Disentangling Inattention to (Energy) Prices from Uncertainty in a Digital World

Jacquelyn Pless[1*]

Last updated: September 2022

Abstract

Digital technologies are enabling increasingly high time resolution in dynamic pricing across many economic settings. This paper studies whether consumers respond to real-time prices and disentangles inattention from price uncertainty using a novel identification strategy and high-frequency electricity flow data from households in the United Kingdom. Exploiting one-minute level fluctuations in household solar energy generation that create exogenous within household-day-hour price variation, I find that demand is very inelastic with respect to real-time prices, and perhaps surprisingly, consumption is increasing in price uncertainty. Access to energy use data through in-home displays readjusts the uncertainty effect but does not change responsiveness to price levels. The findings indicate that consumers are forward-looking and respond to *expected average* prices rather than marginal prices, raising questions about the potential for real-time pricing and behavioral interventions to reduce or shift electricity consumption timing.

Keywords: digitalization; dynamic pricing; uncertainty; inattention; electricity demand

*For helpful comments and feedback, I thank Stephanie Brockman, Tamma Carleton, Harrison Fell, Cameron Hepburn, Katrina Jessoe, Erica Myers, Emily Pakhtigian, David Rapson, Alex Teytelboym, and participants at the 2018 and 2020 Association of Environmental and Resource Economists Annual Meetings, the 13th International Workshop on Empirical Methods in Energy Economics 2020 (ETH Zurich), the 2020 Workshop in Environmental Economics (London School of Economics, Grantham Research Institute on Climate Change and the Environment, and Imperial College London), and members of the Oxford Martin Programme on Integrating Renewable Energy. I especially thank Eoghan McKenna of University College London for assistance with data gathering during the early stages of the project and Austin Zaccor for excellent research assistance. I am also grateful for financial support from the Oxford Martin School. The author remains fully responsible for the paper's contents and any mistakes herein.

[1]MIT Sloan School of Management, 100 Main Street, Cambridge, MA, 02142. Email: jpless@mit.edu.

1 Introduction

The widespread proliferation of digital technologies for buying and selling goods and services is enabling the use of dynamic pricing in many industries and with increasingly granular time resolution. Some forms of dynamic pricing have a long history—like seasonal and time-of-use pricing in tourism—but those that fluctuate rapidly according to market conditions have emerged more recently as modern advancements in digitalization and "big data" are becoming ubiquitous. Prices now can adjust in real time, such as in transportation (e.g., Uber), hospitality (e.g., Airbnb), and e-commerce (e.g., Amazon). In theory, prices changing along with shifts in supply and demand should improve efficiency. But this assumes that consumers respond optimally, and behavior often deviates from the economist's ideal. For instance, consumers may mis-optimize when prices and taxes are not salient (Chetty, Looney and Kroft 2009; Finkelstein 2009; Bradley and Feldman 2020) or when input quantities are complex (Jessoe and Rapson 2014).[1] Such behavioral biases are generally considered to be different forms of inattention (Gabaix 2019).

When prices fluctuate so rapidly and unpredictably, though, and consumers make decisions based on their expectations about future prices, uncertainty also plays an important role. Uncertainty is sometimes implicitly conflated with inattention but it is fundamentally different, as they come with different economic foundations and lead to different conclusions for both policy and management. The distinction has three key implications. First, inattention implies a (rational or irrational) behavioral bias whereas errors due to price uncertainty suggest that consumers form imperfect expectations about future prices. Second, consumers respond to *expected prices* under uncertainty (Borenstein 2009) and sometimes *average* prices when prices are complex (Ito 2014), whereas attributing optimization errors to inattention assumes that consumers respond to marginal prices. Knowing the price to which consumers respond is important for accurately estimating the demand elasticity. Third, the policy and management implications are different. Interventions that increase price transparency might work when inattention is the driver of "mistakes." But if it is price uncertainty, policymakers and managers may want to focus on reducing such uncertainty by smoothing prices, or targeting the same objectives using tools beyond behavioral interventions.

In this paper, I study how households respond to real-time electricity prices using a novel identification strategy, and I disentangle whether inattention or uncertainty is the culprit of the low

[1]Additional examples include Busse, Silva-Risso and Zettelmeyer (2006) (rebates for car purchases), Goldin and Homonoff (2013) (cigarette taxes), Feldman, Goldin and Homonoff (2018) (sales and excise taxes), Bradley (2017) (property taxes), and Chetty and Saez. (2013) and Feldman, Katuscak and Kawano (2016) (tax credits).

demand elasticity that I find. I use one-minute level electricity consumption and solar production data from households in the United Kingdom that have solar panels, exploiting random within-household-hour fluctuations in solar generation that create exogenous variation in prices. Increases in solar generation essentially lead the marginal price of electricity to drop to zero. The high frequency nature of my data allows me to estimate the price elasticity of demand with respect to price levels (i.e., solar generation at a within-household level with granular time resolution) and to construct various measures of price volatility within-hour as well as over the course of the day to capture uncertainty.

Responding to prices in real time requires paying attention at every moment, which is cognitively taxing, so consumers may rationally choose to not respond to small and rapid price changes. At the same time, changes in real-time prices driven by solar are almost perfectly salient—prices basically drop to zero (or at least a price much lower than the prevailing retail rate) when the sun is shining as long as the household's current consumption does not exceed that which their solar system is generating. Knowing that prices are lower therefore does not require devoting a significant amount of attention. On the other hand, changes in prices can be very volatile, and households must form expectations about future prices under uncertainty to inform their decisions. Responses both in a current period and in future periods are driven by expectations.

I conduct three main sets of analyses to disentangle these underlying mechanisms. First, I test how households respond to various measures of price volatility (the standard deviation and range of solar generation within-hour and within-day) as proxies for uncertainty. Second, to explore whether households plan their consumption based on price expectations, I examine whether they shift their consumption timing based on the average prices and average price volatility throughout the days. Theory would suggest that both gross consumption and shifting should decrease with uncertainty. Third, I test whether in-home displays (IHDs) that provide clear price signals change responses to either price levels or price volatility. If IHDs adjust how consumers respond to price levels in real time, consumers may be inattentive. On the other hand, if the degree to which they adjust consumption shifting changes—particularly in response to price volatility throughout the day—uncertainty is likely at play, as the additional information might be improving their price forecasts and planning.

While understanding how consumers respond to real-time prices is increasingly important in many economic settings, it is especially important in the energy context. It has long been known that static electricity pricing creates inefficiencies (Borenstein 2005; Borenstein and Holland 2005;

Joskow and Wolfram 2012; Puller and West 2013; Borenstein and Bushnell 2019). Renewable electricity can play a major role in reducing greenhouse gas emissions, but its variability generates frequent changes in the marginal cost of production. Dynamic prices might help balance supply and demand and metering technology that enables consumers to respond is also more a ordable and accessible than ever. At the same time, electricity demand is notoriously inelastic in the short run (Reiss and White 2005).[2] Some papers have also found that electricity consumers often make \optimization errors." These realities raise questions about the potential for RTP at the retail level for improving welfare, but studying this is often di cult because true RTP is rarely found in practice.[3]

Using rapid and frequent uctuations in prices driven by solar generation allows me to overcome this barrier. I obtained household electricity consumption and solar generation high-frequency data from the largest smart grid demonstration project in the United Kingdom to date, the Community Led Network Revolution (CLNR) eld trial. One-minute level data on power ows were gathered for 300 households with solar panels between 2012 and 2014. While the number of households is relatively low, there is great depth in the data given the time resolution, as it contains more than 300 million observations and provides important identi cation bene ts. By exploiting variation within household-day-hour in some cases and household-day variation in others, I can exibly control for unobservable di erences in behavior and other drivers of consumption pattern di erences across households, as well as how households may be changing di erentially over time, with household-day xed e ects. About half of the households were given in-home displays (IHDs) that provide real-time information. The IHDs also included a lighting system that glows green when households are exporting solar and red when they are importing electricity from the supplier, which draws attention to prices and reduces complexity associated with calculating them.[4]

My central result is that electricity demand is very inelastic with respect to real-time prices, and while IHDs do not increase attention, they correct mistakes that consumers make due to uncertainty. The price elasticity ranges from -0.022 to -0.057 depending on the speci cation. I also nd a small load-shifting e ect|peak gross consumption decreases by 1.3% for every one kWh increase in daytime solar generation. However, consumption increases in uncertainty, o setting a

[2]Demand is more elastic in the long run but it can still take several years before consumers respond to price increases (Deryugina, MacKay and Reif 2020).

[3]Hourly pricing and critical peak pricing are often referred to as RTP, but prices do not change on the spot in these cases as real-time prices do. Studies of hourly pricing so far also typically provide information in advance.

[4]Sub-optimal responses due to complexity are often rational given the information acquisition costs but are still a form of inattention.

third of the load-shifting that does occur. Real-time feedback mitigates the error by 80 percent but does not affect responses to price levels at all. This indicates that IHDs do not increase attention and that the effects operate entirely through the uncertainty channel. I corroborate the results with a machine learning approach, whereby I match households based on their consumption profiles and build household-level prediction models to form a new control group.

I conduct a series of tests to probe further. First, I show that all of the effects when using 15-minute and hourly data occur only on weekend mornings whereas load-shifting effects occur only on weekdays. This suggests that consumers respond to expected average prices, as they are not responding when they are most likely home. Uncertainty mistakes are also not made in real time. They occur on weekdays only and are driven by households over-estimating solar generation when scheduling appliances in advance to run later in the day. I show that this is not because consumers are myopic or present-biased—there are no effects of the previous day's prices or morning prices. Rather, consumers appear to be (rationally) forward-looking, perhaps consulting weather forecasts to inform their decisions. Such errors are only made when there is an upswing in the day's total solar generation relative to the previous day's, suggesting that errors are due to excitement.

If consumers do not respond to marginal prices and they do not engage with their IHDs in real time, it is surprising that IHDs correct most of the uncertainty error, since they do not provide information on upcoming price changes. I investigate this and provide evidence that households likely use the aggregated data collected over the course of the day or a few weeks to learn, either forming new habits or investing in energy-saving durable goods. Households with IHDs adjust quickly, but those without them also learn, albeit at a slower pace (within about three months). I interpret these findings as suggesting that households learn how to better take advantage of lower-priced power from bills but IHDs facilitate slightly faster learning.

This paper's primary contribution is to the economics of digitalization literature, and especially to the research studying dynamic pricing in the "sharing economy" (Cohen, Hahn, Hall, Levitt and Metcalfe 2016; Cachon, Daniels and Lobel 2017; Hall, Horton and Knoepfle 2020). There is also a growing body of work exploring supply and demand in peer-to-peer markets (Einav, Farronato and Levin 2016; Cullen and Farronato 2020) and the effects of peer production on incumbents (Seamans and Zhu 2014; Farronato and Fradkin 2018). While dynamic pricing should, in theory, improve efficiency, the findings from this paper suggest that there may be limits to the benefits because consumers do not perfectly optimize when prices are extremely volatile.

With this in mind, this paper also contributes to the behavioral economics literature, demonstrating the importance of disentangling inattention from uncertainty. This applies in many familiar settings, as real-time pricing is increasingly commonplace. Taxes and subsidies also change over time and can be nonlinear depending on opaque factors. This work is most closely related to the behavioral papers examining how information provision affects electricity demand, such as with IHDs (Faruqui and Sergici 2011; Jessoe and Rapson 2014; Harding and Lamarche 2016; Bollinger and Hartman 2020; Prest 2020; Bollinger and Hartman 2020) and more frequent or automatic billing (Gilbert and Graff Zivin 2014; Sexton 2015).[5] A common explanation for information effects is that consumers do not pay attention.[6] My findings highlight the importance of distinguishing inattention from uncertainty. This paper also contributes to the related wider literature on non-pecuniary strategies aiming to induce energy conservation and purchases of energy efficient durable goods (Allcott 2011b; Allcott and Rogers 2014; Allcott and Taubinsky 2015; Ito, Ida and Tanaka 2018; Myers 2019; Myers and Souza 2020).

Third, this paper contributes to the literature studying how consumers behave under different electricity pricing regimes, such as block pricing (Ito 2014), time-of-use pricing (Harding and Lamarche 2016), critical peak pricing (Jessoe and Rapson 2014; Faruqui and Sergici 2011), and hourly pricing (Taylor, Schwarz and Cochell 2005; Wolak 2011; Allcott 2011a; Gillan 2018; Fabra, Rapson, Reguant and Wang forthcoming). To my knowledge, this paper is the first to study prices that truly fluctuate in real time, unexpectedly, and without information in advance. [7]

Lastly, the findings have important implications for policy and energy management, as they highlight the limitations of RTP and behavioral interventions for reducing energy consumption or shifting its timing. Policymakers and energy managers may want to consider alternative mechanisms, such as automation, or support innovation in technologies like energy storage that could enable all consumption to be met with clean electricity.[8] This is not to say that dynamic pricing has no benefits, but rather that it is likely insufficient for achieving meaningful change.

The paper proceeds as follows. Section 2 provides background on the institutional context and household decision-making incentives. Section 3 details the data and empirical strategies. I present the main findings in Section 4 and Section 5, and the paper concludes in Section 6.

[5] These questions are also studied in the water (Wichman 2017) and fuel cost contexts (Busse et al. 2006; Allcott 2013; Sallee, West and Fan 2016; Allcott and Knittel 2019).

[6] One exception is Jessoe and Rapson (2014), who rule out imperfect salience. Bollinger and Hartman (2020) also diverges from the literature by showing that automation achieves greater demand response.

[7] The two closests papers to ours are Fabra et al. (forthcoming) and Gillan (2018). However, in both studies, prices change hourly (rather than frequently within-hour) and either advance information or notifications are provided.

[8] Even automation has its limits. While it has been shown to be more effective than IHDs (Bollinger and Hartman 2020), it does not resolve the challenges associated with lighting and electronic use in the evenings.

2 Background and Institutional Details

2.1 The CLNR Smart Grid Demonstration Project

I obtained data from the Customer-Led Network Revolution (CLNR), the United Kingdom's largest socio-technical smart grid demonstration to date. The major £54 million collaboration was headed by Northern Powergrid in partnership with British Gas, Durham University, Newcastle University, and EA Technology Limited and implemented in the Northeast region of England. Their objective was to gather and analyze high-frequency electric power monitoring data to develop a better understanding of how UK households consume electricity, with a particular focus on their use of onsite solar generation and the role of information provision (Sidebotham and Northern Power Grid, 2015). The team also conducted follow-up surveys to ask participants more about how and why they used the technologies or information that they were provided and how they thought it informed their decision-making.

Households were recruited for the field trial by British Gas. The only requirement for participation in the solar-specific trials was that solar PV systems had been already installed and that they owned their solar systems.[9] Households were approached and offered a smart meter to participate, which tracks electricity consumption as well as exports (of solar generation) and imports (of power provided by the supplier when solar generation does not meet time-coincident demand). To provide an additional incentive for participation, they were offered £50 vouchers upon joining and finishing the trial. All homes were fitted with smart meters that track consumption and import/export meters that track solar generation. These are usually placed in a utility cupboard or garage, and whether a household is currently importing electricity from the grid or exporting their zero-priced solar electricity must be calculated manually if they are to use this information.

Half of the participants were also given enhanced in-home displays (IHDs) that send transparent price signals. The devices monitor solar generation and electricity consumption in real time. While providing data in this format already greatly reduces information acquisition costs, the IHDs also use a lighting system whereby a light glows green when the household is exporting their self-generated solar (and the price is zero) and red when the household is importing electricity from their supplier (and the price is equal to their typical retail rate). One can imagine that the IHDs affect decision-making depending on how consumers use them. If the IHD remains in sight while

[9] A number of those originally recruited were found to be part of a \rent a roof" schemes, such that they rented/leased their systems. These households were discarded and additional participants were recruited.

household members are home, the lighting system could draw their attention and motivate them to turn o lights when importing or turn on the kettle when exporting. If the household does not engage with the IHD in real time but rather examines the data collected at the end of the day, they could learn from their experiences and adjust over time, such as by changing their behavior or purchasing new electricity-consuming appliances that have delay timers allowing them to schedule appliance run times in advance. Although the eld trial implementation team aimed to randomly distribute the IHDs, complete randomization was not achieved.[10] Those with IHDs appear to consume slightly more electricity and generate more solar on average, and we describe how we address this in the next section.

2.2 Electricity Consumption Decision-Making Incentives

Electricity prices for households with solar panels depend on their self-generated solar, how their consumption aligns with such generation, and the electricity rate tari structures that they face. Absent energy storage, consumption must be precisely time-coincident with self-generated solar in order to consume it, and households otherwise import power from the incumbent supplier at the retail rate that they typically pay for electricity. They also export their excess generation to the grid. When households are compensated for each unit of solar that they export via a feed-in tari (FiT), the marginal price of consuming ones own solar electricity is the di erence between the retail rate that they otherwise would pay and the \export rate" (FiT). [11]

In this setting, households have a nancial incentive to maximize solar self-consumption. The UK uses net metering and FiTs but the retail price exceeds the compensation received for each kWh of solar power exported to the grid for all households during the entire sample period. The electricity prices are also at during the day. Households would not face this same incentive structure if net metering was based on aggregated power ows or if the export rate exceeded the import rate. La Nauze (2019) discusses some of these more complex incentive structures in detail.

There are various ways in which consumers may make \mistakes" in their consumption decisions, given how complex it is to calculate prices and given the uncertainty associated with solar generation uctuations. Sunshine is a reasonably salient heuristic that signals when prices are low, but it can uctuate quickly and frequently, such as when clouds pass through. It's also a function of the size

[10] Conversations with those that ran the trials and my reading of the documentation indicate that there was some di culty with uptake. It's unclear what could have driven this issue. Once this became problematic, each household approached was o ered an IHD until the minimum number set in advance was met.

[11] Although self-generated solar is actually free, there is an opportunity cost of consuming it given how the household would otherwise getting paid for it.

and design of the system as well as local air pollution, making it cognitively complex to calculate, especially in real time. The marginal price of electricity at any point in time is also determined by current consumption levels, as prices are only lower than the retail rate if consumption does not exceed solar generation. Knowing the price thus requires understanding precisely how much electricity household appliances and lighting require as well as how much solar is being generated. Information acquisition costs are high, as calculating prices in real time is a cognitive burden, time-consuming, and requires constant attention.

Absent energy storage, the primary way in which households can substitute electricity imported from the incumbent supplier with their own self-generated solar power is by shifting their electricity load from hours when the sun is not shining to daytime hours in which onsite lower-priced power can meet those needs. For instance, this can be done by using electricity-consuming appliances during the day rather than the evening. The households in the sample are all located in the Northeast region of Great Britain, where it is extremely uncommon to have or use an air conditioner and electric heating is rare. As such, load-shifting mostly entails running one's dishwasher or washing machine when the sun is shining. Households could do this in real time if they are home during the day or by using delay timers and scheduling appliances in advance to run later.

3 Data and Empirical Strategies

3.1 Data

Each household in the eld trial that I study was tted with a meter to record onsite solar PV generation and net power ows at 1-minute intervals. I construct several other variables for the analysis. Gross demand is equal to the sum of net demand (which can be negative or positive) and solar PV generation. Exports are equal to the absolute value of net consumption when it is negative and zero otherwise. Imports are equal to net consumption when it is greater than zero, capturing consumption that exceeds time-coincident onsite solar generation. I also calculate solar self-consumption|the amount of solar generated onsite that the household consumes|since the source of power that meets changes in demand can shed light on the environmental and market implications of behavior change. Self-consumption is electricity demand that is precisely time-coincident with solar PV generation but which does not exceed solar output.[12]

[12] This is simply equal to PV generation when no excess PV generation is exported to the grid. When exports are positive, self-consumption is the di erence between PV generation and exports, since total PV generation captures power that is consumed on-site plus the excess power exported to the grid (McKenna, Pless and Darby 2018).

Using extremely high-frequency data is critical to accurately calculate self-consumption (Widen, Wackelgard, Paatero and Lund 2010). Since solar generation tends to be \spikey"|either from sudden cloud coverage or local air pollution|and demand can be as well due to household appliance use, granular time-resolution allows one to observe whether usage coincides perfectly with generation (in the absence of onsite electricity storage).[13] McKenna et al. (2018) show that using time-averaged data produce substantial errors, ranging up to 15 percent when using 30-minute time-averaged data as opposed to 1-minute resolution data.

Figure I illustrates the importance of accounting for these frequent and short-lived spikes. The electricity demand pro le follows the expected pattern, with demand being highest during the UK's peak period between 4pm and 8pm, and solar PV generation follows a reasonable bell-shaped curve from sunrise to sunset. There are spikes in PV generation throughout the day as well as spikes in the use of self-generated power that capture household behavior such as appliance use.

FIGURE I HERE

I follow a few standard steps to address meter reading errors and data abnormalities that I detail in Appendix A. The nal dataset includes a total of 294 households, 144 of which have IHDs and 150 do not. The data cover a long time period relative to others in the literature that examine household electricity use. I observe behavior for more than a year for all but 20 of the households, and I have data for an average of 510 days per household overall, providing the bene t of being able to study longer-run behavior and across all four seasons. This results in about 14.3 million observations at 15-minute intervals.

I calculate gross consumption, imports, and self-consumption using the 1-minute data to preserve accuracy but then aggregate the variables to more realistically capture household behavior and appliance use for the regression analyses. I start by aggregating to the 15-minute level data given how many household activities take more than a minute but less than an hour, like the frequent use of kettles to make tea many times throughout the day in the UK. I also estimate the model using hourly data to account for how the use of some household appliances require more time, such as running a washing machine or dishwasher. Finally, I use daily level data to examine load-shifting behavior. Testing the e ects at each of these levels of aggregation will also allow us to explore whether households respond to marginal or average prices and the ways in which they

[13] For instance, if hourly data are used, separate meters for generation and consumption could both show readings of 1 kWh. Without knowing their time-coincidence, we cannot be sure that self-consumption was 1 kWh and instead 0 kWh (if all solar generation was exported).

engage with their IHDs, as I describe later.

Appendix Table B.1 presents summary statistics for the 15-minute, hourly, and daily derived and observed variables. The means of each variable are about what I would expect for the average UK household in the region that I are studying. I also see that households with IHDs consume more electricity and generate more solar than those without, which could be indicative of having larger homes. This isn't problematic for the question that I study—when I examine the role of IHDs, I are most concerned with comparing how IHDs change the responsiveness to price levels and uncertainty within household and even within household-day—but it highlights the importance of controlling for unobservables with a rich set of fixed effects.

3.2 Estimating Demand Response to Real-Time Prices

My first objective is to estimate the price elasticity of demand with respect to real-time prices. I measure prices by solar electricity generated by the household. An increase in solar PV generation by one kWh, Q_{it}, can be thought of as a reduction in the marginal price of electricity. When using the 15-minute and hourly level data, I use within household-day variation for identification and estimate linear models of the following form:

$$Y_{it} = \quad + \quad Q_{it} + \quad_{id} + \omega_{moh} + X_{it} \quad + \quad_{it}; \tag{1}$$

where Y_{it} is kWh of either gross consumption, imports (net consumption when it is positive), or self-consumption of electricity generated by onsite solar for household i in time t. Using solar generation as the price variable introduces a substantial amount of cross-sectional and temporal variation within-day and within-household that is driven by household-level solar shocks. Solar generation in total each day or in the long-run is a function of endogenous factors like system size, design, and location, but I control for these factors.

The fixed effects control for the many other factors that affect electricity demand. Household-day fixed effects, $_{id}$, control for unobservable differences in consumption patterns across households and how those household-specific patterns may change each day (d). Since energy consumption tends to vary systematically by hour (h) and differentially for each day of the week (o) depending on preferences and when household members tend to be home, as well as by month (m) given how weather drives electricity demand, I also include month-hour-day of week fixed effects, ω_{moh}.[14] The

[14] For example, the tendency to be home more in the winter or on the weekends affects demand, and consumption varies by hour based on when household members are home and the weather.

matrix X_{it} includes hourly air temperature and humidity to control for how such variables a ect a household's use of certain appliances.[15]

I aggregate the data to the daily level to observe whether responses to uctuations in prices change the day's consumption behavior as a whole and whether households engage in load-shifting, moving consumption away from the peak demand period (4pm to 8pm) towards earlier hours in the day when prices are lower and they're more likely to be consuming solar electricity. I estimate variations of the following model when using the daily-level data:

$$Y_{id} = \quad + \quad Q_{id} + \quad_{i} + \quad_{d} + !_{im} + \quad_{io} + \quad_{id}; \qquad (2)$$

where Y_{id} is daily gross consumption or, when examining load-shifting, peak period gross consumption or imports for household i on day d. When the day's total gross consumption is the dependent variable, Q_{id} is the household's total daily solar generation. It is the household's \daytime" solar generation (total solar between the hours of 9am and 4pm) when estimating load-shifting e ects. Importantly, I include each household's peak period solar generation in the load-shifting regressions as well to control for peak period prices. Household ($_i$) and day ($_d$) xed e ects are included as before. Since household-day xed e ects would be absorbed when using daily-level day, I include household-month ($!_{im}$) and household-day of the week ($_{io}$) xed e ects to still control for how households' consumption norms vary seasonally and throughout the week.[16] I estimate some variations of this model in the robustness checks section, especially to be sure that the exclusion of household-day xed e ects does not bias the results.

3.3 Disentangling Inattention and Uncertainty

Next, I explore whether IHDs a ect behavior, and if so, whether this is due to an increase in attention or improvements in forming expectations about future prices. Solar generation each period, which I have been using thus far, captures price levels. The availability of IHDs may change responsiveness to price levels by drawing attention to whether the household is currently exporting or importing power with its very visible lighting system (and thus the price is either zero or the retail rate), but as discussed in previous sections, prices are also highly uncertain. Rather than engaging with their IHDs in real time, households might learn from them and adjust the ways

[15] Although these are hourly, there is just a single measure that applies to each household, since I do not observe household location. I know that all households are within a certain electricity service region, though, so I take the average of ve representative locations within that region. More detail is provided in the Appendix.

[16] I do not include the weather control variables here since they are absorbed by daily xed e ects.

in which they forecast future prices.

Knowing whether IHDs change behavior due to attention or price uncertainty requires 1) a measure of price variability that is distinct from price levels, and 2) knowing whether households are engaging with their IHDs in real time. To begin, I construct household-day-hour and household-day measures of solar generation variability to measure price uncertainty, allowing us to estimate how demand responds to price levels and variability simultaneously and how IHDs change the responsiveness to each. I estimate the following model to see whether IHDs affect within household-day behavior with respect to price levels versus variability when using the 15-minute and hourly data:

$$Y_{it} = \alpha + \beta_1 Q_{it} + \beta_2 Q_{it} \times IHD_i + \beta_3 U_{it} + \beta_4 U_{it} \times IHD_i + \gamma_{id} + \omega_{moh} + X_{it}\theta + \epsilon_{it}; \qquad (3)$$

and the following model when using daily data:

$$Y_{id} = \alpha + \beta_1 Q_{it} + \beta_2 Q_{it} \times IHD_i + \beta_3 U_{it} + \beta_4 U_{it} \times IHD_i + \gamma_i + \delta_d + \omega_{im} + \eta_{io} + \epsilon_{id}; \qquad (4)$$

where U_{it} is price uncertainty faced by household i in time t and IHD_i is an indicator equal to one if household i has an IHD and zero otherwise. All other variables are the same as before. I measure price uncertainty as the standard deviation of solar generation within household-day-hour when estimating Equation 3 and within household-day using just "daytime" hours when estimating Equation 4. All other variables are the same as before. The coefficients β_1 and β_3 capture how Y_{it} responds to changes in price levels and uncertainty, respectively, and β_2 and β_4 capture how IHDs adjust those responses.

Although households with IHDs tend to consume more electricity and generate more solar, solar generation and its variability are within-household and within-day (or within-hour) measures. Variation in these measures is still driven strictly by exogenous changes in factors like the weather. Therefore, even if the IHD effects relative to households without IHDs are biased in some way despite the rich set of fixed effects, the difference between the coefficients associated with the two IHD interaction effects can be thought of as capturing whether those with IHDs change behavior in relation to price levels or uncertainty.

One might be concerned that IHDs may change responses to price variability simply because they draw attention to the price fluctuations, particularly in the 15-minute and hourly-level cases. In later sections, I conduct additional tests to explore whether households appear to be using their IHDs in real time or learning from them and improving their forecasting abilities. This helps confirm that this is indeed capturing different concepts and that the responses to variability are associated with price uncertainty as opposed to attention.

3.4 Matching + Machine Learning Approach

To corroborate the findings on the role of IHDs, I combine machine learning and matching methods in the spirit of Burlig et al. The approach can be thought of as an enhanced matching approach, whereby I predict outcomes for households with IHDs as though they did not have them using household-specific prediction models. The Appendix provides more detail but I discuss the intuition here.

I start by matching households with and without IHDs based on their electricity consumption and solar generation profiles.[17] I then build household-specific prediction models for all households that were not given IHDs using random forests, which outperformed other approaches like LASSO. Although I explored the predictive power benefits of including fully interacted fixed effects, the most successful models were those that used hour, day of the week, month, and year fixed effects, as well as solar generation, hourly solar uncertainty, range of hourly solar, and mean daily gross consumption as predictors. Lastly, I take the original data of those with IHDs and predict their outcomes using the prediction models of their matched non-IHD household, such that I have household-specific predictions of how households with IHDs may have behaved without them at the 15-minute level. I estimate Equation 3 using the original IHD households as the treated group, and the "new" households that have the inputs of IHD households but the predicted values as outcomes as the control group.

4 Main Results

4.1 Demand Response to Real-Time Pricing

I begin by estimating Equation 1 when using 15-minute and hourly data and Equation 2 when using daily data, including only observations for which solar PV generation is greater than zero.[18]

[17] I used optimal pair matching. See appendix for details on variables used.

[18] This essentially includes all observations for times between sunrise and sunset for the 15-minute and hourly data, as even if there is almost no sun, the slightest bit of light will generate some very small amount of solar. I omit

The results are presented in Table I, providing the effects of solar generation on gross consumption (Column 1), imports (net consumption when it is positive) (Column 2), and solar self-consumption (Column 3). I find that households increase gross consumption by about 0.07 kWh for every one kWh increase in generation. This is consistent across all three levels of time granularity, although statistical significance is lost when using the daily data. The estimates translate into price elasticities ranging from -0.022 and -0.056 when using the mean values.

[TABLE I HERE]

From an environmental perspective, one might be concerned about increased emissions associated with more electricity consumption, but the increase in gross demand is entirely met by self-consumption whereas imports decrease substantially. For every one kWh increase in solar, imports decrease by somewhere between 0.17 kWh and 0.21 kWh. This suggests that there is substantial intraday substitution of imported power with one's own solar power. Whether this actually reduces emissions, though, depends on the energy sources supplying electricity from the incumbent grid at the times in which imports are displaced. If the grid is using power generated from solar or another clean energy source, then the net effect on emissions is zero, but emissions go down if the household would have otherwise imported electricity generated by conventional sources.

These findings are likely over-estimates of the true effects. Solar households tend to be more engaged with their electricity consumption than the average household in the first place, leaving less room for adjustment, and they are often very interested in shifting the timing of their consumption to take advantage of their clean self-generated power.[19] While one might ask whether the estimates are actually lower bounds since households with solar tend to have higher incomes and may be less sensitive to price changes, evidence in the literature so far indicates that electricity demand is actually only weakly correlated with income (citations). I suspect this is particularly true for load-shifting, since lower income households are less likely to have the most modern appliances, like those with delay timers.

4.2 Load-Shifting

To examine whether households shift consumption from peak hours when demand is highest (4pm to 8pm) to daytime hours when the most solar is likely to be generated (9am to 4pm), I estimate the

observations when solar generation is zero because there are no price changes during this time. All observations are included for the daily data.

[19] Survey responses following the CLNR trial confirm this assumption, as respondents indicated that their motivation for participating was to learn how to take better advantage of their self-generated solar power (CLNR 2013).

e ect of total daytime solar on peak gross and net demand (imports) at the daily level. An increase in daytime solar generation should decrease both gross consumption and imports during peak hours if load-shifting occurs. I control for peak solar generation given how this could mechanically a ect peak imports and self-consumption, and I provide further robustness checks that peak generation does not drive the ndings in Section 4.3.

The main results are presented in Table II and show that households do indeed change the timing of their consumption in response to more solar generation during the day, but the e ects are small. Columns 1 and 2 provide the e ects of daytime solar generation on peak gross consumption and net imports, respectively, and evening gross consumption and imports (4pm to midnight). I nd that a one kWh increase in daytime solar decreases both peak gross consumption and imports by -0.027 kWh and evening gross consumption and imports by about -0.035 kWh.

To put these numbers into perspective, the e ects on peak imports and gross consumption represent 0.1% changes relative to the dependent variable means and price elasticities of -0.039 and -0.037, respectively. I interpret these ndings as upper bounds again for the same reasons as before. Furthermore, the identical e ects for gross and net demand assures us that all reductions during peak hours are driven by reductions in imported power from the grid rather than reductions in self-consumption.

[TABLE II HERE]

4.3 Robustness Checks

Not a mechanical relationship|A key complication to rule out is that the e ects are not driven by a mechanical relationship rather than intentional household decision-making. For instance, an increase in solar generation will mechanically increase self-consumption and reduce imports, all else equal. Although the e ects on gross consumption and load-shifting already indicate that there is something more going on|daytime solar generation that does not align with peak period hours cannot mechanically a ect self-consumption, for example, and gross consumption should not change if consumers do not perceive a reduction in prices|I conduct a series of robustness checks to be sure.

First, we estimate the e ects for only the sunniest times of day. Since base-load demand alone can be enough to account for consumption all available solar when the sun is just rising or setting, the results for self-consumption could be biased upwards and for imports they could be biased downwards. Ensuring that the e ects hold when solar generation far exceeds base-load can provide con dence that the e ects are driven by behavioral change.

The results are presented in Appendix Table B.2 using the 15-minute (Panel A) and hourly (Panel B) data. In Column 1, I include only \daytime" hours, which I de ne as between 9am and 4pm. Columns 2 and 3 use variations of this de nition, including hours 10am to 4pm in Column 2 and hours 12pm to 3pm in Column 3. In Column 4, I include only observations for which solar generation exceeds the median value of solar generation, and in Column 5, I do the same but nding the median for only \daytime" hours. The e ects range from 0.058 to 0.060 when using the 15-minute level data and 0.059 and 0.068 when using the hourly data, barely changing from the baseline and remaining statistically signi cant at the 1% or 5% level in all cases. Taken together, these ndings for the impact on gross consumption provide con dence that the average e ects capture household decision-making.

Second, I probe whether the load-shifting results are driven by anything mechanical. This is highly unlikely since daytime solar cannot be consumed in peak or evening hours (absent onsite storage). That said, if a day is particularly sunny, daytime solar generation may be correlated with solar in other hours, and more solar during peak hours reduces imports mechanically (but not gross consumption). I already control for peak solar generation in all of the load-shifting regressions, but to show how little this matters for the results, I estimate the e ect of peak generation alone without daytime generation included just to be sure. In Appendix Table B.3, I show that the e ect of peak generation on peak gross consumption and imports (Columns 1 and 2) and full evening (4pm to midnight) gross consumption and imports does not change the results. Peak generation does not have a statistically signi cant e ect on the outcomes when included on its own (Panel A) or when daytime generation is included (Panel B).

Other Robustness Checks. I conduct a few more robustness checks to be sure that the e ects are not driven by any of my modelling choices and provide the results in Appendix Table B.4. First, since I are not able to include household-day xed e ects in the daily-level regressions, I use household and day xed e ects separately as opposed to interacted and estimate 1 while using the 15-minute and hourly data to be sure that the results are not driven strictly by within-household variation. Columns 1 and 2 show that the results are nearly identical to the baseline ndings.

Similarly, the daily-level estimates could be biased if households learn or adjust di erentially over time due to the exclusion of household-day xed e ects. To account for this, I control for the number of months that have passed since each household entered the eld trial. The results are provided in Columns 3 and 4 of Appendix Table B.4 when using daily peak gross and peak imports respectively. The e ects do not change relative to the baseline. This con rms that the household and day e ects on their own are su cient. Lastly, in Columns 5 and 6, I include household-month-day of week xed e ects fully interacted for the load-shifting regressions and the results again do not change.

5 In-Home Displays and Inattention vs. Uncertainty

Providing useful guidance on whether real-time pricing might be welfare-enhancing requires knowing whether households are actually responding to marginal prices in real time as well as whether they are fully optimizing, and if not, why. I tackle these questions in this section by disentangling whether in-home displays a ect behavior by increasing attention as opposed to reducing uncertainty, and applying the approach in various ways to shed light on the prices to which consumers respond, the adjustments consumers make to load-shift, how and why they might make optimization \mistakes", and whether they form new habits or make long-term adjustments. Doing so allows me to assess what interventions might have the most promise for achieving policy and load-shifting objectives.

5.1 Main E ects of In-Home Displays

Consumers may be (rationally or irrationally) inattentive to prices, whereby \inattention" can be thought of as a proxy for most behavioral biases discussed in the literature (Gabaix 2019). If this is the case, salience of prices and forms of information that do not impose a high cognitive burden are important. On the other hand, consumers might (rationally) reduce consumption due to price uncertainty, or they could increase consumption due to price uncertainty because the risk is minimal or they are excited by potential increases in solar generation. To test whether households respond in ways that align with their preferences, and if not, disentangle the underlying source of errors, I estimate variations of Equation 3 and Equation 4.

Table III provides the main baseline results. In Panel A, I provide the e ects on gross consumption when estimating Equation 3 and using the 15-minute and hourly data. Panel B shows the load-shifting results when using daily data and estimating Equation 4. The key takeaway is that consumers make mistakes with respect to price uncertainty and IHDs adjust these errors almost entirely but do not a ect responsiveness to price levels at all.

In Columns 1 and 3, I estimate the models when including only price levels and their interaction with the IHD dummy. The IHD effects are zero in all cases. Once I also include price uncertainty and its interaction with the IHD dummy (Columns 2 and 4 of Table III), however, I can see that households increase consumption with respect to uncertainty and IHDs reduce this effect. A one kWh increase in the standard deviation of within-hour solar generation increases gross consumption by 0.032 kWh and 0.122 kWh when using the 15-minute and hourly data, respectively (Panel A). The provision of IHDs, however, more than offsets these effects. In all of these regressions, the main effects of price levels are stable (although statistical significance is just barely lost in Column 4), and the effect of IHDs on the responsiveness to price levels is still zero, indicating that IHDs are changing behavior only through the uncertainty channel.

These findings also carry over to load-shifting. The results in Columns 1 and 3 of Table III show that IHDs do not change load-shifting behavior at all with respect to price levels. However, once adding the uncertainty measures in Columns 2 and 4, we can see that households increase consumption during peak hours when uncertainty during daytime hours is high, offsetting some of their load-shifting. The IHDs again help to offset this effect, though, and have no effect on how consumers respond to price levels (daytime solar). The correction is imperfect but substantial—about 80% of the error.

The load-shifting results begin to suggest that IHDs are changing the response to price variability because households are getting better at planning as opposed to simply drawing their attention to fluctuations. If so, this implies that households are using their IHDs to inform their expectation formation and start making better forecasts. Understanding this distinction is critical, as the latter would indicate that these effects are not actually related to price uncertainty at all, and I explore this more below.

5.2 Machine Learning Results

When estimating Equation 3 using the data constructed from the machine learning and matching approach, the results are very similar to the main econometric results and do not qualitatively change. Appendix Table B.5 provides the estimates when using the 15-minute and hourly data.[20] I find that a one kWh increase in solar generation increases gross consumption by somewhere between 0.064 kWh and 0.073 kWh. Households again make errors with respect to uncertainty.

[20] I do not estimate daily effects using the machine learning-generated data because the predictions are poor during hours when generation is equal to zero. This is because generation is a strong determinant of the outcome variables, and this leads to poor predictions in evening hours when generation is equal to zero.

Gross consumption increases when uncertainty increases but the magnitudes of the effects are smaller. The IHDs again offset the errors while not affecting responses to price levels at all.[21] These help provide more confidence in the results from the baseline estimates. The fact that these two methods produce similar results also suggests that the machine learning and matching approach might be useful in other settings in which a true treatment group does not exist.[22]

5.3 Policy Implications: Behavioral Intervention or Innovation?

Understanding how and why households increase consumption with respect to price uncertainty—and why IHDs mitigate this—is important for informing policy. The IHDs do not predict future prices, raising questions about how consumers are actually using their IHDs and whether IHDs are the appropriate tool for addressing price uncertainty. Understanding the ways in which they use their IHDs can also confirm that the results are indeed capturing price uncertainty errors as opposed to increased attention to price fluctuations. In this section, I examine the prices to which consumers respond, test whether they are present-biased or forward-looking, examine whether effects are symmetric, and explore whether they learn or adjust over time. Taken together, the findings point to the need for price-smoothing innovation—with technological or operational innovations—to meet popular policy objectives.

5.3.1 To What Prices Do Consumers Respond?

Our analyses so far have assumed that responses to price changes are responses to observed marginal prices in real time. However, under uncertainty, rational agents respond to *expected* marginal prices (Borenstein 2009), and recent evidence suggests that consumers actually might respond to *average* electricity prices when prices are complex (Ito 2014). Given how prices in this context are both complex and uncertain, I propose another type of price to which electricity consumers might respond: *expected average* prices. The reasoning behind this is that, if consumers are rationally forward-looking but unable to respond to expected marginal price changes in real time, they must form expectations about future average prices over the course of a few hours in order to consume when prices are low. Consumers in this setting thus may be responding to either expected marginal or expected average prices.

[21] Statistical significance is lost on the IHD effect when using the 15-minute level data, but the magnitude of the effect is consistent with previous findings in terms of more than offsetting the mistake, the IHD effect also fully offsets the error and remains statistically significant when using the hourly data.

[22] A caveat, however, is that the method suffers from selection on unobservables.

Understanding the prices to which consumers respond can shed light on how econometricians should estimate electricity demand elasticities. It also can provide insight on the ways in which consumers use their IHDs to inform their consumption choices. That is, if IHDs change behavior with respect to expected marginal prices, this implies that they are engaging with their IHDs in real time and that real-time information is helpful. If IHDs change behavior with respect to expected average prices, this implies that they learn from the IHDs better than they do without them, but this would be in the form of studying the data collected each day rather than engaging in real time.

I carry out a series of tests to explore this. First, I examine whether the e ects di er on weekends versus weekdays. If consumers respond to expected marginal prices, the uncertainty error and IHD e ects when using the 15-minute and hourly data should be largest on the weekends when household members are home. On the other hand, if households use expected average prices to inform consumption, I expect the uncertainty error and IHD e ects to be larger on the weekdays. This is when household members are more likely to form expectations about daytime average prices in advance, since they may then need to be out of the house most of the day for work, and schedule appliances to run while they are away.

Table IV presents the results. The ndings are consistent with households responding to expected average prices. They appear to make mistakes due to imperfect forecasting and appliance scheduling that does not precisely align with the timing of solar generation. I draw this conclusion from the following logic. Consumers respond sub-optimally to uncertainty within short time increments (Panel A) on the weekends but not the weekdays. In other words, when household members are more likely to be home, they over-react to sunshine uctuations as they occur. Having an IHD does not x this error, suggesting that they are not engaging with their IHDs in real time. On the other hand, load-shifting uncertainty errors occur on the weekdays.[23] This suggests that \mistakes" are due to not perfectly aligning appliance scheduling with solar during the day. Households may form expectations about average daily prices based on weather forecasts but such forecasts are imperfect. They learn how to improve their predictions with IHDs, but this must occur in the form of learning from data gathered each day.

To provide more con dence in these conclusions, I also use the 15-minute level data to estimate the e ects for di erent times throughout the day. I expect the e ects to be highest when household members tend to be home and have the time to engage with their IHDs. Appendix Table B.7

[23] Statistical signi cance is lost when examining peak gross consumption (Panel B, Column 2), but when examining peak imports (which matters more for policy objectives), I can see that the \mistakes" occur on weekdays, and IHDs greatly reduce this error.

shows that uncertainty mistakes and the IHD correction associated with expected marginal prices are entirely explained by weekend mornings. As household members are likely to be home and to have more free time during these hours—not rushing to work during the weekday mornings but also likely out and about during the day on the weekends—it appears as though households engage with their IHDs in real time in these select hours. However, this is clearly a very small part of the week, providing further suggestive evidence of the limitations of IHDs for helping consumers better respond to expected marginal prices.

These results also confirm that the approach to disentangling inattention from uncertainty is indeed capturing two distinct concepts. If the uncertainty error and IHD correction were made in real time (i.e., in response to marginal prices), then I couldn't be sure that IHDs weren't just drawing households' attention to the price variability as opposed to informing their expectations. Finding that the effects appear to be made in relation to expected average prices suggests that errors are made with respect to forecasting. Indeed, in Section 6.3.2, I show that consumers are forward-looking in their electricity consumption decision-making.

Finally, I show how estimating the demand elasticities at the inappropriate level of granularity can lead to misleading conclusions about whether load-shifting occurs. A handful of recent papers test whether households load-shift in response to time-of-use pricing and conclude either that no load-shifting occurs or that IHDs do not impact load-shifting (Gilbert and Graff Zivin 2014; Martin and Rivers 2018; Prest 2020). They do so by estimating the heterogeneous effects of prices at different times of the day and interpret constant effects across all hours as indicating that there is no load-shifting. However, this approach assumes that households respond to marginal prices. Furthermore, testing for load-shifting should entail estimating the effects of lower prices during daytime hours on peak consumption (or the impact of higher peak prices on daytime consumption), which also then enables using average prices.

I show the importance of this distinction by estimating just the effects of solar generation using Equation 1 at the 15-minute level. The effects are plotted in Figure II. I can see that, as found in a few of the other studies, the effects are flat throughout the day. Yet I demonstrate that households in the sample do indeed decrease consumption in peak hours in response to lower prices earlier in the day, it's just that their response to prices in each time period are constant. The hourly estimates reinforce that households are not responding to marginal prices—otherwise I would expect the effects to increase in the hours when the most solar is generating, such as early to mid-afternoon—while also highlighting how different conclusions are drawn about load-shifting when assuming that households respond to marginal prices.

[FIGURE II HERE]

5.3.2 Are Consumers Present-Biased or Forward-Looking?

Another form of inattention that can lead to undervaluation of bene ts is present bias. In my setting, this would mean that consumers make decisions based on recent or current experiences (such as the current or the last period's weather) rather than being forward-looking and making decisions based on future expected prices (by looking at weather forecasts, for instance). There is an extensive literature showing that consumers and rms often fail to make positive net present value in energy e cient appliances and fuel-saving vehicles because of under-valuing future savings (Allcott and Greenstone 2012; Gillingham, Houde and van Benthem forthcoming).[24] In my setting, understanding whether consumers are present-biased can shed light on whether they make uncertainty \mistakes" because of consumer myopia as opposed to being forward-looking and making forecasting errors. If a present-biased decision-making process explains the uncertainty errors, this not only implies that RTP is ine cient, but also that the uncertainty mistake is due to behavioral bias as opposed to rational behavior and imperfect forecasting.

I provide two sets of results suggesting that households are not myopic. First, I test whether load-shifting can be explained by price levels and variability in the previous period as opposed to the current period (see Appendix Table B.8).[25] Columns 1 and 4 estimate the e ect of lagged generation on the current period's peak gross consumption and peak imports, respectively, and I add the lagged day's uncertainty variable and interaction e ects in Columns 2 and 5. I then add all of the original current day's prices, variability, and information interactions in Columns 3 and 6. The lagged variables are statistically zero across all speci cations while the original results continue to hold. Second, I test whether households make load-shifting decisions based on weather in the morning when they may be scheduling appliances to run later in the day. Appendix Table B.9 shows that there are no statistically signi cant e ects of morning solar generation on load-shifting when including various sets of my baseline daytime solar generation and IHD interactions, and the original results continue to hold. This again suggests that households are forward-looking and forming expectations based on forecasts rather than a heuristic associated with the current weather.

[24] See Gillingham and Palmer (2014) for a review of earlier literature.

[25] I do not do this for the 15-minute and hourly level regressions because of how highly correlated generation is in these short time increments.

5.3.3 Asymmetric Responses and Excitement

Conventional economic theory suggests that increases and decreases in prices should have the same impact on quantity demanded, however evidence from psychology and the behavioral economics indicates that this isn't always the case. If consumers are loss averse, for example, they may respond more to price increases (a reduction in solar generation) relative to decreases (an increase in solar generation) since they experience more loss psychologically from price increases than gains from price decreases (Tversky and Kahneman, 1991). This seems unlikely in my setting since the consumer knows that the highest prices will go is the retail rate, and electricity bills make up a relatively small portion of a household's monthly budget. On the other hand, since these are consumers that are excited about the opportunity to use their own self-generated power, they may make uncertainty errors more frequently when there are upwards swings in solar generation. If so, this suggests that the uncertainty error is driven by \excitement".

I test whether consumers respond symmetrically to changes in solar generation to see if loss aversion or excitement can explain the uncertainty error. To do so, I estimate the baseline speci -cations separately for observations in which the current period's solar generation is higher or lower than the previous period's solar generation. The results are presented in Table V. The evidence consistently shows that the uncertainty error and the IHD mitigation of the over-reaction occurs for solar generation increases (i.e., price decreases) but not at all for solar decreases.[26] This indicates that households make uncertainty mistakes due to excitement as opposed to loss aversion, which is consistent with households in the sample being particularly engaged with their energy use.

5.3.4 Learning and/or Capital Investment Over Time

Whether the provision of ongoing real-time feedback via an IHD is necessary to change behavior as opposed to upfront information depends on whether IHDs facilitate learning and adjustment more so than some sort of upfront information and training. Households could learn about their electricity use and solar generation and change their electricity consumption habits accordingly, or they could invest in physical capital, like more energy-e cient appliances or appliances that have timers. If households learn or invest only if given an IHD, we should see the IHD e ect persist over time. However, if the e ects for those without IHDs converge with those that do have IHDs, this would suggest that households learn and adjust even without an IHD. They could be learning from

[26] Statistical signi cance is just barely lost for the IHD e ect for solar increases when using the hourly data, but the magnitude of the e ect is still much higher than it is for solar decreases.

their electricity bills, for instance.

I test this by estimating the load-shifting e ects for sub-samples of the data using observations for which the household had only been in the trial for less than 90 days, 90 to 180 days, or more than 180 days. The results are presented in Table VI. Households with IHDs load-shift right away, but they also over-react to uncertainty right away, suggesting that having an IHD actually drives this excitement at rst. The uncertainty error is mitigated after about three months. At the same time, those without IHDs also adjust behavior over time, and the price responsiveness of the two sets of households converges. Non-IHD households do not load-shift at all in the beginning but they do and to the same degree as those with IHDs after three months.

These ndings suggest that ongoing information may not be any more e ective than training and feedback upfront. The IHDs help households adjust upfront, an e ect that I showed is not driven by engaging with IHDs in real time, and consumers learn quickly without continuing to adjust further. The learning or adjustments for those without IHDs also indicates that participating in the trial alone helped inform their consumption decisions. This is reasonable, since all households had smart meters installed that could more accurately track their consumption and joining the trial could have prompted them to think more carefully about opportunities to load-shift. Given how the load-shifting itself persists, all households formed new habits and/or made investments in more energy-e cient appliances.

6 Conclusions

While I have shown that households engage in at least some load-shifting under real-time pricing, and that providing real-time information a ects behavior, putting the ndings together sheds light on the limitations of RTP and behavioral interventions when the objective is to reduce energy use or shift consumption timing. There are ve key pieces of evidence that lead to this conclusion: 1) demand is very inelastic; 2) households respond to expected average prices, not marginal prices, and real-time information only provides marginal prices; 3) while IHDs a ect behavior, the e ect occurs entirely through the price uncertainty channel by informing expectation formation, but IHDs are not forecasting devices; 4) consumers do not engage with their IHDs in real time; 5) all households learn and adjust relatively quickly, either by way of new habits or investing in new appliances, although IHDs facilitate slightly faster learning.

This is not to say that real-time pricing has no bene ts|its impact on demand response just appears to be limited. A broad takeaway is that technological and operational innovation is needed to achieve substantial reductions in emissions. For example, cost-effective electricity storage could enable consumption in the evenings to be met with clean electricity. Even automation cannot achieve this. Policymakers have various options to consider, such as taxing carbon at the right level to induce innovation in clean technology and providing subsidies to help firms overcome financing constraints associated with investing in R&D.

This paper also emphasizes the importance of disentangling whether information affects behavior by increasing attention or improving price expectation formation, a challenge that characterizes many familiar settings. If consumers do not respond to marginal prices or engage with real-time feedback, estimating the elasticity of demand using marginal prices can be misleading, and behavioral interventions aiming to increase attention might not be welfare-enhancing. Resolving price uncertainty requires a different policy approach.

As with most studies of this nature, there are certainly limitations associated with the external validity of the findings. The implications are unlikely to apply in all dynamic pricing contexts, for instance, particularly when decisions are being made about much larger purchases or more expensive activities (e.g., airline tickets or vacation accommodations). Markets in which prices are highly volatile but purchases or activities make up a smaller portion of a consumer's budget are more similar to ours (e.g., Uber rides for wealthy riders), or when there are very few available substitutes. Nonetheless, the effects found in this paper are likely over-estimated if anything, and thus the general conclusions are likely to apply in most electricity markets.

References

Allcott, H. , \Rethinking Real-Time Electricity Pricing," Resource and Energy Economics, 2011a, 33 (4), 820{42.

_ , \Social Norms and Energy Conservation," Journal of Public Economics, 2011b, 95 (9), 1082{1095.

_ , \The Welfare E ects of Misperceived Product Costs: Data and Calibrations from the Automobile Market," American Economic Journal: Economic Policy, 2013, 3 (3), 30{66.

_ and C. Knittel , \Are Consumers Poorly Informed about Fuel Economy? Evidence from Two Experiments," American Economic Journal: Economic Policy, 2019, 11 (1), 1{37.

_ and D. Taubinsky , \Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market," American Economic Review, 2015, 105 (8), 2501{2538.

_ and M. Greenstone , \Is There an Energy E ciency Gap?," Journal of Economic Perspectives, 2012, 26 (1), 3{28.

_ and T. Rogers , \The Short-Run and Long-Run E ects of Behavioral Interventions: Experimental Evidence from Energy Conservation," American Economic Review, 2014, 104 (10), 3003{37.

Bollinger, B. and W. Hartman , \Information vs. Automation and Implications for Dynamic Pricing," Management Science, 2020, January, 1{501.

Borenstein, S. , \The Long-Run E ciency of Real-Time Electricity Pricing," The Energy Journal, 2005, 26 (3), 93{116.

_ , \To What Electricity Price Do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing," Paper presented at NBER Summer Institute, 2009.

_ and J. Bushnell , \Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and E ciency," Energy Institute at Haas, Working Paper 294R, 2019.

_ and S. Holland , \On the E ciency of Competitive Electricity Markets with Time-Invariant Retail Prices," RAND Journal of Economics, 2005, 36, 469{493.

Bradley, S. , \Inattention to Deferred Increases in Tax Base: How Michigan Homebuyers are Paying for Assessment Limits," Review of Economics and Statistics, 2017, 99 (1), 53{66.

_ and N. Feldman , \Hidden Baggage: Behavioral Responses to Changes in Airline Ticket Tax Disclosure," American Economic Journal: Economic Policy, 2020, 12 (4), 58{87.

Busse, M., J. Silva-Risso, and F. Zettelmeyer , \$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions," American Economic Review, 2006, 96 (4), 1253{70.

Cachon, G., M. Daniels, and R. Lobel , \The Role of Surge Pricing on a Service Platform with Self-Scheduling Capacity," Manufacturing & Service Operations Management, 2017, 19 (3), 368{384.

Chetty, R., A. Looney, and K. Kroft , \Salience and Taxation: Theory and Evidence," American Economic Review, 2009, 99 (4), 1145{77.

_ and E. Saez. , \Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients," American Economic Journal: Applied Economics, 2013, 5 (1), 1{31.

CLNR , \Insight Report: Domestic Solar PV Customers," 2013.

Cohen, P., R. Hahn, J. Hall, S. Levitt, and R. Metcalfe , \Using Big Data to Estimate Consumer Surplus: The Case of Uber," NBER Working Paper 22627, 2016.

Cullen, Z. and C. Farronato , \Outsourcing Tasks Online: Matching Supply and Demand on Peer-to-Peer Internet Platforms," Management Science, 2020.

Deryugina, T., A. MacKay, and J. Reif , \The Long-Run Dynamics of Electricity Demand: Evidence from Municipal Aggregation," American Economic Journal: Applied Economics, 2020, 12 (1), 86{114.

Einav, L., C. Farronato, and J. Levin , \Peer-to-Peer Markets," Annual Review of Economics, 2016, 8, 615{635.

Fabra, N., D. Rapson, M. Reguant, and J. Wang , \Estimating the Elasticity to Real Time Pricing: Evidence from the Spanish Electricity Market," American Economic Association Papers & Proceedings, forthcoming.

Farronato, C. and A. Fradkin , \The Welfare E ects of Peer Entry in the Accommodation Market: The Case of Airbnb," NBER Working Paper 24361, 2018.

Faruqui, A. and S. Sergici , \Dynamic Pricing of Electricity in the mid-Atlantic Region: Econometric Results from the Baltimore Gas and Electric Company Experiment," Journal of Regulatory Economics, 2011, 40 (1), 82{109.

Feldman, N., J. Goldin, and 2018 Homono , \Raising the Stakes: Experimental Evidence on the Endogeneity of Taxpayer Mistakes," National Tax Journal, 2018, 71 (2), 201{30.

_ , P. Katuscak, and L. Kawano , \Taxpayer Confusion: Evidence form the Child Tax Credit," American Economic Review, 2016, 106 (3), 807{35.

Finkelstein, A. , \E-ZTax: Tax Salience and Tax Rates," Quarterly Journal of Economics, 2009, 124 (3), 969{1010.

Gabaix, X. , \Behavioral Inattention," Handbook of Behavioral Economics: Applications and Foundations 1, 2019, 2, 261{343.

Gilbert, B. and J. Gra Zivin , \Dynamic Salience with Intermittent Billing: Evidence from Smart Electricity Meters," Journal of Economic Behavior & Organization, 2014, 107, 176{90.

Gillan, J. , \Dynamic Pricing, Attention, and Automation: Evidence from a Field Experiment in Electricity Consumption," working paper, 2018.

Gillingham, K. and K. Palmer , \Bridging the Energy E ciency Gap: Policy Insights from Economic Theory and Empirical Evidence," Review of Environmental Economics and Policy, 2014, 8 (1), 18{38.

_ , S. Houde, and A. van Benthem , \Consumer Myopia in Vehicle Purchases: Evidence from a Natural Experiment," American Economic Journal: Economic Policy, forthcoming.

Goldin, J. and T. Homono , \Smoke Gets in Your Eyes: Cigarette Tax Salience and Regressivity," American Economic Journal: Economic Policy, 2013, 5 (1), 302{36.

Hall, J., J. Horton, and D. Knoep e , \Ride-Sharing Markets Re-Equilibriate," Working Paper, 2020.

Harding, M. and C. Lamarche , \Empowering Consumers through Data and Smart Meter Technology: Experimental Evidence on the Consequences of Time-of-Use Electricity Pricing Policies," Journal of Policy Analysis and Management, 2016, 35 (4), 906{31.

Ito, K. , \Do Consumers Respond to Marginal or Average Prices? Evidence from Nonlinear Electricity Pricing," American Economic Review, 2014, 104 (2), 537{63.

_ , T. Ida, and M. Tanaka , \Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand," American Economic Journal: Economic Policy, 2018, 10 (1), 240{267.

Jessoe, K. and D. Rapson , \Knowledge is (less) power: Experimental evidence from residential energy use," American Economic Review, 2014, 104 (2), 537{563.

Joskow, P. and C. Wolfram , \Dynamic Pricing of Electricity," American Economic Review: Papers and Proceedings, 2012, 102 (3), 381{385.

Martin, S. and N. Rivers , \Information Provision, Market Incentives, and Household Electricity Consumption: Evidence from a Large-Scale Field Deployment," Journal of the Association of Environmental and Resource Economists, 2018, 5 (1).

McKenna, E., J. Pless, and S. Darby , \Solar Photovoltaic Self-Consumption in the UK Residential Sector," Energy Policy, 2018, 118, 482{491.

Myers, E. , \Are Home Buyers Inattentive? Evidence from Capitalization of Energy Costs," American Economic Journal: Economic Policy, 2019, 11 (2), 165{188.

_ and M. Souza , \Social Comparison Nudges without Monetary Incentives: Evidence from Home Energy Reports," Journal of Environmental Economics and Management, 2020, 101.

Nauze, A. La , \Power from the People: Rooftop Solar and a Downward-Sloping Supply of Electricity," Journal of the Association of Energy and Resource Economists, 2019, 6 (6), 1135{1168.

Prest, B. , \Peaking Interest: How Awareness Drives the E ectiveness of Time-of-Use Electricity Pricing," Journal of the Association of Environmental and Resource Economists, 2020, 7 (1), 203{143.

Puller, S. and J. West , \E cient Retail Pricing in Electricity and Natural Gas Markets," American Economic Review: Papers and Proceedings, 2013, 103 (3), 350{355.

Reiss, P. and M. White , \Household Electricity Demand, Revisited," Review of Economic Studies, 2005, 72 (2), 853{83.

Sallee, J., S. West, and W. Fan , \Do Consumers Recognize the Value of Fuel Economy? Evidence from Used Car Prices and Gasoline Price Fluctuations," Journal of Public Economics, 2016, 135, 61{73.

Seamans, R. and F. Zhu , \Responses to Entry in Multi-Sided Markets: The Impact of Craigslist on Local Newspapers," Management Science, 2014, 60 (2), 476{493.

Sexton, S. , \Automatic Bill Payment and Salience E ects: Evidence from Electricity Consumption," Review of Economics and Statistics, 2015, (2), 229{241.

Taylor, T., P. Schwarz, and J. Cochell , \24/7 Hourly Response to Electricity Real-Time Pricing with up to Eight Summers of Experience," Journal of Regulatory Economics, 2005, 27 (3), 235{62.

Wichman, C. , \Information Provision and Consumer Behavior: A Natural Experiment in Billing Frequency," Journal of Public Economics, 2017, 152, 13{33.

Widen, J., E. Wackelgard, J. Paatero, and P. Lund , \Impacts of Di erent Data Averaging Times on Statistical Analysis of Distributed Domestic Photovoltaic Systems," Solar Energy, 2010, 84 (3), 492{500.

Wolak, F.A. , \Do Residential Customers Respond to Hourly Prices? Evidence from a Dynamic Pricing Experiment," American Economic Review: Papers and Proceedings, 2011, 101 (3), 83{7.

MAIN TEXT TABLES

Table I: Average E ects of Solar Generation on Electricity Consumption

Dep. Var.:	Gross Consumption (1)	Imports (Net) (2)	Self-Consumption (3)
Panel A: 15-Minute Level			
Solar Generation (kWh)	0.071***	-0.172***	0.243***
	(0.027)	(0.014)	(0.019)
Obs.	4,700,124	4,700,124	4,700,124
Mean Dep. Var.	0.148	0.093	0.055
Household-Day FEs	x	x	x
Month-Hour-DoW FEs	x	x	x
Panel B: Hourly Level			
Solar Generation (kWh)	0.075**	-0.169***	0.245***
	(0.031)	(0.017)	(0.021)
Obs.	1,218,229	1,218,229	1,218,229
Mean Dep. Var.	0.590	0.382	0.209
Household-Day FEs	x	x	x
Month-Hour-DoW FEs	x	x	x
Panel C: Daily Level			
Solar Generation (kWh)	0.073	-0.206***	0.280***
	(0.055)	(0.033)	(0.034)
Obs.	104,692	104,692	104,692
Mean Dep. Var.	12.457	10.267	2.190
Household FEs	x	x	x
Day FEs	x	x	x
Household-Month FEs	x	x	x
Household-DoW FEs	x	x	x

Notes: Estimates report the e ect of a 1 kWh increase in solar generation on gross consumption, imported (net) consumption, and self-consumption when using 15-minute, hourly, and daily level data. An increase in solar corresponds to a decrease in price. Hourly weather controls are included for Panel A and Panel B regressions. Standard errors are clustered by household. Asterisks denote * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table II: Load-Shifting Main Results

Dep. Var.:	Peak Gross (1)	Peak Imports (2)	Evening Gross (3)	Evening Imports (4)
Daytime Solar (kWh)	-0.027**	-0.027***	-0.035**	-0.034***
	(0.013)	(0.009)	(0.015)	(0.012)
Obs.	104,692	104,692	104,692	104,692
Mean Dep. Var.	3.107	2.920	5.253	5.077
Household FEs	x	x	x	x
Day FEs	x	x	x	x
Household-Month FEs	x	x	x	x
Household-DoW FEs	x	x	x	x

Notes: Estimates report load-shifting results whereby we estimate a household's total daytime solar generation (9am to 4pm) on peak (4pm to 8pm) and total evening (4pm to midnight) gross consumption and imports (net consumption) using daily-level data. Peak solar generation is included as a control variable in all regressions. Standard errors are clustered by household. Asterisks denote * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table III: Main E ects of IHDs on Gross Consumption and Load-Shifting

	(1)	(2)	(3)	(4)
Panel A: Gross Consumption	15-Minute Level		Hourly Level	
Solar Generation (kWh)	0.060*	0.057*	0.065*	0.061
	(0.033)	(0.033)	(0.038)	(0.038)
Solar (kWh) * IHD	0.018	0.022	0.018	0.022
	(0.037)	(0.037)	(0.038)	(0.038)
Uncertainty		0.032*		0.122**
		(0.017)		(0.061)
Uncertainty * IHD		-0.035*		-0.141*
		(0.019)		(0.073)
Obs.	4,700,124	4,695,932	1,218,229	1,218,229
Mean Dep. Var.	0.148	0.148	0.590	0.590
Household-Day FEs	x	x	x	x
Month-Hour-DoW FEs	x	x	x	x
Panel B: Daily Load-Shifting	Peak Gross		Peak Import	
Daytime Solar (kWh)	-0.027**	-0.041**	-0.027**	-0.042***
	(0.013)	(0.020)	(0.011)	(0.015)
Daytime Solar * IHD	-0.001	0.011	0.000	0.012
	(0.007)	(0.009)	(0.007)	(0.008)
Uncertainty		0.126*		0.129**
		(0.075)		(0.063)
Uncertainty * IHD		-0.101*		-0.098*
		(0.056)		(0.050)
Obs.	104,692	104,692	104,692	104,692
Mean Dep. Var.	3.107	3.107	2.920	2.920
Household FEs	x	x	x	x
Day FEs	x	x	x	x
Household-Month FEs	x	x	x	x
Household-DoW FEs	x	x	x	x

Notes: Estimates report baseline ndings for how IHDs change behavior entirely through the uncertainty channel. Panel A provides results for when using 15-minute and hourly data and Panel B provides daily load-shifting results. Hourly weather controls included for Panel A and peak generation included as a control in Panel B. Standard errors are clustered by household. Asterisks denote * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table IV: Weekend vs. Weekday E ects and Imperfect Appliance Scheduling

	Weekends (1)	Weekdays (2)	Weekends (3)	Weekdays (4)
Panel A: Gross Consumption	15-Minute Level		Hourly	
Solar Generation (kWh)	0.061*	0.055*	0.063*	0.060
	(0.032)	(0.033)	(0.037)	(0.039)
Solar (kWh) * IHD	0.027	0.019	0.030	0.019
	(0.036)	(0.038)	(0.037)	(0.039)
Uncertainty	0.055**	0.021	0.240**	0.069
	(0.028)	(0.016)	(0.113)	(0.057)
Uncertainty * IHD	-0.039	-0.032	-0.183	-0.121
	(0.029)	(0.020)	(0.120)	(0.075)
Obs.	1,345,905	3,350,027	348,628	869,601
Mean Dep. Var.	0.158	0.144	0.629	0.575
Household-Day FEs	x	x	x	x
Month-Hour-DoW FEs	x	x	x	x
Panel B: Daily Load-Shifting	Peak Gross		Peak Imports	
Daytime Solar (kWh)	-0.047**	-0.037*	-0.049***	-0.036**
	(0.019)	(0.021)	(0.016)	(0.015)
Daytime Solar * IHD	0.004	0.012	0.006	0.013
	(0.012)	(0.009)	(0.012)	(0.008)
Uncertainty	0.047	0.141	0.061	0.140*
	(0.074)	(0.087)	(0.070)	(0.071)
Uncertainty * IHD	-0.018	-0.125*	-0.028	-0.119**
	(0.071)	(0.066)	(0.068)	(0.058)
Obs.	29,864	74,724	29,864	74,724
Mean Dep. Var.	3.115	3.104	2.920	2.919
Household FEs	x	x	x	x
Day FEs	x	x	x	x
Household-Month FEs	x	x	x	x
Household-DoW FEs	x	x	x	x

Notes: Estimates report the e ect of a 1 kWh increase in solar generation on gross consumption, imported (net) consumption, and self-consumption when using 15-minute, hourly, and daily level data separately for weekdays and weekends. An increase in solar corresponds to a decrease in price. Hourly weather controls included for Panel A and peak generation included as a control in Panel B. Standard errors are clustered by household. Asterisks denote * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table V: Asymmetry Suggests Errors are Due to \Excitement"

	Solar Decrease (1)	Solar Increase (2)	Solar Decrease (3)	Solar Increase (4)
Panel A: Gross Consumption	15-Minute Level		Hourly	
Solar Generation (kWh)	0.054	0.057*	0.061	0.070
	(0.033)	(0.033)	(0.040)	(0.045)
Solar (kWh) * IHD	0.016	0.032	0.016	0.030
	(0.038)	(0.036)	(0.039)	(0.038)
Uncertainty	0.014	0.049**	0.106	0.192**
	(0.019)	(0.023)	(0.083)	(0.081)
Uncertainty * IHD	-0.008	-0.057**	-0.047	-0.142
	(0.028)	(0.025)	(0.108)	(0.090)
Obs.	2,265,456	2,388,606	557,721	651,013
Mean Dep. Var.	0.155	0.142	0.633	0.554
Household-Day FEs	x	x	x	x
Month-Hour-DoW FEs	x	x	x	x
Panel B: Daily Load-Shifting	Peak Gross		Peak Imports	
Daytime Solar	-0.010	-0.058***	-0.016	-0.057***
	(0.017)	(0.021)	(0.017)	(0.017)
Daytime Solar * IHD	0.014	0.014	0.014	0.015
	(0.015)	(0.011)	(0.014)	(0.011)
Uncertainty	-0.001	0.209**	0.015	0.212**
	(0.075)	(0.097)	(0.074)	(0.084)
Uncertainty * IHD	-0.074	-0.170**	-0.073	-0.160**
	(0.068)	(0.086)	(0.067)	(0.077)
Obs.	51,067	53,376	51,067	53,376
Mean Dep. Var.	3.132	3.082	2.970	2.870
Household FEs	x	x	x	x
Day FEs	x	x	x	x
Household-Month FEs	x	x	x	x
Household-DoW FEs	x	x	x	x

Notes: Estimates report e ects of solar generation separately for observations for which solar generation increased or decreased relative to the previous period. The ndings indicate that the main e ects are driven by cases when solar increased. Hourly weather controls included for Panel A and peak generation included as a control in Panel B. Standard errors are clustered by household. Asterisks denote * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table VI: Learning and Adjustment Happens Quickly With and Without IHDs

Dep. Var.:	Peak Gross	Peak Import	Peak Gross	Peak Import	Peak Gross	Peak Import
Days in Trial:	< 90 Days		90-180 Days		> 180 Days	
	(1)	(2)	(3)	(4)	(5)	(6)
Daytime Solar	0.007	0.003	-0.035	-0.047*	-0.040*	-0.038**
	(0.018)	(0.018)	(0.025)	(0.025)	(0.022)	(0.016)
Daytime Solar * IHD	-0.035**	-0.037**	0.007	0.013	0.011	0.011
	(0.018)	(0.018)	(0.019)	(0.019)	(0.009)	(0.008)
Uncertainty	-0.062	-0.052	0.143	0.182	0.092	0.089
	(0.079)	(0.079)	(0.121)	(0.119)	(0.077)	(0.062)
Uncertainty * IHD	0.207**	0.209**	-0.076	-0.106	-0.089	-0.077
	(0.086)	(0.088)	(0.111)	(0.110)	(0.064)	(0.055)
Obs.	14,887	14,887	18,348	18,348	70,975	70,975
Mean Dep. Var.	3.211	3.150	3.275	3.150	3.033	2.802
Household FEs	x	x	x	x	x	x
Day FEs	x	x	x	x	x	x
Household-Month FEs	x	x	x	x	x	x
Household-DoW FEs	x	x	x	x	x	x

Notes: Standard errors are clustered by household. Asterisks denote * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

MAIN TEXT FIGURES

Figure I: Power Flows for a UK Household with Solar PV on June 6^{th}; 2013

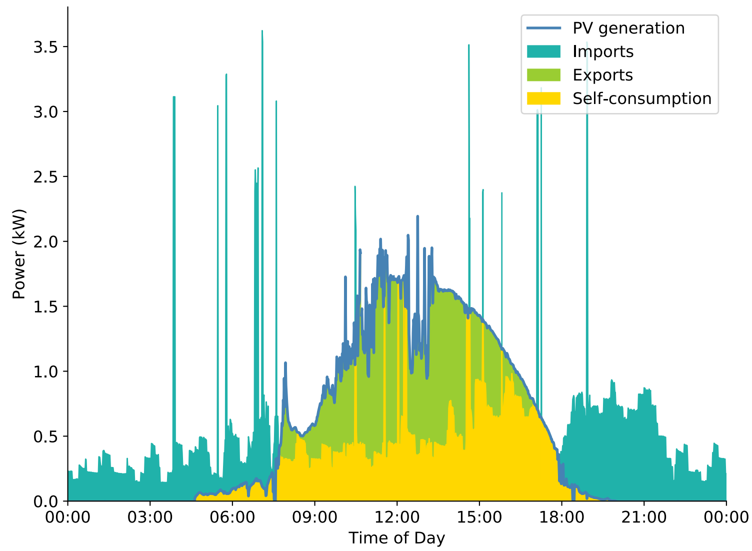

Figure II: Hourly Effects of Solar Generation on Gross Consumption (kWh)

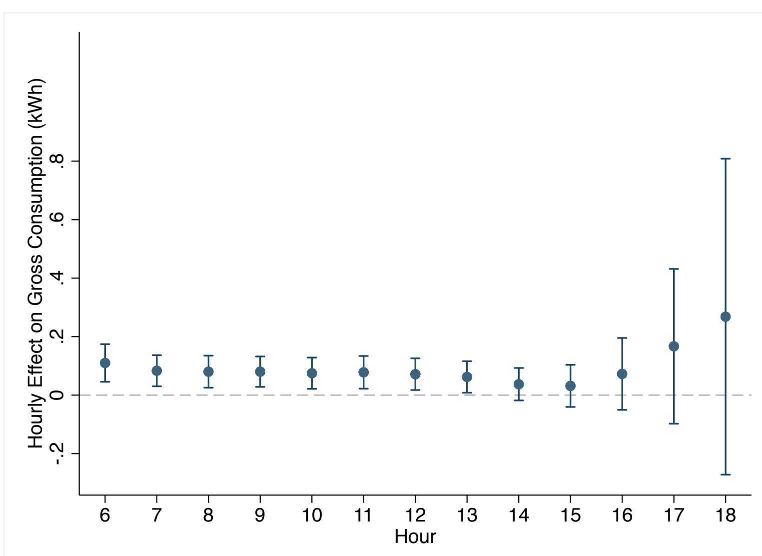

A Appendix: Data Preparation—Online Only

Several steps were taken to prepare the data for analysis. We begin with household-level 1-minute resolution data measuring solar PV generation and net electricity demand, which we matched for each household by the observation timestamp. Imports and exports were calculated as described in Section II at the 1-minute resolution and recorded as missing if any variables used in the calculations were missing data for that household-minute combination. About 1.7 percent and 1.2 percent of the original solar generation and net demand readings were missing data, and the split of missing data was relatively even between IHD and non-IHD households.

We aggregated solar PV generation, net demand, imports, and exports to the 15-minute time-resolution so that it was more manageable for regression analysis and converted the units from power to energy. If data were missing at the 1-minute level for any of the minutes within a 15-minute interval, we marked this 15-minute interval as missing data so that the aggregate data were only populated for intervals where 1-minute data were complete. Since imports and exports were calculated with the 1-minute disaggregated data, matching time-coincident power flows at a very granular level, the remaining derived variables—gross demand and self-consumption—could be calculated accurately using the 15-minute data. This process of starting with the 1-minute data eliminates inaccuracies that otherwise are accrued by using time-averaged data.

We conducted a final check for potential errors in the 15-minute data. Solar PV generation readings of less than -0.05 kW were interpreted as errors in terms of sign (and not magnitude) and thus we converted these into positive power flows. Those that read between 0 and -0.05 kW are reasonable due to parasitic losses—energy losses that occur due to the operation of inverters in the system. Solar PV and export readings in excess of 5 kW also likely reflect errors, considering all PV systems in the dataset are less than 4 kW in size. We omitted these observations as well as any cases for which (positive) exports exceeded solar PV generation.

We dropped a small number of observations that were duplicates in all variables, including household ID, date, and time. There was one household ID with meter readings of zeros across all variables, which we dropped as well. Overall, this process resulted in omitting less than 1 percent of the 15-minute level observations, and those omitted were reasonably split between IHD and non-IHD households. The final sample consists of about 14.3 million observations.

B Appendix: Additional Tables—Online Only

Table B.1: Summary Statistics for Households With and Without IHDs

	Means			Standard Devs.		Observations	
	IHDs	Non-IHDs	Difference	IHDs	Non-IHDs	IHDs	Non-IHDs
	(1)	(2)	(3)	(4)	(5)	(6)	
Panel A: 15-Minute Level							
Solar Generation	0.122	0.114	0.008***	0.139	0.125	2.4m	2.3m
Gross Consumption	0.154	0.141	0.013***	0.184	0.189	2.4m	2.3m
Imports (Net Consumption)	0.097	0.090	0.007***	0.169	0.173	2.4m	2.3m
Self-Consumption	0.058	0.051	0.007***	0.066	0.057	2.4m	2.3m
Panel B: Hourly Level							
Solar Generation	0.465	0.436	0.029***	0.527	0.477	630k	588k
Gross Consumption	0.614	0.565	0.049***	0.601	0.631	630k	588k
Imports (Net Consumption)	0.394	0.368	0.026***	0.545	0.576	630k	588k
Self-Consumption	0.220	0.196	0.024***	0.243	0.207	630k	588k
Panel C: Daily Level							
Solar Generation	4.98	4.58	0.40***	4.51	4.01	54k	51k
Daytime Solar (9am-4pm)	4.41	4.07	0.34***	3.75	3.34	54k	51k
Gross Consumption	13.0	11.9	1.10***	7.06	7.40	54k	51k
Imports (Net Consumption)	10.7	9.82	0.88***	6.75	7.00	54k	51k
Self-Consumption	2.34	2.03	0.31***	2.07	1.70	54k	51k
Peak Gross Consumption	3.23	2.98	0.25***	2.09	2.13	54k	51k
Peak Imports	3.04	2.79	0.25***	2.14	2.14	54k	51k

Notes: Reports summary statistics for households with and without IHDs. Means and standard deviations are in kilowatt hours (kWh) unless otherwise noted. Peak hours are from 4pm to 8pm. In Panels A and B, only observations for which solar generation is positive are included. Statistical significance of the differences in the means are reported. Asterisks denote $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Table B.2: 15-Minute and Hourly Effects Are Not Driven by a Mechanical Relationship

Dep. Var.:	Gross 9am-4pm (1)	Gross 10am-4pm (2)	Gross 12pm-3pm (3)	Gross >Median (4)	Gross >Daytime Median (5)
Panel A: 15-Minute Level					
Solar Generation (kWh)	0.059***	0.061***	0.059***	0.063***	0.066**
	(0.020)	(0.021)	(0.020)	(0.024)	(0.026)
Obs.	3,404,114	2,622,292	1,349,928	2,348,884	1,911,218
Mean Dep. Var.	0.152	0.156	0.155	0.151	0.152
Panel B: Hourly Level					
Solar Generation (kWh)	0.060***	0.065***	0.068***	0.059**	0.059**
	(0.023)	(0.024)	(0.025)	(0.025)	(0.025)
Obs.	863,412	658,694	335,489	928,891	849,306
Mean Dep. Var.	0.611	0.626	0.619	0.597	0.598
Household-Day FEs	x	x	x	x	x
Month-Hour-DoW FEs	x	x	x	x	x

Notes: Estimates report the effect of a 1 kWh increase in solar generation on gross consumption hours in the day when it's unlikely that baseload demand alone accounts for all consumption. In Columns 1 through 3, different ranges of hours are included. In Columns 4 and 5, observations above the median of solar generation within the estimation sample or within daytime hours, respectively. Hourly weather controls are also included in all regressions. Standard errors are clustered by household. Asterisks denote $^*p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Table B.3: Load-Shifting Effects Are Not Driven by Mechanical Relationship

Dep. Var.:	Peak Gross (1)	Peak Imports (2)	Evening Gross (3)	Evening Imports (4)
Panel A: E ect of Peak Solar Only				
Peak Solar (kWh)	0.331	-0.141	0.374	-0.119
	(0.299)	(0.203)	(0.332)	(0.219)
Obs.	104,692	104,692	104,692	104,692
Mean Dep. Var.	3.107	2.920	5.253	5.077
Panel B: When Including Daytime Solar				
Peak Solar (kWh)	0.355	-0.117	0.405	-0.089
	(0.305)	(0.206)	(0.338)	(0.222)
Daytime Solar (kWh)	-0.027**	-0.027***	-0.035**	-0.034***
	(0.013)	(0.009)	(0.015)	(0.012)
Obs.	104,692	104,692	104,692	104,692
Mean Dep. Var.	3.107	2.920	5.253	5.077
Household FEs	x	x	x	x
Day FEs	x	x	x	x
Household-Month FEs	x	x	x	x
Household-DoW FEs	x	x	x	x

Notes: Estimates show that the findings are not mechanically driven by peak sola generation. Standard errors are clustered by household. Asterisks denote $^*p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Table B.4: Additional Robustness Checks

Dep. Var.:	15-min Gross (1)	Hourly Gross (2)	Daily Peak Gross (3)	Daily Peak Imports (4)	Daily Peak Gross (5)	Daily Peak Imports (6)
Solar Generation (kWh)	0.073**	0.075**				
	(0.029)	(0.032)				
Daytime Solar (kWh)			-0.027**	-0.026***	-0.026**	-0.026***
			(0.013)	(0.009)	(0.013)	(0.009)
Obs.	4,700,168	1,218,553	104,692	104,692	103,417	103,417
Mean Dep. Var.	0.148	0.590	3.107	2.920	3.106	2.918
Household-Day FEs	x	x				
Household FEs			x	x	x	x
Day FEs			x	x	x	x
Month-Hour-DoW FEs	x	x				
Household-Month FEs			x	x	x	x
Household-DoW FEs			x	x	x	x
No. of months in trial			x	x	x	x

Notes: Estimates the effect of peak solar generation on peak and evening gross consumption and imports. There are zero statistically significant effects, indicating that the main load-shifting results are not driven by peak solar generation. Weather controls are included for the regressions in Columns 1 and 2 and peak generation included as a control for the regressions in Column 3-6. Standard errors are clustered by household. Asterisks denote *p <0.10, **p <0.05, ***p <0.01.

Table B.5: Machine Learning Results: Effects on Gross Consumption

	(1)	(2)	(3)
Panel A: 15-Minute Level			
Solar Generation (kWh)	0.067***	0.067***	0.065***
	(0.022)	(0.022)	(0.021)
Solar Generation * IHD	0.006	0.006	0.009
	(0.036)	(0.036)	(0.035)
Uncertainty			0.025*
			(0.014)
Uncertainty * IHD			-0.027
			(0.018)
Obs.	4,225,640	4,225,640	4,221,742
Mean Dep. Var.	0.163	0.163	0.163
Household-Day FEs	x	x	x
Month-Hour-DoW FEs	x	x	x
Panel B: Hourly Level			
Solar Generation (kWh)	0.073***	0.067***	0.064***
	(0.022)	(0.023)	(0.022)
Solar Generation * IHD		0.012	0.016
		(0.037)	(0.036)
Uncertainty			0.092*
			(0.050)
Uncertainty * IHD			-0.118*
			(0.065)
Obs.	1,081,924	1,081,924	1,081,924
Mean Dep. Var.	0.654	0.654	0.654
Household-Day FEs	x	x	x
Month-Hour-DoW FEs	x	x	x

Notes: Estimates report results from when estimating the baseline model of Equation 3 but while using the data generated through the matching and machine learning method. Hourly weather controls are included in all regressions. Standard errors are clustered by household. Asterisks denote *$p<0.10$, **$p<0.05$, ***$p<0.01$.

Table B.6: IHD Effects on Load-Shifting for All Evening Hours

	Evening Gross (1)	Evening Gross (2)	Evening Imports (3)	Evening Imports (4)
Daytime Solar (kWh)	-0.037**	-0.051**	-0.037**	-0.051***
	(0.017)	(0.024)	(0.014)	(0.020)
Daytime Solar * IHD	0.003	0.015	0.004	0.016
	(0.009)	(0.012)	(0.009)	(0.012)
Uncertainty		0.121		0.126*
		(0.081)		(0.071)
Uncertainty * IHD		-0.106		-0.105*
		(0.065)		(0.060)
Obs.	104,692	104,692	104,692	104,692
Mean Dep. Var.	5.253	5.253	5.077	5.077
Household FEs	x	x	x	x
Day FEs	x	x	x	x
Household-Month FEs	x	x	x	x
Household-DoW FEs	x	x	x	x

Notes: Load-shifting results for when considering all evening hours (4pm to midnight) as the dependent variable rather than just peak hours as done in the main text regressions. Peak generation is included as a control in all regressions. Standard errors are clustered by household. Asterisks denote *p <0.10, **p <0.05, ***p <0.01.

Table B.7: Effects on Gross Consumption Based on Time of Day and Week

	Weekend 8am-11am (1)	Weekend 12pm-4pm (2)	Weekend 4pm-9pm (3)	Weekday 8am-11am (4)	Weekday 12pm-4pm (5)	Weekday 4pm-9pm (6)
Solar Generation (kWh)	0.057**	0.026	0.191*	0.040**	0.041	0.218
	(0.024)	(0.025)	(0.112)	(0.018)	(0.027)	(0.167)
Solar (kWh) * IHD	0.030	0.062	-0.051	0.039	0.035	-0.130
	(0.039)	(0.038)	(0.103)	(0.034)	(0.041)	(0.144)
Uncertainty	0.075**	-0.019	-0.316	0.030	-0.007	-0.232
	(0.032)	(0.027)	(0.249)	(0.022)	(0.018)	(0.223)
Uncertainty * IHD	-0.071*	0.050	0.243	-0.026	0.001	0.309
	(0.039)	(0.036)	(0.282)	(0.027)	(0.025)	(0.263)
Obs.	480,267	492,437	193,379	1,201,143	1,227,065	474,729
Mean Dep. Var.	0.163	0.174	0.168	0.141	0.151	0.168
Household-Day FEs	x	x	x	x	x	x
Month-Hour-DoW FEs	x	x	x	x	x	x

Notes: Estimates report the effects for weekends and weekdays during different times of the day when using the 15-minute level data. Uncertainty and information effects at this time resolution are driven by weekend mornings. Weather controls are included in all regressions. Standard errors are clustered by household. Asterisks denote *p <0.10, **p <0.05, ***p <0.01.

Table B.8: (No) Effects of Lagged Solar Generation on Load-Shifting

Dep. Var.:	Peak Gross	Peak Gross	Peak Gross	Peak Imports	Peak Imports	Peak Imports
	(1)	(2)	(3)	(4)	(5)	(6)
Lagged Daytime Solar	0.001	0.004	0.007	0.001	0.003	0.005
	(0.004)	(0.006)	(0.006)	(0.003)	(0.006)	(0.006)
Lagged Daytime Solar * IHD		-0.001	-0.002		0.000	-0.002
		(0.006)	(0.006)		(0.006)	(0.005)
Lagged Uncertainty		-0.019	-0.030		-0.008	-0.020
		(0.038)	(0.040)		(0.036)	(0.037)
Lagged Uncertainty * IHD		-0.006	0.004		-0.012	-0.002
		(0.038)	(0.039)		(0.036)	(0.037)
Daytime Solar			-0.040**			-0.041***
			(0.020)			(0.015)
Daytime Solar * IHD			0.011			0.011
			(0.008)			(0.008)
Uncertainty			0.124			0.126**
			(0.075)			(0.063)
Uncertainty * IHD			-0.100*			-0.097*
			(0.057)			(0.050)
Obs.	104,402	104,402	104,402	104,402	104,402	104,402
Mean Dep. Var.	3.108	3.108	3.108	2.921	2.921	2.921
Household FEs	x	x	x	x	x	x
Day FEs	x	x	x	x	x	x
Household-Month FEs	x	x	x	x	x	x
Household-DoW FEs	x	x	x	x	x	x
No. of months in trial	x	x	x	x	x	x

Notes: Estimates report load-shifting results when including the previous day's daytime solar generation and uncertainty, showing that all load-shifting effects are occurring based on prices in the current period. The dependent variable is peak gross consumption in Columns 1-3 and peak imports in Columns 4-6. The current period's peak generation is included as a control in all regressions. Standard errors are clustered by household. Asterisks denote $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Table B.9: (No) Effects of Morning Solar Generation on Load-Shifting

	Peak Gross (1)	Peak Gross (2)	Peak Gross (3)	Peak Imports (4)	Peak Imports (5)	Peak Imports (6)
Morning Solar	0.005	0.005	0.059	-0.025	-0.033	0.022
	(0.054)	(0.068)	(0.085)	(0.043)	(0.057)	(0.064)
Morning Solar * IHD		0.001	-0.007		0.011	-0.002
		(0.045)	(0.042)		(0.044)	(0.041)
Daytime Solar			-0.043*			-0.042**
			(0.023)			(0.017)
Daytime Solar * IHD			0.011			0.012
			(0.009)			(0.009)
Uncertainty			0.129*			0.130**
			(0.078)			(0.064)
Uncertainty * IHD			-0.100*			-0.098*
			(0.055)			(0.050)
Obs.	104,692	104,692	104,692	104,692	104,692	104,692
Mean Dep. Var.	3.107	3.107	3.107	2.920	2.920	2.920
Household FEs	x	x	x	x	x	x
Day FEs	x	x	x	x	x	x
Household-Month FEs	x	x	x	x	x	x
Household-DoW FEs	x	x	x	x	x	x

Notes: Estimates report the load-shifting effects when including the current day's morning solar generation, confirming that all of the effects are driven by daytime solar generation and uncertainty. Peak generation is included as a control in all regressions. Standard errors are clustered by household. Asterisks denote $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.