The Distributional Impacts of Real-Time Pricing

Michael Cahana Natalia Fabra Mar Reguant Jingyuan Wang*

August 2023

Abstract

When designing electricity tariffs, efficiency considerations favor switching from time-invariant to real-time prices (RTP). Using hourly electricity consumption data, we analyze its distributional implications through a novel approach that infers individual household income from zip-code-level income and household-level electricity consumption. We find that RTP is slightly regressive, i.e., it increases the bills of low-income households by 0.5% on average, and 5% among those who lose from the switch, who are overrepresented in the lowest quintiles. This finding is explained by the correlation between income and household consumption patterns, electric appliances, and locations. Although the distributional impacts are economically small, counterfactuals show that they might worsen as high-income groups adopt demand-response technologies in the future. We propose a re-design of electricity tariffs that overcomes the regressive effects of RTP while preserving its efficiency properties.

Keywords: dynamic pricing, electricity, distributional effects, generalized method of moments, clustering.

JEL Classification: L94, H23, C55.

*Michael Cahana: Quora. Natalia Fabra: Department of Economics, Universidad Carlos III and CEPR, natalia.fabra@uc3m.es. Mar Reguant: Department of Economics, Northwestern University, CEPR and NBER, mar.reguant@northwestern.edu. Jingyuan Wang: Department of Economics, Northwestern University, jingyuan-wang@u.northwestern.edu. For their comments and feedback, we thank seminar participants at Yale, UPenn, TSE, UPF, HKS, Siences Po, the EARIE conference, the Bank of Spain, ECOBAS, ETH Zurich, Imperial College, Delaware, UMass Amherst, ITAM, the CBO, the ESEM-EEA conference, the TSE energy conference, and the Atlantic Workshop on Energy and Environmental Economics. Comments by Mateus Souza are also appreciated. Fabra and Reguant have generously benefited from the support of the BBVA Foundation. Fabra acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 772331-ELECTRIC CHALLENGES) and Comunidad de Madrid (MIMA-CM). Reguant acknowledges the support of NSF grant SES-1455084 and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101001732-ENECML).

1 Introduction

Economists are increasingly paying attention to the distributional impacts of energy and climate policies.[1] Beyond concerns over environmental justice, equity issues can become a bottleneck to passing efficient policies as they might undermine the political and social support required to complete the energy transition successfully. This paper studies the distributional implications of Real-Time Pricing (RTP) for electricity.[2] While economists have traditionally favoured RTP as an efficient policy tool (see, e.g., Borenstein (2005), and Borenstein and Holland (2005)), its broader implementation has been hindered by the widespread fear that it might create adverse distributional impacts across households (Joskow and Wolfram, 2012).[3]

The benefits of RTP are well known. Dynamic prices induce efficient consumption patterns as they more accurately reflect the changing marginal costs of generating electricity, which can fluctuate substantially during a day or a week and across the year. RTP thus creates incentives for energy conservation during high-priced hours, as well as for load shifting from high-priced to low-priced hours (see Jessoe and Rapson (2014), Burger et al. (2019), Faruqui et al. (2009), Wolak (2011), Allcott (2011), among others). This improves productivity and investment efficiency (Borenstein, 2005), mitigates market power (Poletti and Wright, 2020) and might have positive environmental impacts (Holland and Mansur, 2008). The energy transition will strengthen the benefits of dynamic pricing since the intermittency of renewable energy will enlarge the marginal cost and price savings that can be achieved through demand response. The same applies to the deployment of electric vehicles, batteries, or other forms of electricity storage for which dynamic pricing will provide more efficient signals. However, despite these benefits, the use of RTP has mostly been limited to industrial consumers (Blonz, 2022). Extending it to residential households has often been highly controversial, given the fear that low-income households could suffer from higher and more volatile bills, particularly so at times of extreme price events.[4] This is compounded by their limited flexibility to benefit from real-time price changes, at least in the short-run when their household equipment remains fixed.

We construct a framework to assess these equity concerns. In particular, we quantify the distributional impacts of dynamic pricing by analyzing the heterogeneous bill impacts from a switch from time-invariant prices to RTP across a large population of households. Our analysis uses hourly smart meter electricity consumption data of over one million households in Spain for eighteen months. Access to the Spanish data is precious because Spain is the only country where RTP

[1]See, e.g., Sallee (2019) and Shapiro and Walker (2021), and the papers included in Deryugina et al. (2019), among others.

[2]Pigouvian taxes, such as carbon pricing, provide another example of the same phenomenon, namely, the fact that efficiency-improving policies may fail to be adopted due to distributional concerns that affect the political economy of getting the policy passed (Douenne and Fabre, 2022).

[3]Levinson and Silva (2022) show how preferences over income redistribution affect electricity rate design.

[4]In Europe, these fears have been amplified during the 2021-2022 energy crisis following the conflict in Ukraine, as energy prices have reached record-high levels, increasing the share of energy-poor households (see, for instance, the European Commission's energy poverty observatory https://www.energypoverty.eu/). Similar equity concerns often arise in the US as well (Wang et al., 2021).

has been broadly rolled out as households' default option.[5] Methodologically, we propose a novel approach for constructing accurate measures of households' expected income, which is instrumental in uncovering the distributional effects of the policy.

Approach To analyze the distributional impacts of RTP, we combine estimates of (i) the households' bill impacts of the switch from time-invariant prices to RTP, and (ii) their income.

We observe bills under RTP, but not under time-invariant prices. To compute the former, we set the time-invariant prices at the average real-time price[6] and use the observed hourly quantities demanded under RTP. The implicit assumption, based on our previous work (Fabra et al., 2021), is that the short-run elasticity of electricity demand to changes in real-time prices is not significantly different from zero. This does not exclude the possibility that the medium-run elasticity is positive, as we consider in a counterfactual exercise.

To estimate income data at the household level, we first follow the standard approach of assuming that all households within a zip code have the same income distribution, based on the observed share of national income quintiles at the zip-code level. We find that the policy's impacts are not correlated with income, or just very modestly, with low-income households benefiting on average from the switch to RTP. However, these zip-code level regressions miss substantial within-zip-code heterogeneity, potentially biasing the analysis, as also shown by Borenstein (2012).[7]

To better capture within-zip-code income heterogeneity, we propose a novel method that combines the household electricity consumption data and the zip code income distributions for inferring the individual household unobserved income. As the first step, we use flexible classification algorithms to assign households to representative types according to their electricity consumption profiles. In addition to these flexible types, we also classify households depending on their amount of contracted power capacity, which is highly correlated with income. Once we have classified households into types, we estimate the income distribution of each household as a function of their type, their individual characteristics, and zip code characteristics. We do so by imposing that the inferred distribution of income based on our household types, aggregated at the zip-code level, matches the observed income distribution at the zip-code level using a generalized method of moments (GMM).

The critical assumption to identifying the impact of types on income is that the set of potential types is shared across zip codes (within a group of nearby zip codes).[8] This allows us to estimate the income distribution of types that rationalizes the observed zip-code income distributions. The household-level income distribution can then be obtained by combining the type income distribution

[5] In some countries, such as Norway or New Zealand, RTP is offered by competitive retailers, but as far as we know, it is not the default option. For instance, Borenstein (2013) states that "I'm aware of no place in the U.S. that time-sensitive rates are the default for residential customers."Also, according to the European Commission (2009), "The case of Spain with a regulated default dynamic price contract is unique."Pébereau and Remmy (2022) examine the barriers to adopting RTP in New Zeland.

[6] We compute monthly and annually flat prices to capture different sources of seasonal variation.

[7] We build on the ecological inference literature explored in Borenstein (2012) with non-parametric and semi-parametric methods that help to discern zip-code level aggregate effects from individual ones.

[8] Not all types need to be present across all zip codes within a group since the probability of a given type at a given zip code might equal zero.

with a household's type. While this identification argument is non-parametric, in practice we also implement a semi-parametric version that leverages the functional form assumptions on the impact of types, individual characteristics, and zip code demographics on the distribution of income.

Because the household classification algorithm is sensitive to choices made by the researcher, we perform three types of analyses confirming the robustness of our results to those choices: we perform Monte Carlo simulations, we cross-validate our predictions by leaving some zip-codes out of our estimation and then predicting them, and we compare our inferred consumption-income patterns to consumption expenditure survey data (CEX).[9] As an additional reality check, we also show that, as one would expect, contracted power capacity is strongly positively correlated with our household income estimates. In contrast, assigning each household to the observed zip code's income distribution would significantly mask such a positive correlation.[10]

Our proposed method should prove useful in many settings, beyond electricity markets, where researchers have access to detailed individual-level data on socio-economic variables but they lack precise information about household income.[11] For instance, in studying the distributional implications of carbon taxes (Chanut, 2021), scanner data could be used to classify consumers into types according to their consumption bundles and expenditure, and our proposed approach would then deliver their estimated incomes given those types. Our approach could also prove useful in other fields in economics, such as Public Economics (Chetty et al., 2020), Education (Bleemer and Mehta, 2022), Finance (Gross et al., 2021), or Labor Economics (Gustman and Steinmeier, 2000), to name just a few.[12]

Main Findings Accounting for income heterogeneity within zip codes is important to uncover the distributional impacts of real-time pricing. In fact, and in contrast with the predictions made with aggregate income distributions at the zip-code level, we find that real-time pricing is regressive

[9] Ideally, one would like to have income data at the individual level to compare the performance of our proposed method with only zip-code level distribution data versus the performance using the individual-level data. However, as in most cases in practice, individual income data is not available, which justifies our contribution.

[10] In the Appendix, we also consider an alternative parametric approach to describe the link between a household's income and its electricity consumption, in the spirit of Berry et al. (1995) and Berry et al. (2004). We show that this approach faces some limitations in our setting, given the difficulties in summarizing the heterogeneity in electricity consumption data without making it too computationally intensive. Our approach finds a data-driven compromise to handle these limitations.

[11] As in our case, some databases contain individual-level data (e.g., on consumption, health status, education, etc.) but, because of privacy issues, only contain the zip code where the household is located. This makes it impossible to match the household with income at lower levels of aggregation (e.g., at the census tract). Similarly, household characteristics tend to be well documented in Census data, while some countries only provide detailed income statistics by zip code.

[12] Chetty et al. (2020) study the heterogeneous impacts of COVID-19 by analyzing household spending. They proxy for cardholders' incomes using the median household income in the zip code in which they live. Bleemer and Mehta (2022) quantify the wage return to majoring in economics. They proxy family income by the mean adjusted gross income in the student's home zip code. Gross et al. (2021) measure how the generosity of the consumer bankruptcy system affects the cost of credit. To measure the income distribution of bankruptcy filers, they use the median income in the filer's zip code. Last, Gustman and Steinmeier (2000) study the joint retirement decisions of dual-career couples. They use survey data on wages to infer household-level income. However, for those years for which survey data is unavailable, they use tenure, experience and health to impute wages and thus income. In all these cases, our approach could help uncover income heterogeneity within zip codes and thus deliver a potentially richer distributional analysis.

compared to an annual flat price. In particular, the switch to RTP increases the bills of low-income households by 0.5% on average, and 5% among those who lose, which are over-represented in the lowest income quintiles.[13] Nevertheless, since the overall effects are economically small, the switch to RTP in the Spanish market did not lead to concerning levels of redistribution across income groups, at least during our sample period.

This finding is explained by differences in household consumption profiles over time. While high-income households consume disproportionately more at peak times within the day/month, low-income consumers tend to consume more during the winter months. Therefore, switching from annual flat prices to monthly flat prices is regressive as low-income households face higher prices during winter. In contrast, switching from monthly flat prices to hourly prices is progressive as high-income households face higher prices during the peak hours within the day/month. Overall, the former regressive effect dominates given that price differences are wider across months than across hours of the day or month in our sample.

We explore the main channels that explain these different consumption patterns: heating, ventilation, and air conditioning mode (HVAC) and household locations. Electric heating (EH) and air conditioning (AC), which account for almost 30% of an average household's annual consumption, vary widely across regions depending on their average weather conditions and the availability of gas infrastructure. Furthermore, EH and AC are negatively and positively correlated with income, respectively.[14] Since electricity prices in Spain are significantly higher during winter and lower during summer, the use of EH by low-income households and the use of AC by high-income households explain the adverse distributional implications of exposing households to the monthly price variation across the year.

We also conduct two counterfactual experiments to understand whether the distributional effects of RTP could change going forward. We consider two likely scenarios: an increase in the incidence of extreme price events across the year, as experienced during the 2021-22 European energy crisis, and an increase in the demand elasticity of high-income households as investments in batteries, solar panels or electric vehicles would allow them to better respond to price spikes and benefit from price volatility. Under both scenarios, the magnitude of the regressive impacts of RTP would be enlarged.

Last, we show that the equity-efficiency trade-off could be partially avoided through an adequate tariff design that mitigates the regressive effects of RTP while preserving its efficiency properties. In particular, we propose two-part tariffs with (i) volumetric charges that reflect real-time-prices, and (ii) fixed-fees that are a function of contracted power and/or electric heating status. The fixed-fees can be designed to partly overcome the regressive effects of RTP given that, as we have shown, contracted power and electric heating status are strongly correlated with income.

[13] These estimates are net of variable taxes (e.g., VAT). Introducing them in our analysis would enlarge the magnitude of the distributional implications but would not change the sign of the impact (i.e., whether a household loses or gains from RTP).

[14] In Spain, older buildings might not have a formal heating system, and heating is provided with highly inefficient electric heaters, driving the negative correlation with income.

The structure of our paper is as follows. We next discuss the related literature. Section 3 describes the background of the Spanish RTP system and provides an overview of the data. Section 4 describes the methodology used to infer individual household income. Section 5 provides the results of our analysis, and Section 6 explores the channels. Section 7 performs counterfactual analyses, and Section 8 proposes an efficient and equitable electricity tariff design. Last, Section 9 concludes.

2 Related Literature

There is increasing policy and academic interest regarding the distributional impacts of electricity tariff design. A hotly debated issue is whether the fixed costs of electricity supply should be recovered through fixed fees or volumetric charges. For instance, Burger et al. (2019) analyze the distributional impacts of moving towards two-part tariffs in which the fixed costs of electricity supply are recovered through the fixed fee instead of volumetric charges. They find that this would hurt low-income households more but argue that two-part tariffs can be designed to mitigate such adverse impacts while preserving most of the efficiency gains.[15]

Most of the studies analyzing the distributional impacts of dynamic pricing have focused on the effects of Critical Peak Pricing (CPP), probably the most commonly used form of dynamic pricing. CPP combines standard fixed rates (or TOU) during most of the year, with occasional price increases (e.g. 10-15 over a year) when the supply/demand margin is particularly tight. Borenstein (2012) shows that CPP would have a modest impact on most residential bills. In particular, low consumption households would see their bills decline, high consumption households would see them rise, and low-income households would see almost no bill changes. Instead, using results from pilot programs with voluntary participation, Faruqui et al. (2010) find that low-income households benefit from CPP because they tend to have flatter household consumption profiles and are more responsive to dynamic prices. In a recent paper, Schittekatte et al. (2022) compare the efficiency of TOU and CPP under simulated prices in a market with high renewables penetration.

Beyond differences in the efficiency impacts, the distributional effects of CPP and RTP can also be quite different. First, the distributional impact of a switch from time-invariant rates to CPP is limited to differences in consumption during the critical peaks but has no differential effects across households outside these events, even across households with very different consumption profiles. Furthermore, the distributional impacts of CPP also depend on the household's ability and incentives to adjust its consumption after a price increase. This result is less relevant in the case of RTP given that price changes are less salient, thus reducing the household's ability to avoid the potential adverse bill impacts of RTP. As households tend to be more aware of price changes under CPP, they are typically better equipped to mitigate such potential adverse effects by reducing their load at critical times.

Instead of analyzing CPP, Horowitz and Lave (2014) use hourly load data from Commonwealth Edison residential households to determine which households would save money when moved from

[15] Borenstein (2012), Borenstein (2013), and more recently Brolinson (2019), have also analyzed the distributive implications of increasing block pricing, which is often used to promote energy conservation.

a time-invariant rate price to RTP. Larger households save money under RTP, while smaller households, and disproportionately low-income households, lose money under RTP. On the contrary, Burger et al. (2019) find that transitioning towards more time-varying rates tends to make low-income households better off. More recently, Leslie et al. (2021) have analyzed the distributional implications of a move to RTP in Victoria (Australia). They match substation electricity consumption data with demographic data to identify the characteristics of households that would benefit from RTP. They find that RTP would primarily help households in areas with low house prices, high levels of renters and elderly residents.

Another set of papers simulate the distributional impacts of an RTP system with opt-in. Borenstein (2007) addresses this question in an analysis of industrial and commercial households in Southern California. His research shows that if households switched to RTP and exhibited price elasticities of -0.1, their surplus would increase. Yet, a substantial share (38-44%) would still be worse off. Only with much higher elasticities would such households be better off under RTP.

Our analysis thus contributes to the study of dynamic pricing by performing a detailed analysis of the distributional implications of RTP at the household level. It has the advantage of relying on actual data of a broad population of users who were defaulted into RTP. Our results align with those of previous studies. We highlight that potential harm to low-income households depends greatly on several channels: type of flat prices (monthly vs. annual), consumption profiles, HVAC status, and geographical locations. By identifying the channels that drive these results, our analysis can be informative about the potential effects of RTP in other jurisdictions, which can help mitigate the adverse distributional impacts of RTP before it is implemented.

Methodological contribution. From a methodological point of view, our procedure to refine individual household income using zip-code level income distribution data and *kmeans* clustering is novel. As further discussed below, the first step of our approach contributes to the literature that has developed methods of clustering observations. We combine this with the idea of backing out individual primitives (income) from outcome variables, as in demand estimation models (Berry et al. (2004), Fox et al. (2011), and Bajari et al. (2007)). The energy engineering literature (e.g., Haben et al. (2015), Al-Wakeel et al. (2017), Melzi et al. (2015), and Tureczek and Nielsen (2017)) has used machine learning models to classify electricity load curves but, as far as we are aware of, it has not used this approach to infer household income.

The literature on finite mixture models proposes a series of methods for inference and clustering. Most papers use a parametric approach with normal distribution assumptions (McLachlan et al. (2019)). Bonhomme et al. (2016) propose a nonparametric approach in which they use repeated data to obtain identification power. The literature on discrete heterogeneity using clustering methods models grouped-heterogeneity differently from the finite-mixture literature. Bonhomme and Manresa (2015) propose a grouped fixed-effect (GFE) estimator which uses clustering methods to capture time-varying grouped heterogeneity. They show the consistency properties of estimators using clustering methods like *kmeans* when the true heterogeneity is discrete.

Bonhomme et al. (2022) propose a two-step grouped fixed-effects estimator that uses *kmeans* clustering methods in the first step and then estimates a model with group fixed-effects in the second step. They show the advantage of GFE over the classic FE model and prove the asymptotic properties of the GFE estimator with panel data. They show that the asymptotic properties also hold with continuous heterogeneity. For this reason, using discrete heterogeneity is a dimension reduction device rather than a substantive assumption about population unobservables. While our first step is similar to Bonhomme et al. (2022), as we both use individual-level moments to identify groups, our second step differs from theirs. Our second step matches the aggregate moment of the primitives (the zip-code level income distribution) and identifies group-specific parameters (the group-specific income distribution). Theirs is a systematic approach to heterogeneity in a fixed-effects model regression, which maximizes the likelihood of observing individual-level outcome variables in the second step.

We also contribute to the literature on demand system estimation with partial microdata. While we have rich individual-level consumption data and repeated observations for every individual, we do not have income information at the individual level, only the zip-code level income distribution. This prevents us from using micro-moments, as in Berry et al. (2004).

A classic structural approach would be to estimate a parametric individual consumption function with income as a covariate. Market-level demographics could help identify the relationship between income and consumption, as first shown in Berry et al. (1995). However, there are two limitations. First, the inversion over individual hourly consumption data is computationally burdensome. We simplify the methodology using our two-step estimator, as the second step's computational burden is almost the same as that of a constrained OLS estimator. Similar to the the fixed-grid nonparametric approach (Fox et al. (2011) and Bajari et al. (2007)), our purpose of discretizing household types is to simplify the computation and convert a big structural model into an OLS-style constrained estimator.

Specifying a parametric model for electricity consumption can be complex as the parametric relationship between income and electricity consumption is quite heterogeneous. Such simplification would load many variations of the impacts of RTP on the error term.[16] Therefore, a more parsimonious parametric approach might not fully utilize the high-dimensional repeated observations for each individual. Our proposed two-step estimator allows for household-grouped heterogeneity, without restricting the parametric functional form of the group-outcome variable relationship (e.g., electricity consumption). We allow income covariates to flexibly affect household types with the type/group distribution differing across zip codes.

[16] Appendix G presents the evidence and provides details on the parametric estimation.

3 Background and Data

3.1 Dynamic Pricing in the Spanish Electricity Market

In 2015, the Spanish regulator made real-time pricing (RTP) the default option for all households that had previously not switched away from their default provider.[17] Instead of paying a traditionally flat retail price, most residential household contracts were changed by default to a retail tariff that varies hourly according to the changes in wholesale electricity prices. Households that had previously switched away from their default provider were given the choice to opt into RTP, while households switched into RTP were given a choice to opt-out to a competitive retail supplier, most of which offer time-invariant tariffs. Given the high inertia in retail choice (Fowlie et al., 2021; Hortaçsu et al., 2017), the fact that RTP was introduced as the default option (with the possibility to opt-in and out) implies that it affected a significant fraction of the residential sector.

The default Spanish electricity tariffs comprise two components: the price of electricity in the wholesale market, which varies hourly, and a regulated access charge that covers other system costs (such as transmission and distribution, among others). Since the wholesale electricity market operates at the national level, all Spanish households under RTP face the same hourly prices, which the System Operator publishes one day-ahead on its web page. Regarding the access charge, households are defaulted to a time-invariant rate. However, they can opt-out and choose a Time-of-Use tariff for the access charge, which implies lower charges at night and on weekends. In our sample, a modest share of households are subject to the night-time tariff ($\sim$13.68%).

To be subject to RTP, households must have a smart meter installed. By the end of 2015, almost 12M smart meters had been installed in Spain, of which around 10.19 million were successfully integrated into the electricity suppliers' information and telecommunication systems. By 2018, all residential households in Spain (28.02 million) had a smart meter installed.

Figure 1 shows how real-time prices moved over time during our sample period. The daily, monthly and seasonal variation in prices due to changes in wholesale demand and supply conditions is immediately transmitted to the retail prices. This is unlike other retail tariffs that also fluctuate over time, but for which the pass-through rate tends to be lower. One would expect households to be able to shift their consumption across the hours of the day, but not much beyond that. Within the day, the peak vs non-peak price differences are relatively modest in the context of dynamic electricity pricing. Peak prices exceed non-peak prices by approximately 30% (excluding the first part of the sample period, when price differences reached 80%). This within-day price variation is much smaller than the one analyzed in the experimental literature, in which peak prices are increased by 200-600% (Harding and Sexton, 2017). The variation in monthly prices is more pronounced, with prices during the winter that can be 60-70% higher on average.

[17] Also, the new pricing scheme only applied to households with peak demand below 10kW, which only excludes households with very high consumption that have to contract with a competitive retailer.

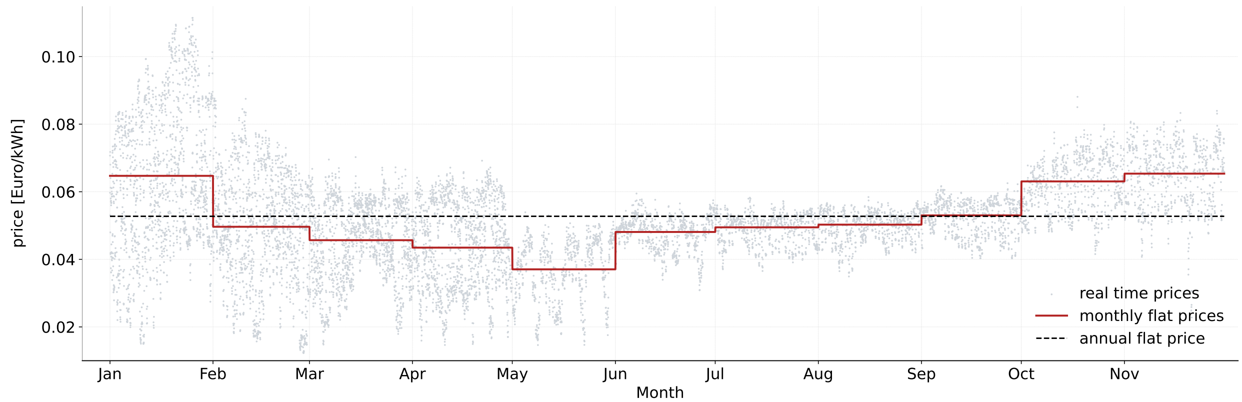

Figure 1: Price fluctuations over time (real-time, monthly and annual prices)

3.2 Hourly Electricity Consumption

Our dataset contains information for nearly four million Spanish households from January 1st, 2016, to May 31st, 2017. It was provided to us by Naturgy, one of the largest Spanish utility companies. Households in our sample mostly reside in Madrid, although they are scattered throughout Spain.[18] After treating for outliers with excessive zero consumption observations or missing zip code data,[19] as well as households outside of the regulated utility's territories,[20] we are left with 1,246,783 households covering 750 zip codes. We further drop December 2016 and May 2017 observations for data quality reasons, which leaves 15 months in our sample period (January 2016 to November 2016 and January 2017 to April 2017).[21] We thus have 17,371,296 household-month pairs in total.

The data include hourly consumption information (in kWh) for each household served by the utility, leading to more than 13 billion data points of hourly consumption data. The data also specify the type of tariff each household has, the hourly prices corresponding to the tariff identifier (in €/kWh), each household's contracted power, and its postal code information.

3.3 Annual Bills under RTP and Time-Invariant Prices

We compare household electricity bills under RTP versus time-invariant prices. To do so, we construct a revenue-neutral alternative to dynamic pricing, assuming zero price elasticity.[22]

[18] The geographic distribution of households is shown in the Appendix in Figure A.1.

[19] The algorithm for cleaning outliers drops a household from the sample if more than 25% of its consumption observations are zero, or if more than 5% are null.

[20] The default geographic provider is the one in charge of offering the default RTP tariff. Hence, households outside the utility's regional regulated territory can never be part of the RTP scheme.

[21] Our smart meter data lacks virtually all consumption data for December 2016, and is very incomplete for May 2017.

[22] Assuming that households do not change their consumption depending on their tariff might bias the distributional impacts if households have a significant demand response to short-run price changes and if such elasticities vary across income groups. We analyze this issue in Section 7.2.

For each individual household i and each month m, we compute the following expressions:[23]

$$Bill_i^{RTP} := \sum_{hdm} p_{hdm} \cdot kWh_{i,hdm}, \tag{1}$$

$$\overline{Bill}_i := \overline{p} \sum_{hdm} kWh_{i,hdm}, \tag{2}$$

$$\overline{p} := \frac{\sum_{hdm} p_{hdm} \cdot (\sum_i kWh_{i,hdm})}{\sum_{hdm}(\sum_i kWh_{i,hdm})}, \tag{3}$$

where p_{hdm} is the real-time electricity price in hour h of day d in month m and $kWh_{i,hdm}$ is the consumption of individual household i in that hour. Equations (1) and (2) give the bills under RTP and under an annual time-invariant price, respectively, which is defined in equation (3).[24] We will use the difference $Bill_i^{RTP} - \overline{Bill}_i$ to compute the bill change from being at RTP relative to time-invariant prices.

Last, we compute another measure for the time-invariant prices but at the monthly level, i.e., a constant revenue-neutral price for each month, and the resulting annual bill,

$$\overline{Bill}_i^m := \sum_{hmd} \overline{p}_m \cdot kWh_{i,hdm}, \tag{4}$$

$$\overline{p}_m := \frac{\sum_{hd} p_{hdm} \cdot (\sum_i kWh_{i,hdm})}{\sum_{hd}(\sum_i kWh_{i,hdm})}, \quad \forall m. \tag{5}$$

This allows us to decompose the bill changes due to the switch from an annual time-invariant rate to RTP as the sum of the "across months" and the "within month" effects. These are respectively captured by the first and second terms of the equation below:

$$\Delta Bill = \underbrace{[\overline{Bill}_i^m - \overline{Bill}_i]}_{across-month} + \underbrace{[Bill_i^{RTP} - \overline{Bill}_i^m]}_{within-month}. \tag{6}$$

The across-months effect reflects the bill impacts due to the monthly price variation across the year, while the within-month effect reflects the bill impacts due to the hourly price variation within the month.

Table 1 reports summary statistics for our main variables. The average annual consumption of a household is around 2,572 kWh, for which on average they pay 135 €/year.[25] There is heterogeneity in household energy bills due to differences between consumption levels and the timing of their consumption. Moving from time-invariant prices to RTP implies that some households lose (on average, losers face 4.54% higher bills) while other households gain (on average, winners enjoy

[23] In practice, electricity bills also include other cost components, such as network charges, which are independent of consumption and/or energy prices.

[24] Since the sample period includes the months of January, February, March, and April for 2006 and 2007, the observations for those months are each weighted by 0.5 to get a measure of the annual average bills.

[25] Recall that these amounts do not include other cost components, such as network costs or taxes. Depending on the household's contracted power and tariff choice, these additional costs can multiply the household's annual electricity bill by approximately 2.

Table 1: Summary statistics (household-annual level)

	Mean	Std	25%	50%	75%
$Consumption_i$ [kWh]	2616.40	1946.02	1366.84	2119.10	3244.14
$Bill_i^{RTP}$ [€]	137.80	101.59	71.99	111.98	171.37
$\overline{Bill}_i$ [€]	137.98	102.63	72.08	111.75	171.09
$\overline{Bill}_i^m$ [€]	137.85	102.05	71.90	111.67	171.12
$\Delta Bill$ [%]	0.02	5.18	-2.70	-0.71	2.18
$\Delta Bill$ (losers) [%]	4.54	4.28	1.13	3.32	6.76
$\Delta Bill$ (winners) [%]	-3.12	2.93	-4.24	-2.26	-1.12
Decomposition:					
$\Delta Bill$ within month [%]	0.07	1.93	-0.80	0.09	0.97
$\Delta Bill$ across months [%]	-0.05	4.81	-2.22	-0.94	1.60

Notes: This table reports household-level statistics. There are 1,246,783 observations. All units are measured in €, except for $Consumption_i$, which is measured in kWh. Annual bills (annual consumption [kWh]) are 11-month bills (consumption) from January to November because we do not observe December data. All percentages are computed with the bills under an annual time-invariant price, $\overline{Bill}_i$, in the denominator. By construction, the mean changes under RTP are 0 when expressed in Euros but differ from zero when expressed in %. The reason is that we first compute the bill change for each household, and we then take the average across households.

a 3.12% bill reduction). The bill impacts due to the across-months price variation are highly heterogeneous across households, ranging from a 1.60% bill increase for the 75% percentile to a 2.22% bill reduction for the 25% percentile. The within-month effect is smaller and more homogeneous across households.

3.4 Demographic Data

To examine whether demographics explain differences in the socio-economic impacts of RTP, we have also collected demographic data from the Spanish National Institute of Statistics (INE) and a private data provider, MB Research. The former provides demographics at the census district level (population, age, sex, education, dwelling types, and income distribution data).[26] In contrast, the latter provides income distribution data at the zip-code-level.[27]

As a first step towards understanding the distributional impacts of RTP, we regress the logs of average electricity consumption, average peak electricity consumption, and the bill impacts on income, using a cross-sectional sample at the zip-code level:

$$Y_j = \ln(Median\ Income)_j + HH\ size_j + \phi_j + \epsilon_j, \tag{7}$$

where j indexes the zip code and ϕ_j are province-age group-income group fixed effects. These regressions measure the correlation between median income per household at the zip-code level

[26] As we know the zip code of each household, but not its census, we match census districts and postal codes and then aggregate the census district data at the postal code level.

[27] Appendix A provides a more detailed description of these data sources.

Table 2: Zip code monthly-level regressions

	ln(kWh)	ln(kWh peak)	$\Delta Bill$ [%]
ln[IncPerHH]	0.059	0.085	0.438
	(0.057)	(0.067)	(0.443)
HHsize	0.319***	0.330***	-2.678*
	(0.045)	(0.040)	(0.831)
R-squared	0.552	0.675	0.317
N	538	538	538

Notes: All regressions include province-age group-income group fixed effects. $IncPerHH$ stands for median income per household, and $HHsize$ gives the mean number of people in the household.

and consumption, consumption at peak times (11 am-10 pm) and the bill change from the switch to RTP. Positive (negative) coefficients would reflect that households in higher (lower) income zip codes consume more in total, consume more at peak times and pay more under RTP.

Table 2 reports the results. Intuitively, column (1) suggests that household electricity consumption is positively correlated with income after controlling for household size. Column (2) suggests that peak electricity consumption positively correlates with income. Column (3) shows a positive correlation between income and bill changes, thus suggesting that households in lower-income zip codes are better off under RTP relative to the higher-income zip codes. In all cases, the relationship with income is noisy and statistically insignificant, partly due to the limited signal in aggregate zip code data.

4 Inferring the Household-Level Distribution of Income

The results from the reduced-form analysis face an important limitation as they overlook the existing income heterogeneity within zip codes. Therefore, the aggregate results are likely to underestimate the distributional impacts of the policy.

To get a more precise estimate of who loses and who wins from RTP, we develop a structural methodology to infer the individual household income distributions. Let us assume that household allocation of their hourly electricity consumption during the day (denoted kWh_{ih}, suppressing day index) is determined by a set of variables, such as temperature and seasonal components at the zip-code level (denoted x_{ih}) and their lifestyle (represented by their type θ_i), plus some random shocks ϵ_{ih},

$$kWh_{ih} = f(x_{ih}, \epsilon_{ih}|\theta_i). \tag{8}$$

Allowing the household's type θ_i to be correlated with its income helps us identify how income correlates with electricity consumption, and therefore study the distributional impacts of RTP.

The proposed methodology follows two steps. In the first step, we classify households into different types based on their contracted power capacity,[28] their electricity consumption patterns, and their HVAC ownership, which we infer from their hourly electricity consumption. Based on these results, we construct the aggregate probabilities of types at the zip-code level. In the second step, we assume that each type has a fixed distribution of income, which is unknown. We estimate the probability distribution by exploiting aggregate moments: the implied income distribution from the types within a zip code should match the observed zip code income distribution. These aggregate moments help us identify the probability that each household type belongs to a national quintile.

More formally, our objective is to uncover the income distribution of discrete household types, $\theta \in \Theta = \{\theta_1, \ldots, \theta_N\}$. To define the income distribution, we partition the income domain into K bins, $inc_k \in \{1, .., K\}$. We use national income quintiles of the household income distribution, so $K = 5$. Let $\eta_k^n = Pr(inc_k|\theta_n)$ denote this discrete probability of income conditional on household type θ_n. The goal is to estimate η_k^n for each income bin k and type θ_n, which we then apply to each household based on their types to infer their expected unobserved distribution of income.

The estimation assumes that the income distributions of the same type θ_n from different zip codes are the same, equal to $\eta^{\theta_n} = \{\eta_k^n\}_{k=1}^K$. This assumption would be too strong if we estimated the model combining all the zip codes in Spain. We instead assume that each type's income distribution is the same across zip codes within a province and estimate a set of $\{\eta^\theta\}_{\theta\in\Theta}$ for each province separately. We relax this assumption in our robustness section, in which we allow for additional types within a province for urban vs rural areas.

We next explain each step in more detail.

4.1 Step 1: Identifying Household Types

We define household types based on their contracted power, which we observe; their HVAC ownership status, which we infer from the correlation of their hourly consumption and temperature across seasons; and their hourly consumption patterns, which we construct from the smart meter data.

4.1.1 Classification by contracted power

Households face a fixed monthly charge for their contracted power, which is the maximum consumption allowed at any point in time. Since it is a function of the household's size and installed electrical equipment, it tends to be highly correlated with income. Contracted power can vary from 1 to 10kW (with 0.1 increments), but most households in our sample chose 2.5-5 kW. We classify households into two groups, depending on whether their contracted power is below or above 4 kW. 52% of the households in our sample belong to the low contracted power group (L), and the

[28] Contracted power capacity is the maximum consumption allowed at any point in time. Since households pay a fixed monthly fee as a function of their contracted power, they have incentives to contract it according to their actual electricity needs.

remaining 48% belong to the high contracted power group (H). Classifying households according to their contracted power is powerful because we observe it at the household level.

4.1.2 Classification by heating and air conditioning (HVAC) status

As detailed in Appendix B, we identify HVAC status (electric heating and/or air conditioning) by testing the seasonal correlation of hourly consumption and hourly temperature. Intuitively, we infer that a household has electric heating if it uses a relatively high amount of electricity during cold spells. Similarly, we infer that a household has air conditioning if it uses a disproportionately high amount of power during hot days. To calibrate the thresholds, we use a GMM estimator that matches the macro moments of the HVAC ownership rate at the regional level. This algorithm follows and complements the engineering literature that uses high-frequency data to identify HVAC status.[29]

Because the classification is based on individual patterns, the output of the procedure is a household-level indicator on whether the household used AC, electric heating, or both, creating a generated variable that allows us to classify households individually. Because our sample covers mostly the northern part of Spain where people rarely use AC, and given that we are limited in the number of types that we can allow, we focus on electric heating (EH) for the household classification in the estimation.

4.1.3 Classification by consumption patterns

We conduct the estimation separately for each province in our data (nine provinces in total). Within each province, we classify households based on their observable characteristics and consumption patterns.

We use a *kmeans* clustering algorithm to classify households based on moments of their hourly electricity consumption data. In total, 198 variables are generated to capture daily and seasonal consumption patterns for each household. We then apply a *kmeans* clustering algorithm to all households in the same province. Our 198 variables include:

- weekday average daily consumption and weekend average daily consumption in kWh;
- mean and standard deviation of hourly market share for each of the 24 hours by weekday and weekend;
- four variables capturing seasonal patterns in consumption: the ratio of winter consumption to annual consumption, the ratio of summer consumption to annual consumption, standard deviation of monthly consumption, and correlation of monthly consumption and the monthly flat price.

The first two sets of variables (194 variables in total) reflect household electricity consumption patterns within the day-month, while the remaining 4 variables reflect the seasonality across months.

[29]See Westermann et al. (2020) and Dyson et al. (2014).

Figure 2: An example of *kmeans* types in Madrid with high contracted power and no electric heating

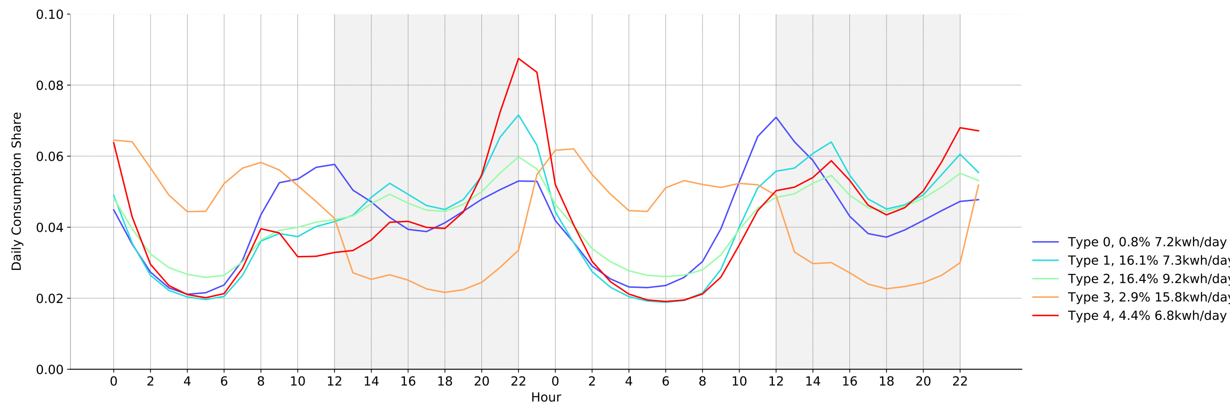

Notes: This figure provides an example of the *kmeans* classification of households in Madrid with high contracted power and no electric heating. The five clusters group households according to their electricity consumption profiles over the day. The first 24 hours are for weekdays and the last 24 hours are for weekends.

The former mainly depend on the household's lifestyle, while the latter are very much affected by the HVAC ownership.

We first classify households according to their individually observable contracted power and their individually inferred electric heating ownership: (L, EH), $(L, NoEH)$, (H, EH) and $(H, NoEH)$. We then use the *kmeans* clustering algorithm based on the above 198 variables to further classify households within each of these categories. To avoid small sample issues, we only allow for further heterogeneity in provinces and categories in which we have a sufficiently large number of households.[30]

Figure 2 illustrates an example of the *kmeans* classification. It shows the average daily consumption patterns for weekdays and weekends of households with high contracted power and no electric heating in Madrid. One can see that the algorithm picks up a variety of consumption patterns: households who consume mostly during weekends at brunch time and in the evening (cluster 0), households who consume at lunch time and in the evening (cluster 1 and 2), households who consume in the evening (type 3), and households who consume mostly at night (type 4).

4.2 Step 2: Identifying the Income Distribution of each Type

From step 1, we get distinct household types, θ_i^g, assigned to each household in a province g. The type space for each province g is $\Theta^g \equiv \{\theta_1^g, \theta_2^g, ...\theta_{N^g}^g\}$, where θ_n^g contains information on whether the household's contracted power is low (L) or high (H), on whether it owns electric heating (EH) or not, and on its *kmeans* type. In our main specification we set the number of types to be $N^g = 12$

[30] In practice, we set reduce the *kmeans* clustering types if a type contains less than 1,000 households. For example, in a province in which electric heating is rare, we reduce the number of types within that category.

for all provinces, with 3 *kmeans* types within each contracted power-EH category.[31] We estimate the types and income distribution for each province separately. From now onward, we suppress the superscript g for clarity.

We denote the share of type n households in zip code j as $P^j(\theta_n)$, and compute it as follows,

$$P^j(\theta_n) = \frac{1}{HH_j} \sum_i \mathbb{1}(\theta^i = \theta_n), \tag{9}$$

where HH_j is the total number of households in zip code j.

Non-parametric estimator Once we have a distribution of types at the zip-code level, we can uncover the unknown probabilities of types having a certain income by using across-zip-code restrictions in the share of types. For example, if the income at a certain zip code is relatively high, and if there are relatively many households in that zip code with high contracted power, the algorithm will conclude that the likelihood of high income for the high contracted power type is larger. Assuming that the underlying income distribution of a type θ_n is the same across zip codes within a province, we get the following moment conditions by matching the observed and predicted zip-code-level income distributions:

$$\min_{\eta} \quad \sum_j \omega_j \sum_{k=1}^{K} (Pr_k^j - \sum_{\theta_n \in \Theta} \eta_k^n P^j(\theta_n)), \tag{10}$$

$$s.t. \quad \sum_{k=1}^{K} \eta_k^n = 1 \qquad \forall \theta_n \in \Theta, \tag{11}$$

where ω_j is a weight representing the population of zip code j, Pr_k^j is the share of households in income quintile k in zip code j, and η_k^n is the probability that type θ_n belongs to quintile k.

The above objective function (10) uses a set of $(K-1)\times$ *Number of zip codes within the group* moments to identify the $(K-1) \times N$ unknown probabilities of income, η, where K is the number of income bins and N is the number of types. Thus, we need at least N zip codes to identify η. In practice, a larger number of zip codes can help reduce noise, which can otherwise lead to an inaccurate classification of consumer types and $P^j(\theta_n)$.

In our application, the number of zip codes that can be naturally grouped together is limited (e.g., a given geographical area), and thus we are constrained in the number of types that we can accommodate. In our main classification, we only have 12 types per province. Therefore, in our main implementation we rely on a semi-parametric estimator that can provide additional flexibility at the cost of some functional form assumptions.

[31] $N^g = 12$ is for provinces with sufficient population. As explained above, we make sure that the number of households within each type is greater than 1,000. When there are too few households of a given type, we merge it with other types.

Semi-parametric estimator For this reason, we also consider a semi-parametric extension that allows types to have some differences across zip codes based on individual characteristics and zip-code level demographics. More concretely, we specify the probability of a given household of type θ_i to have a probability distribution that depends on the individual (x_i) and zip-code level (z_j) demographics. This makes the method computationally more intensive, as we need to keep track of individual characteristics. However, it has the great advantage of allowing individuals classified into the same type to have distinct income distributions.

We use the following moment conditions by matching the observed and predicted zip-code-level income distributions:

$$\min_{\eta,\alpha,\beta} \quad \sum_j \omega_j \sum_{k=1}^{K} (Pr_k^j - \sum_{i \in \mathcal{I}_j} Pr_k(\theta_i, x_i, z_j)), \tag{12}$$

$$s.t. \quad Pr_k(\theta_i, x_i, z_j) = \frac{exp(\delta_{ijk})}{\sum_{k'=1}^{K} exp(\delta_{ijk'})}, \qquad \forall k \in [1, \cdots, K], \tag{13}$$

$$\delta_{ijk} = \alpha_k + \beta_0^{\theta_i} \times k + \beta_1 x_i \times k + \beta_2 z_j \times k, \tag{14}$$

where ω_j is a weight representing the population of zip code j, Pr_k^j is the share of households in income quintile k in zip code j, and $Pr_k(\theta_i, x_i, z_j))$ is the predicted probability of household i from zip code j belonging to income quintile k. θ_i is the household's kmeans type, x_i includes the household's contracted power capacity (continuous variable) and binned monthly electricity consumption (dummy variables), and z_j represents the zip code's demographic variables.[32]

Different than in our previous specification (10), the probability of income Pr_k is now a function of these variables. We use classic discrete choice logit formulas to parameterize the relationship between observed variables, household types, and income quintile probabilities, as shown in equations (13) and (14). α_k is a common income-bin dummy and ($\beta_0^{\theta_i}$) are type-specific coefficients and ($\beta_1^{\theta_i}, \beta_2^{\theta_i}$) are the same for all types. The first captures if a given type is correlated with higher income (($\beta_0^{\theta_i} > 0$). The second captures if higher contracted power and higher electricity consumption within a type is correlated with higher income ($\beta_1 > 0$). The last one is a vector for each type, and captures the correlation of zip-code characteristics and income. We would expect lower sociodemographics in the zip-code to be negatively correlated with the distribution of income, holding constant the type and other characteristics (e.g., for unemployment, we would expect $\beta_2 < 0$). Both $\beta_1^{\theta_i}$ and $\beta_2^{\theta_i}$ contribute to the distribution of income of a given type to be different across zip codes, as shown in Appendix F. Additionally, two households within the same zip code with the same type can have different income distributions due to the effect of β_1.

Using the estimated income distribution for each type, we calculate the implied income distribution for each household. As a sanity check, we show the aggregate distribution of income by contracted power in Figure 3. As expected, in Panel (a), the green line shows a positive correlation between income levels and contracted power, with higher (lower) income households being

[32] We test alternative parametric specifications as well as type classifications in the empirical implementation.

Figure 3: Contracted power and income quintiles

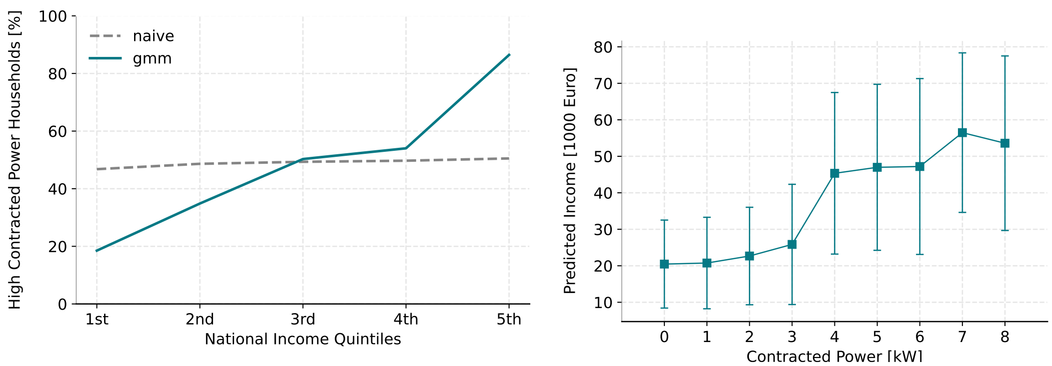

(a) Share of consumers with high contracted power by income

(b) Estimated average income as a function of contracted power

Notes: Panel (a) depicts the share of high contracted power households by estimated income quintiles. The green line depicts the share using our two-step method while the gray dashed line uses the zip-code-level income distribution alone. In our dataset, 52% of households are classified as having low contracted power, i.e., below 4 kW. Panel (b) depicts the estimated average income of households as a function of the contracted power. Both figures show that contracted power is strongly positively correlated with income, while this correlation would be lost under the naïve approach.

more likely to have high (low) contracted power. We predict that over 80% of households in the 5th quintile have high capacity vs. around 20% in the first quintile. In line with this, Panel (b) shows that the estimated income is strongly correlated with contracted power. The naïve estimator misses these patterns and only features a slight increase over the income gradient. This highlights the value of our approach.

4.3 Validation of inferred income

Our estimator provides a refined probabilistic assignment of income to households that is more granular than the zip code income distribution. In our methodology, we make several assumptions and choices to improve the probabilistic assignment of income to households.

Our method aims to better infer a household's expected income distribution. To understand the added value and performance of our estimator in small samples, we perform three checks to our method. First, in Appendix C, we perform a Monte Carlo simulation in which we assume that we know each individual household's income. We show that the estimator correctly recovers household income in expectation and examine what happens when some of our assumptions and choices differ from the true data-generating process. Second, in Appendix D, we perform an out-of-sample validation of the methodology by assessing how well our model predicts the distribution of income of out-of-sample zip codes. Third, also in Appendix D, we examine how predicted income-consumption patterns match annual records from the consumption expenditure survey microdata (CEX).

Monte Carlo validation For the Monte Carlo simulations, we consider the non-parametric estimator with a finite number of types in the data-generating process. We add noise to the distribution of each individual household, to reflect the fact that types might not exactly capture all income differences. In our Monte Carlos, we know the true type of households and their expected income (including the noise). We can compare the true income distribution to the inferred one. In the case of the naïve approach, the inferred income distribution is the same for all consumers in a zip code. In the case of our method, the imputation will be estimated by type. Figure C.1 shows that the naïve distribution of income tends to be much flatter (i.e., homogeneous) than the true distribution. Using our method, the inferred expected distribution lines up much better with the truth. This fit is naturally improved as we allow for more types and data.

Another way to show this result is by showing the inferred distribution of income for households belonging to a given quintile. In our Monte Carlo, we simulate a household's quintile. A household that belongs to the fifth quintile should have an underlying expected distribution with higher income. However, neither of these objects is known to the econometrician. We find that the naïve approach fails to estimate that households belonging to the high quintile have a higher income distribution. Instead, the probability of having a certain income level is very similar across households along all quintiles, as shown in Panel C.2a. As we allow for more types, the distribution of household income differs more along quintiles, as shown in Panel C.2c. The method cannot infer the realized income distribution of any given individual, which depends on additional realization draws, but it is classifying them probabilistically in a more accurate fashion.

Finally, we examine if our inferred income is still more correlated with true income than with the naïve approach in the presence of misspecification (Figure C.3). We find that misclassifying zip codes as sharing common types still leads to an improved correlation between imputed income and the true expected income. This is true as long as the zip codes have common shared patterns. For example, in our simulations, richer households tend to consume more at night in all zip codes. Therefore, misclassification across zip codes can still unveil useful patterns. In the empirical implementation, we avoid mixing consumers from very different areas (provinces) and test the robustness of our results with alternative groupings (rural vs. urban).

Overall, the Monte Carlo simulation helps highlight the value of our approach. With enough flexibility, we can unveil within-zip-code heterogeneity that would be muted using a naïve approach. As long as we allow for sufficient flexibility and have enough data, this classification appears to improve the inferred household income in expectation.

Out-of-sample validation We conduct cross-validation by including a subsample of zip codes in each province and predicting the out-of-sample income distribution for the other zip codes. Figure D.1 reports the distribution of in-sample and out-of-sample prediction errors. The two distributions are similar, which suggests that our approach captures the true relationship between income and consumer types. It can therefore handle out-of-sample predictions.

CEX validation For the CEX check, we create measures of annual electricity consumption by income quintile and region. In Figure D.2, we compare the estimated results to the CEX survey data. In all provinces, the naïve approach captures only across zip code income variation and cannot explain the relationship between income and consumption. On the contrary, the GMM approach performs clearly better even though the relationship between income and electricity consumption is flatter than in the CEX survey data.[33]

This is in line with the results from our Monte Carlo simulation. Our methodology cannot identify individual income perfectly, only the expected distribution. Therefore, it should not be seen as a perfect substitute for micro data, but as an improvement to otherwise very aggregate data. The implication of this attenuation bias when it comes to estimating the policy impacts are discussed further in the next section.

5 Quantifying the Distributional Impacts of RTP

We aim to identify the winners and losers from the move to RTP, using our estimated income distribution at the household level. We explore two dimensions of the distributional impacts: across and within income groups.

5.1 Impacts across Income Bins

We start by analyzing the heterogeneity in bill impacts across income groups. Figure 4 classifies households in five national income quintiles and plots the bill impacts following a switch from a time-invariant annual price to real-time prices. Results depend on whether one uses our estimated household-level income distribution (GMM approach) or the zip-code level income distribution (naïve approach). Under our proposed approach, the move toward real-time pricing is regressive, as it benefits high-income households while making low-income households worse off. Neglecting the within-zip-code income heterogeneity would deliver the opposite conclusion, with low-income groups paying slightly less under RTP. Furthermore, the naïve approach would also miss an essential part of the distributional implications, as the predicted bill impacts would be almost flat across income groups.

Figure 5 depicts the percentage of losers by income quintiles and average bill changes for losers and winners separately. According to our estimation, the lowest income bin has 7% more losers than the highest income bin, and an average loser's annual bill increases by 5%. The naïve approach underestimates (overestimates) the share of losers (winners) among the low (high) income quintiles. Furthermore, the naïve approach underestimates (overestimates) the loss of the losers among the low (high) income quintiles.

Figure 6 decomposes the bill impacts in changes within and across months. Panel (a), which relies on our estimated income distribution, uncovers fundamentally different distributional impli-

[33] An exception to the better fit is Madrid. Our utility data only covers select parts of the city and region, and thus the household subsample is not as comparable to the CEX region subsample.

Figure 4: Bill changes due to the switch to RTP [%]

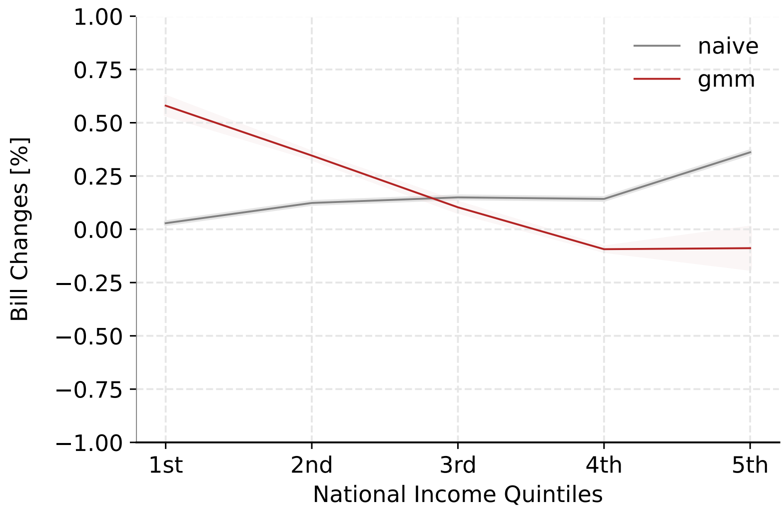

Notes: This figure represents the bill increase in % when moving from an annual time-invariant price to RTP. Results are reported for the five national income quintiles, with household income classified according to our estimated income (GMM) or to the zip code income (naïve).

Figure 5: Winners and Losers from the switch to RTP [%]

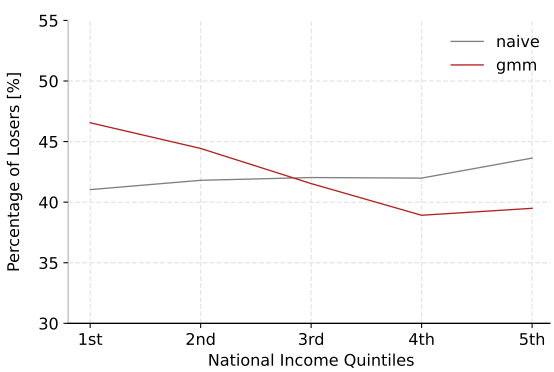

(a) Percentage of losers by income

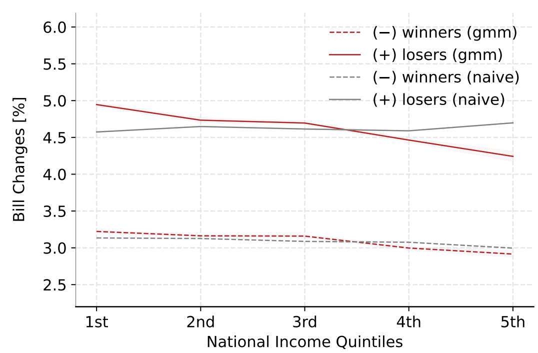

(b) Average bill changes for Winners and losers [%]

Notes: These figures report the bill impacts for winners and losers. Panel (a) reports the percentage of losers by income quintiles, while Panel (b) reports average bill changes for losers and winners separately. These figures reflect large within-income bin heterogeneity of the impact. See Figure 8 for a detailed discussion on within-income bin heterogeneity.

Figure 6: Decomposition of the distributional impact

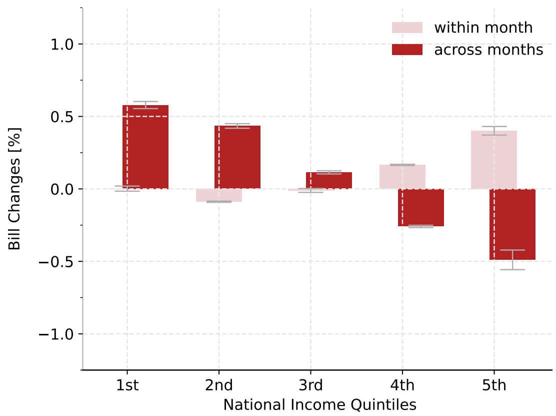

(a) Bill change [%]: GMM approach

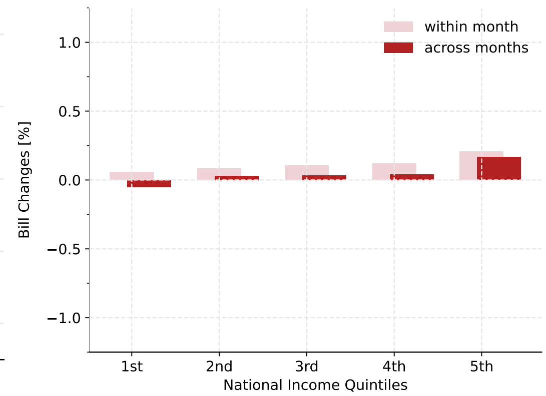

(b) Bill change [%]: naïve approach

Notes: These figures decompose the bill change in % when moving from an annual time-invariant price to monthly prices (pink bars) and from monthly prices to RTP (red bars), for the five national income quintiles. Panel (a) classifies households according to our estimated income (GMM approach), while Panel (b) relies on the zip-code level income (naïve approach). The bars would sum up to zero if expressed in Euro, but this is not the case when expressed in %. Also, note that these figures represent the national average, which hides the heterogeneity in the bill impacts across regions. See Figure 13 for the regional decomposition.

cations depending on the source of price variation. While a move from an annual price to monthly prices is regressive (across-months channel), the switch from monthly to hourly prices is progressive (within-month channel). The larger magnitude of the former explains why the move from an annual price to RTP is regressive overall. In the next section, we explore the channels that explain these patterns. As shown in Panel (b), using the zip-code level income distribution rather than our estimated household-level income distribution would hide these effects, and both bill impacts would appear to be slightly progressive and very small in magnitude.

5.1.1 Monte Carlo assessment of income impacts

Given that our method from inferring households' income might be subject to limitations, in which cases are our estimates more likely to better capture the full distributional impacts of the switch to RTP relative to the the naïve approach?

We perform a Monte-Carlo simulation to inform this discussion. Assume that the true data generation process behind the distributional impacts is governed by the following equation:

$$impact_{i,z} = t \times \theta_i + k \times inc_i + \sigma_z \times (\phi_z + \overline{\phi}_{zipgroup}) + \sigma_e \times \epsilon_{iz}. \tag{15}$$

The impact on household i in zip code z is $impact_{i,z}$, which is a function of household i's type θ_i, income inc_i, its zip code's fixed effect ϕ_z, and the zip code group's fixed effect $\overline{\phi}_{zipgroup}$. The error term ϵ_{iz} is orthogonal to all other variables, and it is normally distributed $\sim N(0, 1)$. The red coefficients capture the scale of each component: t represents how the individual type affects the final impact, k represents how individual income affects the impact, and σ_z and $sigma_e$ are the scales of regional fixed effects and unobservables, respectively. In our simulations, we explore the importance of each of these components in driving our results.

The true distributional impact is the summation of the direct impact of income (k) and the impact through the correlation between income and other variables in equation (15), including location impacts σ_z and individual impacts t. The zip-code level variation is observed, and types and income are inferred. Figure 7 reports the true distributional impact compared to the naïve and the GMM approach under different parameter assumptions. Panel (a) indicates that the GMM approach can capture the full distributional impact through individual types and the households' locations.

For our method to work perfectly, we need types, individual characteristics, and zip codes to be a sufficient statistic to determine the policy impacts. Thus, it is important to capture type heterogeneity based on policy impacts, thus the focus on consumption profiles and appliance ownership. We explore what happens when our approach does not work perfectly when $k \neq 0$. In this case, realizations of household income (not the expected distribution) directly enter the policy impacts.[34] As shown in panel (b), if the policy impact is correlated with the remaining unobserved income, we will only partially identify the effect.

In contrast, the naïve approach can only capture the impact through the locations, as shown in panel (c). When the true causal impact of individual types and income is 0, $t = 0, k = 0$, the GMM and naïve approaches give the same results, both of which are consistent with the true impact, as all variation is at the zip-code level. Otherwise, the naïve estimator features bias that can even go in the opposite direction of the true effects.

5.2 Impacts within Income Bins

We now turn to assessing the heterogeneous bill impacts within income groups, which are hidden by the relatively small average bill impacts across income groups (Figure 6). As highlighted by Cronin et al. (2019), energy use within income bins can still be highly heterogeneous. This can be due to an inability to better measure each household income or long-term wealth, but also due to genuine heterogeneity in consumer preferences and choices holding income constant. Panel (a) in Figure 8 plots the distribution of the percentage bill impacts from moving from an annual time-invariant price to RTP for the 1st, 3rd, and 5th quintiles. Whereas most consumers gain or lose at most 2%, the gains or losses can reach +/- 6% for some households. As can be seen, the right tail of the 1st quintile shows higher bill changes than those of the 5th quintile. However, there is a large share of within-quintile heterogeneity.

The large differences in impacts within each income quintile hide other sources of heterogeneity, such as location and its implications for heating and air conditioning (HVAC) use, an issue on

[34] Our method improves the assigned distribution of income to a given household but cannot predict the exact realized income of a household.

Figure 7: Assessing the method with a Monte-Carlo simulation

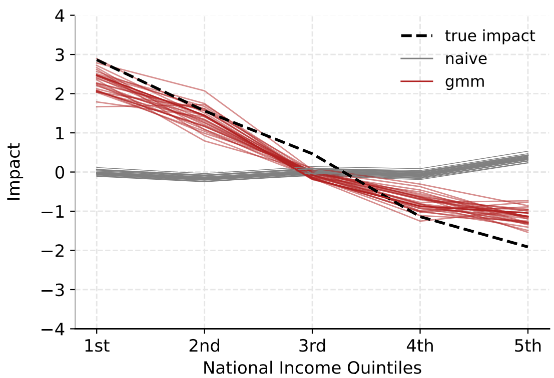

(a) Full bias correction

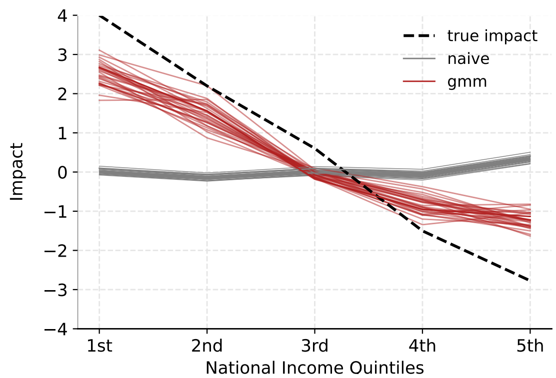

(b) Partial bias correction

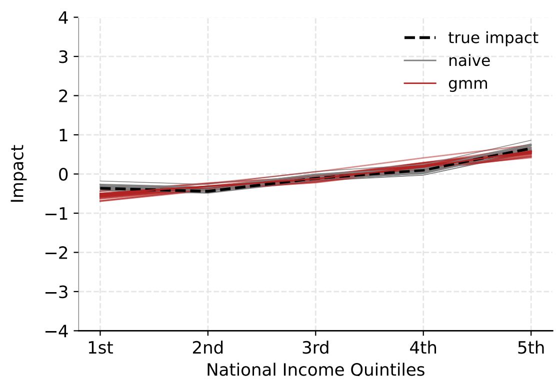

(c) No naïve bias

Notes: The figure show Monte Carlo results for the estimation of the policy impacts. The true policy impacts are depicted with the dashed line. Case (a) "Full bias correction" shows a simulation in which our type and zip code are sufficient to fully recover to effects, $t = 1, k = 0, \sigma_z = 1, \sigma_e = 1$. Case (b) "Partial bias correction" shows that the method recovers only part of the effect if the unobserved income realizations are correlated with the policy impacts, $t = 1, k = 1, \sigma_z = 1, \sigma_e = 1$. The method still provides a substantial improvement when compared to the naïve estimator. Case (c) shows that the naïve estimator provides the correct policy impacts only if the zip code variation is driving the effect, $t = 0, k = 0, \sigma_z = 1, \sigma_e = 1$.

Figure 8: Bill changes due to the switch to RTP [%]

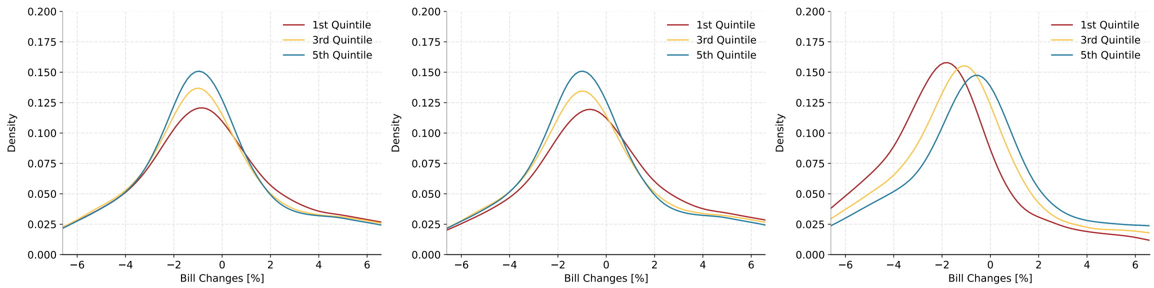

(a) All regions (b) Electric heating regions (c) Non-electric heating regions

Notes: These figures show the distribution of the bill changes due to the switch to RTP in the first, third and fifth income quintiles. Panel (a) shows the distributions at the national level, while Panels (b) and (c) distinguish between regions with a high and a low prevalence of electric heating, respectively. Together, they show that (i) there are large heterogeneities within income groups, and (ii) the low-income households are particularly hurt in the electric heating regions.

which we will elaborate further below. Panels (b) and (c) split the distributions between those regions where electric heating is prevalent ("electric heating regions"), from those where it is not ("non-electric heating regions"). The comparison of both plots shows that low-income households are relatively more negatively impacted in the electric heating regions, while the reverse applies to the non-electric heating regions. This finding suggests that the distributional impacts are not only driven by income differences but also by household locations and HVAC status. The following section is devoted to disentangling these channels.

5.3 Robustness

Our method is subject to several researcher choices that can impact the results. In Appendix E, we provide estimates of Figure 4 under alternative specifications. We consider the impact of the number of clustering types, the inclusion of observable or inferred characteristics (contracted power and electric heating), and the impact of urban vs rural heterogeneity, allowing for separate heterogeneity within each province. We find that our estimator is robust to these modifications, as explained in more detail in the appendix.

We find that allowing for few *kmeans* types (or even none) delivers consistent results with our main findings. Intuitively, heating and contracted power already explain individual income heterogeneity that the naïve estimator ignores. Similarly, we find that allowing for observable types can provide a useful categorical classification at the household level, but the *kmeans* alone already report the reversal in heterogeneity that we see in our main results. Intuitively, the *kmeans* also capture heating and utilization patterns. For some provinces with a substantial rural presence in the northwest, allowing for rural vs urban heterogeneity can change the results, but the main patterns do not vary.

Figure 9: Load curve by income quintiles

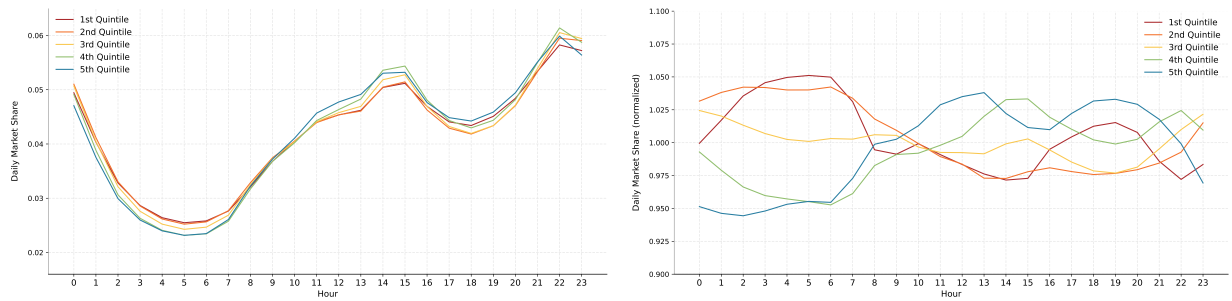

(a) Hourly consumption shares [%]

(b) Normalized consumption shares [%]

Notes: Panel (a) shows the average consumption patterns over the day for the five national income quintiles. Panel (b) depicts the normalized hourly consumption shares, defined as the share of the household's daily consumption at a given hour, over the average share in the sample. This shows that while consumption levels are not very different across income groups, their distribution across time is highly heterogeneous.

6 Channels

This section uncovers how income affects the bill impacts of RTP. We focus on differences across households regarding their consumption profiles, their HVAC status, and their locations.

6.1 Consumption Profiles

Our previous results show that moving from time-invariant monthly prices to RTP is progressive, i.e., low-income groups gain from this switch. This result is explained by household daily consumption patterns, as documented below. Figure 9a plots the average hourly consumption profiles of households in each of the five income groups. It shows that high-income households consume relatively more electricity during peak hours. While the differences across income groups seem small, there is significant heterogeneity in their consumption shares across the day. To uncover these, Figure 9b plots the daily consumption share for a given hour relative to the sample average by income group.[35] A number above one implies that the household concentrates a greater share of its consumption at that hour relative to the average share. This figure shows that the high-income group consumes more at peak times than the sample average, while the low-income group consumes relatively more at off-peak times. In other words, the consumption profiles of high (low) income households tend to be procyclical (countercyclical).

Beyond this graphical evidence, we can use simple regressions to understand how income affects electricity consumption patterns during the day and how that leads to within-month gains and losses from RTP.

We start by computing the price coefficient for each household ($\textit{price coeff}_i$) by regressing the household's hourly consumption on hourly prices, plus a constant. This coefficient captures whether

[35] For example, we compute the share of daily consumption at noon for a given income group, and we compare it to the average.

Table 3: Income, hourly consumption patterns, and within month bill changes

	Price coeff.	$\Delta Bill^m$ [%]	$\Delta Bill^m$ [%]	$\Delta Bill^m$ [%]
Price coeff.		5.479***		5.423***
		(0.039)		(0.039)
1st quintile	0.072*		0.384	-0.009
	(0.037)		(0.235)	(0.044)
2nd quintile	0.310***		1.745***	0.064
	(0.054)		(0.341)	(0.061)
3rd quintile				
4th quintile	0.439***		2.524***	0.145**
	(0.063)		(0.395)	(0.071)
5th quintile	0.381***		2.379***	0.312***
	(0.030)		(0.179)	(0.024)
R^2	0.292	0.769	-0.045	0.771
N	1139876	1139876	1139980	1139876
FE	zip code	zip code	zip code	zip code

Notes: This table reports regression results for equations (16), to (19) in columns (1) to (4), respectively. A zip code fixed effect is included in all regressions. The number of observations is slightly lower than in Table 1 because we have dropped some households for which a good identification cannot be obtained.

a household's consumption pattern is positively or negatively correlated with real-time prices. We then measure the correlation between the price coefficients and the income levels (regression (16))[36] as well as the extent to which the price coefficient explains the within-month bill effects (regression (17)):

$$price\ coeff_i = \sum_{k=2}^{5} \beta_k \mathbb{1}(Inc_k)_i + Z_i + \alpha_z + \epsilon_i, \tag{16}$$

$$\Delta Bill_i^m = \gamma\ price\ coeff_i + \alpha_z + \epsilon_i, \tag{17}$$

where α_z is a zip-code level fixed-effect, and Z_i are control variables, including household tariff choices and HVAC status.[37]

Table 3, columns (1) and (2), reports the estimated results. As can be seen, higher-income households tend to have higher price coefficients after controlling for zip code fixed-effects and a set of individual-level control variables. In turn, the price coefficient explains a large share of the variation in within-month bill changes, which is highly significant, as shown in column (2). These two pieces of evidence explain why moving from monthly prices to RTP tends to benefit the low-income groups.

Our two last regressions confirm that income affects bill changes mainly through the correlation

[36]Note that in equation (16), the first income bin is omitted. Hence, β_k reflects how much more correlated income and the price coefficients in group k are, relative to the lowest income group.

[37]Details about the HVAC status variables are explained in the next subsection.

Figure 10: Impact of price volatility on within month redistribution

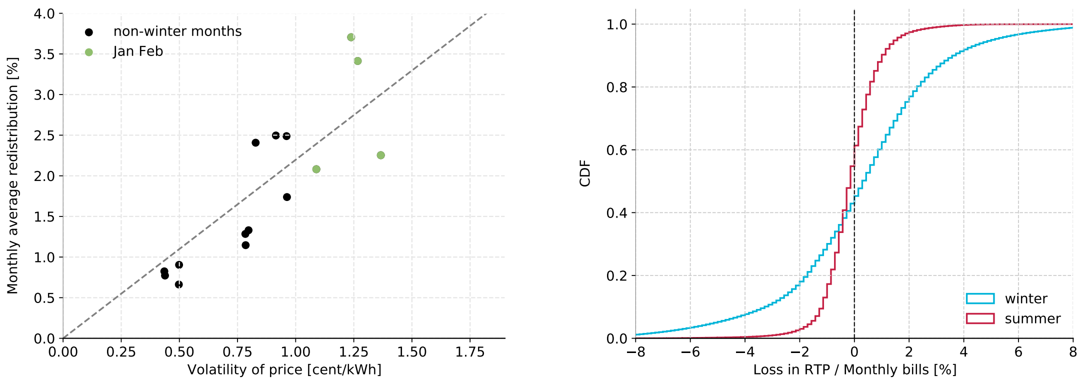

(a) monthly price volatility and RTP redistribuion

(b) winter and summer within month bill change in RTP

Notes: Panel (a) shows more redistribution during months of higher price volatility. The three dots with the highest price volatility correspond to January 2017, January 2016, and February 2017. The Y-axis of Panel (a) is the average absolute within-month bill changes at the household level, which is highly correlated with the monthly price volatility. The higher price volatility during the winter months explains why the CDF of the bill changes during winter is flatter than in the summer, as shown in Panel (b).

between consumption and real-time prices, as captured by the price coefficients:

$$\Delta Bill_i^m = \sum_{k=2}^{5} \beta_k \mathbb{1}(Inc_k)_i + \alpha_z + \epsilon_i, \tag{18}$$

$$\Delta Bill_i^m = \gamma \ price \ coeff_i + \sum_{k=2}^{5} \beta_k \mathbb{1}(Inc_k)_i + \alpha_z + \epsilon_i. \tag{19}$$

Column (3) of Table 3 shows that income is correlated with the within-month bill increase. However, if controlled for the price coefficient, the direct effect of income on the within-month impact becomes minimal, as shown in column (4). Furthermore, the estimates for the price coefficient are very similar when controlling or not controlling for income, i.e., the first line in columns (2) and (4). This finding highlights that the channel for the distributional impact runs through the correlation of household consumption patterns and real-time prices, which differs across income levels.

Price volatility amplifies this channel. Figure 10 shows the relationship between price volatility, defined as the standard deviation of hourly prices within a month, and the monthly redistribution effect, defined as the sum of bill changes (in absolute value) including all households. In months with more price volatility, bill changes can go up to 2.5-3.5%, but the changes remain low at many other times of the year. Since winter months depict higher price volatility, the distributional impact becomes greater, as shown in Panel (b) of Figure 10.

6.2 HVAC status

Domestic electric heating (EH) and air conditioning (AC) strongly impact electricity consumption, both regarding the levels and consumption patterns over time. Panels (a) and (b) in Figure 11 plot the average consumption patterns of households with and without electric heating during the day and across the year, respectively; Panels (c) and (d) do the same for AC. As can be seen, there are substantial differences in the consumption patterns of households with different HVAC status. Households with electric appliances consume significantly more across all hours of the day than those households without them. Also, their consumption tends to be peakier, particularly so in the case of heating.

Furthermore, there are strong seasonal effects. As expected, households with electric heating consume more during the winter months (October through April), while households with AC consume more during the summer months (June through September). In the case of heating, these effects are more pronounced for high-income households compared to low-income households. In contrast, in the case of AC, the effects are pretty similar across income groups.

In general, higher-income households are more likely to have AC, while lower-income households are more likely to have electric heating. This fact is explained by the high costs required to install other heating systems (e.g., gas or central heating) relative to electric heating, which commonly relies on low-cost plug-in radiators. Indeed, 23% of the 5th quintile and only 18% of the 1st have AC. In contrast, 32% of the 1st quintile and 11% of the 5th have electric heating.[38] Since prices tend to be higher during the winter when electric heating is used, a move from an annual price to RTP tends to hurt low-income households relatively more. The across-months effect is quantitatively strong and offsets the within-month effects we documented in the previous subsection.

The following regressions, which capture the impact of HVAC status on either electricity consumption or the bill changes due to switching to RTP, report similar evidence:

$$Y_i = \beta^{AC}\mathrm{AC}_i + \beta^{EH}\mathrm{EH}_i + Z_i + \alpha_z + \epsilon_i, \tag{20}$$

where Y_i is either kWh_i or $\Delta Bill_i$, α_z is a zip code fixed-effect, and Z_i includes household-level control variables. The coefficients β^{AC} and β^{EH} capture the effect of AC and electric heating ownership on either the household's electricity consumption or on the bill changes.

The estimates show that AC increases annual electricity consumption by 312 kWh, i.e., 15% of a median household's annual consumption. For electric heating, the increase in annual consumption is five times higher, i.e., 1017 kWh, representing 52% of a median household's annual consumption. Through their effects on household consumption patterns, electric heating increases bills under RTP by 3.0%, but AC leads to a 0.7% lower bill. These opposite signs have a simple explanation: electric heating (AC) increases consumption during winter (summer) months when prices are higher (lower).

[38] These results are reported in Figure B.3a in the Appendix. For AC, these differences are stronger conditional on location. In Spain, the lower-income regions tend to be warmer, implying that lower-income households have more AC. Indeed, there are only minor differences within regions.

Figure 11: Load curves by HVAC status and income

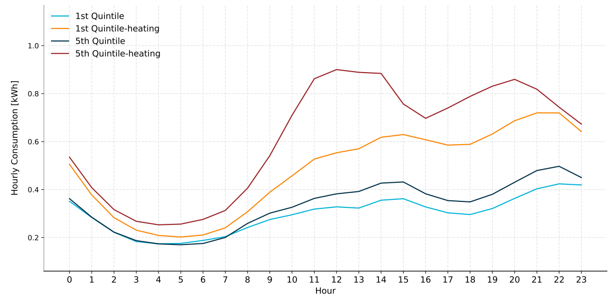

(a) Hourly consumption pattern [kWh] - Heating

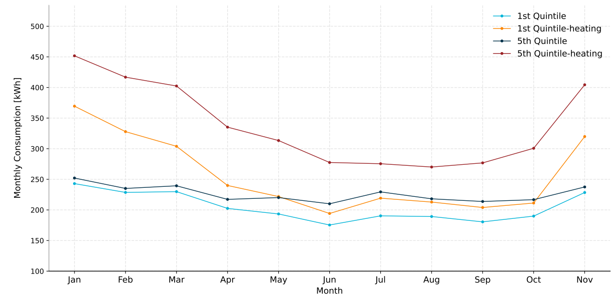

(b) Annual consumption pattern [kWh] - Heating

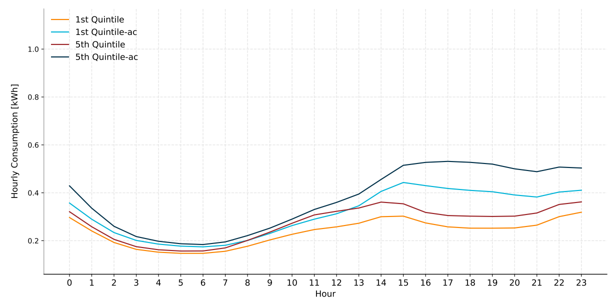

(c) Hourly consumption pattern [kWh] - AC

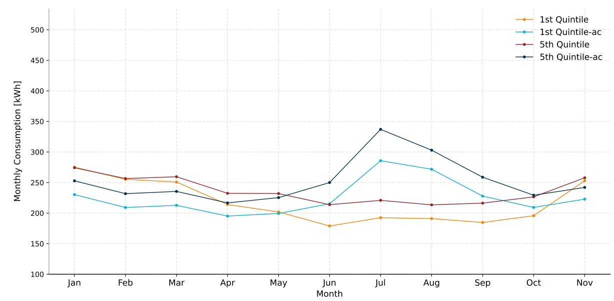

(d) Annual consumption pattern [kWh] - AC

Notes: These figures show consumption profiles over the day (the left panels) and the year (the right panels) for households with electric heating (the upper panels) and AC (the lower panels). Results are reported for low (1st quintile) and high-income households (5th quintile). The lines are mean hourly consumption for each group of consumers, truncating the top 1 percentile kWh observations.

Table 4: HVAC status and bill changes due to the switch to RTP

	kWh	$\Delta Bill$ [%]	$\Delta Bill^m$ [%]	$\Delta Bill^a$ [%]
AC	312.425***	-0.668***	0.217***	-0.885***
	(5.181)	(0.014)	(0.005)	(0.013)
Heating	1017.145***	3.024***	1.307***	1.717***
	(5.035)	(0.014)	(0.005)	(0.013)
R^2	0.172	0.163	0.108	0.151
N	1131430	1131430	1131430	1131430
FE	zip code	zip code	zip code	zip code

Notes: Column (1) reports the regression results from estimating equation (20) for consumption as the dependent variable, and columns (2)-(4) for the total bill change, the within month bill change, and the across months bill change, respectively. A zip code fixed effect is included in all regressions. One can see that households with AC (electric heating) tend to pay less (more) under RTP. This gain (loss) is mainly driven by the across months effect.

Figure 12: Bill changes [%] by electric HVAC status

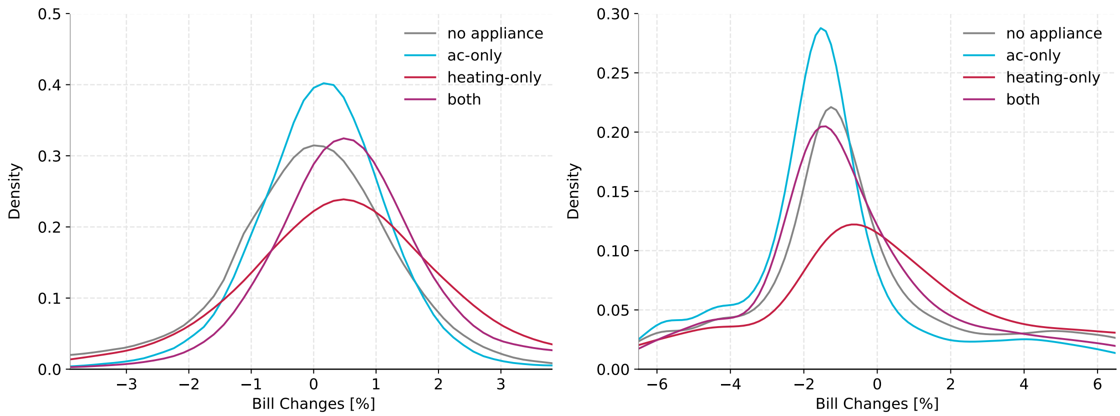

(a) Within month effects

(b) Across months effects

Notes: These figures plot the distribution of the bill changes due to the switch to RTP for households with no electric HVAC, with AC only, with electric heating only, or with both. The within-month and across-months effects are shown in Panels (a) and (b), respectively. The bigger bill increases are suffered by households with electric heating due to the across months effect.

These results are consistent with the evidence reported in Figure 12, which decomposes the bill impacts in the within and across months channels, distinguishing according to HVAC status. Regarding the within-month effects shown in Panel (a), both AC and electric heating owners would lose on average as they consume more. The winter price volatility is ten times larger than during the summer, amplifying the bill impacts of electric heating. Regarding the across months effects shown in Panel (b), AC users gain from being under RTP while electric heating users lose for reasons explained above.

6.3 Location

Another key driver of the distributional implications of RTP is location heterogeneity. Consumption patterns have much to do with local weather conditions, affecting HVAC status, even when controlling for income. Moreover, regional differences in the availability of heating infrastructure, mainly gas, affect the prevalence of electric heating in the region. For instance, whereas the availability of heating systems reaches 90.4% in Madrid, it is only 59.9% in the more rural Galicia. Castilla y Leon is the region where electric heating is least common (where only 8.6% of households have electric heating, as compared to the national average, 18.6%, because they rely more on gas and oil heating).[39]

Figure 13 decomposes the distributional effects of RTP in three dimensions: across regions (represented by the four lines), within months in Panel (a) and across months in Panel (b). As can be seen, the across-month price variation is the primary driver of the distributional implications

[39] See Table A.1 in the Appendix for details.

Table 5: Average bill increase by region

	$\Delta Bill$ [%]	$\Delta Bill^m$ [%]	$\Delta Bill^a$ [%]
Castilla y Leon	-1.269***	-0.214***	-1.055***
	(0.016)	(0.006)	(0.015)
Castilla-La Mancha	-0.700***	0.104***	-0.803***
	(0.011)	(0.004)	(0.010)
Galicia	0.630***	-0.009***	0.639***
	(0.007)	(0.003)	(0.007)
Madrid	0.044***	0.236***	-0.193***
	(0.009)	(0.003)	(0.008)
R^2	0.015	0.005	0.017
N	1139980	1139980	1139980

Notes: The reported coefficients result from regressing the bill changes on the regional dummies, without additional controls. The coefficients thus represent the mean bill increase at each region. Households in Castilla y León pay less (across months) under RTP because they are less likely to have electric heating than other regions.

of RTP, across income groups and regions. Furthermore, whereas these seasonal effects make RTP regressive in the electric heating regions (in the figure, Castilla la Mancha, Galicia, and Madrid), they make them progressive in the non-electric heating region (Castilla y Leon). The within-month channel is slightly progressive, but its magnitude is small. This evidence is consistent with the results from regressing the bill changes on regional dummies, with the coefficients capturing the mean bill increase under RTP for each region. Results are reported in Table 5. Overall, we conclude that HVAC status is a crucial driver of the distributional implications of RTP due to its effects on the levels and patterns of consumption across the year.

7 Counterfactual Experiments

We assess the counterfactual implications of two recent phenomena: (i) the increased incidence of price spikes and price volatility and (ii) changes in household equipment due to investments in demand response devices, batteries, or solar panels.

7.1 Commodity Risk and Energy Poverty

The reported distributional impacts are small since there was little price variation during our sample period. However, the impacts would be enlarged if prices within or across months became more volatile, as has been the case after our sample period. Indeed, in 2021 the average price in the Spanish electricity market tripled relative to the average price in previous years. Several reasons made this price shock particularly harmful for low-income households. First, there are more low-income households under the default real-time pricing policy relative to high-income households since they are entitled to a social tariff as long as they do not opt-out. Second, the price shock

Figure 13: Geographical heterogeneity and decomposition of the distributional impact

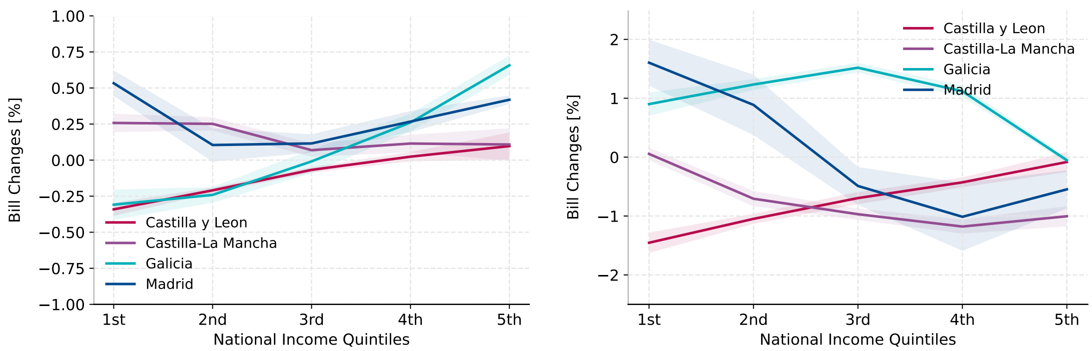

(a) Estimated bill change: within month (b) Estimated bill change: across months

Notes: These figures decompose the distributional effects of RTP in the within-month and the across-month effects in Panel (a) and (b), respectively, for four regions. The within-month channel is slightly progressive, but its magnitude is small relative to the across-months channel, which is regressive in all regions (except for Castilla y Leon, where there is little electric heating).

was particularly strong during the winter, which hit low-income households harder because they have relatively more electric heating. And third, price levels increased overall without affecting the within-day price variation much, which remained limited. The reason is that CCGTs set prices in the Spanish electricity wholesale market, which depend on the daily gas and CO2 prices. However, in the future, as renewable energies start setting the market price during some hours of the day, the within-day price volatility might become larger, potentially allowing the low-income households to benefit from their less peaky consumption patterns.

We consider three scenarios with low, medium and high-price trajectories in Panel (a) and with low, medium and high price volatility in Panel (b). To quantify the distributional implications of higher and more volatile prices, we simulate market prices from August 2021 to March 2022 using actual prices from August 2020 to March 2021, as shown in Figure 14. We simply add the same constant to all prices in the month for the mean price increases. For the volatility increases, we enlarge the departure of each price from the monthly mean. Specifically, we simulate prices according to the following equation:

$$\widehat{p}_{hdm}^{21} = (p_{hdm}^{20} - \overline{p}_{m}^{20}) \times \frac{\sigma_m^{21}}{\sigma_m^{20}} + p_m^{21},$$

where $\widehat{p}_{hdm}^{21}$ is the 2021 simulated price for hour h in day d and month m; p_{hdm}^{20} is the actual price in 2020 at that same date; and p_m^{20} and σ_m^{20} are the mean prices and standard deviation of prices in month m in 2020. Last, p_m^{21} and σ_m^{21} are the factors by which we scale prices and the standard deviation of prices, as plotted in Figure 14. We simulate the distributional impacts under nine scenarios with high, medium, and low monthly average prices and high, medium, and low monthly

Figure 14: Simulating a large price shock

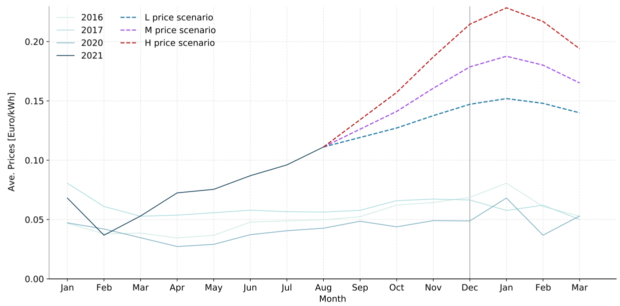

(a) Simulated prices (mean changes)

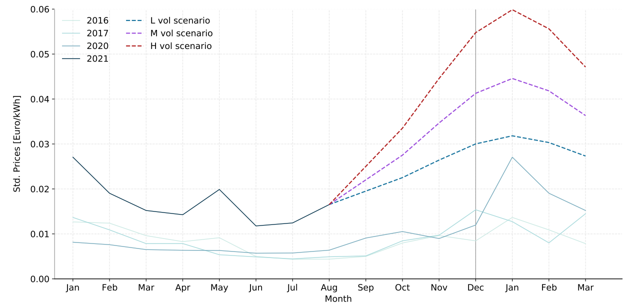

(b) Simulated prices (volatility changes)

Notes: Panel (a) plots the actual prices in the Spanish electricity market from 2016 to August 2021 and the simulated prices from August 2021 until March 2022 for the low, medium and high-price scenarios and the middle volatility. Panel (b) plots the actual price volatility (measured by the standard deviation) from 2016 to August 2021 and the simulated volatility from August 2021 until March 2022 with low, medium and high volatility assumptions for the middle price scenario. The actual monthly mean prices from September to December 2021 are close to our high-price scenario, and the actual volatility is even higher than our high volatility scenario.

price volatility.[40] The actual monthly mean prices from September to December 2021 are close to our high-price scenario, and the actual volatility is even higher than our high volatility scenario.

Keeping household demand patterns fixed, Figure 15 reports the distributional effects of switching to RTP using the actual and simulated prices for 2021-2022 under the nine scenarios described above. As seen in Panel 15a, the effects of RTP are regressive, and the magnitude of the distributional impact is greater than the one reported in the previous section. On average, a low-income household bill increases 8 percentage points more than a high-income household bill. The within-month effect is progressive, in line with our previous results, as shown in Panel 15b. However, it is so small that it does not matter for the overall effect, which is almost fully explained by the across-months effect, leading to bill increases of over 80% for the average households. These large increases make the overall impact of the price shock even more regressive. Not only are low-income households paying a higher increase in bills, but these large bill increases are also a substantial portion of their budgets.

In Section 6 we emphasized the importance of taking into account within-quintile heterogeneity and concluded that location and HVAC status are two main drivers behind the distributional impacts. We explore this further by analyzing the heterogeneous impacts due to the various price shocks. Figure 16 shows the average bill impact of the switch to RTP across regions for households with and without electric heating under the three price scenarios. As it can be seen, the difference between the light and dark red dots (representing the difference in the bill impact for households with and without electric heating under the high-price scenario) is wider than the difference between the light and dark blue dots (representing the difference in the bill impact for households with and

[40] Note that we do not have any December data. Therefore, we exclude December from our counterfactual simulation, which might lower our estimated impacts.

Figure 15: Distributional implications of RTP under a large price shock

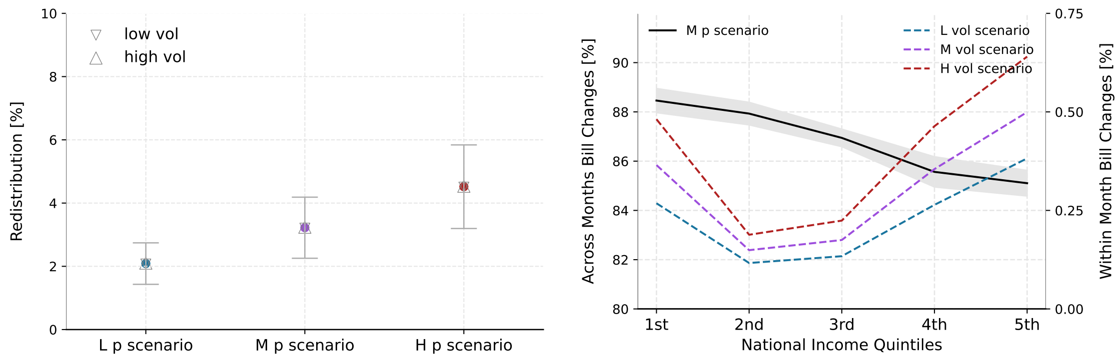

(a) Redistribution

(b) Within-month and across months effects

Notes: This figure illustrates the distributional implications of price increases and increased price volatility. Panel (a) shows that there is more redistribution with higher prices (as this enhances the across-months channel, which makes the low-income households relatively worse off) and lower volatility (as this mitigates the within-month channel, which would otherwise benefit the low-income households). Redistribution is defined as the additional bill increases (in percentage points) for the lowest quintile vs the highest quintile. The volatility effect is nevertheless much weaker than the price effect. Panel (b) shows the distributional impact due to the across-months channel (solid line; scale on the left axis) and due to the within-month channel (dashed lines; scale on right axis) for the middle price scenario.

without electric heating under the low price scenario). Also, for a given price shock, the difference between the most impacted region, which is also the lowest income region in our sample (Galicia), and the other regions is greater under the high-price scenario. It follows that large price shocks increase the heterogeneity in bill impacts across households through the two channels: location and HVAC status.

7.2 Accounting for Potential Responses: Demand Elasticity

So far, we have assumed that demand is inelastic, i.e., household electricity consumption remains unchanged after the move to RTP. However, it is reasonable to suspect that electricity demand will depict some price elasticity to short-run price changes in the future. For instance, this will be the case if households install devices that allow them to automatically adjust their consumption in response to price changes (Bollinger and Hartmann, 2020). Another source of demand response could well come through the deployment of electric vehicles and batteries, which typically allow households to benefit from arbitraging within-day price differences. The deployment of rooftop solar installations–which can be understood as a medium-run response to price increases–could also have important distributional implications to the extent that they allow households to reduce their consumption from the grid and thus have more stable energy costs.

These investments, together with using more electricity-intensive equipment (e.g., electric heat pumps), enhance the possibility and incentives for more active demand management. However, it is likely that these investments, and hence the scope for demand elasticity, will be positively

Figure 16: Heterogeneous impacts of the price shocks

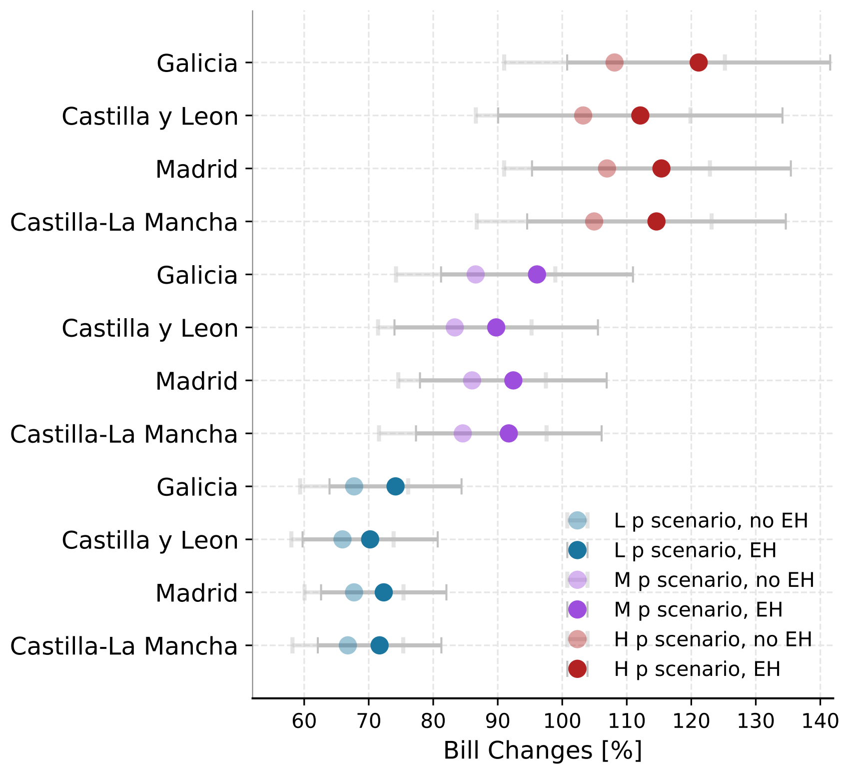

Notes: This figure shows the impact of price shocks on real-time pricing bill changes in % under low, medium, and high-price scenarios. Different colors represent different price scenarios, as shown in Figures 14 and 15. Price volatility is set at the medium level. The dots represent the average impacts of each HVAC-location group, and the grey lines represent the standard deviations. Regions are ordered from North to South in three blocks for the high, medium and low price scenarios. The darker dots represent households with electric heating, while the lighter dots are for the remaining households. As expected, the differences between the red dots become much larger than between the blue ones, i.e., stronger price shocks enlarge the differences in the bill impacts across the heating and non-heating households and locations.

correlated with household income.[41] Furthermore, this equipment mostly provides flexibility to respond to short-run price changes, not the seasonal price fluctuations that mainly affect low-income households. Consistent with this, Doremus et al. (2022) find that energy spending by US low-income households is half as responsive to extreme temperatures as compared to high-income households.

To explore these issues, we recompute household electricity bills under the assumption that they adjust their consumption to price changes using the following parametrization:

$$kWh^e_{i,hd} = kWh_{i,hd} \times \left[1 + \frac{p_{hd} - \overline{p}_d}{A + \overline{p}_d} \times \tau_i\right], \tag{21}$$
$$\tau_i = -\alpha \hat{inc}_i,$$

where $kWh^e_{i,hdm}$ and $kWh_{i,hdm}$ denote household i's adjusted and actual consumption, respectively; p_{hd} is the real-time price, $\overline{p}_d$ is the daily average price, A is the bill's fixed fee,[42] τ_i is a negative parameter indicating the household's elasticity, α is a scale factor to adjust the elasticity to a reasonable magnitude, and $\hat{inc}_i$ is the household's estimated income. Because total consumption during the day tends to be relatively inelastic, we only allow households to adjust the timing of their consumption within the day by moving elastic activities to low price periods. In other words, households reduce (increase) their consumption when the real-time price is higher (lower) than the daily average. The magnitude of the change depends on the value of τ_i, which is positively correlated with income (in absolute terms).

Figure 17 shows the distributional impacts of a switch to RTP with elastic consumers under the assumption that income is positively correlated with the demand elasticity. As expected, as compared to our baseline results (solid red line in the figure), demand elasticity reduces the bills of the high-income households under RTP relatively more, given that they can adjust their consumption to the price changes. In the figure, we have considered two assumptions regarding the price elasticity. The dashed line shows the results assuming an elasticity between 0.05-0.3. The dotted line shows the results assuming that the elasticity allows for maximum savings of 10% of the total bill.[43] Under the latter assumption, the within-month impact (shown in panel b of Figure 17) also becomes regressive as high-income households adopt smart devices that allow them to also respond to the within-day price changes.

[41] Using tax credit data, Borenstein and Davis (2016) find that higher-income Americans households invest relatively more than the lower income in installing solar panels, buying hybrid and electric vehicles, and other "clean energy" assets.

[42] As explained in the background section, the energy prices account for about 50% of the hourly prices, while the volumetric fixed fee accounts for the other half. This implies that hourly price fluctuations (the term before τ_i in equation (21)) generally do not exceed 50%. In our data, hourly price fluctuations usually vary from -5% to 5% around the daily average prices, and they can be as large as +/- 30%.

[43] This assumption is motivated by the fact that smart devices like thermostats allow for average savings of 10%-15%.

Figure 17: Distributional implications of RTP under demand elasticity

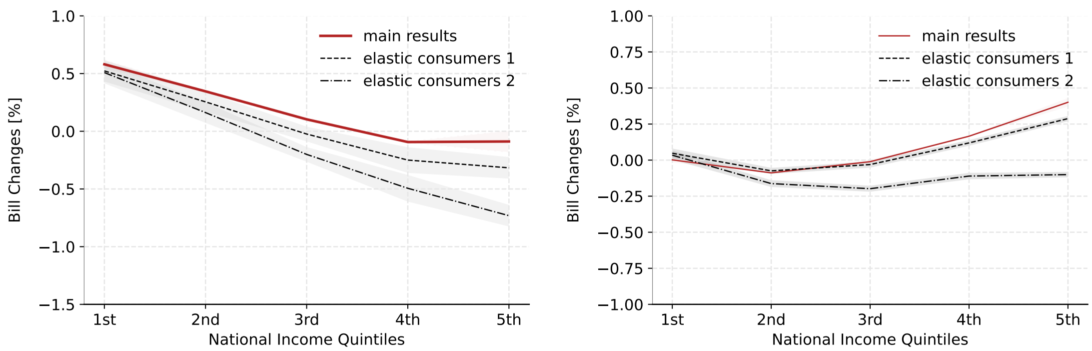

(a) Aggregate impact (b) Decomposition: within month impact

Notes: This figure illustrates the distributional implications of RTP when rich households have higher price elasticity. Panel (a) shows the aggregate bill change for each income bin, and Panel (b) decomposes it and shows the within-month impact. The regressive effect of RTP is now larger.

Table 6: Alternative tariffs

	(i) EH subsidy	(ii) Contracted power Tariff	(iii) Both
Heating subsidy [€]	-3.55	–	-2.16
Contracted power tariff [€/kW]	–	0.33	0.19
Fixed fee [€]	0.54	-1.39	-0.47

Notes: These tariffs are revenue neutral.

8 Designing Efficient and Equitable Electricity Tariffs

There are several ways to design pricing schemes to retain the efficiency properties of RTP while making it socially acceptable.[44] Understanding the channels by which RTP affects the various consumer groups is a necessary first step to do so. Our results highlight that the within-day/month price signal does not give rise to distributional concerns. This price signal is relevant for households to adjust their electricity consumption within a day/month, which is the most plausible source of demand elasticity. In contrast, the regressive effects stem from the across months price variation, as low-income households are faced with higher real-time winter prices when they consume the most.

Such regressive effects due to the across-months price variation could be avoided through an appropriate choice of the fixed fees, which should be more sensitive to income differences.

Figure 18 illustrates the distributional effects of moving from time-invariant prices to two part tariffs made of real-time prices plus fixed fees, where the latter are a function of observable characteristics that are positively correlated with income. An advantage of the Spanish context is that,

[44] For instance, as proposed by Borenstein (2007), the equity-efficiency trade-off could be partially avoided by setting some baseline fraction of the bill at a time-invariant price and letting the remaining fraction be a function with real-time prices. See also Borenstein et al. (2021) for an analysis of the efficiency-equity trade-off in electricity tariff design.

Figure 18: Distributional Implications of Alternative Tariff Designs

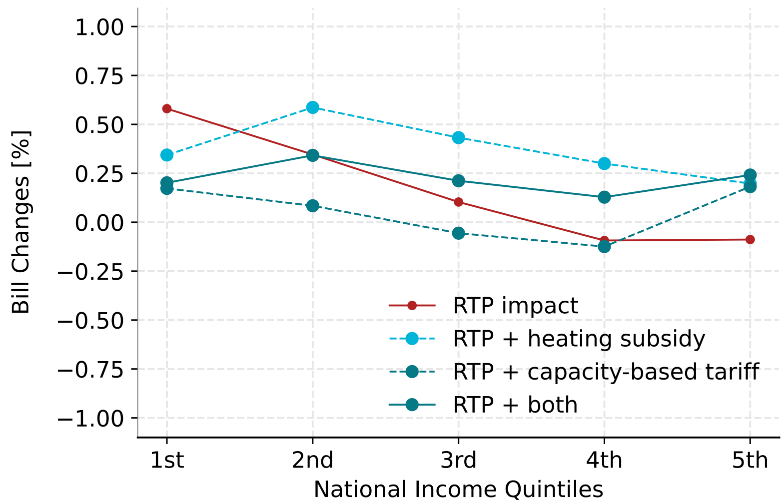

Figure 19: Impact under alternative tariffs [%]

Notes: This figure represents the bill increase in % when moving from an annual time-invariant price to RTP (red), or to RTP plus various revenue-neutral policies: a heating subsidy (blue), a fixed fee which is a function of contracted power (dotted green), or both (solid green). Results show that the regressive effects of RTP found in our baseline specification can be offset through an appropriate design of the fixed fees.

as we have shown, contracted power is strongly correlated with income (Figure 3). Hence, it would suffice to make the fixed fees a steeper function of contracted power.

For this reason, we first consider fixed fees that are a function of contracted power. We compute the optimal increase to contracted power fees that minimize disparities in average bill increases across quintiles. The revenues from such increase are spread equally across households to make this a revenue-neutral policy. Charging an additional 0.33€/kW for contracted power shifts the distributional impact of RTP from the red curve to the dotted green curve in panel (b) of Figure 18. As it can be seen, this tariff almost fully offsets the regressive effects of RTP.

Since electric heating is negatively correlated with income, we also consider adding an electric heating subsidy to households, which is paid by a uniform increase in fixed fees to all households, also making the policy revenue neutral. Again, we minimize differences across quintiles. Paying a 3.47€ subsidy to households with electric heating shifts the distributional impacts of RTP from the red to the dotted blue curve. As it can be seen, the subsidy lowers electricity bills for the lower income households.

Finally, combining the two gives an almost flat distributional impact. The green solid line in Figure 18 shows results from a 2.16€ electric heating subsidy and a 0.19€/kW contracted power tariff increase.

9 Conclusions

We have evaluated the distributional implications of the switch from time-invariant to real-time electricity prices in the Spanish electricity market, which became the first country to implement RTP as the default option for residential households broadly. While Fabra et al. (2021) show that this regulatory change had little impact on household electricity consumption, the question of whether households were asymmetrically impacted by it remains unanswered. This issue is fundamental as the fear of adverse distributional implications might have delayed a broader implementation of RTP elsewhere.

Access to hourly electricity consumption data at the household level for a large sample of representative Spanish households has allowed us to obtain meaningful conclusions about how their electricity bills have changed under RTP. Access to detailed socio-demographic data has further allowed us to understand the distributional implications of those changes.

An important step of our analysis is the estimation of household income. Working with our estimated income distribution, rather than with the zip-code level income distribution, has allowed us to uncover distributional effects of RTP that would have otherwise remained hidden. The electricity consumption data has also allowed us to infer the household ownership of electric heating or air conditioning, which are key determinants of the gains and losses from RTP.

The analysis reveals that, in the context of the Spanish electricity market, the move to RTP has been slightly regressive as lower income groups have been made worse off relative to the higher income groups. Interestingly, this overall effect can be decomposed into two channels: the bill impacts due to the within month and the across months price variation. We have found that the daily consumption patterns of low-income households tend to be relatively countercyclical, i.e., they consume relatively more when prices are lower, which implies that the move from time-invariant prices to RTP tends to benefit them. However, low-income households consume relatively more during winter when prices are higher because of their dependence on electric heating. The magnitude of this latter channel explains the overall regressive effect of RTP. However, the overall impact of RTP remained small and not of concern during our study period, thanks to the relatively stable prices and limited volatility. An increase in price levels and price volatility (as experienced in Europe during the 2021-2022 energy crisis) can further worsen the distributional implications, as we show in our counterfactual analysis.

These findings are not a criticism of real-time pricing as a useful policy tool. Instead, they convey its potential distributional effects in ways that should allow for the design of an equitable real-time pricing system. Indeed, we have also shown that the potentially adverse effects of RTP can be avoided with an adequate redesign of two-part tariffs, whose fixed fees should be more sensitive to income differences, as reflected in households' contracted power.

References

Al-Wakeel, A., J. Wu, and N. Jenkins (2017). K-means based load estimation of domestic smart meter measurements. *Applied energy 194*, 333–342.

Allcott, H. (2011). Rethinking real-time electricity pricing. *Resource and Energy Economics 33*(4), 820–842. Special section: Sustainable Resource Use and Economic Dynamics.

Bajari, P., J. T. Fox, and S. P. Ryan (2007). Linear Regression Estimation of Discrete Choice Models with Nonparametric Distributions of Random Coefficients. *American Economic Review 97*(2), 459–463.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile Prices in Market Equilibrium. *Econometrica 63*(4), 841.

Berry, S., J. Levinsohn, and A. Pakes (2004). Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market. *Journal of Political Economy 112*(1), 68–105.

Bleemer, Z. and A. Mehta (2022). Will studying economics make you rich? a regression discontinuity analysis of the returns to college major. *American Economic Journal: Applied Economics 14*(2), 1–22.

Blonz, J. A. (2022). Making the best of the second-best: Welfare consequences of time-varying electricity prices. Volume forthcoming.

Bollinger, B. K. and W. R. Hartmann (2020, January). Information vs. Automation and Implications for Dynamic Pricing. *Management Science 66*(1), 290–314.

Bonhomme, S., K. Jochmans, and J. Robin (2016). Non-parametric estimation of finite mixtures from repeated measurements. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 78*(1), 211–229.

Bonhomme, S., T. Lamadon, and E. Manresa (2022). Discretizing Unobserved Heterogeneity. *Econometrica 90(2)*, 625–643.

Bonhomme, S. and E. Manresa (2015). Grouped Patterns of Heterogeneity in Panel Data. *Econometrica 83*(3), 1147–1184.

Borenstein, S. (2005). The long-run efficiency of real-time electricity pricing. *The Energy Journal 26*(3).

Borenstein, S. (2007). Wealth transfers among large customers from implementing real-time retail electricity pricing. *The Energy Journal 28*(2).

Borenstein, S. (2012, August). The Redistributional Impact of Nonlinear Electricity Pricing. *American Economic Journal: Economic Policy 4*(3), 56–90.

Borenstein, S. (2013). Effective and equitable adoption of opt-in residential dynamic electricity pricing. *Review of Industrial Organization 42*(2), 127–160.

Borenstein, S. and L. Davis (2016). Will studying economics make you rich? a regression discontinuity analysis of the returns to college major. *Tax Policy and the Economy 30*(1).

Borenstein, S., M. Fowlie, and J. Sallee (2021, February). Designing electricity rates for an equitable energy transition. Working Paper WP 314.

Borenstein, S. and S. Holland (2005, Autumn). On the Efficiency of Competitive Electricity Markets with Time-Invariant Retail Prices. *RAND Journal of Economics 36*(3), 469–493.

Brolinson, B. (2019, September). Does Increasing Block Pricing Decrease Energy Use? Evidence from the Residential Electricity Market. Working paper, Georgetown University, Department of Economics.

Burger, S. P., C. R. Knittel, I. J. Pérez-Arriaga, I. Schneider, and F. v. Scheidt (2019, February). The efficiency and distributional effects of alternative residential electricity rate designs. Working Paper 25570, National Bureau of Economic Research.

Chanut, N. (2021, December). The Heterogenous Carbon Content of Food Consumption and Its Implication for Policy. Technical report, London School of Economics.

Chetty, R., J. N. Friedman, N. Hendren, M. Stepner, and T. O. I. Team (2020). The economic impacts of covid-19: Evidence from a new public database built using private sector data. NBER Working Papers 27431, National Bureau of Economic Research, Inc.

Cronin, J. A., D. Fullerton, and S. Sexton (2019). Vertical and horizontal redistributions from a carbon tax and rebate. *Journal of the Association of Environmental and Resource Economists 6*(S1), S169–S208.

Deryugina, T., D. Fullerton, and W. Pizer (2019). An introduction to energy policy trade-offs between economic efficiency and distributional equity. *Journal of the Association of Environmental and Resource Economists 6*(S1), S1 – S6.

Doremus, J. M., I. Jacqz, and S. Johnston (2022). Sweating the energy bill: Extreme weather, poor households, and the energy spending gap. *Journal of Environmental Economics and Management 112*(C).

Douenne, T. and A. Fabre (2022, February). Yellow vests, pessimistic beliefs, and carbon tax aversion. *American Economic Journal: Economic Policy 14*(1), 81–110.

Dyson, M. E., S. D. Borgeson, M. D. Tabone, and D. S. Callaway (2014). Using smart meter data to estimate demand response potential, with application to solar energy integration. *Energy Policy 73*, 607–619.

European Commission (2009, 11). Energy prices and costs in europe. https://ec.europa.eu/transparency/regdoc/rep/10102/2019/EN/SWD-2019-1-F1-EN-MAIN-PART-5.PDF.

Fabra, N., M. Reguant, D. Rapson, and J. Wang (2021). Estimating the Elasticity to Real Time Pricing: Evidence from the Spanish Electricity Market. *AEA Papers and Proceedings, American Economic Association 111*, 425–29.

Faruqui, A., R. Hledik, and J. Tsoukalis (2009). The power of dynamic pricing. *The Electricity Journal 22*(3), 42–56.

Faruqui, A., S. Sergici, and J. Palmer (2010). The impact of dynamic pricing on low income customers. *Institute for Electric Efficiency Whitepaper*.

Fowlie, M., C. Wolfram, C. A. Spurlock, A. Todd, P. Baylis, and P. Cappers (2021, June). Default Effects and Follow-On Behavior: Evidence from an Electricity Pricing Program. Technical Report forthcoming, Review of Economic Studies.

Fox, J. T., K. i. Kim, S. P. Ryan, and P. Bajari (2011). A simple estimator for the distribution of random coefficients. *Quantitative Economics 2*(3), 381–418.

Gross, T., R. Kluender, F. Liu, M. J. Notowidigdo, and J. Wang (2021, July). The Economic Consequences of Bankruptcy Reform. *American Economic Review 111*(7), 2309–2341.

Gustman, A. L. and T. L. Steinmeier (2000, July). Retirement in Dual-Career Families: A Structural Model. *Journal of Labor Economics 18*(3), 503–545.

Haben, S., C. Singleton, and P. Grindrod (2015). Analysis and clustering of residential customers energy behavioral demand using smart meter data. *IEEE transactions on smart grid 7*(1), 136–144.

Harding, M. and S. Sexton (2017, October). Household Response to Time-Varying Electricity Prices. *Annual Review of Economics 9*(1), 337–359.

Holland, S. P. and E. T. Mansur (2008). Is real-time pricing green? the environmental impacts of electricity demand variance. *The Review of Economics and Statistics 90*(3), 550–561.

Horowitz, S. and L. Lave (2014). Equity in residential electricity pricing. *The Energy Journal 35*(2).

Hortaçsu, A., S. A. Madanizadeh, and S. L. Puller (2017). Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market. *American Economic Journal: Economic Policy 9*(4), 192–226.

Jessoe, K. and D. Rapson (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review 104*(4), 1417–38.

Joskow, P. L. and C. D. Wolfram (2012, May). Dynamic Pricing of Electricity. *American Economic Review 102*(3), 381–385.

Leslie, G., A. Pourkhanali, and G. Roger (2021, January). Identifying consumption profiles and implicit cross-subsidies under fixed-rate electricity tariffs. Working paper, Monash University.

Levinson, A. and E. Silva (2022, October). The electric gini: Income redistribution through energy prices. *American Economic Journal: Economic Policy* (14(2)), 341–65.

McLachlan, G. J., S. X. Lee, and S. I. Rathnayake (2019). Finite Mixture Models. *Annual Review of Statistics and Its Application 6*(1), 355–378.

Melzi, F. N., M. H. Zayani, A. B. Hamida, A. Same, and L. Oukhellou (2015). Identifying daily electric consumption patterns from smart meter data by means of clustering algorithms. In *2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA)*, pp. 1136–1141. IEEE.

Poletti, S. and J. Wright (2020). Real-time pricing and imperfect competition in electricity markets. *The Journal of Industrial Economics 68*(1), 93–135.

Pébereau, C. and K. Remmy (2022, March). Barriers to Real-Time Electricity Pricing: Evidence From New Zealand. CRC TR 224 Discussion Paper Series crctr224_2022_339, University of Bonn and University of Mannheim, Germany.

Sallee, J. M. (2019, May). Pigou Creates Losers: On the Implausibility of Achieving Pareto Improvements from Efficiency-Enhancing Policies. NBER Working Papers 25831, National Bureau of Economic Research, Inc.

Schittekatte, T., D. S. Mallapragada, P. L. Joskow, and R. Schmalensee (2022, October). Electricity retail rate design in a decarbonized economy: An analysis of time-of-use and critical peak pricing. Working Paper 30560, National Bureau of Economic Research.

Shapiro, J. S. and R. Walker (2021, May). Where is pollution moving? environmental markets and environmental justice. *AEA Papers and Proceedings 111*, 410–14.

Tureczek, A. M. and P. S. Nielsen (2017). Structured literature review of electricity consumption classification using smart meter data. *Energies 10*(5), 584.

Wang, Q., M.-P. Kwan, J. Fan, and J. Lin (2021). Racial disparities in energy poverty in the united states. *Renewable and Sustainable Energy Reviews 137*, 110620.

Westermann, P., C. Deb, A. Schlueter, and R. Evins (2020). Unsupervised learning of energy signatures to identify the heating system and building type using smart meter data. *Applied Energy 264*, 114715.

Wolak, F. A. (2011). Do residential customers respond to hourly prices? evidence from a dynamic pricing experiment. *American Economic Review 101*(3), 83–87.

Online Appendix

"The Distributional Impacts of Real-Time Pricing"

A Data sources

A.1 Income data

In this appendix, we provide further details about the demographic data that we use in our analysis. These data are provided by the Spanish National Institute of Statistics, Instituto Nacional de Estadística (INE), and correspond to the most recent census (2011). The data contain information at the census level on population, age, sex, education, dwelling types (main dwelling, secondary dwelling, or empty dwelling), number of rooms per dwelling, and net surface area of dwellings for each census district in Spain. We have also collected detailed information on the distribution of income at the district (and sometimes section) level.[45] We only include places from which we have electricity consumption data. This limits our analysis to four regions: Galicia, Gastilla y Leon, Madrid, and Gastilla-La Mancha. Figure A.1 plots the location of these provinces.

We complement the data from the INE with data from MB Research at the postal code level. INE data reports median and mean income per household for each census. MB Research reports the distribution of household income, where the cutoffs are representative of the quintiles in the national distribution of income. Therefore, these two distributions of income complement each other at different parts of the support.

We know the zip code of each household, but not its census. To create a crosswalk between postal codes and census districts, we use shapefiles of Spanish postal codes and census districts provided by the INE. Census districts are matched to postal codes with which they have significant intersection.[46] On average, postal codes are matched to around seven census districts. Once census districts and postal codes are matched, census district data are aggregated at the postal code level. We find that some zip codes are not present in the shapefiles. To complement the map between zip codes and districts, we use data with latitude and longitude for the universe of street addresses in the postal code system ("callejero").[47] A district section and a zip code are matched if the latitude and longitude of the address lays inside that section.

A.2 Smart meter data

As explained in the main text, we partner with a large distribution utility in Spain to obtain de-identified smart meter data at the household level. Our dataset contains information for close to

[45] For confidentiality reasons, sections are often not reported as they are a fairly small geographical units. For small to medium sized municipalities, data are often only available at the municipality level, which often coincides with the postal code. Very small municipalities might not have their data reported.

[46] The matching algorithm is as follows: if 90% or more of a census district's area is contained within a postal code, or if 90% or more of a postal code's area is contained within a census district, then the census district is matched to the postal code.

[47] This information can be obtained at https://www.ine.es/prodyser/callejero/.

four million Spanish households from January 1st, 2016 to May 31st, 2017. It was provided to us by Naturgy, which is one of the largest Spanish utility companies. Its households tend to reside most densely in Madrid, although they are also scattered throughout Spain.[48] After treating for outliers with overly zero consumption observations or missing zip code data, [49] as well as households outside of the regulated territories of the utility. [50] The final sample contains 1,246,783 households, covering 750 zip code regions. We further drop December 2016 and May 2017 observations for data quality reasons, which leaves 15 months in our sample period (January 2016 to November 2016, and January 2017 to April 2017). We thus have 17,371,296 household-month pairs in total.The data include hourly consumption information (in kWh) for each household served by the utility, leading to more than 13 billion data-points of hourly consumption data.

A.3 HVAC statistics by province

We obtain province-level statistics about mode of heating and air conditioning to discipline our algorithm to infer appliance ownership. These moments are obtained from the Spanish National Statistics Institute (INE) and are displayed in Table A.1.

[48]The geographic distribution of households is shown in the Appendix in Figure A.1.

[49]The algorithm for cleaning outliers drops a household from the sample if more than 25% of its consumption observations are zero, or if more than 5% are null.

[50]The default geographic provider is the one in charge of offering the default RTP tariff. Hence, households outside of the utility's regional regulated territory can never be part of the RTP scheme.

Figure A.1: Geographic distribution of households

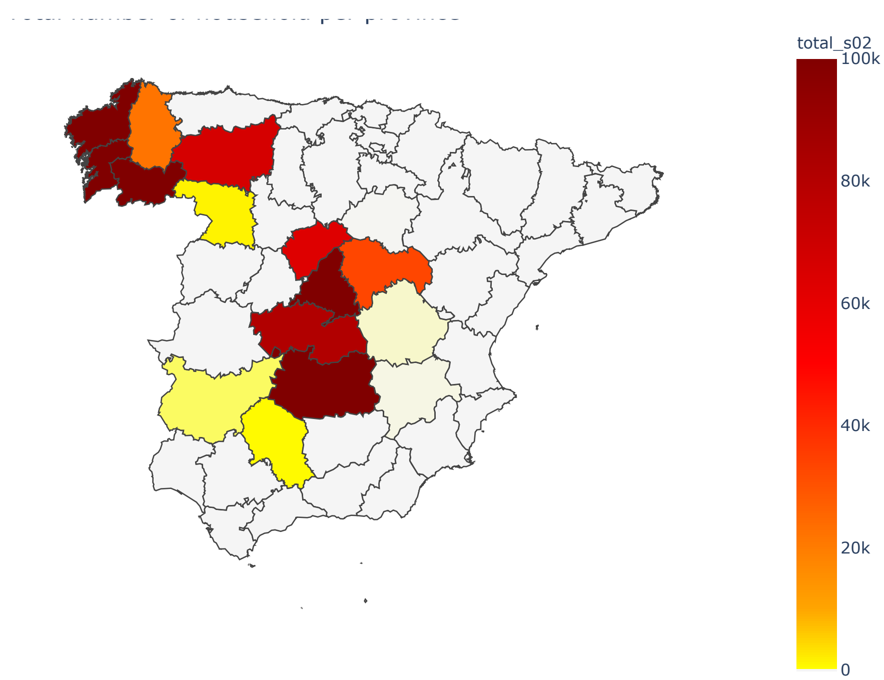

Table A.1: Statistics on the availability of heating systems

State	Heating availability	Electric heating			
		Total	(1)	(2)	(3)
Castilla y León	90.8	8.6	2.0	6.9	0.4
Castilla -La Mancha	86.2	15.3	1.9	13.5	–
Galicia	59.9	14.8	4.1	10.9	0.4
Madrid	90.4	15.6	8.3	8.3	0.5

Notes: (1) Individual electric boiler (2) Electric radiators and accumulators (3) Radiant wire. Source: Spanish National Statistics Institute (INE), Household and Environment Survey 2008 (https://www.ine.es/dynt3/inebase/index.htm?type=pcaxis&path=/t25/p500/2008/p01/&file=pcaxis&L=0).

B Inferring HVAC status

In this appendix, we infer household HVAC status by exploiting the richness of the smart meter data. The idea of using high-frequency data to infer HVAC status has been applied to engineering papers like Westermann et al. (2020) and Dyson et al. (2014). For each household, we first run the following regression to obtain the correlation between its electricity consumption with temperature in winter and summer:

$$\begin{aligned} kWh_{i,hdm} =& \beta^i temp_{hdm} + \beta^i_s temp_{hdm} \times \mathbb{1}(summer \times daytime) + \beta^i_w temp_{hdm} \times \mathbb{1}(winter \times daytime) \\ &+ \alpha^i_{fe} \mathbb{1}(hour \times month \times weekends) + \epsilon_{i,hdm} \end{aligned} \tag{B.1}$$

where $kWh_{i,hdm}$ is hourly consumption of household i in hour h on day d in month m, and $temp_{hdm}$ is the corresponding temperature at that time; α^i_{fe} are hour-month-weekends fixed effects, which we include to control for unobserved consumption heterogeneity across time. The coefficients of interest are β^i_w and β^i_s, which measure how much more a household consumes in response to a temperature increase in winter (summer) relative to other times of the year. We only account for daytime responses because according to the Household and Environment Survey 2008 carried out by the Spanish National Statistics Institute (INE), around 95% of households turn off their AC at night (and around 80% of households turn off their heating at night). We also include the term $\beta^i temp_{hdm}$ to control for the general trend of each household.

First, regarding the response to temperature, one would expect that a household with AC would consume more in the summer as temperature increases. Therefore, β^i_s should be positive. Similarly, households with electric heating are expected to increase their consumption as temperature decreases in winter. Therefore, β^i_w should be negative. Second, regarding the mean consumption level for different seasons across households, one would expect that households with AC (electric heating) would consume more in summer than on average, i.e., $\alpha^i_s > \alpha^i$ for households with AC and $\alpha^i_w > \alpha^i$ for households with electric heating.

After getting the seasonal consumption differences in slope and in levels for each household, we estimate the criteria for HVAC status by matching the estimated state level market shares with the surveyed market share shown in Table A.1. To identify household AC ownership, we perform the following optimization:

$$\min_{\underline{\beta}_s, \underline{\alpha}_s} \quad \sum_s (s^{AC}_s - \hat{s}^{AC}_s(\underline{\beta}_s, \underline{\alpha}_s))^2 + \lambda(\hat{s}^{AC}_L(\underline{\beta}_s, \underline{\alpha}_s) - \hat{s}^{AC}_H(\underline{\beta}_s, \underline{\alpha}_s)) \tag{B.2}$$

$$\text{s.t.} \quad \hat{s}^{AC}_s(\underline{\beta}_s, \underline{\alpha}_s) = \sum_{i \in s} \mathbb{1}(\alpha^i_s - \alpha^i > \underline{\alpha}_s) \times \mathbb{1}(\beta^i > \underline{\beta}_s + 1.96\sigma^i_s) \tag{B.3}$$

$$\hat{s}^{AC}_L(\underline{\beta}_s, \underline{\alpha}_s) = \sum_{i \in L} \mathbb{1}(\alpha^i_s - \alpha^i > \underline{\alpha}_s) \times \mathbb{1}(\beta^i > \underline{\beta}_s + 1.96\sigma^i_s) \tag{B.4}$$

$$\hat{s}^{AC}_H(\underline{\beta}_s, \underline{\alpha}_s) = \sum_{i \in H} \mathbb{1}(\alpha^i_s - \alpha^i > \underline{\alpha}_s) \times \mathbb{1}(\beta^i > \underline{\beta}_s + 1.96\sigma^i_s) \tag{B.5}$$

Table B.1: Estimated threshold values

AC thresholds		EH thresholds	
$\underline{\beta}_s$	$\underline{\alpha}_s$	$\underline{\beta}_w$	$\underline{\alpha}_w$
0.50	0.01	-4.00	0.05

The first term in the objective function captures, for each state s, the difference between the surveyed AC market share, s_s^{AC}, and the estimated one, $\hat{s}_s^{AC}(\underline{\beta}_s, \underline{\alpha}_s)$, for given thresholds $(\underline{\beta}_s, \underline{\alpha}_s)$ for having or not having AC. The second term is a penalizing term that prevents the low contracted power group from having a higher share of AC ownership than the high contracted power group. We add this penalty because an AC or an electric heating system requires higher power usage. Thus, households with electric HVAC are expected to contract more power capacity. We follow a similar procedure to estimate electric heating status. Figure B.1 shows 4 example households and depicts the correlation of household consumption and temperature in our data. The patterns are similar the results from Dyson et al. (2014). Estimation threshold values are reported in Table B.1.

To show that our classification of household types is informative both about their consumption behavior, we plot consumer daily load curves by identified HVAC status in Figure B.2. EH owners have relatively higher consumption during both day and night because electric heating devices are in general more energy consuming than AC, as show in Panel (a). We also observe that high consumption is particularly high during winter months for households that have electric heating, while it is peaking in the summer for those households with air conditioning, as shown in Panel (b).

We also use the individual EH status variable in the GMM procedure to infer income. This allows us to infer the income distribution of households as a function of HVAC status. As shown in Figure B.3, we find that electric heating is particularly concentrated on the low-income bins, while air conditioning is positively correlated with income, which is intuitive. We also show that the patterns of HVAC status and income can change depending on the region. While air conditioning tends to be associated with high income (for regions with a meaningful share of air conditioning), electric heating is negatively correlated with income particularly in the most urban regions (Madrid), as newer building tend to rely on city gas for heating.

Figure B.1: Example: household hourly consumption and temperature

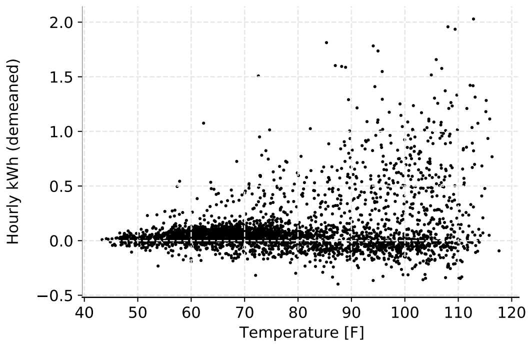

(a) Identified AC owner 1 hourly kWh (demeaned)

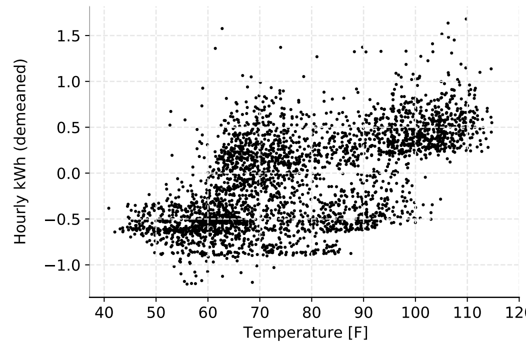

(b) Identified AC owner 2 hourly kWh (demeaned)

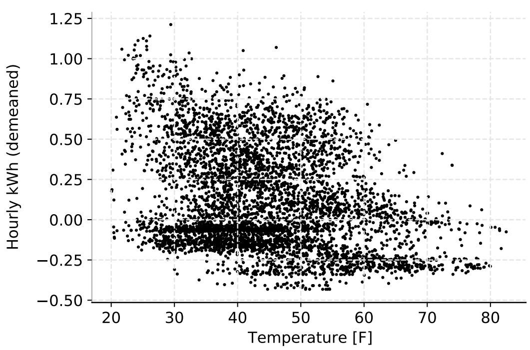

(c) Identified EH owner 1 hourly kWh (demeaned)

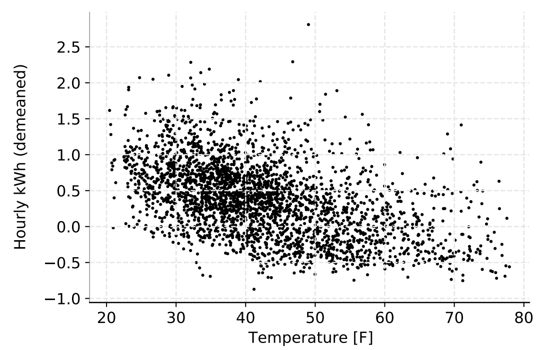

(d) Identified EH owner 2 hourly kWh (demeaned)

Notes: These figures show hourly consumption for one representative household. These figures show the correlation between household consumption and temperature and the variations that help us identify HVAC status. All data points are demeaned at household-hour level. For identified AC owners (the upper panels), we plot data points during June, July, August, and September. For identified EH owners (the lower panels), we plot data points during the November, January, February, and March. The lower panels have more data points because we have January-March data for 2016 and 2017. The four panels are: (a) an AC owner that responses partially; (b) an AC owner that responses in all hours; (c) an EH owner that responses partially; (d) an EH owner that responses in all hours.

Figure B.2: Load curves by HVAC status

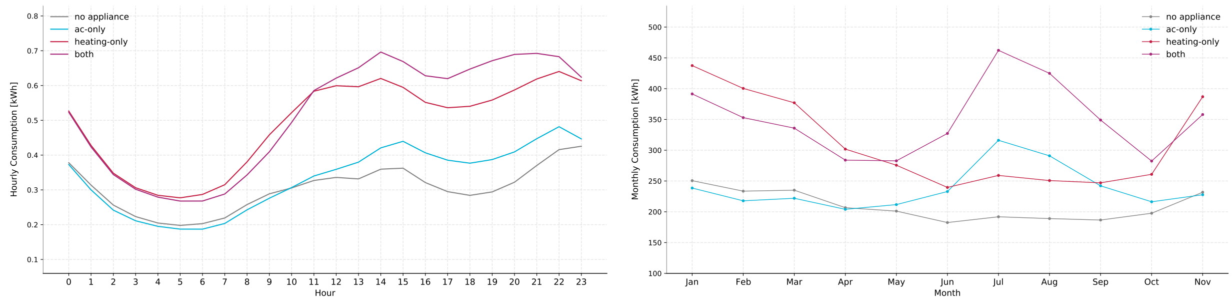

(a) Hourly consumption pattern [kWh]

(b) Annual consumption pattern [kWh]

Notes: These figures show consumption profiles over the day (the left panels) and the year (the right panels) for households with electric heating, AC, or both.

Figure B.3: HVAC status and Income

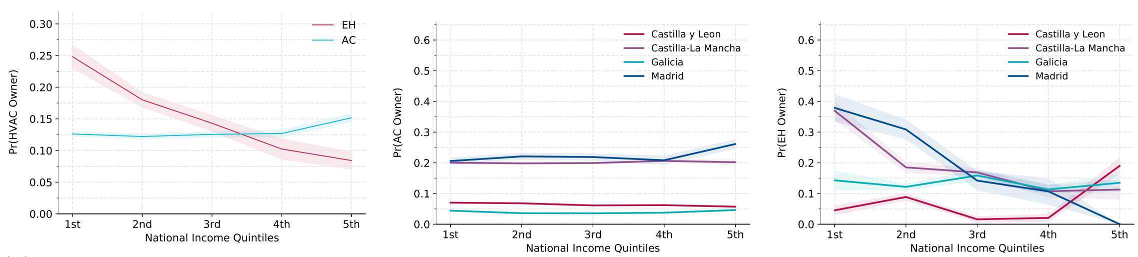

(a) Share of electric heating owners and AC owners

(b) AC ownership by state

(c) EH ownership by state

C Monte Carlo for inferring household income

Our estimator provides a refined probabilistic assignment of income to households that is more granular than the income distribution at the zip-code level. In order to understand the performance of our estimator in small samples and under misspecification, we perform a Monte Carlo simulation. We use the smart-meter household-level consumption data from our sample and create a data generating process in which we know each individual's income. We assign types to individuals based on their consumption profiles, which we then use to assign them to a certain income bin, respecting an assumed distribution of income, and zip code. We then aggregate the randomly-assigned incomes to the zip-code level, so that we can compute the distribution of income at the zip-code level, which is what the econometrician can observe.

The more detailed steps are as follows:

1. We summarize smart-meter household-level data (one entry per household) to reduce the dimensionality of the data and assign them to a group.

2. We classify households into five types for each group using a kmeans algorithm based on their hourly market shares and total consumption.

3. We sort types based on their peak market share (hours 8 - 23). We assign the more "peaky" types to a higher income distribution, reflecting the within-month correlation of peak consumption and income.

4. This distribution is fixed conditional on a type and is the same across all zip codes within a group, but we introduce some noise to capture unmodelled randomness in the data.

5. Using the assigned types, we assign a zip code number to each household based on a pre-established probability that type θ belongs to zip code z, $Pr(z|\theta)$.

6. These zip codes and the zip-code-level income distribution, together with the household-level consumption patterns, is what is observed for the estimation.

These steps allow us to create an individual and zip-code-level distribution of income that is consistent with the underlying types and assumptions. It also allows us to create an aggregate version of the income data at the zip-code level.

We then compare two methods:

- the naïve approach, which assigns the zip-code distribution of a zip code to all households within a zip code,

- our two-step approach, which classifies first households within a region into N *kmeans* types and then fits the aggregate distribution of income via GMM.

As explained in the main text, our method's goal is to better infer the income distribution of a given household. In our Monte Carlos, we know the true type of households and can compute their

Figure C.1: Simulation results: Imputed income by household type

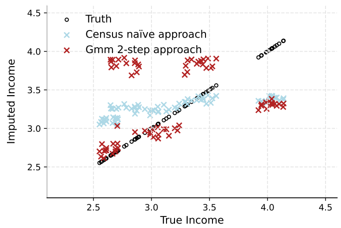

(a) 2 types, sample size 200

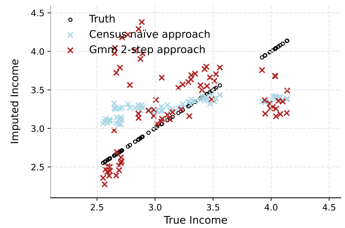

(b) 3 types, sample size 200

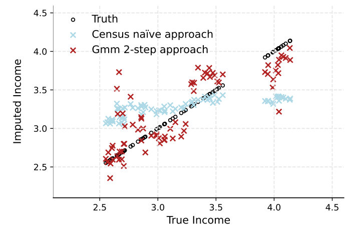

(c) 4 types, sample size 200

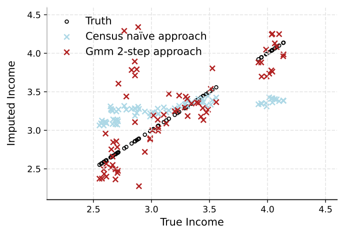

(d) 5 types, sample size 200

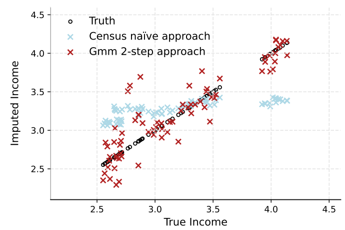

(e) 5 types, sample size 500

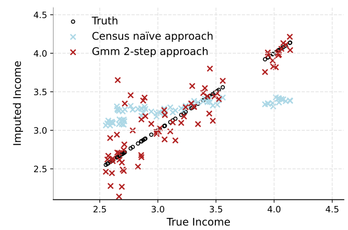

(f) 5 types, sample size 1000

expected income. We compare it to our inferred income. In the case of the naïve approach, this amounts to imputing the same expected income to all consumers in a zip code. In the case of our method, the imputation will be by estimated type. Figure C.1 shows that the naïve distribution of income tends to be much flatter (i.e., homogeneous) than the true distribution. Using our method, the inferred expected distribution lines up much better with the truth. This fit is naturally improved as we allow for more types and data.

Another way to show this result is by showing the inferred distribution of income of households belonging to a given quintile. In our Monte Carlos, we simulate a household's quintile. A household simulated to belong to the fifth quintile should have an underlying expected distribution with higher income. However, neither of these objects are known to the econometrician. We find that the naïve approach fails to estimate that households belonging to high quintile have a higher distribution of income. Instead, the probability of having a certain level of income is very similar across households along all quintiles, as shown in Panel C.2a. As we allow for more types, the distribution of income of households becomes more different along quintiles, as shown in Panel C.2c.

Finally, we examine if our inferred income is still more correlated with true income than with the naïve approach in the presence of errors in Figure C.3. We find that misclassifying zip codes into heterogeneous groups still leads to an improved correlation between imputed income and the true expected income.

Overall, the Monte Carlo simulation is useful to highlight the value of our approach. With enough flexibility, we are able to unveil within-zip-code heterogeneity that would be muted using a naïve approach. As long as we allow for sufficient flexibility and have enough data, this classification appears to improve the inferred household income in expectation.

Figure C.2: Simulation results: Distribution of imputed income conditional on true income quintile

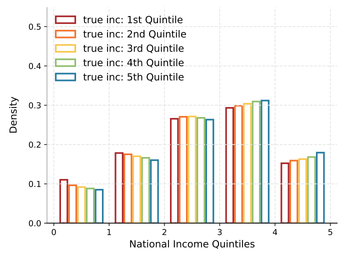

(a) Naïve approach, sample size 500

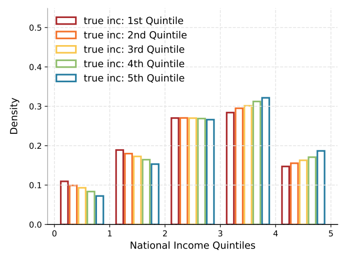

(b) 2 types, sample size 500

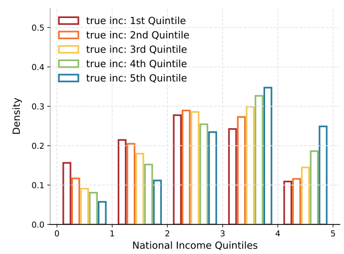

(c) 5 types, sample size 500

Figure C.3: Simulation results: Imputed income (5 types, wrong zip code group, sample size 1000)

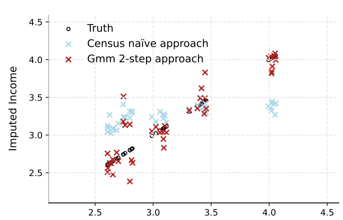

(a) Imputed type-average income vs True income

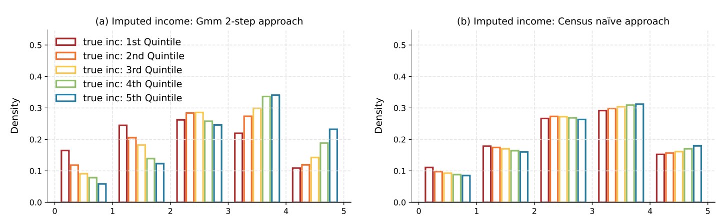

(b) Distribution of imputed income for each true income quintile

Figure D.1: Distribution of prediction errors

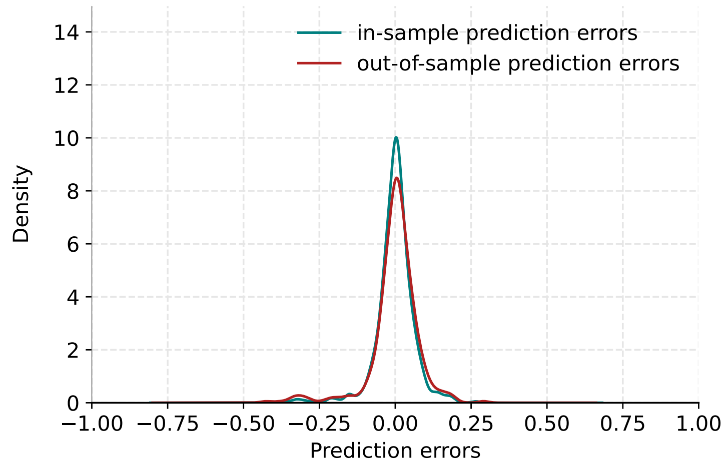

Notes: These figures illustrate the distribution of zip-code-level prediction errors.

D Cross-validation and CEX out-of-sample moments

We have performed two sets of tests to uncover the advantages of our GMM approach relative to the naïve approach. First, we conduct cross-validation by including a subsample of zip codes in each province and predicting the out-of-sample income distribution for the other zip codes. Figure D.1 reports the distribution of in-sample and out-of-sample prediction errors. The two distributions are similar, which suggests that our approach captures the true relationship between income and consumer types. It can therefore can handle out-of-sample predictions.

Second, we report the estimated relationship between income and electricity consumption using the GMM approach and the naïve approach. In Figure D.2, we compare the estimated results to the CEX survey data. In all provinces, the naïve approach captures only across zip code income variation and cannot explain the relationship between income and consumption. On the contrary, the GMM approach performs better even though the relationship between income and electricity consumption is flatter than in the CEX survey data.

Two reasons explain this departure. First, our approach has limitations. As explained in Section 5.1.1 and 7, when income has a direct impact on the variable of interest (e.g., the bill change impact in our main context or electricity consumption in this section), our approach does not capture the full relationship. We believe there is no direct impact of income on RTP bill changes other than through the correlation of income with consumer types (e.g., lifestyle, ...). However there might be a direct impact of income on electricity consumption. Thus, in Figure D.2, the GMM results have a lower slope than the CEX survey results, as expected in the simulation 7. Second, our data covers only a subset of zip codes, especially in Madrid where we only have data in relatively low income zip codes. The different coverage between our sample and the CEX sample accounts for the differences across the two figures.

Figure D.2: Relationship between income and electricity consumption: Comparing survey results, the imputed income results, and the naïve income results

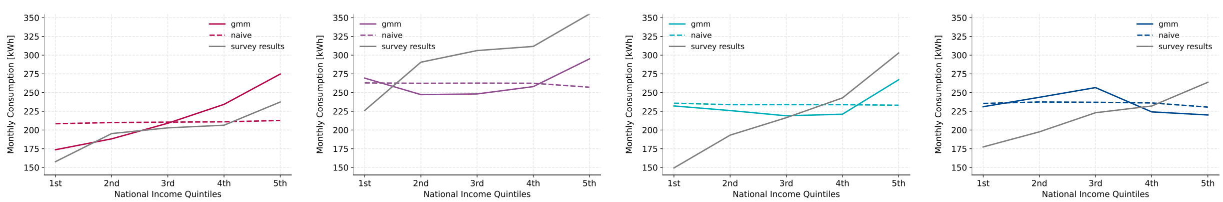

(a) Castilla y Leon (b) Castilla-La Mancha (c) Galicia (d) Madrid*

Notes: These figures show that the imputed income from our approach is closer to the actual income from the survey data as compared to the naïve approach. The figures depict the relationship between electricity consumption and income for the four regions in our sample. In each figure, the gray line respresents the CEX survey results, the blue line represents the results using the GMM approach, and the dashed blue line those from using the naïve approach.

*Our utility data for Madrid only covers parts of the region's territory.

E Robustness to alternative specifications

E.1 Details about the main specifications

Table E.1: Choice of specifications

State	Prov. ID	# HHs (# zips)	Total # types	number of kmeans types within category			
				$(noEH, L)$	$(noEH, H)$	(EH, L)	(EH, H)
Galicia	15	227,247 (95)	12	3	3	3	3
	32	83,315 (114)	10	3	3	3	3
	36	160,585 (64)	8	3	3	1	1
Castilla	24	43,584 (105)	8	3	3	1	1
y León	40	57,106 (131)	12	3	3	3	3
Madrid	28	340,409 (30)	8	3	3	1	1
Castila-	13	132,968 (47)	8	3	3	1	1
La	19	23,262 (36)	8	3	3	1	1
Mancha	45	80,809 (50)	10	3	3	1	1

Notes: EH ($noEH$) means households with(out) electric heating, H and L represent household contracted power choice. We mostly choose fewer types within the EH categories because the size of the population is smaller. We group all zip codes in the same province and estimate them together. Robustness checks show that our results are not sensitive to the zip code classification.

E.2 Robustness

We conduct sensitivity analysis of our estimates regarding several choices in our estimating procedure. We assess robustness with respect to the following dimensions:

- Number of k-means types.
- Inclusion of contracted power and/or heating categories.
- Urban-rural zip groupings within a province.

Number of k-means types Figure E.1 shows the results of setting different numbers of types. The gray line shows the naïve approach without any within zip code heterogeneity. The green line shows the results when including the contracted power type and the HVAC status type, and no flexible type (*kmeans* type) is allowed. The yellow dashed, red solid, and dark red dashed lines correspondingly show results from specifications with 2, 3, and 4 *kmeans* types ($N^g = 4, 12$, and 16). All specifications include contracted power type and HVAC status. Only the main specification's standard error is included to keep the figure easy to read. The standard errors of other specifications are similar to the ones in the main specification.

As expected, Figure E.2 shows that the estimated distributional impacts are more pronounced with more types. Once we include contracted power type and HVAC status, the result is already correcting some of the bias of the naïve approach, showing that households in the low-income

Figure E.1: Estimated Bill Changes [%] from Alternative Specifications (I)

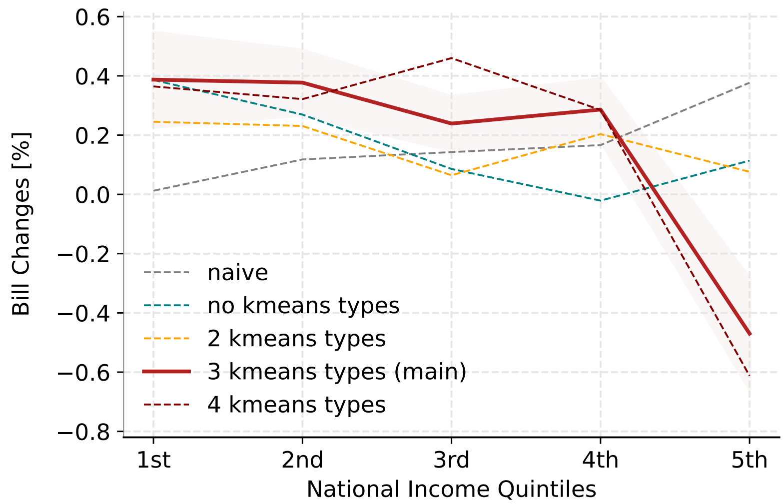

Notes: This figure represents the estimated bill increase in % when moving from an annual time-invariant price to RTP under different specifications. Different lines represent different numbers of types set in the estimation. Only the main specification's standard error is included to keep the figure easy to read.

quintile lose from the switch to RTP. The reason is that contracted power is highly correlated with household income, making it powerful in identifying the distributional impact. Moreover, the result from the 4 *kmeans* type ($N^g = 16$) specification falls in the confidence interval of the result from the main specification (3 *kmeans* types, $N^g = 12$).

Observable and inferred characteristics Figures E.2a and E.2b show what happens when we ignore the individual characteristics of households to classify them into types. These characteristics are either directly observed at the household level (contracted power) or inferred by exploiting the rich patterns in smart meter data (presence of electric heating). Ignoring these characteristics could potentially reduce our ability to recover income precisely.

We find that ignoring contracted power does not affect our findings. Even though contracted power has high predictive power when it comes to income, our *kmeans* clusters and electric heating dummies already replicate our main results, as shown in Figure E.2a.

We find that removing the electric heating classification affects some of the patterns in our findings. In particular, we find that the highest quintile is negatively affected by the policy change. However, the impacts remain regressive on average.

Rural vs urban consumers To account for type differences between rural and urban customers within a province, we use population to account heterogeneity more flexibly. The 750 zip codes in our sample are located in 9 provinces. For the provinces with more than 100 zip codes, we further classify the zip codes into more groups based on population. If estimation results suggest little heterogeneity across the within-province groups, we keep the whole province as one group in our

Figure E.2: Estimated Bill Changes [%] from Alternative Specifications (II)

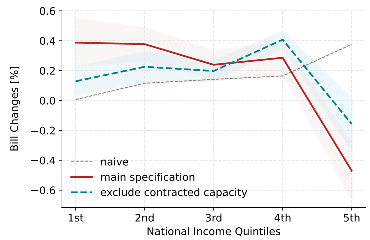

(a) Exclude household contracted power

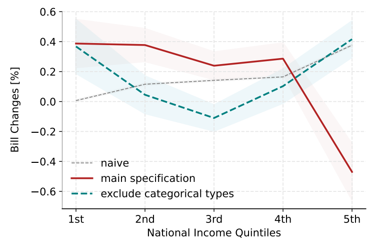

(b) Exclude all categorical types (electric heating and contracted power)

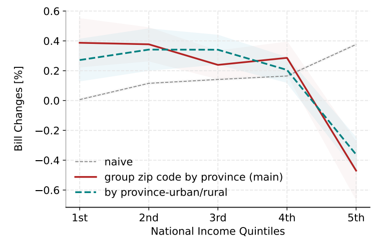

(c) Group zip codes by province-urban/rural

main specification. Unfortunately, we can only allow for urban-rural differences in the provinces with the most significant number of zip codes.

Figure E.2c shows that our main results are mainly unaffected by allowing for additional rural-urban heterogeneity. The one exception is for the province with postal code 32, Ourense, which has many very small zip codes in rural areas.

F Robustness to semi-parametric approach

We report results for 4 different semi-parametric specifications. Specification 1 only includes households' contracted power capacity (x_i) and does not include any zip code-level demographic variables (z_j). Specification 2 includes households' contracted power capacity and zip codes' population. Specification 3 includes households' contracted power capacity and a zip codes' age distribution variable (i.e., the probability of being older than 65). Specification 4 includes all three variables. Figure F.1 reports the estimated impact from our main non-parametric specification, the naïve approach, and all the four semi-parametric specifications.

Figure F.1: Estimated Bill Changes [%] from Semiparametric Specifications

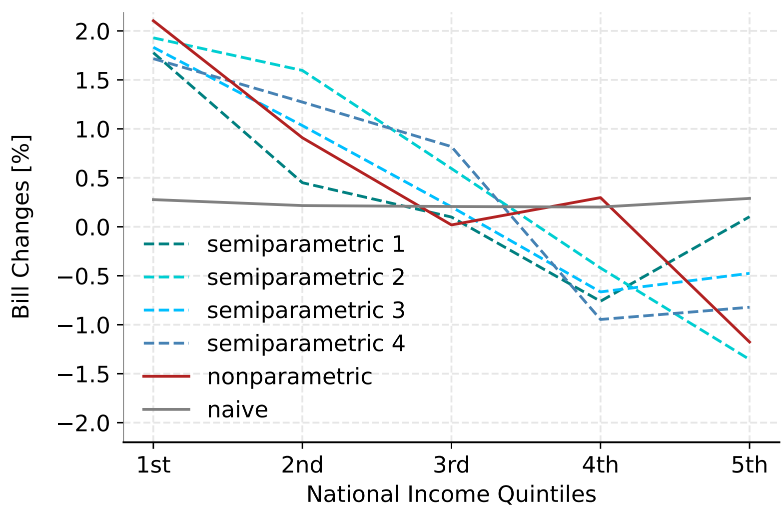

Notes: This figure represents the estimated bill increase in % for consumers in Madrid when moving from an annual time-invariant price to RTP under different specifications. Different dashed lines represent different semi-parametric specifications in the estimation. The red line represents the main specification's results and the gray line represents the naïve approach's results.

Figure F.2: Estimated Income Distribution of Two Example Types from Semiparametric Specifications

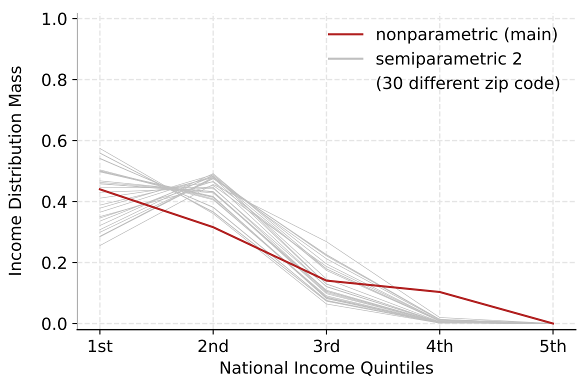

(a) Income distributions of type 0

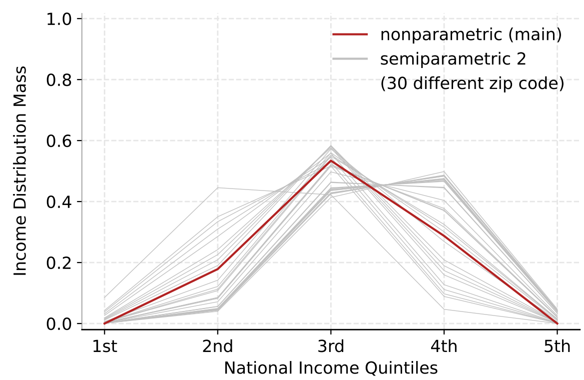

(b) Income distributions of type 7

Notes: This figure represents the imputed income distribution for two example types in Madrid. The red lines report the estimated income distribution for the two types from the non-parametric specification (the main specification). In our main specification, we assume that within a province, households' income distribution is the same for the same type across different zip codes. The grey lines report the estimated income distributions from the second semi-parametric specification. In the semi-parametric specifications, we allow households' income distribution to be different according to their contracted power capacity and zip-code-level demographics. Each gray line represents type 0 (type 7) in a different zip code.

G An alternative estimator

In this Appendix we consider an alternative estimator to our two-step approach. In particular, we rely on a more parametric approach to describe the relationship between household income and electricity consumption. As explained in the main text, we assume that household allocation of their hourly electricity consumption during the day (denoted kWh_{ih} and suppressing day index) is determined by a set of variables, such as hourly electricity prices and temperature, their type θ_i, and some random shocks ϵ_{ih}. We reproduce equation (8) here:

$$kWh_{ih} = f(x_h, \epsilon_{ih}|\theta_i).$$

Allowing the household's type θ_i to correlate with its income helps us identify how income affects electricity consumption. This section uses a parametric model to describe the function f and identify the relationship between household income and consumption, following the spirit of Berry et al. (1995) and Berry et al. (2004).

Just as we did in the main text, we estimate the demand system for each province separately because of the large heterogeneity across regions; see Section 6. All the parameters are province-specific. To save on notation, we suppress the province subscripts.

G.1 An alternative parametric approach

Let us assume that each household is optimizing its daily consumption across 24 hours. Thus, we can think of each household-day as a "market", which we repeatedly observe every day. A household's consumption is correlated with its income $inc_i \in [1, 2, 3, 4, 5]$, indicating each quintile of the Spanish national income distribution. Let $K = 5$ denote the total number of income bins.

Our income data contains zip-code level average income per household and the distribution of household income within each zip code. We know the proportion of households in each national income quintile. Therefore, unlike most papers in the literature that assume that inc_i follows a normal distribution, we assume that inc_i follows a discrete distribution with $inc_i \in [1, 2, 3, 4, 5]$, representing the national income quintiles. We observe the PMF of income in each zip code.

In day d, the consumption of household i in zip code z at hour h is given by:

$$kWh_{ih,d} = g(inc_i) \times s_h(inc_i, \{p_{hd}\}_{h=1}^{24}, \{temp_{zhd}\}_{h=1}^{24}|\beta^i) \times \xi_{zhd} \times \varepsilon_{ihd}, \tag{G.1}$$

$$g(inc_i) = exp(\eta^\alpha \mathbb{1}(inc_i = k) + \sigma^\alpha \nu_t^i). \tag{G.2}$$

We use $g(inc_i)$ to denote the daily household's scale factor, and s_h to denote the share of kWhs consumed in this hour. We allow income to affect both of these processes. ξ_{zhd} is a zip code-hour-day level error term, and ε_{ihd} is a household-hour-day level error term, both of which follow a normal distribution with a zero mean.

We assume that the factor $g(inc_i)$ follows a log-Normal distribution. η^α is a vector with length K, where the kth element is the mean electricity consumption for income bin k, and σ^α is the scale parameter for this log-Normal distribution.

We assume that the share s_h is a function of hourly prices p_{hd}, hour-zip-code-specific temperature $temp_{zhd}$, and household income inc_i. The detailed form is as follows:

$$s_h(inc_i, \{p_{hd}\}_{h=1}^{24}, \{temp_{hd}\}_{h=1}^{24} | \beta^i) = \frac{exp(u_{ih,d})}{1 + \sum_{h=1}^{21} exp(u_{ih,d})} \qquad \forall h \in [4, 23], \quad \text{(G.3)}$$

$$u_{ih,d} = [p_{hd} \quad temp_{zhd} \quad 1]\beta_t^i + \epsilon_{ih,d} \qquad \forall h \in [4, 23], \quad \text{(G.4)}$$

$$\beta_t^i = \begin{bmatrix} \beta_{t,1}^i \\ \beta_{t,2}^i \\ \beta_{t,0}^i \end{bmatrix} \qquad \forall t \in [1, 2, 3, 4, 5], \quad \text{(G.5)}$$

$$= \begin{bmatrix} \beta_{t,1}^0 \\ \beta_{t,2}^0 \\ \beta_{t,0}^0 \end{bmatrix} + \begin{bmatrix} \eta_{t,1}^\beta \\ \eta_{t,2}^\beta \\ \eta_{t,0}^\beta \end{bmatrix} inc_i + \begin{bmatrix} \sigma_{t,1}^\beta \nu_t^i \\ \sigma_{t,2}^\beta \nu_t^i \\ \sigma_{t,0}^\beta \nu_t^i \end{bmatrix}, \qquad \text{(G.6)}$$

where h is an index for an hour, and t is an index for a time window, with $t = 1, ..., 5$ indicating the time intervals 4am-7am, 8am-11am, 12pm-3pm, 4pm-7pm, and 8pm to 11pm, correspondingly. 0-4 am are defined as outside options given that electricity consumption is mostly passive during that time window.

We assume that a household's utility of allocating 1 kWh into a certain hour h is (G.4), where β_t^i is explained by both a set of random draws μ_t^i and income inc_i. The first, second, and third elements of β_t^i are coefficients for prices, temperature, and a constant, respectively. We assume these random coefficients follow normal distributions, with variances σ^β and means $\beta^0 + k \times \eta^\beta$ $\forall k \in [1, 2, 3, 4, 5]$. We allow the coefficients to be different across time windows. The coefficients for different time windows are correlated only through household income.

To simplify the notation, let $\theta^\alpha = (\sigma^\alpha, \eta^\alpha)$ denote all consumption level parameters, and let $\theta^\beta = (\{\beta_t^0, , \eta_t^\beta, \sigma_t^\beta\}_{t=1}^{t=5})$ denote the parameters related to the consumption allocation within a day. Notice that the level of household consumption is independent from the kWh allocation parameters, $\theta^\beta = (\beta_p, \beta_t, \beta_0, \eta^\beta)$. Therefore, we can identify consumption level parameters $\theta^\alpha = (\sigma^\alpha, \eta^\alpha)$ separately.

We use a simulated method of moments to estimate θ^α and θ^β in this demand system. We draw inc_i from the income distribution of a zip code and draw random draws μ. We then aggregate the implications for $\overline{kWh}_{zh}$ in each zip code, and match the moments to fit the kWh patterns at the zip-code level:

$$\sum_d \left(kWh_{zh,d} - kWh_{hd}(\theta^\alpha, \theta^\beta, inc_z) \right)^2 \quad \forall z, h, \qquad \text{(G.7)}$$

$$kWh_{hd}(\theta, inc_z) = \int_{inc_i} \int_{\nu^i} kWh_{hd}(\theta, \nu, inc_i) dF_\nu dF_{inc_i}^z \quad \forall z, h, d. \qquad \text{(G.8)}$$

The second equation above is the integral of equation (G.1). As househhold income only affects the mean of the parameters, the above moment condition is equivalent to:

$$\sum_d \left(\overline{kWh}_{zh,d} - \sum_k Pr_z(inc_k)kWh_{hd}(inc_k)\right)^2 \quad \forall z, h, \tag{G.9}$$

$$kWh_{hd}(inc_k) = \int_{\nu^i} kWh_{hd}(\theta^i)dF_\theta^k, \quad \forall k, h, d, \tag{G.10}$$

where F_θ^k is the distribution of the demand parameters $\theta^i = (\beta_1^i, \beta_2^i, \beta_0^i)$.

In the parametric approach, the correlation between income and types is parameterized by η^β, and it is the same across all zip codes. Also, the distribution of types $\theta^i = (\beta_1^i, \beta_2^i, \beta_0^i)$ is the same for all zip code regions, conditionally on zip code demographics. We need these assumptions to give us identification power.

To make use of the repeated data for each household, we add a set of covariance moments to capture the relative attractiveness of different hours to the same consumer. These covariance moments allow households with a higher coefficient in the morning time window to have a lower (or higher) coefficient in the noon time window. They connect utility for different hours for the same household and help identify the distribution of the random coefficients, β_t^i. Because we assume that the β_t^i for different time windows can only be correlated through inc_i, these covariance moments help identify the income coefficient η^β.

Because data suggest considerable heterogeneity in household consumption patterns across months, we allow households to have month-specific utilities, i.e., equation (G.1) and all the coefficients above are month-specific. Therefore, similar to the covariance moments above, we add one more set of across-month covariance moments to make use of the relationship between summer and winter coefficients, which gives us identification power from the seasonal patterns.

Because there are thousands of covariance moments by combining all the time windows and all the months, it is hard to use all the across-time windows and across-months variations (which is why we propose our two-step estimator). Due to the computational burden, we only use data from January and August 2016 and only use the following two sets of covariance moments:

$$\left[Cov\Big(kWh_{zh1,i}^m, kWh_{zh2,i}^m\Big) - Cov\Big(kWh_{h1,d}^m(\theta^\alpha, \theta^\beta, inc_z), kWh_{h2,d}^m(\theta^\alpha, \theta^\beta, inc_z)\Big)\right]^2 \quad \forall m, (h1, h2) \text{ pair}, \tag{G.11}$$

$$\left[Cov\Big(kWh_{zh,i}^{m1}, kWh_{zh,i}^{m2}\Big) - Cov\Big(kWh_{h,d}^{m1}(\theta^\alpha, \theta^\beta, inc_z), kWh_{h,d}^{m2}(\theta^\alpha, \theta^\beta, inc_z)\Big)\right]^2 \quad \forall h, (m1, m2) \text{ pair}, \tag{G.12}$$

where m1 indicates January 2016 and m2 indicates August 2016.

The estimated distributional impact of the parametric approach can be computed by predicting the bill impacts implied by different pricing schemes at different income levels. The results are presented in Table G.1. One can see that the results are in between those of the naïve approach and our two-step estimator. Still, the estimated bill impacts underestimate the distributional impact across income bins. Several reasons may lead to the poor performance of these parametric

Table G.1: Estimated distributional impact of Madrid (Bill changes [%] by income)

	Naïve	Parametric approach	Two-step estimator
1st Quintile	0.54	0.93	1.73
	–	(0.82)	(0.12)
2nd Quintile	0.46	0.72	0.77
	–	(0.68)	(0.13)
3rd Quintile	0.38	0.42	0.04
	–	(0.11)	(0.11)
4th Quintile	0.35	0.13	0.09
	–	(0.05)	(0.09)
5th Quintile	0.49	-0.05	-1.47
	–	(0.06)	(0.19)

approaches to electricity consumption, as we discuss below.

G.2 Limitations of the alternative estimator

Unobserved Heterogeneity within region The parametric approach uses demographics and the normally distributed error term $\sigma\mu^i$ to explain heterogeneity across types. It assumes that the distribution of types (β^is) is the same across zip codes, conditionally on demographics. These parametric assumptions ignore important unobserved heterogeneity within a region.

According to our data, different zip codes have different distributions of the random coefficients and different correlations of each household's consumption and income, as shown in Figure G.1. The figure shows the estimated distribution of individual price coefficients (Panels (b) and (c))[51] and the income distribution (Panel (a)) of four zip codes in the same region. One can see that the correlation between prices and electricity consumption can be substantially different among zip codes with similar distributions of income. The parametric approach has to use the σ^β coefficient to explain these differences. This suggests that including demographics as a way to classify households might be too limited to understand the heterogeneous impacts of real-time pricing. Because there is substantial unobserved heterogeneity within zip codes, our approach with discrete types attempts to allow for greater flexibility.

One might think that we can make the utility function more flexible to better suit the actual electricity consumption patterns. However, as shown in the following equations, being fully flexible does not work, as we would underidentify the parameters. The only parametric way that might work is to compute the market shares at the household-day level, treating each household as a "market" and each "household-day" as a "household or choice maker". However, this is computationally burdensome.

If we try to modify the moment conditions (G.7) and (G.8) in a more flexible way, a natural way

[51] Notice that, with individual-level data, we can estimate β^i for each individual using individual-level logit regressions. The estimated β^i helps reduce the distributions of the high dimensional outcome variable kWh into the distribution of β^i.

Figure G.1: Evidence of heterogeneous distribution of beta, even conditional on zip code income distribution

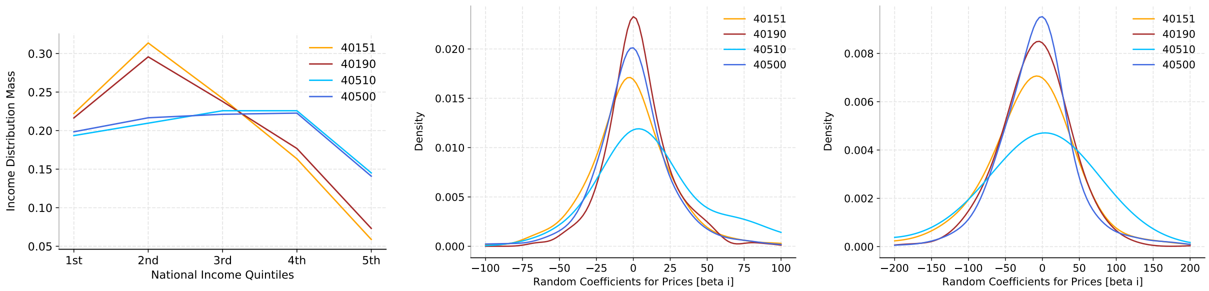

(a) Income distributions

(b) Distribution of individual price coefficients $\beta^i_{t=1,1}$, 4-7am

(c) Distribution of individual price coefficients $\beta^i_{t=3,1}$, , 12-3pm

Notes: This figure illustrates the distribution individual coefficients for different zip codes. The distributions are heterogeneous, even when conditioning on zip code income distributions. The two blue zip codes have a similar income distribution and the two red zip codes have a similar income distribution. However, the distributions of the β^i coefficients are different for each pair of zip codes, regardless of the time windows. This evidence violates the classic assumption in the parametric approach.

would be to consider the following two equations, which allow different zip codes to have different distributions of type $r \in [1, 2, .., R]$, denoted by $Pr_z(r)$, in a flexible manner. A type r can indicate a cluster of coefficient patterns (e.g., in the fixed grid nonparametric approach), or kWh patterns (e.g., in our main specification), or both. The relationship between income and household types is still characterized by η^k:

$$\sum_h \left(\overline{kWh}_{zh} - \sum Pr_z(r) kWh_{zh}(r)\right)^2, \quad \forall z, \tag{G.13}$$

$$\sum_k \left(Pr_z(inc_k) - \sum_{i \in z} \sum_r \eta^k_{zr} Pr_z(r)\right)^2, \quad \forall z, \tag{G.14}$$

$$\text{s.t.} \quad \sum_k \eta^k_{zr} = 1, \quad \forall z, \forall r. \tag{G.15}$$

Allowing fully flexible income-type distributions for each zip code does not work, because when type and income are both unknown, the above system of equations (G.13) and (G.14) is greatly underidentified without imposing more structure. Allowing for zip-code-specific types, i.e., assigning only one type to a zip code with probability one, can perfectly match the aggregate moments. The result becomes equivalent to the “naïve” approach of assuming that all households within a zip code have the same income distribution. A natural modification to this is our two-step estimator, which allows for flexible types and can easily identify income-type distributions under the assumption that there is sufficient overlap in types across zip codes.

Limits to parameterization The second problem with using a parametric approach is that electricity consumption is a highly dimensional choice. Price or temperature variation are thus

Figure G.2: Evidence of agnostic consumption patterns of different types of households

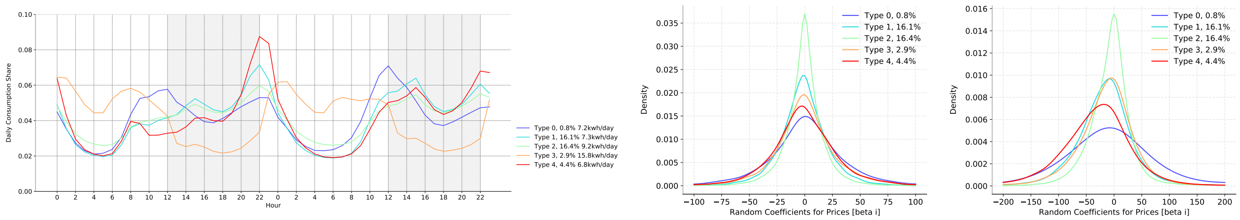

(a) Kmeans types for Madrid

(b) $\beta^i_{t=1,1}$, 4-7am

(c) $\beta^i_{t=3,1}$, 12-15am

Notes: This figure illustrates the agnostic consumption patterns by Kmeans types (Panel (a)) and the distribution of β^is for each Kmeans type (Panel (b) and (c)). The β^is are defined in the parametric model and are estimated through individual-level logit regressions. The Kmeans types are clearly distinct, while the distributions of β^is are very similar across Kmeans types. This indicates that the β^is, although capturing some aspects of the consumption patterns, are not informative enough for classifying consumer types.

Figure G.3: Estimated load curve from the parametric estimation

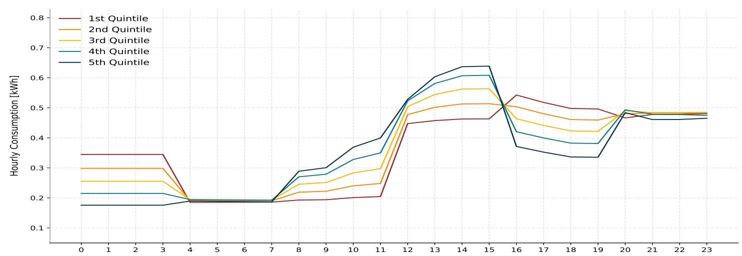

Notes: This figure depicts the predicted load curve from the parametric model. Compared to Figure G.2 Panel (a), this predicted load curves are less flexible.

insufficient to explain it. The vector β might not be a proper primitive to model electricity consumption, and equation (G.3) might not be sufficiently detailed for our purposes.

Using an over-simplified model would ignore unobserved heterogeneity in consumption patterns. Two households of different consumption types might have very similar price correlations. If we ignored this, our income estimates would be biased. Figure G.2 provides an example showing that types with similar price coefficients may still have substantially distinct consumption patterns. The distributions of the individual price coefficients $\beta^i_{t=1,1}$ are homogeneous, as can be seen in Panel (b). However, as shown in Panel (a), the types are clearly heterogeneous even within the same time window.

Moreover, we can compare Figure G.2 Panel (b) and Figure G.1 Panel (b). These two figures plot the distributions of the same coefficients. The heterogeneity across types is even smaller than the heterogeneity across zip codes. Thirdly, both Panel (b) and (c) from Figure G.2 imply that the blue type (type 1) and the orange type (type 3) are similar to each other. However, they might have different lifestyles (types) and probably have different income distributions. The parametric approach will ignore this heterogeneity and therefore give biased results.

The parametric model can capture some aspects of the existing heterogeneity. Figure G.3 indicates that higher income household consumption is more procyclical, consistently with our main model findings. However, the model does not allow for much further heterogeneity (other than via noise in the random coefficients). We miss the heterogeneity present in Panel (a) of Figure G.2, even though we have 45 parameters per month-province. Thus, it may be hard to use any parametric functional form to describe the market share choices of different types. Comparing the consumption patterns and the price variation, we know that β in equation (G.6) is not a sufficient statistic for all household types.

Overall, we need some tools to simplify the high-dimensional heterogeneity in this setting and cluster households into a smaller number of types. In short, we need a dimension reduction tool. As explained in Bonhomme et al. (2022), allowing for discrete heterogeneity (clustered by *kmeans*) is an efficient dimension-reduction device to deal with an agnostic electricity consumption model. A parametric model, which describes the data variations using a small number of parameters, can serve a similar role. It reduces the distributions of the high-dimensional outcome variable kWh into the distributions of the β^i. However, as shown above, a parametric approach would be subject to limitations, making the more agnostic *kmeans* clustering approach preferable.

Repeated observations of each individual Panel data with repeated observations from each individual provide essential variations to identify the joint distribution of income and electricity consumption. The variation across individuals can be explained by income (or other individual features) variation, while other variables can explain the variation across time for the same individual, e.g., prices and temperature.

As mentioned above, the parametric approach uses repeated observations from the same individual through the covariance moments. These covariance moments are essential for identifying the income coefficients η because they connect household income with their across hour and month consumption patterns. We have hourly electricity consumption data from each household for more than one year, i.e., around 10,000 observations per individual, giving rise to thousands of candidate covariance moments. The large number of candidate moments and the complicated format of the income-β moments are due to the lack of an explicit dimension reduction device. Ideally, one would want to include many moments to remain agnostic about how to select them, but this is computationally not feasible.

In our approach, we simplify the highly-dimensional nature of smart meter consumption in our first step. We then connect household income with their consumption patterns in our second step. The second step serves the same role as the covariance moments but it is computationally simpler. We avoid the thousands of candidate covariance moments because we have reduced the dimensions non-parametrically in the first step. The *kmeans* clustering method has helped discretize the highly-dimensional household heterogeneity. Therefore, the relationship between income and types can be more agnostic and more explicit than the income-β relationship in the parametric approach, which is very much tied to the price coefficients.

Overall, these limitations highlight that it is difficult to summarize the heterogeneity in electricity consumption data using a parametric approach. These limitations also highlight that oftentimes the problem can become computationally expensive. Our approach tries to find a data-driven compromise to handle these two difficulties.