



Time-varying pricing may increase total electricity consumption: Evidence from Costa Rica

Tabaré Capitán^{a,*}, Francisco Alpizar^b, Róger Madrigal-Ballesteró^c, Subhrendu K. Pattanayak^d

^a Department of Economics, Swedish University of Agricultural Sciences, Sweden

^b Department of Social Sciences, Wageningen University, Netherlands

^c EJD – Central America at the Tropical Agricultural Research and Higher Education Center (CATIE), Costa Rica

^d Sanford School of Public Policy & Department of Economics, Duke University, United States

ARTICLE INFO

Article history:

Received 15 July 2020

Received in revised form

21 September 2021

Accepted 28 September 2021

Available online 30 September 2021

JEL classification:

Q41

Q47

Q50

Keywords:

Dynamic pricing

Energy

Behavioral adjustments

LMICs

ABSTRACT

We study the implementation of a time-varying pricing (TVP) program by a major electricity utility in Costa Rica. Because of particular features of the data, we use recently developed understanding of the two-way fixed effects differences-in-differences estimator along with event-study specifications to interpret our results. Similar to previous research, we find that the program reduces consumption during peak-hours. However, in contrast with previous research, we find that the program increases total consumption. With a stylized economic model, we show how these seemingly conflicted results may not be at odds. The key element of the model is that previous research used data from rich countries, in which the use of heating and cooling devices drives electricity consumption, but we use data from a tropical middle-income country, where very few households have heating or cooling devices. Since there is not much room for technological changes (which might reduce consumption at all times), behavioral changes to reduce consumption during peak hours are not enough to offset the increased consumption during off-peak hours (when electricity is cheaper). Our results serve as a cautionary piece of evidence for policy makers interested in reducing consumption during peak hours—the goal can potentially be achieved with TVP, but the cost is increased total consumption

© 2021 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Economists contend that basic economic principles can help attain energy efficiencies. For example, the cost of producing electricity varies both during the day and between days, but most households pay a constant price per kilowatt hour (kWh) of electricity consumption. Because that price does not reflect the cost of production at the time of consumption, there is no incentive for households to consume less electricity when the cost of production is high. In response to this problem, the use of time-varying pricing (TVP) schedules—under which the price is higher when the cost of production is higher—can be a mechanism for allocative efficiency (Allcott, 2011; Wolak, 2011; Jessoe et al., 2014).

* Corresponding author at: Department of Economics, Swedish University of Agricultural Sciences, Box 7013, 750 07 Uppsala, Sweden.

E-mail addresses: Tabare.Capitan.Jimenez@slu.se, Tabare.Capitan@gmail.com (T. Capitán), Francisco.Alpizar@wur.nl (F. Alpizar), rmadriga@catie.ac.cr (R. Madrigal-Ballesteró), subhrendu.pattanayak@duke.edu (S.K. Pattanayak).

Such effectiveness comes from two ways in which households can respond to the implementation of TVP: technological changes (e.g., to replace an inefficient heater for a more efficient one or to automate air conditioners to respond to the time-varying prices of electricity) and behavioral changes (e.g., to manually turn off lights when rooms are not in use or to do laundry late at night instead of during the evening). Our ability to predict the effects of implementing TVP in different contexts depends on our understanding of the responses that lead to changes in electricity consumption. For example, implementing TVP with a population that already uses modern and efficient appliances would leave little room for technological changes, and reductions would need to rely on behavioral changes. The challenge is to isolate the role of each type of changes to understand the net effect of households' responses to the implementation of TVP.

In this paper we tackle part of the challenge by partially isolating the portion of the effect of implementing TVP that comes from behavioral changes. To do so, we study the implementation of a voluntary time-of-use (TOU) pricing program by a major electricity utility serving San José, Costa Rica. Critical to our study, because of year-round mild temperatures, <2% of the households in San José use heating or cooling devices (Ministerio de Ambiente y Energía, 2019), leaving little room for technological responses to the TOU program. Specifically, we use a differences-in-differences research design—i.e., we compare outcomes for units whose treatment status changes to units whose treatment status does not change—to estimate the effect of the implementation of TVP on total consumption and interpret the estimate as the effect of behavioral changes. Contrary to the common finding in this literature (Faruqui and Sergici, 2010; Mina Badtke-Berkow et al., 2015; Dutta and Mitra, 2017), we find that households increase their total consumption in response to the implementation of TOU pricing.

Since participation in the program is voluntary and households self-select into the program, our results are not informative about the effect of mandatory TOU pricing.¹ However, similar to Allcott (2011), we argue that estimates from a voluntary program are of policy interest because TVP is likely to be optional over the next few years, given that a mandatory shift to TVP would increase the bill for consumers that use more electricity than average at times when market prices are high (Borenstein, 2007). Therefore, our results are helpful to understand the early phases of these potential voluntary programs. For example, assuming a Roy model of selection on gains (Heckman and Vytlacil, 2007), our estimates provide an upper bound for the effect of mandatory TOU pricing coming from behavioral responses.

We add to an empirical literature on the effect of TVP that has largely relied on data from natural and field experiments, often in collaboration with electricity utilities (typically in the US and the EU). In this literature, researchers consistently find that TVP schedules, in different forms, reduce both consumption during peak hours and total consumption (Faruqui and Sergici, 2010; Mina Badtke-Berkow et al., 2015). For example, Wolak (2011) find this in the context of hourly pricing, critical peak pricing, and critical peak pricing with rebate; Allcott (2011) in the context of Real-Time pricing; and Jessoe et al. (2014) in the context of Time-of-Use pricing.

These reductions reflect the net effect on electricity consumption of three potential household responses to the TVP: substitution of consumption over time (e.g., doing laundry late at night instead of during the evening), substitution away from comfort (e.g., turning off the heater and feeling cold), or substitution toward a more efficient energy capital stock (e.g., using a fan instead of air conditioner).² We claim that most of the empirical evidence on the effect of TVP comes from contexts in which substitution toward a more efficient energy capital stock—associated with technological changes—plays a major role in the response of households to the implementation of TVP. The reason is that most of the data previously used to study the effect of TVP comes from rich countries in which heating and cooling devices drive consumption (e.g., Canada, United States, Australia, Japan, and countries in western Europe).³ While none of these studies explicitly identify the specific changes in the household that contribute to the reduction in consumption, it is reasonable to assume that changes related to cooling and heating devices (e.g., installing a more efficient device or adjusting the settings of the device) are likely the main drivers of the reductions in electricity consumption.⁴

In the absence of separate estimates of the role of technological and behavioral changes in reducing electricity consumption, the mechanism behind the findings in this literature—i.e., reduction of total consumption and consumption during peak hours—remains a conjecture (Allcott, 2011; Jessoe et al., 2014). The typical conjecture is that both behavioral and technological changes reduce consumption during peak hours. And any potential increase in consumption during off-peak hours (e.g., due to substitution over time) is more than offset by the reduction in consumption during peak hours and the effect of technological changes in reducing consumption at all times during the day (e.g., from using a more efficient air conditioner). Our empirical result—i.e., an increase in total consumption—is consistent with the typical conjecture once the differences in the context are accounted for. Specifically, in the near absence of technological changes (as is expected in Costa Rica), behavioral changes to reduce consumption during peak hours are not enough on their own to offset the increased consumption, due to substitution over time, during off-peak hours. Further, the consumption of every kWh of electricity

¹ This is a common feature in most of the literature (Faruqui and Sergici, 2010; Mina Badtke-Berkow et al., 2015).

² Note that these categories are not exclusive. For example, using a fan instead of air conditioner might also imply substitution away from comfort.

³ Allcott (2011) used data from an experiment with households who had recently purchased an energy efficient air conditioner. While his estimates of the effect are likely low relative to what it could have been without the new air conditioners, there was still plenty of room for technological changes, since the experiment was conducted in Chicago, where heating devices are widely used.

⁴ Some studies use ex post surveys to ask participants what they think they did in response to the TVP. For example, participants in the experiment studied by Allcott (2011) say they turned off lights, used fans instead of air conditioners, turned down air conditioners, and washed clothes during low price hours instead of during the afternoon. Still, these data do not allow us to infer the contribution of each type of response to the change in consumption.

Block pricing	Time-varying pricing
0 – 30 kW h	0–300 kW h
CRC 2,219.40 (flat rate)	Peak CRC 156.90 (per kW h)
31 – 200 kW h	Mid-peak CRC 65.05 (per kW h)
CRC 73.98 (per kW h)	Off-peak CRC 26.78 (per kW h)
201 – 300 kW h	301-500 kW h
CRC 113.53 (per kW h)	Peak CRC 178.60 (per kW h)
Over 300 kW h	Mid-peak CRC 72.72 (per kW h)
CRC 117.37 (per kW h)	Off-peak CRC 30.62 (per kW h)
	Over 500 kW h
	Peak CRC 211.75 (per kW h)
	Mid-peak CRC 85.47 (per kW h)
	Off-peak CRC 39.55 (per kW h)

Fig. 1. Pricing schedules.

that a household can substitute over time—moving away from the peak hours—is cheaper.⁵ Hence, by the law of demand, households will consume more than what they would have consumed during peak hours for every kWh substituted over time.⁶ For example, if members of a household can do laundry at 9 P.M. instead of 6 P.M., they would face a lower price per washing load under TVP than under non-TVP; with a downward-sloping demand curve, they might choose to do laundry three times per week instead of two times per week. We further develop this intuition (with a stylized model in Section 5) after presenting our empirical results.

Note, however, that the mechanism discussed above to explain how the implementation of TVP schedules can lead to a decrease (as in previous studies) or an increase (as in our study) in total consumption is contingent upon the voluntary nature of the TVP program. Due to self-selection, the households that join the program are precisely the very households that can benefit from TVP (i.e., for whom the average price is lower under TVP than what it would have been under non-TVP), either because of their prior consumption patterns or because of their ability to adjust their consumption. Of course, with a mandatory TVP program, which precludes self-selection, there will likely be some winners and some losers among the households (see Appendix A). As such, the evaluation of a (potential) mandatory TVP program might introduce new challenges and priorities, including the importance of measuring heterogeneous effects and its associated distributional consequences.

We make four contributions. First, we expand the literature and refine our understanding of the effect of TVP by partially isolating the effect of the behavioral response of residential consumers on total consumption. Second, we use recent econometric developments on the understanding of the two-way fixed effects differences-in-differences estimator to overcome data limitations and interpret our results. Third, by showing that the common finding of an increase in total consumption in response to a TVP schedule might not hold in our context, we encourage researchers to revisit short-run and long-run expected outcomes of TVP in alternative settings (Borenstein, 2005), and add a cautionary piece of evidence for decision makers interested in reducing peak-hour consumption. Fourth, to our knowledge, we provide the first empirical impact evaluation of a TVP schedule in the residential sector of a low- or middle-income country (LMIC).⁷

⁵ Fig. 1 shows that the average price under the non-TVP schedule is higher than the price during off-peak hours under the TVP schedule.

⁶ This is true both for households in Costa Rica and for households in other countries (studied in previous research). The key difference is that, in our case, there is not as much room for technological changes. Without the action of technological changes reducing consumption during the day, the increased consumption of cheaper electricity during off-peak hours leads to an increase in total consumption.

⁷ While common for the commercial and industrial sectors, few electric utilities offer TVP for the residential sector in LMICs (Dutta and Mitra, 2017). Thus, the first challenge to evaluate the impact of TVP in LMICs is that there are few TVP programs to study. The second challenge is that the few TVP programs available are often implemented independently of potential research designs that would enable the same rigorous evaluation of the effect of the programs that is often conducted in rich countries. We were able to tackle the first challenge by developing an informal relationship with an electric utility in Costa Rica over 4 years, which led to a formal partnership in which we were able to use their data to analyze the effect of their TVP program. We tackle the second challenge by leveraging recent developments in econometrics to interpret estimates from observational data. Admittedly, due to data limitations, our paper does not hold to the standard of previous research on the effect of TVP based on data from rich countries (e.g., Allcott, 2011). However, we argue that the importance of our paper is precisely to give a first empirically grounded answer to the question “What is the effect of TVP in

The rest of the paper proceeds as follows. First, we describe the context of the study—the climate and the characteristics of energy consumption in Costa Rica—and the program that introduced a voluntary TVP schedule. Second, we discuss our data and methods. Third, we present results about the effect of TVP on total electricity consumption. Fourth, we use a stylized model to discuss how our finding may not be at odds with previous findings. Finally, we summarize the findings of our research, discuss limitations, and consider potential policy implications.

2. Context: location and program

Similar to many countries in the tropics, Costa Rica has a mild tropical climate and most households do not need to use devices to heat or to cool their houses.⁸ Our data were provided by the *Compañía Nacional de Fuerza* (CNFL), a major electric utility that serves the greater metropolitan area in Costa Rica (roughly the center of the country). Specifically, we mostly use data from households in San José—the capital of the country (monthly average temperatures range from 21.8 to 23.7 °C).

As opposed to the countries in which the effect of TVP has been previously studied—where heating and cooling devices are the main drivers of electricity consumption—most households in San José do not use either. Instead, technological changes in our setting come from changes toward more efficient capital stock, such as replacing electric stoves for gas stoves or replacing old fridges—none of which might come necessarily as a response to TVP.⁹ In any case, such replacements are likely to be less frequent in our context than in rich countries with more disposable income.¹⁰ With little room for technological changes, changes in consumption ought to be driven mostly by behavioral changes (e.g., showering with lukewarm water instead of hot water, turning off lights when not in use, or doing laundry later at night instead of during the evening).

A survey of a representative sample from the population of 1,524,414 residential customers in the country, commissioned by the Ministry of Energy and Environment, helps us characterize electricity use of CNFL customers ([Ministerio de Ambiente y Energía, 2019](#)).¹¹ The average consumption of electricity in the urban sector country-wide is 214 kWh per month and 228 kWh per month in San José. This consumption is explained by food refrigeration (30%), electric cooking (7.3%) and use of other cooking appliances (8%), water heating (16.4%), lighting (11.5%), entertainment (16.3%), laundry (4.2%), and other categories of energy consumption.¹² Most households have a fridge (96.4%), a washer (96%), and fewer than 15 lightbulbs (90%). About half of the households use an electric stove (49%) and an electric shower head (42.5%). Very few households have an air conditioner (1.8%), dishwasher (0.47%), or a dryer (0.2%).

Most of the electricity produced in Costa Rica is relatively clean and inexpensive, because it is generated using renewable sources (mostly hydropower). However, the demand during peak hours sometimes cannot be met with renewable sources. During these times, relatively expensive and environmentally detrimental fossil fuels are used. In response, the *Compañía Nacional de Fuerza y Luz* (CNFL) introduced a voluntary program in 2007 intended to reduce electricity consumption of residential customers during peak hours.

Before the program, the CNFL only offered a block pricing schedule to residential customers. Under this pricing schedule, the CNFL defines four ranges of monthly consumption: 0–30, 31–200, 201–300 kWh, and more than 300 kWh. Consumers in the first group pay only a flat rate and the rest pay the flat rate plus a price per additional kWh. This price is higher for higher levels of consumption. For example, the left panel in [Fig. 1](#) shows the rates of the block pricing schedule during June 2019.¹³ Here, a consumer who uses 550 kWh pays a flat rate of CRC 2219.40, CRC 73.98 per kWh for consumption from 31 to 200 kWh, CRC 113.53 for consumption from 201 to 300 per kWh, and CRC 117.37 for each kWh consumed above the 300-kWh threshold.

With the introduction of the time-varying pricing voluntary option, the CNFL offered a hybrid pricing schedule with characteristics of both block pricing and time-of-use pricing. Here, in addition to the block pricing based on the amount of consumption, the CNFL charges a differentiated rate depending on the time of the consumption: peak hours (10:00–12:30 and 17:30–20:00), mid-peak hours (06:01–10:00 and 12:30–17:30), and off-peak hours (20:00–06:00).¹⁴ For the consumption in each of these three periods, instead of the total consumption, the CNFL uses block pricing with the following blocks of consumption: 0–300, 301–500 kWh, and over 500 kWh. As an example, the right panel in [Fig. 1](#) shows the rates for the TVP schedule during June 2019.

The program is voluntary, and the only requirement to join is for households to have a monthly consumption above 200 kWh—which is similar than the average consumption in San José (228 kWh). While the CNFL promoted the program

LMICs?” Our hope is that our research will spark interest for collaboration between researchers and electric utilities, to get the necessary data to answer the question with the same level of technical rigor used in current studies in rich countries (e.g., estimate elasticities of demand).

⁸ In addition, we use a dataset that does not include regions in which the temperature is high enough to justify the use of cooling devices (the northern side of the country and coastal regions).

⁹ Although half of the population already owns refrigerators that are 5-year-old or newer ([Ministerio de Ambiente y Energía, 2019](#)).

¹⁰ Costa Rica’s GDP per capita (PPP), as reported by the International Monetary Fund (IMF) in September 2019, is \$18,182. For perspective, the IMF’s analytical group ‘Advanced economies’ has a GDP per capita (PPP) of \$48,610.

¹¹ The survey follows up on a survey conducted in 2012 ([Ministerio de Ambiente y Energía, 2012](#)). Because the findings relevant to our study in both surveys are similar, we limit our discussion to the most recent one.

¹² The composition was estimated using self-reported data on device usage and engineering data on electricity consumption by device ([Ministerio de Ambiente y Energía, 2019](#)).

¹³ During the period of the study, consumers faced a schedule with the same structure and incentives as this one.

¹⁴ These periods are defined by the CNFL based on historical aggregate consumption data.

Table 1
Number of households in each group.

Description	Group	Number of households
Only out of the program	<i>O</i>	436,865
Only in the program	<i>I</i>	6826
Joined the program	<i>J</i>	106
Left the program	<i>L</i>	555
	Total	444,352

on their branches and through social networks when the program was introduced (circa 2007), all efforts to promote the program ceased soon thereafter to avoid the cost of replacing the single-phase meter of new volunteers for a three-phase meter.¹⁵ This replacement is necessary to implement the TVP, because most consumers have a single-phase meter that only records the total electricity use since installation, but to price consumption at different times the CNFL needs to see consumption at different times. Because the CNFL stopped promoting the program early on, most of the people who signed up for the program did it at the beginning.

There is no data on why customers chose to join the program. Given that it was not promoted during the period of consumption that we can observe, our guess is that most customers have never heard about the program and that those who joined after the CNFL stopped promoting the program found out either by word of mouth from a customer in the program, asking explicitly about the existence of such program to the CNFL, or perusing the CNFL's website. Similarly, there is no data on why customers leave the program. Thus, unlike many papers that carefully model and explain how consumers self-select into TVP programs in other settings (e.g., Ito et al., 2021), we cannot explain selection. Instead, we focus on the still relevant local average treatment effect—for households that self-select, what is the effect of the TVP program on total consumption?

3. Data and methods

We use monthly data of electricity consumption from 2011 to 2015 of all 444,352 CNFL residential electricity contracts.¹⁶ Assuming each unique contract corresponds to one unique household, our unit of analysis is the household. This assumption is realistic and only leaves out a few special cases such as households with multiple homes. Unfortunately, we do not have data about any characteristics of the households.¹⁷ For months when a household is not in the TVP program, we only observe one value corresponding to their total monthly consumption. For months when a household is in the TVP program, we observe three values corresponding to monthly consumption during each of the three periods (i.e., peak, mid-peak, and off-peak hours).

There are two main issues with the data. First, the TVP program was introduced in 2007 and most of the households joined soon thereafter. Because we only observe consumption data starting in 2011, we are left with only 106 households that joined the program between 2011 and 2015, as well as 555 households that left the program in that period (see Fig. 1).¹⁸ Thus, although we have monthly consumption data for about half a million households, only 7487 of them joined the program, and from these 7487, we only observe treatment status variation for 661 households. Second, because the program is voluntary and historical consumption over 200 kWh is required to join, households that joined the program might differ from the ones that did not join the program. We cannot address this self-selection into participation with our data. Instead, we identify groups of comparable households and obtain the local average treatment effect for each of these groups.

First, we group households according to their treatment status between 2011 and 2015: households that were not in the TVP program at any point in time (any month) between 2011 and 2015 (*O*), households that were in the program every month between 2011 and 2015 (*I*), households that were not in the program to start with, but joined the program between 2011 and 2015 (*J*), and households that were in the program to begin with but left the program between 2011 and 2015 (*L*). All households in groups *I*, *J*, and *L* were in the time-varying pricing (TVP) program for at least one period, which implies both that they met the requirement of minimum historical consumption over 200 kWh and that they chose to join the program. Because households in these three groups are likely different from households in group *O*, comparing the consumption of households in groups *I*, *J*, or *L* to the consumption of households in group *O* to estimate the effect of the TVP program would likely lead to an estimate that reflects both the effect of the program and a self-selection bias.

Indeed, as shown in Fig. 2, the average consumption of households that are only out of the program (group *O*) is about half of the average consumption of households in any of the other groups (see Fig. B.1 in Appendix B for the difference between these groups over time). Because we do not have data to address the self-selection bias we would introduce by using households in group *O* in our analyses, we exclude from the analyses the data from all 436,865 households that were

¹⁵ The financial health of the company changed for worse after the introduction of the TVP, for reasons unrelated to the TVP program.

¹⁶ After dropping data from 19 contracts of households that joined and left the program between 2011 and 2015.

¹⁷ The CNFL does not collect such data.

¹⁸ We were only able to obtain data starting in 2011 because the CNFL uses a system that was implemented back in 2011 and does not have access to consumption prior to 2011. We were informed that such data exist in a physical storage device in an undisclosed format, but we were unable to get these data.




Fig. 2. Mean consumption by group and treatment status.

never in the TVP program. Therefore, we rely on comparisons between households in the remaining groups—*I*, *J*, and *L*—to estimate the effect of the implementation of TVP on total consumption for households that self-selected into the program. For households from groups *J* and *L*, the higher mean consumption when they were in the TVP program compared to when they were out of the program suggests that households increased their consumption due to the TVP schedule (Figs. 2 and B.2 in Appendix B).

To estimate the effect of the time-varying pricing program (i.e., the treatment) on total consumption, we use a differences-in-differences (DD) research design and specify a two-way fixed effects differences-in-differences (TWFE) model. Explicitly, we use OLS to estimate

$$y_{it} = \alpha_i + \lambda_t + \beta^{DD} D_{it} + \epsilon_{it}, \quad (1)$$

where y_{it} is the electricity consumption of household i in month t , α_i and λ_t are household and time fixed effects, D_{it} is a treatment dummy that is equal to 1 if household i is in the program in month t and 0 otherwise, and ϵ_{it} is the error term of the econometric model. The treatment effect is given by the coefficient β^{DD} . We present separate estimates for households that joined the program (group *J*) and households that left the program (group *L*). Thus, the impact is estimated by comparing the households before and after they joined the program, or before and after they left the program.

As opposed to the canonical differences-in-differences model with two periods and two groups, the interpretation of these estimates is not simply the difference between the difference from one period to another in the outcome variable of the treatment group and the difference from one period to another in the outcome variable of the control group. For example, consider a case with three periods. In the first period all households are not yet treated. In the second period, one group of households gets treated and the effect is estimated by comparing them to all the other households (i.e., same as in the canonical DD). In the third period, another group of households gets treated and the effect is estimated by comparing them to (1) the group of households that were already treated in the second period, and (2) the group of not-yet-treated households.

This three-period example shows that when multiple groups and periods are considered, the estimates of the treatment effect come from multiple comparisons between treatment and control groups, and these treatment and control groups vary in each period. More generally, Goodman-Bacon (2018) shows that the TWFE estimator when there is variation in when the treatment status changes is a weighted average of all possible two-group/two-period DD estimators in the data. For each of these 2×2 DD estimators, there is one control group whose treatment status does not change and one treatment group whose treatment status changes; and the weight is proportional to group sizes—size of the treatment and control groups—and the variance of the treatment variable. We use the exposition in Goodman-Bacon (2018) to explore and interpret our estimates.¹⁹

Furthermore, we presume the treatment effect changes over time (i.e., it takes time to adjust behavior or make technological changes). This heterogeneity complicates the interpretation of the estimates from any TWFE DD model that relies on households that have already been treated acting as controls. To shed some light on the dynamics of the treatment effect to

¹⁹ We interpret our results following Goodman-Bacon (2018)'s work, but the literature on two-way fixed effects with heterogeneous treatment effects is active; that is, recent advances are yet-to-be-published working papers. Relevant to this work, Athey and Imbens (2018) show that, with staggered adoption, the standard DD estimator is unbiased. But the results rely on random assignment of the adoption date from a design perspective. de Chaisemartin and D'Haultfoeuille (2020) explore the two-way fixed effects model with heterogeneous effects in general, and their special case of DD with staggered adoption is consistent with the exposition in Goodman-Bacon (2018). Baker et al. (2021) provide a relatively gentle overview of the recent literature.

Table 2
Effect of the TVP program on total consumption.

Data from groups:	Join		Left	
	I, J	J	I, L	L
β^{DD}	24.8	67.9	6.8	16.3
Two-sided p -value	0.19	0.025	0.209	0.008
Number of households	6932	106	7381	555
Number of observations	289,917	4571	310,694	25,348

TWFE DD estimated with OLS. Dependent variable is monthly consumption in kWh.

aid the interpretation of our estimates, we use an event-study specification.²⁰ Explicitly, for the households that join (group J), we estimate:

$$y_{it} = \alpha_i + \lambda_t + \sum_{\tau=1}^m \delta_{-\tau} D_{i,t-\tau} + \sum_{\tau=0}^q \gamma_{-\tau} D_{i,t+\tau} + \varepsilon_{it}, \quad (2)$$

where y_{it} is electricity consumption of household i in month t , α_i and λ_t are household and time fixed effects, the first summation term represents m leads or ‘before-treatment’ effects and the second summation represents the effect in the period of the treatment and q lags or ‘during-treatment’ effects, and ε_{it} is the error term of the econometric model. For the households that left (group L), we also estimate this equation, but the leads represent the ‘during-treatment’ effects of the treatment and the lags represent the ‘after-treatment’ effects of treatment.

Finally, because the three-phase meters are only installed when a household is in the TVP program, we cannot observe consumption by time block for households that are not in the program.²¹ Without those data, it is not possible to estimate the effect of the TVP program on peak consumption without very strong assumptions. In [Appendix E](#), we engage in a thought experiment to offer some suggestive evidence of the effect of TVP on peak consumption. That is, we use the fact that the time blocks of consumption (i.e., peak, mid-peak, and off-peak hours) were defined based on historical aggregate consumption data and assume that the consumption of the households in the program was consistent with the definition of the time blocks when they were not in the program. Although our results are only suggestive, they are in line with previous studies finding a decrease in electricity consumption during peak hours. In the results section we focus on the estimates of the effect of the TVP on total consumption because that is the answer to our research question.

4. Results

[Table 2](#) shows the (TWFE DD) estimates of the effect of time-varying pricing (TVP) on total electricity consumption. All estimates are positive, with sizes ranging from 2% (6.8 kWh) to 16% (67.9 kWh) of the average total consumption. However, these estimates are an average that reflects both the individual 2×2 DD estimators that can be constructed from the data and the weights of the TWFE estimator ([Goodman-Bacon, 2018](#)). A detailed discussion follows.

First, we consider the two TWFE estimates of the treatment effect for the households that joined the program between 2011 and 2015 (first two columns in [Table 2](#)). In the first column we present an estimate obtained using data both from households that were only in the program and households that joined the program (groups I and J), and in the second column an estimate obtained using data from households that joined the program (group J).

In the TWFE estimate using only data from households in group J , there are only two types of 2×2 DDs: comparisons between ‘not-yet treated’ households and ‘newly treated’ households, and comparisons between ‘already treated’ households and ‘newly treated’ households. Comparisons between ‘not-yet treated’ households and ‘newly treated’ households are akin to the comparison of pre and post periods in the canonical DD model. Comparisons between ‘already treated’ households and ‘newly treated’ households introduce a source of bias not present in the canonical DD model. Explicitly, when the treatment effect changes over time, the ‘already treated’ households are no longer good controls and comparisons between them and ‘newly treated’ households lead to biased DD estimates. In our case, there seems to be a declining trend of the treatment effect over time (see [Figs. 3, C.1 and C.4](#) in [Appendix C](#)). The bias reduces the DD estimate, because ‘already treated’ households—whose treatment effect declines over time—are compared to ‘newly treated’ households. In addition, the treatment status for most of the households changes within the first 4 periods of the observed data (see [Fig. C.2](#) in [Appendix C](#)). The implication is that most of the 2×2 DDs are estimated using ‘already-treated’ households (with a declining treatment effect over time) as control. Further, the low variance of the treatment variable for the group cohorts in the first

²⁰ [Sun and Abraham \(2020\)](#) show that the event-study specification is appropriate even with time-varying treatment effects if all cohort groups share the same dynamics, which we believe is a reasonable assumption for our context.

²¹ [Jesoe et al. \(2014\)](#) overcame the lack of consumption data by period in their control group by using a complementary dataset of hourly consumption from a small sample of households in the utility. A similar dataset is not available for Costa Rica. [Allcott \(2011\)](#) randomized the TVP pricing over a population in which every household had a meter capable of keeping track of consumption by hour, therefore, he could observe consumption data by hour for both the control and experimental groups.




Fig. 3. Event-study specification for households that joined the program.

four periods leads to a lower weight for the DD estimate than what it is implied by their sample size (see Fig. C.3 in Appendix C).

Including data from households in group *I*—only in the program—changes both the DD terms and the weights of the terms. The DD terms change because the new group adds a new type of 2×2 DD: comparisons between ‘only in’ and ‘newly treated’ households. Thus, this case differs from the exposition in Goodman-Bacon (2018) in that he uses an ‘only out’ instead of an ‘only in’ group. The mechanics of the estimator remain the same—it is just a comparison between households for which treatment status changes and households for which treatment status does not change. However, by using ‘only in’ instead of ‘only out’ households, we again introduce a source of bias when treatment effects vary over time. What can be said about the bias in this case? Given a downward-sloping trend in the treatment effect, it is reasonable to assume that, on average, the treatment effect of the ‘only in’ households is lower than the treatment effect of the ‘already treated’ households. However, it is also reasonable to assume that households have had more time to adjust to the change in pricing schedule and the treatment effect is likely not changing over time as much as immediately after joining the program. As a result, either the treatment effect of the DDs from comparisons with ‘only in’ households is lower or the negative bias from downward sloping treatment effect changes is stronger. In any case, the bias in the TWFE estimate is aggravated when including ‘only in’ households. On the other hand, the weights of the 2×2 DD terms change because, while the variance of the treatment status remains the same, the absolute size of the subsample used in the DD and the relative share of control and treatment groups in each pair changes. As a result, the DDs comparing ‘only in’ households—with a now lower treatment effect—to ‘newly treated’ households get more weight and further reduce the TWFE estimate relative to the estimate obtained using only households from group *J*. Our estimates are likely to be a lower bound of the treatment effect of the program for the households that joined between 2011 and 2015.

Because it is not ideal to summarize a (potentially) time-varying treatment effect with a single coefficient, we try to capture the dynamics of the treatment effect with an event-study specification (see Eq. (2)). The results using data of households that join (group *J*) are summarized in Fig. 3. The point estimates suggest no treatment effect before joining the program (the *p*-value from an *F*-test is 0) and a positive but declining treatment effect after joining the program (the *p*-value from an *F*-test is 0.52). However, we are unable to get precise estimates for most periods. For example, the 90% confidence interval for the treatment effect 6 months before joining the TVP program ranges from about –250 and 250 kWh. This lack of precision is because there are very few households treated 6 months (or later) after the start of the observed period (see Fig. C.2 in Appendix C).

Let us now return to the TWFE estimates for the households that left the program between 2011 and 2015 (Fig. 2, last two columns). In column 3 we present an estimate obtained using data both from households that were only in the program and households that left the program (groups *I* and *L*), and in column 4 an estimate obtained using only data from households that left the program (group *L*). Relative to when households joined the program, the analysis for households that left is flipped. Instead of estimating the treatment effect by using the change in consumption when households join, we estimate the treatment effect by using the change in consumption when households leave. Alternatively, the interpretation is that the change in treatment status goes from ‘join the program’ to ‘leave the program’. That is, we say that when a household changes its treatment status—leaves the program—it becomes ‘newly untreated’.

In the TWFE estimate using only data from households in group *L*, there are only two types of 2×2 DDs: comparisons between ‘not-yet untreated’ households and the ‘newly untreated’ households, and comparisons between ‘already untreated’ households and ‘newly untreated’ households. Unlike the case with data from group *J*—in which the time-varying treatment effect affects households’ consumption *after* they change their treatment status to treated—here the change in treatment status to untreated implies that households’ consumption is affected by the time-varying treatment effect *before* their change in treatment status. Therefore, we get the opposite result from the case when households join the program: the second type




Fig. 4. Event-study specification for households that left the program.

of DDs is akin to the canonical DD model with two periods and the first type of DDs introduces a source of bias when treatment effects vary over time.²²

From the previous case, it seems that the treatment effect declines over time for households that joined (see Figs. 3, C.1 and C.4 in Appendix C). We infer that households ‘only in’ the program likely have a lower treatment effect or that changes—if any—over time are smaller than when households recently joined the program. And data from households that left suggest that the treatment effect increases over time (see Figs. 4, D.1 and D.4 in Appendix D). Therefore, the estimate with only data from group *L* is likely to have a positive bias that inflates the treatment effect. While still reflecting the positive bias from the upward-sloping time-varying treatment effect, including data from group *I* is likely to reduce the effect in two ways. First, the 2×2 DD estimators of comparisons with the ‘always treated’ are likely to be, on average, lower than the average 2×2 DD estimator from the TWFE using only households from group *L*. Second, while the variance of the treatment status remains the same, with the new data, the absolute size of the subsample used in the DDs and the relative share of control and treatment groups in each pair changes. The implication is that the weight of the DDs biased by the upward-sloping time-varying treatment effect is lower than before. Because there are opposite forces at play, the interpretation of the bias in the estimates in this case is ambiguous. The TWFE estimate using only data from group *L* is likely higher than the true effect, but we cannot infer the net result of the opposing effects at play when including the data from group *I*.

Of course, the upward-sloping time-varying treatment effect for households that left might be the reason households left the program (see Figs. 4 and D.1 in Appendix D). In fact, households in group *L* have the highest hourly consumption during peak hours relative to the hourly consumption during mid-peak hours of all groups that self-selected into the program (see Fig. E.1 in Appendix E). They also seem to have the highest hourly consumption during peak hours (see Fig. E.2 in Appendix E), and yet have the lowest total consumption of all groups (see Fig. H.1 in Appendix H). In the absence of direct information from the households that left, the discussion remains largely speculative, but it seems that these households are less able to adjust their consumption during the day to exploit the benefits of the TVP pricing and, given their lower total consumption, pay a higher price as a share of the total bill for their lack of flexibility.²³ Still, one can redefine the treatment as ‘leaving the program’ and, through the lens of a Roy model of selection on gains (Heckman and Vytlačil, 2007), interpret the results as an upper bound of the average effect on total consumption of changing the pricing from TVP to a non-TVP schedule.


Given the previous discussion, the appropriateness of combining data from groups *I* and *L* to get the TWFE estimate is not undisputable. If households that left the program are different than households that were always in the program, combining them would lead to an estimate biased due to self-selection into leaving the program.

Again, we try to capture the dynamics of the treatment effect with an event study specification, but this time using data from group *L*. The results are summarized in Fig. 4. The point estimates suggest a positive treatment effect when households are in the program (the *p*-value from an *F*-test is 0) and no treatment effect when households are not in the program (the *p*-value from an *F*-test is 0.15). As opposed to the case using data from group *J*, here we obtain more precise estimates because there is a similar number of households leaving the program each month (see Figs. D.2 and D.3 in Appendix D). In other words, the sizes of the control and treatment groups are more equal.

²² The analysis becomes more complicated if we do not assume that household leaving the program do not have adjust to the new pricing gradually over time. In other words, we assume that once they leave the program, households immediately get to a new level of consumption, and any further changes are unrelated to leaving the TVP program. This might be a reasonable assumption considering that block pricing is easier to understand and most households should have experience with such pricing schedule.

²³ Based on our calculations of monthly billing, those in group *L* pay a higher average price under TVP (CRC 73.64) than those in group *J* (CRC 69.94) and those in group *I* (CRC 71.23). See more details on the calculation in Appendix F.

Behavioral changes without substitution over time



Behavioral changes with substitution over time




Fig. 5. Behavioral changes in response to TVP.

In summary, given the data constraints, we have not provided unbiased estimates of the effect of a time-varying pricing program on total consumption. However, we argue that we have provided evidence suggesting that, contrary to findings from previous research, the effect of the implementation of TVP on total consumption might not be negative. Instead, it seems that when there is little to no room for technological change, behavioral changes largely determine the net effect on total consumption, and households that self-select into the TVP program may increase their total consumption.

5. Reconciling seemingly conflicting results

While previous empirical research finds that households decrease total electricity consumption in response to Time-Varying Pricing (TVP), we find the opposite. How to reconcile this seemingly mixed evidence? We do not have the data to provide a precisely estimated empirical answer.²⁴ Instead, in this short section, we use a stylized model to show how our findings—an increase in total consumption—may not be at odds with previous findings—a decrease in total consumption.

First, consider a world with no technological changes and a two-period TVP program (i.e., peak and off-peak). A household in this program faces two separate demand curves, one during peak hours and one during off-peak hours.²⁵ Without substitution over time (i.e., cooking late at night instead of during the evening), consumption decreases during peak hours and increases during off-peak hours (Fig. 5, top panel). With substitution over time, the household’s demand curve during peak hours shifts to the left by the amount of consumption that is possible to substitute over time, with a corresponding shift to the right in the demand curve during off-peak hours. Critically, every kWh ‘transferred’ from peak hours to off-peak hours is transformed into more than 1 kWh because of the lower price. For simplicity and because it does not affect the

²⁴ Because of the lack of the necessary data (which plagues much of the research in this field), we use theory to consider one potential mechanism underlying our empirical results (and its consistency with results found by others). Establishing this mechanism empirically would require data on precise actions by households in response to the implementation of a TVP program, and the impact of each action on their electricity consumption. While we would love to have access to and work with such data, collecting such data is practically challenging in the types of settings we study, and potentially very expensive.

²⁵ The price during peak hours is higher than the constant price before the TVP, and the price during the off-peak hours is lower than the constant price before the TVP.




Fig. 6. Only behavioral changes vs behavioral and technological changes.

narrative, we do not show this effect in Fig. 5 nor Fig. 6.²⁶ In both cases, the effect of the TVP on total consumption is simply the net effect $(q_0^{peak} - q_1^{peak}) + (q_1^{off} - q_0^{off})$.²⁷

Second, assume that households respond to an average price²⁸

$$p^{TVP} = p_H w^{peak} + p_L w^{off}, \tag{3}$$

where p_H is the price of the TVP during peak hours, p_L is the price of the TVP during off-peak hours, and w^{peak} and w^{off} are weights describing the relative amount of consumption of the household in each of the two periods, with $w^{peak} + w^{off} = 1$. By revealed preferences (households self-selected into the program), we have that $p^{TVP} < p_0$, where p_0 is the constant price of electricity before the implementation of the TVP program.^{29, 30}

Third, a household responding to p^{TVP} with only behavioral changes increases total consumption simply because $p^{TVP} < p_0$ (Fig. 6, left panel). Instead, a household responding to p^{TVP} with both technological and behavioral changes may either reduce or increase total consumption (Fig. 6, right panel). Here, the technological changes shift the demand curve to the left (i.e., because of a new fridge that uses less electricity all day). Along the new demand curve (D_1), households increase consumption because $p^{TVP} < p_0$. Whether the net effect is an increase or decrease depends on the relative size of the effect coming from technological changes and behavioral changes. We reconcile the seemingly mixed evidence by noting that our results come from a setting with almost no room for technological changes (Fig. 6, left panel), and previous results come from a setting in which there is plenty room for technological changes (Fig. 6, right panel).

6. Conclusion

Empirical research has shown that using time-varying pricing (TVP) helps to solve the allocative inefficiency in the residential electricity sector (Allcott, 2011; Wolak, 2011; Jessoe et al., 2014). Explicitly, researchers consistently find that TVP schedules reduce both total consumption and consumption during peak hours. However, none of these studies have explicitly addressed the separate effects of technological and behavioral changes in consumption when a TVP schedule is implemented.

²⁶ Specifically, to show such effect, both figures should show a new demand curve in off-peak hours that is pivoted (with a smaller slope) to reflect that each kWh transferred from peak hours to off-peak hours is cheaper. This effect makes it more likely to obtain our finding of an increase in total consumption, but it is not critical to show how the mechanism underlying our finding (of an increase in total consumption) might be consistent with the mechanism underlying previous findings (of a decrease in total consumption).

²⁷ Here we discuss only the consequences of a change in price. However, there are alternative mechanisms that could lead to changes in the amount of electricity consumed when activities that use electricity are shifted from peak hours to off peak hours. For example, in response to the TVP pricing, one could take a bath late at night instead of in the morning. Without the need to leave home to work, the bath might be longer (resulting in more electricity use) simply because now there is more time to enjoy it (beyond the increased use that can be explained by the lower price per kWh). We thank Jo Albers for bringing up the potential role of mechanisms beyond the response to lower prices.

²⁸ Ito (2014), for example, show that households do not seem to respond to marginal electricity prices. Instead, they respond to some average price.

²⁹ Recall that this stylized model is appropriate for our setting with a voluntary TVP program. The relationship between average prices might not be the same for a mandatory TVP program.

³⁰ We can explain the existence of people sophisticated enough to choose a TVP schedule but not sophisticated enough to make day-to-day electricity consumption choices based on marginal prices in at least two ways. First, they might be choosing the TVP under the influence of some optimism bias. Second, they might be rationally inattentive; it is worth to spend the time thinking about which pricing schedule to choose, but it might not be worth to spend too much time thinking about everyday electricity consumption choices such as turning off a lightbulb. We thank Sarah Jacobson for raising this issue.

To partially isolate the effect on total consumption of behavioral changes in response to TVP, we study a voluntary program that introduced a TVP schedule to households in San José, the capital of Costa Rica. We chose San José because it has a mild climate and very few households use heating or cooling devices. With little room for technological changes (relative to a rich country), consumption changes are largely driven by behavioral changes. Contrary to previous research, we find that TVP may increase total electricity consumption. However, we argue that this seemingly conflicting result may not be at odds with previous findings. Instead, we show how the difference between our result and previous findings can be explained by the differences in the context of the studies.

Despite variation in the timing of the treatment and a potential time-varying treatment effect, we are able to present suggestive evidence that households may increase consumption in response to TVP. To do so, we exploit recent econometric developments on the understanding of the two-way fixed effects differences-in-differences estimator (Goodman-Bacon, 2018) along with event-study specifications to interpret our results. Given data limitations and the general issues with capturing a potentially dynamic effect with one coefficient, we believe that the recent econometric developments expand the scope of what is possible to learn using a differences-in-differences research design.

We also emphasize the need for better data to study time-varying prices in alternative contexts. While in this paper we work with what was reasonable to collect, we had to rely on somewhat strong assumptions to provide suggestive evidence of the effects of implementing a TVP schedule (e.g., our discussion of biases in the TWFE DD estimates). In contrast, most of the previous studies have relied on rich data from large randomized control trials (Faruqui and Sergici, 2010; Badtke-Berkow et al., 2015). Researchers and utilities in low- or middle-income countries (LMIC) might have to work together to generate the data required for a comprehensive analysis of TVP in their own context. There are still open questions that can only be answered with better data. For example, (Ito, 2014) shows that households do not respond to marginal prices of electricity; instead, they respond to *some* average price. In our context, we find suggestive evidence along the same lines (see Fig. G.1 in Appendix G), but little can be said about the average prices to which households respond.

Finally, our research could help policy makers better anticipate the likely effects of implementing time-varying prices in different contexts. For example, when implementing a TVP schedule in a context with low usage of heating and cooling devices, policy makers would know not to expect reductions in consumption during peak hours as large as those found by previous research. Further, our finding of an increase in total electricity consumption due to the TVP schedule challenges previous research on the short-run and long-run expected outcomes of TVP (Borenstein, 2005). In addition, we are not aware of any previous study on the effect of TVP in a LMIC setting. Our hope is that policy makers in contexts closer to ours than to a rich country will be able to use this study to motivate policy making using empirical evidence relevant to their own context. Specifically, our results serve as a cautionary piece of evidence for policy makers interested in reducing consumption during peak hours with TVP—the goal can potentially be achieved, but the cost is increased total consumption. Therefore, policy makers should weight both policy targets carefully before opting for TVP.

Data

We obtained consumption data directly from the *Compañía Nacional de Fuerza y Luz* (CNFL), an electric utility in Costa Rica. As requested by the CNFL, we signed up a confidentiality agreement (available upon request) that prevents us from sharing the data used in this research. We direct interested researchers to contact the CNFL if interested in the same or similar data (www.cnfl.go.cr). All the code used in the analysis is available in <https://osf.io/jrmaq/>.

Funding

Economic support from the Swedish Agency for International Development Cooperation (SIDA) through the Environment for Development (EfD) initiative is gratefully acknowledged

Conflict of interest

None declared.

Acknowledgements

Our work was improved thanks to useful comments by David Aadland, Jo Albers, Hunt Allcott, Sara Capitán, Benjamin Gilbert, Sarah Jacobson, Marc Jeuland, Esteban Méndez-Chacón, Alexandre Skiba, Linda Thunström, Klaas van't Veld, and Donald Waldman.

Our work also reflects improvements in response to comments by participants at the Third Annual Meeting of the Sustainable Energy Transitions Initiative, the 2018 Colorado State University and University of Wyoming joint Graduate Student Symposium, the 20th Colorado University at Boulder Environmental and Resource Economics Workshop, and the 2020 Online Summer Workshop in Environment, Energy, and Transportation (Economics).

Finally, we thank the *Compañía Nacional de Fuerza y Luz* for sharing the data used in this research.

Appendix A. Average prices under alternative pricing

In Fig. A.1 we use a Marschak-Machina triangle to show all possible consumption profiles. The average price a consumer pays for electricity under TVP depends on the share of consumption in each of the three relevant periods (peak, mid-peak, and off-peak). The y -axis in the figure represents the share of consumption during peak hours (p), the x -axis represents the share of consumption during mid-peak hours (m), and the residual $1 - p - m$ represents the share of consumption during off-peak hours. For example, the xy -coordinate (0.40, 0.22), marked by a black dot, represents the average consumption profile of households in the TVP of our study. On average the share of electricity used in peak, mid-peak, and off-peak hours is 22%, 40%, and 38%. The shaded area (top of the figure) shows all the consumption profiles in which the average price under block pricing is higher than the average price under TVP. Further, the color green at the bottom of the graph (and its intensity) indicates low average prices under TVP and the color red at the top of the graph (and its intensity) indicate high average prices under TVP.³¹

From Fig. A.1 we learn that (i) the average pricing under TVP is not always lower than the average price under block pricing, (ii) there is significant variation in the average price a household would pay under TVP depending on their consumption profile, and (iii) households from our study have an average consumption profile for which the average price under TVP is lower than the average price under block pricing. As discussed in the introduction, we believe that because the TVP program is voluntary (as most TVP programs around the world), the location of the black dot in the non-shaded area (where TVP price < block price) is likely because of self-selection on future gains. In other words, we expect the first households to voluntarily join a TVP program to be precisely the households that stand to gain the most. It is still an open question, however, whether these households stand to gain because their consumption profile was already in the non-shaded area or because these households were able to adjust their consumption.³²




Fig. A.1. Average TVP prices for different consumption profiles. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

³¹ To make this graph we used the actual prices from our study, both under block pricing and under TVP. We normalized total consumption to 100 and calculated the price a household would pay under block pricing (i.e., a constant price based on the total consumption) and the price a household would pay under TVP under all possible consumption profiles (i.e., the price changes depending on when the household uses the electricity).

³² Analyzing this would require better data to observe consumption during the day before and after the TVP program. We do not have this data, and it is extremely difficult to acquire or rare to find such data in settings we study.

Appendix B. Self-selection into the TVP program




Fig. B.1. Consumption over time by type of contract.

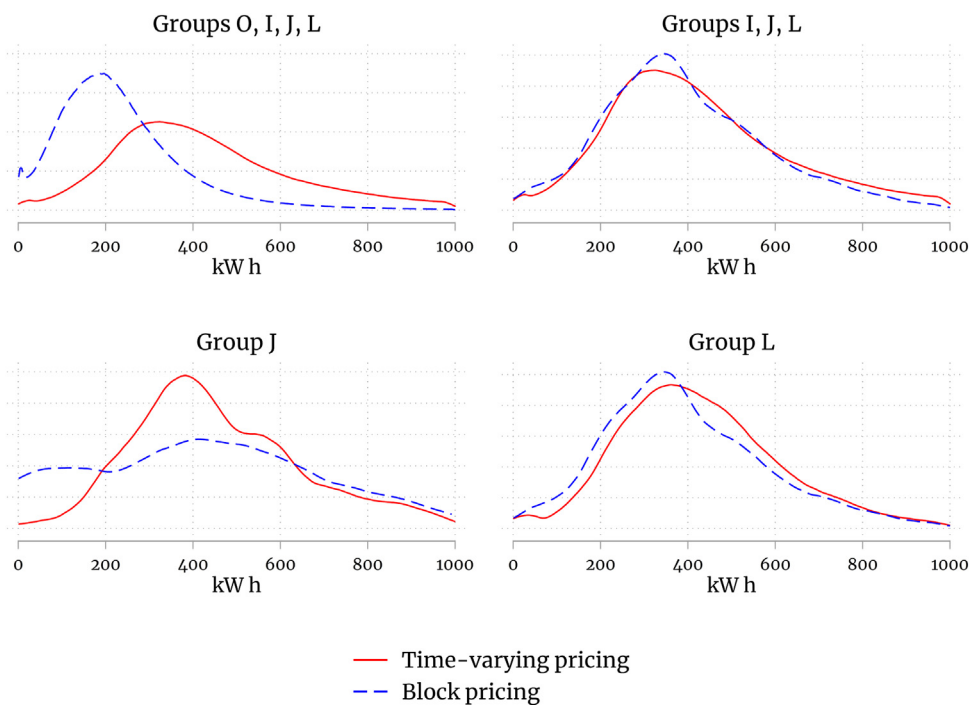


Fig. B.2. Consumption by group and treatment status. Epanechnikov kernel densities comparing the consumption of different groups of households when they are in the TVP program and when they are not in the TVP program.

Appendix C. Group J: Households that joined between 2011 and 2015




Fig. C.1. Consumption trend after joining the program.




Fig. C.2. Number of households joining in each period.




Fig. C.3. Number of observations in each relative period.




Fig. C.4. Demeaned consumption by relative period.

Appendix D. Group L: Households that left between 2011 and 2015




Fig. D.1. Consumption trend before leaving the program.




Fig. D.2. Number of households leaving in each period.




Fig. D.3. Number of observations in each relative period. The dip in the middle is explained by observations we dropped because the consumption data only covered a portion of the month.




Fig. D.4. Demeaned consumption by relative period.

Appendix E. Consumption during peak hours

Estimating the effect of Time-Varying Pricing (TVP) on consumption during peak hours has been the focus of previous research. Here we provide suggestive evidence consistent with previous results: TVP is effective at reducing consumption




Fig. E.1. Disaggregated consumption by type of household.

during peak hours. Because we do not observe consumption data disaggregated by time of consumption when a household is not in the TVP program, we rely on a very strong assumption to provide graphical evidence of the effect of TVP on peak consumption. Explicitly, we use the fact that the time blocks (peak, mid-peak, and off-peak hours) were defined based on historical aggregate consumption data, and assume that the disaggregated consumption of household that were in the program for at least one period- groups *I* (only in the program), *J* (joined the program), and *L* (left the program) was consistent with the definition of the time blocks *when they were not in the program*. Since we assume the counterfactual, we do not have data to compare the actual disaggregated consumption. Instead, we rely on graphical evidence showing the disaggregated consumption over time by each type of household. In practice, under our identifying assumption, we say that if the consumption per hour during the peak-hour time block is not the highest, the program was successful in reducing consumption during the peak hours. This result holds unless the differences in consumption patterns between the general population and the self-selected group that participated in the TVP program are such that the self-selected groups do not consume the most during peak hours when not in the program.

Fig. E.1 shows our graphical evidence. In the top panel, for households that were always in the program, the hourly consumption during mid-peak hours is the highest, closely followed by the hourly consumption during peak hours, and the consumption during off-peak hours is the lowest. Also, there seems to be a seasonal pattern around the end and beginning of the year in which hourly consumption during peak hours is higher than during mid-peak hours. Overall, the consumption for all time blocks for this group of households is smooth, which might imply that households have been in the program for enough time and are not adjusting their consumption any further in response to the earlier change of pricing schedule. We observe the same patterns for households that joined during 2011 and 2015, but hourly consumption is not as smooth as in the previous case. Still, the gap from mid and peak hourly consumption is increasing in a way that might eventually converge to the patterns observed for the households that were always in the program. Finally, for the households that left the program between 2011 and 2015, it is hard to distinguish between the hourly consumption during peak and mid-peak hours. Under our identifying assumption, we conclude that the graphical evidence is suggestive of success of the policy in reducing consumption during peak hours.

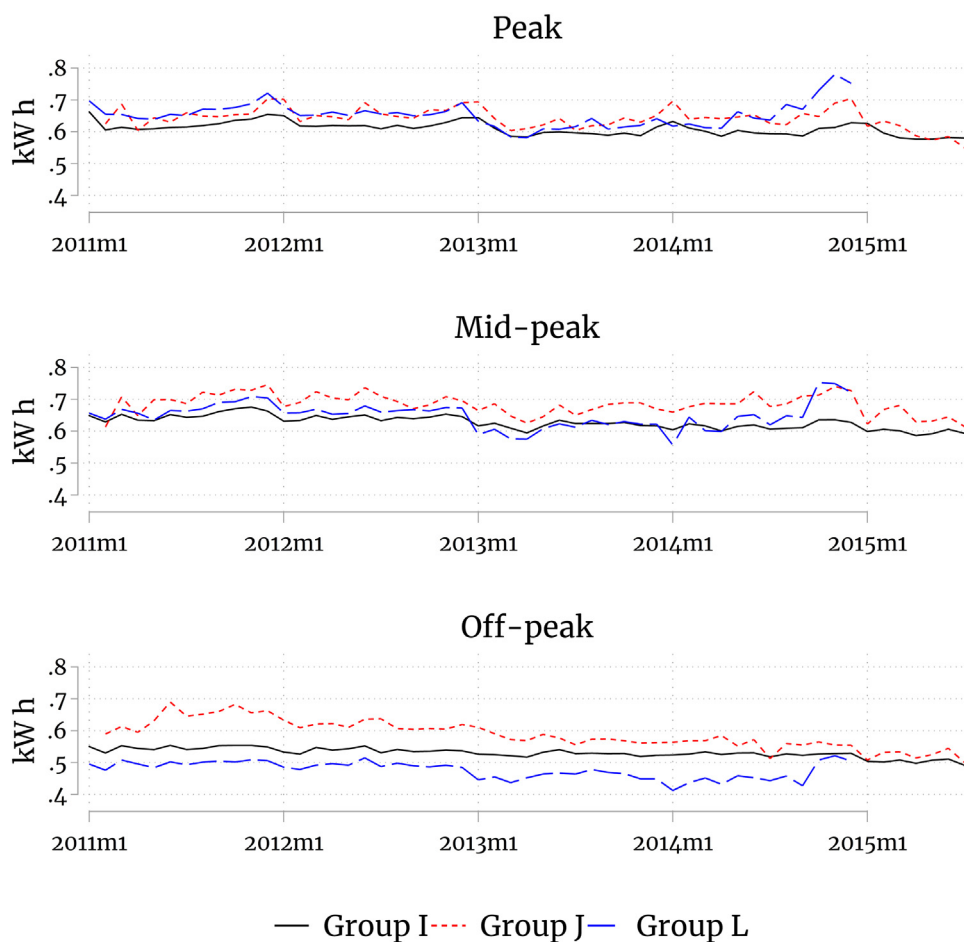


Fig. E.2. Consumption over time by time block and type of contract.

Appendix F. Selection on gains

Households in the TVP program used 22.28%, 40.35%, and 37.37% of their electricity in peak, mid-peak, and off-peak hours. Given this load profile, the average price per kWh under TVP is CRC 71.29, but it would have been CRC 92.36 under block pricing.³³ In other words, the average price under TVP is much lower than the hypothetical average price that households in the TVP program would have paid had they been under block pricing.

We argue that our estimates must be interpreted as a local treatment effect that is “helpful to understand the early phases of these potential voluntary programs”. Since most TVP programs in the residential sector are voluntary, the scope of previous research on the topic is similar to ours. Our key contribution is to shed some light on the mechanism underlying the effect of TVP in total consumption. Specifically, we argue that the lower average price under TVP reduces total consumption both in our case and in previous work (a change along the demand curve). However, previous work has been conducted in settings in which technological changes in response to the TVP are much more likely to occur than in our case, which might lead to an inward shift of the demand curve. In the presence of a shift in the demand curve, households might consume less even if the new average price is lower.

Overall, this selection on gains creates large savings for those households that can take advantage of the TVP pricing (structural winners). In fact, households in the TVP program saved money (relative to counterfactual billing under block pricing) in about 93% of the 315,265 monthly bills we calculated (see the distribution of the difference between billing under TVP and billing under block pricing in Fig. F.1. And these savings are economically significant. On average, each household

³³ We do not have billing data. Instead, we calculated monthly billing based on our consumption data. Specifically, we calculated monthly billing data based on the pricing described in Table 1.




Fig. F.1. Distribution of the difference between billing under TVP and billing under block pricing.

saved CRC 148,836 per year relative to counterfactual billing under block pricing. Perhaps more striking, the aggregate savings for the households under TVP were CRC 3,910,537,980. These aggregate savings can also be interpreted as potential income lost by the electric utility. Our work shows the potential for significant losses if the benefits from shifting some consumption away from peak hours are not enough to offset the lost income from structural winners joining the TVP program.

Appendix G. Bunching




Fig. G.1. Bunching. This figure reproduces the key figure in Ito (2014). It shows evidence in support of the hypothesis that households do not respond to marginal prices, because there is no bunching to the left of increases in the price per kWh.

Appendix H. Consumption by group




Fig. H.1. Total consumption over time by type of households.

Appendix I. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.reseneeco.2021.101264>.

References

- Allcott, H., 2011. Rethinking real-time electricity pricing. *Resour. Energy Econ.* 33 (4), 820–842.
- Athey, S., Imbens, G.W., 2018. Design-Based Analysis in Difference-In-Differences Settings With Staggered Adoption. Technical Report w24963. National Bureau of Economic Research.
- Badtke-Berkow, M., Centore, M., Mohlin, K., Spiller, B., 2015. A primer on time-variant electricity pricing. *Environ. Defense Fund*.
- Baker, A., Larcker, D.F., Wang, C.C.Y., 2021. How Much Should We Trust Staggered Difference-In-Differences Estimates?
- Borenstein, Severin, 2005. The long-run efficiency of real-time electricity pricing. *Energy J.* 26 (3).
- Borenstein, S., 2007. Wealth transfers among large customers from implementing real-time retail electricity pricing. *Energy J.* 28 (2).
- de Chaisemartin, C., D'Haultfœuille, X., 2020. Two-Way Fixed Effects Estimators With Heterogeneous Treatment Effects. arXiv:1803.08807 [econ].
- Dutta, G., Mitra, K., 2017. A literature review on dynamic pricing of electricity. *J. Oper. Res. Soc.* 68 (10), 1131–1145.
- Faruqui, A., Sergici, S., 2010. Household response to dynamic pricing of electricity: a survey of 15 experiments. *J. Regul. Econ.* 38 (2), 193–225.
- Goodman-Bacon, A., 2018. Difference-in-Differences With Variation in Treatment Timing. Technical Report w25018. National Bureau of Economic Research.
- Heckman, J.J., Vytlačil, E.J., 2007. Econometric evaluation of social programs. Part I: causal models, structural models and econometric policy evaluation. In: Heckman, J.J., Leamer, E.E. (Eds.), *Handbook of Econometrics*, Vol. 6. Elsevier, pp. 4779–4874 (Chapter 70).
- Ito, K., 2014. Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *Am. Econ. Rev.* 104 (2), 537–563.
- Ito, K., Ida, T., Tanaka, M., 2021. Selection on Welfare Gains: Experimental Evidence From Electricity Plan Choice. Technical Report w28413. National Bureau of Economic Research.
- Jessoe, K., Rapson, D., Smith, J.B., 2014. Towards understanding the role of price in residential electricity choices: evidence from a natural experiment. *J. Econ. Behav. Organ.* 107, 191–208.
- Badtke-Berkow, Mina, Centore, Michael, Mohlin, Kristina, Spiller, Beia, 2015. A Primer on Time-Variant Electricity Pricing.
- Ministerio de Ambiente y Energía, 2012. Encuesta de consumo energético nacional en el sector residencial de Costa Rica.
- Ministerio de Ambiente y Energía, 2019. Estudio para la caracterización del consumo energético en el sector residencial.
- Sun, L., Abraham, S., 2020. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *J. Econometr.*
- Wolak, F.A., 2011. Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment. *Am. Econ. Rev.* 101 (3), 83–87.