Means-Tested Subsidies and Market Power: Evidence from a Heat Pump Program

Paul Dutronc-Postel* Etienne Fize† Aurel Melard‡

September 21, 2026

Abstract

Anticipating supply-side responses to demand-side programs is particularly challenging when the program is means-tested. We study these responses for heat pump subsidies, the largest environmental subsidy targeting low-income French households. Using rich administrative data, we show that poorer households face consumer prices that are discontinuously lower than richer households around eligibility thresholds. This difference in prices is larger than the difference in subsidies by 31%. We develop a model of pricing under imperfect competition to characterize the incidence of means-tested subsidies. The means-test provides firms with information on household income. This combines with market power to generate price discrimination and income-dependent pass-through of the subsidy. To quantify each mechanism's contribution to equilibrium prices and take-up, we leverage linked household and firm data to estimate a structural demand-and-supply model. Counterfactual simulations show that price discrimination increases the share of public spending going to the poorest eligible category by 11% at no efficiency cost.

Keywords: Incidence, Environmental subsidies, Price discrimination, Market power, Energy efficiency.

*IPP (Paris School of Economics), paul.dutronc[at]ipp.eu.

†IPP (Paris School of Economics), etienne.fize[at]ipp.eu.

‡CREST, École Polytechnique, IPP, aurel.melard[at]ipp.eu.

§The authors thank the French National Housing Agency for their help in accessing and navigating the subsidy data, and the Statistical Service of the Ministry of the Environment for their help in accessing tax credit data. We thank Pierre Boyer, Xavier D'Haultfoeuille, Louis-Gaëtan Giraudet, Florian Grosset-Touba, Marion Leroutier, Hugo Molina, Will Rafey, Mathias Reynaert, seminar participants at CREST and PSE, conference participants at EPP, EAERE, MWEET, MCEE, ZEW Public Finance, AFSE for their helpful comments and suggestions. All remaining errors are our own. The authors gratefully acknowledge financial support from the EDF-funded Chaire Développement Durable.

1 Introduction

The ability of demand-side programs to reach their objectives depends on the degree of competition on the supply side. This is particularly the case for means-tested programs (e.g., EITC, housing subsidies or school vouchers), because they may reveal information on consumers' income to imperfectly competitive firms. Specifically, we highlight two ways in which market power can distort the effect of a means-tested subsidy on equilibrium prices. First, information on consumers' income allows firms to resort to price discrimination across income groups. Second, market power lets firms retain part of the subsidy. Under price discrimination, this pass-through is set category by category rather than once for the market as a whole, so only a fraction of the subsidy differential is passed on to relative consumer prices. Both mechanisms impact the take-up and price, so that the policy's distributional effects can go beyond those intended by the subsidy schedule.

We analyze these mechanisms in the context of the largest subsidy program to heat pumps in France, called *Ma Prime Rénov'* (hereafter, MPR). Each year from 2020 to 2024, around 150,000 households received an MPR subsidy for a heat pump of €3,000 to €5,000 depending on income, for a total annual public spending of 600 million euros. This targeted design is not specific to MPR: means-testing has become a standard way of subsidizing large household investments, particularly environmental ones (EV tax credits with income thresholds in the transportation sector, WAP or HEAR in US residential energy programs). Such schemes first circumvent the inability of Pigouvian taxation to reach progressivity (Sallee, 2019; Bonnet et al., 2025). Second, they are able to trigger additional take-up among constrained households (Allcott et al., 2015), while limiting windfall effects at the top (Borenstein and Davis, 2024). In the case of heat pumps, each additional adoption that a subsidy triggers translates into large energy savings and avoided emissions (Bernard et al., 2024).

The ability of a means-tested subsidy to reach its efficiency and equity objectives however depends on how prices react to it. The risk that market power may arise and alter prices is particularly high in markets for environmental goods because of constrained supply of skilled labor (Jagger et al., 2013). Ignoring imperfect competition may therefore bias the evaluation of the trade-off between externality abatement and redistribution of such schemes. Ours is the first paper to document and quantify this bias.

We do this in three steps. First, we build a model of optimal firm pricing under imperfect competition to study how prices react to the introduction of a means-tested subsidy. We derive a necessary and sufficient condition on demand and supply primitives for the consumer price gap to exceed the subsidy gap across income categories. Second, leveraging rich administrative data on MPR beneficiaries and eligibility discontinuities in the subsidy schedule, we identify this consumer price gap in the context of the heat pump market and provide evidence of price discrimination. Third, we leverage additional administrative data on all eligible households and installing firms to estimate a structural model of demand and supply. We build counterfactual policies to assess the effect of discrimination on the equity-efficiency trade-off of the policy.

Our theoretical model extends Weyl and Fabinger (2013) to means-tested subsidies. A continuum of consumers, split into a low- and a high-income group, chooses between two products offered by symmetric firms competing in Bertrand-Nash. Both groups differ in their demand elasticity. Only the low-income group receives the subsidy. *Transaction* prices can differ across groups only if the means-test reveals income categories to firms.[1] When it does, two forces act in opposite directions on relative consumer prices: discrimination itself lowers transaction prices for the more elastic low-income group, while incomplete pass-through of the subsidy, now set group by group, returns part of the subsidy to firms.[2] Which force dominates depends on the elasticity gap between groups and on the curvature of low-income demand. Imperfect competition can therefore either amplify or dampen the redistribution intended by the means-test.

Using exhaustive administrative data on households, subsidies and firms we then show that, in the French heat pump market, means-testing redistributes beyond its schedule: the consumer price gap exceeds the difference in the heat pump subsidy across income categories defined by the policy. Two reduced-form exercises help us identify this result. First, we leverage longitudinal variation: the means-tested scheme was introduced in 2020, when it replaced a flat-rate tax credit. We show that all households paid the same transaction price prior to 2020, regardless of their income category. After the introduction of the means-test, however, high-income households, who see their average subsidy decrease, face higher transaction prices, and low-income households, who see their subsidies increase, face lower transaction prices. Second, we exploit the thresholds between the four different eligibility categories in the post-2020 period, in a regression discontinuity design exercise. Slightly richer households who are discontinuously eligible for *less* subsidy face a *premium* of €494 in transaction price, which corresponds to 31% of the €1,591 difference in subsidy value. This premium averages two cases: a larger premium (57%) when firms file subsidy applications on households' behalf and thus perfectly observe their income category, and a smaller one when households apply by themselves (14%). These effects and differences survive controls for local market conditions, household characteristics and heat pump quality, a pattern we read as evidence of price discrimination.[3]

Our reduced-form estimates identify jointly the effect of price discrimination and of differential pass-through on relative prices. Disentangling them, and measuring the effect of the scheme on price levels and aggregate take-up, requires estimating demand and supply. We do so with a

[1] We refer to prices charged by firms as "transaction prices." "Consumer prices" are what consumers effectively pay after the subsidy is deducted. Under the MPR scheme, the household only pays this "consumer price" as the subsidy is paid directly to the firm by the administration.

[2] Conditional on the demand function not being log-convex. Log-convexity implies consumer pass-through above 100% and increasing in market power, a phenomenon known as "overshifting".

[3] We also show that quality is not discontinuous at the thresholds, therefore the gap in price cannot come from a difference in quality.

nested logit model estimated with micro-moments, building on Berry et al. (2004) and leveraging exhaustive micro data on eligible households, firms, take-up and prices. On the demand side, taste for heat pumps depends on housing characteristics, distance to the installer and price. The price term follows Miravete et al. (2025), whose specification accommodates income effects and allows for flexible combinations of elasticity and curvature, the two primitives that govern the effect of discrimination. On the supply side, firms compete in prices and may charge different prices across income categories for the same product. In counterfactual simulations, we quantify the contribution of price discrimination to the equity-efficiency trade-off of the policy. Income effects are large and push take-up among low-income households below what their housing and preferences alone would imply. Means-testing without discrimination would have closed this gap. Allowing discrimination amplifies the redistribution in take-up operated by means-testing by 27%, with no effect on aggregate prices and take-up, at constant budget. This amounts to shifting an additional 1.3 percentage points of public spending toward the lowest-income category, or 11% of the baseline.

Related literature In this paper, we study the pass-through of a subsidy that, being means-tested, may itself enable price discrimination. Three strands speak to this question, but each considers at most one of these two ingredients. Following Weyl and Fabinger (2013), the first studies the pass-through of a uniform subsidy, which depends on market power and on the convexity of demand: firms may raise prices, as with federal aid among for-profit colleges (Cellini and Goldin, 2014), leave them unchanged, as for the Toyota Prius tax credit (Sallee, 2011), or cut them by more than the subsidy, as for residential solar panels (Pless and van Benthem, 2019). The second studies how means-testing distorts pricing while assuming discrimination away (Polyakova and Ryan, 2021; Decarolis, 2015; Tebaldi, 2025; Fack, 2006): subsidizing low-income consumers tilts the composition of demand toward its elastic segment, which pushes prices down, while incomplete pass-through attenuates the decrease. The third studies the welfare effects of price discrimination itself (Adachi, 2023) without its potential interaction with policies. Studying means-testing and discrimination jointly complements what each strand concludes separately. Means-testing generates an additional channel of incidence: by letting firms discriminate between the elastic and the inelastic parts of the income distribution, it can deliver more than full consumer incidence even when the curvature condition for overshifting fails. Discrimination can amplify the redistribution intended by the means-test, by alleviating rationing in the low-income market and increasing it in the high-income one. Closest to us but in a very specific market, Fillmore (2023) shows that giving colleges access to students' financial aid eligibility leads them to lower tuition for low-income students. We differ in two respects. First, we study third-degree discrimination on income in a standard business-to-consumer market, rather than individual-level discrimination in an auction-like mechanism between colleges competing for students. Second, we separate the price-discrimination channel from the pass-through channel, theoretically and empirically. That decomposition transfers to any means-tested program operating in an imperfectly competitive market.

By applying this theoretical framework to heat pump subsidies, we contribute to the evaluation of environmental policies and their interaction with market power. In the large literature on residential environmental subsidies, we are, to our knowledge, the first to study how such subsidies interact with their supply side. Empirical work on these programs has studied the shortfall between engineering and realized energy savings on the demand side (Allcott and Greenstone, 2024; Fowlie et al., 2015; Gillingham et al., 2013; Chan and Gillingham, 2015; Gillingham et al., 2016; Fowlie et al., 2018), attributing it to rebound effects, distorted investment or overstated theoretical gains.[4] A smaller theoretical literature studies targeting as a way to reach consumers the subsidy would otherwise miss (Allcott et al., 2015; Colas and Reynier, 2024). All assume complete pass-through of the subsidy. We relax this assumption and show it matters: firms retain part of the subsidy, and unequally across income groups.

Imperfect competition is more often included in the study of green vehicle policies, such as fuel-economy standards (Goldberg, 1998) or EV adoption subsidies (Barwick et al., 2024; Muehlegger and Rapson, 2022). Closest to our question, Muehlegger and Rapson (2022) studies a means-tested EV subsidy and finds close to complete pass-through, but does not allow for price discrimination across income groups, despite such discrimination being a known feature of the car market (D'Haultfœuille et al., 2019). Our model and evidence suggest this mechanism can matter quantitatively for the efficiency and equity of such policies.

More broadly, we relate to the literature on the interaction between market power and environmental policy in polluting markets: corrective taxes interact with production already below the competitive level (Buchanan, 1969; Fowlie et al., 2016), and with oligopolists' optimal timing of rent extraction (Sinn, 2008; Van Der Ploeg and Withagen, 2015; Kellogg, 2024; De Cannière, 2025). Rationing enters with the opposite sign in a market that mitigates emissions: imperfect competition holds adoption below its efficient level, and we show that a targeted subsidy can counteract that rationing precisely where it binds the most, among low-income consumers.

The rest of the paper proceeds as follows. Section 2 provides background on the French residential energy efficiency market and the MPR policy. Section 3 describes the data. Section 4 develops a model of incidence of a means-tested subsidy under market power. Section 5 evaluates the effect of the subsidy on relative prices across income categories. Section 6 outlines the structural model, and presents the result of estimation and counterfactual simulations.

[4]See Giandomenico et al. (2022) for a survey.

2 Institutional framework

2.1 Demand-side policies – Energy efficiency programs

In 2020, the French housing agency, (*Agence Nationale de l'Habitat*, hereafter ANAH), implemented the "My Renov' Subsidy" (*Ma Prime Rénov'*, MPR) subsidy scheme to increase the number of energy-efficient residential housing renovations. Homeowners of housing units built more than 15 years ago are eligible for the subsidy for 26 types of investments as of 2021, such as external insulation or heat pumps.[5] The MPR subsidy is targeted toward low-income households, which makes it similar to the US Weatherization Program. MPR subsidy levels are lump-sum amounts that depend on the type of investment chosen and on the income category of beneficiaries. Four income categories are defined based on three thresholds of taxable income. These eligibility thresholds are a function of household size and of whether the household lives in the Paris region or outside of it. They are updated every year with inflation. We restrict the population of interest to owner-occupiers of houses built more than 15 years ago.

For a typical air/water heat pump, "Very Poor", "Poor", "Intermediate" and "Rich" households are eligible for €4,000, €3,000, €2,000 and 0 euros respectively in 2021 - and €5,000, €4,000, €3,000 and 0 euros starting 2022. Figure 1a represents the value of the subsidy whereas Figure 1b represents the realized subsidy rate 2022, that is, the ratio of value of the subsidy to the total price of renovation works as a function of taxable income; the subsidy represents at least 40% of the total price of a heat pump – installation costs included – for "Very Poor" households, 29% for "Poor" households, and about 20% for "Intermediate" households.

The application process works as follows: households typically first visit ANAH local or online services to be informed about eligible actions and subsidy levels. Other types of housing retrofit subsidies coexist with MPR, and can be cumulated with the MPR subsidy. Households then select an eligible firm and apply to MPR. Importantly, households can delegate the application process to the selected firm. This option, called "mandating", gives the firm access to some private information on households, including the amount of subsidy they are entitled to, but not the precise value of their taxable income. This allows firms to back out the income category of their clients, should they be mandated. Upon completion of the work and final invoicing, the agency transfers the subsidy directly to the firm and the household pays the remainder.

From 1999 to 2019, prior to the MPR scheme, an income tax credit was available to all French tax households, named the Ecological Transition Income Tax Credit (*Crédit d'Impôt Transition Écologique*, CITE).[6] Households could claim a percentage of the total amount spent on eligible

[5]From 2020 to 2021, housing units built more than 2 years ago at the time of application were eligible; this was then extended to 15 years since January 2022.

[6]In 2020, the MPR lump-sum subsidy only existed for "Very Poor" and "Poor households". The same year, the CITE tax credit switched away from a flat tax rate and adopted the same lump-sum format as the subsidy, and continued to be available only for "Intermediate" households. Thus, in 2020 the combination of the tax credit and the subsidy closely resembled the ultimate structure of the subsidy starting 2021. "Rich" households lost eligibility for

renovation works carried out on their main residence, up to a cap. Other types of housing retrofit subsidies could also be cumulated with CITE, just like they could with MPR. While the scheme had undergone several reforms since its creation, it was stable in the years immediately preceding the reform, and covered similar types of renovation works as the MPR scheme, including heat pumps. At the time the tax credit disappeared, its rate stood at 30%.

Figure 1: MPR eligibility and rates in 2022

(a) Value of the subsidy (b) Subsidy rates by taxable income level

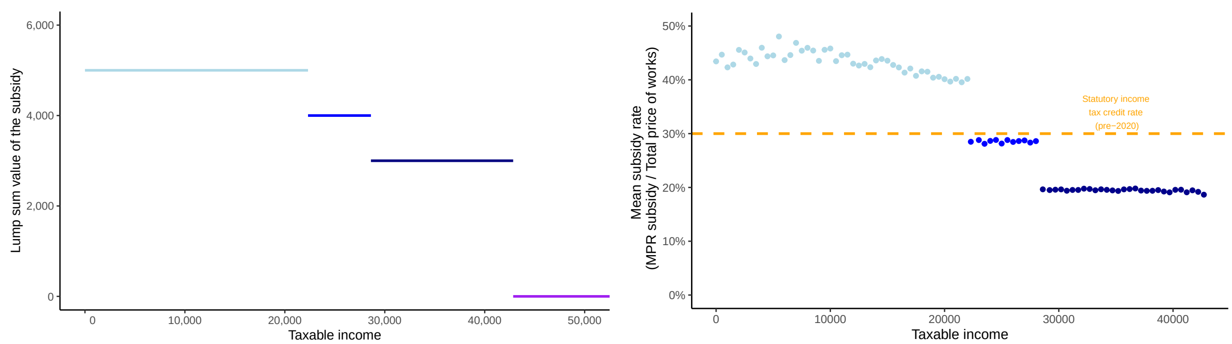

Note: This Figure presents the MPR subsidy schedule for heat pumps in 2022 as a function of taxable income, for 2-persons households living outside of the Paris region. Eligibility thresholds depend on where the household lives (within or outside the Paris region), on household size, and on household taxable income. The four colored categories are those defining eligibility under the MPR scheme: "Very Poor", "Poor", "Intermediate", and "Rich" categories. Panel (a) represents the theoretical value of the subsidy by income category. Panel (b) represents the effective average subsidy rate by bins of €500 of taxable income, in 2022, for heat pumps and for the three subsidized categories. The subsidy rate is defined as the amount of Ma Prime Rénov' subsidy received as a fraction of total cost of renovation work.

2.2 Supply – Certified energy efficiency market

In 2015, a certification, called RGE, was created to guarantee the quality of subsidized retrofit works.[7] Since then, households who wish to benefit from public subsidy programs – including the CITE income tax credit before 2020, and MPR subsidy after – are required to opt for an RGE-certified company. We therefore consider RGE firms as the relevant supply for our analysis.

RGE supply is locally concentrated and shows signs of market power. The density of eligible RGE companies and the concentration of supply, as measured by distance to five closest firms, are heterogeneous across local markets. Figure 2 shows that the operating margin of RGE firms is increasing with local market concentration, controlling for *département* fixed effects – the main administrative subdivision of France (NUTS 3 level). This suggests that companies in concentrated markets are indeed able to price above marginal cost.

There are several reasons for this relatively low level of competition on the certified energy

either scheme for heat pumps, but kept being eligible for insulation renovation works.

[7] RGE stands for "Certified as Guarantor of the Environment" (*Reconnu Garant de l'Environnement*)

Figure 2: Operating margin of certified firms, as a function of local market concentration

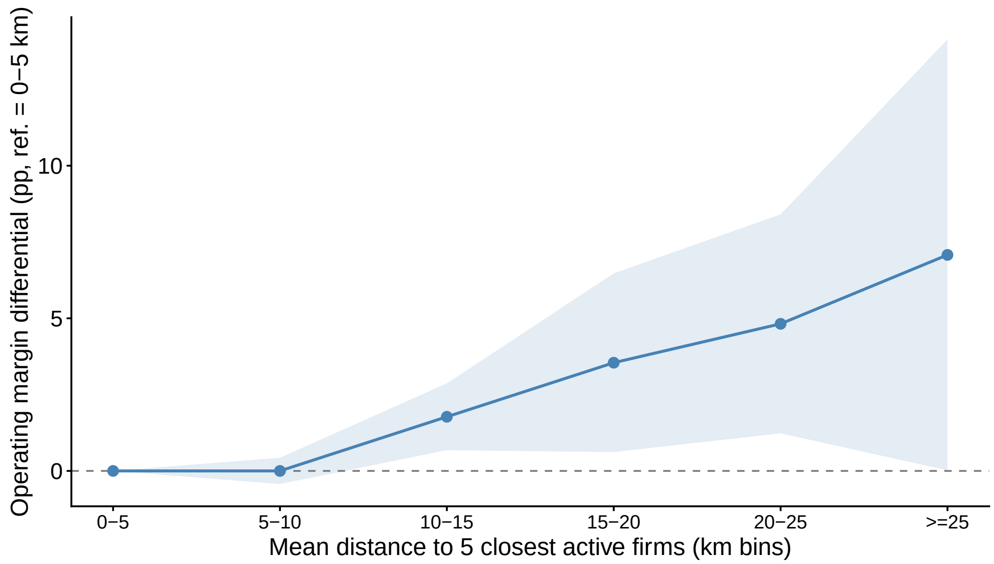

NOTE: This figure presents the result of the regression of operating margin of active certified firms on the concentration of their local market. An active company is a company that has been certified for heat pumps during the past year and was selected by at least one household to perform an MPR work. Concentration is defined at the household level as the mean distance to the five closest active firms. Concentration at the firm level is defined as the average concentration around households that opted for the company. Regression controls for *département* fixed effects. Mean operating margin in the baseline category is 8.8%.

efficiency market. The first one is an institutional barrier to entry: certification is costly and needs to be renewed every four years. It comprises an administrative fee and costs associated with training the firm's craftsmen and monitoring audits made by the certifying body during the four-year span of the certification. The precise annual amount therefore depends on the number of trained professionals and on the certifying body. It is estimated to be around €2,000 for a company where only one professional is trained, or 6% of the median MPR turnover of RGE companies in 2022.[8]

A second competitive constraint on this market is that it is local. Around 50% of heat pumps are installed by firms that are located less than 25 km away.[9]

Finally, entry and capacity might be limited on the retrofit market by constraints on the labor market. On the labor demand side, retrofit companies are indeed in competition with the market for new construction, that recruits similar profiles. On the supply side, it might not be profitable or feasible for employees of the construction sector to seek training to acquire retrofit skills. One piece of suggestive evidence of such entry constraints is that the number of certified firms

[8]Estimated using the 2024 tariff grid of Qualibat, one of the main RGE certification bodies. RGE certification and the associated costs are retrofit action-specific. For example, a company that wishes to enter the subsidized market both for heat pumps and water heaters needs to pay for two certifications. Appendix Figure A.12 shows that in practice, firms mostly focus on a small range of actions.

[9]Appendix Figure A.14 shows the full distribution of household-firm distance for subsidized heat pumps in 2022.

actually decreased between 2019 and 2020, as displayed in Appendix Figure A.13; in 2023, it still lay below its 2019 level.

3 Data

Geolocalized household tax data The Fideli database is produced by the French National Statistics Office (Insee) and is based on tax records from the housing tax and the personal income tax. It is an exhaustive, individual-level dataset of French households and their residence. It contains information on the demographic composition of households, several earnings and income measures —notably taxable income and disposable income— and on the primary and secondary residence of households. Residences are geolocalized, both with physical addresses and geocoordinates. The data are available annually, for years 2017 to 2023. Residence is measured on January 1 of a given year, and annual earnings and income variables are available for the two years before that.

MPR subsidy data These data consist of administrative records from the ANAH, the agency in charge of the MPR scheme. It contains information on household taxable income and composition, which are necessary to compute the level of the subsidy to which the household is eligible, and geolocalization. It also contains information on all retrofit actions carried out by subsidized households, at the level of each retrofit action: type of retrofit action, amount of subsidy and total price paid, identifier of the firm that was chosen to perform the retrofit work. This dataset contains 2.0 million applications between 2020 and early 2024, of which around 150,000 per year heat pump applications.

We match MPR and Fideli data using geographical coordinates, as well as other common variables: type of housing (apartment or house), ownership status (owner- or tenant-occupied housing), number of occupants, household income. This matching procedure has two main objectives: first, retrieve all information on housing characteristics and household composition from Fideli for each MPR application. Second, identify, among all eligible Fideli households, those that take up the subsidy. An MPR application is successfully matched with a Fideli housing unit in 92% of cases. The matching procedure and sample composition are detailed in Appendix B.1.

We also use the firm identifier present in MPR applications in order to match this data with firm-level data.

Firm-level exhaustive data We use the FARE database to obtain information on firms. This dataset is maintained by Insee and is based on various administrative sources – mostly corporate income tax declarations. It gathers information on firms' activity and balance sheets. The latest vintage available so far is 2023. We complete this dataset with the official registry of all

RGE certified firms, which is produced by the Ministry of the Environment. For each firm, this additional dataset contains information on the types of renovation for which it is certified, and the periods of certification. Finally, we use the complete geolocalized registry of French firms produced by Insee in order to locate plants and firms.

CITE income tax credit data We use information on the former tax credit directly from the exhaustive POTE income tax database, in order to document the shift from the single rate tax credit to the means-tested subsidy. The POTE data contains a panelized version of income tax declarations for all French tax households. It contains all items that are filled in a typical tax filing. It notably contains information on the value of each renovation step subsidized via the tax credit.

Energy Efficiency Certificates data The Energy Efficiency Certificate scheme, known as CEE, was introduced by the French government in 2006 to help fund the energy transition. Under this scheme, energy suppliers are required to encourage consumers to reduce their energy consumption, with specific targets to meet over defined periods. Among other variables, CEE data contains information on CEE date, amount, retrofit action type, technical properties of the installed equipment, and information on housing, including location. We use it to document quality of installed heat pumps.

4 A model of incidence for means-tested subsidies

We model heat pump pricing under a subsidy targeted at low-income households and imperfect competition. We consider two cases: uniform pricing and third-degree price discrimination. We show a necessary and sufficient condition for price discrimination between the subsidized and unsubsidized groups to have progressive effects. This condition depends on three demand and supply primitives: demand curvature in the subsidized group, differential in demand elasticities across groups and the level of market power.

The environment We consider a model of unit demand for heat pumps, in a market characterized by imperfect competition. A continuum of mass one of consumers is split into two groups, low- and high-income, respectively denoted l and h. Each of these markets is made of a continuum of consumers that choose between two products $j \in \{1, 2\}$ offered by two separate firms, and an outside option.[10] These firms compete in symmetric Bertrand-Nash. Given prices, aggregate demands in each sub-market are denoted $s_{jm}(p_{1m}, p_{2m}), m \in \{l, h\}$. Note that at equilibrium,

[10]The number of products can be extended at the expense of clarity of exposition, but without any impact on the drivers of the results.

prices will always be equal in the two firms and we denote $s_m(p) = s_{jm}(p) = s_{-jm}(p)$ the demand for any of the products in market m.

We make four assumptions on demand functions and on the cost structure of firms.

Assumption 1. *Firms face constant marginal cost, equal across products and income groups:*

$$\forall m \in \{l, h\}, \forall j \in \{1, 2\} \; c_{jm} = c$$

Assumption 2. *For any level of price, products* 1 *and* 2 *are partial or perfect substitutes:*

$$\forall m \in \{l, h\}, \forall j \in \{1, 2\}, 0 < \frac{\partial s_{jm}}{\partial p_{-jm}}(p_1, p_2)$$
$$\frac{\partial s_{-jm}}{\partial p_{jm}}(p) \leq \left| \frac{\partial s_{jm}}{\partial p_{jm}}(p) \right|$$

Assumption 3. *The curvature of aggregate demand in each sub-market m, denoted $\rho_m(p)$, respects the following condition:*

$$\rho_m(p) \equiv \frac{s_m \left[\frac{\partial^2 s_{1m}}{\partial p_1^2} + \frac{\partial^2 s_{1m}}{\partial p_1 \partial p_2} \right]}{\frac{\partial s_{1m}}{\partial p_1} \left[\frac{\partial s_{1m}}{\partial p_1} + \frac{\partial s_{1m}}{\partial p_2} \right]} < 2$$

Assumption 4. *At uniform price $\overline{p}$ prevailing before the subsidy, demand is more elastic among low-income consumers:*

$$\epsilon_l^{own}(\overline{p}) > \epsilon_h^{own}(\overline{p})$$
$$\epsilon_m^{own}(p) \equiv -\frac{\partial s_{1m}}{\partial p_1} \frac{p}{s_{1m}}$$

Assumption 1 can be relaxed and is used for clarity of exposition of the results.[11] Assumptions 2 and 3 are assumptions on preferences that ensure that the profit maximization problem is well behaved.[12] Following the price discrimination literature, Assumption 4 defines the low-income group as the weak market. This is a testable restriction on preferences: we estimate own-price elasticities by income category in Section 6.3 and find that they decrease with income.

[11] Allowing marginal costs to differ across groups, $c_\ell \neq c_h$, leaves Lemma 2 and the pass-through term of Proposition 1 unchanged, and modifies only the first term of condition (1), which becomes $c_h \frac{\epsilon_h^I(p_h^*)}{\epsilon_h^I(p_h^*) - \theta_h(p_h^*)} - c_\ell \frac{\epsilon_\ell^I(p_\ell^*)}{\epsilon_\ell^I(p_\ell^*) - \theta_\ell(p_\ell^*)}$.

[12] Assumption 2 accommodates the case where products are perfect substitutes as a limit case, which implies perfect competition and $\theta = 0$. Assumption 3 accommodates cases of log-convexity, that requires $\rho_m(p) \geq 1$, and are a necessary and sufficient condition for cost pass-through above 1, i.e. overshifting, following Weyl and Fabinger (2013).

An important feature of this framework is that it is agnostic about the sources of product differentiation: it makes no assumption on why consumers have heterogeneous preferences over suppliers. Geographic proximity, brand loyalty, search frictions, quality differences, and switching costs are all compatible with the model.

Introduction of a means-tested subsidy – no price discrimination We consider the introduction of a subsidy $b > 0$ to the low-income group and assume that each firm prices the two groups uniformly. Since the two firms are symmetric, the unique equilibrium price $\overline{p}$ before the introduction of the subsidy verifies the following market share-weighted average Lerner index:

$$\frac{\overline{p}-c}{\overline{p}} = \frac{1}{\sum\limits_{m\in\{l,h\}} \frac{s_m(\overline{p})}{s_l(\overline{p})+s_h(\overline{p})}\epsilon_m^{own}(\overline{p})}$$

In this baseline case, the more elastic market implicitly subsidizes the less elastic one, a fact already noted by Polyakova and Ryan (2021). The unique transaction price after the subsidy, $\bar{q}(b)$, then solves

$$\frac{\bar{q}-c}{\bar{q}} = \left[\frac{s_\ell(p_\ell)}{s_\ell(p_\ell)+s_h(p_h)}\frac{p_h}{p_\ell}\epsilon_\ell^{own}(p_\ell) + \frac{s_h(p_h)}{s_\ell(p_\ell)+s_h(p_h)}\epsilon_h^{own}(p_h)\right]^{-1},$$

where $p_\ell = \bar{q} - b$ and $p_h = \bar{q}$.

The effect on aggregate prices is *a priori* ambiguous. On the one hand, firms can capture part of the subsidy by increasing prices. On the other hand, the subsidized more elastic group represents a higher share of total consumption, which should push prices down. Ultimately, incidence depends on relative elasticities and the way they vary with price. The effect on relative prices, however, is certain: since listed prices do not differ across categories, the differential in consumer prices cannot be different from the amount of subsidy b.

Adding third-degree price discrimination We now assume that the means-tested subsidy indirectly provides firms with information on the sub-market to which each consumer belongs. Implicitly, it assumes that firms have access to the subsidy level consumers are eligible for, for example through the quotation process. Therefore, firms now systematically perform third-degree price discrimination across the two sub-markets.

There are two effects of the subsidy on prices under third-degree discrimination. First, before taking into account the subsidy, discrimination itself makes firms deviate from initial price $\overline{p}$ separately for each sub-market. The elastic sub-market l stops implicitly subsidizing the inelastic market h and now faces lower price p_l^*, while consumers of h face a higher price p_h^*.

Lemma 1. *Under Assumptions 1 to 4, third-degree price discrimination lowers the price faced by the low-income group and raises the price faced by the high-income group:*

$$p_l^* < \overline{p} < p_h^*$$

Second, the means-tested subsidy is passed through to consumers of the subsidized sub-market l. Because of constant marginal costs of firms, the strong sub-market is unaffected. We denote by $p_l^*(b)$ the consumer price post-subsidy in sub-market l. The following lemma states that consumer pass-through of the subsidy is positive:

Lemma 2. *Under Assumptions 1 to 4, consumer price in market l as a result of subsidy b decreases:*

$$p_l^* - p_l^*(b) = \int_0^b \varrho_l(t)dt > 0$$

where $\varrho_m(t) \equiv \left[1 + \theta_m\big(p_m(t)\big)\big(1 - \rho_m(p_m(t))\big)\right]^{-1}$ denotes consumer pass-through at point t in market m and $\theta_m(p)$ is the conduct parameter that measures market power. Specifically, it measures the share of overall change in demand following a price change that is internalized by a single firm. As such, it is related to the elasticities of own- and industry- demands (ϵ_m^{own} and ϵ_m^I). The latter takes into account both the effects of own and competitors' price change on demand:

$$\epsilon_m^I(p) \equiv -\frac{ds_{1m}}{dp}\frac{p}{s_{1m}} = -\left[\frac{\partial s_{1m}}{\partial p_1} + \frac{\partial s_{1m}}{\partial p_2}\right]\frac{p}{s_{1m}}$$
$$\theta_m(p) \equiv 1 + \frac{\frac{\partial s_{1m}}{\partial p_2}}{\frac{\partial s_{1m}}{\partial p_1}}$$

The result of Lemma 2 comes from noticing that the conduct parameter is always between 0 and 1, and the curvature of demand always below 2. Pass-through at any given t is therefore always positive, *ie* l consumers always benefit from part of the subsidy. Note that it can also be above 1, as soon as demand curvature is above 1, a phenomenon known as "overshifting" (Weyl and Fabinger, 2013). An interpretation for the effect of demand curvature is the following: for a given level of elasticity, a more convex demand means that as a result of a shift in marginal cost (or tax), the firm faces a different customer type, which allows it to pass through a higher share of this change.

In the following proposition, we characterize a condition under which third-degree price discrimination increases the redistributive impact of the means-tested subsidy:

Proposition 1. *Under Assumptions 1 to 4, the redistributive effect of the means-tested subsidy is higher under price discrimination if and only if:*

$$c\left[\frac{\epsilon_h^I(p_h^*)}{\epsilon_h^I(p_h^*)-\theta_h(p_h^*)}-\frac{\epsilon_l^I(p_l^*)}{\epsilon_l^I(p_l^*)-\theta_l(p_l^*)}\right]+\int_0^b \varrho_m(t)dt > b \tag{1}$$

Customers in l (resp. h) benefit (resp. lose) from allowing price discrimination if and only if this condition holds.

The two terms of the left-hand side of this condition correspond respectively to the differential in consumer price between l and h explained by price discrimination alone and pass-through of the subsidy in market l. The first term is derived by noticing that optimal pricing under price discrimination is given by the usual Lerner index:

$$\begin{aligned}\frac{p_l^*-c}{p_l^*}&=\frac{1}{\epsilon_l^{own}(p_l^*)}=\frac{\theta_l(p_l^*)}{\epsilon_l^I(p_l^*)}\\ \frac{p_h^*-c}{p_h^*}&=\frac{1}{\epsilon_h^{own}(p_h^*)}\end{aligned}$$

The two channels of the subsidy under price discrimination, characterized by Lemmas 1 and 2, might have opposite consequences on this price differential. On the one hand, price discrimination in itself has clear progressive consequences toward the more elastic low-income group. On the other hand, as long as there is no overshifting, the second term implies that part of the subsidy differential is captured by firms, thus reducing progressivity. The proposition states that if elasticities are sufficiently different from one another, and consumer pass-through sufficiently close to one, allowing price discrimination on top of the subsidy increases the consumer price of h, decreases the consumer price of l and increases redistribution.

Conditional on the level of market power, the key drivers of the condition in Proposition 1 are the elasticity differential across income groups and the curvature of low-income demand. In the following sections, we show that the French heat pump market satisfies this condition: price discrimination amplifies redistribution, and we estimate the structural primitives that explain why.

5 Reduced-form evidence of price discrimination

The model predicts that, under price discrimination, transaction prices should diverge across income categories following the introduction of means-testing, while remaining homogeneous before. Prices should also jump discontinuously at income thresholds. We exploit both sources of variation in this section. The evidence consistently points toward third-degree price discrimination with respect to income categories, with the transaction price differential representing 30–50% of the subsidy differential across groups.

5.1 Motivating evidence

Figure 3 plots average heat pump subsidy and transaction price by year and income category around the 2020 reform. Before 2020, the scheme offered a homogeneous subsidy and transaction prices were near-identical across income categories. Starting in 2021, a clear pattern emerges: transaction prices are consistently higher for higher-income categories, despite those categories receiving a lower subsidy.[13] This co-movement between the introduction of income-based targeting and the divergence in transaction prices across income groups is a first piece of evidence consistent with price discrimination. This figure also illustrates the inability of a simple reduced-form exercise to disentangle a potential price discrimination channel from the pass through. Price discrimination is likely to affect the transaction prices faced by each eligibility group differently, notably as a function of their demand elasticity. The pass-through channel is likely to affect prices differently in each group as well, as a function of subsidy differentials between groups, but also of demand curvature and of the structure of the supply faced by consumers in each group.

Zooming in on the post-2020 period, Figure 4 plots average transaction prices by fiscal income relative to the two main income thresholds, separately for mandated and non-mandated files. Price differences appear sharply at the thresholds rather than trending smoothly with income, and are visibly larger among mandated files. In mandated files, the firm handles the subsidy application and therefore has direct access to the household's income category at the outset of the quotation process. If price discrimination operates through this information channel, the estimated gap should be substantially larger among mandated files, which is the case as illustrated in Figure 4. However, we do not think of this information transmission as being completely absent in the non-mandated case: indeed, exchanges with ANAH officers suggest that some firms may acquire information on the eligibility levels of their clients through informal discussion during the quotation process. We thus interpret non-mandated and mandated applications as differing in the likelihood that firms obtain that information.

5.2 Estimation strategy

We estimate transaction price differentials at income thresholds using the non-parametric locally linear RDD of Calonico et al. (2014), with bias-corrected and robust inference. Details of the methodology are provided in Appendix B.2. Under standard RDD assumptions, a positive and

[13] In 2020, the Intermediate category remained eligible only to the former tax credit, while the Very Poor and Poor categories gained access to MPR, which may explain the reversed gap in that year. The High Income category lost eligibility for any subsidy starting in 2020.

Figure 3: Heat pump subsidy and transaction price by year and income category

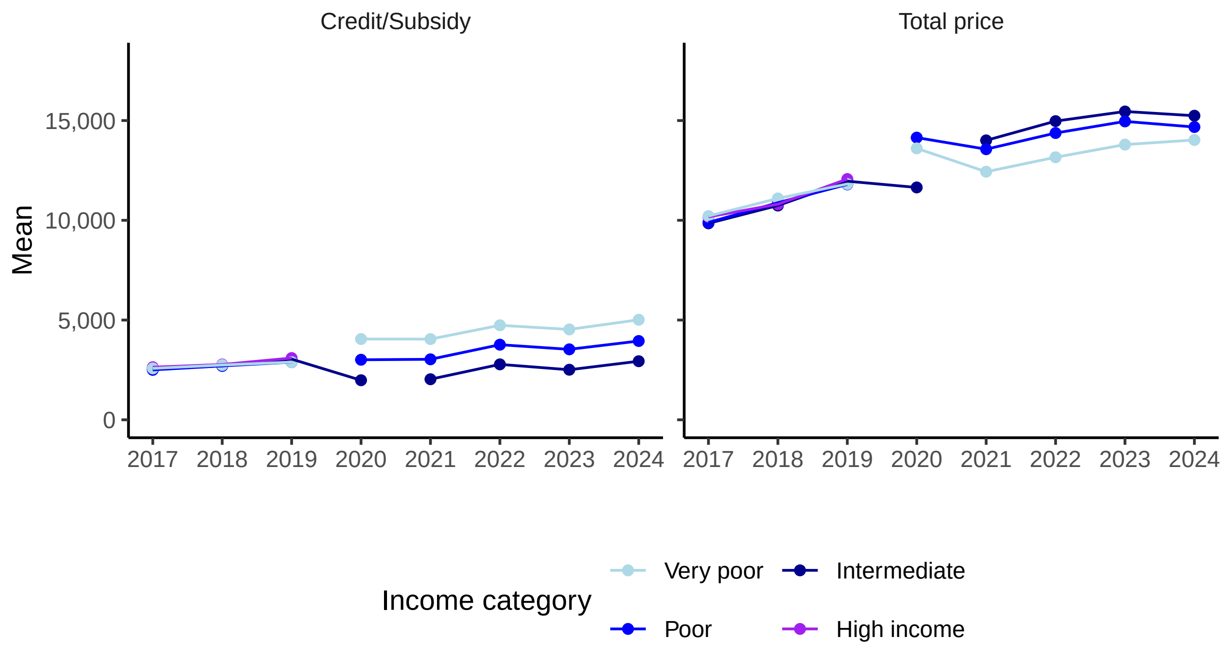

Note: Panel (a) plots the average credit or subsidy received for installing a subsidized heat pump, for four categories of income. Panel (b) plots the transaction price paid by beneficiary households on subsidized heat pump installations. Prior to 2020, these correspond to installations subsidized by the CITE; post 2020, by MPR. The four income categories are the ones defining eligibility for MPR. For 2020 and before, the income categories were not defined. We use the "Very Poor"/"Poor" and "Poor"/"Intermediate" income thresholds that the ANAH was applying in other subsidy schemes (for heavily deteriorated housing, and for disability adaptation). These thresholds follow the IPC inflation rate. We retropolate the Intermediate/high income threshold for years 2017–2020 using the same rate.
In Panel (b), all data are observed, either in tax records, for years 2017 to 2019 and for 2020 for "Intermediate" households, or in ANAH records for years 2020 to 2024. In Panel (a), data for years 2020 to 2024 are observed in MPR records. For years 2017 to 2019, the tax credit is not directly observed; we simulate it based on the observed, declared expenses of Panel (b), applying the 30 % nominal rate and some additional computation rules. Most importantly, a cap applies to the total amount of expenses for which a household is eligible for a tax credit, and this cap has been fixed in nominal terms since 2005. This explains why the average value of the tax credit increases at a slower rate over 2017–2019 than the corresponding total expenses. Source: MPR, CITE.

Figure 4: Heat pump transaction price by fiscal income relative to income thresholds

(a) Non-mandated

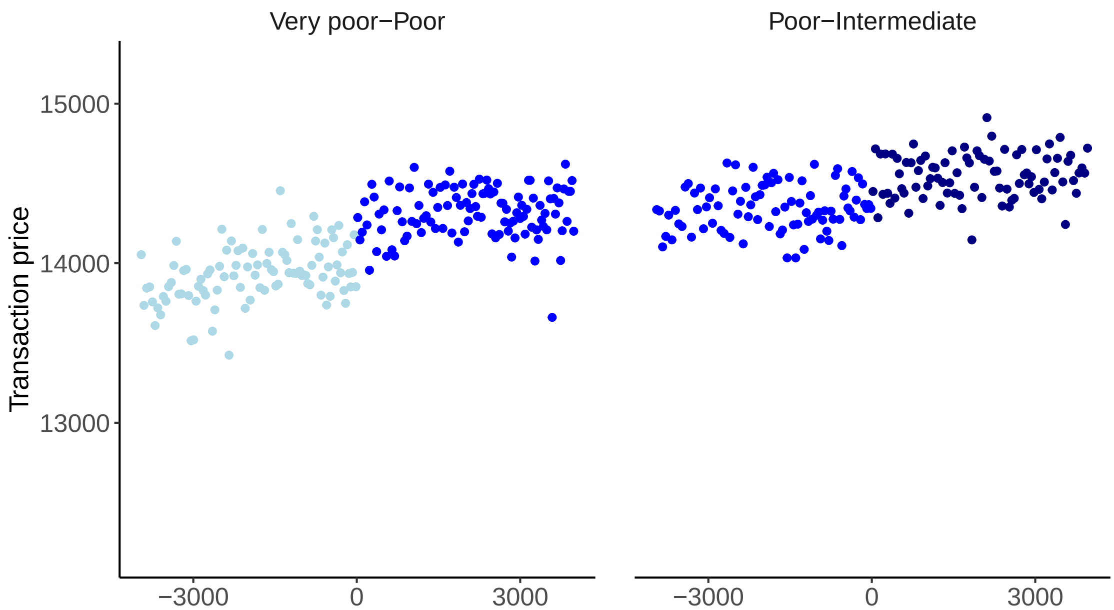

(b) Mandated

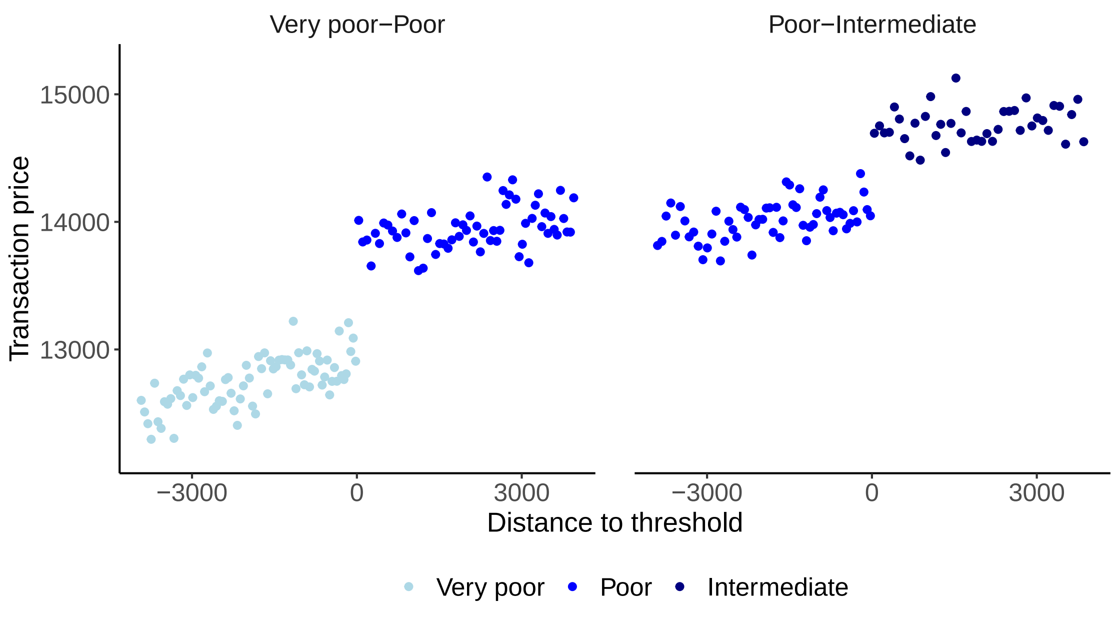

Note: This figure presents a scatterplot of mean transaction price of MPR heat pumps by distance to the Very Poor–Poor and Poor–Intermediate income thresholds. Panel (a) covers non-mandated files and Panel (b) mandated files; both panels share a common vertical scale. Each dot corresponds to 500 observations. Source: MPR.

significant estimate at the threshold is consistent with third-degree price discrimination across income categories. It would imply that the demand and supply primitives of the French heat pump market satisfy the condition of Proposition 1.

This strategy faces two main threats. First, we only observe transaction prices for households that take up the subsidy. Since take-up falls discontinuously on the right of the threshold — the lower subsidy reduces the incentive to invest — the composition of the sample changes at the cutoff (Figure 5). This density discontinuity raises concerns about the causal interpretation of the RDD estimate. Marginal consumers who drop out just above the threshold may face systematically different market conditions or sort toward lower-priced firms. In Appendix Figure A.2, we confirm that some observable characteristics are imbalanced at the threshold: consumers on the right perform fewer co-investments and live in smaller households.

Second, residual consumers on the right of the threshold might choose higher-quality heat pumps. For the observed price gap to reflect quality rather than discrimination, these households — who receive a lower subsidy — would need to systematically select higher-quality equipment, which is unlikely *a priori*.

Figure 5: Number of MPR observations by distance to the MPR income thresholds

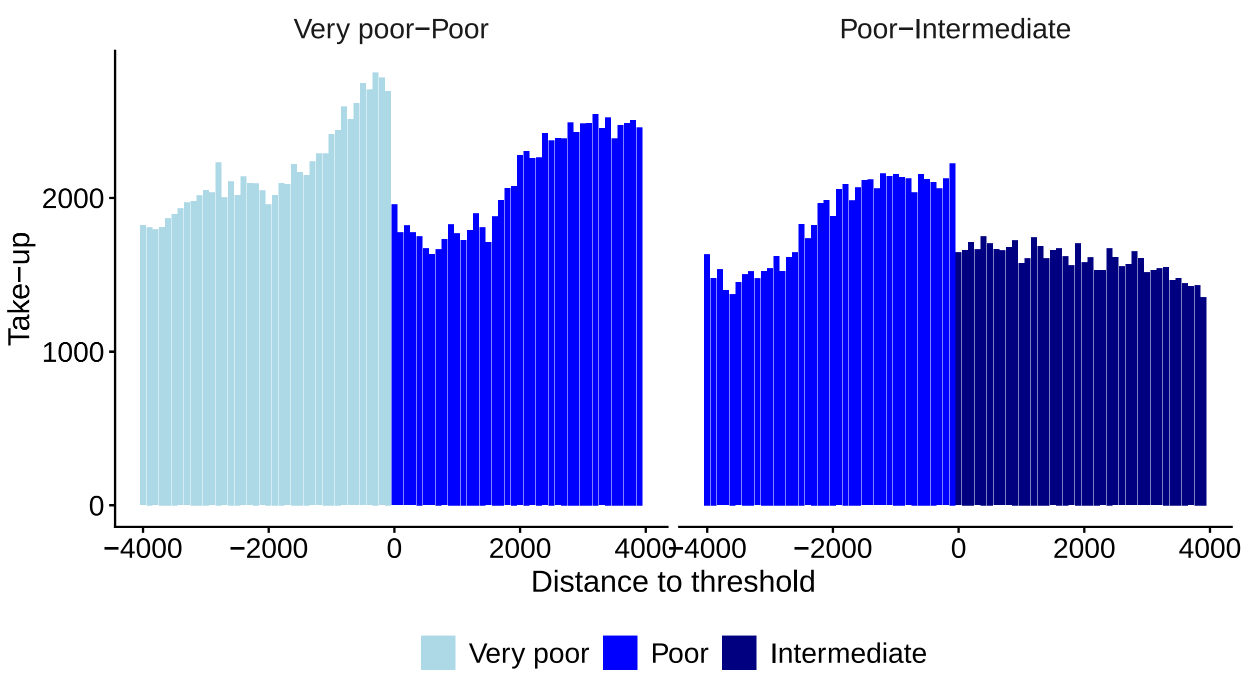

NOTE: This figure presents the number of MPR files by 100-euro bin of distance to the income threshold. An observation corresponds to an MPR heat pump file between 2020 and 2024. Income threshold is normalized to zero. SOURCE: MPR.

To address the first concern, we residualize transaction prices on municipality fixed effects before estimation, removing very local market-level confounders.[14] We then add individual-level controls following Calonico et al. (2019). As that paper establishes, adding controls in the rdrobust estimator improves efficiency under covariate balance but does not restore causal identification when the density is discontinuous; we use it here to verify that observable imbalances do not drive the estimated gap. We present each specification separately for mandated and non-mandated files.

To address the second concern, we apply the same rdrobust estimation to the coefficient of performance (COP) of heat pumps, a quality measure available for the subsample of MPR files matched to CEE administrative data.

5.3 Results

5.3.1 Transaction prices.

Figure 6 displays RDD estimates on transaction prices at both thresholds, separately for non-mandated files, mandated files, and all files pooled, under the three specifications. The pattern is consistent across thresholds: the estimated price differential is larger both economically and statistically among mandated files, where firms systematically access household income information. At the Very Poor–Poor threshold, the municipality-FE estimate reaches approximately €640 for mandated files, compared to €230 for non-mandated files, a factor of roughly three. At the Poor–Intermediate threshold, the mandated estimate is around €450, while the non-mandated estimate is statistically indistinguishable from zero. This heterogeneity supports the interpretation that firm access to income category information is the mechanism enabling discrimination. The partial attenuation with individual controls indicates that some sorting accounts for part of the raw differential, but a substantial gap persists after conditioning on observable characteristics and local market conditions. Pooling the two thresholds, the subsidy falls by €1,591 as a household crosses into the next income category, while the consumer price rises by €2,084, that is, by 131% of the subsidy differential. The residual 31% is the transaction-price premium itself (€494), which represents 57% of the subsidy differential among mandated files and 14% among non-mandated ones.

5.3.2 Quality.

Figure 7 shows that RDD estimates of the coefficient of performance are statistically indistinguishable from zero at both thresholds. The transaction price differentials documented above are therefore not driven by quality differences across income groups. Most firms indeed specialize in a single performance tier and thus cannot discriminate on this margin (Appendix Figure A.4).

[14]The municipality level is very fine grained in France, it is most likely that a local market spans across several municipalities. There are around 35,000 different municipalities in France compared to around 19,500 in the USA

Figure 6: RDD estimates on transaction prices

(a) Very Poor–Poor

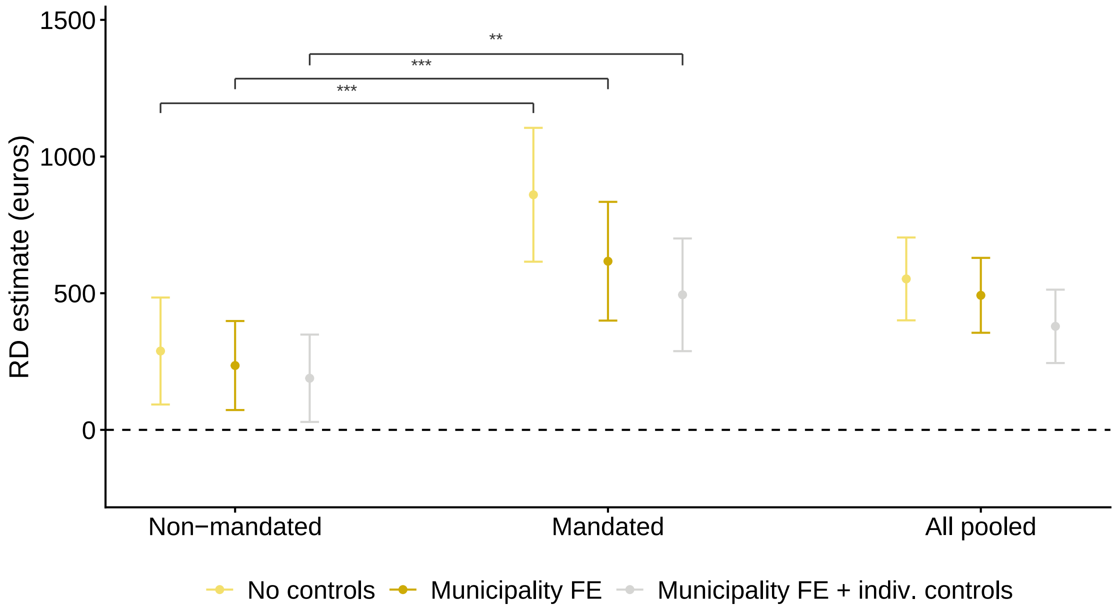

(b) Poor–Intermediate

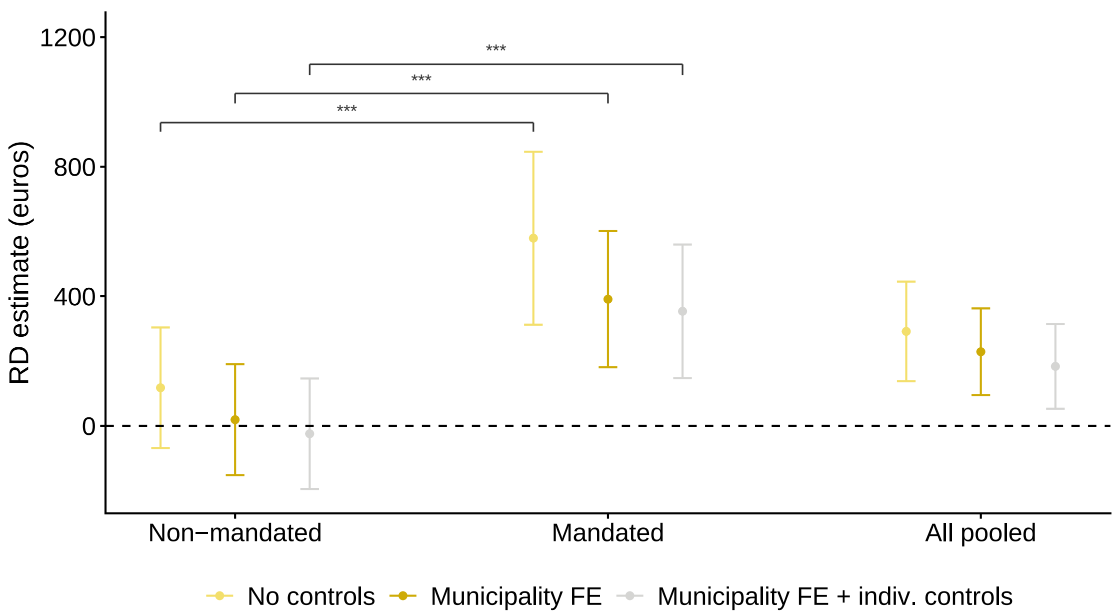

NOTE: This figure presents the estimated transaction price differential at MPR income thresholds, separately for non-mandated files, mandated files, and all files pooled, at the Very Poor–Poor and Poor–Intermediate thresholds. Three sets of estimates are presented for each specification: without controls (No controls), residualized on municipality fixed effects (Municipality FE), and residualized while controlling for individual characteristics (Municipality FE + indiv. controls). Sample is made of one observation per MPR file involving a heat pump. Brackets indicate the significance of the difference between the two point estimates they connect, with stars denoting significance at the 10%, 5% and 1% level. SOURCE: Fideli, MPR.

Figure 7: RDD estimates on coefficient of performance

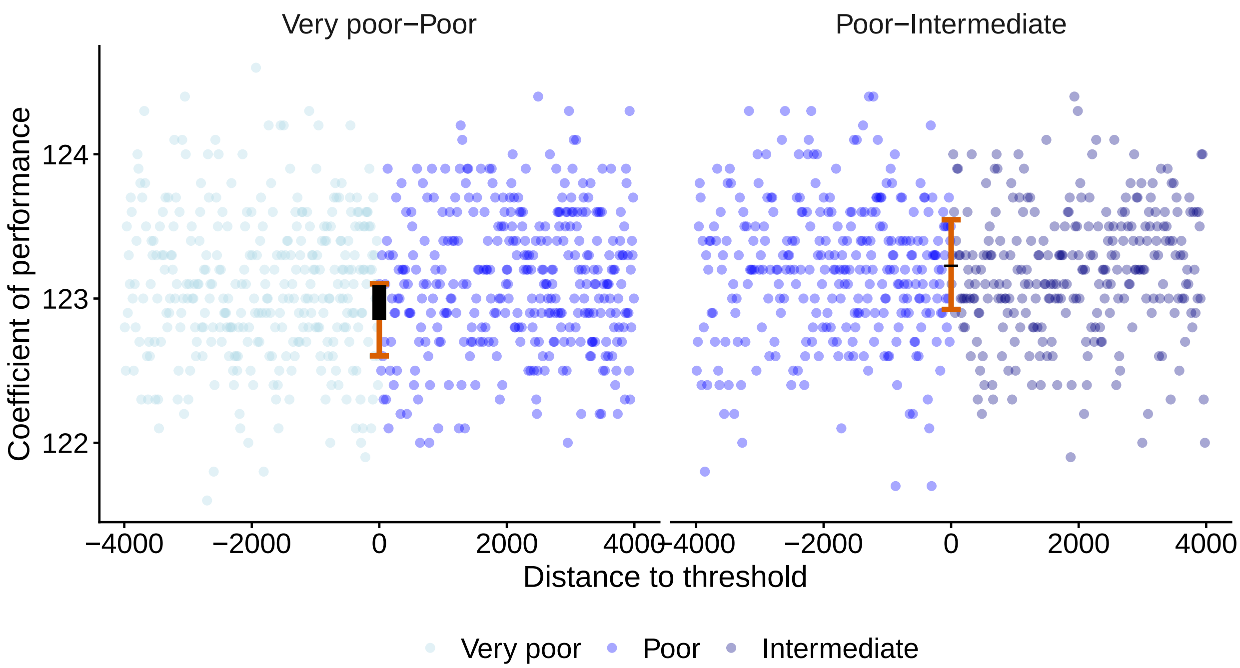

Note: This figure presents a scatterplot of mean coefficient of performance of MPR heat pumps by distance to the Very Poor–Poor and Poor–Intermediate income thresholds, and the estimated differential in performance at each threshold. Each dot corresponds to 100 observations. The RD coefficient is displayed as a black bar with associated 95% confidence intervals. Source: Fideli, CEE, MPR.

5.4 Robustness

5.4.1 Firm fixed effects.

Controlling for firm fixed effects instead of municipality fixed effects reduces the effect but still yields statistically significant estimates (Appendix Figure A.3). The attenuation is nonetheless consistent with part of the discrimination operating through firms sorting selectively into one income category. This specification therefore shows that, even within firms, we still observe price discrimination for firms serving adjacent income categories.

5.4.2 Matched design.

The contrast between mandated and non-mandated files compares two self-selected populations, so the difference between their discontinuities is not by itself an estimand. We therefore match application files on seven predetermined characteristics across the four cells defined by mandate status and by position relative to the threshold, and estimate the mandated vs. non-mandated differential directly, as a difference in discontinuities on common support (Appendix Table A.2).[15] Pooling both thresholds, the transaction-price discontinuity is €645 among mandated files against €157 among non-mandated ones, a differential of €488 (t = 10.4) on cells of 93,715 files each. Matching attenuates the unmatched differential by roughly a quarter, from €667, but leaves it large and precisely estimated. Because mandated files also face a larger subsidy drop at the threshold (€-338, $t = -13.7$), we normalize each segment by the subsidy differential it actually faces: mandated files convert 41% of it into transaction price against 13% for non-mandated ones, preserving the factor of three.

5.4.3 Heterogeneity

Heterogeneity by housing surface (Appendix Figure A.5) yields consistent positive transaction price gaps at the thresholds, among households with more comparable home sizes: selection based on housing size is therefore unlikely to drive the results. Heterogeneity by local market concentration provides suggestive evidence that the price gap increases with concentration (Appendix Figure A.7). Extending the analysis to other renovation actions finds positive price differentials in several cases and negative ones in others (Appendix Figure A.6), consistent with the theoretical result that the direction of the effect depends on market-specific primitives.[16] Across

however the USA has five times more inhabitants than France.

[15] Matching is done within a caliper on the number of occupants in the household and the household fiscal size, the standing of the housing unit, the date of construction, total surface, the date at which the subsidy is awarded, a Paris region vs. elsewhere indicator.

[16] Importantly, we do not have data at hand that document the quality of other renovation actions, which prevents us from running the same analysis as we do on heat pumps. Quality being a crucial adaptation margin to take into account, we only include results on other actions' transaction prices as robustness checks.

actions, the mandated/non-mandated asymmetry recurs, reinforcing the information channel.

5.4.4 Bounding composition effects.

Households just above and just below a threshold need not be identical, and any residual difference in characteristics translates mechanically into a price difference. We bound that channel in two ways. We first estimate a hedonic regression of the unit heat pump transaction price on the same seven predetermined characteristics, with municipality fixed effects (Appendix Table A.3, column 1). An additional square meter raises the transaction price by €9.4, a heat pump in a dwelling built one year later is €3.4 cheaper, and an additional room adds €26. Multiplying each gradient by the corresponding covariate discontinuity at the threshold and summing gives a total composition effect of €13 (s.e. 12), or 2.9% of the €450 discontinuity estimated on the same sample, with the largest individual terms largely offsetting one another. Secondly, using a non-parametric variant, which reweights the eight province $\times$ household-size strata by their estimated discontinuity in sample share rather than pricing individual characteristics, gives €4.8, close to 1% of the estimate. Composition therefore accounts for at most a few percent of the discontinuity we attribute to discrimination.

Consumers in the same municipality, with similar observable characteristics, choosing heat pumps of equivalent quality, face substantially different transaction prices depending on income category, and this gap is systematically larger when firms access household income information. While the density discontinuity at the threshold prevents a fully causal interpretation, the convergence of evidence across specifications, thresholds, and subsamples is consistent with third-degree price discrimination enabled by policy design.

The observed gap is, however, the joint outcome of price discrimination and differential pass-through of the subsidy and reduced-form estimates alone cannot separate these two channels.

6 Structural model: price discrimination under means-testing

This section disentangles the two channels, so as to measure the effect of the means-test, and of the price discrimination it enables, on price levels and aggregate take-up. Section 6.1 defines and estimates a discrete-choice model of demand and supply that recovers the primitives needed to separate them: substitution patterns identify product differentiation and market power, price sensitivity parameters the distribution of elasticities and the curvature of demand, and firms' first-order conditions their marginal costs and markups. Section 6.2 builds counterfactuals isolating the contribution of price discrimination to prices and take-up, at constant public spending. Section 6.3 reports the results.

6.1 The model

Demand for heat pumps We consider M local markets $m \in \{1, ..., M\}$, in which consumers $i \in \{1, ..., n_m\}$ choose among a set of $J_m + 1$ products $\{0, ..., J_m\} \equiv \mathcal{J}_m$, where each $j > 0$ corresponds to a heat pump installer and $j = 0$ is the outside good – not purchasing any heat pump. $J = \sum_m J_m$ is the total number of inside products across markets.

We assume that firms only observe the income category of each consumer i, denoted $c(i) \in \{1, 2, 3\}$ and thus offer category-specific prices p_{cjm} for good j and income category c.[17] p_{cjm} should be understood as consumer price after income category-specific subsidies b_c. Transaction price is denoted q_{cjm}, so that $p_{cjm} = q_{cjm} - b_c$. We therefore assume that price discrimination is universal: firms systematically know the income category of consumers. This assumes away the distinction between mandated and non-mandated applications, on which part of our reduced-form exercises rely. We come back to the reasons why we make this assumption, and to its implications, in Section 6.4.

Consumers choose the product that maximizes their indirect utility. Indirect utility derived from a given product $j > 0$ is heterogeneous across consumers of a given income category and given by V_{ijm}:

$$U_{ijm} = \delta_{c(i)jm} + \alpha \frac{(I_i - p_{c(i)jm})^\lambda - 1}{\lambda} + \sum_{r=1}^{R} \beta_y^r y_i^r + \beta_d d_{ijm} \tag{2}$$

$$V_{ijm} = U_{ijm} + \epsilon_{ijm}, \epsilon_{ijm} = \zeta_{i1}(\sigma) + (1 - \sigma)\tilde{\epsilon}_{ijm}, \forall (i, j, m) \tag{3}$$

where U_{ijm} denotes the deterministic part of indirect utility and ϵ_{ijm} an idiosyncratic demand shock, extreme-value distributed, with an individual-specific term common to all inside products. δ_{cjm} are mean utilities for product j in income category c. They decompose as $\delta_{cjm} = X_{cjm}\theta_1 + \chi_{cjm}$, where X_{cjm} collects product characteristics, and χ_{cjm} is the product-income-category demand residual. y^r are consumer-specific characteristics. d_{ijm} corresponds to the distance between the product's firm and consumer's housing. The price sensitivity term $\alpha \frac{(I_i - p_{c(i)jm})^\lambda - 1}{\lambda}$ follows the specification proposed by Miravete et al. (2025). It allows for flexible demand elasticity-curvature combinations, the two demand primitives that drive the effect of price discrimination as is apparent in 4. It is also a flexible representation of the income effects that households face. As renovation investments can represent a substantial part of annual disposable income, such income effects might be material.[18]

[17] The fourth income category is not eligible for an MPR subsidy for a heat pump investment. We therefore exclude it from the analysis.

[18] In particular, if $\lambda = 0$, it yields the Berry et al. (1995) specification $\alpha \log(I_i - p_{c(i)jm})$. It implies strong liquidity constraints: a household will never opt for a heat pump if its price equals or exceeds disposable income. Conversely, if $\lambda = 1$, preferences are quasi-linear and there is no income effect.

This specification of demand yields the typical nested logit form for individual choice probabilities s_{ijm}:

$$s_{ijm} = \underbrace{\frac{\exp(\frac{\tilde{U}_{ijm}}{1-\sigma})}{\sum_{j'=1}^{J_m} \exp(\frac{\tilde{U}_{ij'm}}{1-\sigma})}}_{s_{ijm}|in} \cdot \underbrace{\frac{\left(\sum_{j'=1}^{J_m} \exp(\frac{\tilde{U}_{ij'm}}{1-\sigma})\right)^{1-\sigma}}{1+\left(\sum_{j'=1}^{J_m} \exp(\frac{\tilde{U}_{ij'm}}{1-\sigma})\right)^{1-\sigma}}}_{s_{im}^{in}} \tag{4}$$

$$\tilde{U}_{ijm} = U_{ijm} - U_{i0m} = U_{ijm} - \alpha \frac{I_i^{\lambda} - 1}{\lambda} \tag{5}$$

where $s_{ijm}|in$ denotes the "within" choice probability, *ie* the probability for i of choosing j conditional on choosing any inside product of market m, and s_{im}^{in} denotes the probability of choosing any inside product.

Heat pump supply In each market m, F_m firms f propose a set of goods $j \in \mathcal{J}_{fm}$ (such that $\bigcup_f \mathcal{J}_{fm} = \mathcal{J}_m \backslash \{0\}$) to three categories of consumers c at price $p_{cjm} = q_{cjm} - b_c$, where $b_c > 0$ is the subsidy level for income category c and q_{cjm} is the transaction price. A product corresponds to a firm-establishment combination. For simplicity of the estimation, we make the assumption that all firms are single-product firms.[19] Firms therefore maximize the following profit function:

$$\pi_{jm} = \sum_{c=1}^{3} (p_{cjm} + b_c - mc_{cjm}) s_{cjm}(p_m) \tag{6}$$

Note that we allow *a priori* for marginal cost to vary for the same product across income categories.[20] There are therefore as many first-order conditions as marginal costs to estimate. However, in the choice of moments used for estimation, we assume that expected marginal cost is equal across income categories.

$$FOC_{cjm} : \frac{\partial \pi_{jm}}{\partial p_{cjm}} = s_{cjm} + \frac{\partial s_{cjm}}{\partial p_{cjm}} (p_{cjm} + b_c - mc_{cjm}) = 0 \tag{7}$$

Inverting the system of first-order conditions allows us to recover estimated marginal costs

[19] In practice only 3% of our observations are multi-establishment firms. We find that relaxing this assumption has no effect on the signs and magnitude of our results.

[20] Since we allow for price discrimination, setting transaction (pre-subsidy) prices q_{cjm} is equivalent to setting consumer prices p_{cjm}.

and markups in the heat pump market. Details and assumptions of this estimation can be found in Appendix C.2.

Data and parameters We estimate the above model on metropolitan France in 2022, the year with highest take-up of MPR for heat pumps. The construction of the estimation sample, that restricts the Fideli-MPR matched dataset to 20 high-take-up markets and to the eligible population excluding extreme income values, is detailed in Appendix C.1. A market m corresponds to a *département*. We focus on the $M = 20$ *départements* with most heat pumps subsidized by MPR and installed in 2022. Within each market, a consumer i is a housing unit eligible for an MPR subsidy for a heat pump.[21] We further restrict our analysis to houses and owner-occupiers. We also exclude some household types that appear to be outliers in terms of take-up and could bias estimates.[22] Details of sample selection can be found in Appendix C.1. A product j is a firm that was chosen at least once by consumers of the market during the year.[23]

Product characteristics entering X_{cjm} include: the age of the company (years since founding), its size proxied by the number of FTE employees, market fixed effects (one per *département*), and income-category fixed effects. Individual characteristics are dummies for construction period (before 1949, between 1949 and 1975, between 1975 and 2000, after 2000), the share of winter days spent in extremely cold temperatures, housing quality index, housing surface, and income. The income linear parameter in indirect utilities should be thought of as a proxy of any determinant of take-up that is unobserved and correlated with income. It is measured by standard of living. Except for construction dummies and income, all product and individual characteristics are normalized to avoid scale effects in estimation.

Identification challenges The first identification challenge concerns potential correlation between firms' characteristics and residual demand shocks. We assume exogeneity of firms' characteristics, especially firm size and location, in the short run. One might be concerned about the fact that firms adjust capacity and size to demand, and self-select into some markets by adjusting location. We assume that in the short run, both characteristics are constrained and do not adjust to demand idiosyncratic shocks. Regarding endogeneity of price, we use appropriate instruments for price variations in our choice of moments, as detailed in the estimation procedure below.

A second challenge for identification is the large number of products with low market shares – the median firm performs two installations per year and market. For those firms, the mean indirect utility of the product would be estimated imprecisely. We deal with this by aggregating products with similar characteristics in the estimation into groups of products $g \in 1, ..., G_m$,

[21] These are owners of a principal residence, built before 2007, who are part of the first three income categories.

[22] Bottom and top 5% of the income distribution, houses below 25m^2, or with very low or very high standing.

[23] Note that we therefore allow firms located outside of the *département* m bounds to be in the choice set $\mathcal{J}_m$, as long as they were chosen by at least one housing in m.

$\sum_{m=1}^{M} \sum_{c=1}^{3} G_m = 3 \cdot G$. Groups are defined by geographical proximity – 10 km radius hexagons – and by firm characteristics. The precise methodology is detailed in Appendix C.2.

Estimation procedure The set of parameters to be estimated is $\theta^d = \{\theta_1^d, \theta_2^d\}$, where θ_1^d is the vector of mean utility parameters and $\theta_2^d = \{\beta_y^T, \beta_d, \alpha, \lambda, \sigma\}^T$ captures taste heterogeneity across consumers and the nest parameter. We estimate this model in a two-step BLP procedure following Berry et al. (2004). The main aspects are explained below; the precise estimation procedure is detailed in Appendix C.2.

For any given value of θ_2, we recover via a nested fixed-point algorithm the unique vector $\hat{\delta}(\theta_2)$ such that predicted market shares match observed shares, that is $s_{cjm}(\hat{\delta}(\theta_2), \theta_2) = s_{cjm}^d, \forall c, j, m$, where:

$$s_{cjm}(\delta, \theta_2) = \frac{1}{n_m} \sum_{i, c(i)=c} s_{ijm}(\delta, \theta_2) \tag{8}$$

We then estimate θ_2^d by GMM using moments defined below. As a second step, given $\hat{\delta}(\hat{\theta}_2^d)$, we estimate θ_1^d by OLS of $\hat{\delta}$ on X.

For the first step of the estimation, we retain three sets of moments. First, as is typical in the micro-BLP literature and following Conlon and Gortmaker (2025), we retain the mean of some individual (and product) characteristics among individuals who take up an inside product to compute micro-moments of the type:

$$h_k^{\mu} = \frac{1}{N} \sum_{m=1}^{M} \sum_{i=1}^{n_m} \sum_{g=1}^{G_m} z_{igm}^{\mu,k} (s_{igm} - \mathbb{1}_{igm}) \tag{9}$$

where $\mathbb{1}_{igm}$ is a dummy equal to 1 if individual i chose group of products g and $N = \sum_{m=1}^{M} n_m$ is the total number of eligible housing across markets. Targeted means include individual characteristics $y_i^r, r \in [1, ..., R]$, distance between housing and product d_{ijm}, as well as standard of living, below and above the median of eligible housing within each income category. Because price enters mean utility δ_{cgm} only as a cell-specific constant, the saturated demand-side inversion absorbs any variation in price across (c, g, m) cells, and cross-category price differences induced by the subsidy schedule carry no information about α, β on their own. What identifies price sensitivity is instead the dispersion in disposable income among takers within a cell: since eligibility is assigned on reference taxable income while utility depends on standard of living (equivalized disposable income), households in the same eligibility category – and thus facing the same price – differ in disposable income, and this within-cell variation, interacted with price, identifies the curvature of the price term. The level of α is in turn disciplined jointly with the supply-side moments, through the markup equation. This variation exists because we use standard of living, i.e., equivalized disposable income, as the relevant measure for income in the utility maximization problem on the demand side, while eligibility is based on reference taxable income, a different concept of income which is essentially gross income. Because of their unequal exposure to taxes and benefits, households with exactly the same disposable income may be part of as many as all four different eligibility levels.[24]

The second set is made of supply-side moments targeting equal mean marginal costs across income categories for the same product group:

$$h_c^S = \frac{1}{\sum_{m=1}^{M} G_m} \sum_{m=1}^{M} \sum_{g=1}^{G_m} (mc_{cgm} - mc_{c+1,gm}) \tag{10}$$

A third set consists of two aggregate IV moments that interact exogenous instruments z^Z for the within-nest share $s_{cgm|in}$ with the demand residual $\chi_{cgm} = \delta_{cgm} - X_{cgm}\theta_1$:[25]

$$h_k^Z = \frac{1}{3G} \sum_{m=1}^{M} \sum_{c=1}^{3} \sum_{g=1}^{G_m} z_{cgm}^{Z,k} \chi_{cgm} \tag{11}$$

These moments identify the nest parameter σ, which governs within-market substitution between inside products. The two instruments ($K^Z = 2$) are the share of group g's full-time equivalents weighted by inverse mean distance of firms to eligible housing, and the share of eligible housing within 50 km of the firm to the number of competing firms in the same geographic hexagon. Both are constructed from 2021 firm characteristics, predetermined with respect to 2022 demand residuals.

In total, we have $K^\mu = 14$ micro-moments, $K^S = 2$ supply moments, and $K^Z = 2$ aggregate IV moments to estimate the 11 parameters of θ_2. We define $H = \begin{pmatrix} h^Z(\delta(\theta_2), \theta_2) \\ h^S(\delta(\theta_2), \theta_2) \\ h^\mu(\delta(\theta_2), \theta_2) \end{pmatrix}$ the column vector of these 18 moments. Our GMM estimator of θ_2^d is therefore:

$$\hat{\theta}_2^d = \arg\min_{\theta_2} H^\top W H \tag{12}$$

where W is the optimal weighting matrix as defined in Conlon and Gortmaker (2025). For

[24]Figure A.9 in Appendix C.1 represents the distribution of eligibility categories by levels of standard of living in 2022.

[25]Relevance and exclusion restriction of the two instruments are discussed in Appendix C.2.

inference, the two steps are stacked into a single GMM system. The OLS normal equations $X^\top(\delta - X\theta_1) = 0$ are treated as additional aggregate moment conditions, so that standard errors for both $\hat{\theta}_1^d$ and $\hat{\theta}_2^d$ correctly account for the sampling uncertainty of the other step (see Appendix C.2).

6.2 Counterfactuals

In our counterfactuals, we isolate the contribution of price discrimination to the effect of the means-tested subsidy on overall and relative take-up and prices. To do so, we depart from the observed schedule, by removing successively the possibility for firms to price discriminate – one single transaction price per product –, the subsidy differential across firms – one single subsidy level per income category – and the difference in income distribution across income categories. This last counterfactual is used to motivate the existence of a means-testing policy: absent financial constraints and related income effects in take-up, what would be the distribution of take-up? How much of the observed income-driven differential in take-up is compensated by the means-tested subsidy?

For all counterfactuals, the average subsidy level is adjusted so that total budgetary spending is kept constant and equal to the observed one. The simulation procedure of prices and subsidy levels in all counterfactuals can be found in Appendix C.4. These exercises make two assumptions. First, we assume constant marginal costs within each income category. In particular, it implies constant returns to scale.[26] Second, we assume that the outside option $j = 0$ does not react to policy changes, in terms of relative prices.

6.3 Results

Table 1 displays the main estimated demand parameters. The main parameters of interest for our analysis are the ones driving elasticities: α, λ and σ. In particular, it appears that some income effects are at play ($\lambda < 1$), although they are less strong than what would be implied by the price sensitivity specification of Berry et al. (1995) ($\lambda > 0$). The value of the nest parameter implies that inside products – all heat pump providers – are relatively highly substitutable options for customers, although not perfectly. Finally, one can note a strong negative contribution of housing-firm distance to indirect utility of inside products, confirming that markets are indeed local and allowing for local market power of firms.

Taken together, these parameters imply own- and aggregate elasticities of take-up with respect to price that are given in Figure 8. As expected, at any given price, the Very Poor and

[26] Constant marginal cost of heat pump installation might seem reasonable for small-scale changes in quantities. An interpretation of our assumption is that the changes in demand implied by the counterfactual policies are small enough for firms not to change their size and cost structure.

Intermediate categories, i.e., the eligible categories with the lowest and highest income, are respectively the most and least elastic groups.

On the supply side, and as implied by the moments we leverage for estimation, we find marginal costs that are very close to one another on average and close to €5,000, as Figure 9 shows. These marginal costs are found to be below subsidy levels on average, which implies markups above 100% and rationalizes own-product elasticities below 1. The presence of market power also results in incomplete pass-through of subsidies within each income category, as Figure 10 shows.[27] In particular, consumer pass-through is smallest among the Very Poor income category and around 85% at observed prices. Two forces are thus at play, in opposite directions. Firms capture part of the subsidy: within each income category, a one-euro increase in b_c lowers the consumer price by less than one euro ($\partial p_{cgm}/\partial b_c \approx 0.85$). Yet transaction prices q_{cgm} are themselves lower for Very Poor households, so the gap in out-of-pocket prices p_{cgm} between income categories is larger than the corresponding gap in subsidies b_c.

[27]Figure 9 shows that Very Poor households face higher markups $\mu_{cgm} = (p_{cgm} + b_c - mc_{gm})/p_{cgm}$ — a margin per euro paid out of pocket, not a Lerner index — despite lower transaction prices $q_{cgm} = p_{cgm} + b_c$. A higher subsidy b_c effectively acts as a cost reduction from the firm's perspective. Consumer price p_{cgm} does not fall proportionally, leaving a larger gap between effective cost and consumer price.

Table 1: Demand parameter estimates

	Parameter	Std-err
Price sensitivity		
α	0.159^{*}	0.075
λ	0.701^{**}	0.082
Individual characteristics		
Built 45–74	0.326^{**}	0.023
Built 75–2000	-0.091^{**}	0.023
Built >2000	-0.014	0.028
Share extr. cold days	0.138^{**}	0.030
Standing	-0.068^{**}	0.008
Surface	0.213^{**}	0.007
Income	0.059^{*}	0.024
Firm characteristics		
Firm age	-0.233^{**}	0.025
Firm size (FTE)	0.303^{**}	0.050
Firm-housing characteristics		
Distance	-1.096^{**}	0.040
Nest parameter		
σ	0.340^{**}	0.101
Mean utility		
$\overline{\delta}$ (cat. 1)	-7.021^{**}	1.152
Cat. 2 − Cat. 1	-0.042	0.057
Cat. 3 − Cat. 1	-0.208^{*}	0.087

Note : Coefficient estimates and standard errors from the micro-BLP demand model. Best run selected as the minimum-GMM estimate among runs with 20 markets. $^{*}p < 0.05$, $^{**}p < 0.01$.

Figure 8: Simulated demand elasticities by income category

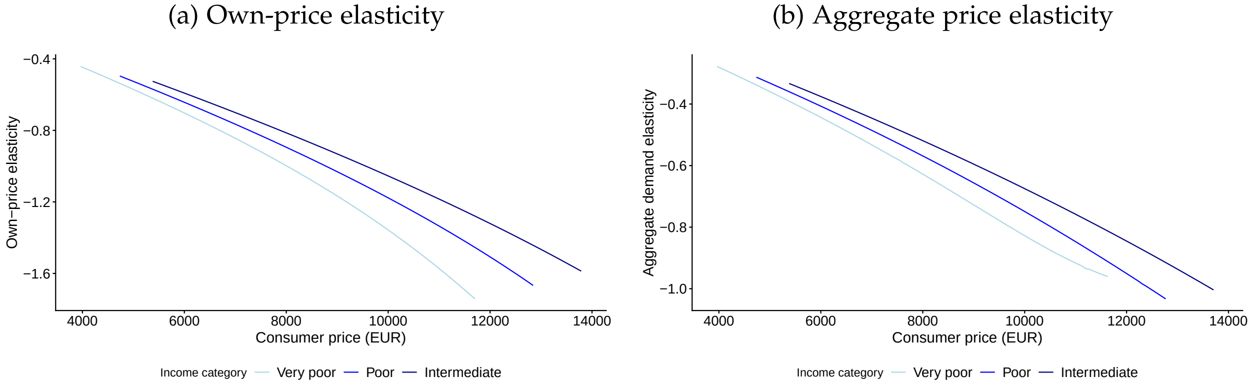

NOTE: This figure presents mean own- and aggregate price elasticities by price level and income category. Elasticities are obtained by simulating counterfactuals in which marginal costs are shifted exogenously, prices are re-optimized and corresponding take-up is computed. Own-price elasticities are computed as weighted means across products, weight corresponding to the product's market share. SOURCE: Fideli, MPR.

Figure 11 shows simulated consumer prices, take-up, take-up rates, and estimated CO_2 savings across the four counterfactual scenarios and all three income categories.

As a benchmark, the first column equalizes both income and subsidy levels across income categories. Consumer prices are near-identical across groups (€7,456 for Very Poor, €7,558 for Intermediate, €7,602 for Poor households). In these conditions, Very Poor households exhibit the highest take-up rate (3.6%), above Poor (3.4%) and Intermediate (3.1%) households, reflecting housing characteristics and preferences that make heat pump adoption slightly more desirable for this group.

Restoring income heterogeneity while keeping subsidies equal – column 2 – reverses the take-up rate pattern. Income constraints reduce Very Poor take-up from 12,700 to 11,700 units, while Intermediate households, facing lower liquidity constraints, expand from 17,700 to 19,000 units. Consumer prices are largely unaffected at this stage. These patterns are reminiscent of the income tax credit period, with no means-test and no detectable differences in transaction prices (see Figure 3).

Introducing means-tested subsidies – column 3 – more than offsets income effects for low-income households. Higher subsidies reduce the Very Poor consumer price by €1,650 (22%), restoring and exceeding their baseline take-up (13,300 units, above the 12,700 homogeneous benchmark). Meanwhile, the reduced subsidy for Intermediate households raises their consumer price by €1,600 (21%), bringing their take-up back below the benchmark. The Poor category, which receives a medium subsidy, is moderately affected throughout.

Allowing firms to price discriminate across income categories – last column – amplifies redistribution further, while leaving aggregate take-up virtually unchanged (40,400 units in total). Price discrimination reduces the Very Poor consumer price by an additional €575 and raises the Intermediate price by €340, adding 520 units to Very Poor take-up and removing 420 from Inter-

Figure 9: Distribution of marginal costs and markups, by income category

(a) Marginal costs

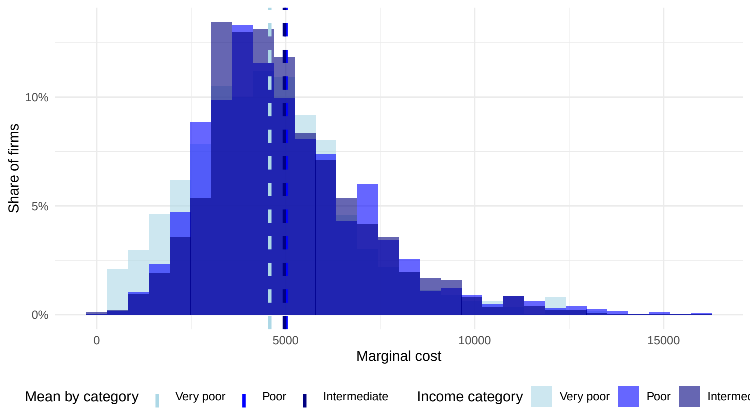

(b) markups

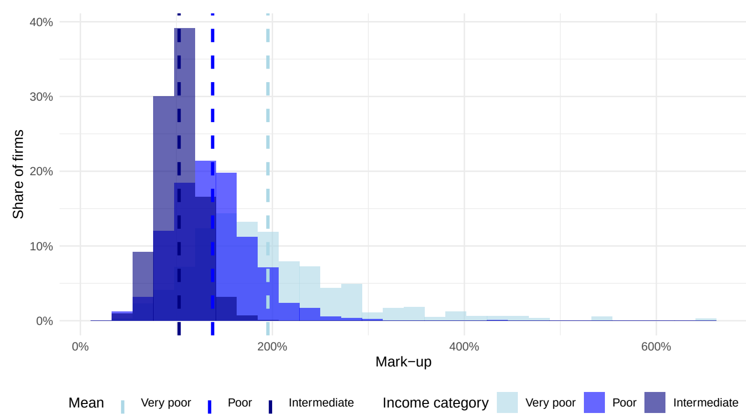

NOTE: This figure presents the distribution of marginal costs and markups by income category, as estimated by the micro-BLP model. An observation corresponds to a product-income category combination, weighted by its predicted market share. SOURCE: Fideli, MPR.

Figure 10: Simulated pass-through by income category

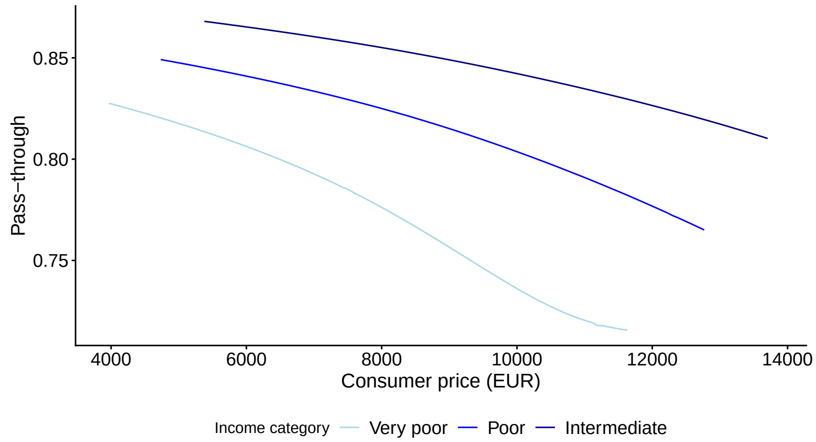

NOTE: This figure presents simulated marginal cost pass-through by price level and income category. Pass-through is obtained by simulating counterfactuals in which marginal costs are shifted exogenously by €100 and prices are re-optimized. Pass-through is the ratio between the resulting change in price and the change in marginal cost. SOURCE: Fideli, MPR.

mediate take-up compared to column 3. Means-testing closes the take-up gap between the Very Poor and Intermediate categories by around 3,460 units, and price discrimination closes it further by 940 units, which we interpret as an increase in redistribution of the policy by 27%. Strikingly, this redistribution comes at no efficiency cost: total take-up, and therefore estimated CO_2 savings, are virtually identical across all four scenarios (we assume around 1,800 kt CO_2 avoided over the heat pump lifetime, using 2.21 tCO_2/year per heat pump over a 20-year lifetime; see Appendix C.6 for methodology). If anything, there is a slight uptick in aggregate take-up following price discrimination. Its effect on environmental efficiency is however negligible: at a constant budget of €263 million, it corresponds to an additional 7 tCO_2 abated per million euros of public spending. This reflects the balance between two forces: price discrimination allows firms to expand take-up the most in the most elastic market, where subsidies are more additional; but Very Poor households simultaneously represent a larger share of total take-up, and each inframarginal heat pump is more costly for public finances in that category. Both effects balance each other, leaving the number of heat pumps financed per million euros virtually unchanged at 154 across all scenarios. The same shift can be read on the budget. At constant public spending, price discrimination raises the share of the MPR budget accruing to Very Poor households from 41.6% to 42.9%, a gain of 1.3 percentage points, which represents a further 11% of the 13.7 points that means-testing itself redistributes toward that category.

Figure 11: Consumer prices, take-up and CO_2 savings by income category and counterfactual scenario

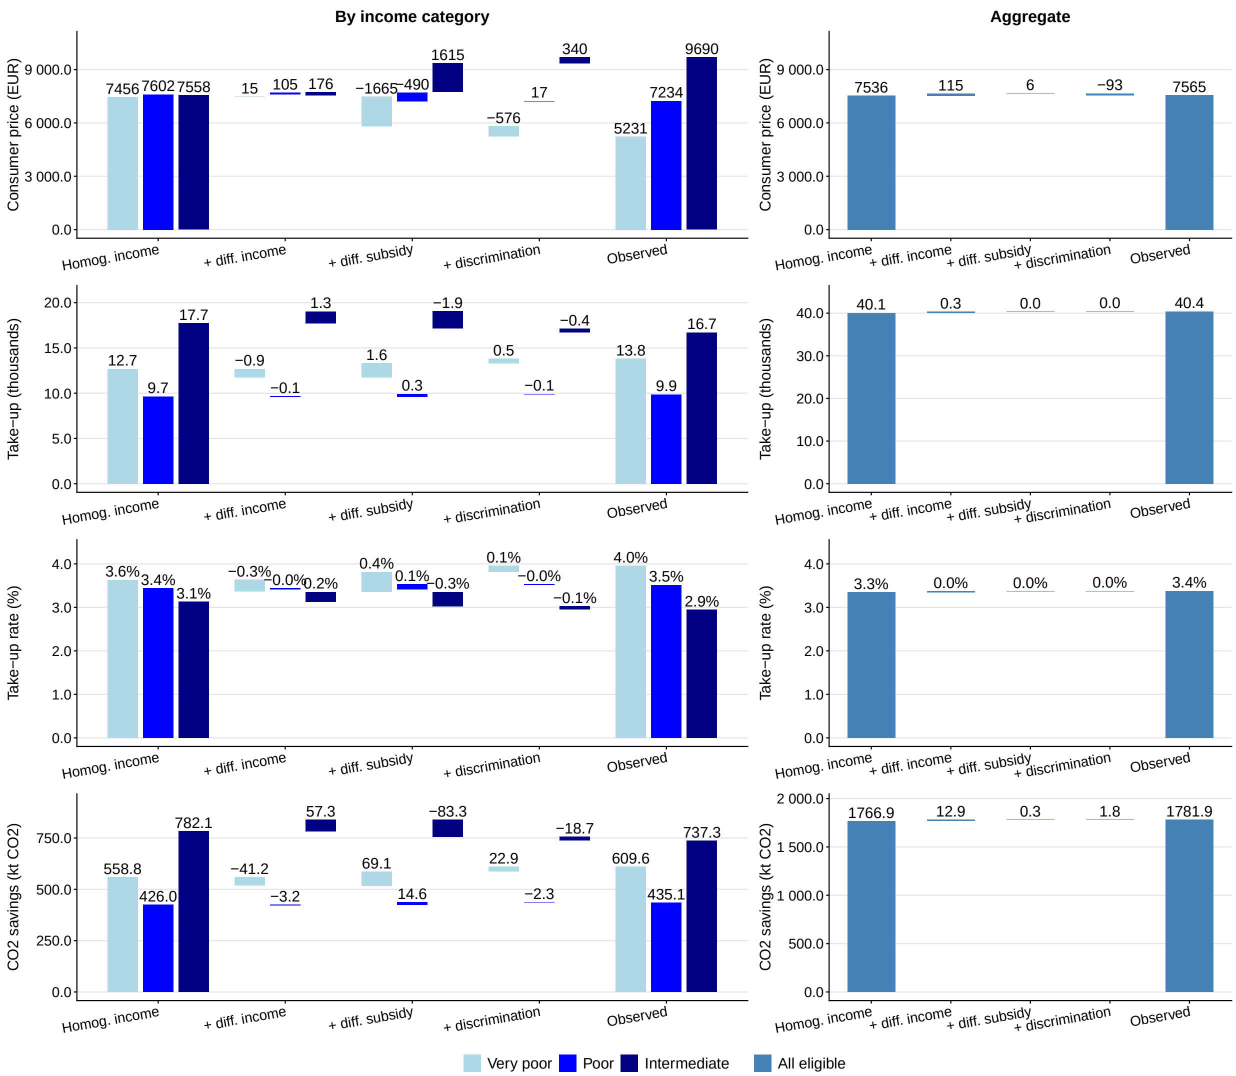

NOTE: This figure presents simulated consumer prices, take-up levels (in thousands), take-up rates and CO_2 savings (kt CO_2, lifetime) by income category, for four scenarios: equal income and subsidy (column 1), differentiated income and equal subsidy (column 2), differentiated income and subsidies with a common transaction price (column 3), differentiated income and subsidies allowing for price discrimination (column 4). Left panels show the three income categories separately; right panels show the aggregate across all three. Results are simulated on the 20 estimation markets. Government spending is held constant at €263 million across all scenarios. SOURCE: Fideli, MPR, Ademe, Bernard et al. (2024).

6.4 Robustness simulation exercise

One strong assumption we make in this structural model is that price discrimination happens systematically. Our reduced-form results, specifically those summed up in Figure 4, suggest on the contrary that price discrimination may actually mostly happen for households who mandate firms to fill out the subsidy application on their behalf ("mandated" applications), thereby revealing information on their income category. Modeling that choice endogenously is beyond the scope of our paper: the supply side's maximization problem ceases to be separable, and mandate status is by construction observed only for applicants. We therefore examine the consequences of the simplification by simulation, and report the exercise in full in Appendix C.5.

We simulate data-generating processes calibrated on our results, taking our estimated demand parameters and the marginal costs recovered from the first-order conditions as the truth and holding the observed market structure and subsidy schedule fixed, in which only a share π_c of households get assigned their income-category-specific prices and the rest a uniform price common to all categories. We then estimate our initial model, which is blind to the mandating step, on each simulated dataset, and compare the contribution of price discrimination it recovers to the DGP's truth.

Two results matter here. At the empirical mandated shares $[48\%, 39\%, 33\%]$, the procedure recovers the contribution to take-up of the Very Poor category almost exactly (223 against 218 installations) while overstating that of the Intermediate category by 17%; on prices, it recovers 87% of the true contribution to the Very Poor–Intermediate price differential. Setting $\pi_c = [0\%, 0\%, 0\%]$, where no discrimination takes place and any estimated contribution is a pure artifact, the procedure returns €37 against a true contribution of zero, an order of magnitude below what we estimate on the data. The bias does not grow with the share of non-discriminated transactions.

7 Conclusion

Subsidies targeted at low-income households are an attractive policy instrument to tackle environmental externalities, compared to first-best Pigouvian taxation, especially on grounds of political acceptability and distribution. The equity-efficiency trade-off that they represent, compared to Pigouvian taxation, is often evaluated assuming perfect competition on the supply side.

In this paper, we argue that not accounting for potential imperfect competition leads to a biased evaluation of such programs. We find, in the case of a French program of support to residential energy efficiency investments, that means-tested schemes interact with market power in two main ways: first, by providing information on subsidy levels to firms, means-tested schemes can allow firms to price discriminate between low-income households, with high price elasticity, and high-income households who are relatively less elastic. Second, market power allows for imperfect pass-through of subsidy differentials to consumers. In our context, these two effects occur and as the discrimination channel outweighs the regressive pass-through effect, it results in an increase in redistribution of the policy at no efficiency cost: conditional on existing market conditions, means-testing can be an efficient way to limit rationing due to imperfect competition among low-income consumers. A striking feature of our results is that this amplification of redistribution requires neither overshifting of the subsidy nor a deliberate pro-competitive design: it emerges naturally from the interaction of a convex enough demand curve and a sufficiently large elasticity differential across income groups, both of which we estimate to hold in the French heat pump market. Our evidence from other renovation actions suggests that the sign of the effect can reverse when these conditions do not hold, arguing for systematic evaluation of supply-side conditions when designing or assessing means-tested subsidies, rather than assuming competitive incidence. Where market conditions are favorable, as in our setting, means-testing can do more distributional work than its subsidy schedule alone suggests.

References

Adachi, Takanori, "A sufficient statistics approach for welfare analysis of oligopolistic third-degree price discrimination," *International Journal of Industrial Organization*, January 2023, *86*, 102893.

ADEME, "Base Carbone, version 23.6 — Facteurs d'émission : gaz naturel," 2024.

Allcott, Hunt and Michael Greenstone, "Measuring the Welfare Effects of Residential Energy Efficiency Programs," NBER Working Paper 23386, National Bureau of Economic Research January 2024.

_ , Christopher Knittel, and Dmitry Taubinsky, "Tagging and Targeting of Energy Efficiency Subsidies," *American Economic Review*, May 2015, *105* (5), 187–191.

Barwick, Panle Jia, Hyuk-Soo Kwon, and Shanjun Li, "Attribute-based Subsidies and Market Power: an Application to Electric Vehicles," Technical Report w32264, National Bureau of Economic Research, Cambridge, MA March 2024.

Bernard, Louise, Andy Hackett, Robert Metcalfe, and Andrew Schein, "Decarbonizing Heat: The Impact of Heat Pumps and a Time-of-Use Heat Pump Tariff on Energy Demand," Technical Report w33036, National Bureau of Economic Research, Cambridge, MA October 2024.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, July 1995, *63* (4), 841.

_ , _ , and _ , "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," *Journal of Political Economy*, 2004, *112* (1), 68–105. Publisher: The University of Chicago Press.

Bonnet, Odran, Étienne Fize, Tristan Loisel, and Lionel Wilner, "Compensating against fuel price inflation: Price subsidies or transfers?," *Journal of Environmental Economics and Management*, January 2025, *129*, 103079.

Borenstein, Severin and Lucas W Davis, "The Distributional Effects of U.S. Tax Credits for Heat Pumps, Solar Panels, and Electric Vehicles," July 2024.

Buchanan, James M., "External Diseconomies, Corrective Taxes, and Market Structure," *The American Economic Review*, 1969, *59* (1), 174–177. Publisher: American Economic Association.

Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik, "Robust Data-Driven Inference in the Regression-Discontinuity Design," *The Stata Journal: Promoting communications on statistics and Stata*, December 2014, *14* (4), 909–946.

— , — , Max H. Farrell, and Rocío Titiunik, "Regression Discontinuity Designs Using Covariates," *The Review of Economics and Statistics*, July 2019, *101* (3), 442–451.

Cannière, Charlotte De, "Pump it? Market power and the energy transition in the global oil market," 2025.

Cellini, Stephanie Riegg and Claudia Goldin, "Does Federal Student Aid Raise Tuition? New Evidence on For-Profit Colleges," *American Economic Journal: Economic Policy*, November 2014, *6* (4), 174–206.

Chan, Nathan W. and Kenneth Gillingham, "The Microeconomic Theory of the Rebound Effect and Its Welfare Implications," *Journal of the Association of Environmental and Resource Economists*, March 2015, *2* (1), 133–159. Publisher: The University of Chicago Press.

Colas, Mark and Emmett Reynier, "Means-Tested Solar Subsidies," *CESifo Working Paper Series*, 2024. Number: 11378 Publisher: CESifo.

Conlon, Christopher and Jeff Gortmaker, "Incorporating Micro Data into Differentiated Products Demand Estimation with PyBLP," *Journal of Econometrics*, January 2025, p. 105926.

Decarolis, Francesco, "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?," *American Economic Review*, April 2015, *105* (4), 1547–1580.

D'Haultfœuille, Xavier, Isis Durrmeyer, and Philippe Février, "Automobile Prices in Market Equilibrium with Unobserved Price Discrimination," *The Review of Economic Studies*, October 2019, *86* (5), 1973–1998.

Fack, Gabrielle, "Are housing benefit an effective way to redistribute income? Evidence from a natural experiment in France," *Labour Economics*, December 2006, *13* (6), 747–771.

Fillmore, Ian, "Price Discrimination and Public Policy in the US College Market," *The Review of Economic Studies*, May 2023, *90* (3), 1228–1264.

Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan, "Market-Based Emissions Regulation and Industry Dynamics," *Journal of Political Economy*, February 2016, *124* (1), 249–302.

— , Michael Greenstone, and Catherine Wolfram, "Are the Non-monetary Costs of Energy Efficiency Investments Large? Understanding Low Take-Up of a Free Energy Efficiency Program," *American Economic Review*, May 2015, *105* (5), 201–204.

— , — , and — , "Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program*," *The Quarterly Journal of Economics*, August 2018, *133* (3), 1597–1644.

Giandomenico, Lauren, Maya Papineau, and Nicholas Rivers, "A Systematic Review of Energy Efficiency Home Retrofit Evaluation Studies," *Annual Review of Resource Economics*, October 2022, *14* (Volume 14, 2022), 689–708. Publisher: Annual Reviews.

Gillingham, Kenneth, David Rapson, and Gernot Wagner, "The Rebound Effect and Energy Efficiency Policy," *Review of Environmental Economics and Policy*, January 2016, *10* (1), 68–88. Publisher: The University of Chicago Press.

_ , Matthew J. Kotchen, David S. Rapson, and Gernot Wagner, "The rebound effect is overplayed," *Nature*, January 2013, *493* (7433), 475–476. Publisher: Nature Publishing Group.

Goldberg, Pinelopi Koujianou, "The Effects of the Corporate Average Fuel Efficiency Standards in the US," *The Journal of Industrial Economics*, March 1998, *46* (1), 1–33.

Grigolon, Laura and Frank Verboven, "Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation," *The Review of Economics and Statistics*, 2014, *96* (5), 916–935. Publisher: The MIT Press.

Jagger, Nick, Tim Foxon, and Andy Gouldson, "Skills constraints and the low carbon transition," *Climate Policy*, January 2013, *13* (1), 43–57. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/14693062.2012.709079.

Kellogg, Ryan, "The End of Oil," Technical Report w33207, National Bureau of Economic Research, Cambridge, MA November 2024.

Ministère de la Transition écologique, "Arrêté du 4 août 2021 relatif aux exigences de performance énergétique et environnementale des constructions de bâtiments en France métropolitaine (RE2020)," 2021. tex.howpublished: Journal officiel de la République française.

Miravete, Eugenio J, Katja Seim, and Jeff Thurk, "Elasticity and Curvature of Discrete Choice Demand Models," 2025.

Muehlegger, Erich and David S. Rapson, "Subsidizing low- and middle-income adoption of electric vehicles: Quasi-experimental evidence from California," *Journal of Public Economics*, 2022, *216*, 104752.

Nevo, Aviv, "A practitioner's guide to estimation of random-coefficients logit models of demand," *Journal of economics & management strategy*, 2000, *9* (4), 513–548.

Newey, Whitney K. and Daniel McFadden, "Large sample estimation and hypothesis testing," in "Handbook of Econometrics," Vol. 4, Elsevier, 1994, pp. 2111–2245.

Pless, Jacquelyn and Arthur A. van Benthem, "Pass-Through as a Test for Market Power: An Application to Solar Subsidies," *American Economic Journal: Applied Economics*, October 2019, *11* (4), 367–401.

Ploeg, Frederick Van Der and Cees Withagen, "Global Warming and the Green Paradox: A Review of Adverse Effects of Climate Policies," *Review of Environmental Economics and Policy*, July 2015, *9* (2), 285–303.

Polyakova, Maria and Stephen P Ryan, "Market Power and Redistribution: Evidence from the Affordable Care Act," Technical Report 26367, National Bureau of Economic Research 2021.

Sallee, James, "Pigou Creates Losers: On the Implausibility of Achieving Pareto Improvements from Efficiency-Enhancing Policies," Technical Report w25831, National Bureau of Economic Research, Cambridge, MA May 2019.

Sallee, James M., "The surprising incidence of tax credits for the Toyota Prius," *American Economic Journal: Economic Policy*, 2011, *3* (2), 189–219.

Service des données et études statistiques (SDES), "TREMI 2020 — Travaux de rénovation énergétique des maisons individuelles," Technical Report, Observatoire national de la rénovation énergétique (ONRE) 2022.

Sinn, Hans-Werner, "Public policies against global warming: a supply side approach," *International Tax and Public Finance*, August 2008, *15* (4), 360–394.

Tebaldi, Pietro, "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," *The Review of Economic Studies*, January 2025, *92* (1), 586–620.

Weyl, E. Glen and Michal Fabinger, "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition," *Journal of Political Economy*, June 2013, *121* (3), 528–583. Publisher: The University of Chicago Press.

A Model of incidence for means-tested subsidies: proofs

Marginal profit. We first show that under Assumptions 1 to 4, the profit maximization problem is well behaved, that is, the marginal profit of firms is decreasing. The first-order condition of the profit maximization problem is the partial derivative of profit with respect to own price, evaluated at symmetric Bertrand-Nash price. We then differentiate it with respect to market price.

$$\begin{aligned}
\frac{\partial \pi_{jm}}{\partial p_j}(p) &= s_{jm}(p) + (p-c)\frac{\partial s_{jm}}{\partial p_j}(p) \\
\frac{d}{dp}\left[\frac{\partial \pi_{jm}}{\partial p_j}(p)\right] &= 2\frac{\partial s_{jm}}{\partial p_j}(p) + \frac{\partial s_{jm}}{\partial p_{-j}}(p) + (p-c)\left[\frac{\partial^2 s_{jm}}{\partial p_j^2}(p) + \frac{\partial^2 s_{jm}}{\partial p_j \partial p_{-j}}(p)\right] \\
&= 2\frac{\partial s_{jm}}{\partial p_j}(p) + \frac{\partial s_{jm}}{\partial p_{-j}}(p) - \frac{s_{jm}}{\frac{\partial s_{jm}}{\partial p_j}}\left[\frac{\partial^2 s_{jm}}{\partial p_j^2}(p) + \frac{\partial^2 s_{jm}}{\partial p_j \partial p_{-j}}(p)\right] \\
\frac{d}{dp}\left[\frac{\partial \pi_{jm}}{\partial p_j}(p)\right] < 0 &\iff \rho_m(p) < 2 + \left(\frac{1}{\theta_m(p)} - 1\right)
\end{aligned}$$

Noticing that the conduct parameter $\theta_m(p) \in [0,1]$ under Assumption 2 yields the result. □

Lemma 1. Lerner index derived from first-order condition under uniform pricing considers a weighted average of weak- and strong-market elasticities. It yields the following equivalence:

$$\begin{aligned}
\epsilon_l^{own}(\overline{p}) > \epsilon_h^{own} &\iff \frac{1}{\epsilon_l^{own}(\overline{p})} < \frac{\overline{p}-c}{\overline{p}} < \frac{1}{\epsilon_h^{own}(\overline{p})} \\
&\iff \frac{\partial \pi_{jw}}{\partial p_j}(p) < 0 < \frac{\partial \pi_{js}}{\partial p_j}(p)
\end{aligned}$$

Monotonicity of $\frac{\partial \pi_{jm}}{\partial p_j}(p), \forall m \in \{s, w\}$ obtained from the assumptions yields the result. □

Lemma 2. Total differentiation of firms' first-order conditions with respect to t yields cost pass-through at any given level of t

$$\frac{\partial p}{\partial c}(t) = \frac{1}{1 + \theta_l(p(t))(1 - \rho_l(p(t)))} > 0$$

Assumptions 2 and 3 ensure that pass-through is positive. Integrating it between 0 and full subsidy (negative tax $-b$) yields the result. □

Proposition 1. Total price differential under discrimination: The first term is the price discrimination effect: $p_h^* - p_l^*$. It is obtained using the Lerner index under discriminated pricing. The second term is minus the total pass-through of the subsidy in the weak market derived in Lemma 2, $p_l^* - p_l^*(b)$.

To show that when the condition holds, h consumers lose and l consumers win, let us consider a case where redistribution level (consumer price differential across sub-markets) is constrained to be $t \geq 0$. Profit maximization under a constrained redistribution level t can be written as:

$$\pi_{1w}(p-t) + \pi_{1s}(p) = (p-t-c+b)s_{1w}(p-t) + (p-c)s_{1s}(p)$$

The first-order condition can be written $F(\tilde{p}_h(t), t) = 0$. Applying the implicit function theorem to it and noticing that $\frac{\partial \tilde{p}_l}{\partial t} = \frac{\partial \tilde{p}_h}{\partial t} - 1$ yields:

$$\frac{\partial \tilde{p}_h}{\partial t} = \frac{-\frac{\partial F}{\partial t}}{\frac{\partial F}{\partial p}} = \frac{\pi_l''}{\pi_l'' + \pi_h''} > 0$$
$$\frac{\partial \tilde{p}_l}{\partial t} = -\frac{\pi_h''}{\pi_l'' + \pi_h''} < 0$$

where we denote $\pi_m'' = \frac{d}{dp}\left[\frac{\partial \pi_{jm}}{\partial p_j}(p)\right] < 0$.

Let us denote t^* the optimal level of price differential between the weak and strong market, from the point of view of firms. It corresponds exactly to the one derived in the condition of this proposition. If $t^* < b$, uniform pricing can be seen as forcing t to go beyond t^*, thus increasing $\tilde{p}_h$ and decreasing $\tilde{p}_l$, from what we showed on monotonicity of $\tilde{p}$ with respect to t. The reverse is true for $t^* \geq b$. $\square$

B Reduced-form evidence on transaction prices

B.1 Sample definition

Matching pipeline The analysis is based on MPR administrative records, matched to Fideli housing units through the Arel database, which stores the housing identifier of each MPR file. Matching is done in three successive steps. First, each Arel record carries a Fideli housing identifier that allows a direct deterministic join to the Fideli dwelling-level file, attaching household income, composition, and housing characteristics. This step recovers 100% of Arel records (Step 1). Second, non-missing grant dates and investment type codes are required for matching with MPR applications (Step 2, 94.4% retained). Third, Arel records are matched to the corresponding MPR application dossier using a sequence of progressively relaxed exact joins on transaction price and subsidy amount, date of commission and end of work, and geodesic distance, with spatial proximity as a tie-breaker (Step 3, 91.8% retained). Table A.1 documents each step and its retention rate, overall and by year.

For the RDD exercise on all renovation actions, Step 4 further requires that transaction prices and subsidy amounts are not missing (94% of Step 3). Step 5 restricts to heat pump actions. Steps 6 to 8 match MPR heat pump files to their related CEE subsidy account, when it exists, to recover the coefficient of performance of the installed heat pump. This analysis focuses on slightly less than 50% of MPR heat pump files.

Table A.1: Sample selection in the MPR matching pipeline

	Total			2020			2021			2022			2023		
	N	%i.	%p.	N	%i.	%p.	N	%i.	%p.	N	%i.	%p.	N	%i.	%p.
Initial population	1,999,788	100.0%	—	140,907	100.0%	—	699,029	100.0%	—	689,556	100.0%	—	470,296	100.0%	—
Step 1: Fideli housing	1,999,755	100.0%	100.0%	140,907	100.0%	100.0%	699,026	100.0%	100.0%	689,554	100.0%	100.0%	470,268	100.0%	100.0%
Step 2: date & type	1,888,426	94.4%	94.4%	140,533	99.7%	99.7%	690,512	98.8%	98.8%	678,528	98.4%	98.4%	378,853	80.6%	80.6%
Step 3: MPR app. file	1,836,038	91.8%	97.2%	138,360	98.2%	98.5%	673,996	96.4%	97.6%	657,604	95.4%	96.9%	366,078	77.8%	96.6%
RDD — all renovation actions															
Step 4: amounts > 0	1,724,140	86.2%	93.9%	131,208	93.1%	94.8%	637,202	91.2%	94.5%	614,322	89.1%	93.4%	341,408	72.6%	93.3%
RDD — heat pumps only															
Step 5: heat pump only	369,150	18.5%	21.4%	27,341	19.4%	20.8%	129,912	18.6%	20.4%	130,131	18.9%	21.2%	81,766	17.4%	23.9%
Step 6: CEE subsidy > 0	317,993	15.9%	86.1%	23,028	16.3%	84.2%	115,543	16.5%	88.9%	112,518	16.3%	86.5%	66,904	14.2%	81.8%
Step 7: CEE in Arel	145,950	7.3%	45.9%	13,537	9.6%	58.8%	67,905	9.7%	58.8%	56,992	8.3%	50.7%	7,516	1.6%	11.2%
RDD — heat pump quality															
Step 8: CEE registry	145,755	7.3%	99.9%	13,511	9.6%	99.8%	67,807	9.7%	99.9%	56,931	8.3%	99.9%	7,506	1.6%	99.9%

Note: Each row corresponds to a successive restriction applied to the MPR–Fideli matched dataset. *N*: number of housing units remaining after the restriction. *%i.*: percentage of the initial population (Step 1). *%p.*: percentage of the previous step. Steps 1–4 define the population used for the RDD on all renovation actions; Steps 5–8 define the subsamples used for RDD analyses restricted to heat pumps.

Match quality across the income distribution. Figure A.1 plots the share of MPR records successfully matched at Step 4, by 1,000-euro bin of household standard of living, for all renovation actions and for heat pumps separately. Match quality is broadly homogeneous across the income distribution: roughly 85–90% of records are matched at all income levels. The exception is households with very low standard of living (below approximately €12,000), where the match rate falls somewhat.

This pattern raises no concern for the RDD exercise: both the Very Poor–Poor and Poor–Intermediate income thresholds lie well above the range where match quality deteriorates. For the structural demand model, the relevant sample is the 20-market estimation dataset restricted to households above the 5th percentile of income (roughly €12,000), which explicitly excludes the low-match region. The structural model's observed take-up rates are therefore unlikely to be systematically biased by differential matching quality.

Figure A.1: Match quality by standard of living

(a) All renovation actions (b) Heat pumps only

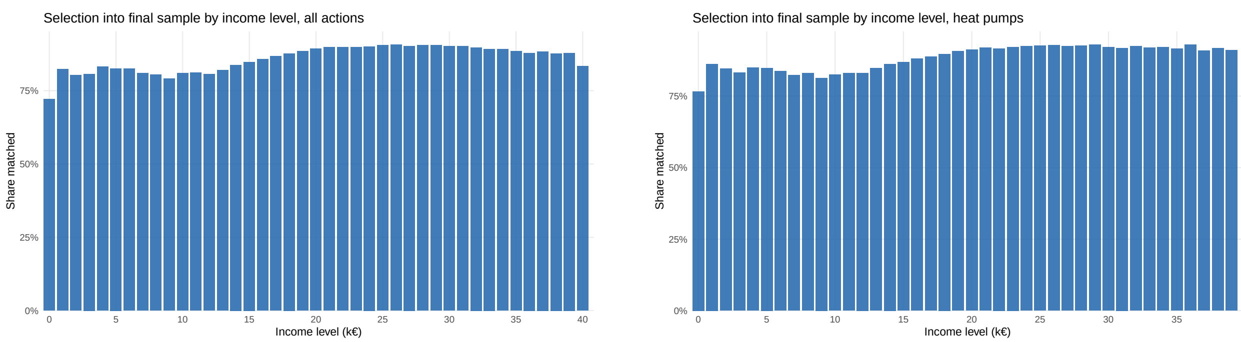

NOTE: Each panel shows the share of MPR records successfully matched at Step 4 of the matching pipeline, by €1,000 bin of household standard of living (*niveau de vie*). Left panel: all renovation actions. Right panel: heat pump actions only. SOURCE: Fideli, MPR.

B.2 Estimation

The estimation strategy adapts the robust non-parametric RDD of Calonico et al. (2014) to account for municipality-level confounders and individual heterogeneity. It proceeds in three steps.

Step 1: RD robust estimation

We estimate the transaction price discontinuity at each income threshold $\bar{x}$ using the rdrobust estimator of Calonico et al. (2014). The treatment is the sharp change in subsidy level at the threshold: every household is assigned to the income category corresponding to its fiscal income, and the design is sharp since ANAH cross-checks declared fiscal income with administrative records. The estimand is:

$$\tau_{\text{SRD}} = \lim_{x \downarrow \bar{x}} \mathbb{E}[Y_i \mid X_i = x] - \lim_{x \uparrow \bar{x}} \mathbb{E}[Y_i \mid X_i = x]$$

where X_i is fiscal income normalized to zero at the threshold. Estimation is by local linear regression on each side of the cutoff:

$$\hat{\mu}^{\pm}_{(0)}(h) = \arg\min_{\boldsymbol{\beta}^{\pm}} \sum_{i:\,\pm(X_i-\bar{x})\geq 0} \big(Y_i - \boldsymbol{r}_1(X_i-\bar{x})^\top \boldsymbol{\beta}^{\pm}\big)^2 K\Big(\tfrac{X_i-\bar{x}}{h}\Big)$$

with $\boldsymbol{r}_1(x)$ the polynomial basis of degree 1, triangular kernel $K(\cdot)$ and MSE-optimal bandwidth h selected by the procedure of Calonico et al. (2014). The point estimate is $\hat{\tau}_p(h) = \hat{\mu}^{+}_{(0)}(h) - \hat{\mu}^{-}_{(0)}(h)$, and inference relies on bias-corrected and robust ("robust bias correction") confidence intervals that remain valid under the MSE-optimal bandwidth.

Step 2: Residualization on municipality fixed effects

Transaction prices Y_i are first regressed on a full set of municipality fixed effects. The residuals $\tilde{Y}_i$ abstract from local market-level variation in prices, ensuring that the estimated treatment effect at the threshold captures income-based differences rather than geographical variation in market conditions. Since firms operate locally and market power varies substantially across municipalities, this step is essential to remove location-driven price differences from the comparison. The residualized outcome $\tilde{Y}_i$ is used as the dependent variable in all subsequent RDD specifications.

Step 3: Adding individual controls

Following Calonico et al. (2019), a third specification augments the local linear regression with a vector of individual-level controls Z_i:

$$\hat{\mu}^{\pm}_{(0)}(h) = \arg\min_{\boldsymbol{\beta}^{\pm},\boldsymbol{\gamma}^{\pm}} \sum_{i:\,\pm(X_i-\bar{x})\geq 0} \big(\tilde{Y}_i - \boldsymbol{r}_1(X_i-\bar{x})^\top \boldsymbol{\beta}^{\pm} - Z_i^\top \boldsymbol{\gamma}^{\pm}\big)^2 K\Big(\tfrac{X_i-\bar{x}}{h}\Big)$$

Controls Z_i include: housing construction period dummies (pre-1949, 1949–1974, 1975–1999, post-2000), housing surface area, household size, household standard of living, a dummy for mandated files, number of co-investment retrofit actions in the same file, and indicators for housing quality standing. As established in Calonico et al. (2019), adding controls in this way improves estimation efficiency under covariate balance but does not restore causal identification when the density of the forcing variable is discontinuous at the cutoff. We use these controlled estimates to assess whether the residual price gap can be attributed to observable imbalances across the threshold.

B.3 Additional results

This subsection collects robustness checks for the main RDD results. They address four concerns in turn: (i) whether the observed characteristics of households are balanced at the income thresholds; (ii) whether firm-level sorting confounds the estimated price gap; (iii) whether the quality of installed heat pumps changes at the threshold; and (iv) whether similar price effects are found in other renovation categories and whether they vary with local market concentration.

Balance tests (Figure A.2). The density discontinuity documented in Figure 5 means that the composition of takers changes at the threshold, preventing a fully causal interpretation of the price gap. To assess the severity of this concern, we apply the RDD estimator to a set of predetermined household and housing characteristics as pseudo-outcomes. Panel (a) reports p-values; Panel (b) reports estimates as a share of the left-side mean. Several characteristics are mildly imbalanced, notably co-investment count and household size. Their inclusion as controls in the main specification leaves the price gap estimate essentially unchanged.

Firm fixed effects (Figure A.3). When residualizing on firm rather than municipality fixed effects, estimates are small or close to zero at both thresholds. This specification is very demanding: the median heat pump installer performs only two MPR-subsidized installations per year, leaving little within-firm variation to identify a price differential. The attenuation is nonetheless consistent with part of the discrimination operating through selective firm sorting into income categories. The mandated/non-mandated pattern survives in the firm-FE specification, reinforcing the information-channel interpretation.

Quality of installed heat pumps (Figure A.4). Firms overwhelmingly specialize in a single quality tier: Figure A.4 shows that the majority of active RGE heat pump installers concentrate more than 80% of their sales in their dominant COP category. This suggests that firms cannot or will not discriminate through quality across income categories.

Housing size heterogeneity (Figure A.5). The central threat to identification in our regression discontinuity design is selection into take-up on either side of the threshold. In particular, households taking up heat pumps just above the threshold live in slightly larger homes than those just below (a gap of about 1%). Since larger homes could be paired with larger, more expensive heat pumps, this size differential could mechanically generate the discontinuous price gap we observe. We take this concern seriously. Beyond including housing surface as a linear control in our baseline specification, we estimate the RD coefficient separately by tercile of housing surface. The resulting heterogeneity analysis shows that the positive transaction price gap holds within terciles, in which households are *de facto* more comparable. Selection on housing size is therefore unlikely to be driving our main result.

Other renovation actions and HHI heterogeneity (Figures A.6 and A.7). Extending the analysis to all 16 MPR renovation types finds significantly positive price differentials in several cases and negative ones in one (Figure A.6).[28] This dispersion is consistent with the theoretical framework: the sign of the price discrimination effect depends on market-specific demand curvature and the elasticity differential, which vary across renovation types. In a majority of cases, the mandated and non-mandated estimates have the same sign. Finally, Figure A.7 shows that the estimated price gap at the Very Poor–Poor threshold is broadly increasing with HHI.

Table A.2: Matched two-by-two design: mandated versus non-mandated files

	Non-mandated		Mandated	Difference	
	Estimate	s.e.	Estimate	Estimate	s.e.
Panel A. Prices and subsidies (€)					
Transaction price	156.99	35.04	644.96	487.97	47.01
Consumer price	1396.63	39.24	2222.86	826.24	52.27
Subsidy	–1239.63	19.77	–1577.90	–338.27	24.65
Transaction price, % of subsidy differential	12.7		40.9		
Panel B. Balance on matching covariates					
Surface (m^2)	–2.421	0.334	–2.122	0.298	0.437
Construction year	–4.994	0.653	–5.002	–0.008	0.854
Dwelling type	0.052	0.006	0.049	–0.003	0.008
Number of occupants	–0.036	0.012	–0.031	0.005	0.015
Number of rooms	–0.118	0.015	–0.095	0.023	0.020
MPR Date (days)	–6.556	2.937	3.988	10.544	3.924
Paris region dummy	–0.024	0.002	–0.006	0.019	0.003
Files per cell			93,715		
Total files			374,860		

NOTE : Application files are matched on the seven Panel B covariates across the four cells defined by mandate status and by position relative to the income threshold, so that each cell contains the same number of files. Estimates pool both thresholds. Columns 1–2 report the discontinuity among non-mandated files, column 3 among mandated files, and columns 4–5 the difference in discontinuities. The number of works engaged is determined after filing and is excluded from the matching; it is reported for completeness. SOURCE: Fideli, MPR

[28] MPR subsidizes 26 types of renovation. Only the 16 most prominent ones are present in the Arel database and can therefore be matched with the universe of eligible housing.

Figure A.2: Balance tests at MPR income thresholds

(a) p-values of RDD estimates on pseudo-outcomes

(b) RDD estimates on pseudo-outcomes as % of baseline

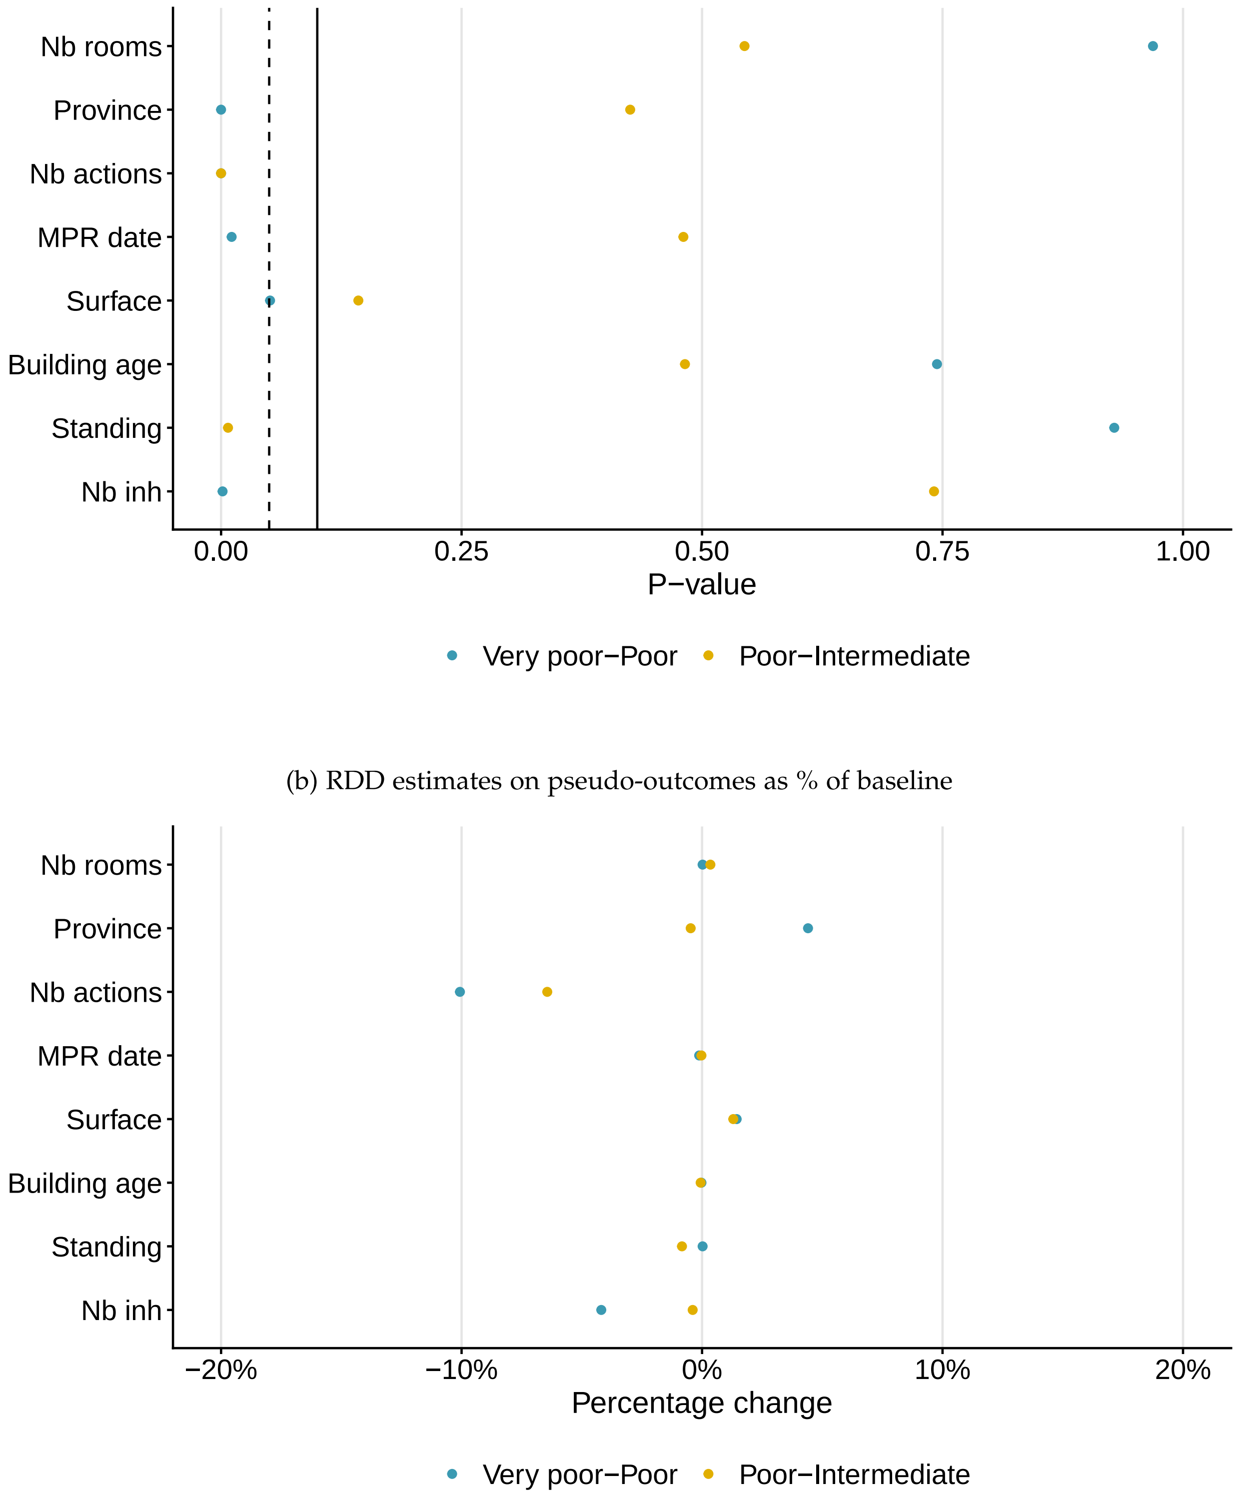

Note: This figure presents balance tests at the MPR income thresholds using a set of predetermined household characteristics as pseudo-outcomes, separately at the Very Poor–Poor and Poor–Intermediate thresholds. Estimation sample is restricted to observations within €4,000 of the threshold. Estimates are computed non-parametrically following Calonico et al. (2014). Panel (a) reports the p-value of the RDD point estimate for each pseudo-outcome; the dashed line marks the 5% significance level. Panel (b) reports the point estimate expressed as a percentage of the baseline value (mean on the left of the threshold). Source: Fideli, MPR.

Figure A.3: RDD estimates on transaction prices — firm fixed effects

(a) Very Poor–Poor

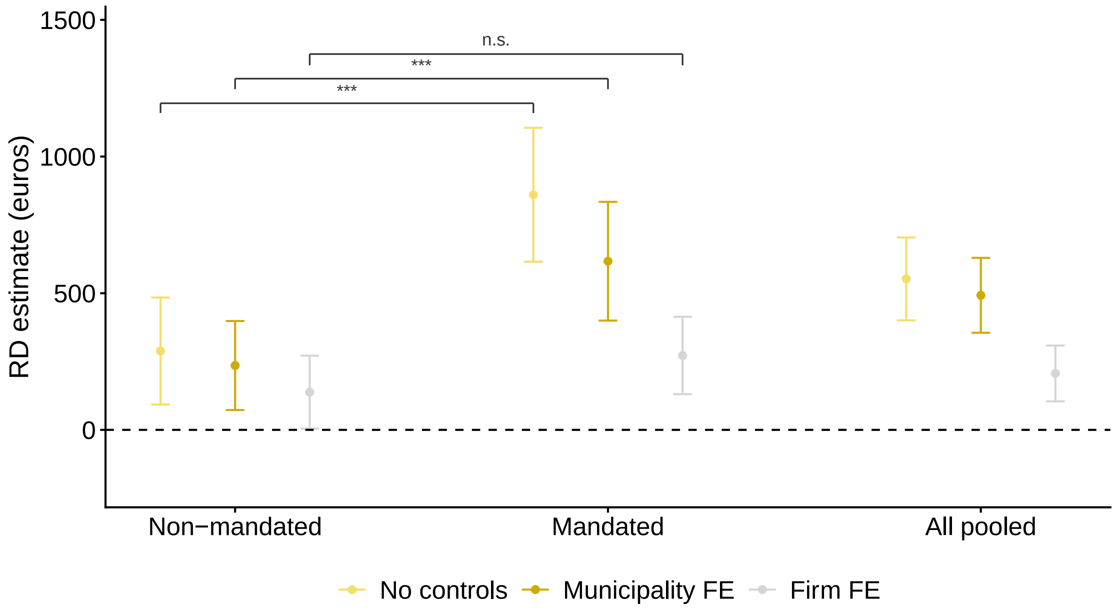

(b) Poor–Intermediate

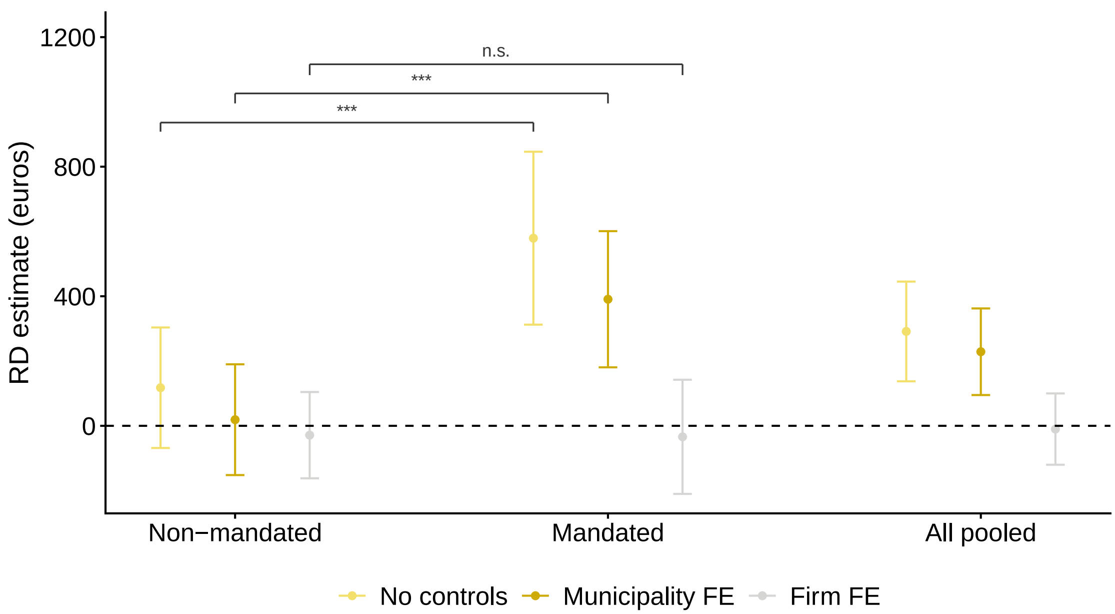

NOTE: This figure presents the estimated transaction price differential at MPR income thresholds, separately for non-mandated files, mandated files, and all files pooled, at the Very Poor–Poor and Poor–Intermediate thresholds. Three sets of estimates are presented for each specification: without controls (No controls), residualized on municipality fixed effects (Municipality FE), and residualized on firm fixed effects (Firm FE). Sample is made of one observation per MPR file involving a heat pump. Brackets indicate the significance of the difference between the two point estimates they connect, with stars denoting significance at the 10%, 5% and 1% level. SOURCE: Fideli, MPR.

Figure A.4: Share of main heat pump quality in total firm activity

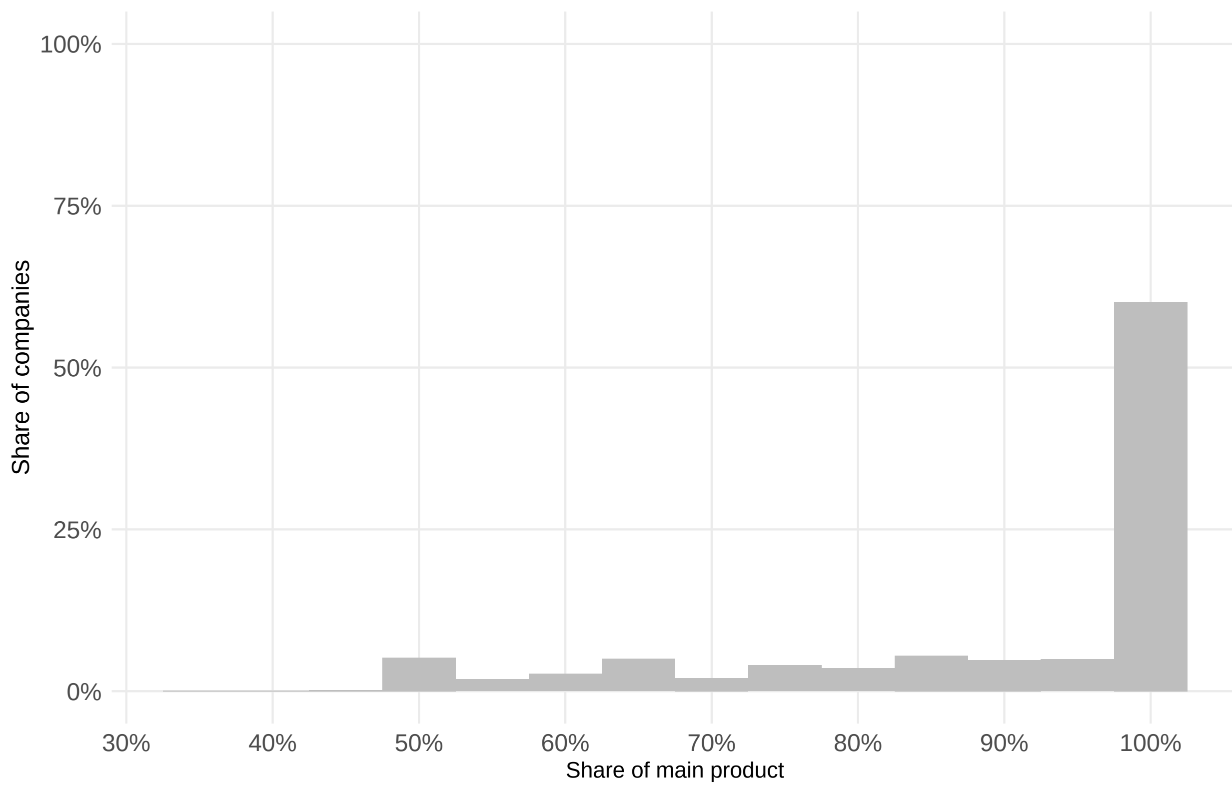

NOTE: This figure shows, for each active RGE heat pump installer, the share of its MPR activity concentrated in its main quality level — the coefficient of performance (COP) category accounting for the largest share of its sales. For each firm, we compute the share of MPR-CEE files by COP level, identify the dominant level, and plot the distribution of that share across all firms. SOURCE: Fideli, CEE, MPR.

Figure A.5: RDD estimates on transaction prices — by housing size

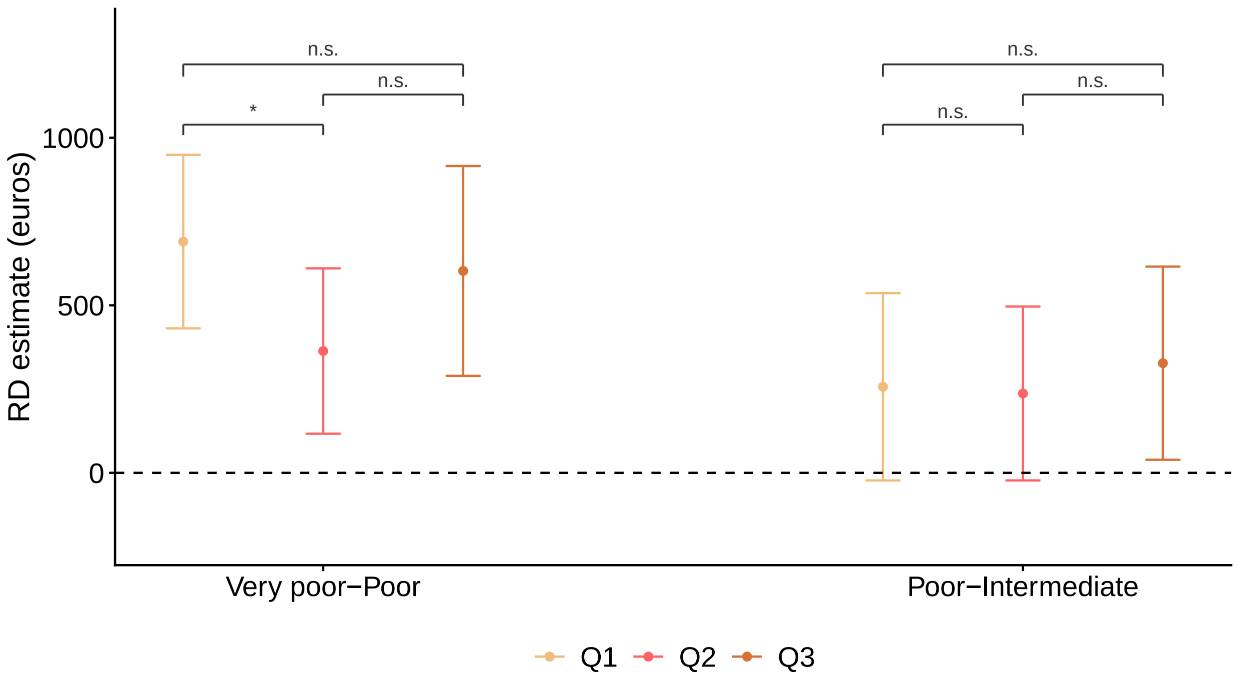

NOTE: This figure presents the estimated transaction price differential at MPR income thresholds, separately at the Very Poor–Poor and Poor–Intermediate thresholds and for three terciles of housing size (Q1 to Q3). Regressions are estimated without controls. Brackets indicate the significance of the difference between the two point estimates they connect, with stars denoting significance at the 10%, 5% and 1% level. SOURCE: Fideli, MPR.

Figure A.6: RDD estimates on transaction prices — all MPR investments

(a) Very Poor–Poor

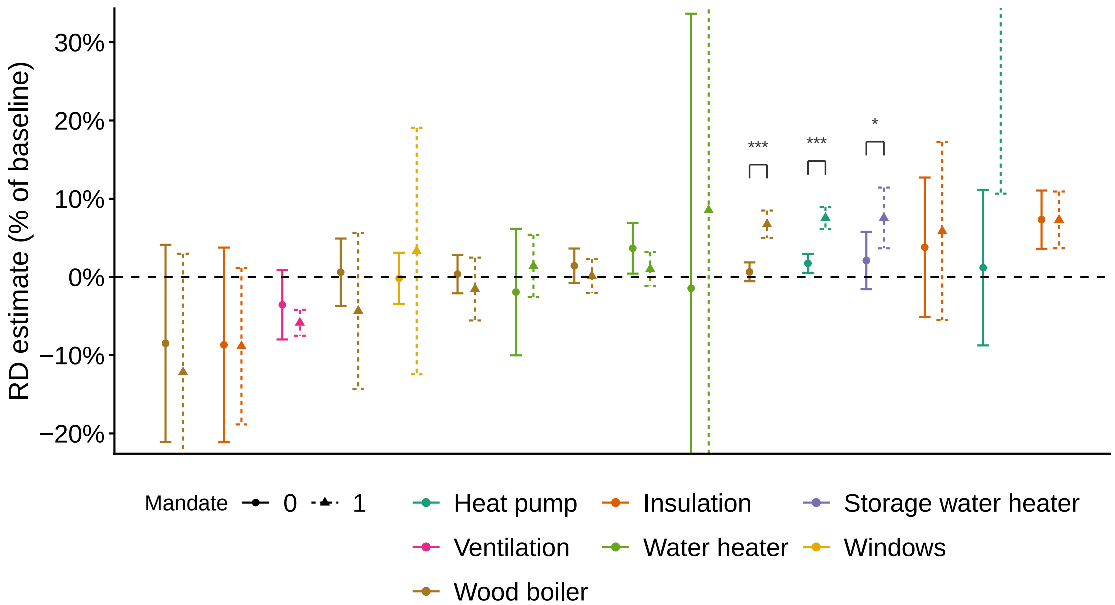

(b) Poor–Intermediate

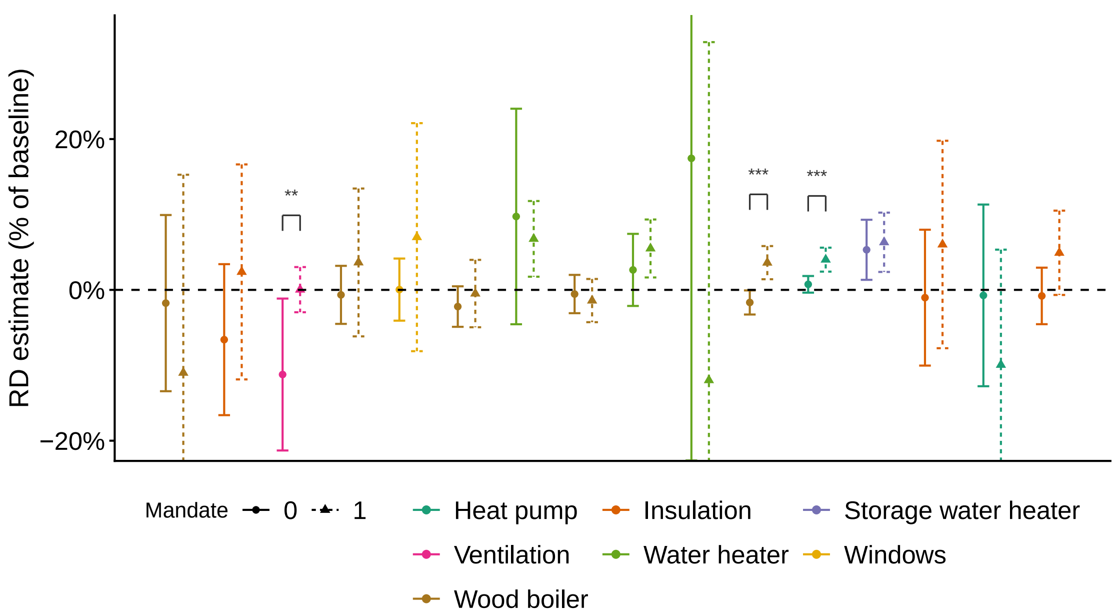

Note: This figure presents the estimated transaction price differential at MPR income thresholds, separately among mandated and non-mandated files, and at the Very Poor–Poor and Poor–Intermediate thresholds, for each type of MPR energy efficiency investment. Regressions are estimated without controls. Point estimates are presented as percentage change with respect to mean price on the left of the threshold. Dot color corresponds to the investment category. Investments are ordered by the pooled RD estimate at the Very Poor–Poor threshold. Brackets indicate the significance of the difference between the two point estimates they connect, with stars denoting significance at the 10%, 5% and 1% level. Source: Fideli, MPR.

Figure A.7: RDD estimates on transaction prices — by local market concentration

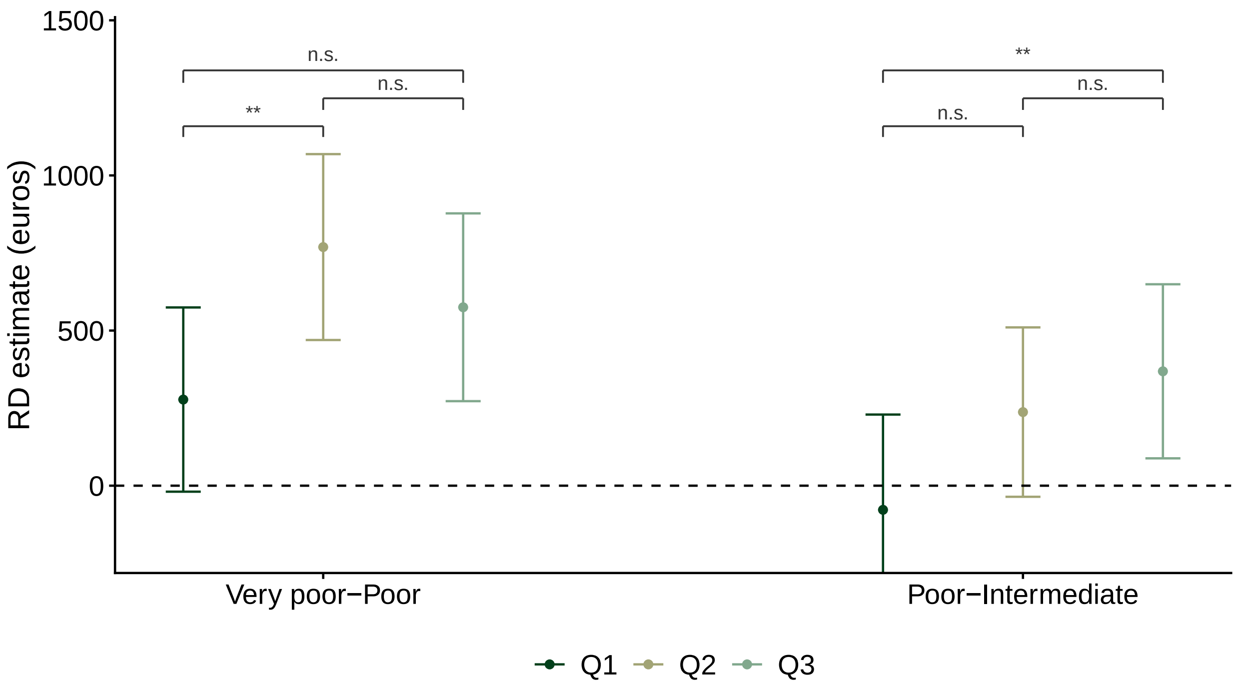

NOTE: This figure presents the estimated transaction price differential at MPR income thresholds, separately at the Very Poor–Poor and Poor–Intermediate thresholds and for three levels of market concentration. Regressions are estimated without controls. Market concentration is defined by the HHI of firms in a 50 km radius around the household. Observations are split into three terciles of concentration (Q1 to Q3) of the same sample size, with increasing HHI. Brackets indicate the significance of the difference between the two point estimates they connect, with stars denoting significance at the 10%, 5% and 1% level. SOURCE: Fideli, MPR.

Table A.3: Bounding the contribution of sample composition

	Hedonic gradient	Discontinuity	Contribution	
	β	Δ	$\beta\Delta$	s.e.
Surface (m^2)	9.431	1.521	14.34	(5.61)
Construction year	–3.391	–0.625	2.12	(3.62)
Dwelling type	–38.836	–0.018	0.70	(0.47)
Number of occupants	–20.393	–0.050	1.02	(0.53)
Number of rooms	26.090	0.024	0.62	(0.63)
MPR Date (days)	1.493	–14.745	–22.01	(9.04)
Paris region dummy	843.822	0.019	16.44	(4.55)
Total composition effect			13.23	(12.16)
RD estimate τ			450.42	(58.61)
Share of τ			*2.9%*	
Stratum-based variant			4.78	

NOTE : Column 1 reports coefficients from a hedonic regression of the unit transaction price on the seven predetermined characteristics with municipality fixed effects (n = 310,419, R^2 = 0.035). Column 2 reports the estimated discontinuity in each characteristic at the income threshold. Their product is the price differential that a change in composition at the threshold would mechanically produce; standard errors on $\beta\Delta$ account for uncertainty in both terms. Estimates pool both thresholds. The stratum-based variant reweights the eight province × household-size strata by their estimated discontinuity in sample share, centering prices within cell.
SOURCE: Fideli, MPR

C Structural model: price discrimination under means-testing

C.1 Sample selection

The estimation sample is constructed from the Fideli 2022 housing database, which covers the full French residential building universe (51.9 million housing units). Table A.4 documents the successive selection steps.

We first retain valid housing units with at least one occupant and a non-missing fiscal income (Step 1). Excluded units are either not a housing unit, or not a primary residence. Step 2 further restricts to units whose income category can be assigned and falls strictly below the highest heat pump eligibility bracket (below income category 4). Step 3 applies the technical eligibility conditions required by the MPR program: the housing unit must be individually owned, must not be social housing, and must have been built before 2007. We also restrict to houses, as take-up in multi-family units is very low. Step 4 adds additional restrictions to exclude extreme observable characteristics: surface area above $25\,\mathrm{m}^2$, income above the 5th percentile (fiscal and disposable income), top and bottom extreme values of housing quality.

Estimating a discrete choice model requires observing both takers and non-takers. MPR takers in the heat pump category are exhaustively observed in the administrative data. Non-takers are drawn from the eligible population using a 42% stratified random sample, stratified by département and income category. The 42% sampling rate is chosen to match the observed response rate in the TREMI 2020 survey of residential energy renovation (Service des données et études statistiques , SDES), which provides the only reliable benchmark for the share of homeowners who actively considered and rejected major heating renovation during the reference period. This leaves 3.1 million housing units.

Finally, the discrete choice estimation requires observing a sufficient number of heat pump choices per product group in each market (*département*). We retain the 20 French *départements* with the highest heat pump take-up among eligible households (Figure A.8). This reduces the sample to 1.2 million units. Restricting to high-take-up markets is a deliberate choice rather than a data limitation: the challenge of this empirical exercise is the small number of observed choices per product group, and concentrating the analysis on the most active markets maximizes the precision of the estimated demand parameters. We discuss the implications of this selection below. Within the 20-market sample, 41,145 households (3.4%) are observed taking up a heat pump through the MPR program in 2022 (Step 7).

Table A.5 compares summary statistics at four stages of the sample selection. Several differences are worth noting.

Between Step 1 (valid housing) and Step 5 (eligible stratified sample), the restrictions shift the sample toward homeowners of houses (by construction, Step 3), larger dwellings (average surface $87\,\mathrm{m}^2$ at Step 1 vs. $103\,\mathrm{m}^2$ at Step 5), older buildings, and slightly lower incomes. The

Table A.4: Sample selection: Model of heat pump demand

Step	N	Share of initial (%)	Share of prev. step (%)
Step 0: All residential building (baseline)	51,962,132	100.0%	—
Step 1: Valid housing	28,933,277	55.7%	55.7%
Step 2: Income category assigned (cat $<$ 4)	20,885,682	40.2%	72.2%
Step 3: Eligible housing	7,863,546	15.1%	37.7%
Step 4: Additional restrictions	7,290,404	14.0%	92.7%
Step 5: 42% stratified sample	3,122,737	6.0%	42.8%
Estimation sample and MPR takers			
Step 6: Final estimation sample (20 markets)	1,196,455	2.3%	38.3%
Step 7: MPR heat pump taker	41,145	0.1%	3.4%

Figure A.8: Map of selected markets

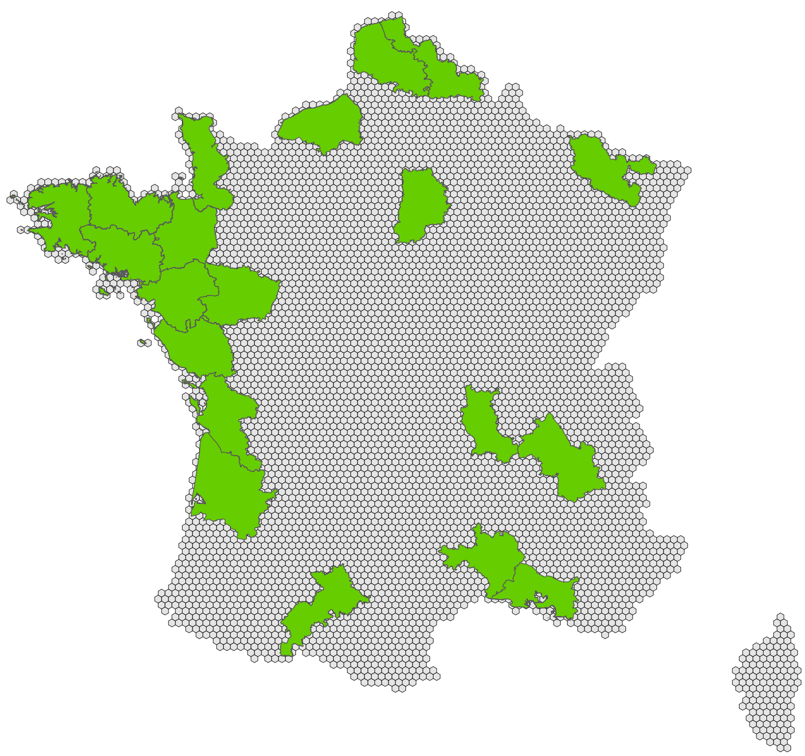

NOTE: This map displays in green the twenty French *départements* used to estimate the structural model. It also shows the 10 km-radius hexagons used to group firms by location in the structural model.

income restriction (Step 4 bottom-coding) excludes few observations.

Figure A.9: Distribution of MPR eligibility categories by quantile of equivalized disposable income

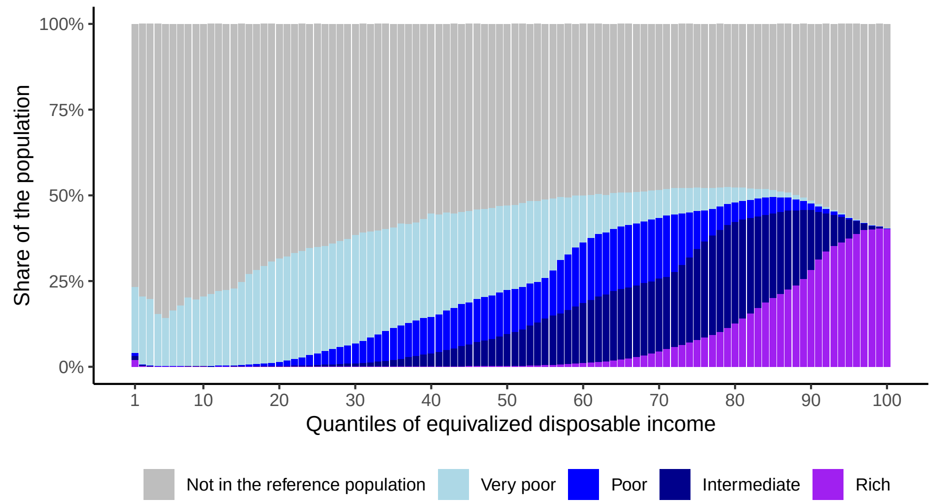

NOTE: This figure displays, for each percentile of equivalized disposable income (standard of living), the share of the Fideli housing stock belonging to each MPR income eligibility category (Very Poor, Poor, Intermediate, Rich) or outside the reference population. The reference population consists of housing units that are individually owned, are occupied as a principal residence, were built before 2007, and have a non-missing fiscal reference income. MPR eligibility categories are assigned on the basis of reference taxable income, which differs from equivalized disposable income. SOURCE: MPR, Fideli.

Between Step 5 (eligible stratified sample) and Step 6 (20-market estimation sample), summary statistics are nearly identical, confirming that the market selection does not introduce strong differences in observable characteristics. The 20 markets are not systematically different in income, housing size, or construction era from the broader eligible population.

Between Step 6 (estimation sample) and Step 7 (MPR heat pump takers), heat pump adopters have slightly larger dwellings (110 m^2 vs. 103 m^2), more occupants (2.5 vs. 2.3), and slightly higher income (€33,700 vs. €32,300 fiscal income; €39,200 vs. €37,300 disposable income). These differences are modest and consistent with heat pump adoption requiring upfront investment and space for the outdoor unit. Importantly, the overlap in income distributions between takers and non-takers remains substantial, ensuring that demand parameters are identified from meaningful within-market variation rather than extrapolation.

C.2 Estimation

Following Berry et al. (2004), the estimation of demand parameters proceeds as follows: first, for a given value of θ_2, a fixed point algorithm is used to concentrate out mean utilities $\delta(\theta_2)$ that rationalize observed market shares. θ_2 and $\delta(\theta_2)$ are used to build GMM moments for estimation of $\hat{\theta}_2^d$ and $\hat{\delta}(\hat{\theta}_2^d)$. $\hat{\theta}_1^d$ is then estimated by OLS from $\hat{\delta}(\theta_2^d)$. The following section presents the

Table A.5: Summary statistics: housing stock

(1) Step 1: Valid housing

	Mean	Min	Median	Max	SD	N
Number of rooms	5.9	0.0	6.0	99.0	1.9	28,849,942
Surface area, m^2	87.7	0.0	82.0	5,310.0	41.4	28,854,582
Year built	1956	1000	1975	2022	64	28,467,688
House	0.57	0.00	1.00	1.00	0.49	28,933,277
Standing (0–9)	5.2	0.0	5.0	8.0	0.7	28,869,626
Number of occupants	2.2	0.5	2.0	466.0	1.4	28,933,277
Fiscal income	35,530	0	28,355	181,718	30,370	28,933,277
Standard of living	25,096	0	22,620	87,300	13,759	28,933,277
Disposable income	39,070	0	32,548	155,153	26,507	28,933,277

(2) Step 5: Eligible sample

	Mean	Min	Median	Max	SD	N
Number of rooms	6.6	0.0	6.0	97.0	1.6	3,122,737
Surface area, m^2	102.5	26.0	97.0	2,987.0	34.6	3,122,737
Year built	1940	1000	1970	2006	68	3,122,737
House	1.00	1.00	1.00	1.00	0.00	3,122,737
Standing (0–9)	5.4	2.0	5.0	7.0	0.6	3,122,736
Number of occupants	2.3	1.0	2.0	15.0	1.2	3,122,737
Fiscal income	32,202	7,669	29,774	181,718	16,107	3,122,737
Standard of living	23,140	10,771	22,699	87,300	6,439	3,122,737
Disposable income	37,200	10,771	34,406	155,153	16,648	3,122,737

(3) Step 6: Estimation sample (20 markets)

	Mean	Min	Median	Max	SD	N
Number of rooms	6.6	0.0	6.0	96.0	1.6	1,196,455
Surface area, m^2	102.8	26.0	97.0	1,184.0	32.9	1,196,455
Year built	1945	1000	1971	2006	63	1,196,455
House	1.00	1.00	1.00	1.00	0.00	1,196,455
Standing (0–9)	5.4	2.0	5.0	7.0	0.6	1,196,455
Number of occupants	2.3	1.0	2.0	14.0	1.2	1,196,455
Fiscal income	32,255	7,669	30,078	181,718	15,679	1,196,455
Standard of living	23,052	10,771	22,670	87,300	6,221	1,196,455
Disposable income	37,320	10,772	34,720	155,153	16,354	1,196,455

(4) Step 7: MPR heat pump takers

	Mean	Min	Median	Max	SD	N
Number of rooms	6.9	1.0	7.0	80.0	1.6	41,145
Surface area, m^2	110.3	26.0	104.0	558.0	33.9	41,145
Year built	1947	1117	1971	2006	62	41,145
House	1.00	1.00	1.00	1.00	0.00	41,145
Standing (0–9)	5.3	2.0	5.0	7.0	0.6	41,145
Number of occupants	2.5	1.0	2.0	14.0	1.3	41,145
Fiscal income	33,696	7,671	31,543	181,718	16,372	41,145
Standard of living	23,079	10,775	22,503	87,300	6,576	41,145
Disposable income	39,167	10,777	36,687	155,153	16,978	41,145

Notes: (1) Step 1: non-missing rfr and at least one occupant. (2) Step 5: 42% stratified non-taker sample (Trémi 2020). (3) Step 6: final estimation sample (20 markets). (4) Step 7: MPR heat pump takers in the estimation sample. Fiscal income, standard of living, and disposable income are winsorized at the 99th percentile of the Step-1 population.

details of this estimation.

Aggregation of products The first challenge for this estimation procedure is the large number of products J_m in each market and high market share of the outside option. Estimation of the δ vector requires a large number of observed choices for each option j for the subsequent estimation of θ_1^d to be precise. That is not the case when observed market shares s_{cjm}^d are small and noisy. We deal with this challenge by aggregating market shares of similar firms called groups of products g. Products j and j' located in the same 10 km radius hexagon and the same quantile of age and size of the firm (two quantiles each) are part of the same group $g = g(j) = g(j')$. Size of firm is defined as the number of full-time equivalents. Both characteristics are 2021 values (year preceding estimation). When resulting groups represent less than 10 heat pumps for the studied period for at least one income category, they are aggregated with the neighboring hexagon of the same size and age quantiles. For a given g, we consider mean product characteristics, including for price and distance to housing, so that all products share the same characteristics and $s_{cgm} = n_{gm}^p \cdot s_{cjm}, g(j) = g$, where n_{gm}^p denotes the number of "identical" products in group g. The δ vector to be estimated therefore has $\sum_m \sum_c G_m = 3 \cdot G$ components, where we denote by G_m the number of groups of products per market and by G the overall number of groups in the sample.

Concentrating out mean utilities For a given value of θ_2, we search for the δ_{cm} vector in each income category-market that ensures that observed market shares equal predicted market shares for all groups of products:

$$\begin{aligned} s^d_{cm} &= s_{cm}(\delta_{cm}, \theta_2) \\ \iff \delta_{cm} &= s^{-1}_{cm}(s^d_{cm}, \theta_2) \end{aligned}$$

Berry et al. (1995) shows that this vector is unique and can be recovered by a fixed point algorithm. For a nested logit, Grigolon and Verboven (2014) show that the following algorithm is a contraction mapping converging toward this unique $\hat{\delta}(\theta_2)$:

1. Initialize $\delta^0 = \log(s^d) - \log(s_0)$.
2. Compute predicted shares $\hat{s}(\delta^{(t)})$.
3. Update $\delta^{(t+1)} = \delta^{(t)} + (1-\sigma)(\log(s^d) - \log(\hat{s}(\delta^{(t)})))$.
4. Repeat until convergence.

In practice, to avoid logs for computational reasons and following Nevo (2000), we iterate over $\tilde{\delta}^{(t)} = \exp(\delta^{(t)})$, so that $\tilde{\delta}^{(t+1)} = \tilde{\delta}^{(t)} \circ (s^d \oslash \hat{s}(\delta^{(t)}))^{1-\sigma}$. Convergence is obtained when $\frac{\|\tilde{\delta}^{(t+1)} - \tilde{\delta}^{(t)}\|}{\|\tilde{\delta}^{(t)}\|} \leq 10^{-10}$.

Recovering marginal costs In matrix form, profit maximization implies the following system of $3G$ first-order conditions (number of groups of products across markets times the number of income categories) that allows to recover marginal costs:

$$S + \Delta(P + B - C) = 0 \tag{A.1}$$

$$C = P + B + \Delta_p^{-1} S \tag{A.2}$$

where $S \in \mathbb{R}^{3G\times 1}$, $\Delta_p \in \mathbb{R}^{3G\times 3G}$, $P \in \mathbb{R}^{3G\times 1}$, $B \in \mathbb{R}^{3G\times 1}$, $C \in \mathbb{R}^{3G\times 1}$ are respectively the matrices of observed market shares, derivative of market shares with respect to prices, consumer prices, subsidies, marginal costs. We assumed that firms are always single-product firms. It means that all non-diagonal terms of Δ are zero. It also explains why we can solve for first-order conditions at the group of product level, as any representative firm of the group behaves in the same way in its pricing decision. Diagonal coefficients of Δ are given by $\frac{\partial s_{cjm}}{\partial p_{cjm}}$, where:

$$\frac{\partial s_{cjm}}{\partial p_{c'j'm'}} = \frac{1}{n_m} \sum_{i,c(i)=c} \phi_{ij'm} s_{ijm} \left[\frac{\mathbb{1}_{j=j'} - \sigma s_{ij'|in}}{1-\sigma} - s_{ij'm} \right] \tag{A.3}$$

$$\phi_{ijm} = -\alpha(I_i - p_{c(i)j'm})^{\lambda - 1} \tag{A.4}$$

where $s_{ij'|in}$ denotes the choice probability conditional on being in the inside nest.

The main assumption we make for identification of marginal costs is that local markets are independent. We argued above that these markets are indeed local. However, their limits are blurry. In particular, a firm might be active at the limit of two markets t, t' with a similar pricing strategy that depends on demand parameters of the two markets. We choose relatively large local markets (*départements*) to limit these frontier effects.

Method of moments estimator To estimate $\hat{\theta}_2^d$, we consider three sets of moments. The first is $K^\mu = 14$ micro-moments h_k^μ that target the mean observed individual characteristics among households that take up the subsidy. The second is $K^S = 2$ supply moments h_c^S that target differentials in marginal costs within products across income categories. The third is $K^Z = 2$ aggregate IV moments h_k^Z that use exogenous instruments $z^{Z,k}$ to identify the within-nest share parameter σ:

$$h_k^\mu = \frac{1}{N} \sum_{m=1}^{M} \sum_{i=1}^{n_m} \sum_{g=1}^{G_m} z_{igm}^{\mu,k} (s_{igm} - \mathbb{1}_{igm}) \tag{A.5}$$

$$h_c^S = \frac{1}{G} \sum_{m=1}^{M} \sum_{g=1}^{G_m} (mc_{cgm} - mc_{c+1,gm}) \tag{A.6}$$

$$h_k^Z = \frac{1}{3G} \sum_{m=1}^{M} \sum_{c=1}^{3} \sum_{g=1}^{G_m} z_{cgm}^{Z,k} \chi_{cgm} \tag{A.7}$$

The two instruments ($k = 1, 2$) are constructed from 2021 firm-level data and are therefore predetermined with respect to 2022 demand residuals. Instrument $z_{cgm}^{Z,1}$ is the share of FTE of group g in market m, weighted by inverse mean distance of housing to the group and normalized to sum to one within each market-income cell. Instrument $z_{cgm}^{Z,2}$ is the ratio of the density of eligible housing within a 50 km radius of the group to the number of competing groups in the same geographic hexagon. These instruments are relevant for σ because they shift the relative attractiveness of groups within a market without directly affecting outside-option utility.[29]

Our method of moments estimator therefore minimizes:

$$\hat{\theta}_2^d = \arg\min_{\theta_2} H^\top W H \tag{A.8}$$

where W is the optimal weighting matrix as defined in Conlon and Gortmaker (2025) and

[29] Both instruments are excludable from the demand equation because they are based on supply-side firm characteristics (firm size proxied by FTE, and spatial concentration) that do not directly enter housing or household utility. The first-stage F-statistic exceeds 1,400, confirming strong relevance.

$H = \begin{pmatrix} h^Z(\delta(\theta_2), \theta_2)) \\ h^S(\delta(\theta_2), \theta_2)) \\ h^\mu(\delta(\theta_2), \theta_2)) \end{pmatrix}$ is the column vector of 18 moments.

In practice, we solve for this minimization problem using the non-linear programming solver "fminunc" in Matlab, using a trust-region algorithm with objective gradient. The number of iterations is limited to 30, step tolerance is set to 10^{-6} and termination tolerance on the function value to 10^{-14}.[30]

Analytical gradient The derivative of micro moments $h^\mu(\delta(\theta_2), \theta_2))$ with respect to θ_2 parameters is given by $\frac{1}{N} Z^\top \frac{d\xi}{d\theta_2} = \frac{1}{N} Z^\top \frac{ds}{d\theta_2}$ where s is the vector made of individual choice probabilities. The derivative of individual choice probabilities with respect to the parameters of interest is given by

$$\frac{d s_{igm}}{d\theta_{2,k}} = \frac{1}{N} \left(\frac{\partial s_{igm}}{\partial \theta_{2,k}} + \sum_{g'=1}^{G_m} \frac{\partial s_{igm}}{\partial \delta_{g'm}} \frac{\partial \delta_{g'm}}{\partial \theta_{2,k}} \right)$$

We denote by $D^s_{\theta_2}$ the $\mathbb{R}^{N^{full} \times (R+2)}$ matrix of partial derivatives of predicted choice probabilities with respect to demand parameters:

$$\frac{\partial s_{igm}}{\partial \theta_{2,k}} = \begin{cases} s_{igm} \left(\dfrac{\tilde{z}^k_{igm}}{1-\sigma} - \dfrac{\sigma}{1-\sigma} \sum_{g'=1}^{G_m} \tilde{z}^k_{ig'm} s_{ig'm|in} - \sum_{g'=1}^{G_m} \tilde{z}^k_{ig'm} s_{ig'm} \right) & \text{if } \theta_{2,k} \neq \sigma \\ s_{igm} \left(\begin{array}{l} \dfrac{\tilde{U}_{igm}}{(1-\sigma)^2} - \dfrac{\sigma}{1-\sigma} \sum_{g'=1}^{G_m} \dfrac{\tilde{U}_{ig'm}}{1-\sigma} s_{ig'm|in} - \sum_{g'=1}^{G_m} \dfrac{\tilde{U}_{ig'm}}{1-\sigma} s_{ig'm} \\ -(1 - s^{in}_{im}) \log\left(\sum_{g'=1}^{G_m} n^p_{g'm} \exp\left(\dfrac{\tilde{U}_{ig'm}}{1-\sigma} \right) \right) \end{array} \right) & \text{else} \end{cases}$$

$$\tilde{z}^k_{igm} = \begin{cases} y^k_i & \text{if } \theta_{2,k} = \beta^k_y \\ d_{igm} & \text{if } \theta_{2,k} = \beta_d \\ \dfrac{1}{\lambda} \left((I_i - p_{c(i)gm})^\lambda - I^\lambda_i \right) & \text{if } \theta_{2,k} = \alpha \\ \dfrac{\alpha}{\lambda^2} \left((\lambda \log(I_i - p_{c(i)gm}) - 1)(I_i - p_{c(i)gm})^\lambda - (\lambda \log(I_i) - 1) I^\lambda_i \right) & \text{if } \theta_{2,k} = \lambda \end{cases}$$

[30] For numerical stability, the three structural parameters $(\alpha, \lambda, \sigma)$ are bounded and optimized over unconstrained logistic transforms $(\tilde{\alpha}, \tilde{\lambda}, \tilde{\sigma}) \in \mathbb{R}^3$: $\alpha = 8 / (1 + e^{-\tilde{\alpha}}) \in (0, 8)$, $\lambda = 3 / (1 + e^{-\tilde{\lambda}}) - 2 \in (-2, 1)$, $\sigma = 0.97 / (1 + e^{-\tilde{\sigma}}) \in (0, 0.97)$. Gradient formulas in this appendix are expressed in terms of the structural parameters $(\alpha, \lambda, \sigma)$ for clarity; the chain-rule Jacobian $\partial\theta_k / \partial\tilde{\theta}_k$ is applied for each transformed parameter.

where $s_{igm|in}$ and s_{im}^{in} respectively denote the probability of i choosing g conditional on choosing an inside option, and of choosing the inside option. D_{δ}^{s}, the $\mathbb{R}^{N^{full}\times(3G)}$ matrix of derivatives of choice probabilities with respect to δ is a block-diagonal matrix given by:

$$\frac{\partial s_{igm}}{\partial \delta_{cg'm'}} = \begin{cases} 0 & \text{if } m \neq m' \text{ or } c(i) \neq c \\ -\frac{1}{1-\sigma} s_{igm}(\sigma s_{ig'm|in} + (1-\sigma) s_{ig'm}) & \text{else if } g \neq g' \\ \frac{1}{1-\sigma} s_{igm}(1 - (\sigma s_{igm|in} + (1-\sigma) s_{igm})) & \text{otherwise.} \end{cases}$$

Nevo (2000) notes, from the implicit function theorem, that the $\mathbb{R}^{(3G)\times(R+2)}$ matrix of derivative of δ with respect to θ_2 has an analytical form given by:

$$D_{\theta_2}^{\delta} = -(D_{\delta}^{\bar{s}})^{-1} D_{\theta_2}^{\bar{s}}$$

where $D_{\delta}^{\bar{s}} \in \mathbb{R}^{(3G)\times(3G)}$ is the matrix of mean derivative of choice probabilities with respect to δ:

$$\frac{\partial s_{cgm}}{\partial \delta_{c'g'm'}} = \frac{1}{n_{cm}} \sum_{i,c(i)=c} \frac{\partial s_{igm}}{\partial \delta_{c'g'm'}}$$

and $D_{\theta_2}^{\bar{s}} \in \mathbb{R}^{(3G)\times(R+2)}$, the matrix of mean derivative of choice probabilities with respect to θ_2:

$$\frac{\partial s_{cgm}}{\partial \theta_{2,k}} = \frac{1}{n_{cm}} \sum_{i,c(i)=c} \frac{\partial s_{igm}}{\partial \theta_{2,k}}$$

The derivatives of the micro moments are therefore given by:

$$\frac{1}{N} Z^{\top} (D_{\theta_2}^{s} - D_{\delta}^{s} (D_{\delta}^{\bar{s}})^{-1} D_{\theta_2}^{\bar{s}})$$

Note that the only change in the way the gradient is computed, compared to Nevo (2000) relates to the presence of price discrimination: choice probabilities s_{igm} are independent of $\delta_{cg'm}$ as soon as $c(i) \neq c$. We therefore aggregate derivative of choice probabilities at the market-income category level, which amounts to considering separate income categories as separate sub-markets in the estimation of the model.

The derivative of supply moments with respect to θ_2 parameters is obtained by noticing that:

$$
\begin{aligned}
mc_{cgm} &= p_{cgm} + b_c + \frac{\sum\limits_{i,c(i)=c} s^u_{igm}}{\sum\limits_{i,c(i)=c} \phi_{igm} s^u_{igm} \left[-s^u_{igm} + \frac{1-\sigma s^u_{igm}|in}{1-\sigma} \right]} \\
&= p_{cgm} + b_c + \frac{\sum\limits_{i,c(i)=c} s_{igm}}{\sum\limits_{i,c(i)=c} \phi_{igm} s_{igm} \left[-s^u_{igm} + \frac{1-\sigma s^u_{igm}|in}{1-\sigma} \right]}
\end{aligned}
$$

where $\phi_{igm} = -\alpha(I_i - p_{cgm})^{\lambda-1}$ is the derivative of indirect utility for any product of g with respect to its own price, $s^u_{igm} = \frac{s_{igm}}{n^p_{gm}}$ and $s^u_{igm}|in = \frac{s_{igm}|in^p}{n_{gm}}$ are respectively the choice probabilities for any representative product of g and this probability conditional on the choice of an inside product. Using the chain rule and former results of derivation of s with respect to θ_2 yields the derivative of $h^S(\delta(\theta_2), \theta_2))$ with respect to θ_2.

The derivative of the aggregate IV moments h^Z with respect to θ_2 follows from the chain rule applied to the demand residual $\chi_{cgm} = \delta_{cgm} - X_{cgm}\hat{\theta}_1$, where $\hat{\theta}_1 = (X^\top X)^{-1} X^\top \delta(\theta_2)$ is the OLS estimator of the linear parameters recovered in the second step:

$$
\frac{dh^Z_k}{d\theta_2} = \frac{1}{3G}(z^{Z,k})^\top \frac{d\chi}{d\theta_2} = -\frac{1}{3G}(z^{Z,k})^\top M_X (D^{\bar{s}}_{\delta})^{-1} D^{\bar{s}}_{\theta_2}
$$

where $M_X = I - X(X^\top X)^{-1} X^\top$.

C.3 Consistency and inference

Following Berry et al. (2004), consistency is obtained as long as $N \to \infty$. Newey and McFadden (1994) states that under i.i.d. observations across households, and provided (i) θ^d is in the interior of Θ which is compact, (ii) $\hat{h}$ is continuously differentiable in a neighborhood $\mathcal{N}$ of θ^d, (iii) $\mathbb{E}[h(\theta^d)] = 0$ and $\mathbb{E}[\|h(\theta^d)\|^2]$ is finite, (iv) $\mathbb{E}[\sup\limits_{\theta \in \mathcal{N}} \|\nabla h(\theta)\|] < \infty$, (v) $\Delta^\top W \Delta$ is non-singular where $\Delta = \mathbb{E}[\nabla h(\theta^d)]$, then:

$$
\sqrt{3G}(\hat{\theta} - \theta^d) \xrightarrow{d} N(0, (\Delta^\top W \Delta)^{-1} \Delta^\top W \Omega W \Delta (\Delta^\top W \Delta)^{-1})
$$

where $3G$ is the number of aggregate observations (groups of products) across submarkets (markets-income categories), and Ω is the covariance matrix of the vector of moments. As usual, we use $W = \hat{\Omega}^{-1}$ as the optimal weighting matrix, so that asymptotic variance is given by:

$$\sqrt{3G}(\hat{\theta} - \theta^d) \xrightarrow{d} N(0, (\Delta^\top W \Delta)^{-1})$$

Conlon and Gortmaker (2025) provide several thought experiments to derive Ω in the case of combined micro and macro moments. We consider the case where the number of submarkets $3M \to \infty$ and is equal to the number of markets for which we have micro data. Since micro data is generated conditional on aggregate data, Conlon and Gortmaker (2025) show that the covariance between aggregate and micro moments is zero. Ω can therefore be written as a block-diagonal matrix:

$$\Omega = \begin{bmatrix} \Omega^{ZS} & 0 \\ 0 & \Omega^{\mu} \end{bmatrix}$$

Ω^{ZS} is the joint covariance matrix of the aggregate moments (IV moments of the first step and OLS moments of the second step) and supply moments.

$$\Omega^{ZS} = \mathbb{E}[G_m] \cdot \mathbb{E}[1/G_m] \begin{bmatrix} \Omega^{Z} & \Omega^{ZS}_{\times} \\ (\Omega^{ZS}_{\times})^\top & \Omega^{S} \end{bmatrix}$$

Sub-block Ω^Z is the sandwich covariance of the aggregate moments:

$$\Omega^{Z}_{kl} = \frac{1}{3G} \sum_{m=1}^{M} \sum_{c=1}^{3} \sum_{g=1}^{G_m} z^{Z,k}_{cgm}\, z^{Z,l}_{cgm}\, \chi^2_{cgm}$$

The supply sub-block Ω^S is the covariance of demeaned marginal costs:

$$\Omega^{S} = \frac{1}{3G} \begin{bmatrix} \sum\limits_{gm} \tilde{\text{mc}}^2_{1gm} + \tilde{\text{mc}}^2_{2gm} & \sum\limits_{gm} -\tilde{\text{mc}}^2_{2gm} \\ \sum\limits_{gm} -\tilde{\text{mc}}^2_{2gm} & \sum\limits_{gm} \tilde{\text{mc}}^2_{2gm} + \tilde{\text{mc}}^2_{3gm} \end{bmatrix}$$

where $\tilde{mc}_{cgm} = mc_{cgm} - \frac{1}{G} \sum\limits_{gm} mc_{cgm}$.

The cross-covariance block $\Omega^{ZS}_{\times}$ is a matrix whose (k, c) entry captures the covariance between the k-th aggregate moment and the c-th supply moment:

$$(\Omega^{ZS}_{\times})_{kc} = \frac{1}{3G} \sum_{m=1}^{M} \sum_{c'=1}^{3} \sum_{g=1}^{G_m} z^{Z,k}_{c'gm}\, \chi_{c'gm}\, \tilde{h}^{S}_{cgm}$$

where $\tilde{h}^S_{cgm} = mc_{cgm} - mc_{c+1,gm} - \frac{1}{G}\sum_{g'm'}(mc_{cg'm'} - mc_{c+1,g'm'})$ is the demeaned product-level contribution to supply moment c.

Ω^μ is the covariance matrix of micro moments. Following Conlon and Gortmaker (2025):

$$\hat{\Omega}^\mu_{pq} = \mathbb{E}[G_m] \cdot \mathbb{E}[1/n_{cm}] \cdot \frac{1}{N}\sum_{igm} \mathcal{s}_{igm}\left(z^p_{igm} - z^p(\hat{\theta})\right)\left(z^q_{igm} - z^q(\hat{\theta})\right)$$

where $z^p(\hat{\theta}) = \frac{1}{N}\sum_{igm} \mathcal{s}_{igm}(\hat{\theta})\, z^p_{igm}$.[31]

C.4 Counterfactuals simulation

With price discrimination In scenarios where price discrimination is allowed, firms implicitly fix consumer prices for each income category and representative product j of a given group, $p_{cjm} = q_{cjm} - b_c$, where $b_c > 0$ is the subsidy level for income category c, q_{cjm} is the transaction (pre-subsidy) price. Each representative firm of its group optimizes profit with respect to prices of its product:

$$\pi_{jm} = \sum_{c=1}^{3}(p_{cjm} + b_c - mc_{jm})\mathcal{s}_{cjm}(p_m) \tag{A.9}$$

First-order conditions are of the form:

$$FOC_{cjm} : \frac{\partial \pi_{jm}}{\partial p_{cjm}} = \mathcal{s}_{cjm} + \frac{\partial \mathcal{s}_{cjm}}{\partial p_{cjm}}(p_{cjm} + b_c - mc_{jm}) = 0 \tag{A.10}$$

This yields the following system of $3G$ first-order conditions:

$$S + \Delta(P + B - C) = 0$$

Note that S and Δ are also non-linear functions of P so that there is no closed-form solution to this system of $3G$ equations. We also impose a government spending constraint $G^\star$ that needs to be matched for all counterfactuals. To solve for equilibrium prices and subsidies, we instead adopt the following nested fixed point algorithm:

[31] Conlon and Gortmaker (2025) recommend weighting each individual observation by $1/n_{cm}$, the inverse of the number of individuals in market-income cell (m, c), so that each cell contributes equally. For consistency with the way micro moments h^μ themselves are computed in the first step (using global weights $w_i = 1/N$), we retain $w_i = 1/N$ throughout.

1. Outer fixed point: initialize $B^{(0)} = B^{baseline}$.

2. Inner fixed point: initialize $P^{(0)} = P^{baseline}$

3. For each iteration t of the inner fixed point, compute individual choice probabilities, resulting aggregate market shares $S(P^{(t)})$ and matrix of derivatives with respect to prices $\Delta(P^{(t)})$ using this price vector.

4. Update $P^{(t+1)} = -B + C - \Delta(P^{(t)})^{-1} \cdot S(P^{(t)})$.

5. For each iteration of the outer fixed point k, repeat until convergence at $t = t(k)$.

6. For each iteration of the outer fixed point k, compute implied government spending $G^{(k)} = S(P^{(t(k))})^{\top} \cdot B^{(k)}$

7. Update $B^{(k+1)} = B^{(k)} + (G^{\star} - G^{(k)})/\tilde{S}^k$ where $\tilde{S}$ is the sum of market shares (overall take-up) across all markets and groups.

8. Repeat until convergence.

Without price discrimination When firms are not allowed to price discriminate, firms fix a set of G transaction prices $q_{jm} = p_{cjm} + b_c$. Note that consumer prices are still income category-specific, since subsidies also are in the general case. Each firm optimizes the sum of profit across income categories with respect to its transaction prices:

$$\pi_{fm} = (q_{jm} - mc_{jm}) \sum_{c=1}^{3} s_{cjm}(q_m, b_c) \tag{A.11}$$

First-order conditions are of the form:

$$FOC_{jm} : \sum_{c=1}^{3} s_{cjm}(q_m, b_c) + (q_{jm} - mc_{jm}) \sum_{c=1}^{3} \frac{\partial s_{cjm}}{\partial q_{jm}} = 0 \tag{A.12}$$

This yields the following system of G first-order conditions:

$$\tilde{S} + \tilde{\Delta}(Q - C) = 0$$

where $\tilde{S} \in \mathbb{R}^{G \times 1}$ and $\tilde{\Delta} \in \mathbb{R}^{G \times G}$ are respectively the matrix of aggregate market shares across income categories $s_{jm} = \sum_{c} s_{cjm}$ and derivative of these aggregate market shares with respect to prices:

$$
\begin{aligned}
\frac{\partial s_{jm}}{\partial q_{j'm'}} &= \sum_c \frac{\partial s_{cjm}}{\partial q_{j'm'}} \\
&= \sum_c \frac{\partial s_{cjm}}{\partial p_{cj'm'}} \frac{\partial p_{cj'm'}}{\partial q_{j'm'}} \\
&= \sum_c \frac{\partial s_{cjm}}{\partial p_{cj'm'}}
\end{aligned}
$$

Again, $\tilde{S}$ and $\tilde{\Delta}$ are non-linear functions of Q so that there is no closed-form solution to this system of J equations. We therefore estimate Q and resulting P from the following fixed point algorithm:

1. Outer fixed point: initialize $B^{(0)} = B^{baseline}$.
2. Inner fixed point: initialize $Q^{(0)} = Q^{baseline}$, where $Q_{jm}^{baseline} = \frac{\sum_c s_{cjm} q_{cjm}^{baseline}}{\sum_c s_{cjm}}$.
3. For each iteration t of the inner fixed point, compute consumer prices, individual choice probabilities, resulting aggregate market shares $S(Q^{(t)})$ and matrix of derivatives with respect to prices $\tilde{\Delta}(Q^{(t)})$ using this price vector.
4. Update $Q^{(t+1)} = -\tilde{\Delta}(Q^{(t)})^{-1}\tilde{S}(Q^{(t)}) + C$.
5. For each iteration of the outer fixed point k, repeat until convergence at $t = t(k)$.
6. For each iteration of the outer fixed point k, compute implied government spending $G^{(k)} = S(Q^{(t(k))})^\top \cdot B^{(k)}$
7. Update $B^{(k+1)} = B^{(k)} + (G^\star - G^{(k)})/\tilde{S}^k$ where $\tilde{S}$ is the sum of market shares (overall take-up) across all markets and groups.
8. Repeat until convergence.

C.5 Robustness: partial price discrimination

Our simplification is motivated by the complexity of modeling selection into mandating. First, it requires taking a stance on the selection process itself: either households choose between two products for each firm — mandated and non-mandated — or firms choose to whom they offer the mandated option. Second, it requires taking a stand on whether firms can still discriminate when they are not mandated.[32] Either choice would be arbitrary. Third, conditional on both,

[32] Our reduced-form estimates do not allow us to reject that price discrimination also occurs for some non-mandated applications: a firm that is not mandated may still learn the income category of the household through the quotation process. Assuming that mandating implies discrimination deterministically, and non-mandating rules it out, therefore understates the share of transactions that are discriminated.

the firm now offers up to four products — three mandated variants, priced by income category, and one non-mandated variant, priced uniformly across categories — which all interact in its profit-maximization problem. Given the supply-side moment we impose on marginal costs, this removes any clean expression for the analytical gradient, while estimation with a numerical gradient would be impossible given our sample size and number of products.

The fact that our model does not allow for households to not reveal their income category to firms has two main consequences for our estimates. First, we attribute to all observed transaction prices a component stemming from price discrimination, even though some transactions are priced uniformly with regard to income. Second, we identify the effect of price discrimination from differences in average prices across income categories; those differences are, in the data, averages between discriminated and non-discriminated transactions, and they are therefore typically smaller than the differences between discriminated transactions alone. Each channel pushes in a different direction, so that we can, overall, either under- or over-estimate the contribution of price discrimination to price differentials.

In addition, selection into "mandating" can bias our results on the effects of price discrimination on price differentials across income categories. If "mandating" is rather something with which firms target households, and if they target households with an offer to be mandated in order to maximize their profits, then they would target the most price-sensitive among those who receive a premium and target the least price-sensitive among those who pay higher prices; this would increase the share of price dispersion attributable to price discrimination. Conversely, mandating could be the result of households' demand. Those who value the most having the firm deal with the paperwork may plausibly be the most liquidity constrained, or those whose opportunity cost of time is the highest. In both cases, selection would happen on income rather than on the expected returns from discrimination, but in a different direction. We do not attempt to sign or quantify the resulting bias, but we bound it instead.

In order to examine the consequences of our simplification, we carry out the following exercise. We simulate data according to data generating processes (DGPs) that we calibrate on our results: we take our estimated demand parameters and the marginal costs recovered from the first-order conditions as the truth. We also use the observed market structure, with the set of firms, the population of households and their respective locations, and the subsidy schedule. Compared to our estimated model, we add in the DGP an additional mandating or price-discrimination step, whereby a share π_c of households get assigned their income-category-specific prices, and the rest get assigned a uniform price common to all categories. Marginal costs are constant, and we assign price discrimination exogenously, so that firms' profits are additive across the price-discriminated and non-discriminated segments. For each DGP, we solve the pricing equilibrium in both segments, and compute the average price for households who actually purchase in either segment for a given product and income category, which corresponds to our treatment of the observed data in the main estimation procedure. For a given DGP, we then estimate our initial model, which is blind to the mandating step and assumes full price discrimination. Similarly to Section 6.3, we compute the contribution of price discrimination recovered in the estimation, and we compare it to the DGP's truth. We use this difference as the relevant measure of bias for our model: in the end, we are interested in whether we mischaracterize the direction in which price discrimination pushes prices, and in the magnitude of this price discrimination contribution to total dispersion.

We study the behavior of this measure of bias under different mandating assignment regimes. We first look at random assignment, with π_c varying from 100% to 0%. The case where $\pi_c = [100\%, 100\%, 100\%]$ corresponds to full price discrimination: in this case, the estimated model is correctly specified, so that we should recover our results from Section 6.3, and a null bias. Setting $\pi_c = [0\%, 0\%, 0\%]$ corresponds to no price discrimination, so that any contribution of price discrimination to price dispersion is a spurious artifact. We let π_c take a few different values, including the triplet $[48\%, 39\%, 33\%]$, which are the observed mandated shares in each of the three categories.[33] We then move on to non-random mandating assignment rules. From our initial results, we derive a household-specific price-sensitivity measure, the marginal disutility of an additional euro in price: $\left|\frac{\partial U_i}{\partial p}\right| = \alpha_i (I_i - p)^{\lambda - 1}$. We rank households according to this score, within market and income category. We study four cases. In the first two, selection into revealing their income category to firms happens for the π_c lowest-score households in each category, and for the π_c highest-score households. These correspond broadly to the selection on income mechanisms of opposite signs described above. In the third, the less price-sensitive households are mandating in the more subsidized categories, and the most price-sensitive in the less subsidized categories; this corresponds to the profit-maximizing targeting case mentioned above.[34] In the fourth, the least price-sensitive households are mandating in the more subsidized categories, and the most price-sensitive in the less subsidized categories.

Two properties of the simulated environments frame the comparison. First, the true contribution is almost exactly linear in π: $906 \times \pi$ reproduces every uniform-π scenario to within 1.1%. Second, the true counterfactual differential is invariant across scenarios ($-3{,}553$ to $-3{,}554$, a spread of one euro), as it must be, since removing discrimination should not depend on how many households were discriminated against in the first place. We measure the contribution in the simulated data exactly as we do in Section 6.3: the baseline is the observed differential, not the model's fitted baseline.

Panel C of Table A.6 reports the bias on the price differential. The estimated counterfactual differential is too wide by €29 to €52 across the entire range, with no trend in π and no maximum at intermediate values; relative to a counterfactual differential of roughly €3,554, the error is close to 1% everywhere. At $\pi = 0$, where the true contribution is exactly zero, the estimated contribu-

[33]These observed shares, ie observed on applicants, are used here as calibration for the population shares.

[34]More precisely, the Very Poor category gets highest, the Poor category is random, the Intermediate gets the lowest.

tion is €37: the artifact the procedure generates out of nothing is an order of magnitude smaller than what we estimate on the data. At the empirical mandated shares the procedure recovers 87% of the true contribution, understating rather than overstating the role of discrimination. Table A.7 reports the same exercise on take-up, the object of the counterfactual in Section 6.3: there the procedure recovers the Very Poor contribution almost exactly (223 against 218 installations) while overstating the Intermediate one by 17% (-192 against -164). The two measures bound the two quantities the paper reports, and they do not err in the same direction.

At $\pi = 1$ the estimated model is correctly specified, yet the estimated counterfactual still differs from the truth by €29. This residual is numerical rather than substantive, and it sets the resolution of the exercise: differences of that order are not interpretable, and the misspecification-attributable component at interior π is correspondingly €8 to €23.

Table A.6: Misspecification exercise: prices and the price differential

	Random assignment, uniform π					Empirical
	1.00	0.75	0.50	0.25	0.00	shares
Panel A. Simulated consumer price (€)						
Very poor	5,220	5,365	5,515	5,666	5,820	5,545
Poor	7,220	7,224	7,228	7,233	7,237	7,251
Intermediate	9,679	9,602	9,524	9,449	9,374	9,499
Panel B. Counterfactual differential, Very poor − Intermediate (€)						
True	−3,553	−3,553	−3,554	−3,554	−3,554	−3,554
Estimated	−3,582	−3,592	−3,593	−3,595	−3,591	−3,606
Panel C. Contribution of price discrimination to the differential (€)						
True	−906	−684	−455	−229	0	−400
Estimated	−877	−645	−416	−188	+37	−348
Bias	+29	+39	+39	+41	+37	+52
Panel D. Mandated − non-mandated price gap (€)						
Very poor	—	−601	−599	−598	—	−637
Poor	—	−19	−19	−12	—	−51
Intermediate	—	+305	+302	+308	—	+273

NOTE : Columns are the data generating processes of Table A.7. Panel A reports the consumer price observed by the econometrician, which averages the discriminated and uniform segments over purchasing households. Panel B compares the true counterfactual differential, computed at the parameters that generated the data, to the one recovered by the estimated model. Panel C measures the contribution of price discrimination as the observed differential minus the counterfactual differential, so that the baseline is the data in both the true and the estimated column. The gap between segments in Panel D is undefined at $\pi = 1$ and $\pi = 0$, where one segment is empty. At $\pi = 1$ the estimated model is correctly specified; the residual €29 bias there is numerical and sets the resolution of the exercise. SOURCE: Fideli, MPR.

Table A.7: Misspecification exercise: bias from assuming universal price discrimination

		Random assignment, uniform π					Empirical shares	
	DGP truth	1.00	0.75	0.50	0.25	0.00	Random	Selected
Panel A. Estimated demand parameters								
α (price sensitivity)	0.1534	0.1533	0.1471	0.1483	0.0697	0.0449	0.1326	0.1110
λ (income effect)	0.6823	0.6821	0.6935	0.7238	0.7602	0.7740	0.7370	0.7280
σ (nest)	0.3317	0.3327	0.3293	0.3422	0.3136	0.1339	0.3424	0.3289
β_d (distance)	−0.578	−0.444	−0.446	−0.435	−0.464	−0.596	−0.430	−0.450
Implied $\lvert\partial U_i/\partial p\rvert$	0.0649	0.0648	0.0641	0.0702	0.0364	0.0244	0.0651	0.0531
Panel B. Recovered marginal cost (€)								
Mean, cell level	4,474	4,476	4,650	5,874	−1,453†	−15,041†	5,480	3,183
Panel C. Contribution of price discrimination to take-up								
True (data generating process)								
Very poor	—	490	368	244	123	0	218	303
Poor	—	−49	−37	−24	−13	0	−34	−58
Intermediate	—	−402	−302	−200	−100	0	−164	−239
Estimated								
Very poor	—	489	365	266	72	−4	223	248
Poor	—	−49	−42	−35	−10	−2	−44	−48
Intermediate	—	−402	−316	−244	−69	−4	−192	−209
Bias (estimated − true)								
Very poor	—	−1	−3	+22	−51	−4	+5	−55
Poor	—	0	−5	−11	+3	−2	−10	+10
Intermediate	—	0	−14	−44	+31	−4	−28	+30

NOTE : Each column is a data generating process calibrated on the estimated model of Section 6.3: demand parameters and marginal costs are taken as the truth, and the observed market structure and subsidy schedule are held fixed. A share π of households is assigned income-category-specific prices; the remainder faces a single price common to all categories. The first five columns assign mandate status at random within each market and income category. The last two use the mandated shares observed in the data $(0.48, 0.39, 0.33)$, assigning them at random or to the most price-sensitive households of each category; Table A.8 reports the full set of assignment rules. The estimated model is then applied to each simulated dataset and is blind to the mandating step. At $\pi = 1$ the estimated model is correctly specified and must return a null bias; at $\pi = 0$ no price discrimination takes place, so any estimated contribution is a pure artefact. The implied marginal disutility of price is $\alpha_i(I_i - p)^{\lambda-1}$ evaluated at the median of $I_i - p$. Marginal costs are constrained to be equal across income categories by the supply moment, so only their level is informative. Contributions are the difference in take-up between the discriminatory and uniform-pricing counterfactuals, in units of installations. † Recovered marginal costs are negative, and therefore economically meaningless: once price sensitivity is sufficiently attenuated, the implied mark-up exceeds the transaction price. Note that the GMM objective barely moves across the whole range, so the overidentifying restrictions give no warning of this; the implausibility of the recovered supply side does. SOURCE: Fideli, MPR.

Table A.8: Selection into mandating: effect on the true contribution of price discrimination

Mandate assignment rule	Contribution to take-up: Very poor	Poor	Interm.	Implied price discontinuity: TMO/MO	MO/INT
Random	217	-34	-163	604	312
Most price-sensitive, all categories	303	-58	-239	616	131
Least price-sensitive, all categories	121	-10	-88	605	472
Most sensitive where subsidised, least elsewhere	527	-60	-502	1,355	1,374
Least sensitive where subsidised, most elsewhere	-102	-16	$+186$	-300	-608
Observed in the data	—	—	—	≈ 640	≈ 450
Full discrimination ($\pi = 1$)	490	-49	-402	—	—

NOTE : Each row is a data generating process in which the mandated shares are held at their observed values $(0.48, 0.39, 0.33)$ and only the *composition* of the mandated group varies. Households are ranked within each market and income category by their marginal disutility of price, $\alpha_i(I_i - p)^{\lambda-1}$. Contributions are the difference in take-up between the realised mixed regime and a counterfactual in which no firm discriminates, at constant government spending, in units of installations. The implied price discontinuity is the transaction-price gap between mandated and non-mandated files, differenced across adjacent income categories, and is directly comparable to the estimates on mandated applications reported in Section 5. No estimation is involved: these are properties of the simulated environments, not of our procedure. SOURCE: Fideli, MPR.

C.6 CO_2 savings: methodology and assumptions

Our CO_2 estimates translate the take-up differences across counterfactuals into metric tons of CO_2 avoided over the lifetime of the installed heat pumps. This appendix details the methodology and its limitations.

From take-up to annual CO_2 savings per heat pump We build on Bernard et al. (2024), who estimate via staggered difference-in-differences that heat pump adoption in the UK causes, on average, a reduction of 9,351 kWh/year in gas consumption and an increase of 3,080 kWh/year in electricity consumption. These causal reduced-form estimates are the only empirical inputs into our carbon accounting. We apply them to French emission factors rather than the UK ones used in the original study.

For natural gas, we use the ADEME Base Carbone factor of 0.262 $kgCO_2$/kWh (PCI, combustion and upstream combined) (ADEME, 2024). This gives a gas-displacement saving of $9{,}351 \times 0.262 = 2.45\,tCO_2/\text{year}$.

For electricity, we use the heating-specific, seasonally-weighted factor of 79.5 gCO_2/kWh adopted in the French RE2020 building-energy regulation (Ministère de la Transition écologique, 2021). This factor is constructed by weighting the carbon content of the French electricity mix by month and by the seasonal consumption profile of electric heating. It captures the same economically meaningful idea as the marginal emission factor used by Bernard et al. (2024): heating load is concentrated in winter and peak hours, when the marginal generation mix is more carbon-intensive than the annual average. The added electricity consumption costs $3{,}080 \times 0.0795 = 0.24\,tCO_2/\text{year}$.

The net annual CO_2 saving per heat pump is therefore $2.45 - 0.24 \approx 2.21\,tCO_2/\text{year}$. Multiplied by a 20-year assumed heat pump lifetime – matched to the physical life of the heating system and consistent with Bernard et al. (2024) – this gives approximately 44 tCO_2 per heat pump over its lifetime. Applied to the 40,400 heat pumps in the observed scenario (column 4), this yields around 1,800 kt CO_2 avoided in total.

Limits of the exercise Several limitations should be borne in mind. First, we import the UK reduced-form ΔkWh estimates rather than estimating French equivalents, because no analogous quasi-experimental evidence exists for France.

Second, we apply the same emission factor regardless of income category. We do not have the data necessary to compute category-specific CO_2 abatement, as this would require matching heat pump takers to individual energy consumption records. In particular, we assume that inframarginal and marginal adopters, and adopters across income groups, avoid the same metric tons of CO_2 per heat pump. The resulting estimate is therefore best interpreted as an approximation.

Third, our 20-year horizon assumes a constant French electricity emission factor. Unlike the

UK grid studied by Bernard et al. (2024), whose decarbonization trajectory over 2024–2043 is a key driver of their lifetime estimate, the French electricity mix is already largely decarbonized (predominantly nuclear) and is not expected to decrease by a comparable multiple over the same horizon. A flat factor is therefore a reasonable simplification for the purpose of ranking counterfactuals, even if it omits potential changes in the French energy mix.

Because the per-unit CO_2 factor is the same across all counterfactual scenarios, differences in total CO_2 savings across columns are exactly proportional to differences in total take-up. This computation is used to give a rough estimate of the differentials in emissions abated across counterfactuals. We find that means-testing and price discrimination have an impact on environmental efficiency of the subsidy that is of a very small order of magnitude. This result is therefore unlikely to differ with alternative and more robust assumptions for energy consumption reduction and emission factors for the French context.

D Additional figures: renovation market

The figures in this section document the structure of the MPR program and the heat pump renovation market, complementing the descriptive evidence in the main text.

The MPR program covers 26 types of retrofit actions (Figure A.10). Heat pumps stand out as one of the most costly actions on average, with total transaction prices around €15,000 and an MPR subsidy covering €4,000 on average. Figure A.11 shows how take-up of the MPR subsidy (and its predecessor tax credit CITE) varies by income bracket. The shift from CITE to MPR in 2020 led to a redistribution of take-up toward lower-income households, consistent with the program's design.

On the supply side, RGE-certified firms are highly specialized: the cluster analysis in Figure A.12 shows that firms concentrate their MPR activity in a narrow set of action types, with heat pump installers rarely performing other renovation categories. The total number of active RGE firms declined between 2019 and 2021 before recovering (Figure A.13), suggesting capacity constraints that can amplify market power. The local nature of the market is documented in Figure A.14: half of all heat pump installations are performed by a firm located within 25 km of the household, supporting the local-market assumption of the structural model.

Figure A.10: Average subsidy and total renovation works costs, per renovation step

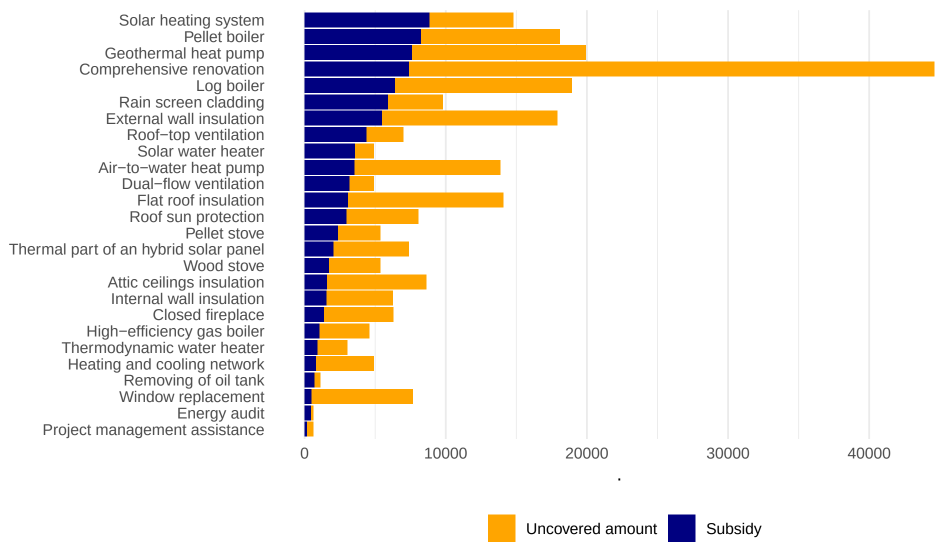

NOTE: This figure presents the average MPR subsidy amount (blue) and uncovered amount (orange) for each of the 26 retrofit actions covered by the scheme. These two amounts sum to the average total cost of a given renovation step. Sample is made of one observation per retrofit action, for all MPR files from 2020 to the first quarter of 2024.
SOURCE: MPR.

Figure A.11: Distribution of CITE and MPR spending by quantiles of equivalized disposable income

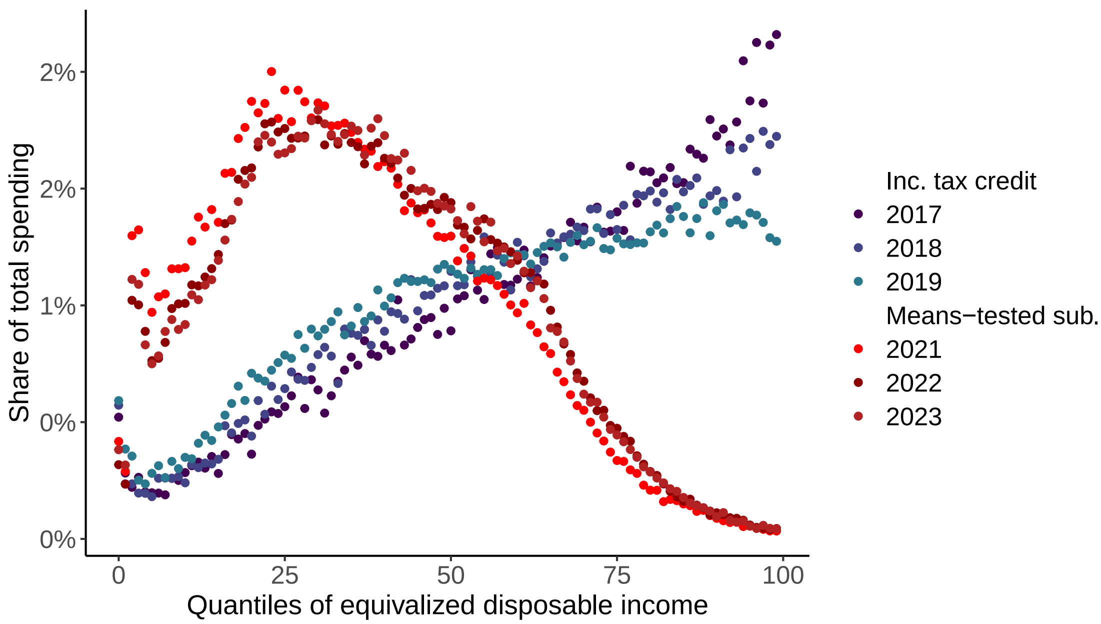

Note: This figure represents how total public spending on subsidized heat pumps is distributed across centiles of the disposable income distribution. In years 2017, 2018 and 2019 this corresponds to the CITE income tax credit spending on heat pumps; in 2021, 2022 and 2023, this corresponds to the MPR subsidy scheme for heat pumps. Disposable income is defined as the income after all direct taxes and transfers; equivalized disposable income is equal to disposable income divided by the number of consumption units, where the first adult in the household counts for one unit, each household member above 14 counts for 0.5 consumption unit, and each member below age 14 counts for 0.3 consumption unit, as per the definition of the French National Statistics Office (INSEE). Centiles are weighted by household size so that each centile contains the same number of individuals. Source: MPR, Fideli.

Figure A.12: Distribution of MPR activity across clusters for RGE companies, by main cluster

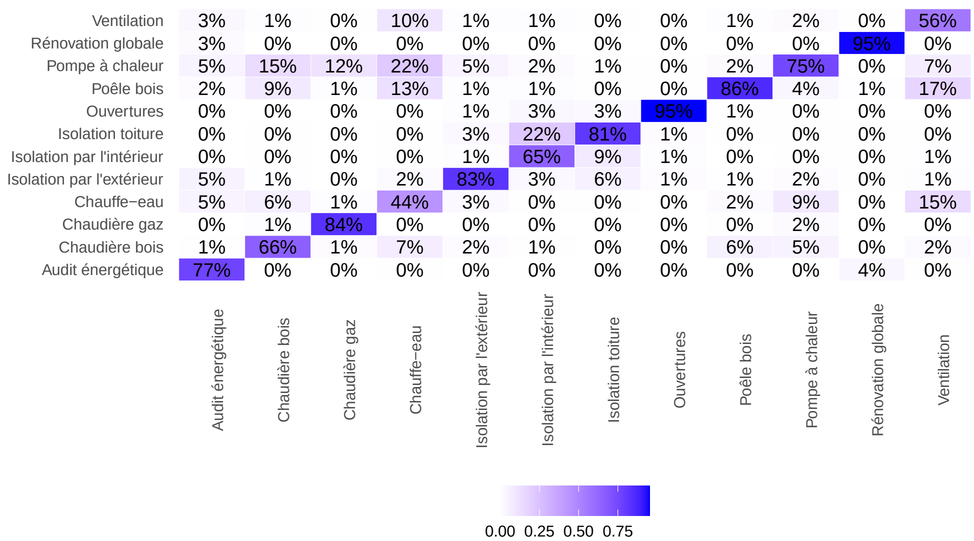

NOTE: This figure presents the distribution of MPR turnover across the 11 clusters of MPR-eligible actions. X-axis corresponds to the main cluster of a given RGE firm (the one in which MPR turnover is highest). On the y-axis, each square corresponds to the average share of total MPR activity in a given cluster. Shares therefore sum to 1 in each column. Sample is made of all RGE companies active in the MPR market in 2022. The color of the square is proportional to the share. SOURCE: MPR.

Figure A.13: Number of RGE firms over time (overall and active in the MPR market)

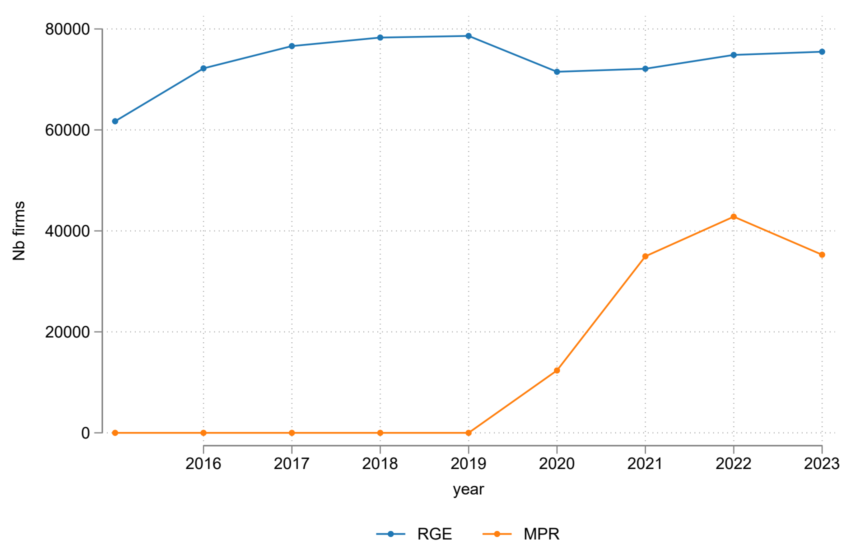

NOTE: This figure presents the number of firms with at least one RGE certification, by year between 2015 (RGE certificate creation) and 2023 (blue line). The number of said firms that performed at least one retrofit action subsidized by the MPR subsidy is indicated by the orange line. SOURCE: MPR, RGE.

Figure A.14: Distribution of distance between housing and selected firms, 2022

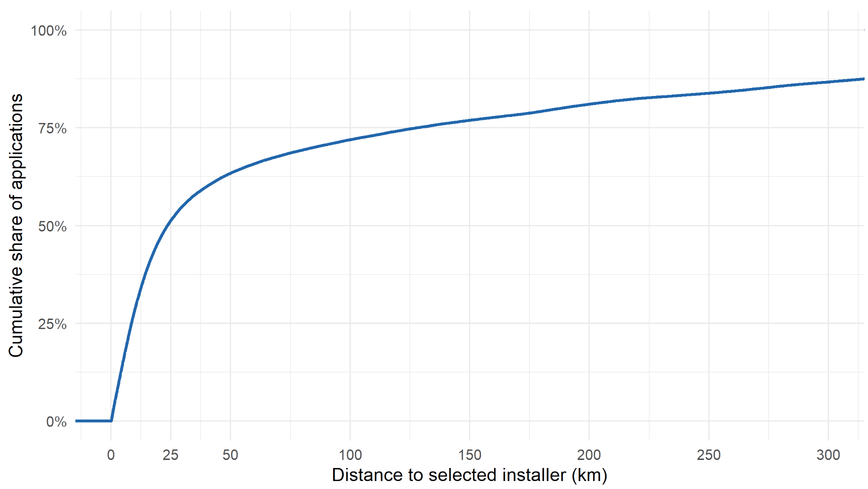

NOTE: This figure presents the cumulative share of MPR files involving a heat pump, as a function of the distance to the selected certified firm, in year 2022. The part of the distribution above 300 km is truncated. SOURCE: MPR, RGE.