The Distributional Impacts of Time-Varying Electricity Pricing:
A Novel Approach to Estimating Household Income

Mar Reguant ⓡ Natalia Fabra ⓡ Jingyuan Wang*

August 2025

Abstract

As interest in the distributional effects of climate policies gains prominence, designing electricity tariffs that are both efficient and equitable becomes critical. Efficiency often favors the shift from time-invariant to time-of-use (ToU) or real-time prices (RTP), but this transition may have distributional implications. We develop a framework to assess the implications and the channels through which distributional impacts manifest. Central to our approach is a novel method that infers individual household income by combining zip-code-level income data with household-level electricity consumption. We demonstrate the value of this method in the Spanish context. First, we show that using more granular estimates of income has an impact on the distributional assessment of electricity tariffs. Second, we find that the potential distributional impact of RTP is very modest compared to the impact of ToU. The most salient effect of RTP is to increase bill volatility, particularly so for the low-income group.

Keywords: electricity pricing, distributional effects, income inference.

JEL Classification: L94, H23, C55.

*Mar Reguant: ICREA-IAE, Northwestern University, CEPR and NBER, mar.reguant@northwestern.edu. Natalia Fabra: CEMFI and CEPR, natalia.fabra@cemfi.es. Jingyuan Wang: Department of Economics, University of Michigan, jingyw@umich.edu. For their comments and feedback, we thank seminar participants at Yale, UPenn, TSE, UPF, HKS, Sciences 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. We thank Michael Cahana for his reseach assistance at the early stages of the project. Comments by Mateus Souza are also appreciated. Fabra and Reguant have generously benefited from funding by the European Union (PI: Reguant, *ENECML*, 101001732; PI: Fabra, *ENERGY-IN-TRANSITION*, 101142583). Views and opinions expressed are however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. Reguant also acknowledges the support of NSF grant SES-1455084.

1 Introduction

Quantifying the distributional effects of climate policies is becoming increasingly central to both academic research (Deryugina et al., 2019) and policy discussions (OECD, 2024). Beyond environmental justice, equity concerns can hinder the political and social support necessary for implementing efficient climate policies. Therefore, understanding how policy impacts vary across income levels is crucial for designing effective and equitable measures that are socially acceptable (Fabra and Reguant, 2024). However, despite growing interest in this area, the lack of individual-level income data often limits researchers' ability to make precise assessments.

In response to this data limitation, we propose a novel method to infer individual-level income in contexts where researchers have access to detailed socioeconomic data but lack precise household income information.[1] Our method fills this gap by offering a robust approach to estimate income at the individual level, which is particularly useful for distributional analyses of climate policies, such as carbon pricing (Chanut, 2021). Beyond climate policy, this method has wider applications in fields such as public economics (Chetty et al., 2023), education (Bleemer and Mehta, 2022), finance (Gross et al., 2021), and labor economics (Gustman and Steinmeier, 2000).[2]

We demonstrate the utility of our method by analyzing the distributional effects of alternative electricity pricing policies, particularly the shift from time-invariant to Real-Time (RTP) and Time-of-Use (ToU) pricing. Economists have long advocated for RTP as an efficient policy tool (see, e.g., Borenstein (2005) and Borenstein and Holland (2005)). Still, its adoption has been slowed by concerns about potential adverse distributional impacts across households (Joskow and Wolfram, 2012).[3] ToU pricing is more widely adopted, though its desirability from an efficiency standpoint remains debated. While it does not deliver all the benefits associated with fully dynamic pricing, its predictability and salience may elicit a stronger demand response (see Fabra et al. (2021) and Enrich et al. (2022)).

Our empirical analysis leverages data from the Spanish electricity market, the only country where RTP has been broadly implemented as the default option for households, and where ToU has become compulsory.[4] Our analysis, based on hourly smart meter data from more than a million

[1]In many cases, databases contain individual-level data (e.g., on consumption, health, education) but, due to privacy concerns, only provide zip code-level information. This prevents researchers from matching households to income data at finer levels of aggregation (e.g., census tracts). Similarly, while household characteristics are well-documented in Census data, some countries only offer detailed income statistics by zip code.

[2]For example, Chetty et al. (2023) examine the heterogeneous impacts of COVID-19 on household spending by using median household income at the zip code level to proxy for cardholder income. Similarly, Bleemer and Mehta (2022) use zip code-based income data to quantify the wage return for majoring in economics. Gross et al. (2021) use zip code median income to analyze how bankruptcy laws affect the cost of credit. Lastly, Gustman and Steinmeier (2000) study retirement decisions using wage data to estimate household income, relying on proxy methods in years when survey data are unavailable. In all these cases, our method could improve the understanding of income heterogeneity within zip codes, allowing for richer distributional analyses.

[3]Levinson and Silva (2022) show how preferences for income redistribution influence the design of the electricity pricing scheme.

[4]In some countries, such as Norway and New Zealand, RTP is offered by competitive retailers, but it is not the default option. For instance, Borenstein (2013) notes: "*I'm aware of no place in the U.S. where time-sensitive rates are the default for residential customers.*" Similarly, the European Commission (2009) states, "*The case of Spain with a regulated default dynamic price contract is unique.* Pébereau and Remmy (2023) explore the barriers to adopting

Spanish households over 18 months, offers a highly representative assessment of the population-wide impacts of RTP and ToU.

Although smart meter consumption data are highly detailed, there is no household-level income data available, which is crucial for assessing the distributional effects of electricity pricing policies. The standard approach (i.e., assuming all households within a zip code share the same income distribution based on national quintiles) would suggest that the impacts of RTP and ToU are either uncorrelated with income or only weakly so. However, this method overlooks significant within-zip-code heterogeneity, potentially introducing a bias (Borenstein, 2012). By accounting for within-zip-code heterogeneity, our method for estimating household income allows us to uncover potential distributional effects that would otherwise remain hidden.

Approach To estimate household-level income data, our proposed method combines household electricity consumption data with zip code income distributions. First, we apply flexible classification algorithms to group households into representative types based on their electricity consumption profiles. Additionally, we categorize households by their contracted power, which correlates strongly with income. Once households are classified into types, we estimate each household's income distribution as a function of its type, individual characteristics, and zip code attributes. 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 for identifying the impact of household types on income is that the set of potential types is shared across zip codes within a group of nearby zip codes.[5] This assumption enables us to estimate the income distribution for each type that rationalizes the observed income distributions at the zip code level. We then derive household-level income distributions by combining the type-level income distribution with the household's assigned type. Although this identification approach is non-parametric, we also implement a semi-parametric version that incorporates functional form assumptions on how types, individual characteristics, and zip code demographics influence the distribution of income.

Since the household classification algorithm is sensitive to choices made by the researcher, we conduct three types of robustness checks to validate our results. First, we perform Monte Carlo simulations. Second, we cross-validate our predictions by leaving some zip codes out of the estimation process and then predicting their outcomes. Third, we compare our inferred consumption-income patterns with data from the Consumption Expenditure Survey (CEX).[6] As an additional validation, we demonstrate that, as expected, contracted power is strongly and positively correlated with our household income estimates. In contrast, assigning each household to the observed income

RTP in New Zealand.

[5]Not all types need to be present in every zip code within a group, as the probability of a specific type in a given zip code could be zero.

[6]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.

distribution of its zip code would mask this correlation, underscoring the value of our method in achieving more accurate predictions.

Main Findings We apply our proposed method to assess the distributional implications of transitioning from time-invariant electricity prices to Real-Time Pricing (RTP) and Time-of-Use (ToU) pricing in Spain. The results indicate that the switch to RTP has very modest – almost negligible – distributional effects, at least within the time frame covered by our analysis. In contrast, the adoption of ToU pricing yields progressive distributional outcomes. Specifically, the shift reduces electricity bills for low-income households by an average of 0.6 €/month, while increasing bills by approximately 0.9 €/month for households in the highest income quintile.[7] Interestingly, the distributional effects of ToU would remain hidden if we did not account for within-zip-code income heterogeneity.

Even though the bill impacts are modest, RTP can pose challenges for lower-income households facing tight monthly budgets (Jack and Smith (2020); Berkouwer and Dean (2022)). Our results show that bill volatility increases for all income groups, with the largest bill volatility affecting the lowest quintile. In contrast, ToU mitigates bill volatility, particularly for the low-income group.

Differences in household consumption patterns explain these findings. High-income households tend to use disproportionately more electricity during peak hours, leading to relatively larger bill increases under both RTP and ToU tariffs, as electricity prices are higher during those hours. In contrast, low-income households consume more electricity during the winter months. While this seasonal pattern has no impact on the distributional effects of ToU – since ToU rates are fixed across months – it does affect the distributional outcomes under RTP, where wholesale prices vary seasonally. However, this seasonal impact is fully offset by within-month price variation, which explains the minimal distributional effects of the shift to RTP.

We identify two primary channels driving household consumption patterns: heating, ventilation, and air conditioning (HVAC) usage and household location. Electric heating (EH) and air conditioning (AC), which together account for nearly 30% of a typical household's annual consumption, vary significantly across regions based on local climate and gas infrastructure availability. Additionally, EH is negatively correlated with income, while AC is positively correlated.[8] Given that electricity prices are substantially higher in winter and lower in summer, the reliance on electric heating (EH) by low-income households and on air conditioning (AC) by high-income households explains the regressive effects of exposing households to monthly price variation under RTP (but not under ToU). This seasonal component offsets the progressive effects of within-month price variation, which are present in both RTP and ToU.

Arguably, these findings are location-specific, as they depend on customer equipment, demand patterns, and the nature of electricity supply, all of which vary across countries and over time.

[7] In percentage terms, these figures correspond to a 1.2% decrease in bills for low-income households and a 1.5% increase for high-income households.

[8] In Spain, older buildings often lack formal heating systems, and inefficient electric heaters are used, contributing to the negative income correlation.

Beyond assessing the distributional implications of RTP and ToU in the Spanish context, our contribution lies in disentangling the underlying channels and proposing a novel methodology to quantify them.

The structure of our paper is as follows. We next discuss the related literature. Section 3 describes the data that we use to estimate household income. Section 4 describes the methodology used to infer individual household income and the income estimation results. Section 5 provides background of the Spanish time-varying electricity pricing and quantifies the distributional implications of RTP and ToU, and Section 6 explores the channels. Last, Section 7 concludes.

2 Related Literature

Our paper contributes to two strands of the literature. First, it contributes to the methodological literature on inference and clustering by proposing a novel approach for inferring household-level income using consumption data and zip code-level information. Second, it advances the study of alternative electricity pricing policies by examining their distributional implications across different income groups.

Income inference and methodology Our method to improve income predictions relates to the ecological inference problem, which precisely focuses on the challenge of drawing conclusions about individual effects from aggregate data, known as the ecological fallacy (Robinson, 1950). The lack of individual-level data hinders our ability to understand the distributional implications of many policies, as highlighted in Banzhaf et al. (2019). A first strand of this literature tries to find bounds to individual effects while remaining flexible, e.g., as in the seminal paper of Goodman (1953). This approach has been extended to electricity pricing, under the bounding assumption that richer households consume more (versus randomly within a group), with applications to non-linear pricing (Borenstein, 2012) and solar net-metering tariffs (Borenstein, 2017). While this method is intuitive in a single-dimensional setting (e.g., total quantity consumed needed to compute monthly non-linear prices), it is less obvious how to make bounding assumptions over the entire consumption vector of households (e.g., quantity consumed over time), which is needed to estimate the impacts of real-time pricing.

Our paper falls into the second strand of the ecological inference literature, which provides a statistical framework to probabilistically infer individual income. King (1997) presents seminal work with a two-by-two problem, inferring non-black and black voting behavior for Democrats versus Republicans, using Bayesian methods. Like our framework, the estimation relies on the observation of several precincts (in our case, zip codes). The original work is limited to categories, although it has been extended, e.g., by Greiner and Quinn (2009). Our method extends this literature by predicting income that is also based on individual-specific covariates. By using constrained GMM approaches and semi-parametric logit-style functions, it is also computationally very tractable.

Methodologically, we combine methods of clustering observations with ideas from industrial organization to back out individual primitives from outcome variables (in our case, income), as in demand estimation models.

The first step of our approach relates to the clustering literature. While most finite mixture models assume parametric (often normal) distributions (McLachlan et al. (2019)), Bonhomme et al. (2016) develop a nonparametric method using repeated data for identification.[9] Closer to our setting, Bonhomme and Manresa (2015) introduce a grouped fixed-effects (GFE) estimator to capture discrete, time-varying heterogeneity via clustering. Bonhomme et al. (2022) extend this idea with a two-step GFE estimator: first clustering with *kmeans*, then estimating group fixed effects. They show that even when true heterogeneity is continuous, treating it as discrete serves as an effective dimension reduction strategy. Our first step is similar, as we use individual moments for clustering to reduce the dimensionality of the data. However, in the second step, instead of estimating a regression model, we recover group-specific income distributions by matching them to the aggregate income distribution at the zip code level.

Our approach is also related to the demand estimation literature using partial microdata. Like Berry et al. (2004), we observe detailed consumption behavior, but we lack individual income data and therefore cannot use micro-moments. While Berry et al. (1995) show how to use market-level demographics to identify demand systems, directly estimating parametric electricity consumption demands is difficult here due to the computational burden of inverting high-frequency consumption data and the complex, heterogeneous link between income and electricity use.Instead, our goal here is to infer their income, rather than characterize their entire demand system. To address this, we use a simplified two-step estimator. Like the fixed-grid approaches in Fox et al. (2011) and Bajari et al. (2007), we discretize household types to reduce complexity. This transforms the second step into a constrained GMM problem. Our method allows for rich group-level heterogeneity without imposing a specific functional form between income and consumption.

Electricity pricing We contribute to the literature assessing the imapct of time-varying pricing. Time-varying pricing improves efficiency by encouraging conservation and load shifting during expensive hours (Jessoe and Rapson (2014), Burger et al. (2019), Faruqui et al. (2009), Wolak (2011), Allcott (2011)), while also supporting investment efficiency (Borenstein, 2005), limiting market power (Poletti and Wright, 2020), and may even generate positive environmental impacts (Holland and Mansur, 2008).[10] Still, real-time pricing (RTP) has been largely confined to industrial customers (Blonz, 2022).

The literature finds a variety of distributional impacts. Borenstein (2013) finds CPP has minimal effects on low-income households, while Faruqui et al. (2010) suggest they may benefit due to flat usage and higher responsiveness. For RTP, Horowitz and Lave (2014) find that smaller, often

[9]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 also used machine learning models to classify electricity load curves but, as far as we are aware, it has not used this approach to infer household income.

[10]Schittekatte et al. (2024) compare ToU and CPP pricing under high renewables, recommending a combined approach.

low-income households pay more, while larger ones save. Leslie et al. (2024) find RTP benefits areas with more renters, older residents, and lower home values. Borenstein (2007) shows that many industrial customers would lose under RTP unless highly price-responsive.

While our findings align with prior research, we emphasize that the effects on low-income households vary significantly based on factors such as the type of time-invariant pricing (monthly vs. annual), consumption patterns, HVAC systems, and geographical location. By identifying these channels, our analysis offers insights that can help other jurisdictions anticipate and mitigate potential concerns about the distributional impacts of RTP before its implementation.

3 Data

The focus of our application is to quantify the distributional effects of alternative electricity pricing rules – specifically, Real-Time Pricing (RTP) and Time-of-Use (ToU) – relative to time-invariant pricing. In this section, we describe the data used to infer household income, which includes hourly electricity consumption, contracted power, and demographic characteristics. We also provide background on the relevant pricing mechanisms within the Spanish electricity market.

3.1 Hourly Electricity Consumption

Our dataset includes information on over three million Spanish households, covering the period from January 1, 2016, to May 31, 2017.[11] They are distributed across various regions of the country, but most of them are located in Madrid and Galicia.[12] Importantly, the data come from a distribution company, which allows us to observe all households in a given geography, rather than a selected sample of customers.

After filtering out outliers (specifically, households with excessive zero consumption observations or missing zip code data)[13] and excluding households outside the regulated utility's service areas,[14] we retain a final sample of 1,303,350 households spread across 750 zip codes. Due to missing data, we also exclude observations from December 2016 and May 2017, leaving us with 15 months of data (January to November 2016, and January to April 2017).[15]

Additionally, the dataset provides information on each household's access tariff, contracted power, and postal code. Contracted power, measured in kW, is the maximum consumption allowed at any point in time.[16] Since households pay a fixed monthly fee as a function of their contracted power, which can be between 20 and 35% of their bill on average, they have incentives to contract

[11] The data were provided by Naturgy, the third largest utility company in Spain.

[12] The geographic distribution of households is illustrated in the Appendix in Figure A.2

[13] The outlier removal algorithm excludes a household if more than 25% of its consumption observations are zero, or if more than 5% of observations are null.

[14] The default provider in each region is responsible for offering the default RTP tariff, so households outside a utility's regulated territory cannot participate in the RTP scheme.

[15] The smart meter data lacks nearly all consumption data for December 2016, and is very incomplete for May 2017.

[16] If households go over the limit, the meter trips and households lose power until they go below the limit.

it according to their electricity needs and their willingness to pay to avoid being tripped if they go over the limit. For the purposes of our estimation, this is an important variable, as it is observed for each household, and it tends to be positively correlated with income. Households with high income tend to live in larger houses and have a higher willingness to pay to avoid the inconvenience of power limits.

3.2 Demographic Data

We obtain demographic data from two sources: the Spanish National Institute of Statistics (INE) and a private data provider, MB Research. The INE provides demographic information at the census district level, including population, age, gender, education, dwelling types, and income distribution data.[17] In contrast, MB Research offers income distribution data specifically for zip codes.[18] We complement our data with individual-level energy consumption survey data (CEX), which includes individual household income decile and region.[19]

The consumption survey data show that higher-income households tend to report higher electricity consumption (see Appendix A.4). However, these basic patterns are not reflected when using only zip code-level income information, raising concerns about the informativeness of relying solely on zip codes. This issue is particularly pronounced in Spain, where zip codes still contain substantial heterogeneity, and is the motivation for our proposed approach to better infer income.[20]

4 Inferring the Household-Level Distribution of Income

To better understand the distributional implications of RTP and ToU, we propose a two-step GMM approach to improve our estimates of household income distributions.

Let us assume that household 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 for their zip code (denoted x_{ih}) and lifestyle (represented by their type θ_i), plus some random shocks ϵ_{ih},

$$\text{kWh}_{ih} = f(x_{ih}, \epsilon_{ih} | \theta_i). \tag{1}$$

The proposed GMM methodology follows two steps. In the first step, we classify households into different types based on their contracted power, 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 for each zip code.

[17] Since we have each household's zip code but not its census district, we match census districts to postal codes and aggregate the data at the postal code level.

[18] Appendix A contains a more detailed description of these data sources.

[19] See Appendix A.4 for a detailed description of the data.

[20] Moreover, the problem is exacerbated by the fact that we are focusing on a single utility, rather than the entire Spanish market.

Allowing the household's discrete type θ_i to be correlated with its income helps us to identify how income correlates with electricity consumption while reducing the dimensionality of the income distribution that we need to estimate. For identification purposes, we allow these types to be shared across similar geographies.

In the second step, we assume that each type has a fixed distribution of income, which is unknown but could be a function of covariates. 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, $\theta \in \Theta = \{\theta_1, \ldots, \theta_N\}$, which can potentially also be a function of covariates. To define the income distribution, we partition the income domain into K bins, $inc_k \in \{1, .., K\}$, using national income quintiles ($K = 5$). For the case without covariates, let $\eta_k^n = Pr(inc_k|\theta_n)$ denote the discrete probability of household type θ_n belonging to quintile k. 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. This gives us an expected income distribution. In practice, we also allow these probabilities to depend on observables.

Next, we 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

As already explained, households pay a fixed monthly fee based on contracted power, which is strongly correlated with income. Contracted power can vary from 1 to 10 kW (with 0.1 increments), but most households in our sample chose 2.5-5 kW. Figure A.1 depicts its distribution. 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 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.[21]

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 perform 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. 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 consumption share for each of the 24 hours by weekday and weekend;[22]
- four variables capturing seasonal patterns in consumption: the ratio of winter consumption to annual consumption, the ratio of summer consumption to annual consumption, the standard deviation of monthly consumption, and the 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 four variables reflect seasonality across months. The former mainly depend on the household's lifestyle, while the latter are greatly affected by HVAC ownership.

4.1.4 Final classification

We first classify households according to their individually observable contracted power and their individually inferred ownership of electric heating: (L, EH), $(L, NoEH)$, (H, EH), and $(H, NoEH)$.

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

[22] Hourly consumption share is defined as the hourly consumption divided by the total consumption for the day.

Figure 1: An example of *kmeans* types in Coruña with low contracted power and no electric heating

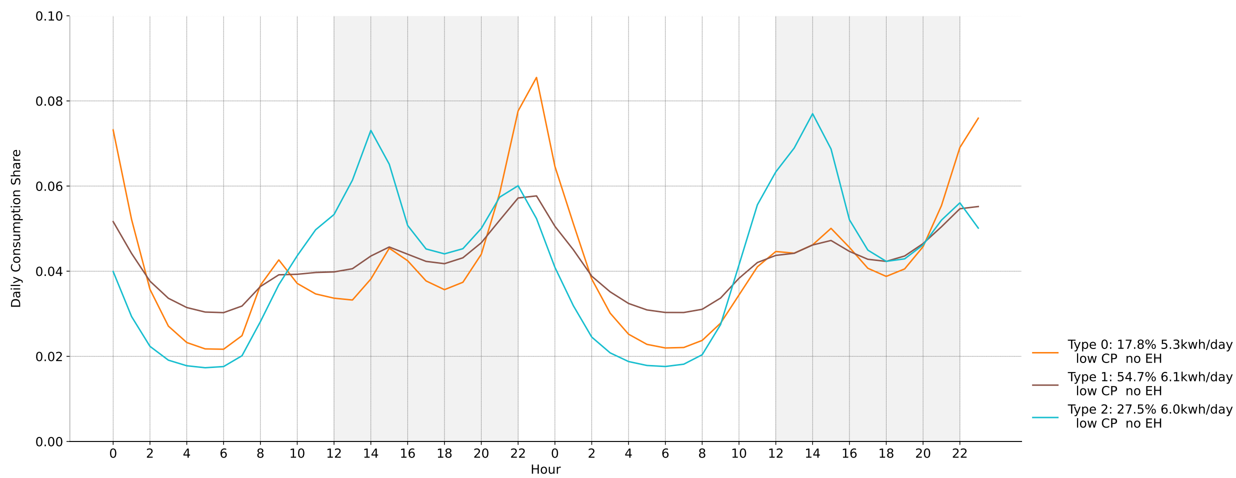

Notes: This figure provides an example of the *kmeans* classification of households in Coruña with low contracted power and no electric heating. The three clusters group households according to their electricity consumption profiles throughout the day (in shares). The first 24 hours are for weekdays, and the last 24 hours are for weekends.

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 for which we have a sufficiently large number of households.[23]

Figure 1 illustrates an example of *kmeans* classification. It shows the average daily consumption patterns for weekdays and weekends of households with low contracted power and no electric heating in Coruña. One can see that the algorithm picks up a variety of consumption patterns: households that consume in the evening (type 0), households that consume slightly more at night (type 1), and households that consume mostly during lunchtime and in the evening (type 2).

4.2 Step 2: Identifying the Income Distribution of each Type

From step 1, we get a type 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 $N^g = 12$ for all provinces, with 3 *kmeans* types within each contracted power-EH category. We estimate the types and income distribution for each province separately. From now on, we suppress the superscript g for clarity.

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

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

$$P^j(\theta_n) = \frac{1}{HH_j} \sum_i \mathbb{1}(\theta^i = \theta_n), \tag{2}$$

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

Conceptual 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:

$$\min_{\eta} \quad \sum_j \omega_j \sum_{k=1}^{K} (Pr_k^j - \sum_{\theta_n \in \Theta} \eta_k^n P^j(\theta_n))^2, \tag{3}$$

$$s.t. \quad \sum_{k=1}^{K} \eta_k^n = 1 \qquad \forall \theta_n \in \Theta, \tag{4}$$

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 (3) 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 have only 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.

Main semi-parametric estimator We implement a semi-parametric estimator that allows the income distribution of types to exhibit differences across individuals and zip codes. It has the advantage of allowing individuals classified into the same type to have distinct income distributions, which otherwise could be too strong of an assumption with a limited number of types.

More concretely, we specify that the probability that a household of type θ_i belongs to the income bin k depends on individual characteristics (x_i) and zip code demographics (z_j). This makes the method computationally more intensive, as we need to keep track of income probabilities at the individual level, rather than at the type level.

We use the following moment conditions by matching the observed and predicted zip code income distributions as above, but integrating over all households rather than using the aggregate shares of types:

$$\min_{\eta,\alpha,\beta} \quad \sum_j \omega_j \sum_{k=1}^{K} (Pr_k^j - \sum_{i\in\mathcal{I}_j} Pr_k(\theta_i, x_i, z_j))^2, \tag{5}$$

$$s.t. \quad Pr_k(\theta_i, x_i, z_j) = \frac{exp(\delta_{ijk})}{\sum_{k'=1}^{K} exp(\delta_{ijk'})}, \qquad \forall k \in [1, \cdots, K], \tag{6}$$

$$\delta_{ijk} = \alpha_k + \beta_0^{\theta_i} \times k + \beta_1 x_i \times k + \beta_2 z_j \times k, \tag{7}$$

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 (continuous variable) and their binned monthly electricity consumption (dummy variables), and z_j represents the demographic variables of the zip code.

Unlike in our previous specification (3), 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 (6) and (7). α_k is a common income-bin dummy. $\beta_0^{\theta_i}$ are type-specific coefficients and (β_1, β_2) are the same for all types. $\beta_0^{\theta_i} \times k$ explains how type θ_i's income distribution differs from the average. If $\beta_0^{\theta_i} > 0$, it means the type is richer than average in a first-order stochastic dominance (FOSD) sense. This ordering means that one type's income distribution cannot have higher probabilities in both the richest and the poorest bins compared to another type.[24]

In our main specification, we focus on individual characteristics as they give the most flexibility to the individual distribution of income, with x_i including a household's contracted power, consumption, and peak consumption, as well as the slope estimates from the HVAC inference step to allow for additional flexibility on heating and cooling behavior.[25] Including these components allows the household-level income distributions to be more flexible within and across zip codes.[26] $\beta_1 x_i$ captures whether higher contracted power within a type is correlated with higher income ($\beta_1 > 0$). $\beta_2 z_j$ captures the correlation of characteristics and income. We would expect lower socio-demographics in the zip code to be negatively correlated with the distribution of income,

[24] The model can also accommodate a quadratic specification, which would relax this parameterization, yielding similar results.

[25] Computationally, we group these characteristics into 15 bins each to make the integration of zip code probabilities faster. We include a range of other controls in the appendix. The results are similar between the non-parametric approach and various alternative parameterizations of the semi-parametric approach.

[26] Consider the following two households: Household A has low contracted power and belongs to a "richer" type, while household B has higher contracted power and belongs to a "poorer" type. Assuming $\beta_1 > 0$, which implies that households with higher contracted power are wealthier, the income distributions of households A and B would not necessarily have a FOSD relationship.

Figure 2: Validation with CEX data

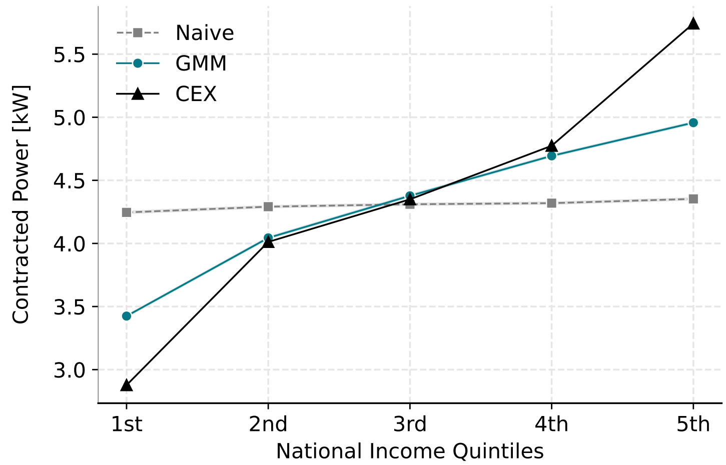

Notes: This figure represents the relationship between household income quintiles (on the x-axis) and contracted power (on the y-axis). The squared green line represents the estimated relationship using our two-step method, while the dashed gray circled line uses only the zip code income distribution. To validate our estimation, the triangle black line depicts the relationship from the CEX survey. Although the CEX survey only reports total annual bills and annual electricity consumed, we infer the contracted power using information about regulated contracted power prices (in euros per KW), taxes (in percent), other fees (in euros) and average annual energy component prices (in euros per KWh) provided by the CNMC (Spanish Competition Authority and Regulator).

keeping 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 being different between zip codes. Furthermore, two households within the same zip code with the same type can have different income distributions due to the effect of β_1.

The final outcomes of interest from the approach are the probabilistic assignments of households to income quintiles, given by the predictions $\eta_{ik} \equiv \hat{P}r_k(\theta_i, x_i, z_j)$.

4.3 Validation of inferred income

Using the estimated income distribution for each type, we calculate the implied income distribution for each household. Our method aims to better infer the expected income distribution of households. To understand the added value and performance of our estimator in small samples, we performed three checks on our method: validation from CEX data, a Monte Carlo simulation, and out-of-sample validation.

First, we use the CEX data to assess the validity of our estimates. Figure 2 compares the relationship between contracted power and income, which are known to be highly correlated. The dashed line shows weak correlation using aggregate zip code demographics, with an average contracted power of around 4 kW across quintiles. In contrast, our estimated income distribution shows a stronger correlation: households in the low quintile average less than 3.5 kW in contracted power, while those in the high quintile average 5 kW. Using the CEX data supports an even steeper relationship, as shown by the circled black line.

We use the same CEX data to compare the relationship between annual electricity consumption and income quintile in Appendix C.1. In Figure C.1, we compare the estimated results with 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.[27]

Second, in Appendix C.2, we perform a Monte Carlo simulation in which we assume that we know each 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. Overall, the Monte Carlo simulation helps highlight the value of our approach. With enough flexibility, we can reveal 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 expected household income.

Finally, in Appendix C.3, 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, with positive results. This is in line with the results from our Monte Carlo simulation.

Our methodology does not allow for precise identification of individual household income, but rather estimates its expected distribution. As such, it should not be viewed as a full substitute for micro-level data, but rather as a refinement over aggregated data sources. The limitations of this approach – and their implications for estimating policy impacts – are discussed in greater detail in the following section.

5 The Impacts of Alternative Electricity Pricing Schemes

Our goal is to identify the winners and losers resulting from the shift from time-invariant electricity pricing schemes to RTP and ToU, using our estimated income distribution at the household level. We analyze the distributional impacts along two dimensions – across income groups and within income groups – based on the changes in monthly electricity bills that are based in changes that the Spanish electricity market underwent.

5.1 Spanish Context

The Spanish electricity market offers a unique natural experiment to assess the distributional effects of RTP and ToU, as the default household tariff includes elements of both. Specifically, electricity bills consist of three elements: a fixed charge based on contracted capacity, a volumetric charge

[27] An exception to the better fit is Madrid. Our utility data only cover select parts of the city and region, and thus the household sub-sample is not as comparable to the CEX region sub-sample. This is an important reason why we do not use the CEX moments explicitly in the estimation, but rather use them for validation purposes. The other reason is that the regions in the CEX survey are at the state level, rather than the province, limiting our number of regions further.

based on consumption, and taxes.[28] Schematically, a household's bill can be written as:[29]

$$Bill_i = \left[p^k k_i + \sum_{h,d,m} p^e_{hdm}\, kWh_{i,hdm} \right] (1+\tau),$$

where p^k is the regulated per-kW capacity price and k_i is the household's contracted capacity; p^e_{hdm} is the volumetric energy price faced in hour h, day d, and month m, and $kWh_{i,hdm}$ is the household's electricity consumption; τ is the energy tax rate.

Capacity charge. The first element, the fixed capacity charge, requires households to pay a regulated price for their contracted power, which reflects the maximum permissible demand at any moment. This contracted power is chosen annually and it represents a substantial share of the energy bill, about one-third of the monthly bill on average.[30] In addition, contracted capacity increases with house size and installed appliances. Thus, it tends to be highly correlated with income, as already shown in Figure 2.

Volumetric charge. The second element, the volumetric charge, depends on the energy price and the household's consumption. The retail energy price consists of two components:

(i) A regulated access charge, which recovers transmission, distribution, and policy costs. During our sample period, it was time-invariant by default, denoted $\overline{f}$, though households could opt into ToU prices, denoted f_{hd}, with cheaper off-peak rates (e.g., nights and weekends) but no monthly variation. Only 14% of households in our data adopt ToU. The access charge represented around 40% of the total energy charge during our sample period (4.4 cents Euro/kWh). A 2021 reform subsequently made ToU tariffs mandatory and more extreme. The new structure divides the day into three periods: peak (10:00–14:00 and 18:00–22:00 on weekdays), shoulder (8:00–10:00, 14:00–18:00, and 22:00–00:00 on weekdays), and off-peak (all other hours, including weekends and nights). These fees differ substantially: shoulder and off-peak tariffs amount to 30% and 3% of the peak tariff, respectively.[31] After this policy change, the access charge during peak times was significant and larger than the average wholesale price.[32]

(ii) The market component of the volumetric price, which is, by default, a direct passthrough of the real-time wholesale prices, p_{hdm}, which are geographically uniform. Although households may opt out of real-time pricing (RTP) by contracting with competitive suppliers—most of whom offered time-invariant tariffs $\overline{p}$ during our study period—RTP remained widespread due to substantial

[28] This tariff structure applies only to households with peak demand below 10 kW. Those exceeding this threshold must contract with a competitive retailer, which typically offers time-invariant rates. Households must also have a smart meter installed. By the end of 2015, nearly 12 million smart meters had been installed in Spain, with approximately 10.19 million successfully integrated into electricity suppliers' information and telecommunication systems. By 2018, all Spanish households (28.02 million) had smart meters installed.

[29] Details on the construction of annual bills follow the weighting procedure described in the original manuscript.

[30] The price is 42 euros per kW of contracted power annually during our sample period.

[31] When considering the switch to ToU, we rely on these tariffs, normalized to ensure revenue neutrality.

[32] See Enrich et al. (2024) for further details on the before and after of the TOU policy change.

Figure 3: Price fluctuations over time (real-time, monthly, and annual prices)

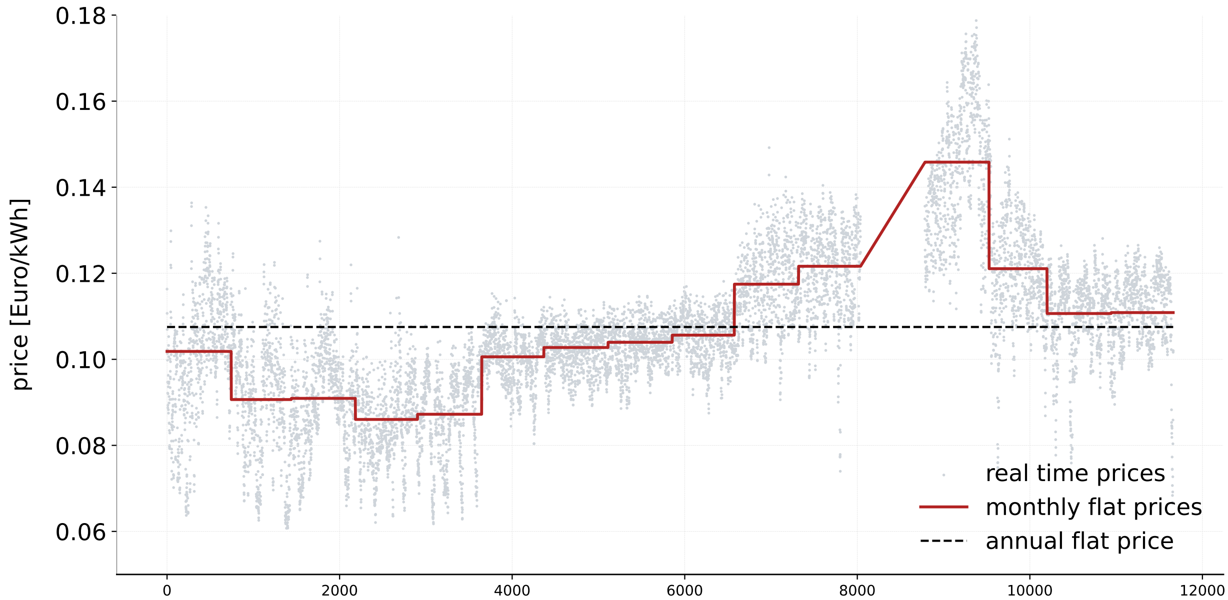

inertia in retail choice (Fowlie et al., 2021; Hortaçsu et al., 2017; Enrich et al., 2022). Figure 3 illustrates the evolution of real-time electricity prices over our sample period. Daily, monthly, and seasonal fluctuations driven by underlying wholesale conditions are passed through one-to-one into retail prices. While within-day variation is present, it is relatively modest compared to experimental settings where peak prices increase by 200–600% (Harding and Sexton, 2017), resulting in limited evidence of demand response (Fabra et al., 2021). By contrast, the monthly variation is considerably larger.

All households also pay a uniform 21% VAT and a 4.864% special tax on electricity consumption.

5.2 Counterfactual bills

To isolate the role of time variation in each volumetric component, we use smart-meter data and observed prices to compute three counterfactual electricity bills: a Real-Time Pricing (RTP) bill with real-time wholesale prices and non-ToU access charges, denoted $Bill_i^{RTP}$; a Time-of-Use (ToU) bill with ToU access charges and time-invariant prices, $Bill_i^{ToU}$; and a "flat" bill with non-ToU access fees and constant prices,$Bill_i^{FLAT}$. The time-invariant components $\overline{f}$ and $\overline{p}$ are calibrated to be revenue-neutral relative to their time-varying counterparts.

The resulting bill shocks are:[33]

$$\Delta Bill_i^{RTP} = Bill_i^{RTP} - Bill_i^{FLAT} = (1+\tau) \sum_{h,d,m} (p_{hdm} - \overline{p})\, kWh_{i,hdm},$$
$$\Delta Bill_i^{ToU} = Bill_i^{ToU} - Bill_i^{FLAT} = (1+\tau) \sum_{h,d,m} (f_{hd} - \overline{f})\, kWh_{i,hdm}.$$

To distinguish hourly from monthly sources of variation under RTP, we also construct a monthly time-invariant price that is revenue-neutral within each month, $\overline{p}_m$. The corresponding bill is denoted $Bill_i^{MONTH}$. This allows us to decompose the RTP effect into a within-month component driven by hourly price variation and an across-month component driven by monthly variation:

$$\Delta Bill_i^{RTP} = (Bill_i^{RTP} - Bill_i^{MONTH}) + (Bill_i^{MONTH} - Bill_i^{FLAT}),$$

which can be written as

$$\Delta Bill_i^{RTP} = (1+\tau) \sum_{h,d,m} \left[(p_{hdm} - \overline{p}_m) + (\overline{p}_m - \overline{p})\right] kWh_{i,hdm}.$$

The across-month component arises only under RTP because ToU tariffs do not vary across months. Therefore, bill changes under ToU depend exclusively on within-month variation.

5.3 Policy shocks along income

We start by analyzing the heterogeneity in bill impacts across income groups. Tables 1 and 2 and Figure 4 classify households into five national income quintiles and report the bill impacts following a switch from a time-invariant annual price to RTP and ToU.

The analysis delivers two main takeaways. First, the transition to real-time pricing appears nearly distributionally neutral, with minimal impact on monthly bills across income groups.[34] In contrast, the switch to ToU pricing exhibits markedly stronger progressive distributional effects. The magnitude of these effects varies substantially depending on the method used to measure income: our estimated household-level income distribution based on the GMM approach yields stronger results than the naïve alternative of relying on zip code-level income data.

Table 1 presents average monthly household electricity bills under the alternative pricing schemes, both for the overall population and across the five income quintiles, with households classified using our GMM-based approach. Variation in household electricity bills arises from differences in both consumption levels and temporal usage patterns, and these effects differ across pricing schemes. On average, households consume approximately 251 kWh of electricity per month, with

[33] The fixed-capacity charge is identical across scenarios and therefore does not affect differences. Taxes scale the differences proportionally without changing signs.

[34] Due to the absence of December data, we may underestimate potential regressive effects of RTP, as lower-income households are more likely to rely on electric heating. We discuss the distributional impact channels in Section 6, and specifically the HVAC ownership channel in Section 6.2.

Table 1: Monthly bills under alternative pricing schemes

	Mean	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
KWh_i	251.4	229.3	246.0	253.0	260.0	268.3
$Bill_i^{RTP}$	55.4	48.4	53.5	55.9	58.2	60.6
$Bill_i^{ToU}$	55.4	47.8	53.0	55.7	58.6	61.5
$Bill_i^{FLAT}$	55.4	48.4	53.5	55.9	58.3	60.6
% losers (RTP)	36.8	36.8	36.6	36.6	36.7	37.4
% losers (TOU)	45.6	38.7	41.5	44.4	49.0	54.4

Notes: This table reports household-level bills for the prices observed in our sample (RTP), for an alternative ToU pricing, and for the flat alternative. There are 1,303,350 households. All units are measured in €/month, except for KWh_i, which is measured in kWh/month and the percent of losers. Average monthly bills include energy, taxes, and other components of the bill such as contracted power fees. The % of losers refers to those households whose bill increase after the switch from the time-invariant tariff. The quintile categories are obtained from the procedure described in Section 4, and used to compute the weighted averages by quintile, reported in columns.

high-income households consuming more on average than low-income households, and thus paying higher monthly bills.

Differences in consumption patterns do not affect the relative bills paid under RTP versus its time-variant alternative, which remain proportional to average consumption. Additionally, the share of households that experience an increase in their electricity bills upon switching from flat-rate pricing to RTP remains relatively stable around 37% across the income distribution, with only a slight increase from 36.6% for the second and third quintiles to 37.4% for the fifth quintile.

In contrast, the transition from flat rates to ToU pricing induces more pronounced distributional impacts. Specifically, the magnitude of bill changes is more substantial: households in the lowest income quintile save 0.6€/month, whereas those in the highest quintile incur an average increase of 0.9€/month. Moreover, the proportion of households with higher bills under ToU is significantly higher in the top income quintile (54.4%) compared to the bottom quintile (38.7%).

Figure 4 plots the predicted bill impacts across income groups, showing that the predictions vary significantly depending on whether we use our estimated household-level income distribution (as in Table 1) or a zip code-level income distribution (the naïve approach). Under our proposed method, the shift to Time-of-Use (ToU) pricing has pronounced distributional effects. Ignoring within–zip code income heterogeneity obscures these progressive effects: the naïve approach yields predicted bill impacts that are nearly flat across income groups. The intuition behind this is that the naïve approach can only capture the impact correlated with geographical factors, missing the impacts that are explained by household types.[35]

To explore the channels, Figure 5 decomposes the bill impacts of the switch to RTP into within-month and across-month components. Panel (a), which relies on our estimated household income

[35] In Appendix C.2.2, we use Monte Carlo simulations to evaluate the extent to which the two-step approach and the naïve approach capture the true impact. The simulation results in Figure C.5 roughly replicate our main policy findings: the two-step procedure provides strong progressive results that align with the true impact, while the naïve approach yields slightly progressive results.

Figure 4: Bill changes due to the switch of tariffs [Euro/Month]

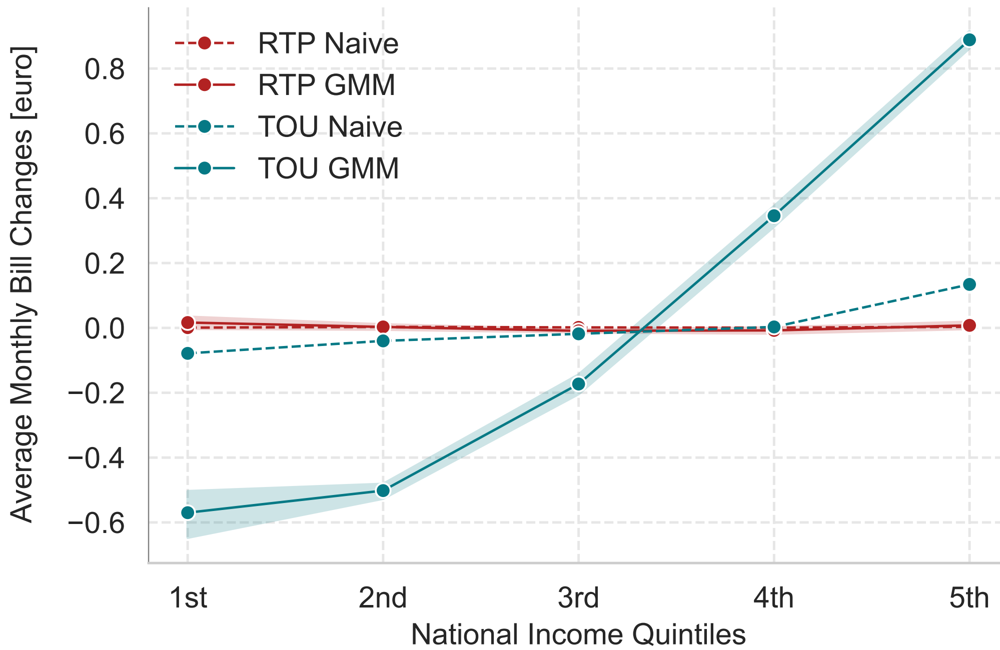

Notes: This figure represents the monthly bill increase when moving from an annual time-invariant price to RTP or ToU. Results are reported for the five national income quintiles, with household income classified according to our estimated income (GMM) or to the zip code income (naïve). The shaded areas represent the confidence intervals that account for uncertainty in our estimated measure of income.

Table 2: Average bill and volatility changes from the switch to RTP

	RTP Bills		ToU Bills		Flat Bills	
	GMM	Naïve	GMM	Naïve	GMM	Naïve
Quintile 1	0.36	0.34	0.30	0.28	0.33	0.29
Quintile 2	0.35	0.33	0.29	0.28	0.31	0.29
Quintile 3	0.33	0.33	0.27	0.27	0.29	0.29
Quintile 4	0.31	0.33	0.26	0.27	0.27	0.29
Quintile 5	0.30	0.32	0.24	0.27	0.25	0.28

Notes: The table reports the average bill changes from the switch to RTP. The first two columns report monthly bill volatility changes. For each household, we calculate the standard deviation of its monthly bills divided by the household's average monthly bill. In this table, we report the average volatility of all households by inferred income quintiles. The middle two columns report the estimated percentage of losers by income quintiles from the GMM and the naïve approaches. The last two columns report average bill changes for losers and winners separately.

Figure 5: Decomposition of the distributional impact of RTP

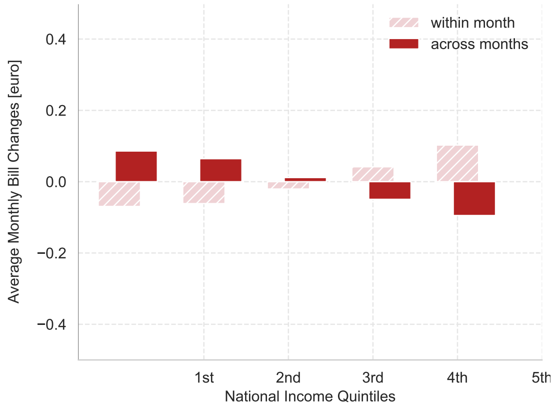

(a) GMM approach

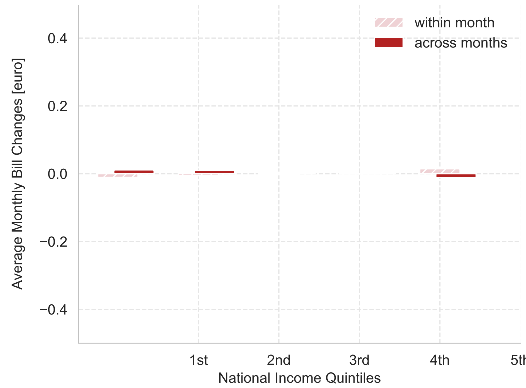

(b) Naïve approach

Notes: These figures decompose the bill change when moving from an annual time-invariant price to monthly prices (pink hashed 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). Note that these figures represent the national average, which hides the heterogeneity in the bill impacts across regions.

distribution, reveals stronger distributional patterns for each source of price variation compared to Panel (b), which is based on the naïve approach using zip code-level income data. Once again, relying on aggregated zip code-level information tends to obscure the underlying distributional effects.

Interestingly, the progressivity of the within-month component is nearly offset by the regressivity of the across-month component, resulting in an overall distributionally neutral effect reported above. However, this decomposition reveals an important nuance: bill volatility increases under RTP, particularly for lower-income households.

Indeed, one concern with time-varying electricity pricing is that lower-income households may struggle to cope with increased month-to-month bill variability, given that they often face binding monthly budget constraints that align with the frequency of their income. The literature on credit constraints in developing countries (*e.g.*, Jack and Smith (2020); Berkouwer and Dean (2022), among many others) suggests that such liquidity constraints can impose significant opportunity costs.

Table 2 investigates this issue by reporting bill volatility under each pricing regime (RTP, ToU, and time-invariant pricing) for all income quintiles, using both the GMM-based and naïve income classification approaches. While the naïve approach indicates only modest differences in bill volatility across the income distribution, these disparities become more pronounced when household-level income is estimated using the GMM approach.

Our results confirm the above concern: bill volatility increases for all income quintiles under RTP relative to time-invariant pricing, with the first (lowest) income quintile experiencing the largest bill volatility. Specifically, bill volatility for the first quintile is 0.36 under RTP and 0.33 under flat bills. This heightened volatility among lower-income households is largely driven by their greater reliance on electric heating, as discussed in Section 5. In contrast, ToU reduces bill volatility relative to time-invariant pricing across all income quintiles, but particularly so for the low-income group.

5.4 Robustness

Our method is subject to several researcher choices that can impact the results. In Appendix D, we provide estimates of Figure 4 under alternative specifications. Figure D.1 suggests that the nonparametric and the semiparametric approaches give similar results. Our results are also robust to different choices of semi-parametric specifications. The different specifications consider alternative numbers of clustering types, the inclusion of observable or inferred characteristics (contracted power and electric heating), and a range of controls. We find that our estimator is robust to these modifications, also at the regional level, as shown in Figure D.2.

5.5 Methodology discussion

In addition to our robustness tests, we explore how to guide the researcher's choices in other settings. Based on the econometric assumptions and Monte Carlo results, we suggest that the choice of consumer types adheres to three principles: (i) the types and explanatory variables should correlate with consumer income and have shared common support across zip codes, (ii) the types should correlate with the policy impact of interest, and (iii) conditional on types, the remaining unobserved income should not directly impact the effects of the policy.

The first principle implies that the income distribution, conditional on types, should either be identical across zip codes (the assumption for our non-parametric approach) or that any relevant differences can be explained by observable factors (the assumption of our semi-parametric approach). The goal is that the remaining unexplained income distribution is uncorrelated with the policy effects (principle (iii)).

The second principle calls for bundling households along the dimensions that are most important for the policy at hand, which also makes principle (iii) more likely to hold. For example, for a policy about non-linear pricing, identifying types based on total monthly consumption would be natural. Rather than assuming that income is fully sorted along consumption to bound outcomes, as in Borenstein (2012), the methodology allows for different income distributions in each consumption bin. It also does not require to assume a particular relationship between the two.

To ensure that the third principle holds, one natural step derving from principle (ii) is to classify consumers based on the impact of the policy itself. If the two are uncorrelated, the method will correctly fail to find an association between the two. As a caveat, in our context and given the small number of zip codes per region and the substantial randomness in policy impacts, one could worry that this would overclassify households spuriously. Instead, we classify households based on general consumption patterns, which are correlated to the bill impacts and contracted power, rather than based on final impacts, which can fluctuate substantially. However, as part of our specifications in the robustness section in Appendix D, we include bill impacts as a predictor of income, and we reassuringly find very similar results.

In Appendix C.2, we provide Monte Carlo exercises that provide further intuition for these principles. We also examine what happens when our assumptions are not fully satisfied. If the types are not accurate enough to predicted household income, then the inferred income will be less correlated with the true income (Figures C.2 and C.3), leading to attenuation. Furthermore, if the policy impact is not fully explained by consumer types, or if true income directly impacts the policy outcome due to bias in our inferred income, our approach cannot fully reveal the true distributional impact (Figure C.5). Nevertheless, it is noteworthy that our approach consistently outperforms the naïve approach. We also use the Monte Carlo exercises to parallel some our robustness checks. When we chose fewer types or classify zip code groups in a sub-optimal, the results from the two-step approach are still closer to the true data than the naïve approach.

6 Channels explaining heterogeneous shocks

This section uncovers the bill impacts of RTP and ToU along several dimensions beyond income. We focus on the differences between households regarding their consumption profiles, their HVAC status, and their locations.

As noted by Cronin et al. (2019) and Douenne (2020), there exists substantial heterogeneity in energy use within income groups. This variation may arise from limitations in measuring household income or long-term wealth, but also reflects genuine differences in consumer preferences and behavioral choices, even when income is held constant.

Panels (a) and (d) of Figure 6 illustrate the distribution of bill impacts resulting from a shift from an annual time-invariant price to RTP and ToU, respectively. These distributions are shown for the 1st, 3rd, and 5th income quintiles. Under RTP, the distributions exhibit relatively modest differences across income groups, but substantial dispersion within each quintile. Still, the maximum monthly gains or losses do not exceed 1 €.

In contrast, ToU pricing results in significantly greater variation both across and within income quintiles. Consistent with the progressive effects of ToU discussed above, the distribution of bill impacts for the 1st quintile is visibly shifted to the left compared to those of the 3rd and 5th quintiles, indicating greater savings for lower-income households. Moreover, the magnitude of gains and losses is larger under ToU, with some households experiencing gains or losses exceeding 4 €per month.

For both pricing alternatives, 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), an issue on 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-

Figure 6: Bill shock due to the switch to RTP and ToU

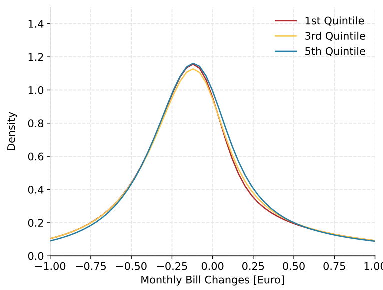

(a) RTP – All regions

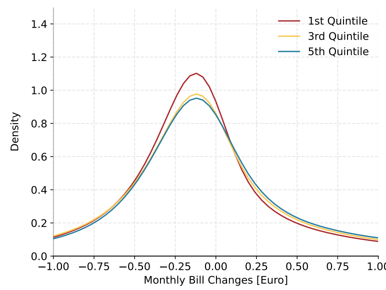

(b) RTP – Electric heating regions

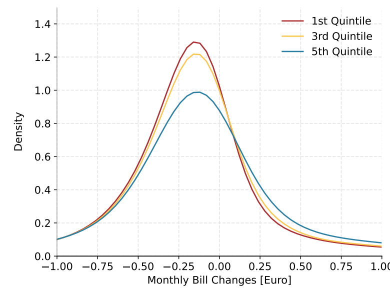

(c) RTP – Non-electric h. regions

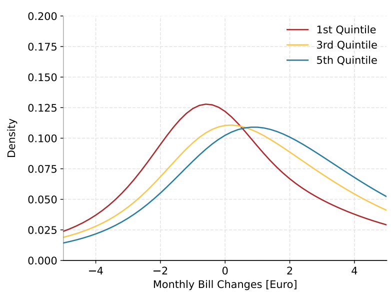

(d) ToU – All regions

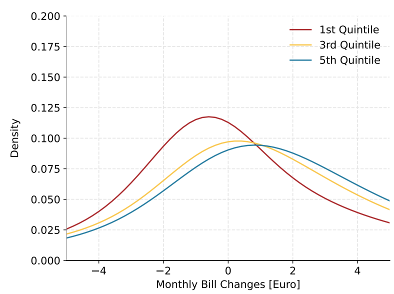

(e) ToU – Electric heating regions

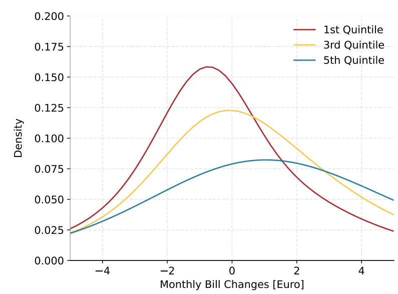

(f) ToU – Non-electric h. regions

Notes: These figures show the distribution of the bill changes due to the switch to RTP (upper panels) and ToU (lower panels) in the first, third and fifth income quintiles. Panel (a) and (d) show the distributions at the national level for RTP and ToU respectively, while Panels (b)-(e) and (c)-(f) distinguish between regions with a high and a low prevalence of electric heating, respectively. Together, they show that there is substantial heterogeneity within income groups, with the highest income group tending to have the most volatility. The differences within and across the quintiles are much more pronounced for ToU than for RTP.

Figure 7: Load curve by income quintiles

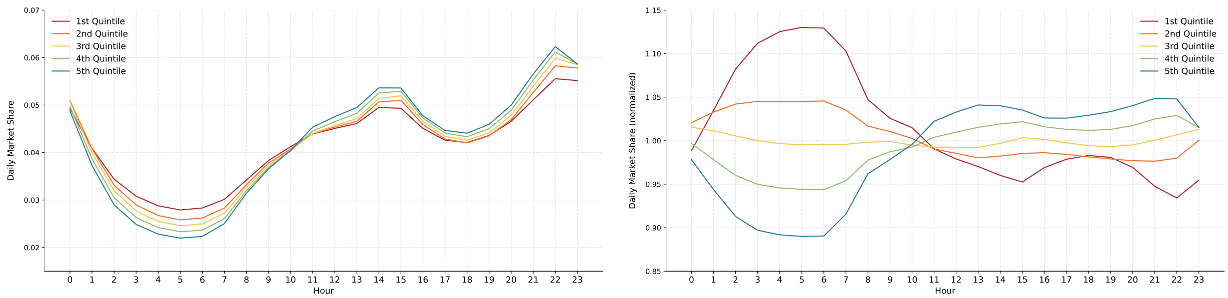

(a) Hourly consumption shares [%]

(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.

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.

6.1 Consumption Profiles

Our previous results demonstrate that transitioning from time-invariant monthly pricing to RTP or ToU has progressive distributional effects, i.e., lower-income households tend to benefit from these changes. These outcomes are primarily driven by differences in daily consumption patterns across households, as we document below.

Panel (a) in Figure 7 shows the average hourly consumption profiles for households across the five income groups. While the overall differences between income groups appear small,[36] there is notable variation in their consumption patterns throughout the day. To highlight this heterogeneity, Panel (b) plots the share of daily consumption for each hour relative to the sample average by income group.[37] Households in the higher-income group tend to consume more during peak hours within the day compared to the sample average, whereas those in the lower-income group consume relatively more during off-peak hours of the day.

6.2 HVAC status

Our previous results also show that transitioning from annual to time-invariant monthly pricing tends to disadvantage low-income households – an effect observed exclusively under RTP. This

[36] These differences would be more pronounced if Madrid were excluded, as it is the only region where high-income households tend to consume less electricity. This may be due to the widespread use of natural gas in Madrid.

[37] For instance, we calculate the share of daily consumption at noon for a given income group and compare it to the overall average.

Figure 8: Load curves by HVAC status and income

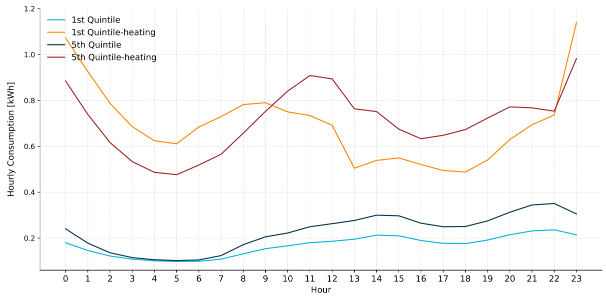

(a) Hourly consumption pattern [kWh] - Heating

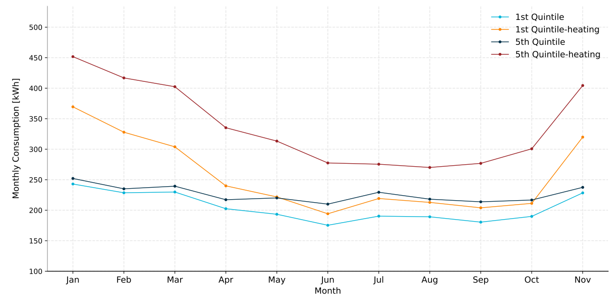

(b) Annual consumption pattern [kWh] - Heating

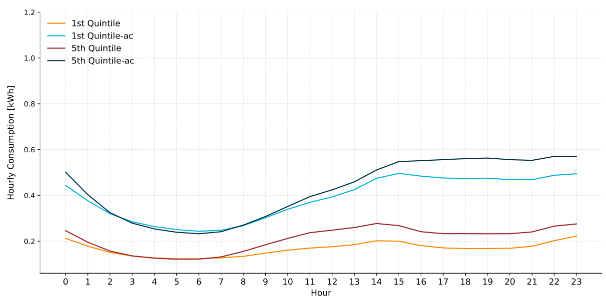

(c) Hourly consumption pattern [kWh] - AC

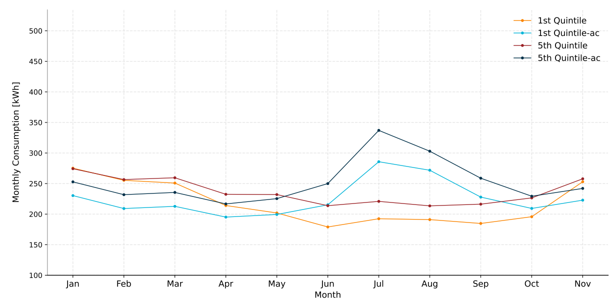

(d) Annual consumption pattern [kWh] - AC

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. Panel (a) on the left presents the hourly consumption profiles for winter months, highlighting the differences between households with and without electric heating. Similarly, panel (c) focuses on the profiles for summer months.

outcome is primarily driven by differences in seasonal consumption patterns, which in turn are largely determined by the presence of electric heating (EH) and air conditioning (AC). These appliances significantly affect both the total volume and the temporal distribution of electricity consumption, as we show next.

Panels (a) and (b) in Figure 8 illustrate the average consumption patterns of households with and without electric heating throughout the day and across the year, respectively. Panels (c) and (d) do the same for AC. As shown, there are notable differences in consumption patterns based on HVAC status. Households with electric appliances consume considerably more electricity at all hours compared to those without, and their consumption patterns tend to be peakier, especially in the case of heating.

Additionally, there are strong seasonal effects. As expected, households with electric heating consume more during the winter months (October through April), while those with AC see higher consumption during the summer months (June through September). For heating, these seasonal effects are more pronounced among high-income households compared to low-income households. In contrast, the seasonal impact of AC usage is relatively consistent across income groups.

Figure 9: Bill changes [Euro] due to RTP by electric HVAC status

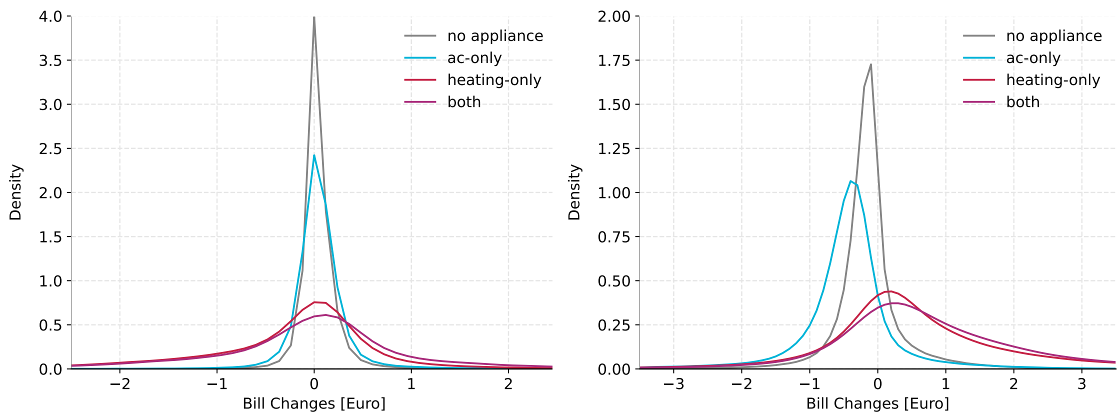

(a) Within month effects

(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.

In general, higher-income households are more likely to have AC, while lower-income households are more likely to rely on electric heating. This disparity is due to the high installation costs of alternative heating systems (e.g., gas or central heating) compared to electric heating, which typically uses low-cost plug-in radiators. Specifically, 13% of households in the 1st income quintile and 15% in the 5th quintile have AC, while 27% of the 1st quintile and 12% of the 5th quintile use electric heating.[38]

Since electricity prices tend to be higher during winter months when electric heating is in use, the shift from an annual price to RTP disproportionately impacts low-income households. The across-months effect offsets the progressive within-month effects discussed in the previous subsection. In the case of ToU, this countervailing effect is muted since ToU tariffs remain unchanged across the year.

This pattern aligns with the evidence presented in Figure 9, which decomposes the bill impacts of RTP into within-month (Panel (a)) and across-month (Panel (b)) effects, disaggregated by HVAC status. As shown in both panels, the greater volatility of winter prices –approximately ten times higher than in summer –amplifies the bill impacts for households using electric heating. Moreover, while the average within-month effect appears largely independent of HVAC status, the across-month effect is significantly influenced by it. Households with air conditioning (AC) tend to benefit under RTP, whereas those relying on electric heating experience losses, consistent with the mechanisms discussed above.

[38] These results are presented in Figure B.2a in the Appendix. For AC, the differences are even more pronounced when accounting for location. Warmer, lower-income regions in Spain have more AC use, with only minor differences within regions.

6.3 Location

Another key factor driving the distributional impacts of RTP and ToU is regional heterogeneity. Consumption patterns are closely tied to local weather conditions, influencing HVAC usage even when controlling for income. Additionally, regional disparities in the availability of heating infrastructure – particularly gas – affect the prevalence of electric heating. For example, while 90% of households in Madrid have access to gas heating systems, this figure drops to 60% in more rural areas like Galicia. Castilla y León has the lowest incidence of electric heating, with only 7% of households using it, compared to the national average of 19%, as the region relies more heavily on gas and oil heating systems.[39]

Figure 10 takes a look at the distributional impacts by region (this figure is analogous to Figure 4, but it reports each region separately). Again, across all specifications, one can see that our GMM-based approach consistently yields more pronounced distributional effects across income groups.

The RTP effects tend to be small and can be progressive or regressive depending on the region, while the ToU effects tend to be progressive throughout (except the fifth quintile in Castilla-y-León, where the effects are noisier). Overall, we conclude that HVAC ownership plays a critical role in shaping the distributional impacts of RTP, as it significantly affects both the level and the seasonal profile of electricity consumption. In contrast, the distributional effects of ToU pricing are less sensitive to HVAC status, since ToU does not depend on seasonal consumption patterns to the same extent.

7 Conclusions

We assess the distributional impacts of transitioning from time-invariant electricity pricing to Real-Time Pricing (RTP) and Time-of-Use (ToU) pricing in the Spanish electricity market, where RTP is the default for households and ToU on access fees has recently become mandatory. While Fabra et al. (2021) found that the adoption of RTP had little effect on aggregate household electricity consumption, its differential impact across income groups remains an open question. This is a critical issue, as concerns about adverse distributional outcomes have hindered broader RTP implementation in other regions. In this context, the distributional effects of ToU serve as a useful benchmark for evaluating those of RTP.

Using hourly electricity consumption data from a large sample of Spanish households, combined with detailed socio-demographic information, we were able to assess how household electricity bills have changed under RTP and ToU and how these changes affect different income groups. Additionally, the consumption data allowed us to infer ownership of electric heating and air conditioning – key factors in determining the distributional effects of the pricing schemes.

Our analysis indicates that the shift to RTP has been largely distributionally neutral. This null effect can be decomposed into two offsetting channels: variation in electricity prices within

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

Figure 10: RTP and ToU Effects by Region

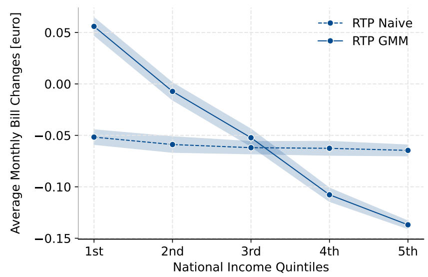

(a) RTP, Madrid

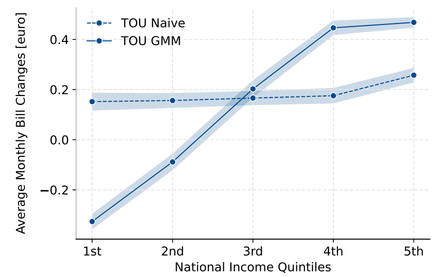

(b) ToU, Madrid

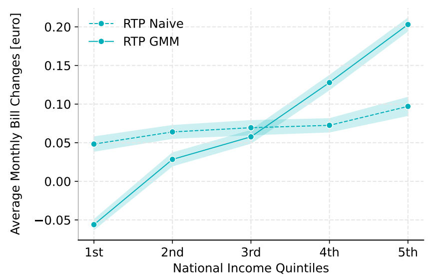

(c) RTP, Galicia

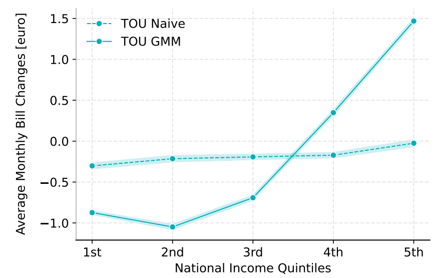

(d) ToU, Galicia

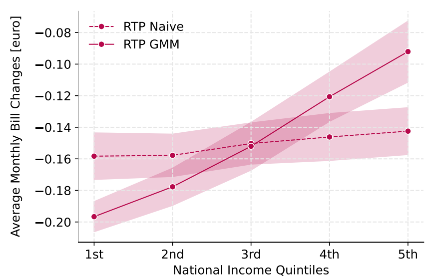

(e) RTP, Castilla y León

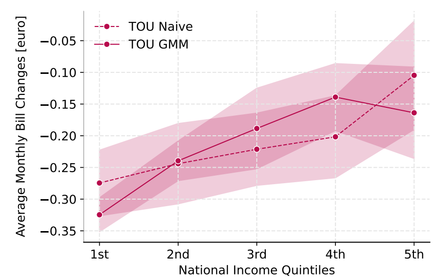

(f) ToU, Castilla y León

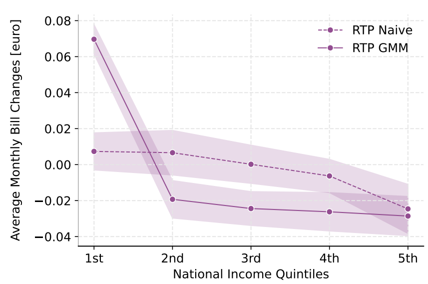

(g) RTP, Castilla-La Mancha

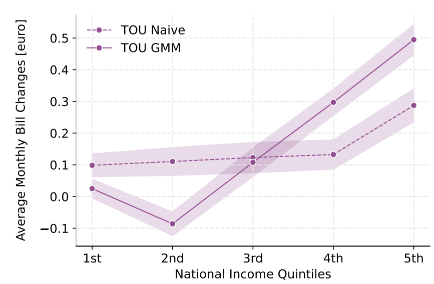

(h) ToU, Castilla-La Mancha

Notes: These figures display the impact on monthly electricity bills (in €) resulting from a shift from time-invariant pricing to either RTP (first column) or ToU (second column), for four representative regions in Spain. Results are presented under two income classification methods: our proposed GMM-based approach and a naïve alternative that ignores within-zipcode heterogeneity. Under ToU pricing, the effects are progressive in all regions analyzed. In contrast, the distributional impact of RTP varies by region – being either progressive or regressive – depending on the prevalence of specific end-use equipment (e.g., electric heating or air conditioning) in each area. The shaded areas represent the confidence intervals that account for uncertainty in our estimated measure of income.

months and across months. We find that low-income households tend to consume more electricity during lower-priced hours within a month, making RTP progressive in that dimension. However, this benefit is offset by a regressive across-month effect: low-income households rely more heavily on electric heating and therefore consume more during the winter, when prices tend to be higher.

In contrast, the shift to ToU pricing has a clear progressive impact. Because ToU tariffs remain constant throughout the year, the gains from low-income households consuming relatively more during off-peak hours are not offset by seasonal variation in electricity consumption. As a result, lower-income households benefit more from ToU than higher-income groups.

These findings underscore the usefulness of quantifying the potential distributional impacts of alternative pricing schemes – and more broadly, of a wide range of public policies for which our proposed method can prove useful. Indeed, a crucial step in this process is uncovering income heterogeneity within zip codes, as failing to do so may obscure key distributional effects. Our proposed methodology offers a path forward by enabling such granularity when actual household-level income data is not available, and it holds promise for evaluating the distributional consequences of policies in other settings as well.

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.

Banzhaf, S., L. Ma, and C. Timmins (2019). Environmental Justice: The Economics of Race, Place, and Pollution. *Journal of Economic Perspectives 33*(1), 185–208.

Berkouwer, S. B. and J. T. Dean (2022). Credit, Attention, and Externalities in the Adoption of Energy Efficient Technologies by Low-Income Households. *American Economic Review 112*(10), 3291–3330.

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 9(6).

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). 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. (2017). Private net benefits of residential solar pv: The role of electricity tariffs, tax incentives, and rebates. *Journal of the Association of Environmental and Resource Economists 4*(S1), S85–S122.

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

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

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

Chetty, R., J. N. Friedman, M. Stepner, and O. I. Team (2023). The economic impacts of covid-19: Evidence from a new public database built using private sector data*. *The Quarterly Journal of Economics 139*(2), 829–889.

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.

Douenne, T. (2020). The vertical and horizontal distributive effects of energy taxes: A case study of a french policy. *The Energy Journal Volume 41*(Number 3), 231–254.

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.

Enrich, J., R. Li, A. Mizrahi, and M. Reguant (2022). Smart Meters and Retail Competition: Trends and Challenges. *AEA Papers and Proceedings 112*, 461–65.

Enrich, J., R. Li, A. Mizrahi, and M. Reguant (2024). Measuring the impact of time-of-use pricing on electricity consumption: Evidence from spain. *Journal of Environmental Economics and Management 123*, 102901.

European Commission (2009). 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., D. Rapson, M. Reguant, and J. Wang (2021). Estimating the Elasticity to Real-Time Pricing: Evidence from the Spanish Electricity Market. *AEA Papers and Proceedings 111*, 425–29.

Fabra, N. and M. Reguant (2024). The energy transition: A balancing act. *Resource and Energy Economics 76*.

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). 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.

Goodman, L. A. (1953). Ecological regressions and the behavior of individuals. *American Sociological Review 18*(6), 663–664.

Greiner, D. J. and K. M. Quinn (2009). R x C Ecological Inference: Bounds, Correlations, Flexibility and Transparency of Assumptions. *Journal of the Royal Statistical Society: Series A 172*(1), 67–81.

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

Gustman, A. L. and T. L. Steinmeier (2000). 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). 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.

Jack, K. and G. Smith (2020). Charging ahead: Prepaid metering, electricity use, and utility revenue. *American Economic Journal: Applied Economics 12*(2), 134–68.

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). Dynamic Pricing of Electricity. *American Economic Review 102*(3), 381–385.

King, G. (1997). *A Solution to the Ecological Inference Problem: Reconstructing Individual Behavior from Aggregate Data*. Princeton, NJ: Princeton University Press.

Leslie, G. W., A. Pourkhanali, and G. Roger (2024). Is the clean energy transition making fixed-rate electricity tariffs regressive? *Journal of Environmental Economics and Management 127*, 103040.

Levinson, A. and E. Silva (2022). 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.

OECD (2024). Distributional impacts of energy transition pathways and climate change. Technical report, OECD, Economics Department.

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 (2023). Barriers to real-time electricity pricing: Evidence from new zealand. *International Journal of Industrial Organization 89*, 102979.

Robinson, W. S. (1950). Ecological correlations and the behavior of individuals. *American Sociological Review 15*(3), 351–357.

Schittekatte, T., D. Mallapragada, P. L. Joskow, and R. Schmalensee (2024). Electricity retail rate design in a decarbonizing economy: An analysis of time-of-use and critical peak pricing. *The Energy Journal 45*(3), 25–56.

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

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 on 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.[40] We only include places from which we have electricity consumption data. This limits our analysis to four regions: Galicia, Castilla y León, Madrid, and Castilla-La Mancha. Figure A.2 plots the location of these provinces.

We complement the INE data with the MB Research data at the postal code level. INE data reports the 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 income distributions complement each other in 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.[41] 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").[42] A district section and a zip code are matched if the latitude and longitude of the address are within 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 about the

[40] 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 may not have their data reported.

[41] 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.

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

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).

hourly electricity consumption for close to four million Spanish households from January 1st, 2016 to May 31st, 2017. It was provided to us by the distribution subsidiary of Naturgy, which is one of the largest Spanish utility companies. One of the advantages of our data is that we have access to the universe of meters in a given geography where Naturgy operates as the single distribution company (regardless of whether they contract with a different retailer).[43] This is mostly the case of Madrid and Galicia, although Naturgy is also present in other zip codes scattered through other parts of Spain.[44]

After treating outliers with overly zero consumption observations or missing zip code data, [45] the final sample contains 1,303,350 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). 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. In addition to hourly consumption, a key part of the identification is the contracted power of the home, which is the maximum consumption allowed at any 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. This is highly correlated with household income and is therefore crucial to our identification strategy. Figure A.1 plots its distribution. The dashed line divides households into two categories: high- and low-contracted power. we use this categorical variable in our estimation.

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.

[43] Most zip codes are only associated with a single distribution company, as Spain is organized in large distribution areas. Municipalities typically only belong to one, with the exception of the city of Madrid. However, in Madrid we have several zip codes, and most of the zip codes belong to only one of two utilities in the city.

[44] The geographic distribution of households is shown in the Appendix in Figure A.2.

[45] 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.

Figure A.1: Distribution of household contracted power

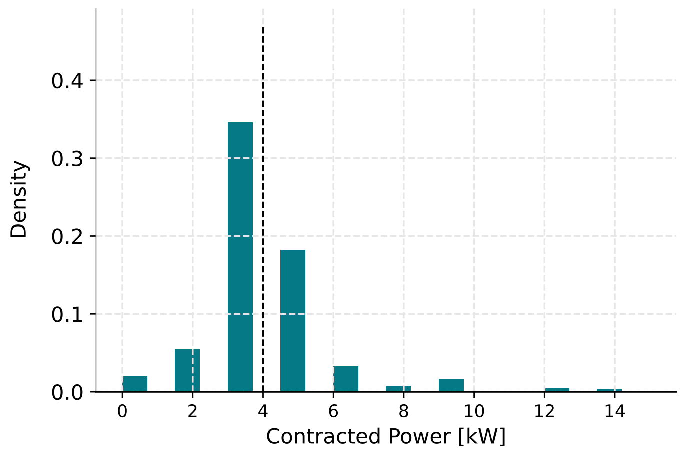

Notes: This figure plots the distribution of contracted power in our data. The dashed line divides households into two categories: high and low contracted power. We use this categorical variable in our estimation.

Figure A.2: Geographic distribution of households

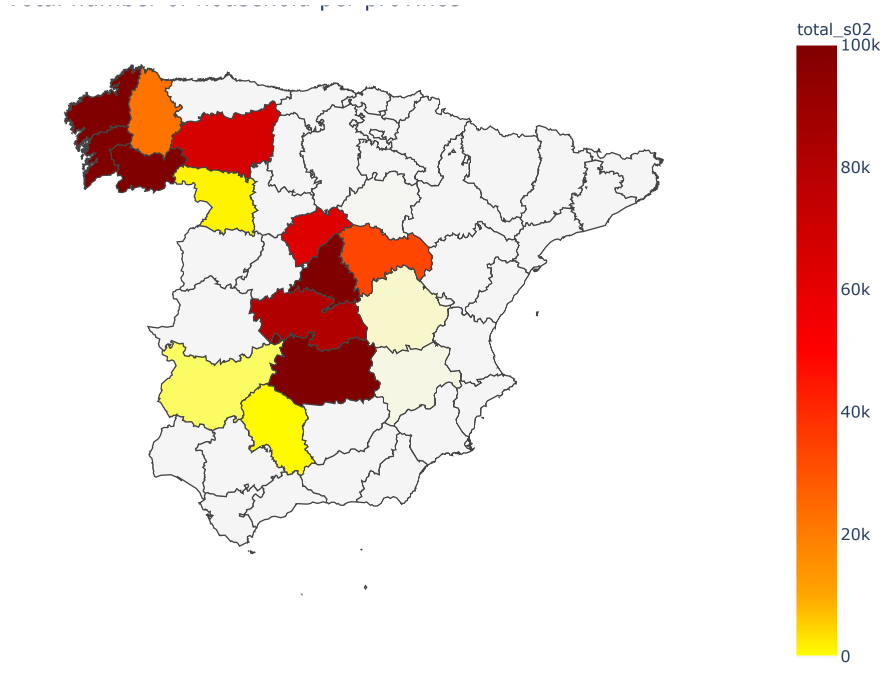

A.4 Consumption expenditure survey

We use data from the Spanish consumption expenditure survey ('Encuesta de Presupuestos Familiares' or EPF in Spanish), provided by the Spanish National Statistics Institute (INE). We use the 2016 consumption survey because it is the one that best matches our sample period.

In the microdata, we observe each individual entry of the survey, which can be geographically linked to the province and contains a survey weight to ensure the representativeness of the survey. We have 5,517 individuals who belong to the geographical areas covered by our utility. The annual nature of CEX data makes it unfit to examine the distributional impact of real-time pricing on households. However, these data allow us to validate the predictions from our methodology in Section 4.

The survey contains detailed expenditure data for electricity consumption. In particular, it asks for annual electricity expenditures and annual electricity consumption, which are then averaged at the monthly level. Additionally, it contains the income decile at the individual level. We make assumptions regarding the components of the electricity bill so that we can additionally infer contracted power at the individual level. Table A.2 provides summary statistics of these variables along the quintiles of the income distribution.

More precisely, to infer the contracted power, we assume the following. During the sample period, we know that electricity taxes are 21% on the total bill, and an additional 5.113% on most of the bill (other than the ad hoc fees), and that the annual price per contracted power is 42 euros per kW, or 3.33 euros per month (plus taxes). These are regulated fees that are observed. We also need to make assumptions about other ad hoc fees contained in the bill (e.g., to pay for the smart meter and other services), which we assume to be 1.5 euros based on the context. We also need to assume a pre-tax price per kWh of 12.45 cents, which we take from the average of a quarterly survey of electricity prices by the competition authority (CNMC).[46] Variations in these assumptions provide comparable correlations of income and contracted power.

Given quantity Q_i and bill amount B_i, contracted power is then given by,

$$CP_i = \frac{B_i/1.21 - 0.1245 * Q_i * (1.05113) - 1.5}{3.33}.$$

We also compute the bill share of the contracted power costs by computing:

$$BS_i = \frac{CP_i * 3.33 * 1.21}{B_i}.$$

We see that contracted power is correlated with income and that contracted power represents a larger share of households' bills for high-income households, given that these are the ones that tend to contract larger amounts than they typically use.

[46] These data are also used in Enrich et al. (2022) and described there in more detail.

Table A.2: Summary statistics from the CEX survey 2016

Quintile	Expenditure	Quantity	Contracted Power	Bill share of CP
1	43.34	181	2.88	0.30
2	54.72	221	4.01	0.33
3	58.52	236	4.35	0.34
4	61.99	246	4.77	0.34
5	71.80	281	5.74	0.35

Notes: Own elaboration based on the Spanish National Statistics Institute (INE), Household Expenditure Survey (https://www.ine.es/dyngs/INEbase/es/operacion.htm?c=Estadistica_C&cid=1254736176806&menu=ultiDatos&idp=1254735976608). Only observations from Castilla-La Mancha, Castilla-y-León, Galicia, and Madrid are included. Expenditure is reported as a monthly average of total electricity bills (in euros), quantity is reported in monthly kWh consumed, and contracted power is reported in kW and is estimated based on the two previous variables and a set of assumptions regarding average prices, taxes, and fees during the period.

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).

We use a daily regression approach based on weather data.[47] For each household, we first run the following regression to obtain the correlation between its electricity consumption and temperature in winter and summer:

$$kWh_d^i = \alpha^i + \beta_{HDD}^i HDD_d^i + \beta_{CDD}^i CDD_d^i + \epsilon_d^i \tag{B.1}$$

where kWh_d^i is the hourly consumption of household i in day d, and HDD_d^i and CDD_d^i are the heating degree days and cooling degree days for that day. The coefficients of interest are β_{HDD}^i and β_{CDD}^i, which measure how much more a household consumes in response to a day in need of heating (HDD, typically in the winter) and a day in need of cooling (CDD, typically in the summer), respectively.

Once we have an estimate of β_{HDD}^i and β_{CDD}^i, we set some thresholds to determine whether a household is assigned a flag for owning electric heating (HDD coefficient) and/or air conditioning (CDD coefficient). First, we only flag them if the estimate we obtain is statistically significant. Second, we assign appliance ownership only to the highest estimates, to match the share of appliance ownership in our data. Finally, and given that this reduces the amount of information provided by our imputed appliance ownership, we preserve the information in the coefficients in our final estimation and consider a robustness check in which these coefficients are directly used to infer income, rather than their discretized "dummy variable" versions.

To show that our classification of household types is informative about their consumption behavior, we plot consumer daily load curves by identified HVAC status in Figure B.1. EH owners have relatively higher consumption during both day and night because electric heating devices are, in general, more energy-intensive than AC, as shown in Panel (a). We also observe that high consumption is particularly high during winter months for households that have electric heating, while it peaks in the summer for those households with air conditioning, as shown in Panel (b).

Once we perform the income estimation, we can also check the correlation between income and HVAC mode. As shown in Figure B.2, 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. Although air conditioning tends to be associated with high income (for regions with a significant share of air conditioning), electric heating is negatively correlated with income, particularly in the most urban regions (Madrid), as newer buildings tend to rely on city gas for heating.

[47] Alternatively, we also implemented an estimator based on hourly data. The two approaches provided similar results, and therefore we use the simpler procedure based on daily outcomes.

Figure B.1: Load curves by HVAC status

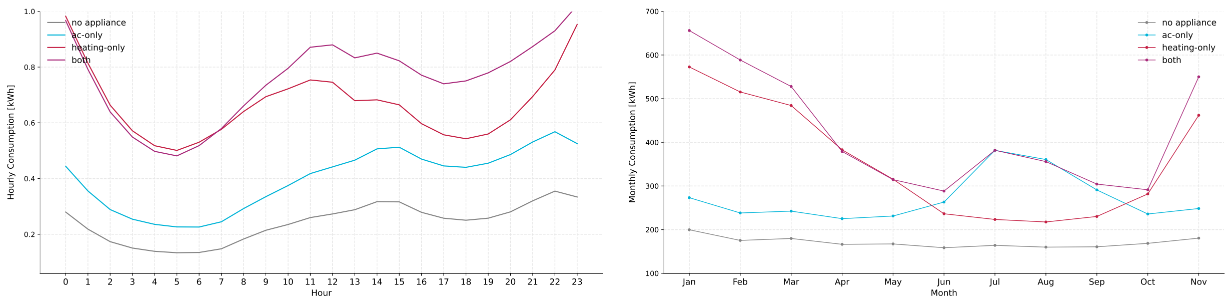

(a) Hourly consumption pattern [kWh]

(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.2: HVAC status and Income

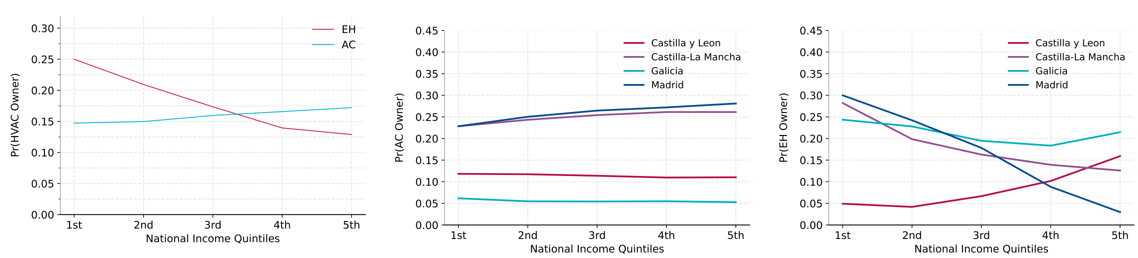

(a) Share of electric heating owners and AC owners

(b) AC ownership by state

(c) EH ownership by state

C Validation of the methodology

C.1 CEX validation

We report the estimated relationship between income and electricity consumption using the GMM approach and the naïve approach. In Figure C.1, 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 C.2.2 and C.5, 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 C.1, the GMM results have a lower slope than the CEX survey results, as expected in the simulation C.5. 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.

C.2 Monte Carlo simulation

We conduct Monte Carlo simulations to evaluate the effectiveness of our two-step procedure compared to the naïve approach and the true data.

In these exercises, we simulate household incomes, types, and policy impacts. We then infer household income and the distributional impact using both the naive approach and our two-step procedure, under the assumption that household income is not directly observed. We compare the inferred household income with the true household income in Section C.2.1 and the estimated distributional impact with the true distributional impact in Section C.2.2.

Section C.2.1 suggests that the two-step procedure performs significantly better than the naïve approach in approximating household income. We also examine the sensitivity of this result to our identification assumptions and assess how misspecification affects the outcomes. Section C.2.2 further demonstrates that the two-step procedure can capture the full distributional impact when income affects policy outcomes only indirectly through its correlation with household types. In contrast, the naïve approach can capture only location-based variation. When income directly influences policy outcomes (e.g., when a policy explicitly targets high- or low-income groups), the two-step method cannot fully recover the impact but still performs significantly better than the naïve approach.

Figure C.1: Relationship between income and electricity consumption: Comparing survey results, the imputed income results, and the naïve income results

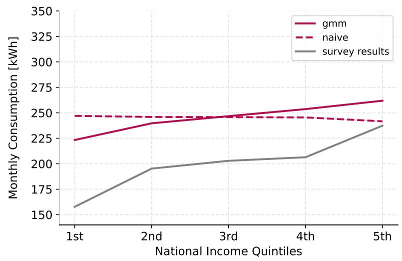

(a) Castilla y León

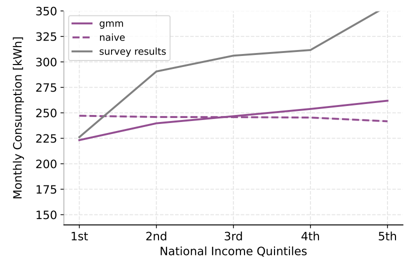

(b) Castilla-La Mancha

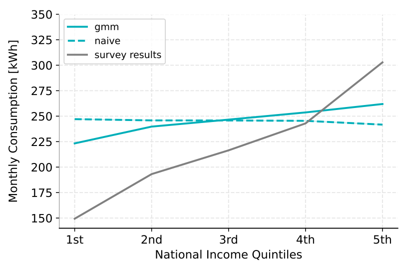

(c) Galicia

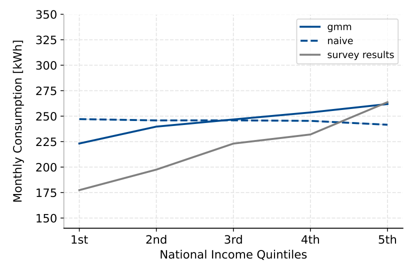

(d) Madrid

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.

C.2.1 Monte Carlo assessment of estimation of 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 joint distribution of household types, 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 detailed steps to create the data underlying the Monte Carlo simulation are as follows:

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

2. We assign a zip code number to each household based on a pre-established probability that type θ belongs to zip code z, $Pr(z|\theta)$. There are 50 zip codes and we classify them into 10 groups.

3. We sort types based on their peak consumption 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. We also allow different zip code groups to have different income-type relationships.

5. Household 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 based on their consumption pattern and then estimate income distribution of each type by fitting 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 expected income as households' true income. We compare it to 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 the two-step method, the imputation will be by estimated type. Figure C.2 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 is much better aligned with the truth. This fit is naturally improved as we allow for more types and data.

Another way to see this result is to show 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.3a. As we allow for more types, the distribution of income of households becomes more different along quintiles, as shown in Panel C.3c.

Finally, we examine if our inferred income is still more correlated with true income than with the naïve approach in the presence of misspecification in Figure C.4. In our non-parametric estimator, a key identification assumption is that the income distribution for each consumer type is identical

Figure C.2: Simulation results: Imputed income by household type

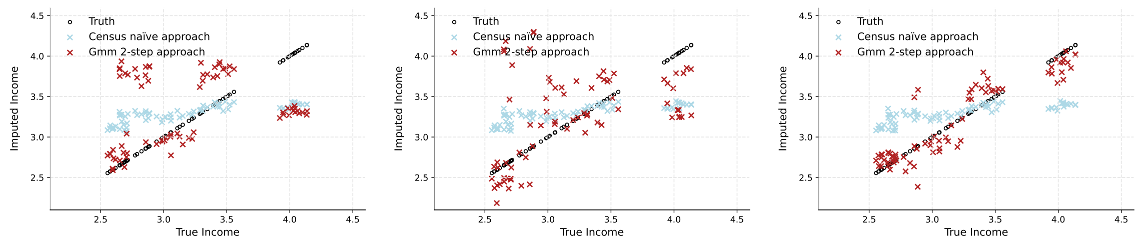

(a) 2 types, sample size 200 (b) 3 types, sample size 200 (c) 4 types, sample size 200

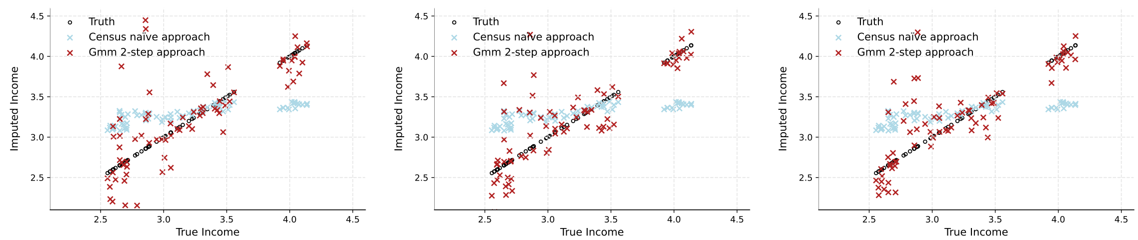

(d) 5 types, sample size 200 (e) 5 types, sample size 500 (f) 5 types, sample size 1000

across zip codes. Importantly, we group zip codes that are similar to each other. What happens when these zip code groups are not in line with the true data generating process? We find that misclassifying zip codes into heterogeneous groups still leads to an improved correlation between imputed income and the true expected income, as long as there is some commonality. While the distribution of income could be improved, it is still much more correlated with the underlying true income distribution than the naïve income distribution, as seen in both Panel (a) and (b) in Figure C.4. This suggests that potentially misspecifying the zip code group would still substantially improve accuracy relative to the naïve approach, yielding an income approximation close to the true distribution.

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.

C.2.2 Monte Carlo assessment of the two-step procedure

Once we obtain an improved estimate of the income distribution, we then use it so summarize the association between income and the impacts of dynamic pricing. Under what conditions can our estimates uncover the full distributional impact of a given policy?

We perform an additional Monte-Carlo simulation to inform this discussion by expressing our policy assessment exercise as a regression framework. Assume that the true data generation process behind the distributional impacts is governed by the following equation, where the key primitives

Figure C.3: Simulation results: Distribution of imputed income conditional on true income quintile

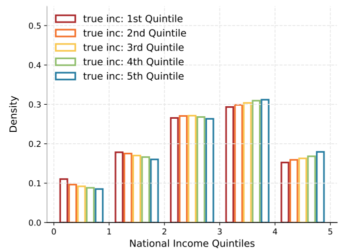

(a) Naïve approach, sample size 500

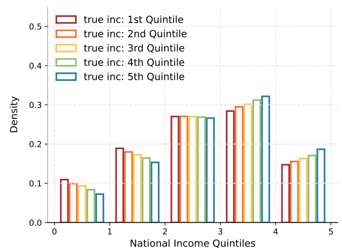

(b) 2 types, sample size 500

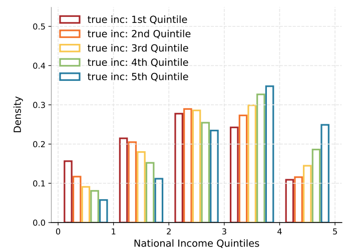

(c) 5 types, sample size 500

Figure C.4: Simulation results: Imputed income (5 types, wrong zip code group, sample size 1000)

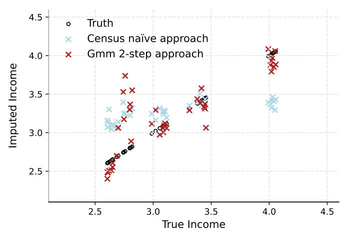

(a) Imputed type-average income vs True income

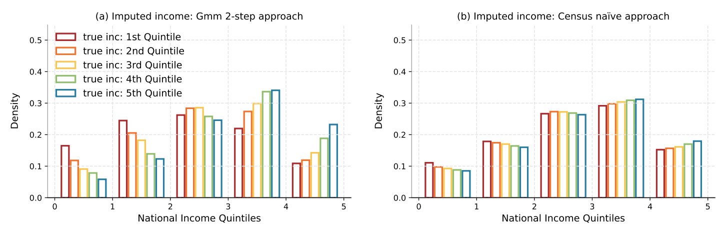

(b) Distribution of imputed income for each true income quintile

are highlighted in red for clarity:

$$impact_{i,z} = t \times \theta_i + k \times inc_i + \sigma_z \times (\phi_z + \overline{\phi}_{zipgroup}) + \sigma_e \times \epsilon_{iz}. \tag{C.1}$$

The policy impact, $impact_{i,z}$, on household i in zip code z 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 σ_e are scale the importance of regional fixed effects and unobservables, respectively.

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 (C.1), including location impacts σ_z and individual impacts t. The zip code level variation is observed, and types and income are inferred.

In our simulations, we explore whether our method can recover the full impact through all these channels. To do so, we calculate the income effects when we use the true distribution of income, our inferred income distribution, and the naïve zip code income distribution. To make the problem interesting, household heterogeneity (types) and income are correlated. Thus, using the wrong measure of income (naïve approach) can substantially bias the results.

Figure C.5: Assessing the method with a Monte-Carlo simulation

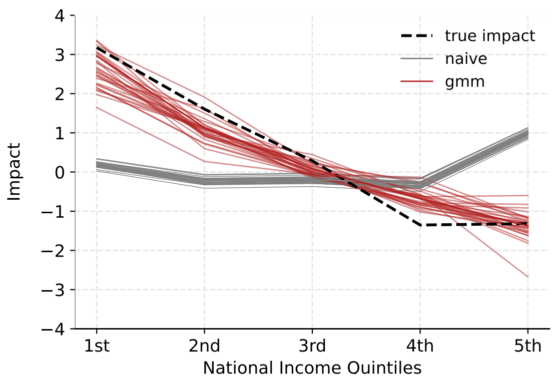

(a) Full bias correction

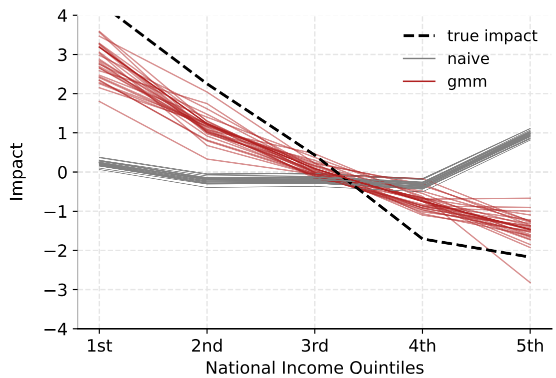

(b) Partial bias correction

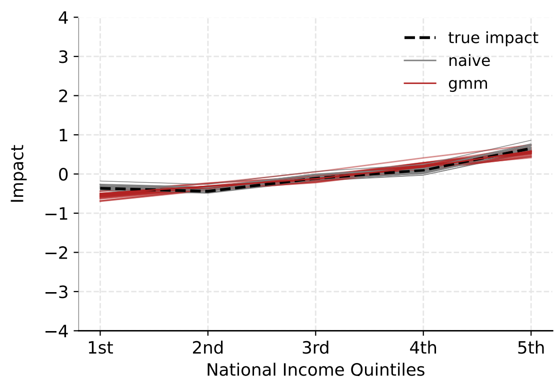

(c) No naïve bias

Notes: The figure shows 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 C.5 reports the estimated policy impacts across these specifications under different parameter assumptions. When $k = 0$, income affects policy outcomes only indirectly through its correlation with household types and other observables. In this case, the two-step approach can fully recover the distributional impact. In other words, types, individual characteristics, and zip codes must together be sufficient statistics for determining policy impacts. Panel (a) shows that in the case, the GMM approach can capture the full distributional impact through individual types and household locations, while the naïve approach can even be biased in the opposite direction of the true effects.

Our method improves the assigned distribution of income to a given household but cannot predict the exact realized income of a household. In this case, realizations of household income (not the expected distribution) directly enter the policy impacts via omitted variable bias. We therefore explore how well our approach captures the distributional impact and how much it improves upon the naïve approach when $k \neq 0$, even if imperfectly specified.

As shown in Panel (b), if the policy impact is correlated with the remaining unobserved income, we will only partially identify the effect: the portion correlated with household types and locations. The naïve approach captures only location-level variation. Because our method additionally captures variation correlated with household types, it achieves a substantial improvement over the naïve approach. The naïve estimator, by contrast, continues to be biased in the wrong direction.

Finally, when both k and t are 0, the two approaches yield the same results, both consistent with the true impact, as shown in panel (c). In this case, the policy affects households only through geographical locations, and the true causal impact of individual types and income is zero. All variation lies at the zip-code level and, therefore, the effects are well captured by the naïve distribution of income. This assumption is probably rejected in most applications, but it is a good benchmark to understand the conditions under which the naïve estimator provides a valid answer.

C.3 Cross-validation

While the Monte Carlo exercises have validated our two-step approach, this section uses real data evidence to conduct further checks. 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. For each province, with Z zip code regions, we repeat the following procedure Z times:

1. Drop one zip code z.
2. Estimate household income with the $Z-1$ zip code using the semi-parametric GMM approach.
3. Predict household income for households in zip code z using observed household types and demographics in zip code z and estimation results from step 2.
4. Aggregate inferred household income to the zip code level and obtain the inferred income distribution for all Z zip codes.
5. Get out of sample error by comparing zip code z's inferred income distribution with observed census income distribution.
6. Get in-sample error by assessing the difference for the rest $Z-1$ zip codes.

Figure C.6 reports the distribution of in-sample and out-of-sample prediction errors for all provinces. 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.

Figure C.6: Distribution of prediction errors

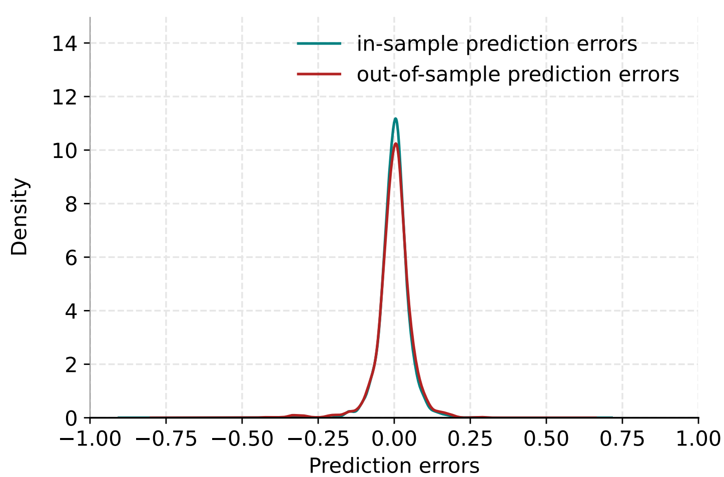

Notes: These figures illustrate the distribution of zip-code-level prediction errors.

D Robustness to alternative specification

We examine the robustness of our results to the number of types, both *kmeans* and discrete categories such as heating mode or contracted power.

Specifications We consider several specifications that aim at examining the robustness to the number of types in the k-means clustering procedure, together with the controls that are used in the semi-parametric estimator.

- **Non-parametric**
  - *kmeans* (2 to 4 *kmeans* groups) combined with heating interacted with low and high contracted power types.
  - *kmeans* (2 to 4 *kmeans* groups) combined with any HVAC interacted with low and high contracted power types.
  - *kmeans* (2 to 4 *kmeans* groups) combined with heating, non-heating low contracted power, and non-heating high contracted power types (fewer types).
- **Semi-parametric** All the different groupings above combined with several parametric controls that allow for individual flexibility along the following covariates:
  - Spec 1: individual contracted power, consumption, peak consumption, and slopes from HVAC estimation [main].
  - Spec 2: individual contracted power, consumption, zip-code unemployment, zip-code share above 50 years old.
  - Spec 3: only slopes from HVAC estimation.
  - Spec 4: consumption, peak consumption, and the loss from RTP.

Figure D.1 shows that the results are consistent across a wide range of specifications. Given that we find that the impacts of RTP are small and not particularly related to income, the results can be somewhat sensitive to the specification, but the qualitative take away is the same. For the ToU results, the qualitative and quantitative results are very robust across all these specifications.

Figure D.2 displays the same sensitivity results by region. We find that the results are very consistent across this wide range of specifications for Madrid and Galicia, the regions with most households. For the other two regions, the results are also qualitatively consistent across specifications but somewhat more noisy.

Figure D.1: Robustness Specifications

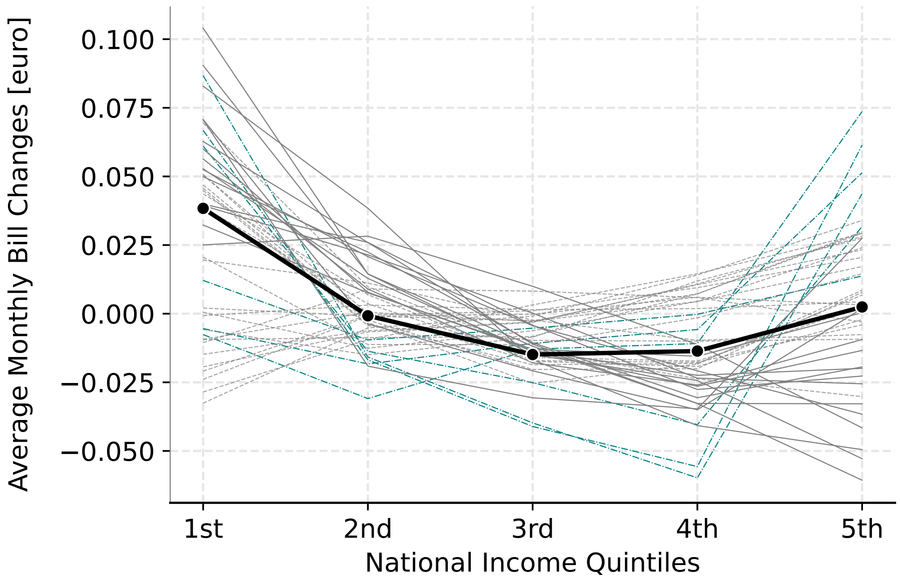

(a) RTP Robustness

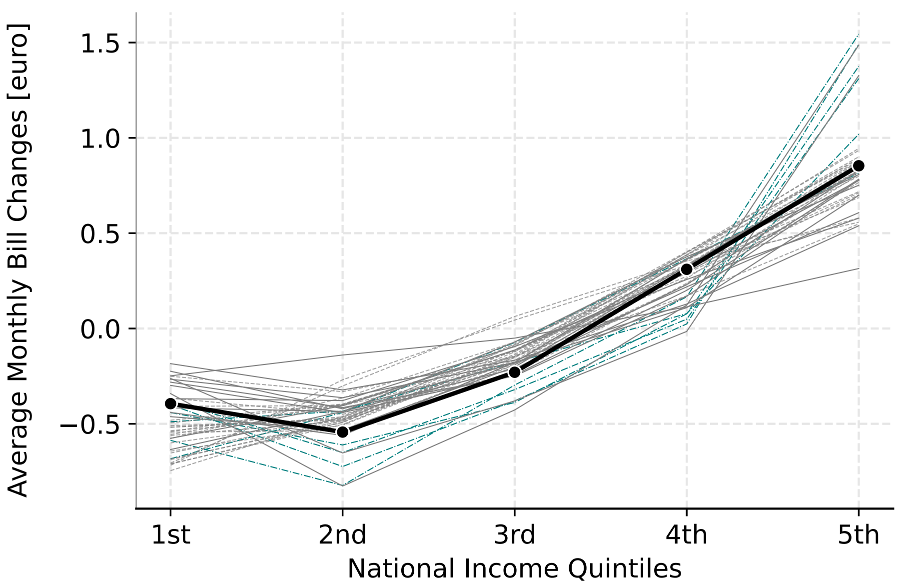

(b) ToU Robustness

Notes: These figures display the robustness of our main results to alternative specifications. Alternative parametric specifications lead to consistent outcomes. The dashed teal lines display results from non-parametric specifications.

Figure D.2: Robustness by Region

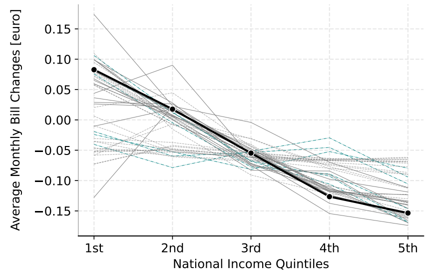

(a) RTP, Madrid

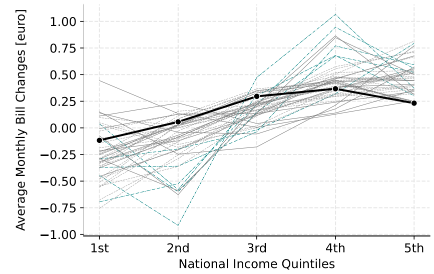

(b) ToU, Madrid

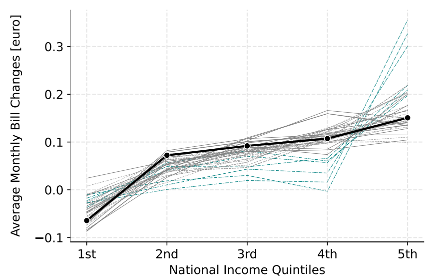

(c) RTP, Galicia

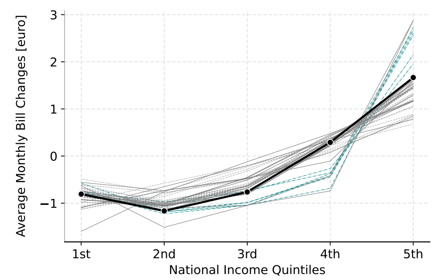

(d) ToU, Galicia

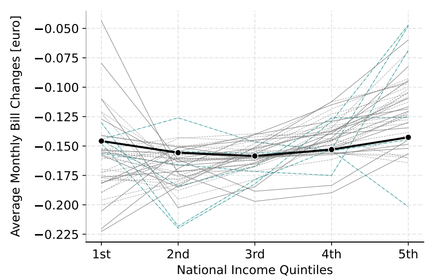

(e) RTP, Castilla y León

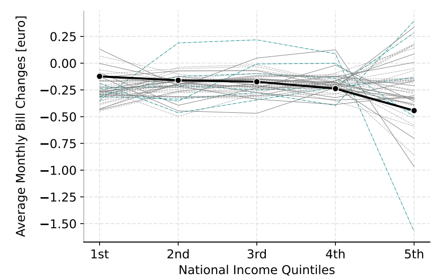

(f) ToU, Castilla y León

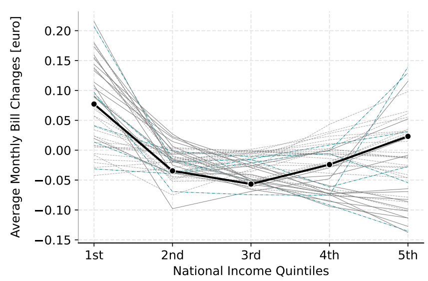

(g) RTP, Castilla-La Mancha

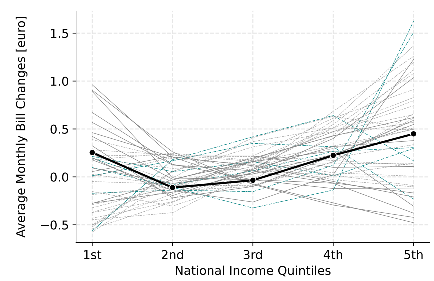

(h) ToU, Castilla-La Mancha

Notes: These figures display the robustness of our main results to alternative specifications. Alternative parametric specifications lead to consistent outcomes. The dashed teal lines display results from non-parametric specifications.