Grantham Research Institute on Climate Change and the Environment

Centre for Climate Change Economics and Policy

Consumer-driven virtual power plants: A field experiment on the adoption and use of a prosocial technology

Quentin Coutellier, Greer Gosnell, Zeynep Gürgüç, Ralf Martin and Mirabelle Muûls

September 2020

Centre for Climate Change Economics and Policy Working Paper No. 379
ISSN 2515-5709 (Online)

Grantham Research Institute on Climate Change and the Environment
Working Paper No. 350
ISSN 2515-5717 (Online)

ESRC Economic & Social Research Council

The Grantham Foundation for the Protection of the Environment

University of Leeds

LSE The London School of Economics and Political Science

The Centre for Climate Change Economics and Policy (CCCEP) was established by the University of Leeds and the London School of Economics and Political Science in 2008 to advance public and private action on climate change through innovative, rigorous research. The Centre is funded by the UK Economic and Social Research Council. Its third phase started in October 2018 with seven projects:

1. Low-carbon, climate-resilient cities
2. Sustainable infrastructure finance
3. Low-carbon industrial strategies in challenging contexts
4. Integrating climate and development policies for 'climate compatible development'
5. Competitiveness in the low-carbon economy
6. Incentives for behaviour change
7. Climate information for adaptation

More information about CCCEP is available at www.cccep.ac.uk

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training. The Institute is funded by the Grantham Foundation for the Protection of the Environment and a number of other sources. It has 11 broad research areas:

1. Climate change adaptation and resilience
2. Climate change governance, legislation and litigation
3. Environmental behaviour
4. Environmental economic theory
5. Environmental policy evaluation
6. International climate politics
7. Science and impacts of climate change
8. Sustainable finance
9. Sustainable natural resources
10. Transition to zero emissions growth
11. UK national and local climate policies

More information about the Grantham Research Institute is available at www.lse.ac.uk/GranthamInstitute

Suggested citation:
Coutellier Q, Gosnell G, Gürgüç Z, Martin R and Muûls M (2020) *Consumer-driven virtual power plants: A field experiment on the adoption and use of a prosocial technology*. Centre for Climate Change Economics and Policy Working Paper 379/Grantham Research Institute on Climate Change and the Environment Working Paper 350. London: London School of Economics and Political Science

Consumer-Driven Virtual Power Plants: A Field Experiment on the Adoption and Use of a Prosocial Technology*

Quentin Coutellier† Greer Gosnell‡ Zeynep Gürgüç§ Ralf Martin¶
Mirabelle Muûls‖

September 14, 2020

Abstract

Demand response strategies that rely on individual behavior change have consistently demonstrated that flexibility in energy demand exists, but that there are constraints to demand shifting such as limited attention or rebound effects. This paper aims to analyze residential energy demand flexibility through a series of field experiments on adoption and usage of WiFi-enabled smart plugs that largely overcome the need for home dwellers' effort and attention altogether. With 169 active participants, we study the adoption and user interaction of devices for automated dynamic energy demand-side management. Our design allows us to explore whether and how devices and incentive systems may be designed and deployed to improve load balancing in the centralized energy grid. Such a system would reduce the need for inefficient and often carbon-intensive back-up power generators during periods of peak demand, and increase the possible share of energy supplied by intermittent renewable energy sources. Our results suggest that users are more inclined to participate if switch-off events are sufficiently long to provide a meaningful probability of winning lottery prizes, but that higher frequency of such events may reduce participation intensity.

*We are grateful to colleagues, in particular Anomitro Chatterjee, for thoughtful feedback that helped us to refine our paper. We also thank Pantelis Koutroumpis, Tom Berman, and Mark Woodbridge for their assistance in providing the technological foundations for the research. We are also grateful to a number of university personnel for their planning logistical support, in particular Caroline Chipperfield, Julie Piesse, Daniel Reeves, James Sandie, Peter Tajok and Matthew Wilkinson. Greer Gosnell acknowledges support from the Grantham Foundation for the Protection of the Environment and from the UK Economic and Social Research Council through the Centre for Climate Change Economics and Policy. Gosnell is a beneficiary of an AXA Research Fund postdoctoral grant as well as a GEMCLIME research secondment funded by the European Union's Horizon 2020 Research and Innovation programme under the Marie Skłodowska-Curie grant agreement No. 681228. The Grantham Institute for Climate Change and the Environment supported the project through pump prime funding. We are grateful to the UK Engineering and Physical Sciences Research Council (EPSRC) for funding through the IDLES grant EP/R045518/1.

†Imperial College Business School, Imperial College London; Department of Geography and Environment, London School of Economics and Political Science (LSE), Houghton Street, WC2A 2AE, UK, United Kingdom.

‡The Payne Institute for Public Policy, Colorado School of Mines, 1500 Illinois St, Golden, CO, 80401, USA; Grantham Research Institute on Climate Change and the Environment, LSE. Corresponding author email: ggosnell@mines.edu.

§Imperial College Business School, Imperial College London

¶Imperial College Business School, Imperial College London; Centre for Economic Performance, LSE.

‖Imperial College Business School, Imperial College London; Centre for Economic Performance, LSE.

1 Introduction

Increasing global attention on the threat posed by climate change has underscored the need for rapid development and deployment of innovative solutions to catalyze a sustainable low-carbon energy transition. Policymakers' ability to reduce overall energy consumption and manage energy market volatility is becoming increasingly important in this context. As decarbonization of electricity supply will rely on large-scale deployment of variable renewable energy sources, smarter management of electricity consumption, i.e. demand-side management (DSM), will prove essential to reach the global targets set out in the Paris Agreement (Gielen et al., 2019).

Internet connectivity and electronic innovation have opened the door for development of DSM systems to accommodate fluctuations in energy supply and demand. This development expands suppliers' and grid operators' toolkits beyond supply-side load balancing mechanisms that often themselves require continuous emissions to maintain ramping capabilities of back-up power plants. Excessive energy demand can be shifted from peak to off-peak hours, when electricity typically costs less, by switching off some appliances for limited periods of time.[1] Thus, residential and commercial consumers have the ability—and, under a growing number of energy plans, the incentive—to contribute to a more flexible energy grid (Carbon Trust and Imperial College London, 2016).

Our research considers the extent to which energy consumers' adoption of an inexpensive Internet of Things (IoT)[2] technology—specifically the smart plug—can allow for increased flexibility in energy demand given an embedded stock of energy-consuming goods and durables in the economy.[3] We posit that scaled utilization of such technology enables the formation of sizable and robust virtual power plants (VPPs). Frequently regarded as important elements of a net zero-carbon future, VPPs aggregate a network of distributed (often renewable) power sources, storage systems, and flexible energy consumers to optimize and dispatch generation or consumption in a smart energy grid. Flexibility on the demand side is particularly crucial during temporary periods of either high energy demand (i.e. peak load) or limited energy supply from variable renewable sources.

On the supply side, technologies such as IoT applications and smart meters can reveal crucial information about patterns and flexibility in households' and businesses' (appliance-specific) energy use. Suppliers and grid operators can use this information—and, given access, use direct load control[5]—to minimize inefficiency and waste while balancing electricity demand with available supply with minimal intervention into consumers' lifestyles. Furthermore, shaving peak load throughout the year can have a significant impact on suppliers' bottom line, providing vast economic motivation for utilities to provide consumers with incentives—to date, usually embedded in rate design—to

[1] Energy consumption from buildings comprises about a fifth of global energy consumption, with the proportion expected to increase as residents of non-OECD countries adopt home electronics and appliances at increasing rates (U.S. Energy Information Administration, 2019). In anticipation of growing energy demand from the residential sector and increased electrification of cooking, heating, and cooling as well as automobiles (Ürge-Vorsatz et al., 2015; Gielen et al., 2019), understanding and managing patterns of electricity use is becoming ever more crucial to meeting global climate change mitigation objectives. Additionally, due to varying carbon intensity of the grid throughout the day, the time of the day at which these growing electricity demands occur will impact their carbon footprint as well as grid reliability (Holland et al., 2016). Analysis of energy efficiency investments related to air conditioning in California demonstrate that the value of energy efficiency upgrades to society increase when time of use is taken into account (i.e. they are characterized by a "timing premium"), as peak savings occur during hours of high value to electricity markets (Boomhower and Davis, 2020).

[2] According to the Oxford Dictionary, IoT refers to "the interconnection via the Internet of computing devices embedded in everyday objects, enabling them to send and receive data."

[3] Our research is pertinent to the literature that provides evidence on "load shifting" behavior (Borenstein et al., 2002). Considering monetary incentives, recent experimental research supports the hypothesis that time-of-use (TOU) electricity pricing and price elasticity signals—that is, reminding consumers how much electricity costs when it is being used—can strongly affect the timing of electricity consumption (Wolak, 2011; Jessoe et al., 2014; Bradley et al., 2016)). The numerous experiments to evaluate consumer responses to dynamic electricity prices and other incentives have been comprehensively reviewed by Harding and Sexton (2017). Non-monetary incentives have also been tested in the context of load shifting[4]. Prest (2020) demonstrated that the effect of TOU electricity pricing, coupled with price signals, leads to a 10% reduction in peak electricity usage, and information provision, using in-home displays, augments these reductions to 15%. Typically, non-pecuniary messages significantly affect high users, while low users are more responsive to financial incentives (List et al., 2017).

[5] Interventions using direct load control, e.g., through smart thermostats, have shown to reduce peak demand by 15% (Ivanov et al., 2013)

reduce consumption during these hours.[6] Thus, smart devices simultaneously create opportunities for cost savings and reduction in negative environmental externalities.

Transition to a clean economy can be facilitated by the widespread adoption of such "prosocial technologies", i.e. technologies whose adoption creates benefits that extend beyond those internalized by the adopter. Prosocial technologies include, for example, residential solar photovoltaic panels or energy-efficient appliances, and they are increasingly being developed to target a variety of social aims, such as tracker apps for Covid-19 exposure, traffic warning devices, and vaccination trackers, among others.

When considering the potential of dynamic pricing in electricity, it has been recently assumed that enabling automation is of crucial importance given the tendency of households to be inattentive to electricity prices (Borenstein and Bushnell, 2019), particularly over time (Houde et al., 2013; Gilbert and Zivin, 2014).[7] This paper contributes to the limited literature that tests this conjecture in the field. IoT technologies—which create the potential to observe, disaggregate, and automate energy consumption—are becoming increasingly pervasive across the economy. In the UK, 17.3 million smart meters were in operation as of the end of Q1 2020, and billions of electronics and appliances are projected to be connected through IoT in the near future (Khajenasiri et al., 2017; Department for Business, Energy and Industrial Strategy, 2020). This rapid deployment of IoT technologies has brought to the fore a number of concerns regarding the performance and reliability of the evolving smart energy system, as issues related to misuse, safety, transparency, and data leakage shake consumers' trust (Alaa et al., 2017). As a consequence, the automation that is potentially necessary for effective implementation of demand response could be difficult to achieve.

We explore these issues by conducting a series of field experiments designed to explore the extent to which user behavior might constrain what is technically feasible—for example, users may have reservations when it comes to adopting "smart" or externally controlled devices—while also exploring how users might be incentivized to overcome such reservations.[8] Additionally, by observing users' tendencies to override our randomly timed smart plug switch-off events, we are able to gain insight into true energy demand flexibility; that is, we can understand which times of day or days of the week individuals may be willing to turn off appliances, and when they may not be willing to do so.

We ran our experiments via a program in university halls, which we call POWBAL (short for "power balancing"). To facilitate the program, we developed a web platform that collects data on how much electricity is consumed through the plugs and allows us to remotely and briefly switch off POWBAL-branded smart plugs, through which users could connect a time-flexible appliance or electronic device. The program incentivized students to participate in an exercise that would help to smooth energy consumption by offering them a free smart plug and the possibility of earning monetary rewards commensurate with plug usage. We also emphasized the prosocial element of participation by informing the user of the environmental benefits of energy demand flexibility.[9] Consenting participants each received a smart plug that allowed our research team (as well as the participants themselves) to control the power to the connected appliance remotely throughout the

[6]As detailed in Pratt and Erickson (2020), a utility's demand during the annual peak demand hour determines the capacity cost it will be required to pay in the following year to ensure the same amount is available to them during the annual peak hour (plus a reserve margin). This cost comprises a significant portion of a utility's capacity costs, which accounts for 25% of utilities' total wholesale market expense, a percentage that is trending upward. Hence, the ability to immediately shave peak load during particular hours of the year that are often difficult to forecast can lead to dramatic cost savings that, if passed to the customer, could provide additional incentive for consumers to accept direct load control mechanisms.

[7]Inattention severely impedes the benefits of dynamic pricing since active response to price changes is cognitively costly. Automation in dynamic pricing aims to enable lower-cost responses to price changes, which promise to achieve at least three times larger reductions for households that do rely on automation (Gillan and Gillan, 2017).

[8]In the context of climate change, any policy proposal to reduce emissions and energy demand will largely rely on an appropriate combination of behavioral, technological, and institutional capabilities. While technologies and institutions are often in the spotlight, influencing behavior is also crucial in tackling climate change (Stern, 2008; **?**). For instance, understanding the adoption behavior surrounding relevant technologies—such as home insulation or smart meters—is arguably as fundamental as developing the technology itself (Allcott and Mullainathan, 2010; Toft et al., 2014; Bugden and Stedman, 2019; Hmielowski et al., 2019).

[9]Several studies on energy flexibility and curtailment have demonstrated the importance of prosocial motivations (e.g., Asensio and Delmas, 2015; Pratt and Erickson, 2020).

study period. Participants then accrued points in proportion to the energy use we displaced, which in turn determined their likelihood of winning monetary prizes in our fortnightly lotteries.

To gain insight into the role of energy service interruption in determining plug usage, participants were randomly assigned to treatments characterized by differing duration and frequency of our randomly-timed “switch-off events”: (i) up to 15 or 30 minutes at a time, and (ii) either once or up to eight times a day (excluding the once per day for 15 minutes cell of the complete 2x2 design). The higher risk of interrupted energy services in the ‘higher intensity’ treatments is compensated through increased probability of winning versatile digital gift cards in our fortnightly lotteries. We assess this tradeoff through treatment differences in plug usage.

From the perspective of internal validity, the chosen context (i.e. university student accommodation) allows us to largely disregard complex irrelevant factors—in particular energy cost considerations—and focus on the question at hand. From an external validity perspective, we acknowledge that the university context may be limiting. On the other hand, we run the experiment at three different campuses comprising diverse student populations, allowing for assessment of generalizability to at least the UK student population. We centered our attention on a developed country context due to (i) these countries’ extensive established centralized energy grid systems that are largely sourced with emissions-intensive power generators (though may accommodate more renewable energy sources, provided flexibility), and (ii) the prevalence of energy-intensive home appliances and electronics amenable to smart plug usage.[10] Nevertheless, these features suggest immense potential of such programs in the developing world—perhaps particularly in rapidly industrializing economies—going forward, as per capita energy use seeks parity with the developed world and the need for rapid energy transitions becomes increasingly urgent, in particular due to air quality and climate concerns.

Our results suggest that users are more inclined to participate if switch-off events are sufficiently long to provide a meaningful probability of winning lottery prizes, but that higher frequency of such events may reduce participation intensity. However, individuals with relatively long but less frequent switch-off events are also more inclined to override the events. We provide evidence that individuals are motivated to consume more energy via the smart plug just after they have been randomly selected to receive a lottery reward.

The remainder of this paper is structured as follows. The next section provides the experimental setup and descriptive statistics. The third section outlines our stylized model and presents our results. Section 4 concludes.

2 Research Design

2.1 Experimental Setup

Recruitment Process

We conducted a series of field experiments focused on measuring individuals’ willingness to adopt POWBAL smart plugs and their subsequent energy consumption via these plugs under three randomly assigned treatments. To this end, we recruited undergraduate and graduate students from five LSE student halls, and from one student hall each at Imperial College London and University of Reading. We only invited individuals living in single-occupancy rooms in order to ensure that usage of the plugs derived from the sole individual exposed to treatment, and to maintain comparability in the characteristics of the residences (e.g., building, management, and events) and residents within a given student hall. The invitation offered participation in

[10] The energy demand displacement potential for large-scale systems of smart plugs in the residential sector is economically significant in developed countries. For example, should plugs be deployed on 30% of domestic refrigerators in the United Kingdom—each of which consumes between 150W to 400W on average—then a large coal-fired power station (2.4GW) could be entirely displaced during peak hours.

a research study aimed at designing systems to manage peak demand and smooth household energy consumption, where the information provided to prospective participants varied only in the frequency and duration with which they would experience switch-off events, as in Table 1 (see Appendix for recruitment materials).

Our sample of prospective participants comprised 1798 residents, where 262, 648, and 888 occupants from Imperial, LSE, and Reading, respectively, received an invitation to participate. We then recruited 231 experimental subjects (30, 93, 108 respectively), who consented to participate in the study and received plugs. Among these, 169 participants (24, 54, 91 respectively) actively used their plugs. We recruited participants via several channels, including email, phone app, flyers, in-person, or welcome packs in their rooms upon move-in. The majority of our participants were recruited in person, using a predetermined script to describe the study when meeting them in the residence.[11]

Students registered to participate in the study by completing a brief consent survey in the Qualtrics survey software. The survey first asked individuals to read a one-page overview of the study's purpose and implications of participation, then to check four boxes confirming that they had read the information sheet, understood their participation rights, and were aware that their plugs could be switched off for periods of time corresponding to their randomly assigned frequency and duration, which were stated explicitly and saliently in both the information sheet and the consent language (see Appendix A).

Our data collection began in April 2019 and lasted until the end of May 2020. Using student halls as the field of study comes with advantages and drawbacks. Students are not typical households and use electricity outside their room, for example in laundry or common dining rooms. They have a smaller range of appliances they can choose to use with their plugs. Even when aggregated for the UK, total electricity consumption of student rooms is a very small proportion of total domestic consumption.[12] However, on the plus side, the share of their overall electricity consumption that happens through the smart plug is likely to be higher than if the plugs were used in a regular household. We therefore capture a larger proportion of their potential behavioral response. In addition, our student participants do not pay for their energy usage, meaning we can test our incentive mechanisms in isolation (i.e. without any confounding of savings from reduced energy costs). The size of the rewards in relation to their total budget is also likely to be greater than in the rest of the population.

2.2 Procedure, Randomization, and Treatments

Our overall procedure is as follows: First, we used simple randomization[13] to randomly assign prospective participants into three groups differing in the frequency and duration of switch-off events in order to discern the likelihood of adoption based on switch-off intensity. We implemented three cells of a 2x2 study design. The first dimension is the number of possible switch-off events in a given 24-hour period. Some invitees were told that the switch-off events may occur once per day, while others were told they may occur once every three hours (i.e., up to eight times in a 24-hour period). The second dimension is the duration of the switch-off events, which may be either 15 or 30 minutes. To bolster group-level sample size and maximize observed switch-off events, we exclude the 15-minutes once per three hours cell of the 2x2 design, as shown in Table 1.

We distributed a (free) smart plug to all consenting participants that we (and they) used to remotely control electricity supply to their selected electronic device or appliance. Over the course of the

[11] Previous research has linked social pressure to technology adoption, adding another dimension to this channel; we do not attempt to measure treatment or welfare effects from our channel of adoption (see Giaccherini et al., 2019), but acknowledge that there may be an effect on adoption that, if anything, would attenuate any treatment effects in our study.

[12] Students in Imperial residences consume on average 18kWh per week. There are roughly 650,000 student rooms in the UK, so in total they would consume 11GWh per week, which under certain assumptions can be produced by 40 wind turbines.

[13] We do not have data on characteristics of the subject pool, hence stratified randomization was not possible.

Table 1: Treatment Group Design

	3 hours intervals	24-hour intervals
15 minutes switch-offs		Treatment 1
30 minutes switch-offs	Treatment 3	Treatment 2

NOTES: This table describes the three treatments with switch-off events varying in both duration and frequency over a 24-hour period. A 3 hours interval means a switch-off event can occur once every 3h, i.e. up eight times in a 24-hour period.

study period, we administered randomly timed switch-off events in accordance with the switch-off intensity to which each participant had been assigned. Participants collected points based on the energy that had been displaced through our switch-off events, as measured by the amount of energy consumed through the plug just prior to the switch-off event multiplied by the duration of the event.[14] Finally, in fortnightly intervals, we performed a lottery where participants with more points had a higher chance of winning monetary prizes.[15] Provided they accumulated at least one reward point during the two week period, the lottery allocated each participant with a probability to win a voucher, in proportion to their reward points.[16] Over the course of the experiment, we distributed £740 worth of versatile digital gift cards. Table 2 shows that among 31 total winners, a large majority - 23 - students won only once. 1 participant won 8 times.

Table 2: Lottery Results

	Vouchers won	Highest voucher won	N winners
One-time winners	£310	£30	23
2-times winners	£60	£10	3
3-times winners	£110	£50	1
4-times winners	£90	£20	2
5-times winners	£60	£20	1
8-times winners	£110	£20	1
Total	£740	£50	31

NOTES: This table reports the outcomes of the lottery, i.e how the vouchers were distributed among lottery winners. The variable "N winners" records how many many participants won repeatedly, the variable "Total won" records the share in £ that was won by x-times winners out of the entire voucher pool.

Additionally, we have another layer of randomization in that we randomly selected the time of day and week during which we administer these events across all participants.

2.3 Data

Our dataset consists of energy consumption readings for all POWBAL plugs that are consuming energy at 5-minute intervals (1,702,547 observations), the timing of the random switch-off events, and the dates and amounts of prize money received for lottery winners.

[14] We computed reward points as the power (in Watts) consumed by a plug just before the beginning of a switch-off event multiplied by the duration of the event (in hours). For example, if the power consumed by the plug prior to switch-off was 2 Watts, and the switch-off event lasted for 30 minutes (i.e. $\frac{1}{2}$ hour), the participant earned 2 Watts x $\frac{1}{2}$ hour = 1 reward point. If a participant overrode the switch-off event (i.e. by pressing the override button on the plug or via the app), we only counted the amount of time that the plug was switched off before the override (not the entire switch-off event).

[15] Our prizes came in the form of digital gift cards, redeemable in a variety of high-street and online retailers as well as hospitality providers.

[16] The lottery worked as follows: the number of sub-lotteries for a £10 voucher was defined by the number of participants with non-zero reward points during a given fortnight, for example 100 sub-lotteries if there had been 100 such participants. Each voucher was given away by drawing the names of participants. The probability to win each of these draws was equal to 15% of that participant's share of the overall number of reward points won in those two weeks.

Figure 1: POWBAL performance over time

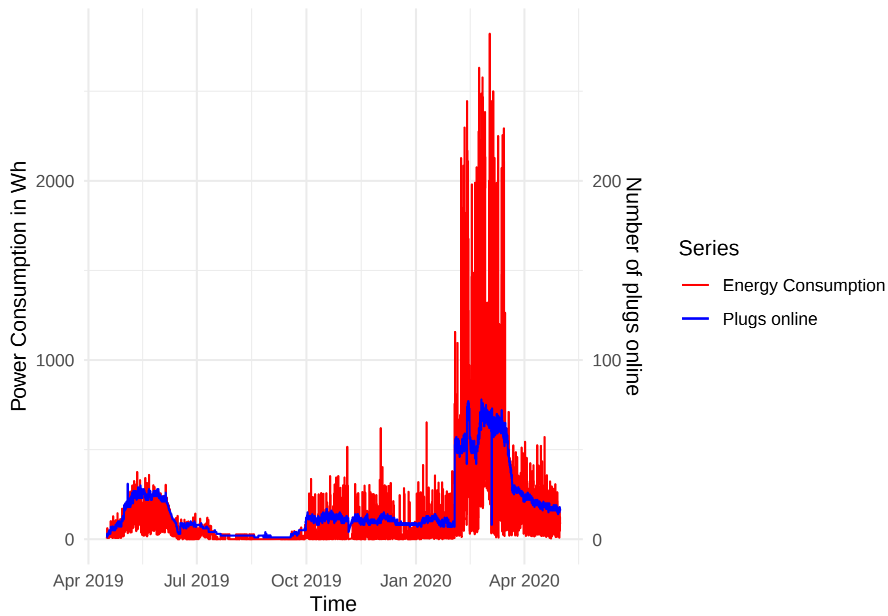

NOTES: This graph reports hourly figures of aggregate power consumption in Watt (red curve) and the number of plugs online (blue curve) over the period of the experiment. Data collection in LSE halls started in spring 2019 and ended during the summer, later trials at Imperial College and Reading University initiated in October 2019 and January 2020 respectively, and ended in June 2020 as the emergence of COVID-19 drew students home from March 2020.

We first provide an overview of our data on electricity consumption via the plugs. Out of the 231 who signed up and received their plug, 169 went on to use it at least once. Figure 1 reports hourly figures of aggregate power consumption. The LSE trial started in spring 2019 and ended over the summer as students successively moved out of the halls, as reflected in the left-hand part of Figure 1. The Imperial trial initiated in October 2019 when students moved in, and the Reading trial began in January 2020 at the start of term. The latter two trials officially ended on June 1, 2020, though the emergence of the COVID-19 lockdown in March 2020 led many students to move out of the halls to be with their families, which is apparent in the figure.[17]

Table 3: Descriptive statistics

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Power consumption	1,702,547	7.19	66.73	0	0	0	3,118
Non-zero power consumption	366,699	33.38	140.71	3.10	6.50	26.70	3,117.90

NOTES: This table reports descriptive statistics of power consumption over the duration of the experiment. Power consumption is recorded in Watts.

Table 3 reports statistics on the power consumption of connected devices (outside of switch-off events) across all measurements throughout the trials. The average power is 7 Watts (akin to that of a phone charger). However, this figure hides considerable heterogeneity both between and within users. Note that most plugs do not consume any electricity most of the time. Of the nearly 1.7 million data measurements we observe, only 22% report non-zero consumption, while in the remaining 78% of observations the plugs were connected to a socket without any devices consuming energy through them . However, there are some users who use up to 3 kW (about the power needed

[17] In our pre-registry, we had indicated that trials would take place in 2018, and simultaneously. Logistical technicalities delayed rollout at two of our sites, and the opportunity to work with University of Reading arose as we were preparing to implement the experiment at Imperial. Given lower than expected uptake in Imperial and LSE, we gladly partnered with Reading to achieve our desired sample size of 200-250, as designated in the pre-registry, which can be found on OSF.

Figure 2: Distribution of Power

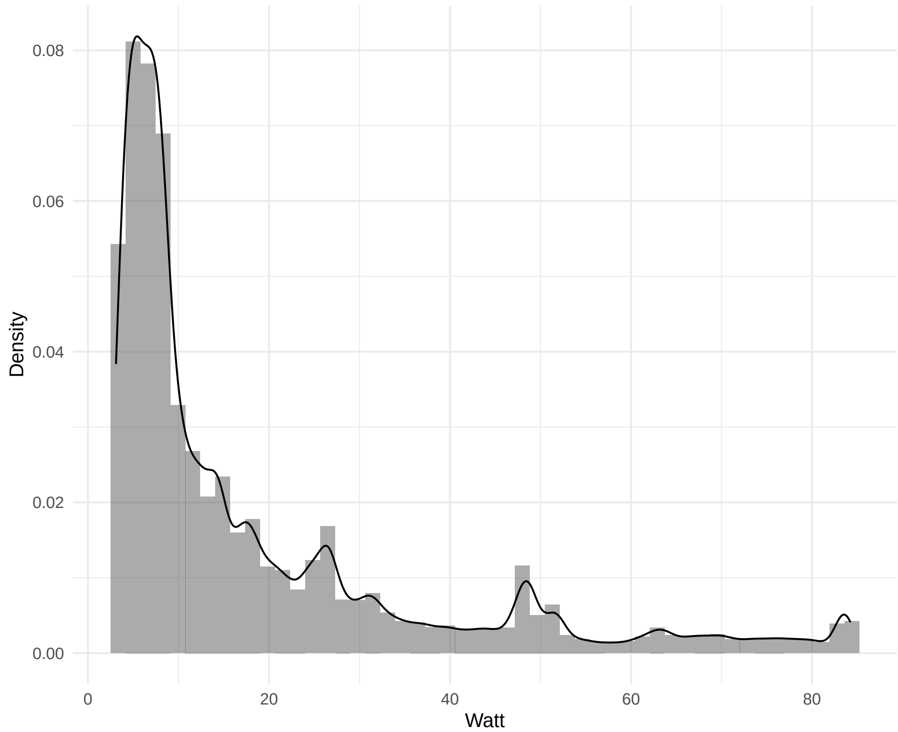

NOTES: This figure shows the fitted density plot of electricity consumption in Watts for non-zero consumption over the trial period, where the cutoff at the 99th percentile is 84.3 Watts.

to run a 3 ton 12 SEER air conditioner) some of the time. Note that individual plugs can consume energy even during a switch-off event if they choose to override the event.

Within-plug consumption changes over time because users connect different devices, switch off the outlet, unplug the POWBAL plug, or toggle between switching a device off or on, which they can do remotely or via a button on the plug itself. Therefore, it also makes sense to examine what power is being drawn when power consumption is non-zero. These data are reported in row 2 of Table 3, which reveals that average consumption is 33 Watts with the 75th percentile at 26.7 Watts, which is on the order of magnitude of a small laptop or a tablet charger, natural usages for a student hall. This coincides with our findings from the debrief survey.[18] Figure 2 reports the distribution of measured power for plugs with non-zero consumption and illustrates that a large proportion of participants used the plug for small amounts of consumption.

2.4 Consumption over the course of the day

Given the importance of optimizing and integrating VPPs in increasingly smart energy systems, demand fluctuations throughout the day need to be well understood. Figure 3 examines the average power usage via the plugs for different hours of the day. The figure shows that the timing of consumption varies substantially. During evening hours from 18:00 (6pm) onward, consumption is more than three times higher than during the lowest consumption period around 5:00 (5am). This pattern is comparable to within-day variations of consumption in typical UK households (Ushakova and Jankin Mikhaylov, 2020), with the notable difference that nighttime consumption is higher through the students' plugs and that the morning peak is not pronounced. In addition, the variation is independent of the day of the week. These different observations signify that the virtual power plant's capacity follows typical power demand to some extent, so that latent and

[18] Please see Appendix

Figure 3: Power consumption throughout the day

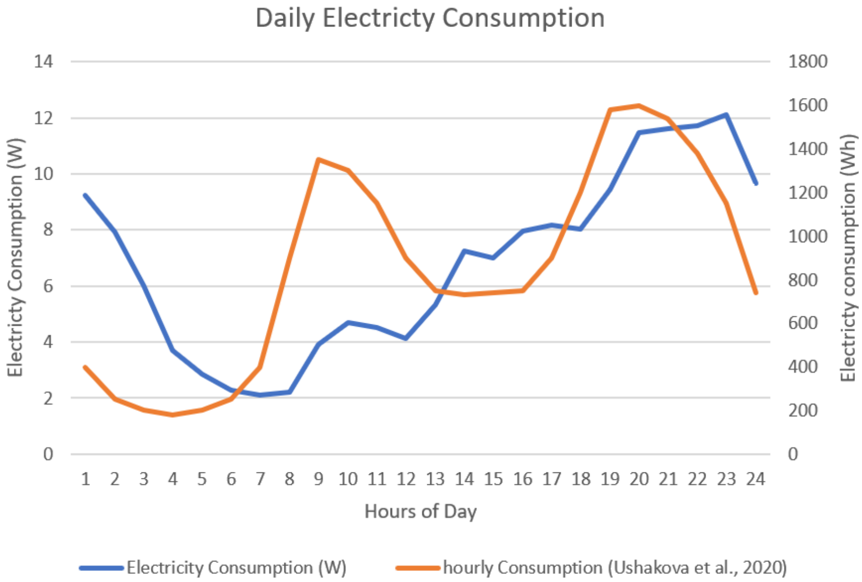

NOTES: This figure depicts average electricity consumed through plugs at each hour of the day, reported in Watts. The hours of the day are noted on the x-axis. The average hourly consumption data for the UK was adapted from ?.

untapped flexibility in the energy system is maximized during periods of peak load. Additionally, we see that the load remains relatively high until around midnight, which could be important if electric vehicle owners plug in their cars to charge at night, as incentivized in many existing electric vehicle energy plans.

3 Empirical Specification and Results

In this section we address our primary experimental research question of how our treatments, and the switch-off intensity they entailed, affected the likelihood of adoption and use. To do so, we compare adoption and use—i.e., whether the plug was used at all (binary) and the magnitude of energy consumption contributed toward the VPP via the plug following adoption (continuous), respectively—across the three groups. We additionally derive insights regarding prosocial motivation to adopt by analyzing the impact of receiving financial rewards on subsequent usage.

3.1 Prosocial Technology Adoption: A Stylized Model

To frame our analysis, we develop a stylized model on prosocial technology adoption. POWBAL differs from most technology adoption experiences—for instance, purchasing and using a smart phone—in that there is no intrinsic utility for the user. The value of adoption arises exclusively as a positive externality; e.g., in our case it facilitates increased efficiency of the electricity grid. However, there is potentially disutility from adoption—for instance, setup costs, data sharing, and relinquished control over select energy services in the home—so users may need to be compensated (Richter and Pollitt, 2018).

For simplicity, we assume that intensity of participation is unidimensional (e.g., frequency or duration of a switch-off event) and that there is a fixed disutility, or 'setup cost', from adoption. Moreover, with increasing switch-off intensity, the disutility of adoption increases at an accelerating rate, as in the red line of Figure 4. We also suppose that transfer payments (e.g., lottery prizes) increase linearly in consumption through the smart plug, as in the blue line of Figure 4. At point B, disutility exceeds benefits and the user will not adopt. At point A, benefits are higher than

Figure 4: A stylized model

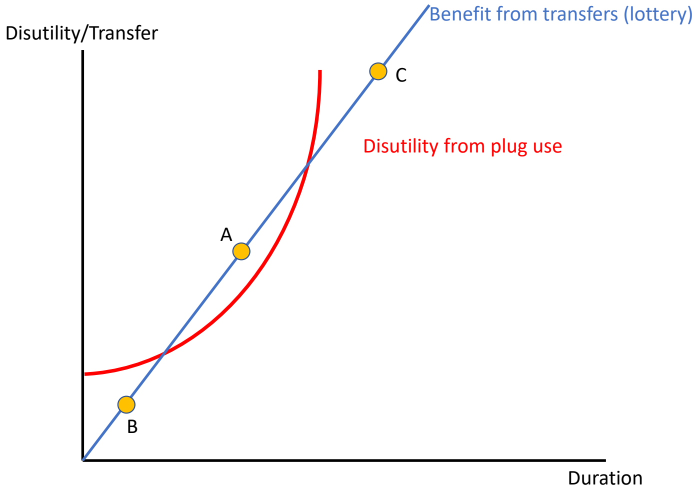

Notes: This figure models disutility from plug use (red curve) combined with benefits from transfers (blue curve) as function of variations in duration of switch-off events. At point A, benefits form the the lottery exceed the disutility from adopting the plug. Conversely, at points B and C, disutility is larger than the potential benefits.

disutility so that a user will adopt. In our current setting the user cannot choose switch-off intensity freely; rather, intensity is exogenously allocated. Hence, depending on the shape of the underlying functions in our experimental settings, individuals will either adopt and use the technology (as at point A) or choose not to adopt (as at points B or C). In the next two sections, we analyze first how the different treatments, and the disutility from plug use that they represent, affect adoption rates. We then explore whether the rewards from the lottery (extrinsic motivation) complement prosocial utility (intrinsic motivation).

3.2 Treatment impacts on adoption

3.2.1 Empirical Specification

Based on the three different treatments, we are able to explore two dimensions that could have an impact on subjects' decisions of whether to adopt the plugs: the duration and the frequency of the switch-off events. All else equal, we would assume that a user prefers less frequency (i.e. a long gap between switch-off events) and shorter events. However, less frequent and shorter switch-off events also translate to lower payoff likelihood, so the expected response to the different treatments is an empirical question.

To answer it, we run regressions of two different measures of adoption on the frequency and the duration of the event. The first measure is the average electricity consumed per week (in kWh) through individual plugs. The second measure is the percentage of weeks that a plug is online, regardless of how much electricity is consumed through it. With the reference group being users with exogenously assigned switch-off events of short duration (15 minutes) and only once every 24 hours, the regressions we run take the following form:

$$Y_i = \beta_0 + \beta_1 Freq + \beta_2 Long + H_i'\gamma + \epsilon_i \tag{1}$$

where Y_i is the measure of adoption, either electricity consumed or percentage of weeks the plug is online at least once. For example, if a user came online for the fist time in week 1, and then only every other week until the end of the trial, the latter outcome would be 50%. $Freq$ is a binary variable set to 1 when a user is subject to frequent switch-off events occurring up to every three hours (as in Treatment Group 3), and to zero otherwise. $Long$ is a variable set to 1 when a user is subject to long switch-off events of up to 30 minutes (as in Treatment Groups 2 and 3), and to zero otherwise. H_i is a vector of residence hall control dummies.

3.2.2 The effect of varying switch-off event frequency and duration

To begin, we graphically explore the effects of the switch-off event characteristics. In Figure 5 we examine the effects of the treatments by observing two measures of adoption—weekly energy consumption and share of users that have their POWBAL plug online at least once during a given week—by treatment. The figure reveals that energy consumption via the plug tends to be highest for Treatment Group 2 (low frequency, long duration of switch-off event) particularly in the early weeks of the treatment. Consumption appears lowest for participants assigned to the lowest switch-off event intensity, with the highest-intensity group in between.

In line with our discussion in section 3.1, one explanation may be that users do not view the disutility from switch-off events as worthwhile when prospective rewards are too low, so infrequent switch-off events lead to user disengagement. On the other hand, frequent switch-off events inconvenience the user, which is why engagement is highest with long duration and low frequency. Explicit regressions of consumption on treatment characteristics, as per equation 1, are presented in Table 4. They are supportive of this interpretation in that they suggest users with long duration switch-off events consume 3.4 kWh more over a week on average than those with short-duration events, whereas high frequency reduces consumption by 2.9 kWh on average, as shown in column 1. The direction of the results is robust to winsorizing at 1% the dependent variable in column 2, which shows the result is not driven by a few outliers; however, the magnitude of the effect is attenuated with winsorization. In column 3, the result holds once we control for the residence in which the user lives.

Some additional results support this interpretation. Note that the gap between Treatment Group 2 and the other groups becomes initially bigger as the trial proceeds. This feature of the data is in line with participants' learning about their rewards: users in group 2 respond to higher rewards whereas users in group 1 lose interest because of their lower chance of receiving rewards. We

Figure 5: Treatment effects on plug usage

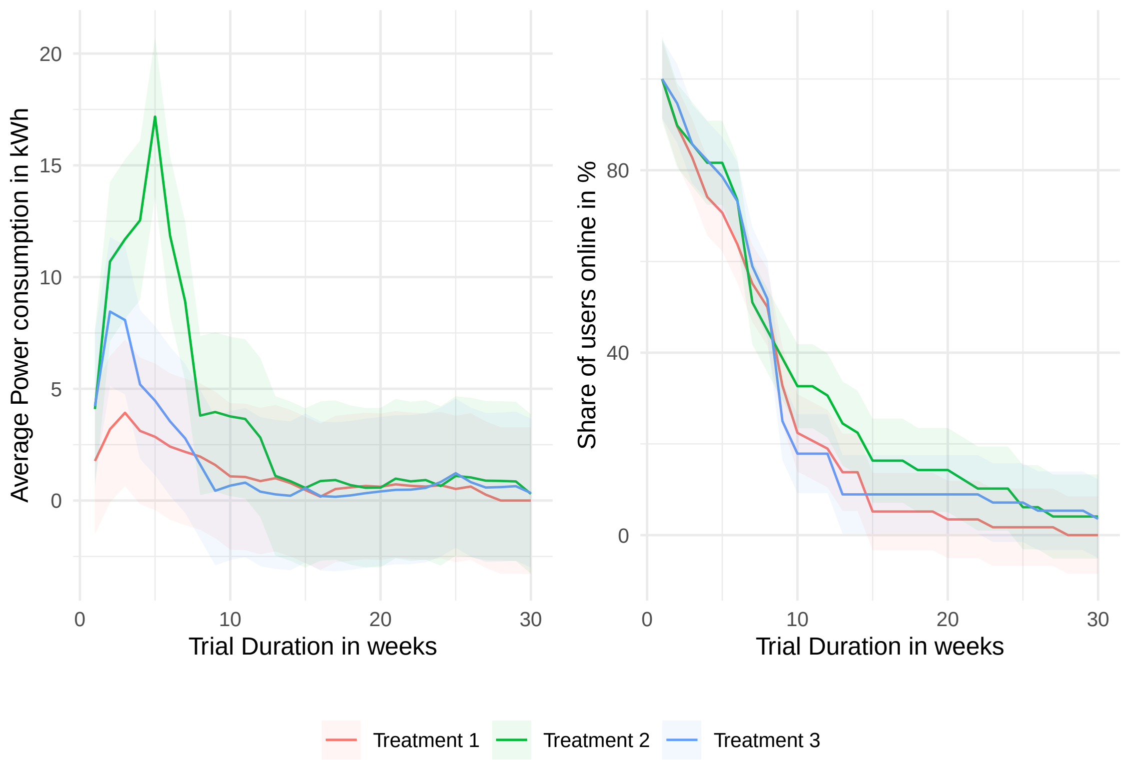

NOTES: The left-hand side graph reports the average energy consumption in kWh consumed through the plugs for each treatment group across all sites over the trial duration measured in weeks where time 0 is when the user first uses the plug. The right-hand side graph reports the share in percent of users that come online at least once, even briefly, during a given week for each treatment group across all sites over the trial duration in weeks.

also see that users in group 2 are, toward the end of the trial period, the most likely to remain engaged, leaving their plugs online for more weeks, as shown in the right panel of Figure 5. This effect is less pronounced than the consumption effect: in the last column of Table 4, we find that users with a longer switch-off duration treatment are 7.7 percentage points more likely to continue participating, though the effect is not statistically significant.

There are a number of caveats to this interpretation. First, Figure 5 illustrates powerfully that there is non-trivial attrition. Come week 15 of the trial most users abandon adoption. It is, however, noteworthy that the observed drop is conflated by the end of term (in the LSE case) and the onset of the COVID-19 lockdown (in the Reading case), as illustrated by Figures 6 to 8 which repeat Figure 5 for our three trial sites separately. Note that for the Imperial site where we had about 27 weeks of data before the lockdown began, attrition is much lower.

Interestingly, the relative performance of the treatment groups is different for the Imperial site compared to the others. There is no advantage for any treatment in the early weeks of the trial. Only toward the end does there appear to be some relative increase for group 3. Hence, we have to be cautious in drawing overly strong conclusions from the aggregated results. However, consistent with earlier results, group 1 comes lower in terms of participation likelihood toward the end of the trial. Hence, we can conclude with some confidence that the need to provide sufficient incentive to participate will dominate—within reason—factors such as the inconvenience of having plugs switch off too often or for too long.

Figure 6: Treatment effects on plug usage (LSE)

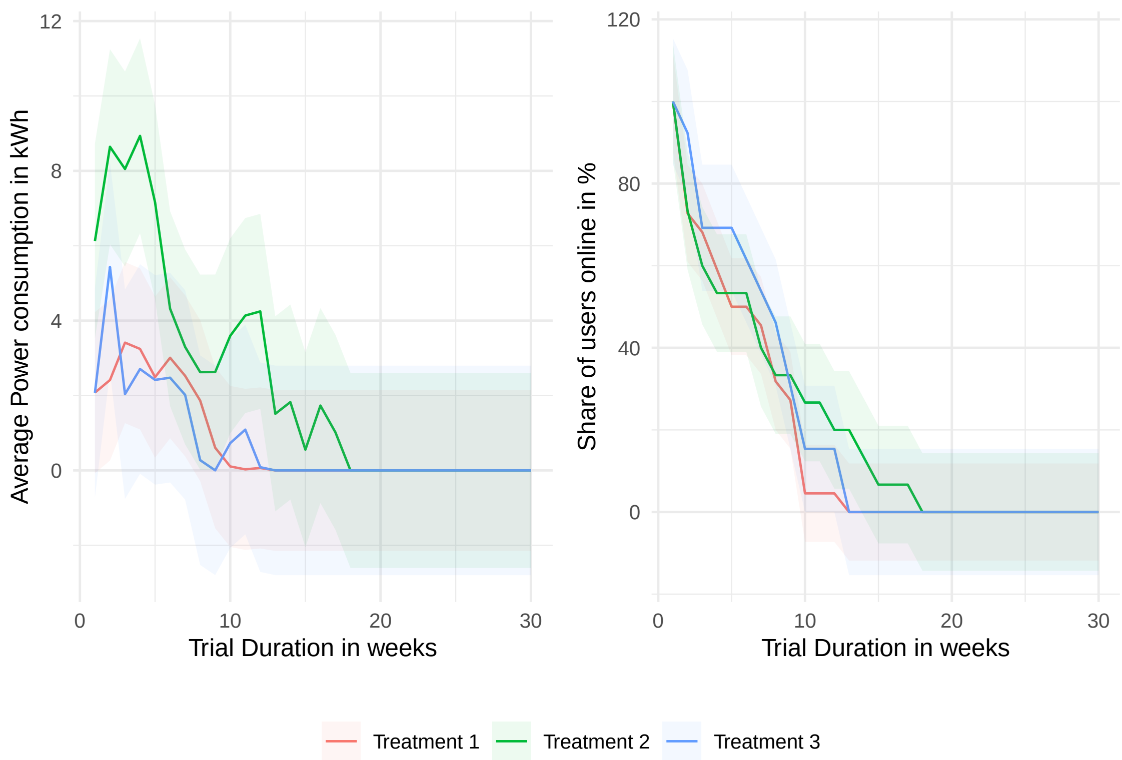

Notes: The left-hand side graph reports the average energy consumption in kWh consumed through the plugs for each treatment group across LSE sites over the trial duration measured in weeks. The right-hand side graph reports the share in % of users that come online at least once during a given week for each treatment group across LSE sites over the trial duration in weeks.

Figure 7: Treatment effects on plug usage (Imperial)

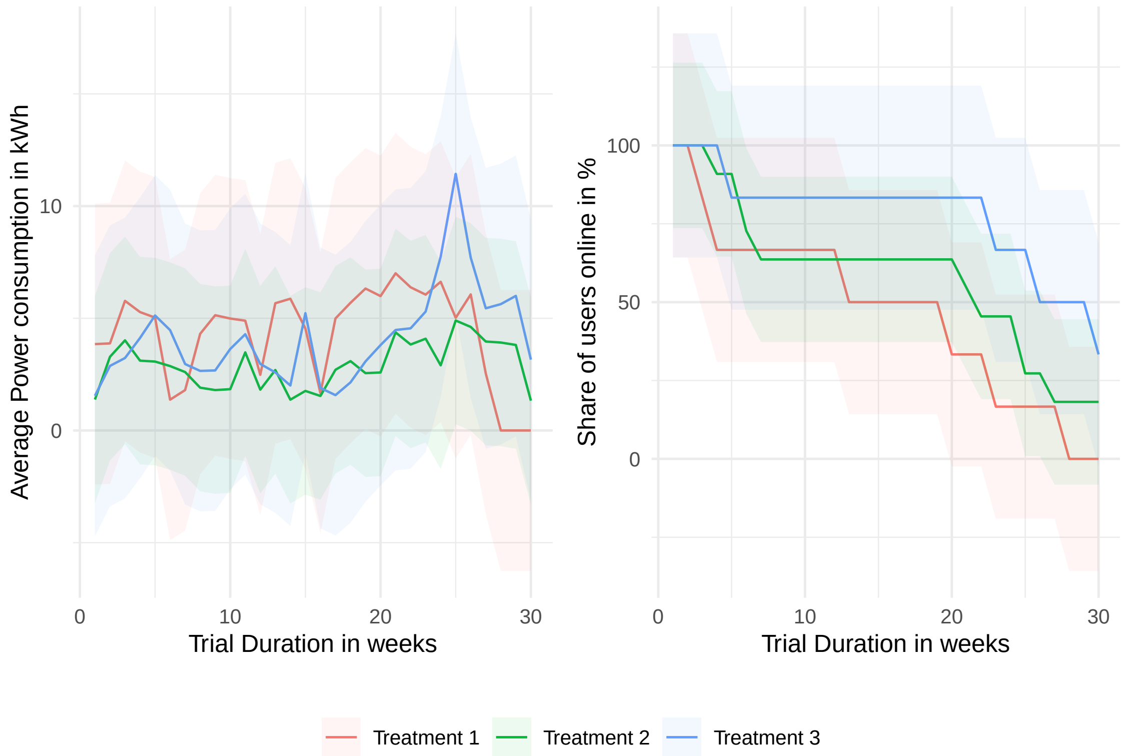

NOTES: The left-hand side graph reports the average energy consumption in kWh consumed through the plugs for each treatment group across Imperial sites over the trial duration measured in weeks. The right-hand side graph reports the share in % of users that come online at least once briefly during a given week for each treatment group across Imperial sites over the trial duration in weeks.

Figure 8: Treatment effects on plug usage (Reading)

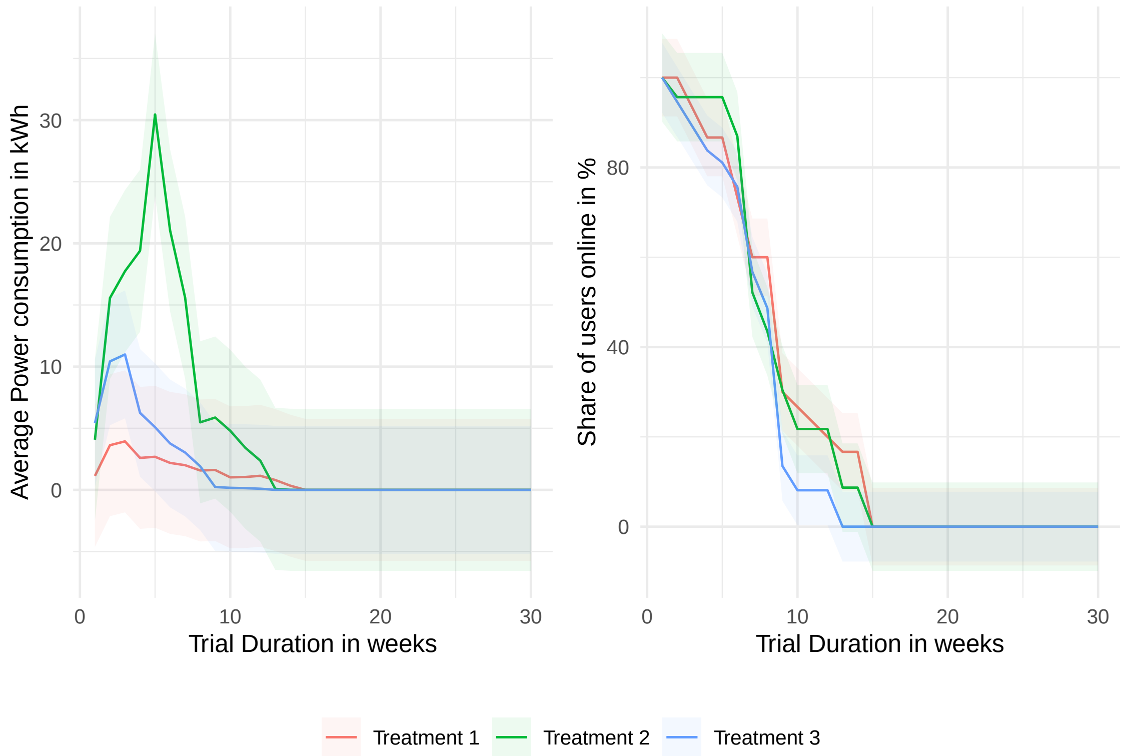

NOTES: The left-hand side graph reports the average energy consumption in kWh consumed through the plugs for each treatment group at the Reading site over the trial duration measured in weeks. The right-hand side graph reports the share in % of users that come online at least once during a given week for each treatment group at the Reading site over the trial duration in weeks.

Table 4: Regressions on Treatment Characteristics

	Dependent variable: Weekly			
	Av. kWh	Av. kWh Wins.	Av. kWh	User online
	(1)	(2)	(3)	(4)
Long duration	3.386^{**}	1.686^{**}	3.370^{**}	7.700
	(1.486)	(0.824)	(1.508)	(5.551)
High Frequency	-2.887^{*}	-1.693^{**}	-3.126^{**}	-7.547
	(1.492)	(0.827)	(1.518)	(5.572)
Reference Group	1.620	1.415^{**}	1.295	39.000^{***}
	(1.002)	(0.556)	(1.862)	(3.742)
Site Controls	No	No	Yes	No
Observations	169	169	169	169
R^2	0.034	0.032	0.045	0.014
Adjusted R^2	0.023	0.020	0.021	0.003
Residual Std. Error	7.762	4.304	7.767	28.987
F Statistic	2.940^{*}	2.713^{*}	1.916	1.218

NOTES: These estimates are the result of OLS regressions. The table reports the effect of varying the duration and frequency of switch-offs events on the weekly average electricity consumption in kWh. Each coefficient represents the difference relative to the control group such as to estimate the treatment effects of long duration and high frequency switch-off events. The dependent variable is average weekly electricity consumption in columns 1 and 3. In column 2, the average weekly electricity consumption is winsorized at 1%. In column 4, the dependent variable is the percentage of weeks that a given user was online at least once after their first connection. In column 3, a set of dummies controls for which student residence the participant resides in. Robust standard errors in parentheses. Significance levels are indicated as *p<0.1; **p<0.05; ***p<0.01

3.3 How prosocial are participants?

In addition to exploring the impact of switch-off characteristics on adoption, we seek to understand the motivation of participants. There could be at least two benefits driving users to adopt, despite disutility from interruption to energy services: first, there could be a purely prosocial motivation, since the POWBAL technology promises to enable clean energy generation. Alternatively, users could be driven by the desire to win lottery prizes.

We can examine the relative importance of these motivations by comparing users' participation intensity—i.e. their energy consumption via the plug—to their reward intensity. We hypothesize that if users are driven by prosocial motives alone, they should not respond to lottery wins by increasing their energy consumption through the plug to increase their chances of future wins.

On the other hand, if they are driven by private rewards *and* if they update their expectations about those rewards according to observed lottery wins, we would expect to find a relationship between receiving rewards and participation intensity. In other words, if we find that consumers respond to rewards being paid out, this is a sufficient condition for the presence of reward effects (although not a necessary one).

We develop a simple model to illustrate this phenomenon. Suppose users' utility is defined as

$$U = (\theta r_i + s_i)\, e_i - e_i^{\eta}$$

where e_i is participation intensity, r_i represents the financial reward, θ is the importance of rewards to users and s_i the prosocial benefit. We assume that $\eta > 1$, which implies increasing marginal disutility from higher participation intensity. We here abstract from fixed costs that were discussed in the earlier model.

Hence optimal participation is defined by

$$U' = 0 = \theta r_i + s_i - \eta e_i^{\eta - 1}$$

so that optimal participation intensity becomes

$$e_i^* = \left(\frac{\theta r_i + s_i}{\eta} \right)^{\frac{1}{\eta - 1}}$$

If there is uncertainty over rewards, consumers will maximize expected utility $E\{U\}$ which implies that optimal adoption becomes

$$e_{it}^* = \left(\frac{\theta E_{it}\{r_i\} + s_i}{\eta} \right)^{\frac{1}{\eta - 1}}$$

with $E_{it}\{r_i\}$ being users' expectation about rewards at time t. If we find that actual reward payouts in the past (e.g. r_{it-1}) drives current participation e_{it}, we can infer that users form their expectations about payouts on the basis of what they have actually received in the past:

$$E_{it}\{r_i\} = f(r_{it-1})$$

Hence, past payouts become an instrument for the effect of rewards on participation (except that we don't directly observe $E_{it}\{r_i\}$). We can then write the reduced form effect in this setting as

$$\frac{\partial e_{it}}{\partial r_{it-1}} = \frac{1}{\eta - 1}\left(\frac{\theta E_{it}\{r_i\} + s_i}{\eta}\right)^{\frac{2-\eta}{\eta-1}} \frac{\theta}{\eta} \times \frac{\partial f(r_{it-1})}{\partial r_{it-1}}$$

Consequently, a regression of e_{it} on r_{it-1} will provide us with the marginal response to rewards, times the response of expectations to rewards. For an entirely prosocial driven user (with $\theta = 0$) this would be zero. For - at least in part - not prosocially driven users it seems reasonable to assume that

$$0 < \frac{\partial f(r_{it-1})}{\partial r_{it-1}} < 1$$

i.e. in response to a very high payment, users would adjust their beliefs upward but not excessively so. Consequently, a reduced form regression of the form

$$e_{it} = \beta r_{it-1} + \epsilon_{it}$$

would give us a lower bound of the response of users to rewards:

$$\frac{1}{\eta - 1}\left(\frac{\theta E_{it}\{r_i\} + s_i}{\eta}\right)^{\frac{2-\eta}{\eta-1}} \frac{\theta}{\eta} > \beta$$

This in turn provides us with an upper bound of the prosocialness of our virtual power plant.

Figure 9 illustrates this by plotting the change in weekly energy consumption against the change in reward payments. We find large variation in both variables. The figure also reveals a positive relationship between the two suggesting that more reward payments will encourage participants to increase the amount of electricity they consume through the plugs.

We examine this response in greater detail by running regressions of the form 3.3 where we also assume that there is a user-specific fixed effect:

$$\epsilon_{it} = \alpha_i + \nu_{it}$$

We deal with unobservable user-specific variation by running the regression in first differences. Table 5 presents the results. Column 1 suggests that a reward payout in the previous week will increase consumption in the current week by 45Wh per £1 paid out. Column 2 allows for further lags, which suggests that there is an additional boost in adoption with an increase of 38Wh two weeks after a payout. However, we also see that this additional boost disappears by the third week after receiving payout. Hence, our first column estimate seems to be a good representation of the net effect.

One concern could be that autocorrelation in a users' consumption could be driving our effects; i.e. if a user consumed more last week she might also consume more (or less) this week. Clearly last week's consumption is driving last week's reward payments. It is easy to account for this autocorrelation however, by including the previous week's consumption levels as further controls (and using a dynamic panel estimator, i.e. Arrelano Bond, for the estimation), which we present in column 3. The results show that autocorrelation has little effect on the qualitative conclusions drawn from columns 1 and 2. Indeed, it leads to a slightly larger impact estimate of rewards.

Based on these findings, what can we conclude about the importance of the financial motive in comparison to the prosocial motive in fueling our virtual power plant? As noted above, our estimates provide a lower bound on the financial motive (and in turn an upper bound on the prosocial motive).

Figure 9: Do rewards improve participation?

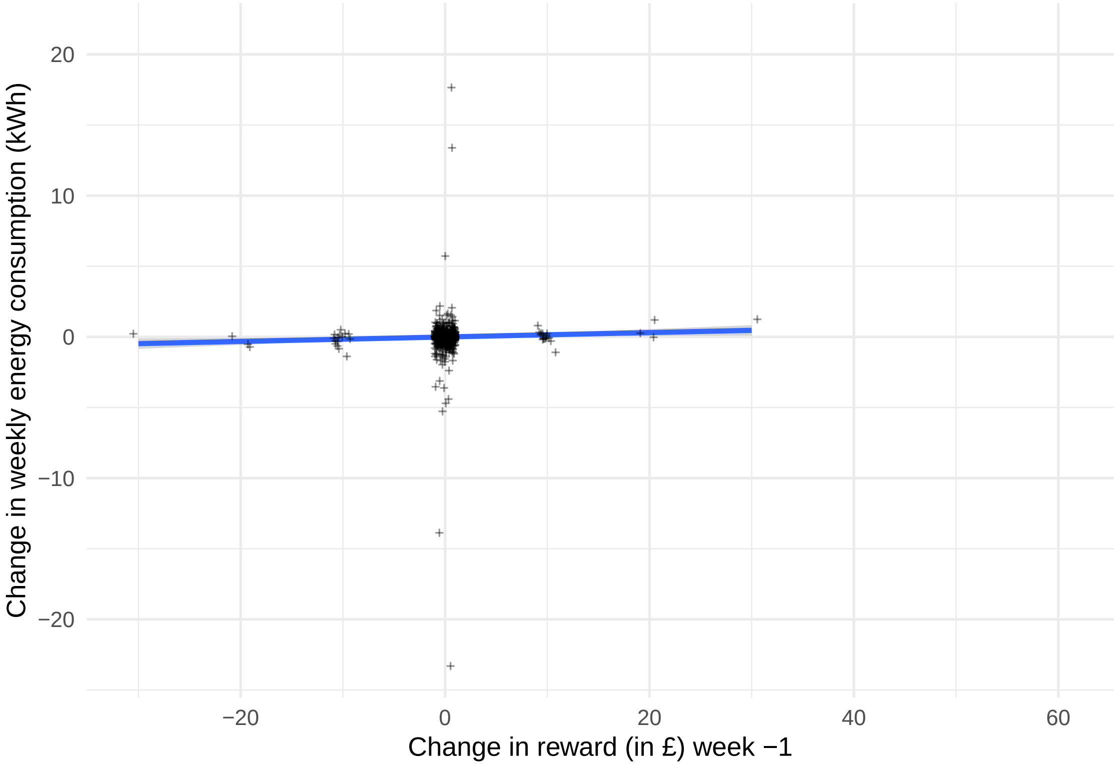

NOTES: This figure plots a fitted line between points that each represent a participant on a given week. On the y-axis, the variable represents the change in weekly energy consumption from week t-1 to week t. On the x-axis, the variable measures the change in reward payments from week t-2 to week t-1. (For example, if a participant wins a 20£ voucher in week t-2 and does not win in week t-1 then the change from week t-2 to t-1 is equal to -20)

Table 5: Regressions of energy consumption on reward payouts

	Dependent variable:		
	Change in weekly energy consumption		
	(1)	(2)	(3)
Reward Week -1	45.361^{**}	52.353^{**}	68.754^{**}
	(19.714)	(22.358)	(30.593)
Reward Week -2		38.224^{*}	52.764
		(20.582)	(33.581)
Reward Week -3		-33.831^{*}	-18.257
		(17.506)	(27.298)
Wh Week -1			-0.074
			(0.348)
Wh Week -2			0.079
			(0.099)
Wh Week -3			0.148
			(0.234)
Week controls	Yes	Yes	Yes
Observations	4742	4742	4742
Users	166	166	166

NOTES: These estimates are the result of OLS regressions. The dependent variable is the change in weekly electricity consumption in Wh from week t-1 to week t. The explanatory variables take the value 1 if the user received a reward in week t-1, t-2 or t-3, and zero otherwise. The table reports the effect of receiving reward in previous weeks on the change in electricity consumption. Column 3 also includes as controls weekly electricity consumption in weeks t-1, t-2 and t-3. Robust standard errors in parentheses. Significance levels are indicated as *p<0.1; **p<0.05; ***p<0.01

Let's consider the estimate in column 1 of Table 5, which suggests that the response to winning £1 is an additional 45 Wh of consumption through the plug. Average weekly consumption is 136 Wh, and average payouts are 75 pence. Hence, not paying the average user would reduce her consumption by $45Wh \times 0.75 = 33Wh$. This amount corresponds to 24% of average consumption, which would seem substantial. Of course, if this amount is worth paying for the corresponding contribution to the VPP, rewards need to be determined in relation to marginal increases in either revenue or carbon reductions.

3.4 Overriding

Our discussion thus far has focused on adoption and amount of energy consumption as the primary outcome variables of our experiment. However, there is a further dimension of how users can respond to the exogenous switch-off schedule, which is by overriding the switch-offs. Table 6 reports regressions where the dependent variable is a binary variable equal to 1 if a switch-off event is overridden by the user, and the independent variable of interest is the treatment to which the user was assigned. In total we observe all 6719 switch-off events that occurred over our trial period.

Overall, treatment assignment (i.e. switch-off intensity) appears to have little impact on overriding behavior, with a rate between 2.1% for treatment group 1 (Reference Group) and 3.4+2.1=5.5% for treatment group 2. However, this is in part because many of the switch-off events occur during periods when users do not consume anything via the plugs. Hence, in column 2 we look at overrides for switch-off events with non-zero load (in Watt) in the five-minute period immediately preceding a switch-off. For this group switch-off events are considerably higher at between 7 and 17%. Comparing the variation of overriding between treatment groups leads to a surprising but robust result: overriding is most prevalent for treatment group 2 with relatively long but infrequent switch-off events.

On the other hand, the amount of overriding in groups 1 and 3 is nearly identical. One might have expected most overriding to occur in group 3, wherein participants experience both long and frequent switch-off events. At this point we can only speculate as to why this is not the case. Perhaps users connect different devices under the various treatments. It could be that the with longer treatments they make certain ex ante that any connected devices are more flexible and therefore suited for better suited for withstanding switch-off events.

Columns 3 and 4 explore whether there is a relationship between pre-switch-off energy consumption and overriding. It would be concerning if there is a positive relationship as this would imply that the overriding behavior could substantially reduce the capacity of the virtual power plant. Column 3 includes the whole sample whereas in column 4 we only include the switch-off events with positive pre-switch-off power usage (as in column 2). While we see a strong association between power usage and overriding in column 3, the association vanishes in column 4. Hence, the effect in column 3 is entirely driven by the extensive margin of power users with positive usage.

4 Discussion and Conclusion

The development and adoption of new and often "prosocial" technologies will be critical to address climate change as well as other prominent societal issues, including economic growth and public health. Smart connected IoT devices powered by artificial intelligence have become a central focus of technological advancement in recent years. In this paper, we introduce POWBAL, a new and original platform for research into the possible impacts of inexpensive smart technologies in facilitating low-carbon energy transitions while increasing electricity reliability and efficiency.

The emergence of IoT applications will continue to increase data availability across a number of sectors. This exponential increase in data collection opportunities will create an environment where

Table 6: Regressions of Override events

	Dependent variable:			
	Override event			
	(1)	(2)	(3)	(4)
Treatment 1	0.021	0.070***	0.021***	0.070***
	(0.016)	(0.005)	(0.006)	(0.020)
Treatment 2	0.055**	0.169***	0.053***	0.167***
	(0.026)	(0.014)	(0.006)	(0.021)
Treatment 3	0.032**	0.085***	0.032***	0.084***
	(0.015)	(0.009)	(0.003)	(0.009)
Pre switch-off power in W			0.0001***	0.0001
			(0.00003)	(0.0001)
Observations	6,719	1,467	6,719	1,467
R^2	0.036	0.103	0.039	0.104
Adjusted R^2	0.036	0.101	0.038	0.101
Residual Std. Error	0.180	0.291	0.180	0.291
F Statistic	83.829***	56.194***	67.927***	42.357***

NOTES: These estimates are the result of OLS regressions. The table reports the propensity of participants to override the switch-off event across treatment groups. The dependent variable is a binary variable equal to 1 when a switch-off event is overridden while each coefficient represents the increased likelihood in % of overrides happening in each of the three treatment groups. In column 1 and 3, regressions are run on the entire data-set while columns 2 and 4 control for periods of non-zero consumption shortly preceding switch-off events. The variable pre-switch-off power in columns 3 and 4 reports the effect of consuming power before a switch-off event on the likelihood of overriding the upcoming switch-off event. Robust standard errors in parentheses. Significance levels are indicated as *p<0.1; **p<0.05; ***p<0.01

harvesting information about the activities and behaviors of individuals, devices, and physical phenomena becomes increasingly commonplace. Understanding which information and incentives are useful in this setting, which IoT applications are likely to be adopted and influence energy consumption behavior, and how users respond to real-time feedback mechanisms is critical.

Via internet-connected power plug adapters, POWBAL allows remote device-level power consumption monitoring as well as remote power switch-offs. We use this capability to conduct a series of field experiments on dynamic demand management. If deployed at scale, such a technology could significantly reduce the requisite capacity to ensure reliable electricity supply while reducing the costs of intermittent renewable technologies by requiring fewer backup generation assets. In other words, the technology allows us to better conceptualize and actualize a consumer-driven virtual power plant (VPP) with instantaneous ramping capabilities. Conventional VPPs consist of the bundling and coordination of a large number of small generation assets (e.g., from solar panels). In our case, these generation assets are small electricity gateway devices that can swiftly be switched off, thereby delivery "negawatts" to the grid.

Our results provide a first step in understanding mechanism design for virtual power plants fueled by smart plugs, an inexpensive IoT technology that can make any appliance or electronic "smart". While our focus on student halls in the United Kingdom may limit the external validity of our results, a number of interesting findings emerge. These include, firstly, that participation is highly skewed with many users connecting very small or no power loads—as might be expected in university housing—despite adopting the plug in principle. However, despite the limited appliances in these residential settings, there are also some superusers connecting loads that could make a substantial contribution to load balancing.

Users appear to respond to different treatment options (with different degrees of switch-off intensity) in a nonlinear way. That is, more intrusion does not necessarily lead to less participation, which may be rationalized via the trade-off between rewards and service interruption. Relatedly, we find that users who experience a more intensive switch-off regimen are not necessarily overriding the load-balancing efforts more frequently, which could indicate that these users are more considerate of the flexibility of devices they connect through their smart plugs. In other words, as with issues related to technology adoption, there are often complementary—and often intangible—investments required of adopters, and the quality and quantity of those investments matter. These complementary investments could be a further driver of non-linear responses to treatment intrusiveness.

Finally, we find that users clearly respond to reward payouts, suggesting that adoption is not purely driven by prosocial motives. Evidence of such varied motivations is good news for efforts to facilitate widespread adoption of such a technology, as prosocially motivated adopters tend to make up only a small fraction of any population. Future research should further rigorously explore the vast array of competing motivations that determine adoption of prosocial technologies that can contribute toward such a virtual power plant, and toward climate change mitigation more broadly.

References

Abrahamse, W., L. Steg, C. Vlek, and T. Rothengatter (2005): "A review of intervention studies aimed at household energy conservation," *Journal of Environmental Psychology*, 25, 273–291.

Alaa, M., A. A. Zaidan, B. B. Zaidan, M. Talal, and M. L. Kiah (2017): "A review of smart home applications based on Internet of Things," .

Alberts, G., Z. Gurguc, P. Koutroumpis, R. Martin, M. Muûls, and T. Napp (2016): "Competition and norms: A self-defeating combination?" *Energy Policy*, 96, 504–523.

Allcott, H. and S. Mullainathan (2010): "Behavior and energy policy," *Science*, 327, 1204–1205.

Asensio, O. I. and M. A. Delmas (2015): "Nonprice incentives and energy conservation," *Proceedings of the National Academy of Sciences*, 112, E510–E515.

Boomhower, J. and L. Davis (2020): "Do Energy Efficiency Investments Deliver at the Right Time?" *American Economic Journal: Applied Economics*, 12, 115–39.

Borenstein, S. and J. Bushnell (2019): "Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency," Working Paper 294R, Energy Institute.

Borenstein, S., M. Jaske, and A. Rosenfeld (2002): "Dynamic pricing, advanced metering, and demand response in electricity markets," Working Paper 105, Centre for the Study of Energy Markets.

Bradley, P., A. Coke, and M. Leach (2016): "Financial incentive approaches for reducing peak electricity demand, experience from pilot trials with a UK energy provider," *Energy Policy*, 98, 108–120.

Buckley, P. (2020): "Prices, information and nudges for residential electricity conservation: A meta-analysis," *Ecological Economics*, 172.

Bugden, D. and R. Stedman (2019): "A synthetic view of acceptance and engagement with smart meters in the United States," *Energy Research and Social Science*, 47, 137–145.

Carbon Trust and Imperial College London (2016): *An analysis of electricity system flexibility for Great Britain.*

Department for Business, Energy and Industrial Strategy (2020): *Smart Meter Statistics in Great Britain: Quarterly Report to end March 2020.*

Giaccherini, M., D. H. Herberich, D. Jimenez-Gomez, J. A. List, G. Ponti, and M. K. Price (2019): "The Behavioralist Goes Door-To-Door: Understanding Household Technological Diffusion Using a Theory-Driven Natural Field Experiment," Working Paper 26173, National Bureau of Economic Research.

Gielen, D., F. Boshell, D. Saygin, M. D. Bazilian, N. Wagner, and R. Gorini (2019): "The role of renewable energy in the global energy transformation," *Energy Strategy Reviews*, 24, 38–50.

Gilbert, B. and J. G. Zivin (2014): "Dynamic salience with intermittent billing: Evidence from smart electricity meters," *Journal of Economic Behavior & Organization*, 107, 176–190.

Gillan, J. and J. M. Gillan (2017): " Dynamic Pricing, Attention, and Automation: Evidence from a Field Experiment in Electricity Consumption," Working Paper 284, Energy Institute.

Harding, M. and S. Sexton (2017): "Household Response to Time-Varying Electricity Prices," .

Hmielowski, J. D., A. D. Boyd, G. Harvey, and J. Joo (2019): "The social dimensions of smart meters in the United States: Demographics, privacy, and technology readiness," *Energy Research and Social Science*, 55, 189–197.

Holland, S. P., E. T. Mansur, N. Z. Muller, and A. J. Yates (2016): "Are there environmental benefits from driving electric vehicles? The importance of local factors," *American Economic Review*, 106, 3700–3729.

Houde, S., A. Todd, A. Sudarshan, J. A. Flora, and K. C. Armel (2013): "Real-time feedback and electricity consumption: A field experiment assessing the potential for savings and persistence," *The Energy Journal*, 34.

Ivanov, C., L. Getachew, S. A. Fenrick, and B. Vittetoe (2013): "Enabling technologies and energy savings: the case of energywise smart meter pilot of connexus energy," *Utilities Policy*, 26, 76–84.

Jessoe, K., D. Rapson, and J. B. Smith (2014): "Towards understanding the role of price in residential electricity choices: Evidence from a natural experiment," *Journal of Economic Behavior and Organization*, 107, 191–208.

Khajenasiri, I., A. Estebsari, M. Verhelst, and G. Gielen (2017): "A Review on Internet of Things Solutions for Intelligent Energy Control in Buildings for Smart City Applications," in *Energy Procedia*, Elsevier Ltd, vol. 111, 770–779.

List, J. A., R. D. Metcalfe, M. K. Price, and F. Rundhammer (2017): "Harnessing Policy Complementarities to Conserve Energy: Evidence from a Natural Field Experiment," Working Paper 23355, National Bureau of Economic Research.

Pratt, B. W. and J. D. Erickson (2020): "Defeat the Peak: Behavioral insights for electricity demand response program design," *Energy Research & Social Science*, 61, 101352.

Prest, B. C. (2020): "Peaking interest: How awareness drives the effectiveness of time-of-use electricity pricing," *Journal of the Association of Environmental and Resource Economists*, 7, 103–143.

Richter, L.-L. and M. G. Pollitt (2018): "Which smart electricity service contracts will consumers accept? The demand for compensation in a platform market," *Energy Economics*, 72, 436–450.

Stern, N. (2008): "The Economics of Climate Change, The Stern Review," *American Economic Review: Papers & Proceedings*, 98, 37.

Toft, M. B., G. Schuitema, and J. Thøgersen (2014): "Responsible technology acceptance: Model development and application to consumer acceptance of Smart Grid technology," *Applied Energy*, 134, 392–400.

Ürge-Vorsatz, D., L. F. Cabeza, S. Serrano, C. Barreneche, and K. Petrichenko (2015): "Heating and cooling energy trends and drivers in buildings," *Renewable and Sustainable Energy Reviews*, 41, 85–98.

U.S. Energy Information Administration (2019): *International Energy Outlook 2019 with projections to 2050.*

Ushakova, A. and S. Jankin Mikhaylov (2020): "Big data to the rescue? Challenges in analysing granular household electricity consumption in the United Kingdom," *Energy Research and Social Science*, 64.

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

A Appendix

Figure 10: Participant Information Sheet

Imperial College London

LSE THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

POWBAL.net

Participant Information Sheet

We are writing on behalf of POWBAL, a research project on energy consumption by Imperial College Business School and the London School of Economics and Political Science. We have organised a platform through which you can win versatile and easy-to-use e-vouchers* through your energy use.

How?

You can sign up to the study by reading this information sheet and agreeing to the consent form (by ticking all of the boxes in the Qualtrics sign-up survey). As a participant in the study, you will receive a smart plug, free of charge, which can be controlled and monitored remotely via the Internet.

You can attach this smart plug to your home electronics or appliances. The more energy you consume through the plug, the higher the likelihood that you will win e-vouchers. A handy Frequently Asked Questions document as well as further details on the project will be emailed to you after you sign up.

What is POWBAL all about?

In short, the plugs can be used to help balance the UK's electricity grid, allowing for cleaner electricity generation. By participating, you will allow us to switch off the plug for **30 minutes no more than once in any given 24-hour period.** Hence you should only connect certain devices for which brief switch-offs are unproblematic. For instance, you could use it to power devices with battery backup or that you regularly leave on standby (note that plugged-in electronics consume energy even when not in use). That said, you **always** remain in control. You will be able to easily override any switch-off if it happens at an inconvenient time via a button on the plug or on your web profile.

Why should you take part?

If we all use smart plugs in this way, it becomes much easier to balance the electricity grid by matching electricity supply and demand more efficiently. Such balancing is important if we want to have cleaner power from renewable sources such as wind and solar in the UK's energy mix. Hence, by taking part you can help to make the power system cleaner. And, of course, you can win rewards. In total, we have **£5000 worth of rewards to give away** over the course of the 2019-2020 academic year, and we anticipate around 250 participants.

For further details on the project or if you have any questions, please email us at powbal@imperial.ac.uk.

Yours faithfully,
The POWBAL team

*Rewards come in the form of Tango gift cards, which are currently redeemable to Amazon, Argos, Caffe Nero, Debenhams, Decathlon, Google Play, Halfords, iTunes, John Lewis, Marks & Spencer, New Look, Nike, Pizza Express, Starbucks, Steam Wallet, Tesco, The Great British Pub (i.e. all participating UK pubs), Ticketmaster, TK Maxx, Uber, and Zalanda.

Figure 11: Participation Consent Form

Thank you for your interest in **POWBAL: Power Balancing for a Sustainable Energy Transition**.

To sign up, please confirm that you have read the POWBAL Information Sheet and understand your rights as a participant by **ticking the boxes below**:

- [ ] I confirm that I have read and understood the POWBAL Information Sheet for the above study and have had the opportunity to ask questions which have been answered fully.
- [ ] I understand that my participation is voluntary and I am free to withdraw any time, without giving any reason, without my legal rights being affected.
- [ ] I understand that I will be asked to use a smart plug that can be switched on and off remotely. I also understand that my electricity usage data may be looked at by responsible individuals from Imperial College London and London School of Economics. I give permission for these individuals to access the data collected during this research and contact me via the email address I provide below.
- [ ] I have been informed about the compensation arrangements.
- [ ] I agree to take part in the above study and allow the POWBAL research team to switch off my smart plug up to 15 minutes no more than once in a given 24-hour period during my participation in this study.

Your preferred email address:

We will send a copy of our study FAQs and brief details on what comes next to the email address listed above. We will also use this email address as the username for your brand new POWBAL account (we'll take care of the setup for you!).

Table 7: Items Connected via Smart Plugs (Debrief Survey Sample)

	Device 1	Device 2
Bluetooth speaker charger	1 (2%)	0 (0%)
Electric heater	1 (2%)	0 (0%)
Games Console	1 (2%)	1 (1%)
Hair electronics (incl. shavers)	1 (2%)	3 (6%)
Kettle	3 (6%)	0 (0%)
LED or computer screen	2 (4%)	0 (0%)
Lamp	6 (11%)	3 (6%)
Laptop charger	19 (36%)	14 (26%)
Phone charger	15 (28%)	11 (21%)
Printer	1 (2%)	0 (0%)
Plug-in diffuser	0 (0%)	1 (2%)
Refrigerator	1 (2%)	1 (2%)
Speaker	0 (0%)	1 (2%)
TV	2 (4%)	2 (4%)
No second device	N/A	15 (28%)
No Response	N/A	1 (2%)

Notes: This table reports the participants answers regarding the items that were connected via smart plugs. The column "Device 1" records the number of participants that stated a particular device. The column "Device 2" records the number of participants that stated a second device if used.