* 57, pp. 119–131.

- Daghar, M. (2021). “Kenya’s charcoal bans have fuelled a smuggling problem”. *ISS Today*. Published under the ENACT project, funded by the European Union.
- De Groote, O. and F. Verboven (2019). “Subsidies and Time Discounting in New Technology Adoption: Evidence from Solar Photovoltaic Systems”. *American Economic Review* 109.6, pp. 2137–2172.
- Dean, J. T. (2024). “Noise, Cognitive Function, and Worker Productivity”. *American Economic Journal: Applied Economics* 16.4, pp. 322–60.
- DellaVigna, S. (2009). “Psychology and Economics: Evidence from the Field”. *Journal of Economic Literature* 47.2, pp. 315–372.
- Department of Finance Canada (2025). *Removing the consumer carbon price, effective April 1, 2025*. Accessed: 2026-03-26. URL: <https://www.canada.ca/en/department-finance/news/2025/03/removing-the-consumer-carbon-price-effective-april-1-2025.html> (visited on 03/22/2025).
- Desbureaux, S., L. Collart, N. Stoop, R. Soubeyran, M. Verpoorten, et al. (2025). *Subsidising Electric Cooking to Protect the Environment: Evidence from a Randomized Trial in the D.R. Congo*. Working paper.
- DRC Ministry of Public Health, DRC Ministry of Planning and Monitoring of the Implementation of the Revolution of Modernity, DRC Ministry of Public Health, and the DHS Program (2021). *Demographic and Health Survey*. Tech. rep.
- Duflo, E., M. Kremer, and J. Robinson (2008). “How High Are Rates of Return to Fertilizer? Evidence from Field Experiments in Kenya”. *American Economic Review* 98.2, pp. 482–488.
- Einav, L., A. Finkelstein, and M. R. Cullen (2010). “Estimating Welfare in Insurance Markets Using Variation in Prices*”. *The Quarterly Journal of Economics* 125.3, pp. 877–921.
- Ethiopian Public Health Institute, Ethiopia Federal Ministry of Health, and the DHS Program (2019). *Mini Demographic and Health Survey*. Tech. rep.
- Experian (2023). *State of the Automotive Finance Market: Q2 2023*. Quarterly Report. Accessed: 2026-03-24. Experian Information Solutions, Inc.
- (2025). *State of the Automotive Finance Market: Q2 2025*. Quarterly Report. Accessed: 2026-03-24. Experian Information Solutions, Inc.
- Fafchamps, M., D. McKenzie, S. Quinn, and C. Woodruff (2014). “Microenterprise growth and the flypaper effect: Evidence from a randomized experiment in Ghana”. *Journal of Development Economics* 106, pp. 211–226.
- Floess, E., A. Grieshop, E. Puzzolo, D. Pope, N. Leach, et al. (2023). “Scaling up gas and electric cooking in low- and middle-income countries: climate threat or mitigation strategy with co-benefits?” *Environmental Research Letters* 18.3, p. 034010.
- Fowlie, M. and R. Meeks (2021). “The Economics of Energy Efficiency in Developing Countries”. *Review of Environmental Economics and Policy* 15.2, pp. 238–260.
- Fowlie, M., M. Reguant, and S. P. Ryan (2016). “Market-Based Emissions Regulation and Industry Dynamics”. *Journal of Political Economy* 124.1.
- (2020). *In Defense of Alternatives to Pollution Pricing*. Policy Brief. Accessed: 2026-03-26. Economics for Inclusive Prosperity (EfIP).
- Gaarder, I., M. Grindaker, T. G. Meling, and M. Mogstad (2026). *Green Waste*. Working Paper 34649. National Bureau of Economic Research.
- Gerster, A. and M. Kramm (2024). “Optimal Internality Taxation of Product Attributes”. *American Economic Journal: Economic Policy* 16.3, pp. 394–419.
- Gertler, P., B. Green, and C. Wolfram (2024). “Digital Collateral”. *The Quarterly Journal of Economics* 139.3, pp. 1713–1766.
- Ghana Statistical Service and the DHS Program (2022). *Demographic and Health Survey*. Tech. rep.

- Ghatak, M. and D. Mookherjee (2025). “Misallocating Misallocation?” *Annual Review of Economics* 17, Volume 17, 2025, pp. 511–538.
- Ghilardi, A. and R. Bailis (2024). *Updated fNRB Values for Woodfuel Interventions*. Tech. rep.
- Gill-Wiehl, A., D. M. Kammen, and B. K. Haya (2024). “Pervasive over-crediting from cookstove offset methodologies”. *Nature Sustainability* 7.2, pp. 191–202.
- Gillingham, K., S. Houde, and A. Van Benthem (2021). “Consumer Myopia in Vehicle Purchases: Evidence from a Natural Experiment”. *American Economic Journal: Economic Policy* 13.3, pp. 207–238.
- Gillingham, K. and K. Palmer (2014). “Bridging the Energy Efficiency Gap: Policy Insights from Economic Theory and Empirical Evidence”. *Review of Environmental Economics and Policy* 8.1, pp. 18–38.
- Gillingham, K. and J. H. Stock (2018). “The Cost of Reducing Greenhouse Gas Emissions”. *Journal of Economic Perspectives* 32.4, pp. 53–72.
- Gould, C. F., M. L. Bejarano, B. De La Cuesta, D. W. Jack, S. B. Schlesinger, et al. (2023). “Climate and Health Benefits of a Transition from Gas to Electric Cooking”. *Proceedings of the National Academy of Sciences* 120.34, e2301061120.
- Goulder, L. H. and I. W. H. Parry (2008). “Instrument Choice in Environmental Policy”. *Review of Environmental Economics and Policy* 2.2, pp. 152–174.
- Grigolon, L., M. Reynaert, and F. Verboven (2018). “Consumer Valuation of Fuel Costs and Tax Policy: Evidence from the European Car Market”. *American Economic Journal: Economic Policy* 10.3, pp. 193–225.
- Hahn, R. W., N. Hendren, R. D. Metcalfe, and B. Sprung-Keyser (2026). *A Welfare Analysis of Policies Impacting Climate Change*. Tech. rep.
- Hausman, J. (1979). “Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables”. *The Bell Journal of Economics* 10.1, pp. 33–54.
- Hendren, N. and B. Sprung-Keyser (2020). “A Unified Welfare Analysis of Government Policies*”. *The Quarterly Journal of Economics* 135.3, pp. 1209–1318.
- Holland, S. P., E. T. Mansur, N. Z. Muller, and A. J. Yates (2016). “Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors”. *The American Economic Review* 106.12, pp. 3700–3729.
- Houde, S. and E. Myers (2025). “Heterogeneous (Mis)Perceptions of Energy Costs: Implications for Measurement and Policy Design”. *Journal of Political Economy Microeconomics*.
- Houde, S. and J. E. Aldy (2017). “Consumers’ Response to State Energy Efficient Appliance Rebate Programs”. *American Economic Journal: Economic Policy* 9.4, pp. 227–255.
- IMARC Group (2025). *East Africa LED Lighting Market: Industry Trends, Share, Size, Growth, Opportunity and Forecast 2025-2033*. Tech. rep.
- Indarte, S., R. Kluender, U. Malmendier, and M. Stepner (2026). *Consumption Wedges: Measuring and Diagnosing Distortions*. Working Paper 34891. National Bureau of Economic Research.
- International Energy Agency (2025). *Universal Access to Clean Cooking in Africa: Progress Update and Roadmap to Implementation*. Accessed: 2026-03-26. Paris: International Energy Agency.
- Ito, K. (2014). “Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing”. *American Economic Review* 104.2, pp. 537–563.
- Ito, K. and S. Zhang (2025). “Do Consumers Distinguish Fixed Cost from Variable Cost? “Schmeduling” in Two-Part Tariffs in Energy”. *American Economic Journal: Economic Policy* 17.2, pp. 194–223.
- Jacobsen, M. R., C. R. Knittel, J. M. Sallee, and A. A. van Benthem (2020). “The Use of Regression Statistics to Analyze Imperfect Pricing Policies”. *Journal of Political Economy* 128.5, pp. 1826–1876.

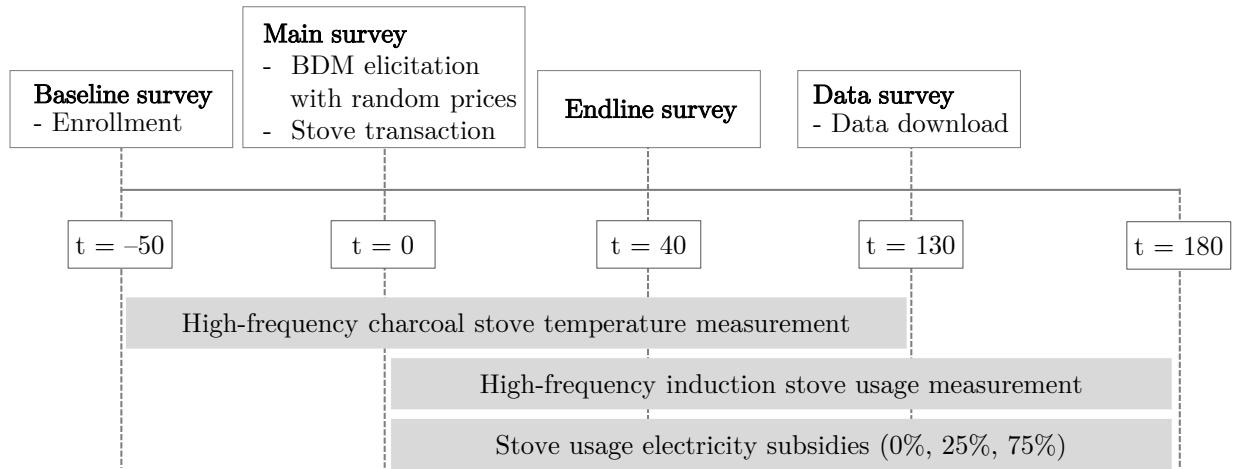
- Kahn, J. A. (1986). “Gasoline Prices and the Used Automobile Market: A Rational Expectations Asset Price Approach”. *Quarterly Journal of Economics* 101.2, pp. 323–339.
- Kahn, M. E. and F. A. Wolak (2013). *A Field Experiment to Assess the Impact of Information Provision on Household Electricity Consumption*. Tech. rep. 08-325. Prepared for the California Air Resources Board and the California Environmental Protection Agency.
- Kenya National Bureau of Statistics, Kenya Ministry of Health, and the DHS Program (2022). *Demographic and Health Survey*. Tech. rep. <https://dhsprogram.com/pubs/pdf/FR380/FR380.pdf>.
- Kenya Power (2023). *Kenya Power Annual Report and Financial Statements for the Year Ended 30 June 2023*.
- (2024). *Kenya Power Annual Report and Financial Statements for the Year Ended 30 June 2024*.
- Klier, T. and J. Linn (2010). “The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly Sales Data”. *American Economic Journal: Economic Policy* 2.3, pp. 134–153.
- Knittel, C. R. and R. Sandler (2018). “The Welfare Impact of Second-Best Uniform-Pigouvian Taxation: Evidence from Transportation”. *American Economic Journal: Economic Policy* 10.4, pp. 211–242.
- Kolstad, C. D. (2011). *Environmental Economics*. 2nd ed. New York: Oxford University Press.
- Kovacova, M. (2026). *Energy efficiency in Kenya: Light bulbs*. illuminem.
- Kremer, M., J. N. Lee, J. M. Robinson, and O. Rostapshova (2013). “The Return to Capital for Small Retailers in Kenya: Evidence from Inventories”. Mimeo, Harvard.
- Leary, J., B. Menyeh, V. Chapungu, and K. Troncoso (2021). “eCooking: Challenges and Opportunities from a Consumer Behaviour Perspective”. *Energies* 14.14, p. 4345.
- Lipsey, R. G. and K. Lancaster (1956). “The General Theory of Second Best”. *The Review of Economic Studies* 24.1, pp. 11–32.
- Madagascar National Institute of Statistics (2022a). *Permanent Household Survey*. Tech. rep.
- (2022b). *Thematic Report on the Results of RGPH-3 - Demographic Projections*. Tech. rep.
- Mahadevan, M., A. Martinez, R. McCord, R. Meeks, and M. Pradhananga (2025). *On the Back Burner: Experimental Evidence for Energy Transitions*. Tech. rep. Working paper.
- Martínez-Marquina, A. and M. Shi (2024). “The Opportunity Cost of Debt Aversion”. *American Economic Review* 114.4, pp. 1140–72.
- McKenzie, D. and C. Woodruff (2006). “Do Entry Costs Provide an Empirical Basis for Poverty Traps? Evidence from Mexican Microenterprises”. *Economic Development and Cultural Change* 55.1, pp. 3–42.
- (2008). “Experimental Evidence on Returns to Capital and Access to Finance in Mexico”. *World Bank Economic Review* 22.3, pp. 457–482.
- Mel, S. de, D. McKenzie, and C. Woodruff (2008). “Returns to Capital in Microenterprises: Evidence from a Field Experiment”. *The Quarterly Journal of Economics* 123.4, pp. 1329–1372.
- Ministry of Energy and Petroleum, Republic of Kenya (2024a). *Kenya National Electric Cooking Strategy (KNeCS)*. National Strategy Report. Accessed: 2026-01-28.
- (2024b). *Kenya National Electric Cooking Strategy: Action Plan*. Action Plan. Accessed: 2026-01-07. Kenya Power and Lightning Company Ltd.
- Mozambique National Institute of Statistics and the DHS Program (2023). *Demographic and Health Survey*. Tech. rep.
- Muehlegger, E. J. and D. S. Rapson (2023). “Correcting Estimates of Electric Vehicle Emissions Abatement: Implications for Climate Policy”. *Journal of the Association of Environmental and Resource Economists* 10.1, pp. 263–282.
- Myanmar Ministry of Planning and Finance (2017). *Living Conditions Survey*. Tech. rep.
- Myers, E. (2019). “Are Home Buyers Inattentive? Evidence from Capitalization of Energy Costs”. *American Economic Journal: Economic Policy* 11.2, pp. 165–188.

- Myers, E. (2020). “Asymmetric information in residential rental markets: Implications for the energy efficiency gap”. *Journal of Public Economics* 190, p. 104251.
- Myers, E., S. L. Puller, and J. West (2022). “Mandatory Energy Efficiency Disclosure in Housing Markets”. *American Economic Journal: Economic Policy* 14.4, pp. 453–487.
- Nigeria National Bureau of Statistics (2024). *General Household Survey - Panel Wave 5*. Tech. rep.
- Nunn, N. (2007). “Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade*”. *The Quarterly Journal of Economics* 122.2, pp. 569–600.
- OPEC Fund for International Development (2024). *Clean Cooking Report 2024*. Knowledge Series Report. Accessed: 2026-01-07. OPEC Fund for International Development.
- Pacific Gas and Electric Company (2025). *Electric Vehicles (EV) Rate Plans*. Accessed October 24, 2025. URL: <https://www.pge.com/en/account/rate-plans/electric-vehicles.html>.
- Palmer, K. and M. Walls (2015). “Can Benchmarking and Disclosure Laws Provide Incentives for Energy Efficiency Improvements in Buildings?” *Resources for the Future Discussion Paper* 15-09.
- Pande, R., R. Burgess, M. Farboodi, A. N. Keiller, and L. Page (2025). *Draft Proposal for a Unified Carbon Market*. Tech. rep. UNFCCC.
- Philippine Statistics Authority and the DHS Program (2022). *Demographic and Health Survey*. Tech. rep.
- Pigou, A. C. (1920). *The economics of welfare*. London: Macmillan.
- Prelec, D. and G. Loewenstein (1998). “The Red and the Black: Mental Accounting of Savings and Debt”. *Marketing Science* 17.1, pp. 4–28.
- Princewill, N., B. Feleke, and L. Madowo (2023). *Anger boils over as Kenya’s cost of living protests shake the nation*. Accessed: 2026-03-26. CNN. URL: <https://www.cnn.com/2023/07/20/africa/kenya-cost-of-living-protests-explainer-intl/index.html> (visited on 07/20/2023).
- Quinn, S. and C. Woodruff (2019). “Experiments and Entrepreneurship in Developing Countries”. *Annual Review of Economics* 11.1, pp. 225–248.
- Republic of Uganda (2022). *Updated Nationally Determined Contribution (NDC) — Republic of Uganda*. Nationally Determined Contribution. Submitted to the UNFCCC on 12 September 2022; Accessed: 2026-01-07.
- Rom, A., I. Günther, and D. Pomeranz (2023). “Decreasing Emissions by Increasing Energy Access? Evidence from a Randomized Field Experiment on Off-Grid Solar Lights”. *Working paper*.
- Rubinstein, F., B. H. Mbatchou Ngahane, M. Nilsson, M. Baame Esong, E. Betang, et al. (2022). “Adoption of electricity for clean cooking in Cameroon: A mixed-methods field evaluation of current cooking practices and scale-up potential”. *Energy for Sustainable Development* 71, pp. 118–131.
- Ryan, N. (2018). “Energy Productivity and Energy Demand: Experimental Evidence from Indian Manufacturing Plants.” *Working Paper*.
- Saha, A., M. A. Razzak, and M. R. Khan (2021). “Electric Cooking Diary in Bangladesh: Energy Requirement, Cost of Cooking Fuel, Prospects, and Challenges”. *Energies* 14.21, p. 6910.
- Sallee, J. M. (2014). “Rational Inattention and Energy Efficiency”. *Journal of Law and Economics* 57.3, pp. 781–820.
- (2025). “The Trouble with Green Subsidies”. *National Tax Journal* 78.1, pp. 171–199.
- Sallee, J. M., S. E. West, and W. Fan (2016). “Do consumers recognize the value of fuel economy? Evidence from used car prices and gasoline price fluctuations”. *Journal of Public Economics* 135, pp. 61–73.
- Segerson, K., S. Polasky, M. Scheffer, U. R. Sumaila, J. C. Cárdenas, et al. (2024). “A cautious approach to subsidies for environmental sustainability”. *Science* 386.6717, pp. 28–30.

- Sejas-Portillo, R., M. Moro, and T. Stowasser (2025). “The Simpler the Better? Threshold Effects of Energy Labels on Property Prices and Energy Efficiency Investments”. *American Economic Journal: Applied Economics* 17.1, pp. 41–89.
- Shapiro, J. S. (2025). “Institutions, Comparative Advantage, and the Environment”. *The Review of Economic Studies* 92.6, pp. 4152–4193.
- Sola, P. and P. O. Cerutti (2021). “Kenya has been trying to regulate the charcoal sector: why it’s not working”. *The Conversation*. Republished by World Agroforestry (ICRAF).
- Sussman, A. B. and E. Shafir (2011). “On Assets and Debt in the Psychology of Perceived Wealth”. *Psychological Science* 23.1, pp. 101–108.
- Tanzania National Bureau of Statistics (2022). *Basic Demographic and Socio-Economic Profile*. Tech. rep.
- Thomas, E. A., S. Tellez-Sanchez, C. Wick, M. Kirby, L. Zambrano, et al. (2016). “Behavioral Reactivity Associated With Electronic Monitoring of Environmental Health Interventions—A Cluster Randomized Trial with Water Filters and Cookstoves”. *Environmental Science and Technology* 50.7, pp. 3773–3780.
- U.S. Energy Information Administration (2024). *How much carbon dioxide is produced per kilowatt-hour of U.S. electricity generation?* Accessed: 2026-01-20. U.S. Energy Information Administration. URL: <https://www.eia.gov/tools/faqs/faq.php?id=74&t=11>.
- U.S. Environmental Protection Agency (2023). *Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances*. EPA Technical Report. Docket ID No. EPA-HQ-OAR-2021-0317. Washington, DC: U.S. Environmental Protection Agency.
- (2025). *Greenhouse Gas Emissions from a Typical Passenger Vehicle*. <https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle>. Last updated 12 June 2025. Accessed 7 April 2026.
- Uganda Bureau of Statistics (2020). *National Household Survey*. Tech. rep.
- Uganda Electricity Distribution Company Limited (UEDCL) (2025). *Customer Information Booklet 2025*. <https://www.uedcl.co.ug/wp-content/uploads/2025/07/Customer-Information-Booklet-2025.pdf>.
- United Nations Framework Convention on Climate Change (2023). *Information note: Development of default values for fraction of non-renewable biomass*. Tech. rep. CDM-MP92-A07.
- United Republic of Tanzania, Vice President’s Office (2021). *National Environmental Policy, 2021*. National Policy. Accessed: 2026-01-07. Vice President’s Office, United Republic of Tanzania.
- Volvo Cars (2025). *Volvo Cars launches free home charging initiative, starting in Sweden*. Press release, Volvo Cars Global Newsroom. ID: 353788.
- Wambui, S. (2010). *Kenya launches energy saving initiative*. Capital FM Business.
- Whitman, T. and J. Lehmann (2011). “Systematic under- and overestimation of GHG reductions in renewable biomass systems: A Letter”. *Climatic Change* 104.2, pp. 415–422.
- World Bank (2024). *Access to electricity (% of population)*. Accessed: 2024-10-12.
- World Bank Group (2020). “The State of Access to Modern Energy Cooking Services”.
- Zambia Statistics Agency, Zambia Ministry of Health, University Teaching Hospital Virology Laboratory, and the DHS Program (2018). *Demographic and Health Survey*. Tech. rep.

I Additional figures

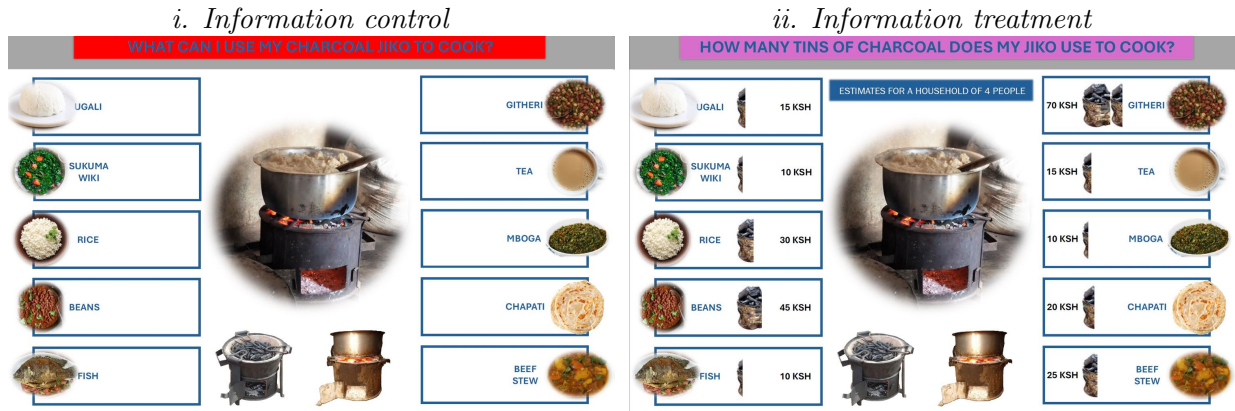
Figure A1: Experiment timeline




Notes: Timeline of study activities. The baseline survey was conducted on average 50 days before the main survey (10th, 90th: 40, 60 days), between August 12 and September 19, 2025. The main survey was conducted between 23 September and 19 November, 2025. The endline survey was conducted on average 40 days after the main survey ([25, 55]), between (November 10 and December 12, 2025). The data download was conducted on average 130 days after the main survey ([110, 150]), between February 18 and March 6, 2026. High frequency induction stove usage monitoring and electricity subsidy transfers were completed remotely.

Figure A2: Randomized information sheets

(a) Charcoal sheet



(b) Electricity sheet



Notes: Each participant receives one charcoal sheet and one electricity sheet. 14% of participants are in the information control group and receive sheets with no cooking cost information (Figure *i* in both panels). 86% of participants are in the information treatment group and receive cost information; respondents in the marginal cost subsidy and information treatment groups saw cooking costs that incorporated the price reduction (Figure *ii* in panel A and Figure *ii*, *iii*, or *iv* in panel B).

II Additional tables

Table A1: Effect of induction stove adoption on charcoal stove usage

	Ordinary Least Squares				Instrumental Variables			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pre-period difference	0.09 (0.08)		0.09 (0.08)		-0.14 (0.18)		-0.14 (0.18)	
Post-period difference	-0.43*** (0.07)		-0.43*** (0.07)		-0.51*** (0.16)		-0.51*** (0.16)	
Pre-post difference		-0.49*** (0.06)		-0.49*** (0.06)		-0.38*** (0.13)		-0.38*** (0.13)
Observations	5296128	5296128	220672	220672	5296128	5296128	220672	220672
Respondents	2007	2007	2007	2007	2007	2007	2007	2007
Control Mean	1.72	1.72	1.72	1.72	1.51	1.51	1.51	1.51
Level of data	Hourly	Hourly	Daily	Daily	Hourly	Hourly	Daily	Daily
Respondent FE?	No	Yes	No	Yes	No	Yes	No	Yes

Notes: Odd columns estimate the treatment-control difference in the pre-period (a baseline balance test) and in the post-period (a treatment effect estimate). Even columns estimate the treatment-control differences in the post period relative to the treatment-control difference in the pre-period, where the latter is absorbed by respondent fixed effects. All regressions cluster standard errors by respondent. All regressions include date and days-since-main-visit fixed effects. Columns 1, 2, 5, and 6 also include hour-of-day fixed effects. Column (1) uses hourly charcoal stove usage data and most closely represents a weighted average of the pre-period and post-period treatment effects shown in [Figure 6](#). The ‘post-period’ coefficient in Column (5) is our preferred specification for the causal impact of stove adoption on daily cooking time.

Table A2: Effect of credit and marginal cost subsidy on willingness-to-pay

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Credit treatment	14.1*** (2.1)	14.2*** (2.1)	14.7*** (1.9)	15.8*** (2.0)	14.1*** (2.1)	14.2*** (2.1)	14.7*** (1.9)	15.8*** (2.0)
Subsidy	-2.0 (1.3)	-2.1 (1.3)	-1.5 (1.2)	-1.5 (1.3)				
Subsidy=25					-2.2 (1.5)	-2.3 (1.4)	-0.8 (1.5)	-0.4 (1.5)
Subsidy=75					-1.7 (1.5)	-1.8 (1.4)	-2.3 (1.4)	-2.5* (1.4)
Credit X Subsidy (any)	2.4 (2.6)	2.3 (2.5)	2.2 (2.4)	1.6 (2.4)				
Credit X 25% Subsidy					0.8 (3.0)	0.9 (2.9)	0.3 (2.8)	-0.6 (2.9)
Credit X 75% Subsidy					4.0 (2.9)	3.8 (2.9)	4.0 (2.7)	3.7 (2.7)
Observations	2134	2134	2134	2134	2134	2134	2134	2134
Control Mean	19.5	19.5	19.5	19.5	19.5	19.5	19.5	19.5
SES controls	None	Some	Many	All	None	Some	Many	All

Notes: The effect of the credit treatment, the electricity subsidy treatment, and their interaction on willingness-to-pay. Columns 1, 2, 3, and 4 (as well as 5, 6, 7, and 8 respectively) vary only in using either no, some, many, or all socioeconomic controls.

Table A3: Impact of credit on daily induction stove usage

	All buyers		Marginal buyers	
	(1) Daily	(2) Respondent	(3) Daily	(4) Respondent
Credit (=1)	-102.9*** (29.8)	-102.8*** (32.2)	-80.9** (37.6)	-83.4* (43.5)
Observations	97952	625	55756	358
Respondents	625		358	
Control Mean	458	454	444	439

Notes: Results from ordinary least squares regressions estimating how credit constraints affect the composition of marginal adopters. Standard errors in parentheses are clustered by respondent. Regressions include date fixed effects and control for randomization strata and other treatments. Columns 1 and 3 use daily data, with standard errors clustered by respondent. Columns 2 and 4 use respondent daily average.

Table A4: Extensive and intensive margin impacts of electricity subsidy

	Willingness-to-pay			Daily induction stove usage		
	(1) USD	(2) USD	(3) USD	(4) Wh	(5) Wh	(6) Log(Wh)
Any subsidy	-0.41 (1.22)			11.66 (29.91)		
=25		-0.51 (1.44)			0.68 (35.21)	
=75		-0.31 (1.39)			23.62 (35.39)	
Subsidy value (USD)			-0.01 (0.08)			
Log(Marginal price)						-0.04 (0.05)
Observations	2134	2134	2134	97952	97952	40966
Respondents				625	625	613
Control mean	28	28	28	399	399	

Notes: Results from ordinary least squares regressions estimating how the electricity subsidy affects outcomes. Columns 1 and 4 pool the two subsidy levels of 25% and 75%; the remaining columns separate the levels. All columns control for randomization strata and other treatments. Columns 4-6 use data at the respondent-by-day level with date fixed effects and standard errors in parentheses clustered by respondent (results are qualitatively indistinguishable when averaging daily usage to the respondent level).

Table A5: Simulated estimates from observed and counterfactual parameters

	Assumed demand distortion (γ)	Assumed correlation (ρ)	Average marginal externality (ϕ)	Extensive margin elasticity (\mathcal{E}_D)	US\$ per tCO ₂ e	
(1) <i>Counterfactual</i>	.11	-1	13.8	-2.2	65.5	4.8
(2) <i>Counterfactual</i>	.11	-.7	10.4	-2.1	65.6	6.3
(3) <i>Counterfactual</i>	.11	-.6	11.1	-2.1	65.5	5.9
(4) <i>Counterfactual</i>	.11	-.2	12.6	-2.2	65.5	5.2
(5) <i>Counterfactual</i>	.11	0	4	-2.2	65.4	16.3
(6) <i>Counterfactual</i>	.11	.2	4.3	-2.2	65.5	15.2
(7) <i>Counterfactual</i>	.11	.6	4.6	-2.2	65.6	14.3
(8) <i>Counterfactual</i>	.11	.7	4.7	-2.2	65.6	13.9
(9) <i>Counterfactual</i>	.11	1	5.5	-2.2	65.4	11.9
(10) <i>Counterfactual</i>	.19	-1	14.8	-1.4	73.6	5
(11) <i>Counterfactual</i>	.19	-.7	10.9	-1.4	73.5	6.8
(12) <i>Counterfactual</i>	.19	-.6	11.5	-1.4	73.6	6.4
(13) <i>Counterfactual</i>	.19	-.2	12.7	-1.4	73.3	5.8
(14) <i>Counterfactual</i>	.19	0	4	-1.4	73.4	18.2
(15) <i>Counterfactual</i>	.19	.2	4.1	-1.4	73.4	17.7
(16) <i>Counterfactual</i>	.19	.6	4.3	-1.4	73.5	17.2
(17) <i>Counterfactual</i>	.19	.7	4.3	-1.4	73.7	17
(18) <i>Counterfactual</i>	.19	1	4.7	-1.4	73.5	15.8
(19) <i>Counterfactual</i>	.5	-1	16.4	-.5	116.6	7.1
(20) <i>Counterfactual</i>	.5	-.7	11.7	-.5	116.9	10
(21) <i>Counterfactual</i>	.5	-.6	12.2	-.5	116.8	9.6
(22) <i>Counterfactual</i>	.5	-.2	13	-.5	117.2	9
(23) <i>Counterfactual</i>	.5	0	4	-.5	117.5	29.2
(24) <i>Counterfactual</i>	.5	.2	3.9	-.5	117.5	30.4
(25) <i>Counterfactual</i>	.5	.6	3.7	-.5	116.9	31.3
(26) <i>Counterfactual</i>	.5	.7	3.6	-.5	116.9	32
(27) <i>Counterfactual</i>	.5	1	3.3	-.5	116.2	35.3
(28) <i>Counterfactual</i>	1	-1	18.8	-.2	306.3	16.3
(29) <i>Counterfactual</i>	1	-.7	12.9	-.2	306.3	23.8
(30) <i>Counterfactual</i>	1	-.6	13.1	-.2	306.3	23.4
(31) <i>Counterfactual</i>	1	-.2	13.5	-.2	306.3	22.7
(32) <i>Counterfactual</i>	1	0	4	-.2	306.3	75.8
(33) <i>Counterfactual</i>	1	.2	3.5	-.2	306.3	87.6
(34) <i>Counterfactual</i>	1	.6	2.9	-.2	306.3	103.9
(35) <i>Counterfactual</i>	1	.7	2.7	-.2	306.3	114.5
(36) <i>Counterfactual</i>	1	1	1.3	-.2	306.3	234.2

Notes: Additional counterfactual estimates using alternative values of demand distortion (γ) and correlation between private and social benefits (ρ)

Table A6: Alternative assumptions

Assumption	Cost per tCO _{2e}
Benchmark	US\$13
(1) One-year durability	US\$25
(2) Three-year durability	US\$8
(3) Fraction non-renewable biomass (fNRB) of 10%	US\$16
(4) Fraction non-renewable biomass (fNRB) of 100%	US\$10
(5) US grid emissions intensity	US\$13

Notes: Subsidy efficiency results under alternative assumptions. The benchmark is this paper’s main estimate. Scenarios (1) and (2) assume that the stove’s expected life span is one year or three years (the main estimate uses 1.9). Scenarios (3) and (4) assume that the fraction of the wood used for charcoal production that is non-renewable is either 10% or 100% (the main estimate assumes 38%). Scenario (5) assumes that Kenya’s average grid emissions were those of the U.S., at 366 grams of CO_{2e} per kWh (the main estimate uses Kenya’s grid average of 154 grams of CO_{2e} per kWh).

Table A7: Households using charcoal as their primary cooking fuel

Country	Households (millions)	Charcoal (millions)	Wood (millions)	Source
China	522.7	17.7	50.1	CNSB (2020)
Madagascar	6.9	4.5	2.2	MNIS (2022; 2022)
Tanzania	13.8	3.6	7.7	TNBS (2022)
Nigeria	40.8	3.2	23.4	NNBS (2024)
Burkina Faso	18.3	3.1	4.4	BFNISD (2021)
Dem. Rep. of the Congo	13.2	3.1	9.9	DMPH (2021)
Ghana	8.4	2.6	3	GSS (2022)
Philippines	23	2	8.5	PSA (2022)
Bangladesh	41	1.7	28.2	BBS (2023; 2023)
Uganda	7.2	1.7	5.1	UBS (2020)
Ethiopia	19.5	1.6	15.7	EPHI (2019)
Mozambique	6.2	1.4	4.3	MNIS (2023)
Kenya	12	1.4	6.4	KNBS (2022)
Zambia	2.8	1.2	1.4	ZSA (2018)
Myanmar	11.2	1.1	6.7	MMPF (2017)
Other		4.7	128.7	
Total		54.3	305.9	

Notes: Charcoal households are those that use charcoal as their primary cooking fuel. Numbers listed are from census data asking households which fuel they use as their primary cooking fuel. Only countries for which census data lists at least 1 million charcoal households are listed. Some countries may be missing.

III Additional derivations

III.1 Proofs: Model set-up

Proof of Lemma 1. Since $\mu'' < 0$, μ' is invertible. Recall that x^* is the unique solution to $\mu'(x^*) = -\pi'$. The FOC then implies $d_{2i}^* \theta_i - \alpha_i = x^*$ for every agent: regardless of an agent's energy demand θ_i or taste α_i , the absolute quantity of clean energy consumed in excess of the agent's taste parameter is the same. Agents with higher θ_i simply use a smaller fraction d_{2i}^* to reach the same level; agents with higher α_i use more clean energy in total but face the same marginal trade-off.

Each agent's optimal fraction is determined by their preference parameters: $d_{2i}^* = \frac{x^* + \alpha_i}{\theta_i}$. Substituting $d_{2i}^* \theta_i = x^* + \alpha_i$ into WTP:

$$wtp_i = (x^* + \alpha_i)\pi' + \mu(x^*) = \underbrace{x^* \pi' + \mu(x^*)}_{\equiv a} + \alpha_i \pi'$$

The constant a captures the common component of WTP shared by all agents — the fuel savings and switching costs evaluated at the common usage margin x^* . Or, put another way, x^* is the usage level for agents with $\alpha_i = 0$ and a is the utility of adoption for such an agent. Utility then scales linearly in $\alpha_i \pi'$. Since heterogeneity in WTP depends only on α_i , heterogeneity in the adoption decision is fully determined by α_i . \square

Proof of Lemma 2. From Section 1.3, $wtp_i = a + \alpha_i \pi'$ where $a = x^* \pi' + \mu(x^*)$. Under δ , the agent perceives fuel price $\delta \pi'$, so define \tilde{x}^* as the unique solution to $\mu'(\tilde{x}^*) = -\delta \pi'$. Since $\delta < 1$, we have $\mu'(\tilde{x}^*) = -\delta \pi' > -\pi' = \mu'(x^*)$, and since μ' is strictly decreasing ($\mu'' < 0$), it follows that $\tilde{x}^* < x^*$: agents underuse the technology when they underperceive fuel savings. The same steps as in Section 1.3 yield:

$$\widetilde{wtp}_i = (\tilde{x}^* + \alpha_i)(\delta \pi') + \mu(\tilde{x}^*) = \underbrace{\tilde{x}^* \delta \pi' + \mu(\tilde{x}^*)}_{\equiv \tilde{a}} + \alpha_i(\delta \pi')$$

Substituting $\alpha_i = (wtp_i - a)/\pi'$:

$$\widetilde{wtp}_i = \delta wtp_i + k \quad \text{where } k = \tilde{a} - \delta a$$

Expanding k :

$$\begin{aligned} k &= \tilde{a} - \delta a \\ &= [\tilde{x}^* \delta \pi' + \mu(\tilde{x}^*)] - \delta [x^* \pi' + \mu(x^*)] \\ &= \delta \pi'(\tilde{x}^* - x^*) + \mu(\tilde{x}^*) - \delta \mu(x^*) \end{aligned}$$

We now show that $k \leq 0$, meaning that marginal distortions reduce WTP by *at least* as much as proportional scaling by δ . The strategy is to treat k as a function of δ on $[0, 1]$ and show it equals zero at both endpoints but is strictly convex in between, so it must be weakly negative throughout.

Step 1. $k(1) = 0$: no distortion, no gap. When $\delta = 1$ there is no distortion, $\widetilde{wtp}_i = wtp_i$ which implies $k = 0$.

Step 2. $k(0) = 0$: complete distortion, both sides vanish. When $\delta = 0$ the agent perceives zero fuel savings and k reduces to $\mu(\tilde{x}^*)$. Intuitively, WTP shifts by the amount of disutility the agent will receive at their optimal level of usage given that there are no savings to be had. The FOC requires $\mu'(\tilde{x}^*) = 0$. Since $\mu'(0) = 0$ by assumption, we have $\tilde{x}^*(0) = 0$. Because $\mu(0) = 0$ we thus have $k = 0$.

Step 3. k is strictly convex: the envelope theorem and diminishing returns.

First derivative. Since $k = \tilde{a}(\delta) - \delta a$, where a does not depend on δ :

$$\frac{dk}{d\delta} = \frac{d\tilde{a}}{d\delta} - a$$

To evaluate $d\tilde{a}/d\delta$, expand $\tilde{a}(\delta) = \tilde{x}^*(\delta)\delta\pi' + \mu(\tilde{x}^*(\delta))$ by the chain rule:

$$\frac{d\tilde{a}}{d\delta} = \frac{d\tilde{x}^*}{d\delta} \delta \pi' + \tilde{x}^* \pi' + \mu'(\tilde{x}^*) \frac{d\tilde{x}^*}{d\delta}$$

Collecting the terms that involve $d\tilde{x}^*/d\delta$:

$$\begin{aligned} &= \frac{d\tilde{x}^*}{d\delta} \underbrace{[\delta \pi' + \mu'(\tilde{x}^*)]}_{= 0 \text{ by the FOC}} + \tilde{x}^* \pi' = \tilde{x}^*(\delta) \pi' \end{aligned}$$

The FOC zeroes out all terms involving re-optimization of usage, following the envelope theorem. Because the agent has already chosen \tilde{x}^* optimally, a small change in δ affects the distorted value function \tilde{a} only through its direct effect on perceived savings, not through the induced change in behavior. Intuitively, the agent is already at a point where the marginal cost of switching equals the marginal benefit, so adjusting usage neither helps nor hurts to

first order. Substituting:

$$\frac{dk}{d\delta} = \tilde{x}^*(\delta)\pi' - a$$

Second derivative. As δ increases, \tilde{x}^* rises—the agent shifts more consumption to clean energy. To find the rate, differentiate the FOC $\mu'(\tilde{x}^*) = -\delta\pi'$ with respect to δ :

$$\mu''(\tilde{x}^*)\frac{d\tilde{x}^*}{d\delta} = -\pi' \quad \implies \quad \frac{d\tilde{x}^*}{d\delta} = \frac{-\pi'}{\mu''(\tilde{x}^*)} > 0$$

(Positive because $\mu'' < 0$.) Therefore:

$$\frac{d^2k}{d\delta^2} = \frac{d\tilde{x}^*}{d\delta}\pi' = \frac{-(\pi')^2}{\mu''(\tilde{x}^*)} > 0$$

Conclusion. Since k is convex with $k(0) = k(1) = 0$, it follows that $k(\delta) \leq 0$ for all $\delta \in [0, 1]$.²³ □

²³For any convex function f with $f(0) = f(1) = 0$: $f(\delta) = f((1-\delta)0 + \delta 1) \leq (1-\delta)f(0) + \delta f(1) = 0$.

Proof of Lemma 3. Under intertemporal distortions (γ), the agent correctly perceives π' and optimizes usage accordingly, but discounts the resulting benefit by γ at the time of adoption. Since $\widetilde{wtp}_i = \gamma wtp_i$, carbon taxes and marginal cost subsidies are equally dampened:

$$\frac{\partial \widetilde{wtp}_i}{\partial \tau} = \gamma \frac{\partial wtp_i}{\partial \tau} = \gamma \frac{\partial wtp_i}{\partial s_2} = \frac{\partial \widetilde{wtp}_i}{\partial s_2}$$

Under marginal distortions (δ), the agent perceives only a fraction δ of any change in π' . Since the envelope theorem implies that the distorted usage fraction \tilde{d}_{2i}^* does not adjust to first order, the effect on WTP operates entirely through the perceived price:

$$\frac{\partial \widetilde{wtp}_i}{\partial \tau} = \delta \tilde{d}_{2i}^* \theta_i < d_{2i}^* \theta_i = \frac{\partial wtp_i}{\partial \tau}$$

where the inequality holds because $\delta < 1$ and because $\tilde{d}_{2i}^* < d_{2i}^*$.

Result: Demand distortions reduce the efficacy of marginal cost instruments. Whether the principal raises p_H (carbon tax), lowers p_L (marginal cost subsidy), or both, the agent's adoption response is dampened by γ , δ , or both. As distortions become severe, changes in the relative fuel price have vanishing effects on adoption.

Result: Fixed cost subsidies bypass the relative fuel price channel. A fixed cost subsidy s_1 reduces the adoption threshold from p_{mc} to $p_{mc} - s_1$, which does not require the agent to respond to fuel price signals. \square

Proof of Lemma 4. Differentiating the adoption condition $\gamma wtp(\underline{\alpha}) = p_{mc} - s_1$ with respect to γ , holding instruments s_1 and s_2 fixed:

$$wtp(\underline{\alpha}) + \gamma \pi' \frac{\partial \underline{\alpha}}{\partial \gamma} = 0 \quad \implies \quad \frac{\partial \underline{\alpha}}{\partial \gamma} = -\frac{wtp(\underline{\alpha})}{\gamma \pi'} < 0$$

since $wtp(\underline{\alpha}) > 0$ (the marginal adopter has positive willingness-to-pay) and $\gamma, \pi' > 0$.

For marginal distortions, the adoption condition (setting $\gamma = 1$) is $\widetilde{wtp}(\underline{\alpha}) = p_{mc} - s_1$, where $\widetilde{wtp}(\underline{\alpha}) = \tilde{a} + \underline{\alpha} \delta \pi'$. Differentiating with respect to δ , holding instruments fixed, and using $d\tilde{a}/d\delta = \tilde{x}^* \pi'$ (envelope theorem, from the proof of Lemma 2):

$$\tilde{x}^* \pi' + \underline{\alpha} \pi' + \delta \pi' \frac{\partial \underline{\alpha}}{\partial \delta} = 0 \quad \implies \quad \frac{\partial \underline{\alpha}}{\partial \delta} = -\frac{\tilde{x}^* + \underline{\alpha}}{\delta} < 0$$

since $\tilde{x}^* + \underline{\alpha} > 0$ (the marginal adopter has positive usage) and $\delta > 0$. □

Proof of Proposition 2. We first show the case with either an intertemporal distortion or a marginal distortion, and then derive the case with both.

Intertemporal distortion ($\gamma < 1, \delta = 1$). Since γ does not affect usage (Section 1.4), the Pigouvian tax $\tau^* = \phi$ still sets the socially efficient usage level. At this tax, the undistorted adoption rule would be $wtp_i \geq p_{mc}$, i.e. $SV_i \geq 0$. But the distorted adoption rule is $\gamma wtp_i \geq p_{mc} - s_1$. At the socially optimal margin, the marginal adopter has $SV = 0$, i.e. $wtp = p_{mc}$, so:

$$\gamma p_{mc} = p_{mc} - s_1 \quad \implies \quad s_1^* = p_{mc}(1 - \gamma)$$

Marginal distortion ($\delta < 1, \gamma = 1$). Under δ , both margins are distorted. To correct usage, the planner must set π' high enough that the perceived price $\delta\pi'$ equals the socially efficient level:

$$\delta\pi' = \pi + \phi \quad \implies \quad \pi' = \frac{\pi + \phi}{\delta}$$

The optimal marginal tax is then given by:

$$\tau^* = \frac{\phi + \pi(1 - \delta)}{\delta}$$

This “super-Pigouvian” tax overshoots the externality to compensate for the agent’s underperception of fuel prices. At this tax, usage is first-best: $\tilde{x}^* = x^{FB}$ where $\mu'(x^{FB}) = -(\pi + \phi)$. A key question is whether the planner also needs $s_1 > 0$ to correct adoption.

The agent’s perceived WTP at this tax is:

$$\widetilde{wtp}_i = (\tilde{x}^* + \alpha_i)(\delta\pi') + \mu(\tilde{x}^*) = (x^{FB} + \alpha_i)(\pi + \phi) + \mu(x^{FB})$$

where the second equality substitutes $\tilde{x}^* = x^{FB}$ and $\delta\pi' = \pi + \phi$. But this is exactly $SV_i + p_{mc}$. So the adoption rule $\widetilde{wtp}_i \geq p_{mc} - s_1$ is equivalent to $SV_i \geq -s_1$, and the socially optimal margin ($SV = 0$) obtains at $s_1^* = 0$.

The super-Pigouvian tax corrects *both* margins simultaneously, and no fixed cost subsidy is needed. The reason is that δ operates purely through the fuel price channel: the agent underperceives the fuel price by factor δ , and inflating the price by $1/\delta$ perfectly undoes this at both the usage and adoption margins. In contrast, γ scales the *entire* future utility—including switching costs μ —which no fuel-price instrument can reach.

Both distortions ($\gamma < 1, \delta < 1$). At the super-Pigouvian tax ($\delta\pi' = \pi + \phi$), the δ distortion is corrected and the agent’s perceived WTP equals $\gamma(SV_i + p_{mc})$, by the same argument as above with an additional γ scaling:

$$\widetilde{wtp}_i = \gamma [(x^{FB} + \alpha_i)(\pi + \phi) + \mu(x^{FB})] = \gamma(SV_i + p_{mc})$$

The adoption rule $\gamma(SV_i + p_{mc}) \geq p_{mc} - s_1$ at the optimal margin ($SV = 0$) gives:

$$\gamma p_{mc} = p_{mc} - s_1 \quad \implies \quad s_1^* = p_{mc}(1 - \gamma)$$

The same expression as under γ alone. □

III.2 Proofs: Social principal

Proof of Equation 1. Since s_1 does not affect usage \tilde{x}^* , it changes social welfare only through the adoption threshold $\underline{\alpha}$.

Marginal social welfare. By the Leibniz rule, each marginal adopter contributes $SV_i(\underline{\alpha}) = (\tilde{x}^* + \underline{\alpha})(\pi + \phi) + \mu(\tilde{x}^*) - p_{mc}$ to social welfare:

$$\frac{\partial W}{\partial s_1} = [(\tilde{x}^* + \underline{\alpha})(\pi + \phi) + \mu(\tilde{x}^*) - p_{mc}] g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_1} \right)$$

Marginal cost. Recall $C = s_1 Q + s_2 U$. A dollar increase in s_1 has three cost effects: (i) \$1 more to each existing adopter, (ii) s_1 paid to new adopters, and (iii) s_2 paid on new adopters' usage. By the product rule and Leibniz rule:

$$\begin{aligned} \frac{\partial C}{\partial s_1} &= \underbrace{Q}_{(i)} + \underbrace{s_1 g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_1} \right)}_{(ii)} + \underbrace{s_2 (\tilde{x}^* + \underline{\alpha}) g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_1} \right)}_{(iii)} \\ &= Q + [s_1 + s_2 (\tilde{x}^* + \underline{\alpha})] g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_1} \right) \end{aligned}$$

FOC. Setting $\partial W / \partial s_1 = \lambda \partial C / \partial s_1$ and dividing both sides by $g(\underline{\alpha})(-\partial \underline{\alpha} / \partial s_1)$:

$$(\tilde{x}^* + \underline{\alpha})(\pi + \phi) + \mu(\tilde{x}^*) - p_{mc} = \lambda \left[s_1 + s_2 (\tilde{x}^* + \underline{\alpha}) + \frac{Q}{g(\underline{\alpha})(-\partial \underline{\alpha} / \partial s_1)} \right]$$

Since $g(\underline{\alpha})(-\partial \underline{\alpha} / \partial s_1)$ is the slope of the demand curve at price $p_{mc} - s_1$, the ratio $Q / [g(\underline{\alpha})(-\partial \underline{\alpha} / \partial s_1)]$ equals $(p_{mc} - s_1) / |\tilde{\mathcal{E}}(p_{mc} - s_1)|$, giving the result. □

Proof of Equation 2. Since s_2 affects both usage \tilde{x}^* (through π') and the adoption threshold $\underline{\alpha}$, the social welfare derivative has an intensive and an extensive margin.

Marginal social welfare. Recall $W = \int_{\underline{\alpha}}^{\infty} [(\tilde{x}^* + \alpha_i)(\pi + \phi) + \mu(\tilde{x}^*) - p_{mc}] g(\alpha) d\alpha$. By the Leibniz rule, s_2 affects W through the integrand (via \tilde{x}^*) and through the lower limit (via $\underline{\alpha}$):

$$\frac{\partial W}{\partial s_2} = \underbrace{[(\pi + \phi) + \mu'(\tilde{x}^*)] \frac{d\tilde{x}^*}{ds_2} Q}_{\text{intensive}} + \underbrace{[(\tilde{x}^* + \underline{\alpha})(\pi + \phi) + \mu(\tilde{x}^*) - p_{mc}] g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_2}\right)}_{\text{extensive}}$$

The intensive term captures the social value of each existing adopter shifting $d\tilde{x}^*/ds_2$ more usage to clean energy: each additional unit saves π in fuel costs and ϕ in emissions, but incurs marginal switching cost $\mu'(\tilde{x}^*)$. The extensive term is the social value of the marginal adopter, $SV(\underline{\alpha})$, times the number of new adopters.

Marginal cost. Applying the product rule to $C = s_1Q + s_2U$:

$$\frac{\partial C}{\partial s_2} = \underbrace{U}_{(i)} + s_2 \underbrace{\frac{d\tilde{x}^*}{ds_2} Q}_{(ii)} + \underbrace{[s_1 + s_2(\tilde{x}^* + \underline{\alpha})] g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_2}\right)}_{(iii)}$$

where (i) is \$1 more on all existing usage, (ii) is s_2 on additional usage by existing adopters, and (iii) is the full subsidy package $\sigma = s_1 + s_2(\tilde{x}^* + \underline{\alpha})$ to new adopters.

FOC. Setting $\partial W/\partial s_2 = \lambda \partial C/\partial s_2$ gives the result. □

Proof of Equation 3. Step 1: Rewrite the s_1 FOC in compact form. Let $\sigma \equiv s_1 + s_2(\tilde{x}^* + \underline{\alpha})$ and $\kappa \equiv (p_{mc} - s_1)/|\tilde{\mathcal{E}}(p_{mc} - s_1)|$. The s_1 FOC is:

$$SV(\underline{\alpha}) = \lambda(\sigma + \kappa)$$

Step 2: Rearrange the s_2 FOC. Group the intensive-margin terms (those with $\frac{d\tilde{x}^*}{ds_2} Q$) separately from the extensive-margin terms (those with $g(\underline{\alpha})(-\partial\underline{\alpha}/\partial s_2)$):

$$(MSV - \lambda s_2) \frac{d\tilde{x}^*}{ds_2} Q + [SV(\underline{\alpha}) - \lambda\sigma] g(\underline{\alpha}) \left(-\frac{\partial\underline{\alpha}}{\partial s_2} \right) = \lambda U$$

Step 3: Substitute from the s_1 FOC. The extensive-margin bracket simplifies: $SV(\underline{\alpha}) - \lambda\sigma = \lambda\kappa$. Substituting:

$$(MSV - \lambda s_2) \frac{d\tilde{x}^*}{ds_2} Q + \lambda\kappa g(\underline{\alpha}) \left(-\frac{\partial\underline{\alpha}}{\partial s_2} \right) = \lambda U$$

Step 4: Simplify the κ term using the derivative ratio. Writing $\kappa = Q/[g(\underline{\alpha})(-\partial\underline{\alpha}/\partial s_1)]$, the $g(\underline{\alpha})$ factors cancel:

$$\lambda\kappa g(\underline{\alpha}) \left(-\frac{\partial\underline{\alpha}}{\partial s_2} \right) = \lambda Q \cdot \frac{-\partial\underline{\alpha}/\partial s_2}{-\partial\underline{\alpha}/\partial s_1}$$

To evaluate this ratio, differentiate the adoption condition for the marginal adopter $wtp(\underline{\alpha}) = \gamma[\tilde{a} + \underline{\alpha} \delta \pi'] = p_{mc} - s_1$ with respect to each instrument. With respect to s_1 (noting $\pi' = \pi + s_2$ does not depend on s_1 , and neither does \tilde{a}):

$$\gamma \delta \pi' \frac{\partial\underline{\alpha}}{\partial s_1} = -1 \quad \implies \quad -\frac{\partial\underline{\alpha}}{\partial s_1} = \frac{1}{\gamma \delta \pi'}$$

With respect to s_2 (because the envelope theorem implies usage does not adjust to the first order and $\tilde{a} = \tilde{x}^* \delta \pi' + \mu(\tilde{x}^*)$, where s_2 only enters through π' , we have that $d\tilde{a}/ds_2 = \tilde{x}^* \delta$):

$$\gamma \left[\tilde{x}^* \delta + \underline{\alpha} \delta + \delta \pi' \frac{\partial\underline{\alpha}}{\partial s_2} \right] = 0 \quad \implies \quad -\frac{\partial\underline{\alpha}}{\partial s_2} = \frac{\tilde{x}^* + \underline{\alpha}}{\pi'}$$

The κ term is therefore:

$$\lambda Q \cdot \frac{-\partial\underline{\alpha}/\partial s_2}{-\partial\underline{\alpha}/\partial s_1} = \lambda Q \cdot \frac{(\tilde{x}^* + \underline{\alpha})/\pi'}{1/(\gamma \delta \pi')} = \lambda Q \gamma \delta (\tilde{x}^* + \underline{\alpha})$$

Substituting:

$$(MSV - \lambda s_2) \frac{d\tilde{x}^*}{ds_2} Q + \lambda Q \gamma \delta (\tilde{x}^* + \underline{\alpha}) = \lambda U$$

Step 5: Divide by Q and substitute $U/Q = \tilde{x}^ + \bar{\alpha}$.*

$$(MSV - \lambda s_2) \frac{d\tilde{x}^*}{ds_2} = \lambda [(\tilde{x}^* + \bar{\alpha}) - \gamma \delta (\tilde{x}^* + \underline{\alpha})]$$

Step 6: Eliminate λ . From the s_1 FOC, $\lambda = SV(\underline{\alpha})/(\sigma + \kappa)$. Dividing both sides of Step 5 by λ :

$$\left[\frac{MSV(\sigma + \kappa)}{SV(\underline{\alpha})} - s_2 \right] \frac{d\tilde{x}^*}{ds_2} = (\tilde{x}^* + \bar{\alpha}) - \gamma \delta (\tilde{x}^* + \underline{\alpha})$$

Step 7: Simplify the LHS. Since $\sigma = s_1 + s_2(\tilde{x}^* + \underline{\alpha})$:

$$\begin{aligned} \frac{MSV(\sigma + \kappa)}{SV(\underline{\alpha})} - s_2 &= \frac{MSV \cdot \sigma + MSV \cdot \kappa - s_2 \cdot SV(\underline{\alpha})}{SV(\underline{\alpha})} \\ &= \frac{MSV \cdot s_1 + s_2[MSV(\tilde{x}^* + \underline{\alpha}) - SV(\underline{\alpha})] + MSV \cdot \kappa}{SV(\underline{\alpha})} \\ &= \frac{MSV(s_1 + \kappa) + s_2 \cdot \Delta}{SV(\underline{\alpha})} \end{aligned}$$

where $\Delta \equiv MSV(\tilde{x}^* + \underline{\alpha}) - SV(\underline{\alpha}) = \mu'(\tilde{x}^*)(\tilde{x}^* + \underline{\alpha}) - \mu(\tilde{x}^*) + p_{mc}$ collects the terms that distinguish the social principal from the environmental principal. Dividing by $d\tilde{x}^*/ds_2$:

$$\frac{MSV \cdot (s_1 + \kappa) + s_2 \cdot \Delta}{SV(\underline{\alpha})} = \frac{(\tilde{x}^* + \bar{\alpha}) - \gamma\delta(\tilde{x}^* + \underline{\alpha})}{\frac{d\tilde{x}^*}{ds_2}}$$

□

Proof of Equation 4. We set $\delta = 1$ throughout to isolate γ , so that $\tilde{x}^* = x^*$ and usage is undistorted. From the adoption condition $s_1^* = p_{mc} - \gamma wtp(\underline{\alpha}^*)$, where $wtp(\underline{\alpha}^*) = a + \underline{\alpha}^* \pi'$ is the marginal adopter's undistorted WTP, differentiating with respect to γ and noting $wtp'(\underline{\alpha}^*) = \pi'$:

$$\frac{ds_1^*}{d\gamma} = \underbrace{-wtp(\underline{\alpha}^*)}_{\text{direct effect}} \underbrace{-\gamma \pi' \frac{d\underline{\alpha}^*}{d\gamma}}_{\text{re-optimization}}$$

The direct effect is unambiguously negative: when γ rises (distortions weaken), agents value the technology more, so the principal can reduce s_1^* while maintaining the same adoption threshold.

The sign of the re-optimization term depends on $d\underline{\alpha}^*/d\gamma$, how the principal's optimal adoption threshold responds to the distortion.²⁴ The proof of Proposition 3 solves for $d\underline{\alpha}^*/d\gamma$ and substitutes back to sign $ds_1^*/d\gamma < 0$. □

²⁴Note that earlier we showed greater distortions increase the adoption threshold, *for a given* s_1 and s_2 . This derivative asks a different question: how does the threshold change after allowing the principal to re-optimize s_1 and s_2 .

Proof of Proposition 3. From Equation 4, $ds_1^*/d\gamma = -wtp(\underline{\alpha}^*) - \gamma\pi' \cdot d\underline{\alpha}^*/d\gamma$, where the direct effect is unambiguously negative. We now solve for $d\underline{\alpha}^*/d\gamma$ using the implicit function theorem and substitute back to sign $ds_1^*/d\gamma$. We set $\delta = 1$ throughout to isolate γ , so that $\tilde{x}^* = x^*$ and usage is undistorted.

Step A: Express s_1 and κ as functions of $(\underline{\alpha}^, \gamma)$.* The adoption condition gives $s_1^* = p_{mc} - \gamma wtp(\underline{\alpha}^*)$. The inframarginal cost κ likewise depends only on $\underline{\alpha}^*$ and γ once expressed in hazard rate form. From Section 3.1, the distorted elasticity satisfies $\tilde{\mathcal{E}}(p_{mc} - s_1^*) = \mathcal{E}((p_{mc} - s_1^*)/\gamma)$, evaluated at $wtp(\underline{\alpha}^*)$. Since $wtp = a + \alpha\pi'$, undistorted demand at price w is $Q(w) = 1 - G(\frac{w-a}{\pi'})$, so the undistorted elasticity at price w is:

$$|\mathcal{E}(w)| = \frac{w}{\pi'} \cdot h\left(\frac{w-a}{\pi'}\right)$$

where $h(\cdot) \equiv g(\cdot)/(1 - G(\cdot))$ is the hazard rate of the α distribution. Evaluating at $w = wtp(\underline{\alpha}^*)$ and substituting into $\kappa = (p_{mc} - s_1^*)/|\tilde{\mathcal{E}}|$:

$$\kappa = \frac{\gamma\pi'}{h(\underline{\alpha}^*)}$$

Step B: Construct the implicit function. The social instrument mix equation (3) is:

$$\frac{MSV(s_1^* + \kappa) + s_2\Delta}{SV(\underline{\alpha})} = \frac{\bar{u} - \gamma u}{\eta}$$

where $\underline{u} = x^* + \underline{\alpha}^*$, $\bar{u} = x^* + \bar{\alpha}$, $MSV = (\pi + \phi) + \mu'(x^*)$, $\Delta = MSV \cdot \underline{u} - SV(\underline{\alpha})$, and $\eta = dx^*/ds_2$. We treat s_2 (and hence $\pi' = \pi + s_2$, x^* , η) as a fixed parameter: if $ds_1^*/d\gamma < 0$ holds for any fixed s_2 , it holds at the equilibrium s_2 in particular. With s_2 held fixed, x^* , η , and MSV are constants determined by π' alone — they depend on neither $\underline{\alpha}^*$ nor γ . The terms \underline{u} , \bar{u} , $SV(\underline{\alpha})$, and Δ depend on $\underline{\alpha}^*$ but not on γ (since γ does not affect the distribution of α or the usage level x^*). The $s_2\Delta$ term is likewise independent of γ . The only γ -dependent terms are s_1 and κ (both explicitly, from Step A) and the γu on the RHS. Since Step A expresses s_1 and κ as functions of $(\underline{\alpha}^*, \gamma)$, the entire equation depends only on $\underline{\alpha}^*$, γ , and s_2 . Cross-multiplying and defining:

$$\Psi(\underline{\alpha}^*, \gamma) \equiv \eta MSV(s_1^* + \kappa) + \eta s_2\Delta - SV(\underline{\alpha})(\bar{u} - \gamma u) = 0$$

Step C: Apply the implicit function theorem. Totally differentiating $\Psi(\underline{\alpha}^*(\gamma), \gamma) = 0$:

$$\frac{d\underline{\alpha}^*}{d\gamma} = -\frac{\partial\Psi/\partial\gamma}{\partial\Psi/\partial\underline{\alpha}^*}$$

We compute the numerator $\partial\Psi/\partial\gamma$ in Step D and the denominator $\partial\Psi/\partial\underline{\alpha}^*$ in Step E, then substitute back into $ds_1^*/d\gamma$ in Step F.

Step D: Compute and simplify $\partial\Psi/\partial\gamma$. From Step A, $s_1^* + \kappa = p_{mc} + \gamma[-wtp + \pi'/h(\underline{\alpha}^*)]$. Differentiating Ψ with respect to γ (holding $\underline{\alpha}^*$ and s_2 fixed):

$$\frac{\partial\Psi}{\partial\gamma} = \eta MSV \left[-wtp + \frac{\pi'}{h(\underline{\alpha}^*)} \right] + SV(\underline{\alpha}) u$$

From Step A, $\kappa = \gamma\pi'/h(\underline{\alpha}^*)$, so $\pi'/h(\underline{\alpha}^*) = \kappa/\gamma$. And from the original definition $\kappa = (p_{mc} - s_1^*)/|\tilde{\mathcal{E}}|$

with $p_{mc} - s_1^* = \gamma wtp$, we have $\kappa/\gamma = wtp(\underline{\alpha}^*)/|\tilde{\mathcal{E}}|$. Substituting into the bracket:

$$\frac{\partial \Psi}{\partial \gamma} = \underline{u} SV(\underline{\alpha}) - \eta MSV wtp(\underline{\alpha}^*) \cdot \frac{|\tilde{\mathcal{E}}| - 1}{|\tilde{\mathcal{E}}|} \quad (11)$$

Step E: Compute $\partial \Psi / \partial \underline{\alpha}^$.* Differentiating Ψ term by term:

The first term, $\eta MSV(s_1^* + \kappa)$, depends on $\underline{\alpha}^*$ through $wtp(\underline{\alpha}^*)$ and $h(\underline{\alpha}^*)$. Since $s_1^* + \kappa = p_{mc} + \gamma[-wtp + \pi'/h(\underline{\alpha}^*)]$, and $\partial wtp / \partial \underline{\alpha}^* = \pi'$ and $\partial[\pi'/h(\underline{\alpha}^*)] / \partial \underline{\alpha}^* = -\pi' h'(\underline{\alpha}^*) / h(\underline{\alpha}^*)^2$, this contributes:

$$-\eta MSV \gamma \pi' \left[1 + \frac{h'(\underline{\alpha}^*)}{h(\underline{\alpha}^*)^2} \right]$$

The second term, $\eta s_2 \Delta$, depends on $\underline{\alpha}^*$ through $\Delta = MSV \cdot \underline{u} - SV(\underline{\alpha})$. Since $\partial \Delta / \partial \underline{\alpha}^* = MSV - (\pi + \phi) = \mu'(x^*)$, this contributes $\eta s_2 \mu'(x^*)$.

The third term, $-SV(\underline{\alpha}) \bar{u}$, depends on $\underline{\alpha}^*$ through both $SV(\underline{\alpha})$ and $\bar{u} = x^* + \bar{\alpha}$. Since $\partial SV / \partial \underline{\alpha}^* = \pi + \phi$ and $\partial \bar{\alpha} / \partial \underline{\alpha}^* = h(\underline{\alpha}^*)(\bar{\alpha} - \underline{\alpha}^*)$,²⁵ the product rule gives: $-(\pi + \phi) \bar{u} - SV(\underline{\alpha}) h(\underline{\alpha}^*)(\bar{\alpha} - \underline{\alpha}^*)$.

The fourth term, $\gamma SV(\underline{\alpha}) \underline{u}$, contributes $\gamma[(\pi + \phi) \underline{u} + SV(\underline{\alpha})]$ by the product rule.

Collecting:

$$\frac{\partial \Psi}{\partial \underline{\alpha}^*} = \underbrace{-\eta MSV \gamma \pi' \left[1 + \frac{h'(\underline{\alpha}^*)}{h(\underline{\alpha}^*)^2} \right]}_{\text{cost-side response}} + \eta s_2 \mu'(x^*) - (\pi + \phi) \bar{u} - SV(\underline{\alpha}) h(\underline{\alpha}^*)(\bar{\alpha} - \underline{\alpha}^*) + \gamma[(\pi + \phi) \underline{u} + SV(\underline{\alpha})] \quad (12)$$

Step F: Substitute back and sign $ds_1^/d\gamma$.* From the decomposition (Equation 4) and Steps C–E:

$$\frac{ds_1^*}{d\gamma} = -wtp(\underline{\alpha}^*) + \gamma \pi' \cdot \frac{\partial \Psi / \partial \gamma}{\partial \Psi / \partial \underline{\alpha}^*}$$

where $\partial \Psi / \partial \gamma$ is given by (11) and $\partial \Psi / \partial \underline{\alpha}^*$ by (12). Since $\Psi = 0$ is the first-order condition and $\underline{\alpha}^*$ is the only remaining choice variable (given s_2), the second-order condition requires $\partial \Psi / \partial \underline{\alpha}^* < 0$ at a welfare maximum.

To sign $ds_1^*/d\gamma < 0$, we rearrange it into an equivalent condition $A > 0$, decompose A into four terms, and sign each under IHR. From the decomposition (Equation 4) and Steps C–E:

$$\frac{ds_1^*}{d\gamma} = -wtp(\underline{\alpha}^*) + \gamma \pi' \cdot \frac{\partial \Psi / \partial \gamma}{\partial \Psi / \partial \underline{\alpha}^*}$$

Multiplying both sides by $\partial \Psi / \partial \underline{\alpha}^* < 0$ (SOC) flips the inequality. So $ds_1^*/d\gamma < 0$ is equivalent to:

$$A \equiv \gamma \pi' \cdot \frac{\partial \Psi}{\partial \gamma} - wtp(\underline{\alpha}^*) \cdot \frac{\partial \Psi}{\partial \underline{\alpha}^*} > 0 \quad (13)$$

²⁵The truncated mean derivative follows from Leibniz's rule: $\partial \bar{\alpha} / \partial \underline{\alpha}^* = \partial \left[\int_{\underline{\alpha}^*}^{\infty} \alpha dG(\alpha) / (1 - G(\underline{\alpha}^*)) \right] / \partial \underline{\alpha}^* = h(\underline{\alpha}^*)(\bar{\alpha} - \underline{\alpha}^*)$.

Substituting $\partial\Psi/\partial\gamma$ from (11) and $\partial\Psi/\partial\alpha^*$ from (12):

$$A = \gamma\pi' \left[\eta MSV \left(-wtp + \frac{\pi'}{h(\alpha^*)} \right) + SV(\alpha) \underline{u} \right] \\ - wtp \left[-\eta MSV \gamma\pi' \left(1 + \frac{h'(\alpha^*)}{h(\alpha^*)^2} \right) + \eta s_2 \mu'(x^*) \right. \\ \left. - (\pi + \phi)\bar{u} - SV(\alpha) h(\alpha^*)(\bar{\alpha} - \alpha^*) + \gamma((\pi + \phi)\underline{u} + SV(\alpha)) \right]$$

Regrouping (the $-wtp$ and $+wtp$ cancel in the $\eta MSV \gamma\pi'$ terms):

$$A = \underbrace{\eta MSV \gamma\pi' \left(\frac{\pi'}{h(\alpha^*)} + \frac{wtp(\alpha^*) \cdot h'(\alpha^*)}{h(\alpha^*)^2} \right)}_{\text{Term I}} \quad (14)$$

$$+ \underbrace{SV(\alpha) [\gamma\pi'\underline{u} - \gamma wtp(\alpha^*) + wtp(\alpha^*) \cdot h(\alpha^*)(\bar{\alpha} - \alpha^*)]}_{\text{Term II}} \quad (15)$$

$$+ \underbrace{wtp(\alpha^*) \cdot (\pi + \phi) [\bar{u} - \gamma \underline{u}]}_{\text{Term III}} \quad (16)$$

$$+ \underbrace{(-wtp(\alpha^*) \cdot \eta s_2 \mu'(x^*))}_{\text{Term IV}} \quad (17)$$

Signing. We now show each term is non-negative. Note that $wtp(\alpha^*) > 0$ at any interior optimum, since $\gamma wtp = p_{mc} - s_1^* > 0$.

Term I > 0 under IHR. $\eta > 0$, $MSV = \phi - s_2 > 0$ (since $s_2 < \phi$ at any interior optimum where the marginal subsidy does not exceed the externality), $\gamma\pi' > 0$. The first summand $\pi'/h(\alpha^*) > 0$ always. Under IHR ($h'(\alpha^*) \geq 0$), the second summand $wtp \cdot h'(\alpha^*)/h(\alpha^*)^2 \geq 0$.

Term II > 0 . $SV(\alpha) \geq 0$ at an interior optimum. Using $wtp(\alpha^*) = \underline{u}\pi' + \mu(x^*)$, the first two terms in the bracket simplify: $\gamma\pi'\underline{u} - \gamma wtp = -\gamma\mu(x^*) \geq 0$ since $\mu(x^*) \leq 0$ (from $\mu(0) = 0$, $\mu'(0) = 0$, $\mu'' < 0$). The remaining term $wtp \cdot h(\alpha^*)(\bar{\alpha} - \alpha^*) > 0$ since $h(\alpha^*) > 0$ and $\bar{\alpha} > \alpha^*$. Both summands in the bracket are non-negative.

Term III > 0 . $\bar{u} - \gamma \underline{u} = (1 - \gamma)x^* + (\bar{\alpha} - \gamma\alpha^*) > 0$, since $(1 - \gamma)x^* \geq 0$ and $\bar{\alpha} > \alpha^* \geq \gamma\alpha^*$.

Term IV ≥ 0 . $\mu'(x^*) < 0$ (from $\mu'(0) = 0$ and $\mu'' < 0$), so $-\mu'(x^*) > 0$. When $s_2 > 0$, Term IV is strictly positive; when $s_2 = 0$, it vanishes.

Since $A = \text{I} + \text{II} + \text{III} + \text{IV}$ with all terms non-negative and Terms I–III strictly positive, $A > 0$. Therefore $ds_1^*/d\gamma < 0$ under IHR. □

III.3 Proofs: Environmental principal

Proof of Equation 5. The budget constraint, adoption condition, and cost structure are identical to the social principal's problem. The only difference is the objective: abatement ϕU rather than social welfare $\int SV_i g(\alpha) d\alpha$:

$$\max_{s_1, s_2} A(s_1, s_2) = \phi \int_{\underline{\alpha}}^{\infty} (\tilde{x}^* + \alpha_i) g(\alpha) d\alpha$$

where ϕ is the per-unit externality reduction. This simplifies the FOCs because ϕ factors out of both the intensive and extensive margins. The s_1 FOC replaces $SV(\underline{\alpha})$ on the LHS with $(\tilde{x}^* + \underline{\alpha})\phi$ —abatement from the marginal adopter—while the cost side (RHS) is unchanged:

$$(\tilde{x}^* + \underline{\alpha})\phi = \lambda \left[s_1 + s_2(\tilde{x}^* + \underline{\alpha}) + \frac{p_{mc} - s_1}{|\tilde{\mathcal{E}}(p_{mc} - s_1)|} \right] \quad (18)$$

The s_2 FOC similarly replaces MSV with ϕ in the intensive term and $SV(\underline{\alpha})$ with $(\tilde{x}^* + \underline{\alpha})\phi$ in the extensive term:

$$\phi \left[\frac{d\tilde{x}^*}{ds_2} Q + (\tilde{x}^* + \underline{\alpha}) g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_2} \right) \right] = \lambda \left[U + s_2 \frac{d\tilde{x}^*}{ds_2} Q + \sigma g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_2} \right) \right] \quad (19)$$

To take derivatives of A and C , we must first specify how $\underline{\alpha}$ and \tilde{x}^* depend on the instruments.

Adoption. The marginal adopter $\underline{\alpha}$ satisfies $\widetilde{wtp}(\underline{\alpha}) = p_{mc} - s_1$. From Section 1.4, $\widetilde{wtp}_i = \gamma(\tilde{a} + \delta \alpha_i \pi')$, where $\tilde{a} = \tilde{x}^* \delta \pi' + \mu(\tilde{x}^*)$. So the adoption condition is:

$$\gamma(\tilde{a} + \delta \underline{\alpha} \pi') = p_{mc} - s_1 \quad (20)$$

Note that s_1 enters the RHS directly, while s_2 enters through $\pi' = \pi + s_2$ and \tilde{a} (which depends on π' through \tilde{x}^*).

Usage. Clean energy usage per adopter is $\tilde{x}^* + \alpha_i$, where \tilde{x}^* is determined by the agent's perceived fuel price:

$$\mu'(\tilde{x}^*) = -\delta \pi' \quad (21)$$

Note that \tilde{x}^* depends on s_2 (through π') but not on s_1 .

How the instruments affect adoption. Both instruments lower the adoption threshold $\underline{\alpha}$, but through different channels: s_1 reduces the upfront cost directly, while s_2 raises the perceived fuel savings. Since the principal is choosing between the two, it is useful to compare their per-dollar effects on $\underline{\alpha}$. We derive each from the adoption condition (20).

Effect of s_1 . Differentiating (20) with respect to s_1 (only $\underline{\alpha}$ depends on s_1):

$$\gamma \delta \pi' \frac{\partial \underline{\alpha}}{\partial s_1} = -1 \quad \implies \quad \frac{\partial \underline{\alpha}}{\partial s_1} = \frac{-1}{\gamma \delta \pi'}$$

Effect of s_2 . Differentiating (20) with respect to s_2 (\tilde{a} , π' , and $\underline{\alpha}$ depend on s_2):

$$\gamma \left(\frac{\partial \tilde{a}}{\partial s_2} + \delta \underline{\alpha} + \delta \pi' \frac{\partial \underline{\alpha}}{\partial s_2} \right) = 0$$

To evaluate $\partial\tilde{a}/\partial s_2$, expand $\tilde{a} = \tilde{x}^*\delta\pi' + \mu(\tilde{x}^*)$ using the chain rule (noting $\partial\pi'/\partial s_2 = 1$):

$$\frac{\partial\tilde{a}}{\partial s_2} = \frac{d\tilde{x}^*}{ds_2}\delta\pi' + \tilde{x}^*\delta + \mu'(\tilde{x}^*)\frac{d\tilde{x}^*}{ds_2} = \frac{d\tilde{x}^*}{ds_2} \underbrace{[\delta\pi' + \mu'(\tilde{x}^*)]}_{= 0 \text{ by (21)}} + \delta\tilde{x}^* = \delta\tilde{x}^*$$

This is the *envelope theorem*: the usage FOC zeroes out all terms involving the re-optimization of \tilde{x}^* , so the effect of s_2 on the common utility level, \tilde{a} , operates only through the direct change in the perceived savings, which is δ times the common usage level \tilde{x}^* .

Substituting back:

$$\gamma \left(\underbrace{\delta(\tilde{x}^* + \alpha)}_{\text{perceived increase in marginal adopter's savings}} + \underbrace{\delta\pi' \frac{\partial\alpha}{\partial s_2}}_{\text{offsetting threshold adjustment}} \right) = 0$$

Dividing by $\gamma\delta$ and solving:

$$\frac{\partial\alpha}{\partial s_2} = -\frac{\tilde{x}^* + \alpha}{\pi'}$$

Comparison. Collecting both results:

$$\frac{\partial\alpha}{\partial s_1} = \frac{-1}{\gamma\delta\pi'} \quad \frac{\partial\alpha}{\partial s_2} = \frac{-(\tilde{x}^* + \alpha)}{\pi'} = \frac{\partial\alpha}{\partial s_1} \gamma\delta(\tilde{x}^* + \alpha)$$

A dollar of s_2 moves the adoption threshold by $\gamma\delta(\tilde{x}^* + \alpha)$ times as much as a dollar of s_1 . The mechanism: a dollar of s_1 reduces the upfront cost by \$1, while a dollar of s_2 raises π' by \$1, which by the envelope theorem raises the marginal adopter's perceived WTP by $\gamma\delta(\tilde{x}^* + \alpha)$ —proportional to their clean energy usage, scaled by the distortions. When $\gamma\delta(\tilde{x}^* + \alpha) < 1$, which is more likely under severe distortions (Section 1.5), s_1 is more effective per dollar at moving the adoption threshold.

Equivalently, since $\widetilde{wtp}_i = \gamma(\tilde{a} + \delta\alpha_i\pi')$ and $\partial\tilde{a}/\partial s_2 = \delta\tilde{x}^*$ by the envelope theorem:

$$\frac{\partial\widetilde{wtp}_i}{\partial s_2} = \gamma\delta(\tilde{x}^* + \alpha_i)$$

For the marginal adopter, $\partial\widetilde{wtp}_\alpha/\partial s_2 = \gamma\delta(\tilde{x}^* + \alpha)$: the derivative ratio above is exactly the marginal adopter's WTP response to a dollar of s_2 . This connects the dampening result of Section 1.5 to the instrument comparison: the dampened WTP derivative is what determines s_2 's ability to move the adoption threshold, and it reappears as the second term in the numerator of the instrument mix equation (3) below.

How s_2 affects usage. Unlike s_1 , the marginal cost subsidy also changes how much each adopter uses. A higher s_2 raises the effective fuel price π' , making clean energy more attractive, so each adopter shifts more consumption toward the clean technology. Differentiating the usage condition (21) with respect to s_2 :

$$\mu''(\tilde{x}^*)\frac{d\tilde{x}^*}{ds_2} = -\delta \quad \implies \quad \frac{d\tilde{x}^*}{ds_2} = \frac{-\delta}{\mu''(\tilde{x}^*)} > 0$$

This is positive (since $\mu'' < 0$) and is the intensive margin that s_2 has but s_1 does not: each

existing adopter generates more abatement. Note, however, that the principal must pay s_2 on every unit of this additional usage.

FOC for s_1 . We begin with s_1 because it is the simpler instrument: since it does not affect usage \tilde{x}^* , it changes both abatement and cost only through the adoption threshold $\underline{\alpha}$.

Marginal abatement. By the Leibniz rule, each marginal adopter uses $(\tilde{x}^* + \underline{\alpha})$ units of clean energy, generating $(\tilde{x}^* + \underline{\alpha})\phi$ in abatement:

$$\frac{\partial A}{\partial s_1} = \phi (\tilde{x}^* + \underline{\alpha}) g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_1} \right)$$

Marginal cost. Recall $C = s_1 Q + s_2 U$. A dollar increase in s_1 has three cost effects: (i) \$1 more to each existing adopter, (ii) s_1 paid to new adopters, and (iii) s_2 paid on new adopters' usage. By the product rule and Leibniz rule:

$$\begin{aligned} \frac{\partial C}{\partial s_1} &= \underbrace{Q}_{(i)} + \underbrace{s_1 g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_1} \right)}_{(ii)} + \underbrace{s_2 (\tilde{x}^* + \underline{\alpha}) g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_1} \right)}_{(iii)} \\ &= Q + [s_1 + s_2 (\tilde{x}^* + \underline{\alpha})] g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_1} \right) \end{aligned}$$

FOC. Setting $\partial A / \partial s_1 = \lambda \partial C / \partial s_1$:

$$\phi (\tilde{x}^* + \underline{\alpha}) g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_1} \right) = \lambda \left[Q + [s_1 + s_2 (\tilde{x}^* + \underline{\alpha})] g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_1} \right) \right]$$

For interpretability, divide both sides by $g(\underline{\alpha})(-\partial \underline{\alpha} / \partial s_1)$ —the number of new adopters per dollar of s_1 —to rescale to a per-marginal-adopter basis:

$$(\tilde{x}^* + \underline{\alpha})\phi = \lambda \left[s_1 + s_2 (\tilde{x}^* + \underline{\alpha}) + \frac{Q}{g(\underline{\alpha})(-\partial \underline{\alpha} / \partial s_1)} \right] \quad (22)$$

Since $g(\underline{\alpha})(-\partial \underline{\alpha} / \partial s_1)$ is the slope of the demand curve at price $p_{mc} - s_1$, this ratio equals $(p_{mc} - s_1) / |\tilde{\mathcal{E}}(p_{mc} - s_1)|$

For later use, define the total subsidy to the marginal adopter:

$$\sigma \equiv s_1 + s_2 (\tilde{x}^* + \underline{\alpha})$$

and define the inframarginal cost per new adopter:

$$\kappa \equiv (p_{mc} - s_1) / |\tilde{\mathcal{E}}(p_{mc} - s_1)|$$

In this notation, (22) is simply $(\tilde{x}^* + \underline{\alpha})\phi = \lambda(\sigma + \kappa)$.

Marginal abatement. Since s_2 affects both the integrand of A (through \tilde{x}^*) and the lower limit (through $\underline{\alpha}$):

$$\frac{\partial A}{\partial s_2} = \phi \left[\underbrace{\frac{d\tilde{x}^*}{ds_2} Q}_{\substack{\text{intensive:} \\ \text{existing adopters} \\ \text{use more}}} + \underbrace{(\tilde{x}^* + \underline{\alpha}) g(\underline{\alpha}) \left(-\frac{\partial \underline{\alpha}}{\partial s_2} \right)}_{\substack{\text{extensive:} \\ \text{new adopters}}} \right]$$

The intensive term is the usage increase per adopter times the number of existing adopters. The extensive term parallels s_1 : $g(\underline{\alpha})(-\partial\underline{\alpha}/\partial s_2)$ new adopters enter, each generating $(\tilde{x}^* + \underline{\alpha})\phi$ abatement.

Marginal cost. Applying the product rule to $C = s_1Q + s_2U$:

$$\frac{\partial C}{\partial s_2} = s_1 \frac{\partial Q}{\partial s_2} + U + s_2 \frac{\partial U}{\partial s_2}$$

By the Leibniz rule, $\partial Q/\partial s_2 = g(\underline{\alpha})(-\partial\underline{\alpha}/\partial s_2)$. For $\partial U/\partial s_2$, both the integrand (\tilde{x}^*) and the lower limit $(\underline{\alpha})$ respond to s_2 :

$$\frac{\partial U}{\partial s_2} = \frac{d\tilde{x}^*}{ds_2}Q + (\tilde{x}^* + \underline{\alpha})g(\underline{\alpha})\left(-\frac{\partial\underline{\alpha}}{\partial s_2}\right)$$

Substituting and regrouping into three cost effects paralleling s_1 —(i) \$1 more on all existing usage, (ii) s_2 on additional usage by existing adopters, (iii) full subsidy package to new adopters:

$$\frac{\partial C}{\partial s_2} = \underbrace{U}_{(i)} + \underbrace{s_2 \frac{d\tilde{x}^*}{ds_2}Q}_{(ii)} + \underbrace{[s_1 + s_2(\tilde{x}^* + \underline{\alpha})]g(\underline{\alpha})\left(-\frac{\partial\underline{\alpha}}{\partial s_2}\right)}_{(iii)}$$

FOC. Setting $\partial A/\partial s_2 = \lambda \partial C/\partial s_2$ and substituting $\sigma = s_1 + s_2(\tilde{x}^* + \underline{\alpha})$:

$$\phi \left[\frac{d\tilde{x}^*}{ds_2}Q + (\tilde{x}^* + \underline{\alpha})g(\underline{\alpha})\left(-\frac{\partial\underline{\alpha}}{\partial s_2}\right) \right] = \lambda \left[U + s_2 \frac{d\tilde{x}^*}{ds_2}Q + \sigma g(\underline{\alpha})\left(-\frac{\partial\underline{\alpha}}{\partial s_2}\right) \right] \quad (23)$$

The parallel with (22) is instructive. The extensive margin terms—those involving $g(\underline{\alpha})(-\partial\underline{\alpha}/\partial s_2)$ —have the same structure as in the s_1 FOC: each new adopter generates $(\tilde{x}^* + \underline{\alpha})\phi$ in abatement and costs the principal σ . The key differences are: (i) the inframarginal cost for s_1 is Q (paying all existing adopters) while for s_2 it is U (paying on all existing usage), and (ii) s_2 has the intensive margin terms $((d\tilde{x}^*/ds_2)Q$ on both sides) that s_1 lacks entirely.

Optimal instrument mix. Eliminating λ between the two FOCs yields a single equation characterizing the optimal mix of s_1 and s_2 .

Recall from above that σ is the total subsidy to the marginal adopter and κ is the inframarginal cost per new adopter. The s_1 FOC (22) is:

$$\phi(\tilde{x}^* + \underline{\alpha}) = \lambda(\sigma + \kappa)$$

Rearranging the s_2 FOC. Expand both sides of (19) and group the intensive-margin terms (those with $d\tilde{x}^*/ds_2 Q$) separately from the extensive-margin terms (those with $g(\underline{\alpha})(-\partial\underline{\alpha}/\partial s_2)$):

$$(\phi - \lambda s_2) \frac{d\tilde{x}^*}{ds_2} Q + [\phi(\tilde{x}^* + \underline{\alpha}) - \lambda\sigma] g(\underline{\alpha})\left(-\frac{\partial\underline{\alpha}}{\partial s_2}\right) = \lambda U$$

Substituting from the s_1 FOC. Both brackets can be replaced using the s_1 FOC. The extensive-margin bracket is immediate: rearranging gives $\phi(\tilde{x}^* + \underline{\alpha}) - \lambda\sigma = \lambda\kappa$. For the intensive-margin bracket, divide the s_1 FOC by $(\tilde{x}^* + \underline{\alpha})$ and expand σ :

$$\phi = \frac{\lambda(\sigma + \kappa)}{\tilde{x}^* + \underline{\alpha}} = \lambda s_2 + \frac{\lambda(s_1 + \kappa)}{\tilde{x}^* + \underline{\alpha}}$$

so $\phi - \lambda s_2 = \lambda(s_1 + \kappa)/(\tilde{x}^* + \alpha)$. Making both substitutions and dividing by λ :

$$\frac{s_1 + \kappa}{\tilde{x}^* + \alpha} \frac{d\tilde{x}^*}{ds_2} Q + \kappa g(\alpha) \left(-\frac{\partial \alpha}{\partial s_2} \right) = U$$

Simplifying. Writing κ equivalently as $Q/[g(\alpha)(-\partial\alpha/\partial s_1)]$, the $g(\alpha)$ factors in the second term cancel, leaving a ratio of derivatives:

$$\kappa g(\alpha) \left(-\frac{\partial \alpha}{\partial s_2} \right) = Q \frac{-\partial\alpha/\partial s_2}{-\partial\alpha/\partial s_1} = Q\gamma\delta(\tilde{x}^* + \alpha)$$

where the last equality uses the derivative ratio derived above. Substituting and dividing by Q :

$$\frac{s_1 + \kappa}{\tilde{x}^* + \alpha} \frac{d\tilde{x}^*}{ds_2} + \gamma\delta(\tilde{x}^* + \alpha) = \frac{U}{Q}$$

Writing $\bar{\alpha} \equiv E[\alpha_i \mid \alpha_i \geq \alpha]$ for the mean taste among adopters, $U/Q = \tilde{x}^* + \bar{\alpha}$. Substituting, rearranging, and dividing by $d\tilde{x}^*/ds_2$:

$$\underbrace{\frac{s_1 + \kappa}{\tilde{x}^* + \alpha}}_{\substack{\text{cost of } s_1 \\ \text{per unit abated}}} = \frac{\overbrace{(\tilde{x}^* + \bar{\alpha}) - \gamma\delta(\tilde{x}^* + \alpha)}^{\text{extensive margin disadvantage of } s_2}}{\underbrace{\frac{d\tilde{x}^*}{ds_2}}_{\substack{\text{intensive margin} \\ \text{advantage of } s_2}}} \quad (24)$$

□

Proof that Proposition 3 holds for the environmental principal. The proof follows the same structure as the proof of Proposition 3 in Section III.2 (Steps A–F). The environmental principal maximizes abatement ϕU rather than social welfare $\int SV_i g(\alpha) d\alpha$. Since abatement is linear in usage—every unit of clean energy generates the same externality reduction ϕ —three objects from the social proof specialize:

- $MSV \rightarrow \phi$: the marginal value of one unit of usage is the externality ϕ , rather than the full social value $(\pi + \phi) + \mu'(x^*)$ which includes private fuel savings and adjustment costs.
- $SV(\alpha) \rightarrow \phi \underline{u}$: the value of the marginal adopter is their abatement $\phi(\tilde{x}^* + \alpha)$, not the full social surplus.
- $\Delta \rightarrow 0$: the gap between marginal and total value vanishes because abatement is linear in usage (no nonlinear μ terms creating a wedge).

The adoption condition and budget constraint are unchanged, so Steps A and C carry over identically.

Step B specialization. Substituting $MSV = \phi$, $SV(\alpha) = \phi \underline{u}$, and $\Delta = 0$ into Ψ from Step B:

$$\eta \phi (s_1^* + \kappa) + 0 - \phi \underline{u} (\bar{u} - \gamma \underline{u}) = 0$$

Dividing by ϕ :

$$\eta (s_1^* + \kappa) - \underline{u} (\bar{u} - \gamma \underline{u}) = 0$$

This is the cross-multiplied form of (5).

Step F specialization. Substituting $MSV = \phi$ and $SV(\alpha) = \phi \underline{u}$ into Terms I–IV of the $A > 0$ condition (13):

- *Term I* = $\eta \phi \gamma \pi' (\pi' / h(\underline{\alpha}^*) + wtp \cdot h'(\underline{\alpha}^*) / h(\underline{\alpha}^*)^2) > 0$ under IHR.
- *Term II* = $\phi \underline{u} [\gamma \pi' \underline{u} - \gamma wtp + wtp \cdot h(\underline{\alpha}^*) (\bar{\alpha} - \underline{\alpha}^*)] > 0$, since $\gamma \pi' \underline{u} - \gamma wtp = -\gamma \mu(x^*) \geq 0$ and $wtp \cdot h(\underline{\alpha}^*) (\bar{\alpha} - \underline{\alpha}^*) > 0$.
- *Term III* = $wtp \cdot \phi [\bar{u} - \gamma \underline{u}] > 0$ (the $(\pi + \phi)$ factor from the social case reduces to ϕ since $\partial[\phi \underline{u}] / \partial \underline{\alpha}^* = \phi$).
- *Term IV* = $-wtp \cdot \eta s_2 \mu'(x^*) \geq 0$, since $\mu'(x^*) < 0$.

All four terms are non-negative under IHR, so $A > 0$ and $ds_1^* / d\gamma < 0$. □

III.4 Non-pecuniary benefits from clean technology

The baseline model assumes $\mu(0) = 0$, making μ a pure switching cost: an agent who uses the clean technology exactly at their taste level ($d_{2i}\theta_i = \alpha_i$) receives no non-pecuniary benefit or cost. We now generalize to $\mu(0) = c \geq 0$, retaining $\mu'(0) = 0$, $\mu'' < 0$, and $\lim_{x \rightarrow \infty} \mu'(x) = -\infty$. The parameter c captures non-pecuniary benefits of clean technology usage such as environmental preferences, social signaling, or intrinsic satisfaction. We assume μ is defined on all of \mathbb{R} with $\mu'' < 0$ everywhere, which ensures the FOC has a unique solution regardless of the sign of π' . Setting $c = 0$ recovers the baseline.

This generalization has two consequences. First, when dirty fuel is more expensive ($\pi' > 0$), the baseline results carry through with minor modifications to the level of WTP and the sign of k . Second, it enables a new case: adoption when clean fuel is more expensive ($\pi' < 0$), driven entirely by non-pecuniary benefits. We treat each in turn.

Dirty fuel more expensive ($\pi' > 0$): generalized results. The FOC $\mu'(x^*) = -\pi'$ depends on μ' , not on $\mu(0)$, so x^* is unchanged. The affine WTP structure is also unchanged:

$$wtp_i = a + \alpha_i \pi' \quad \text{where } a = x^* \pi' + \mu(x^*)$$

Higher baseline WTP: $a \geq c \geq 0$. The constant a is the value of the maximized objective:

$$a = \max_x [x\pi' + \mu(x)]$$

Evaluating at $x = 0$ gives $0\pi' + \mu(0) = c$. Since x^* is the maximizer, $a \geq c$. Economically, even an agent with $\alpha_i = 0$ now receives both the optimized fuel saving and the non-pecuniary benefit c , so their WTP is at least c .

The k proof generalizes. From Section 1.4, $\widetilde{wtp}_i = \delta wtp_i + k$ with $k(\delta) = \tilde{a}(\delta) - \delta a$. We revisit the three-step proof from Section 1.4.

Step 1 ($k(1) = 0$) is unchanged: when there is no distortion, $\widetilde{wtp}_i = wtp_i$ and $k = 0$. Step 3 (strict convexity) is also unchanged:

$$\frac{d^2 k}{d\delta^2} = \frac{-(\pi')^2}{\mu''(\tilde{x}^*)} > 0$$

This expression depends on μ'' and π' , neither of which involves $\mu(0)$. The constant c enters k through $\tilde{a} = \tilde{x}^* \delta \pi' + \mu(\tilde{x}^*)$, but drops out upon double differentiation because c shifts μ by a constant.


Step 2 changes. When $\delta = 0$ the agent perceives zero fuel savings. The FOC $\mu'(\tilde{x}^*) = -\delta \pi'$ becomes $\mu'(\tilde{x}^*) = 0$. Since $\mu'(0) = 0$ by assumption, this gives $\tilde{x}^* = 0$ when $\delta = 0$. Substituting into the definition $\tilde{a} = \tilde{x}^* \delta \pi' + \mu(\tilde{x}^*)$:

$$\tilde{a}(\delta=0) = 0 + \mu(0) = c$$

Since $k(\delta) = \tilde{a}(\delta) - \delta a$, we have $k(0) = c - 0 = c$.

In the baseline ($c = 0$), the three steps gave $k(0) = k(1) = 0$ with strict convexity, implying $k \leq 0$ on $[0, 1]$. With $c > 0$, k is strictly convex with $k(0) = c > 0$ and $k(1) = 0$. The sign on the interior now depends on the slope at $\delta = 1$.

Figure A3: $k(\delta)$ in the quadratic case $\mu(x) = c - x^2$ for different values of c



Notes: The function $k(\delta) = (1 - \delta)[c - \delta(x^*)^2]$ for the quadratic case with $(x^*)^2 = 4$. Blue ($c = 0$): baseline with $k \leq 0$ everywhere. Purple ($c = 2 < (x^*)^2$): k starts positive at c , crosses zero at $\delta^* = c/(x^*)^2 = 0.5$, and returns to zero at $\delta = 1$. Red ($c = 5 > (x^*)^2$): $k \geq 0$ on the entire interval. In all cases k is strictly convex and $k(1) = 0$.

Sign of k when $c > 0$. From Section 1.4, $k'(\delta) = \tilde{x}^*(\delta)\pi' - a$. Evaluating at $\delta = 1$:

$$k'(1) = x^*\pi' - a = x^*\pi' - [x^*\pi' + \mu(x^*)] = -\mu(x^*)$$

If $\mu(x^*) \geq 0$ —the non-pecuniary benefit at optimal usage remains positive—then $k'(1) \leq 0$. Since k is strictly convex ($k'' > 0$), k' is strictly increasing, so $k'(1) \leq 0$ implies $k'(\delta) < 0$ for all $\delta \in [0, 1)$. The function is therefore strictly decreasing on $[0, 1)$ with $k(1) = 0$, giving $k(\delta) \geq 0$ for all $\delta \in (0, 1)$. Marginal distortions reduce WTP by *less* than proportional scaling by δ : $\widehat{wtp}_i = \delta wtp_i + k > \delta wtp_i$.

If $\mu(x^*) < 0$ —switching costs dominate the non-pecuniary benefit at optimal usage—then $k'(1) > 0$, so k turns upward before reaching zero at $\delta = 1$. Since k is convex, starts at $c > 0$, and must dip low enough to return to zero while increasing, it must cross zero at some interior point. There exists a unique $\delta^* \in (0, 1)$ where $k(\delta^*) = 0$: for severe distortions ($\delta < \delta^*$), $k > 0$ as the non-pecuniary benefit dominates; for mild distortions ($\delta > \delta^*$), $k < 0$ as in the baseline.

Quadratic illustration. Let $\mu(x) = c - x^2$, so $\mu(0) = c$, $\mu'(x) = -2x$, $\mu'' = -2$. Then $x^* = \pi'/2$, and the switching cost at optimal usage is $\mu(x^*) = c - (x^*)^2$. We can compute k in closed form:

$$k(\delta) = (1 - \delta)[c - \delta(x^*)^2]$$

This changes sign at $\delta^* = c/(x^*)^2$. When $c \geq (x^*)^2$ —i.e. the non-pecuniary benefit exceeds the switching cost at optimal usage— $\delta^* \geq 1$ and $k \geq 0$ on the entire interval. When $c < (x^*)^2$, k changes sign at $\delta^* = c/(x^*)^2 \in (0, 1)$, matching the general analysis above. [Figure A3](#) illustrates the three cases.

Interpretation. Under severe marginal distortions (δ near 0), the agent perceives nearly zero fuel savings, uses the technology minimally ($\tilde{x}^* \approx 0$), and WTP reduces to approximately c : the non-pecuniary benefit alone. In the baseline ($c = 0$), such an agent has near-zero WTP and drops out of the market. With $c > 0$, the non-pecuniary benefit provides a floor for WTP that marginal distortions cannot erode, because c does not operate through the fuel price channel. This is why k is positive under severe distortions: the floor c exceeds the proportionally scaled WTP δwtp_i when δ is small.

Unchanged results. The following results carry through unchanged when $\mu(0) = c$, because the mechanisms on which they depend do not involve the level of $\mu(0)$.

- **Dampening** (Section 1.5): distortions reduce the WTP response to fuel price instruments ($\partial wtp_i / \partial \tau < \partial wtp_i / \partial \tau$). Unchanged because: the envelope theorem and the inequality $\tilde{x}^* < x^*$ both hold whenever $\pi' > 0$, regardless of $\mu(0)$.
- **Channel 2** (Section 1.6): WTP and abatement both increase in α_i , so distortions exclude the highest-externality agents. Unchanged because: both are affine in α_i with the same-sign coefficient when $\pi' > 0$.
- **Social planner** (Section 1.7): the super-Pigouvian tax corrects δ ; s_1 is needed only for γ . Unchanged because: the algebra showing $SV_i + p = \widetilde{wtp}_i$ at the optimal tax is independent of $\mu(0)$.
- **Instrument-mix FOC** (Section 1.8): the structure of equation (3) is unchanged. It depends on the adoption and usage conditions and the Lagrangian, none of which involve $\mu(0)$.

Channel 1 under δ (Section 1.6) may be affected when $k > 0$. Recall that the distorted elasticity under δ is $|\tilde{\mathcal{E}}(p)| = [p/(p-k)] \cdot |\mathcal{E}((p-k)/\delta)|$ (Section 1.6). When $k < 0$, the prefactor $p/(p-k) < 1$ works against the evaluation-point effect. When $k > 0$, the prefactor $p/(p-k) > 1$ and the two components reinforce each other: the demand curve shifts right, and both the evaluation point and the prefactor increase elasticity. The IHR condition remains sufficient, but is no longer necessary—the result holds more broadly when c is large.

Clean fuel more expensive ($\pi' < 0$): adoption from non-pecuniary benefits. When $c > 0$, agents may adopt even when the clean technology has higher per-unit fuel costs: the non-pecuniary benefit outweighs the fuel premium. This case is relevant when the clean technology is a premium product—for example, an electric vehicle in a market where electricity is more expensive per kilometer than gasoline. The non-pecuniary benefit c captures the satisfaction of owning a clean vehicle (environmental preferences, social signaling, or intrinsic satisfaction), that the owner obtains at the level of usage α_i . We use this EV example throughout to build intuition. Adoption requires $c > p - s_1$; otherwise the non-pecuniary benefit does not cover the technology cost even for the agent with the highest WTP ($\alpha_i = 0$).²⁶

This case differs from the baseline in several important ways. Usage may be zero for some agents (corner solutions), WTP is decreasing rather than increasing in α_i , adoption is determined by an upper threshold rather than a lower one, and both the dampening result and Channel 2 can reverse. We develop each in turn.

Usage and the FOC. The FOC $\mu'(x^*) = -\pi'$ still determines x^* . Since $\pi' < 0$, we need $\mu'(x^*) = -\pi' > 0$. Because $\mu'(0) = 0$ and $\mu'' < 0$, the function μ' is strictly decreasing, so $\mu'(x) > 0$ only for $x < 0$. Thus $x^* < 0$: agents use less clean energy than their taste alone would dictate, because the fuel premium discourages usage. In the EV example, the high cost of electricity relative to gasoline discourages driving the EV, even for agents who prefer it.

Corner solutions. Optimal clean energy usage for agent i is $x^* + \alpha_i$, which is positive only when $\alpha_i > |x^*|$. Agents with $0 \leq \alpha_i \leq |x^*|$ would prefer negative usage, which is infeasible, so they set $d_{2i} = 0$: they adopt but do not use the technology at all. In the EV example, these are buyers who purchase the vehicle for the ownership satisfaction but leave it parked in the garage, driving their gasoline car instead. Their WTP comes entirely from the non-pecuniary benefit evaluated at zero usage:

$$wtp_i = 0\pi' + \mu(0 - \alpha_i) = \mu(-\alpha_i) \quad (\text{corner: } 0 \leq \alpha_i \leq |x^*|)$$

²⁶ $\alpha_i = 0$ has the highest WTP because they obtain the entire non-pecuniary benefit without having to use the device and in turn incur the difference in fuel costs.

For agents with $\alpha_i > |x^*|$, the interior solution gives:

$$wtp_i = (x^* + \alpha_i)\pi' + \mu(x^*) = a + \alpha_i\pi' \quad (\text{interior: } \alpha_i > |x^*|)$$

where $a = x^*\pi' + \mu(x^*) \geq c$ by the same revealed-preference argument as before: a is the maximized value of $x\pi' + \mu(x)$, which is at least $\mu(0) = c$.

WTP is continuous, differentiable, and decreasing. WTP is defined by two different formulas on either side of $\alpha_i = |x^*|$: the concave $\mu(-\alpha_i)$ for corner agents and the linear $a + \alpha_i\pi'$ for interior agents. For the adoption threshold to be well-defined and the demand curve well-behaved, the two pieces must join smoothly. We verify that they agree at $\alpha_i = |x^*|$ in both level and slope.

Continuity. At $\alpha_i = |x^*|$, the interior formula gives:

$$a + |x^*|\pi' = x^*\pi' + \mu(x^*) + (-x^*)\pi' = \mu(x^*)$$

using $|x^*| = -x^*$. The corner formula gives $\mu(-|x^*|) = \mu(x^*)$. Both agree.

Differentiability. The interior slope is $\partial wtp_i / \partial \alpha_i = \pi' < 0$. For the corner slope:

$$\frac{d}{d\alpha_i} \mu(-\alpha_i) = -\mu'(-\alpha_i)$$

At $\alpha_i = |x^*|$: $-\mu'(-|x^*|) = -\mu'(x^*) = -(-\pi') = \pi'$. Both slopes agree, confirming differentiability at the kink.

Since $\mu'' < 0$ implies $-\mu'(-\alpha_i)$ is strictly decreasing in α_i (the corner WTP is concave), and the interior WTP is linear with negative slope π' , WTP achieves its unique maximum of $\mu(0) = c$ at $\alpha_i = 0$ and is strictly decreasing for $\alpha_i > 0$. The agent with $\alpha_i = 0$ has the highest WTP because they receive the full non-pecuniary benefit c without incurring any fuel premium: their optimal usage is zero, so they pay no electricity costs. As α_i increases, the agent needs heavier usage to realize their non-pecuniary benefit, which means paying more of the fuel premium, eroding WTP. In the EV example, the buyer who purchases the car for the status signal but rarely drives it has the highest WTP; the committed daily driver faces the largest electricity bills.

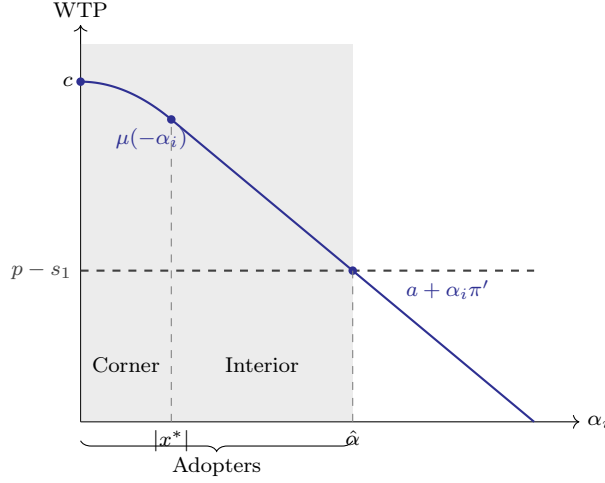
Adoption reverses. Since WTP is decreasing in α_i , the adoption condition $wtp_i > p - s_1$ defines an *upper* threshold $\hat{\alpha}$: agents with $\alpha_i \leq \hat{\alpha}$ adopt. This reverses the baseline, where adoption required $\alpha_i \geq \underline{\alpha}$. Agents who must use the clean technology heavily to obtain the non-pecuniary benefit also incur large fuel premiums that erode willingness to pay. The agents who adopt are those with low to moderate usage requirements to obtain the benefit. In the EV example, it is the casual and status-motivated owners who adopt, not those who would need to be committed daily drivers to obtain the benefit—the opposite of the baseline.

Figure A4 illustrates the WTP function. The concave portion ($\alpha_i \leq |x^*|$) corresponds to corner agents whose WTP comes entirely from the non-pecuniary benefit $\mu(-\alpha_i)$. The linear portion ($\alpha_i > |x^*|$) corresponds to interior agents with WTP $a + \alpha_i\pi'$. The two pieces join smoothly at $\alpha_i = |x^*|$, and the adoption threshold $\hat{\alpha}$ is determined by the intersection with $p - s_1$.

Demand distortions when $\pi' < 0$. The qualitative effects of γ carry over from the baseline, but δ introduces new features because the distortion now operates in the opposite direction: agents underperceive a cost rather than underperceiving a saving.

Intertemporal distortions (γ). As in the baseline (Section 1.4), $\widetilde{wtp}_i = \gamma wtp_i$: the distortion scales WTP proportionally without affecting usage conditional on adoption. The key difference is the direction of selection. Since WTP is decreasing in α_i , the adoption threshold $\hat{\alpha}$ decreases under the distortion, and the agents excluded are the highest- α_i types at the margin. In the baseline

Figure A4: WTP as a function of α_i when $\pi' < 0$



($\pi' > 0$), γ excluded the *lowest*- α_i adopters; here it excludes the highest. This reversal matters for Channel 2 (discussed below).

Marginal distortions (δ): agents underperceive the fuel premium. The agent perceives $\delta\pi'$ instead of π' . Since $\pi' < 0$, the perceived fuel premium is $|\delta\pi'| = \delta|\pi'|$, which is smaller in magnitude than the true premium $|\pi'|$. Define \tilde{x}^* by $\mu'(\tilde{x}^*) = -\delta\pi'$. Since $-\delta\pi' > 0$ but $|\delta\pi'| < |\pi'|$, we have $\mu'(\tilde{x}^*) < \mu'(x^*)$, and since μ' is strictly decreasing, $\tilde{x}^* > x^*$:

$$x^* < \tilde{x}^* < 0$$

The distorted usage margin is closer to zero: under the distortion, agents perceive a smaller fuel premium and consequently use more clean energy than under full information. This is the opposite of the baseline ($\pi' > 0$), where marginal distortions reduce usage.

This also affects corner solutions. Since $|\tilde{x}^*| < |x^*|$, the corner threshold shifts from $|x^*|$ to $|\tilde{x}^*|$: agents with $\alpha_i \in (|\tilde{x}^*|, |x^*|]$ who were at the corner under full information now have positive usage under the distortion. The reduced perceived penalty moves these agents from corner to interior solutions. In the baseline, δ created no such effect because $x^* > 0$ and all agents with $\alpha_i \geq 0$ had interior solutions.

Distorted WTP for interior agents. For agents with interior solutions under both regimes ($\alpha_i > |x^*|$), we can characterize the distorted WTP using the same steps as in Section 1.4. The distorted WTP is:

$$\widetilde{wtp}_i = (\tilde{x}^* + \alpha_i)(\delta\pi') + \mu(\tilde{x}^*) = \underbrace{\tilde{x}^* \delta\pi' + \mu(\tilde{x}^*)}_{\equiv \tilde{a}} + \alpha_i(\delta\pi')$$

Substituting $\alpha_i = (wtp_i - a)/\pi'$ from the undistorted expression:

$$\widetilde{wtp}_i = \tilde{a} + \frac{wtp_i - a}{\pi'}(\delta\pi') = \delta wtp_i + (\tilde{a} - \delta a) = \delta wtp_i + k$$

The k analysis is identical to the generalized baseline: $k(0) = c$, $k(1) = 0$, k strictly convex, with the sign determined by whether $\mu(x^*) \geq 0$. The interpretation differs, however. In the baseline ($\pi' > 0$), the distortion causes agents to underperceive a cost saving, reducing usage; here ($\pi' < 0$), agents underperceive a cost premium, increasing usage relative to the undistorted case. When $k > 0$, the distorted WTP exceeds δwtp_i : the agent's WTP is higher under the distortion than proportional

scaling would suggest, because the non-pecuniary benefit c is unaffected by the distortion while the perceived fuel penalty is smaller.

Distorted WTP for corner agents. Agents who remain at the corner ($\alpha_i \leq |\tilde{x}^*|$) set $d_{2i} = 0$ regardless of the perceived fuel price, so their WTP is:

$$\widetilde{wtp}_i = \mu(-\alpha_i)$$

This is identical to the undistorted corner WTP: marginal distortions cannot affect agents who do not use the technology. The perceived fuel premium is irrelevant when usage is zero. This creates an asymmetry: δ shifts the WTP of interior agents (through k) but leaves corner agents untouched.

Dampening can fail. In the baseline ($\pi' > 0$), marginal distortions unambiguously dampen the WTP response to fuel price instruments (Section 1.5). The envelope theorem gives $\partial \widetilde{wtp}_i / \partial \tau = \delta(\tilde{x}^* + \alpha_i)$ and $\partial wtp_i / \partial \tau = (x^* + \alpha_i)$, and the inequality $\delta(\tilde{x}^* + \alpha_i) < (x^* + \alpha_i)$ holds because both $\delta < 1$ and $\tilde{x}^* < x^*$ work in the same direction.

When $\pi' < 0$, $\tilde{x}^* > x^*$, so the two forces oppose. The distorted usage base $\tilde{x}^* + \alpha_i$ is *larger* than the undistorted base $x^* + \alpha_i$, because the agent underperceives the fuel premium and uses more. Whether the δ scaling is enough to offset this depends on α_i :

$$\delta(\tilde{x}^* + \alpha_i) \stackrel{?}{\lesseqgtr} x^* + \alpha_i$$

For large α_i , the α_i terms dominate both sides and $\delta < 1$ ensures dampening. But for agents near the corner (α_i slightly above $|x^*|$), $x^* + \alpha_i \approx 0$ while $\tilde{x}^* + \alpha_i > 0$, so $\delta(\tilde{x}^* + \alpha_i)$ can exceed $x^* + \alpha_i$.

Quadratic illustration. With $\mu(x) = c - x^2$: $x^* = \pi'/2 < 0$ and $\tilde{x}^* = \delta\pi'/2$. For an agent with $\alpha_i = |x^*| + \epsilon$ (just above the undistorted corner), the usage bases are:

$$\begin{aligned} x^* + \alpha_i &= \epsilon \\ \tilde{x}^* + \alpha_i &= |x^*|(1 - \delta) + \epsilon \end{aligned}$$

where the second line uses $\tilde{x}^* = -\delta|x^*|$ and $\alpha_i = |x^*| + \epsilon$. The dampening condition $\delta(\tilde{x}^* + \alpha_i) < (x^* + \alpha_i)$ becomes:


$$\delta[|x^*|(1 - \delta) + \epsilon] < \epsilon$$

Expanding and rearranging: $\delta|x^*|(1 - \delta) < \epsilon(1 - \delta)$, and dividing by $(1 - \delta) > 0$: $\epsilon > \delta|x^*|$. Dampening fails for agents within $\delta|x^*|$ of the corner. The intuition is that near the corner, the undistorted agent barely uses the technology ($x^* + \alpha_i \approx 0$), so a small increase in τ has almost no effect on their WTP. The distorted agent, by contrast, has a non-trivial usage base ($\tilde{x}^* + \alpha_i > 0$), so the same tax increase generates a meaningful WTP response—the distortion *amplifies* the instrument's effect. [Figure A5](#) illustrates: the distorted response starts earlier (at $|\tilde{x}^*|$ instead of $|x^*|$) and exceeds the undistorted response until α_i is far enough from the corner that the δ scaling dominates.

Channel 2 reverses when $\pi' < 0$. In the baseline ($\pi' > 0$), both WTP and abatement are increasing in α_i (Section 1.6): the agents who value the technology most are also the ones who abate the most, giving a correlation of +1. Distortions exclude these high-value, high-abatement agents from the adopter pool, so the marginal adopter under the distortion has higher α_i and generates more abatement than the marginal adopter absent the distortion.

When $\pi' < 0$, WTP is *decreasing* in α_i (high-taste agents incur large fuel premiums) while abatement $(\tilde{x}^* + \alpha_i)\phi$ remains increasing in α_i for interior agents. The correlation between WTP

Figure A5: Dampening failure near the corner when $\pi' < 0$



Notes: The solid blue line shows the undistorted WTP response to a marginal fuel price instrument ($\partial wtp_i / \partial \tau = x^* + \alpha_i$), which is zero for corner agents ($\alpha_i \leq |x^*|$). The dashed purple line shows the distorted response ($\partial \tilde{wtp}_i / \partial \tau = \delta(\tilde{x}^* + \alpha_i)$), which starts earlier at $|\tilde{x}^*| < |x^*|$ because the distortion moves agents off the corner. For agents with $\alpha_i < |x^*|(1 + \delta)$, the distorted response exceeds the undistorted response: the distortion amplifies rather than dampens the instrument's effect. Illustrated for the quadratic case $\mu(x) = c - x^2$ with $\delta = 0.6$.

and abatement is -1 : the agents who value adoption most (low α_i , who enjoy the non-pecuniary benefit without heavy usage) generate the *least* abatement, while the agents who abate the most (high α_i , heavy users) have the lowest WTP. In the EV example, the status buyers who rarely drive adopt first, but the committed daily drivers who would displace the most gasoline are excluded.

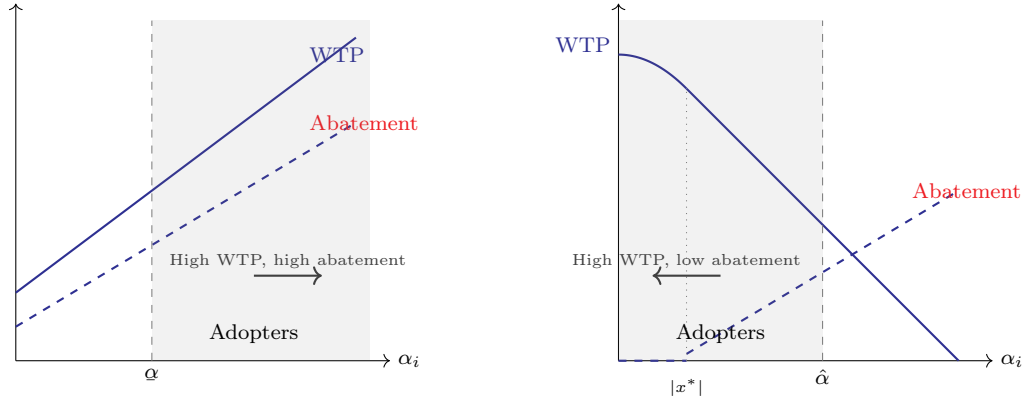
Distortions reduce WTP, lowering the upper threshold $\hat{\alpha}$ and excluding agents at the adoption margin. But these marginal agents—those with the highest α_i among adopters—are precisely the ones who generate the most abatement. The new marginal adopter under the distortion has lower α_i and generates less abatement. Channel 2 now works *against* fixed cost subsidies: a subsidy that brings excluded agents back into the adopter pool reaches marginal adopters who generate less abatement per dollar than they would absent the distortion. This is the mirror image of the baseline result. Figure A6 illustrates the contrast.

Policy implications when $\pi' < 0$. When $\pi' < 0$, fixed cost subsidies are not merely more efficient than marginal instruments—they are necessary. While π' remains non-positive, marginal subsidies generate no adoption incentive through the fuel price channel: the maximum WTP in the population is $c = \mu(0)$ (achieved at $\alpha_i = 0$), which does not depend on s_2 . Adoption therefore requires $c > p - s_1$: the non-pecuniary benefit must exceed the net technology cost.

The case for s_1 rests entirely on Channel 1 (demand elasticity under distortions reduces infra-marginal cost); Channel 2 works in the opposite direction, as the marginal adopter generates less abatement under distortions than without. Whether Channel 1 dominates depends on the distribution of α_i and the severity of the distortions, but the qualitative conclusion is clear: the case for fixed cost subsidies is weaker than in the baseline on the Channel 2 margin.

Regime-flipping. A sufficiently large s_2 can make $\pi' = \pi + \tau + s_2 > 0$, exiting the $\pi' < 0$ regime entirely. At the transition: corner solutions vanish (as $|\tilde{x}^*| \rightarrow 0$ when $\pi' \rightarrow 0$), WTP switches from decreasing to increasing in α_i , and the adoption threshold flips from the upper threshold $\hat{\alpha}$ to the

Figure A6: Channel 2: baseline vs. reversal when $\pi' < 0$
 (a) Baseline ($\pi' > 0$) (b) Reversal ($\pi' < 0$)



Notes: Panel (a): in the baseline, WTP and abatement both increase in α_i . The marginal adopter at α has the lowest WTP and lowest abatement among adopters. Panel (b): when $\pi' < 0$, WTP decreases in α_i while abatement increases. The marginal adopter at $\hat{\alpha}$ has the lowest WTP but the *highest* abatement among adopters. Distortions exclude this high-abatement marginal adopter.

lower threshold α . Once $\pi' > 0$, the full baseline apparatus of Sections 1.4–1.8 applies, with both channels favoring fixed cost subsidies.

This creates a non-convexity in the principal's problem. A moderate s_2 is costly but leaves the economy in the unfavorable regime where Channel 2 works against s_1 . A large s_2 converts the problem to the baseline case where both channels favor s_1 . The principal therefore faces a discrete choice: operate within the $\pi' < 0$ regime using s_1 alone, or spend enough on s_2 to flip the regime. The regime-flipping option may dominate even when s_2 is expensive, because it unlocks the favorable Channel 2 correlation.

A carbon tax is a natural bridge instrument. The tax revenue partially funds s_1 , and the tax itself moves π' toward positive values. Once π' crosses zero, the full baseline apparatus applies without modification. The carbon tax therefore converts the adverse-fuel-price regime into the favorable one, potentially at lower fiscal cost than achieving the same transition through s_2 alone.

III.5 Quadratic adjustment costs

The preceding sections derive all results for a general concave μ . Many of these results simplify dramatically under a specific functional form. We now re-derive the key expressions assuming quadratic adjustment costs:

$$\mu(x) = c - \frac{x^2}{2}$$

where $c \geq 0$. Setting $c = 0$ recovers the baseline of Sections 1.1–1.8; $c > 0$ nests the non-pecuniary benefit generalization of Section III.4. The key primitives under this form are: $\mu(0) = c$, $\mu'(x) = -x$, and $\mu'' = -1$. Cross-references to the general-case sections are provided throughout; the quadratic illustrations already appearing in Sections 1.8 and 1.12 are special cases of the results derived here.

Agent's problem. The FOC $\mu'(x^*) = -\pi'$ becomes $-x^* = -\pi'$, giving:

$$x^* = \pi'$$

The optimal usage excess equals the effective relative fuel price. Each agent's total clean energy usage is $d_{2i}^* \theta_i = x^* + \alpha_i = \pi' + \alpha_i$. Substituting into the WTP expression from Section 1.3:

$$wtp_i = (x^* + \alpha_i)\pi' + \mu(x^*) = (\pi' + \alpha_i)\pi' + c - \frac{(\pi')^2}{2} = \frac{(\pi')^2}{2} + c + \alpha_i\pi'$$

The common component is $a = (\pi')^2/2 + c$: the sum of the optimized fuel saving and the non-pecuniary benefit. WTP is quadratic in the fuel price and linear in taste.

Distorted WTP and $k(\delta)$. Under a marginal distortion δ , the FOC $\mu'(\tilde{x}^*) = -\delta\pi'$ becomes $-\tilde{x}^* = -\delta\pi'$, so:

$$\tilde{x}^* = \delta\pi'$$

The distorted common component is:

$$\tilde{a} = \tilde{x}^* \delta\pi' + \mu(\tilde{x}^*) = \delta^2(\pi')^2 + c - \frac{\delta^2(\pi')^2}{2} = \frac{\delta^2(\pi')^2}{2} + c$$

From Section 1.4, $k(\delta) = \tilde{a} - \delta a$. Substituting:

$$k(\delta) = \frac{\delta^2(\pi')^2}{2} + c - \delta \left[\frac{(\pi')^2}{2} + c \right] = (1 - \delta) \left[c - \frac{\delta(\pi')^2}{2} \right]$$

This closed form makes the sign of k transparent. Three cases arise, depending on the relative magnitudes of c and $(\pi')^2/2$:

1. **Baseline** ($c = 0$): $k(\delta) = -\delta(1 - \delta)(\pi')^2/2 \leq 0$ on $[0, 1]$. Marginal distortions reduce WTP by more than proportional scaling, as in Section 1.4.
2. **Moderate non-pecuniary benefit** ($0 < c < (\pi')^2/2$): k changes sign at $\delta^* = 2c/(\pi')^2 \in (0, 1)$. For severe distortions ($\delta < \delta^*$), the non-pecuniary benefit dominates and $k > 0$; for mild distortions ($\delta > \delta^*$), $k < 0$ as in the baseline.
3. **Large non-pecuniary benefit** ($c \geq (\pi')^2/2$): $\delta^* \geq 1$ and $k(\delta) \geq 0$ on the entire interval. Marginal distortions reduce WTP by *less* than proportional scaling by δ , because the non-pecuniary benefit is experienced in full regardless of the distortion.

These three cases match the general analysis in Section III.4 and are illustrated in Figure A3.

Dampening. Under both distortions, the marginal adopter's WTP response to a dollar of s_2 is $\partial \widetilde{wtp}_i / \partial s_2 = \gamma \delta (\tilde{x}^* + \alpha_i)$ (from Section 1.5). With quadratic μ , the undistorted response is:

$$\frac{\partial wtp_i}{\partial s_2} = x^* + \alpha_i = \pi' + \alpha_i$$

and the distorted response is:

$$\frac{\partial \widetilde{wtp}_i}{\partial s_2} = \gamma \delta (\tilde{x}^* + \alpha_i) = \gamma \delta (\delta \pi' + \alpha_i)$$

The dampening ratio for agent i is:

$$\frac{\gamma \delta (\delta \pi' + \alpha_i)}{\pi' + \alpha_i}$$

This ratio is strictly less than one whenever $\gamma \delta < 1$, and is bounded above by $\gamma \delta$ (tight when α_i is large relative to π' , so the common usage component x^* becomes negligible). When α_i is small, the ratio falls to $\gamma \delta^2$, reflecting the double dampening of the usage component: the agent both underperceives savings (δ) and undervalues them (γ), and usage itself is already reduced ($\tilde{x}^* = \delta x^*$).

Channel 1: demand elasticity under distortions. Since $wtp_i = (\pi')^2/2 + c + \alpha_i \pi'$ is affine in α_i , the shape of demand depends entirely on the distribution $g(\alpha)$, not on the functional form of μ . The Channel 1 results from Section 1.6 therefore carry through without modification. Under γ : $\tilde{\mathcal{E}}(p) = \mathcal{E}(p/\gamma)$, so distortions increase demand elasticity whenever $|\mathcal{E}|$ is increasing in price. Under δ : the elasticity expression $|\tilde{\mathcal{E}}(p)| = p/(p-k) \cdot |\mathcal{E}((p-k)/\delta)|$ is fully explicit because k is now in closed form. The IHR condition on $g(\alpha)$ remains the relevant distributional assumption.

Channel 2: correlation between WTP and abatement. Both wtp_i and abatement $\phi_i = (\tilde{x}^* + \alpha_i)\phi$ are affine in α_i , giving a correlation of +1 for interior solutions. This is exactly the general result from Section 1.6, and the quadratic form does not change it.

Social planner. The first-best requires $\delta \pi' = \pi + \phi$ (the super-Pigouvian tax from Section 1.7). With quadratic μ , the implied policy-inclusive price is:

$$\pi' = \frac{\pi + \phi}{\delta}$$

so the required combination of tax and subsidy is $\tau + s_2 = (\pi + \phi)/\delta - \pi$. The γ correction requires $s_1^* = p(1 - \gamma)$, the same expression as in the general case. When $c = 0$ and $\delta = 1$, no s_1 is needed for the δ correction. The super-Pigouvian tax corrects both margins simultaneously because the agent's underperception operates purely through the fuel price channel, and inflating the price by $1/\delta$ perfectly undoes this at both margins.

Environmental principal: FOCs. From Section 1.8, the intensive margin response is $d\tilde{x}^*/ds_2 = -\delta/\mu''(\tilde{x}^*)$. With $\mu'' = -1$:

$$\frac{d\tilde{x}^*}{ds_2} = \delta$$

Each dollar of s_2 increases every adopter's usage by exactly δ units. When $\delta = 1$, the usage response is one-for-one: $\eta \equiv dx^*/ds_2 = 1$.

The adoption threshold derivatives from Section 1.8 become:

$$\frac{\partial \alpha}{\partial s_1} = \frac{-1}{\gamma \delta \pi'} \quad \frac{\partial \alpha}{\partial s_2} = \frac{-(\delta \pi' + \alpha)}{\pi'}$$

Substituting into the s_1 FOC (1):

$$(\tilde{x}^* + \alpha)\phi = \lambda(\sigma + \kappa)$$

where $\sigma = s_1 + s_2(\tilde{x}^* + \alpha)$ and $\kappa = (p - s_1)/|\tilde{\mathcal{E}}(p - s_1)|$, the same structure as the general case. The s_2 FOC (2) likewise retains the same form, with $d\tilde{x}^*/ds_2 = \delta$ substituted throughout.

Instrument mix equation. Substituting $d\tilde{x}^*/ds_2 = \delta$ into (3), the denominator on the RHS is simply δ :

$$\frac{s_1 + \kappa}{\tilde{x}^* + \alpha} = \frac{(\tilde{x}^* + \bar{\alpha}) - \gamma \delta (\tilde{x}^* + \alpha)}{\delta}$$

The intensive margin advantage of s_2 is exactly δ : each dollar of marginal subsidy generates δ additional units of clean energy use per adopter. The rest of the equation has the same structure as the general case, with the extensive margin disadvantage of s_2 in the numerator.

Formal comparative static: $ds_1^*/d\gamma$ under quadratic costs. We now re-derive the comparative static from Section 1.8, exploiting the quadratic form to obtain an unambiguous sign without any distributional assumption. Set $\delta = 1$ throughout, so $\tilde{x}^* = x^* = \pi'$ and $\eta = dx^*/ds_2 = 1$.

Setup. The marginal adopter's WTP is:

$$wtp(\alpha^*) = \frac{(\pi')^2}{2} + c + \alpha^* \pi'$$

The adoption condition gives $s_1^* = p - \gamma wtp(\alpha^*)$, and the inframarginal cost in hazard rate form is $\kappa = \gamma \pi' / h(\alpha^*)$ (from Section 1.8).

The implicit function Ψ . Substituting s_1^* , κ , and $\eta = 1$ into the cross-multiplied instrument mix equation from Section 1.8:

$$\Psi(\alpha^*, \gamma) \equiv p - \gamma \left[wtp(\alpha^*) - \frac{\pi'}{h(\alpha^*)} \right] - \underline{u} \bar{u} + \gamma \underline{u}^2 = 0$$

where $\underline{u} = \pi' + \alpha^*$ and $\bar{u} = \pi' + \bar{\alpha}$, using $x^* = \pi'$ and $\eta = 1$.

Applying the implicit function theorem. From Section 1.8, Step 5:

$$\frac{\partial \Psi}{\partial \gamma} = \frac{\underline{u} \bar{u} - p}{\gamma}$$

where the simplification uses $\eta = 1$.

Signing the numerator of $d\alpha^/d\gamma$.* From (11) with $\eta = 1$:

$$\frac{d\alpha^*}{d\gamma} = - \frac{\underline{u}^2 - wtp(\alpha^*) \cdot \frac{|\tilde{\mathcal{E}}|-1}{|\tilde{\mathcal{E}}|}}{\partial \Psi / \partial \alpha^*}$$

Under stability ($\partial\Psi/\partial\alpha^* < 0$), the sign of $d\alpha^*/d\gamma$ equals the sign of the numerator. The key step is to show this numerator is positive regardless of the elasticity.

Evaluating $\underline{u}^2 - wtp$. With quadratic μ and $\eta = 1$:

$$\begin{aligned}\underline{u}^2 - wtp(\alpha^*) &= (\pi' + \alpha^*)^2 - \left[\frac{(\pi')^2}{2} + c + \alpha^* \pi' \right] \\ &= (\pi')^2 + 2\alpha^* \pi' + (\alpha^*)^2 - \frac{(\pi')^2}{2} - c - \alpha^* \pi' \\ &= \frac{(\pi')^2}{2} + \alpha^* \pi' + (\alpha^*)^2 - c\end{aligned}$$

When $c = 0$, this is $(\pi')^2/2 + \alpha^* \pi' + (\alpha^*)^2 > 0$. When $c > 0$, the expression $(\pi')^2/2 + \alpha^* \pi' + (\alpha^*)^2 - c$ remains positive at any interior equilibrium where $wtp(\alpha^*) > 0$ (i.e. where adoption is worthwhile), since $\underline{u}^2 = (\pi' + \alpha^*)^2 > wtp(\alpha^*)$ requires only that the marginal adopter's usage squared exceeds their WTP.²⁷

Signing the full numerator. Since $(|\tilde{\mathcal{E}}| - 1)/|\tilde{\mathcal{E}}| < 1$ for all $|\tilde{\mathcal{E}}|$ (and is negative when demand is inelastic), the numerator satisfies:

$$\underline{u}^2 - wtp \cdot \frac{|\tilde{\mathcal{E}}| - 1}{|\tilde{\mathcal{E}}|} > \underline{u}^2 - wtp > 0$$

The first inequality holds because $wtp > 0$ and $(|\tilde{\mathcal{E}}| - 1)/|\tilde{\mathcal{E}}| < 1$; the second is the quadratic-cost result above. Therefore $d\alpha^*/d\gamma > 0$: as distortions weaken, the principal optimally raises the adoption threshold.

Concluding. From (4), $ds_1^*/d\gamma = -wtp(\alpha^*) - \gamma\pi' \cdot d\alpha^*/d\gamma$. Both terms are negative: the direct effect ($-wtp < 0$) and the re-optimization effect ($-\gamma\pi' \cdot d\alpha^*/d\gamma < 0$ since $d\alpha^*/d\gamma > 0$). Therefore:

$$\frac{ds_1^*}{d\gamma} < 0$$

unambiguously. Distortions increase the optimal fixed cost subsidy under quadratic adjustment costs, without any assumption on the distribution of α . This is the key payoff of the quadratic specification: the general result (Section 1.8, Result 3) requires either quadratic μ (condition (a)) or an IHR distributional assumption (condition (b)), but under the quadratic form, condition (a) delivers the result by itself. The uniform illustration in Section 1.8 confirms $ds_1^*/d\gamma < 0$ numerically for a specific parameter set; the result here shows it holds for *any* distribution of α .

Summary. Quadratic adjustment costs buy three things. First, closed-form expressions for x^* , a , $k(\delta)$, and $d\tilde{x}^*/ds_2$ that make the instrument mix equation fully explicit. Second, an unambiguous comparative static $ds_1^*/d\gamma < 0$ that requires no distributional assumption. Third, a unified treatment of the baseline ($c = 0$) and non-pecuniary benefit ($c > 0$) cases from Section III.4, with the sign of $k(\delta)$ depending transparently on the ratio $2c/(\pi')^2$.

²⁷This holds because $\underline{u}^2 - wtp = \underline{u}^2 - a - \alpha^* \pi'$ and $a = (\pi')^2/2 + c \leq (\pi')^2 = (x^*)^2 \leq \underline{u}^2$ whenever $\alpha^* \geq 0$. When $\alpha^* < 0$ but $wtp(\alpha^*) > 0$, the bound still holds because $\underline{u}^2 - wtp = \underline{u}^2/2 + (\alpha^*)^2/2 + \alpha^* \pi'/2 - c/2 + \underline{u}^2/2 - wtp/2$, and the first three terms form a sum of squares after completing the square.

III.6 Marginal distortions also increase demand elasticity

From Section 1.4, $\widetilde{wtp}_i = \delta wtp_i + k$ with $k \leq 0$. Demand at price p is therefore:

$$\tilde{Q}(p) = \Pr(\delta wtp_i + k > p) = Q\left(\frac{p-k}{\delta}\right)$$

Differentiating by the chain rule:

$$\tilde{Q}'(p) = Q'\left(\frac{p-k}{\delta}\right) \frac{1}{\delta}$$

Which gives the elasticity as

$$\tilde{\mathcal{E}}(p) = p \frac{\tilde{Q}'(p)}{\tilde{Q}(p)} = \frac{p}{\delta} \frac{Q'\left(\frac{p-k}{\delta}\right)}{Q\left(\frac{p-k}{\delta}\right)} = \frac{p}{\delta} \frac{\mathcal{E}\left(\frac{p-k}{\delta}\right)}{\frac{p-k}{\delta}} = \frac{p}{p-k} \mathcal{E}\left(\frac{p-k}{\delta}\right)$$

Unlike the γ case, the two components of this expression work in opposite directions. The evaluation point $(p-k)/\delta > p$ increases $|\mathcal{E}|$ (the same force as under γ), but the prefactor $p/(p-k) \leq 1$ works against it. The prefactor arises from the level shift k : every agent's WTP falls by the same absolute amount, but this represents a larger *relative* reduction for low-WTP agents than for high-WTP agents, compressing the demand curve toward the origin.

Whether the net effect increases elasticity depends on the shape of the demand curve. The key condition turns out to be the *hazard rate* of the WTP distribution. Let f denote the density of WTP, F the CDF, and define the hazard rate:

$$h(p) \equiv \frac{f(p)}{1 - F(p)}$$

The hazard rate is the probability of an agent having WTP exactly equal to p , conditional on having WTP at least p . Since $Q(p) = 1 - F(p)$, the demand elasticity can be written:

$$|\mathcal{E}(p)| = \frac{pf(p)}{1 - F(p)} = ph(p)$$

Substituting $|\mathcal{E}(z)| = zh(z)$ into the elasticity expression from above:

$$|\tilde{\mathcal{E}}(p)| = \frac{p}{p-k} \left| \mathcal{E}\left(\frac{p-k}{\delta}\right) \right| = \frac{p}{p-k} \left[\frac{p-k}{\delta} h\left(\frac{p-k}{\delta}\right) \right] = \frac{p}{\delta} h\left(\frac{p-k}{\delta}\right)$$

Now using again the fact that $|\mathcal{E}(p)| = ph(p)$, we can write the condition:

$$\begin{aligned} |\tilde{\mathcal{E}}(p)| &> |\mathcal{E}(p)| \\ \frac{p}{\delta} h\left(\frac{p-k}{\delta}\right) &> ph(p) \\ h\left(\frac{p-k}{\delta}\right) &> \delta h(p) \end{aligned}$$

This holds whenever h is increasing in p : since $(p-k)/\delta > p$, the hazard rate on the left is evaluated at a higher price, and $\delta < 1$ makes the right-hand side even smaller.

An increasing hazard rate means that the higher the price, the more “bunched” the remaining potential adopters are near the margin. When this holds, shifting demand leftward (via $k < 0$) moves us into a region where the remaining buyers are even more concentrated near the cutoff, which more than offsets the $p/(p - k)$ dampening.

This is a well-studied distributional property.²⁸

²⁸Since wtp_i is affine in α_i (Section 1.3), the IHR condition is on the distribution of α_i . Distributions of α_i that satisfy IHR include: normal, uniform, logistic, Gumbel, and Weibull or Gamma with shape parameter ≥ 1 . Distributions that violate IHR include: Pareto (decreasing hazard rate; generates isoelastic demand), lognormal (hazard rate increases then decreases), log-logistic (decreasing or hump-shaped hazard rate), and Weibull or Gamma with shape parameter < 1 . The common feature of the violating distributions is heavy tails: at high α_i , a non-trivial mass of agents remains, so the hazard rate falls. See Bagnoli and Bergstrom (2005) for a comprehensive catalog.

IV Online Appendix Figures and Tables

Table C1: Balance on main treatment variables

	Sample Mean	Credit Treatment	Low BDM Price (<30)	Electricity Treatment	N
	(1)	(2)	(3)	(4)	
Stratification variable: Charcoal spending (USD/month)	19.43 [6.20]	0.23 (0.27)	0.02 (0.27)	-0.04 (0.28)	2134
Stratification variable: Use lpg/gas/ethanol for cooking (=1)	0.58 [0.49]	0.00 (0.02)	0.00 (0.02)	0.00 (0.02)	2134
Stratification variable: Electricity top-up spending (USD, 5 weeks)	4.68 [4.32]	0.39** (0.19)	-0.11 (0.19)	-0.20 (0.20)	2134
Household size	4.95 [1.65]	0.01 (0.07)	0.08 (0.07)	0.03 (0.08)	2134
Age	35.84 [10.15]	0.06 (0.44)	-0.26 (0.44)	-0.18 (0.47)	2134
Male respondent	0.10 [0.30]	0.00 (0.01)	0.02 (0.01)	0.01 (0.01)	2134
Charcoal cooking time (minutes/day)	186.00 [71.08]	-2.02 (3.08)	-3.09 (3.08)	3.00 (3.26)	2134
LPG/ethanol spending (USD/month)	5.23 [10.16]	0.54 (0.44)	-0.49 (0.44)	-0.06 (0.47)	2128
Respondent income (USD/month)	101.44 [115.35]	2.33 (5.00)	1.09 (5.01)	-2.86 (5.29)	2127
Household income (USD/month)	246.99 [204.86]	16.12* (9.61)	16.38* (9.62)	-12.88 (10.18)	1815
Blackout hours (last 30 days)	13.93 [27.49]	2.24* (1.19)	0.50 (1.19)	-0.44 (1.26)	2134
Voltage fluctuation hours (last 30 days)	2.04 [10.88]	-0.70 (0.47)	-0.44 (0.47)	0.20 (0.50)	2134
Joint F-Test		0.09	0.09	0.21	

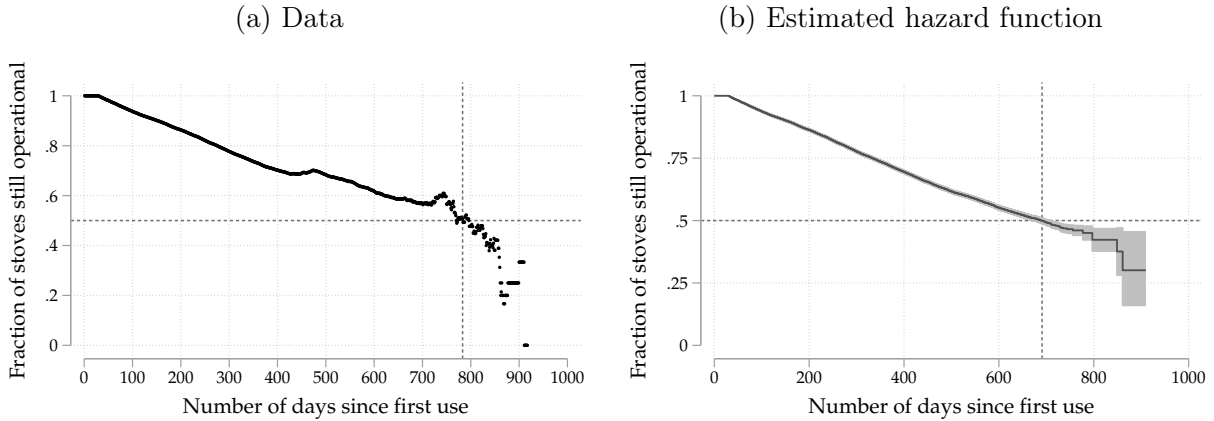
Notes: Test for balance of baseline variables. We only include participants who completed the main survey because they are the only participants whom we assigned to a treatment group.

Table C2: Attrition by treatment assignment and adoption status

	Control Mean	Attrited (Visit 3)
Credit treatment (=1)	0.030	-0.010 (0.007)
BDM price < 30 USD	0.027	-0.005 (0.007)
Electricity treatment > 0	0.021	0.006 (0.007)
Bought induction stove (=1)	0.031	-0.022*** (0.007)
Joint F-Test p-Value		0.00
Sample Mean		0.02

Notes: 2,134 respondents completed the main survey and 2,081 completed the endline survey. Predictors of attrition (defined as not completing an endline visit) among the sample of respondents who completed the main survey.

Figure C1: Stove lifespan estimation



Notes: Stove breakage rates for the 23,444 stoves that had been sold on or before December 31, 2024. The time variable is the number of days between first usage and last usage. A breakage event is the last observed usage if that usage happened more than 3 months before the end of our dataset. We drop the 2% of stoves with less than 24 hours of usage as well as the 3% of stoves with usage between 1 and 29 days since faulty stoves identified within this period are replaced for free by Burn. Median lifespans are marked by the dotted lines. Panel a shows the empirical distribution of survival. Panel b shows the survival curve estimated using a Cox proportional hazard model, according to which the median induction stove sold by Burn has a lifespan of 1.9 years.