Manuscript Last Updated on Tuesday September 12, 2023 at 12:54PM
All python code written to perform our analyses (i.e., data processing, model fitting, and visualisation) is available upon request. As they contain highly sensitive information (e.g., customer postcodes, account identifiers) gathered in collaboration with a commercial entity, we are unable to share our data publicly. However, we welcome all queries about the provenance of our data with regard to any aspect of our research.

Impact of Timing and Type of Requests for Flexible Domestic Energy Use*

British Institute of Energy Economics 2023 Research Conference (Oxford, UK)

Cohen R. Simpson [†]
Centre for Net Zero
London School of Economics, Methodology
cohen.simpson@centrefornetzero.org

Andrew R. Schein
Centre for Net Zero
andrew.schein@centrefornetzero.org

Robert D. Metcalfe
Centre for Net Zero
University of Southern California, Economics
National Bureau of Economic Research
robert.metcalfe@centrefornetzero.org

ABSTRACT Changing the behaviour of private individuals is vital to the transformation of energy systems. Thus, we investigate the potential utility of energy retailers' modulation of their requests to consumers to engage in flexible domestic energy use. We do so by analysing data on total consumption (kWh) during — and formal agreement to participate in — two 60-minute energy-savings events on the part of 666,441 British households. Results from a randomised controlled trial for one event ($N \approx$ 650K) indicate that, conditional on pre-treatment variables (e.g., historical energy usage), being sent a supplementary "heads-up" email the day prior or a day-of "reminder" SMS text message respectively increased the probability of event participation by $\approx$ 6% (95% HDI $= [4.2\%, 7.8\%]$) and $\approx$ 23% (95% HDI $= [14.7\%, 31\%]$) over the control condition (i.e., day-of primary notice [email and/or SMS] only). However, whilst we find clear evidence of a negative association between in-event consumption and earlier, ancillary email-based notice ($\approx$ 4% reduction, ceteris paribus; 95% HDI $= [-10.4\%, 2.4\%]$), models are consistent with especially small versions of this effect that approach zero. And we find no compelling evidence to suggest that levels of consumption were shaped in either direction by ancillary SMS-based contact. Similarly, in a smaller quasi-experiment ($N \approx$ 80-350K) on variation in the broad timing (i.e., day-of vs. day-prior) of primary event notices of the same type with no ancillary contact, we find compelling evidence that being sent day-of (i.e., shorter) email-based notice increased consumption during the second event by $\approx$ 6.7% (95% HDI $= [0.5\%, 12.9\%]$) over the control condition (i.e., day-prior email-based notice). Yet models are again consistent with particularly small versions of this effect that approach zero. Ultimately, results indicate that the broad timing and the general type of appeals matter — where providing consumers with *supplementary* requests to save energy over and above primary messaging can play an important part in converting mere awareness of the need to conserve electricity into formal agreement to *try* to do so.

KEYWORDS Energy Consumer (Domestic); National Grid; Digital Messaging; Causal Inference; Bayesian Regression

** Electronic supplementary material is available at the end of this document.*

Acknowledgements: We are very grateful to Lucy Bourke, Sheng Chai, Daniel Lopez Garcia, Joe Grainger, Maria Jacob, Ryan Jenkinson, and Izzy Woolgar for their research assistance and formative critique — without which this research would not have been possible. We also thank Lucy Yu for her encouragement and for formally facilitating our work, Kieron Stopforth and Simona Burchill for helping to implement our field experiment, and Rebecca Dibb-Simkin, Greg Jackson, Pete Miller, Alex Schoch, Maddie Inglis, Max McShane, and Katrin Weingartner, amongst many others at Octopus Energy Group, for implementing the peak-pricing intervention. All mistakes herein are ours alone.

Author Contributions: CRS: formal analysis, investigation, methodology, project administration, validation, visualisation, writing — original draft, writing — review and editing; AS: conceptualisation, project administration, writing — review and editing; RJ: conceptualisation, methodology; RDM: conceptualisation, resources, validation, writing — review and editing.

Conflicts of Interest: All authors of this research have a financial stake — i.e., formal employment — in the Octopus Energy Group.

[†]Corresponding author.

Introduction

What communication-based mechanisms might energy suppliers use to mould the energy demand of private individuals in domestic contexts? Answering this question is vital to the successful transformation of energy systems — a task in which energy suppliers are expected to play a major role by meeting anticipated growth in demand for low-cost renewables (Bogdanov et al. 2021) and by clarifying gaps in public knowledge about the most effective means of reducing consumption.[1] Nevertheless, mitigation of the adverse effects of anthropogenic climate change is not merely a scientific and business problem. Instead, it involves a complex mix of physical, economic, and *social* dynamics (e.g., see Fankhauser et al. (2022) on threats to the success of "net zero" and Stephenson et al. (2010) on "energy cultures"). Consequently, targeted, scaleable communication with domestic consumers about their behaviour in relation to the environment is integral to efforts to transition to "greener living". Indeed, such social interaction is expected to, for instance, shape one's perception of renewable energy (Gustafson et al. 2022) and one's willingness to engage with climate science (Goldberg et al. 2021).[2] Thus, it is important to clarify the potential efficacy of energy retailers' large-scale communications campaigns to encourage consumers to change how they use electricity — especially as the viability of a request to alter one's behaviour is likely to depend on where, when, how, and by whom it is delivered (see Jost et al. (2022), Bergquist et al.(2023), Przepiorka and Horne (2020), and Zhang et al. (2018)).

To probe this issue, we analyse data from a large messaging campaign designed to alter public demand for power at key moments throughout the Winter of 2022-23. This campaign — known as the Demand Flexibility Service (DFS) — was led by Great Britain's National Grid Electricity Service Operator (NGESO). And it was comprised of 21 events (November 2022 to March 2023) for which British households were asked to reduce their energy consumption (i.e., to "turn down"). DFS events varied in duration — ranging from one to two hours. And consumer turn down was formally remunerated via a price incentive based on £/kWh saved.

Although multiple utility firms took part in the DFS, here we focus on two of 13 DFS-related events delivered by Octopus Energy (OE) — a British renewable energy provider — to its UK-based customers. In particular, we investigate whether OE's requests to participate in its variant of NGESO DFS events — branded by OE as "Saving Sessions" (Supplementary Table 2) — impacted levels of energy consumption (i.e., total kWh of electricity usage) *during events* so as to better understand the implications of the general type (i.e., What channel?) and broad time horizon (i.e., When?) of appeals to engage in flexible domestic energy use.[3] Furthermore, we exploit the two-step nature of OE's delivery of DFS-related events — i.e., customer messaging then customer agreement to participate — to probe whether the general type and broad timing of appeals to engage in flexible domestic energy use impact formal agreement to attempt reduced power consumption.

The structured nature of consumer engagement with DFS events make them an especially apt window through which to explore the efficacy of appeals to restrict power usage. Specifically, OE customers were required to explicitly agree to participate in DFS-related events (hereafter, *one-time* "sign up"). And OE customers were required to explicitly agree to participate in each individual Saving Session (hereafter, "opt in") via digital messages (hereafter, "opt-in notices"). These digital notices typically: (a) communicated the price incentive associated with participation in a specific session (Supplementary Table 2); and (b) provided customers with a hyperlink through which they could opt in to a specific session. In general, opt-in notices took the form of emails to OE account holders (Figure 1a, Figure 1b) and/or notifications to the account holder via the OE mobile application. A customer could receive multiple opt-in notices across the same or different channel (i.e., email, "push" notification to one's mobile device [i.e., phone or tablet]). Nevertheless, all analyses herein relate to the timing of a customer's first, possibly only, opt-in notice for a given Saving Session.[4]

Keeping this in mind, our research questions (RQs) are as follows:

- **RQ1**: To what extent does the broad timing and general type of an appeal to flexibly use electricity shape actual levels of energy consumption?
- **RQ2**: How does willingness to engage in energy-conserving behaviour vary with the broad timing and general type of an appeal to flexibly use electricity?

[1] See "UK energy consumers ready to flick the switch to more sustainable suppliers but green tariffs alone are not enough to stand out." by Ernst & Young Energy (2021).

[2] See also the large body of experimental research and policy reports by Yale University's long-running Programme on Climate Change Communication.

[3] We define "appeals", following Bergquist et al. (2023:4), as communications that "demand and urge people to act more sustainably by targeting their values or responsibilities". In the framework of Bergquist et al. (2023), appeals are distinct from solicitation of environmental commitments, educational messaging, feedback on one's environmental behaviour, financial incentives, and social comparisons around one's environmental behaviour.

[4] The vast majority of OE customers in the data analysed for this research received one opt-in notice via email or one email-based opt-in notice alongside one push-based opt-in notice.

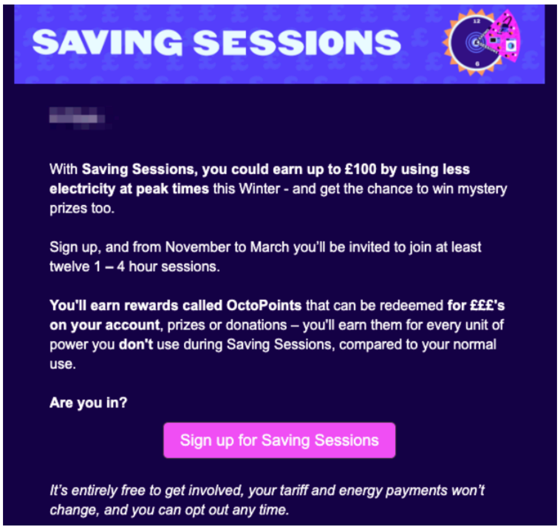

(a) Sign-Up Email

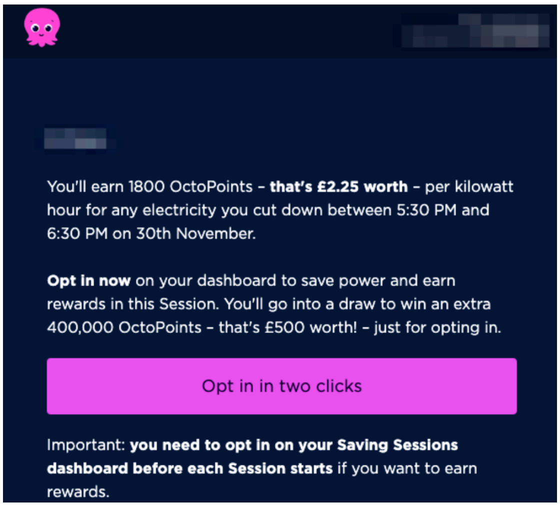

(b) Opt-In Email

Figure 1: Typical email sent to customers encouraging participation in Octopus Energy's implementation of the National Grid's Demand Flexibility Service.

To answers our RQs, we narrowly focus on two 60-minute Saving Sessions during periods of peak electricity consumption — i.e., February 13th, 2023 (17:30PM to 18:30PM) and March 15th, 2023 (16:30PM to 17:30PM).[5] We do so as the customer-messaging set-ups for these two Saving Sessions deviated from typical practice in a manner that allows us to credibly establish causal effects in relation to the broad timing and general type of opt-in notices vis-à-vis session consumption and session participation.[6]

Typically, opt-in notices were sent to OE customers on the day prior to a given Saving Session (hereafter, "day-ahead" notices). However, we use a regression discontinuity design (RDD) to exploit a technical fault that resulted in both the delay and the time-ordered delivery of opt-in notices for the Feb. 13 Saving Session in a manner broadly reflective of the length of time one has used OE as service provider and the magnitude of their account ID. Practically speaking, this allows us to gauge the effect of receiving notice *on the day of the Saving Session itself* (hereafter, "day-of" or "intraday" notices). For the Saving Session on Mar. 15, we instead exploit exogenous variation in the broad time and general type of *supplementary* opt-in notices given alongside a universal, primary day-of notice sent to all DFS-participating OE customers on the morning of Mar. 15. This exogenous variation is the result of a randomised controlled trial (RCT) that we carried out in conjunction with OE.

Our RDD and our RCT are both used to answer RQ1 and RQ2. However, as their technical aspects differ in key ways, we present each as a distinct study. Thus, we first summarise our regression discontinuity (i.e., set-up, assumptions, statistical methods) and follow immediately with results. We then provide a similarly-structured presentation of methods and results for our RCT. After, we conclude with a brief, broader discussion of our findings, their implications for researchers and policymakers, and their limitations.

Study 1: Known Assignment Mechanism, But No Randomisation

Sharp Regression Discontinuity Design (RDD): Basic Setup

Our first study concerns the causal impact of receiving an intraday opt-in notice as opposed to a day-ahead notice. The manner in which opt-in notices were sent for the Feb. 13 Saving Session was not random. However, the unexpectedly-delayed time-ordered dispersal of notices (Figure 2), which we discuss in the *Supplementary Information*, is amenable to a regression discontinuity design (RDD).

Put simply, regression discontinuity (RD) is a quasi-experimental method used to analyse observational data where the mechanism by which a treatment, policy, or exposure was assigned (i.e., administered) is entirely known but there is *no* randomisation. Specifically, given some "assignment" variable A (here, OE customer account ID) used to administer some treatment z (here, intraday opt-in notice) to individuals $i \in N$ (here, all DFS-participating OE customers by Feb. 12), RD is used to compare individuals whose values for the assignment variable A_i that

[5]By "peak consumption", we mean energy usage during times when demand on Great Britain's national grid is highest. These periods are 9:00-11:00AM and 16:30-18:30PM. See "The Big Dirty Turn Down: Deep Analysis of a Domestic Energy Flexibility Trial" by Centre for Net Zero (2021).

[6]Note, we are currently preparing a broader piece of research analysing data from all 13 Saving Sessions tentatively entitled "The Welfare Effects of Demanding Short-term Energy Flexibility".

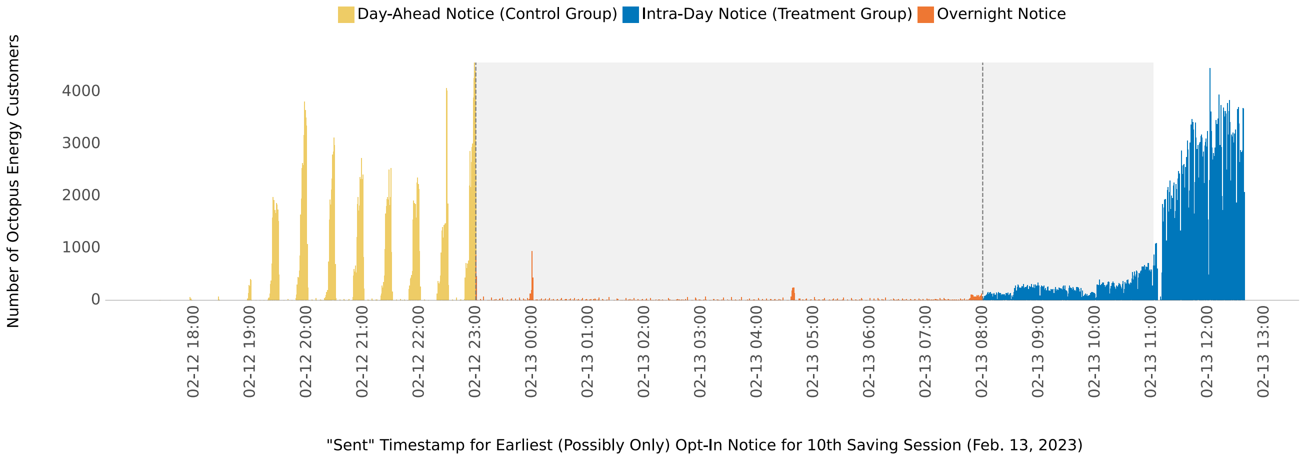

Figure 2: Distribution of the time at which an Octopus Energy customer was sent their first (possibly only) opt-in notice for the Saving Session on Feb. 13, 2023. Shaded region denotes the window of time on which our analysis of energy consumption is concentrated and corresponds to roughly 10:57PM on the 12th to one minute after 11:01AM on the 13th (see Methods (RDD) on "bandwidth"). For our analysis of agreement to participate in the 10th Saving Session we use a window of time corresponding to roughly 10:29PM on the 12th and 10:52AM on the 13th.

fall "just above and just below" a predetermined cut-off C. With some additional assumptions, and if individuals just above the cut-off (Group 1) and just below the cut-off (Group 2) are similar, a causal comparison can be made between the two groups with respect to an outcome of interest (here, session consumption and session participation).

We use a *temporal* cut-off C_{Time} equal to 8:00AM on Feb. 13 — where OE customers sent opt-in notices at or after this time are assigned to our treatment group (i.e., receipt of an intraday notice as opposed to a day-ahead notice). Because our cutoff for treatment is temporal and not a specific account ID, we must map our cutoff to an integer value reflective of the scale and the ordering of OE customers' account IDs. We do this by selecting a window of time around 8:00AM — i.e., *one second* — and identifying the single account ID closest to our temporal threshold when approaching from the left and the single account ID closest to our temporal threshold when approaching from the right based on the timestamp for when OE sent each DFS-participating customer their first (possibly only) opt-in notice.[7] We then sum these account IDs and divide by the value of two to construct an ID-based cutoff for treatment C_{ID}.

Our constructed ID-based threshold C_{ID} = "2,454,839". Account IDs range in size from "2" to "5,863,115" in our sample of 621,204 OE customers who had signed-up to participate in DFS events by Feb. 12 and for whom OE tracked during the Feb. 13 Saving Session. The constructed account ID of "2,454,839" is used as our "sharp" threshold for receipt of intraday notices. This threshold is "sharp" as only account IDs greater than or equal to "2,454,839" are regarded as receiving treatment (i.e., intraday notice). The temporal window around 8:00AM (C_{Time}) was qualitatively chosen to be as narrow as possible to reduce the number of possible account IDs from which to construct the integer based cutoff C_{ID}.

We stress that it would be impossible for OE customers to manipulate their account ID as this would be tantamount to strategically modifying creation of their OE account in relation to our threshold. Indeed, there is no way for OE customers to influence their treatment assignment as C_{Time} and, by extension, C_{ID} were determined and only known by the authors of this research in relation to OE's policy around the times at which customer communication is verboten (i.e., between 8:00PM and 8:00AM). Furthermore, owing to our data on the time at which opt-in notices were sent by OE, we know precisely which customers receive treatment given our threshold — assuming, of course, that *sent* notices are actually *received*.

Note that 11,673 of the 621,204 DFS-participating customers whom OE tracked during the Feb. 13 Saving Session were sent opt-in notices "overnight" (i.e., after 11:00PM on Feb. 12 but before 8:00AM on Feb. 13). We exclude these participants from our models owing to concerns about the stable unit treatment value assumption (i.e., "no hidden versions of treatments" (Gelman, Hill, and Vehtari 2020)). This is done under the assumption that overnight opt-in notices result in a fundamentally distinct treatment compared to the receipt of an intraday notice during working hours.

[7] Notice "sent at" timestamps measured to the millisecond.

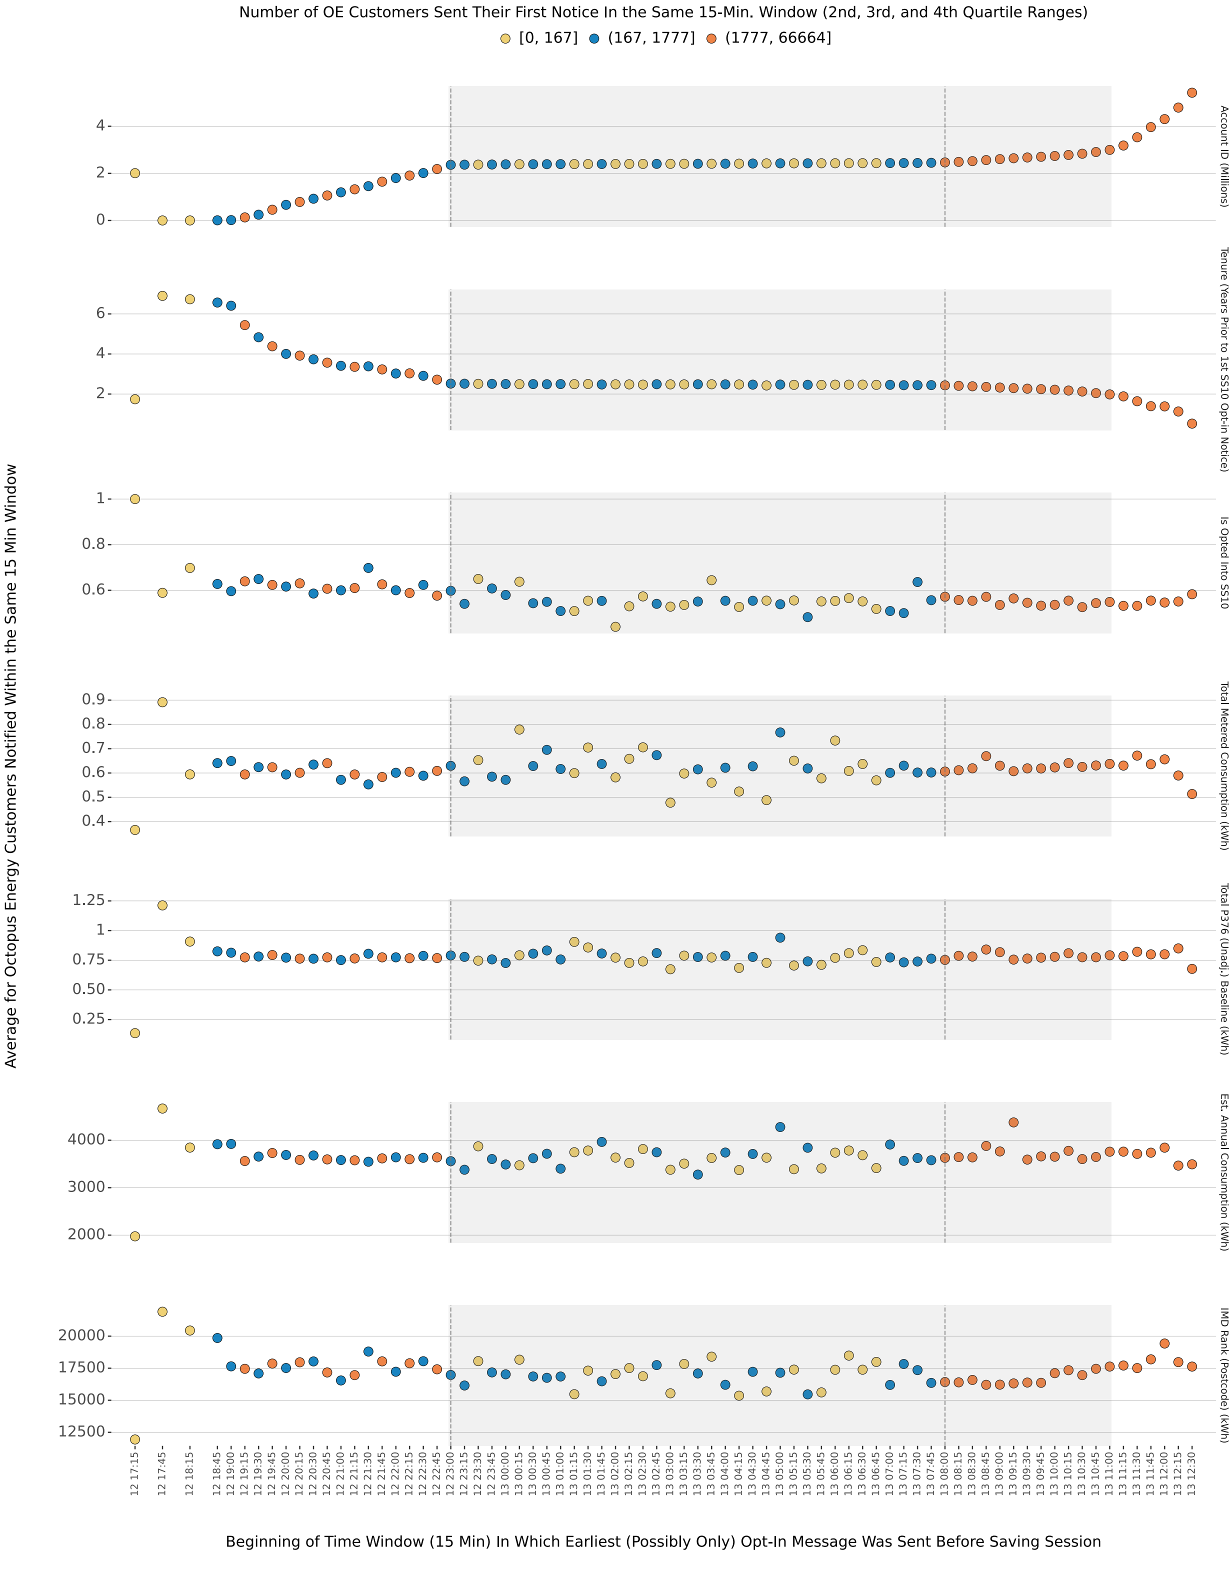

Figure 3: "Dot plot" of "binned" averages for Account ID, our two outcome variables (i.e., total session consumption in kWh and binary indicator for actual session participation or "opt in"), and relevant quantitative pre-treatment covariates. Averages across Octopus Energy customers who were sent their first (possibly only) opt-in notice for the 10th Saving Session (Feb. 13, 2023) in the same 15-minute window or "bin". Windows wherein zero opt-in notices were sent are not shown as averages for these periods are necessarily undefined. The two vertical dashed lines indicate the temporal cutoff for "overnight notices" (left) and for treatment (i.e., intraday notice; right; C_{Time}). And the shaded region denotes the window of time corresponding to the range of account IDs used to fit our models of energy consumption (see Methods (RDD) on "bandwidth"). Note, we use two measure of historic energy usage — i.e., Total P376 (Unadjusted) Baseline (kWh) and Estimated Annual Consumption (kWh). The former is an unweighted average of consumption during the same half-hour of the day during the ten most-recent working days as governed by the the P376 amendment to Great Britain's electricity balancing and settlement code. The latter is OE's predicted customer consumption based on meter readings over one year. For the Index of Multiple Deprivation (IMD), more deprived areas have *lower* postcode ranks.

Removing OE customers who received their first opt-in notice overnight necessarily results in what econometricians call a "donut-hole" regression discontinuity design (Barreca et al. 2011; Barreca, Lindo, and Waddell 2016). In a donut-hole RDD, all observations with scores on the assignment variable A_i within some range immediately around the cutoff are excluded. Typically, excluding study units in this manner is done to address "heaping" (i.e., non-random clustering of observations at points along the observed range or "support" of the running variable). However, donut-hole RDD is an elegant means of handling overnight notices. That said, whilst notice timing *is* clustered due to OE's batched dispersal of messages based on account IDs (see *Supplementary Information*), we *do not* see evidence of heaping in relevant pre-treatment covariates across the range of our running variable using the visual diagnostic (Figure 3) recommended by Barreca et al. (2016).

Finally, because $C_{\text{Time}} = 8{:}00\text{AM}$, our sample is asymmetric around our ID-based cutoff C_{ID} to the left (Figure 2). Put alternatively, we exclude study units with values for the assignment variable A that fall immediately below C_{ID} (i.e., when A_i less than "2,454,839") whilst retaining observations to the immediate right of the cutoff.[8]

Sharp Regression Discontinuity Design (RDD): Methods

Exclusion of the 11,673 OE customers sent overnight opt-in notices results in two groups divided by a sharp, account-ID-based discontinuity at $C_{\text{ID}} = 2{,}454{,}839$. Thus, $Pr(z_i = 1|A_i >= C_{\text{ID}}) = 1$ *and* $Pr(z_i = 1|A_i < C_{\text{ID}}) = 0$. In this respect, treatment status z_i is fully determined by A_i.

Crucially, we assume that OE customers with values for the assignment variable A_i that fall on either side of the cut-off C_{ID} *within a restricted range or "bandwidth"* of account IDs h — i.e., $C_{\text{ID}} - h_{Left}$ and $C_{\text{ID}} + h_{Right}$ — are broadly "similar". That is to say, we assume that OE customers very near to our cutoff have distributions of potential outcomes under treatment (i.e., y^1) and distributions of potential outcomes under no treatment (i.e., y^0) that are nearly equivalent conditional on A_i and possible confounders x (i.e., third variables determinant of both the observed outcome y and assignment A, and thus z).[9]

This assumption is sometimes called "no confounders vary discontinuously across the threshold" (Gelman et al. 2020:438), however, it is a form of conditional ignorability (Gelman et al. 2020:438) in the following style:

$$y^0, y^1 \perp z \mid A,\ x \quad \forall \quad A \in (C_{\text{ID}} - h_{Left}, C_{\text{ID}} + h_{Right}). \tag{1}$$

There is generally a tradeoff between the plausibility of this assumption and the width of the bandwidth h, and thus the number of observations N used for model fitting (see Cattaneo et al. (Forthcoming) on justification of the "local randomisation" RDD). Accordingly, and in keeping with standard practice across the academic literature, we relax the above assumption by narrowly focusing on estimating the causal effect of an intraday notice at the ID-based threshold — i.e., the causal effect when $A_i = C_{\text{ID}}$ — which is a kind of local average treatment effect (LATE).

Specifically, we estimate a (Bayesian) regression model with the following general form:

$$\begin{gathered}
y_i \sim Normal(\mu_{y,i}, \sigma_y) \\
\mu_{y,i} = \beta_0 + \beta_1 z_i + \beta_2(A_i - C_{\text{ID}}) + \beta_3(z_i \times (A_i - C_{\text{ID}})) + \beta_4(A_i - C_{\text{ID}})^2 + \beta_5(z_i \times (A_i - C_{\text{ID}})^2) + X_i\vec{\beta} + \epsilon_i \\
\forall \quad A_i \in (C_{\text{ID}} - h_{Left}, C_{\text{ID}} + h_{Right}) \\
\beta_{1,\ldots,5},\ \vec{\beta} \sim Normal(\mu = 0,\ \sigma = 0.5) \\
\beta_0 \sim Normal(\mu = 0,\ \sigma = 0.5) \\
\sigma_y \sim Gamma(\alpha = 2,\ \beta = 0.5),
\end{gathered} \tag{2}$$

where $C_{\text{ID}} = 2{,}454{,}839$, z is our binary treatment indicator (i.e., intraday notice vs. day-ahead notice) which equals the value of one when $A_i \geq C_{\text{ID}}$, and $(A_i - C_{\text{ID}})$ is the assignment variable (i.e., account ID) *relative to the cutoff*. This relative quantity equals zero when $A_i = C_{\text{ID}}$ such that $z_i = 1$ when $(A_i - C_{\text{ID}}) \geq 0$. Furthermore, X

[8] There have been important methodological advances around regression discontinuity with multiple cut-offs (see Cattaneo et al. (Forthcoming)). However, we opt for a simpler analysis by limiting our attention to a comparison of intraday versus day-ahead opt-in notices. We do so as this comparison is most relevant to practical applications of our research by energy retailers and transmission system operators. This is because these institutions are likely to limit their interaction with consumers during unsociable, non-working hours.

[9] Potential outcomes are simply one's value for the response variable when simultaneously exposed to different experimental conditions. For a given individual i, their potential outcome under treatment (y_i^1) and their potential outcome under no treatment (y_i^0) cannot both be observed. See Gelman et al. (2020, ch. 18) for an introduction to causal inference within the potential outcomes framework.

is a $N \times p$ matrix containing p pre-treatment covariates and/or confounders x believed to jointly determine both the potential outcomes (y^0, y^1) and A for study units $i \in N$ — with h_{Left} and h_{Right} defining the bandwidth of account IDs used for model fitting.

Moreover, (β_0) and $(\beta_0+\beta_1)$ are, respectively, the expected value of the response y at the threshold for customers for whom $z_i = 0$ and for whom $z_i = 1$. Thus, β_1 is the LATE at the threshold — i.e., $(\beta_0 + \beta_1) - (\beta_0)$ or the expected difference or "jump" in the outcome between customers for whom $z_i = 1$ and customers for whom $z_i = 0$. The terms $(\beta_2+\beta_4)$ and $((\beta_2+\beta_3)+(\beta_4+\beta_5))$ respectively summarise the slope (i.e., rate of change in the outcome) in relation to relative account ID $(A_i - C_{\text{ID}})$ for the control and treatment groups ($z_i = 0$ vs. $z_i = 1$). Accordingly, β_3 and β_4 are the difference *in slopes* for the linear and quadratic versions of relative account ID between the treatment and control group. However, recall that we only consider the scenario wherein $(A_i - C_{\text{ID}}) = 0$ such that associated terms "cancel out" for the purposes of interpreting β_0 and β_1 (see Huntington-Klein (2021:514–18) and Brambor et al. (2006) for discussion of multiplicative interactions).

We assume that each of our outcome variables y (i.e., session consumption in kWh, a binary indicator for session participation) adhere to Normal (i.e., Gaussian) distributions. Thus, we map the linear predictor in Equation (2) to each outcome using the identity link in the style of ordinary-least-squares (OLS) regression. In the *Supplementary Information*, we discuss the limited (i.e., restricted) range of our outcome variables in relation to our use of a Gaussian functional form. In the *Supplementary Information*, we also discuss our general model specification and the technique we use to estimate optimal, *asymmetric* bandwidths h_{Left} and h_{Right} given our use of donut-hole RDD. These optimal bandwidths are specific to each of our two outcome variables (i.e., session consumption and session participation). For session consumption, the lower bandwidth bound $(C - h_{Left})$ roughly correspond to 10:57PM on the 12th and the upper bandwidth bound $(C + h_{Right})$ roughly corresponds to 11:01AM on the 13th (Figure 2, Figure 3). For session participation, the lower and upper bounds roughly correspond to 10:29PM on the 12th and 10:52AM on the 13th.

Thus, our RDD has two experimental conditions with the following sizes depending on bandwidth:

- Control Group [Consumption] (N = 15,973): Day-Ahead Opt-in Notice
- Control Group [Opt-in] (N = 43,719): Day-Ahead Opt-in Notice
- Treatment Group [Consumption] (N = 63,216): Intraday Opt-in Notice
- Treatment Group [Opt-in] (N = 55,959): Intraday Opt-in Notice

Finally, we fit the model summarised by Equation (2) within a Bayesian framework. Models were written and estimated with the Python-based probabilistic programming language "PyMC". Bayesian estimation requires one to specify "priors" — i.e., probability distributions that summarise information about each unknown model parameter that is generally external to the data themselves. In the *Supplementary Information*, we discuss why we prefer Bayesian inference over frequentist inference and our choice of prior for each parameter in Equation (2). It suffices to say that we use weakly-informative priors for the coefficients (given the scale of our data) which are Gaussian in shape, centred on zero (i.e., no effect), and that have large probability density at zero (i.e., regularisation). The Gamma prior for σ is diffuse and strictly positive.

Sharp Regression Discontinuity Design (RDD): Results (Main)

Figure 4 is a "bird's-eye view" of our core results. It depicts posterior means for the causal effect (i.e., LATE) of receiving an intraday opt-in notice (vs. a day-ahead notice) from regression discontinuity models of total consumption (kWh) during, and formal agreement to participate in, the Feb. 13 Saving Session. Posterior means are accompanied by 95% highest posterior density intervals. Analogous to frequentist confidence intervals, these "HDIs" indicate the range of values within which 95% of the sampled posterior values with the highest probability fall. Thus, all values within a given HDI are more likely (i.e., more "credible") than those outside. For each outcome, we present posterior means and HDIs from three separate models.[10] These models differ based on specification, specifically the inclusion of potential confounders and quadratic terms (i.e., "Baseline" vs. "Extended"). Supplementary Figure 15 depicts the distribution of posterior values for the LATE under each model in relation to

[10] Results obtained using 48,000 posterior draws and four parallel chains. For each chain, the first 6,000 draws were used for "warm up" and discarded. This resulted in a posterior sample size of 24,000 for each parameter (6,000 draw/chain). Chain convergence was diagnosed using "Bulk" effective sample size (ESS) and "tail" ESS — i.e., estimates of the amount of independent information in the centre and the tails of a posterior distribution or, rather, the posterior sample size in absence of autocorrelation. Vehtari et al. (2021) suggest a per-chain ESS of at least 100. However, Kruschke (2015:184–87) recommends an ESS of 10,000 owing to the general instability of the tails of posterior distributions. To manage computational burden given our large samples, we run chains to achieve an ESS of both types of at least 5,000 (aggregating across chains) for all parameters in all of our models. Chain behaviour was also diagnosed using $\hat{R}$ — i.e., a diagnostic for comparing within- and between-chain variance (Vehtari et al. 2021). Values of $\hat{R}$ less than 1.01 are generally regarded as indicative of chain convergence. And $\hat{R}$ is less than 1.01 for all parameters in all of our models.

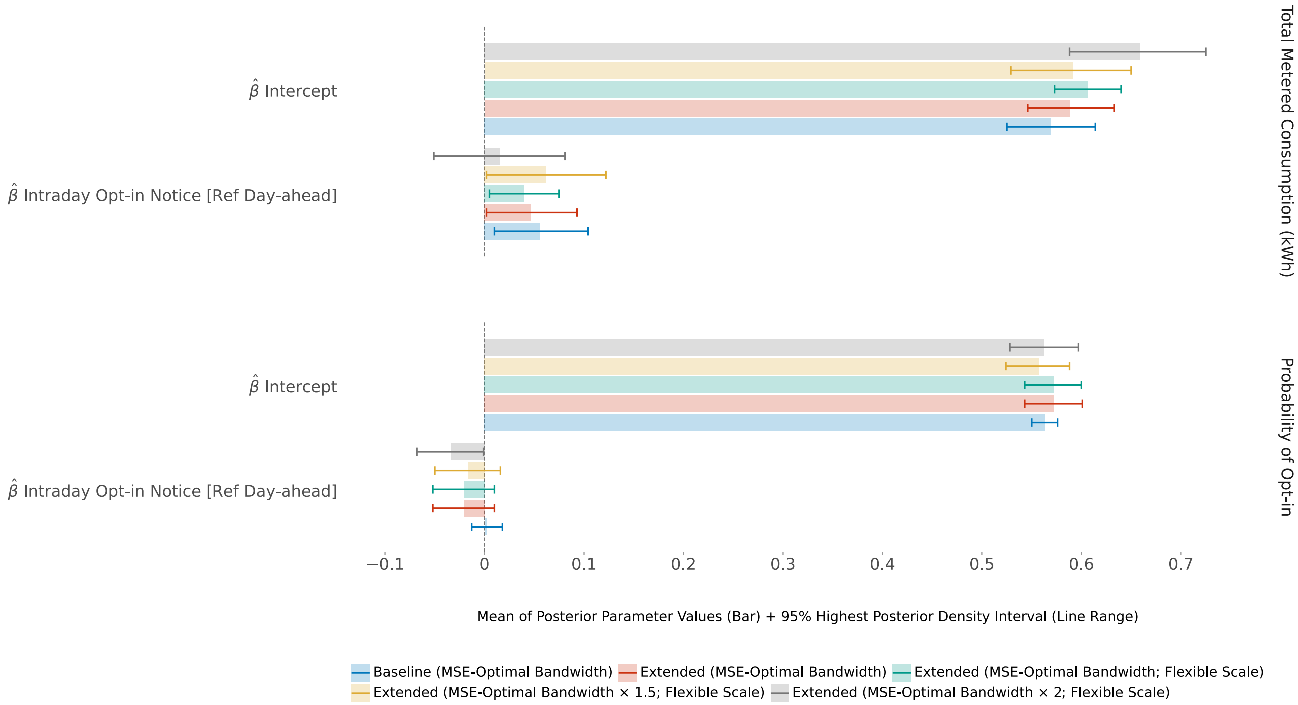

Figure 4: "Forrest plot" depicting posterior means for the LATE ($\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead Notice]) and the expected average outcome in the control group ($\hat{\beta}$ Intercept), holding all other covariates constant, from regression discontinuity models fit to data from the 10th Saving Session (Feb. 13, 2023). Posterior means are accompanied by a 95% HDI. Outcome variable y given on right-hand side of plot. Models are fit to subsets of our Saving Session data using bandwidths optimised to reduce mean-squared error (MSE). Results obtained using three model specifications. The first ("Baseline") only adjusts for the assignment variable, the binary treatment, and their multiplicative interaction (Supplementary Tables 3 and 9). The second ("Extended") adds to the first a quadratic term for the assignment variable and relevant pre-treatment variables such as a customer's historical energy usage, their region of residence, and the degree to which their postcode is deprived (Supplementary Tables 4 and 10). The third is a variant of the "Extended" specification wherein we account for heteroscedasticity using a linear predictor for the scale of the outcome σ. The linear predictor for σ includes only those pre-treatment covariates used to model the conditional mean in the non-flexible "Extended" model (see parameters ϕ in Supplementary Tables 5 and 11; see *Supplementary Information*). Finally, we consider how our results might vary under different ranges of our running variable by taking the "Extended" (Flexible Scale) specification and expanding our MSE-optimal bandwidth by a factor of 1.5 and a factor of two (Supplementary Tables 6, 7, 12 and 13). The number of observation used to fit each of our models appear in Table 1.

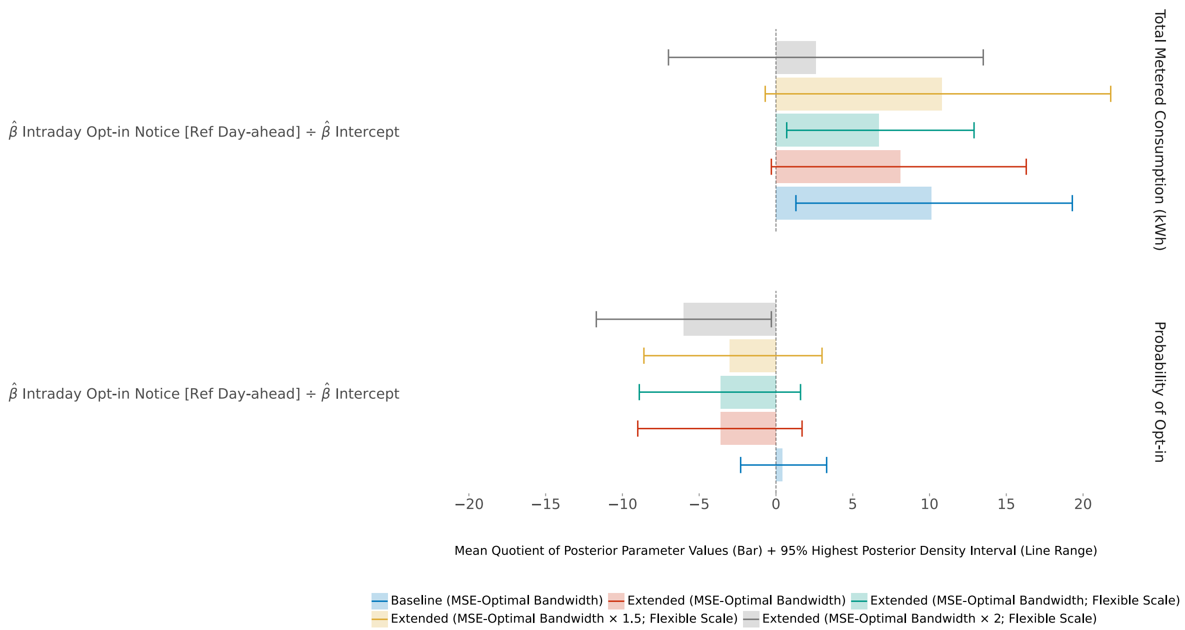

Figure 5: Percentage change in outcome y over the baseline for the control group (i.e., $\hat{\beta}$ Intercept) for the binary treatment given the LATE ($\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead Notice]), holding all other covariates constant. Outcome variable y given on right-hand side of plot. Posterior mean parameter values appear in Figure 4. However, division is performed using *posterior-sample-specific values* for each parameter. This yields a distribution of percentage changes which we summarise using 95% HDIs.

our Region of Practical Equivalence (ROPE), the Probability of Direction (POD), the 95% HDI, and the posterior mode — i.e., *the* most-likely value for a parameter under the fitted model.[11]

Focusing on our simplest model of consumption ("Baseline"), the posterior mean LATE (i.e., $\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead Notice]; Figure 4) indicates that OE customers sent an intraday notice used, on average, 0.056 kWh (95% HDI $= [0.01, 0.104]$) more during the Saving Session compared to OE customers sent a day-ahead notice — the latter of whom had an estimated average consumption of 0.569 kWh (95% HDI $= [0.525, 0.614]$). Put alternatively, consumption on the part of customers sent a day-of notice was higher by $\approx$ 10.1% (95% HDI $= [1.3, 19.3]$, Figure 5, Figure 6). Note, however, that across our three models using the same bandwidth, the posterior mean LATE for intraday notice is sensitive to pre-treatment covariates (e.g., historical energy use) in terms of its magnitude — where it is attenuated to *approx* 0.4 in our "Extended (Flexible Scale)" model (95% HDI $= [0.005, 0.075]$), holding all else constant (typically at the sample average). Furthermore, across the three models using the same bandwidth, $\approx$ 2-5% of the LATE's posterior distribution falls within our ROPE — a range of values we regard as null for practical purposes.

Strictly speaking, a conservative interpretation (Kruschke 2018) would see this posterior behaviour classified as inconclusive give our ROPE. Nevertheless, the LATE's POD for being greater than zero is $\approx$ 97-99% regardless of covariates (Supplementary Figure 15). Thus, under our three models using the same bandwidth, there is near certainty that the LATE "exists" using the decision-making scheme of Makowski et al. (2019) and that the LATE is positive (i.e., increased consumption under short notice). Accordingly, there is uncertainty around the LATE's "significance". As we discuss in the *Supplementary Information* under "Adjudication on Evidence", there is great value in even very small effects for our particular research application. And the clear POD ultimately leads us to conclude that there is compelling evidence to suggest that being sent an intraday notice had a positive causal impact on energy consumption during the Feb. 13 Saving Session for the customers in our bandwidth-specific

[11] Bayesian inference yields *entire distributions of possible values for some parameter of interest*, not just a single point estimate. We judge whether there is evidence of an effect using properties of posterior distributions in line with the existence-significance decision-making framework of Makowski et al. (2019) and Kruschke (2018). By "existence" we mean "the consistency of an effect in one particular direction (i.e., positive or negative), without any assumptions or conclusions as to its size, importance, relevance, or meaning" (Makowski et al. 2019:9). Crucially, this is conceptually distinct from "significance" — i.e., is an effect "worthy of attention" such that its magnitude is not "likely to be too small to be of high importance in real-world scenarios or applications" (Makowski et al. 2019:9)? We judge existence of a treatment effect using the Probability of Direction (POD). The POD is the proportion of the posterior distribution either side of zero (i.e., no effect). We judge significance of an effect using the proportion of the entire posterior distribution contained within a Region of Practical Equivalence (ROPE) — i.e., a range of values around zero that are, for practical purposes, *all* considered to be null. In the *Supplementary Information*, we further discuss these continuous metrics of evidence and how we set our ROPEs.

sample with two provisos. The first is that there is a trivial possibility that the LATE is so tiny as to be practically unimportant. And the second is that the LATE may not apply to broader time periods as it is sensitive to our widest bandwidth (PMEAN = 0.016; 95% HDI = $[-0.051, 0.081]$. Put simply, the broad timing of notices of the same type appears to have mattered in this particular scenario (RQ1), albeit possibly to a small degree and only for a narrow set of customers.

As for participation, our linear probability models are inconclusive. Specifically, there is no compelling evidence to suggest that the broad timing of notices of the same type impacted the likelihood of opting into the Feb. 13 Saving Session (RQ2) in either direction. Amongst our three models using the same bandwidth, the posterior mean LATE (i.e., $\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead Notice]) switches from 0.002 to -0.02 when we adjust for pre-treatment covariates, the latter of which indicates a decrease in the probability of opt-in of $\approx 3.6\%$ (95% HDI = $[-9, 1.7]$, Figure 5, Figure 6) for those receiving intraday notice compared to the opt-in probability of those in the control group (PMEAN = 0.225; 95% HDI = $[0.224, 0.227]$). Nevertheless, under all three linear probability models using the same bandwidth, the LATE has a non-trivial POD for having a sign opposite to that of the posterior mean (Supplementary Figure 15) and, depending on model, $\approx$ 10-45% of the LATE's posterior distribution includes our ROPE. These results do not appear to be strongly dictated by bandwidth.

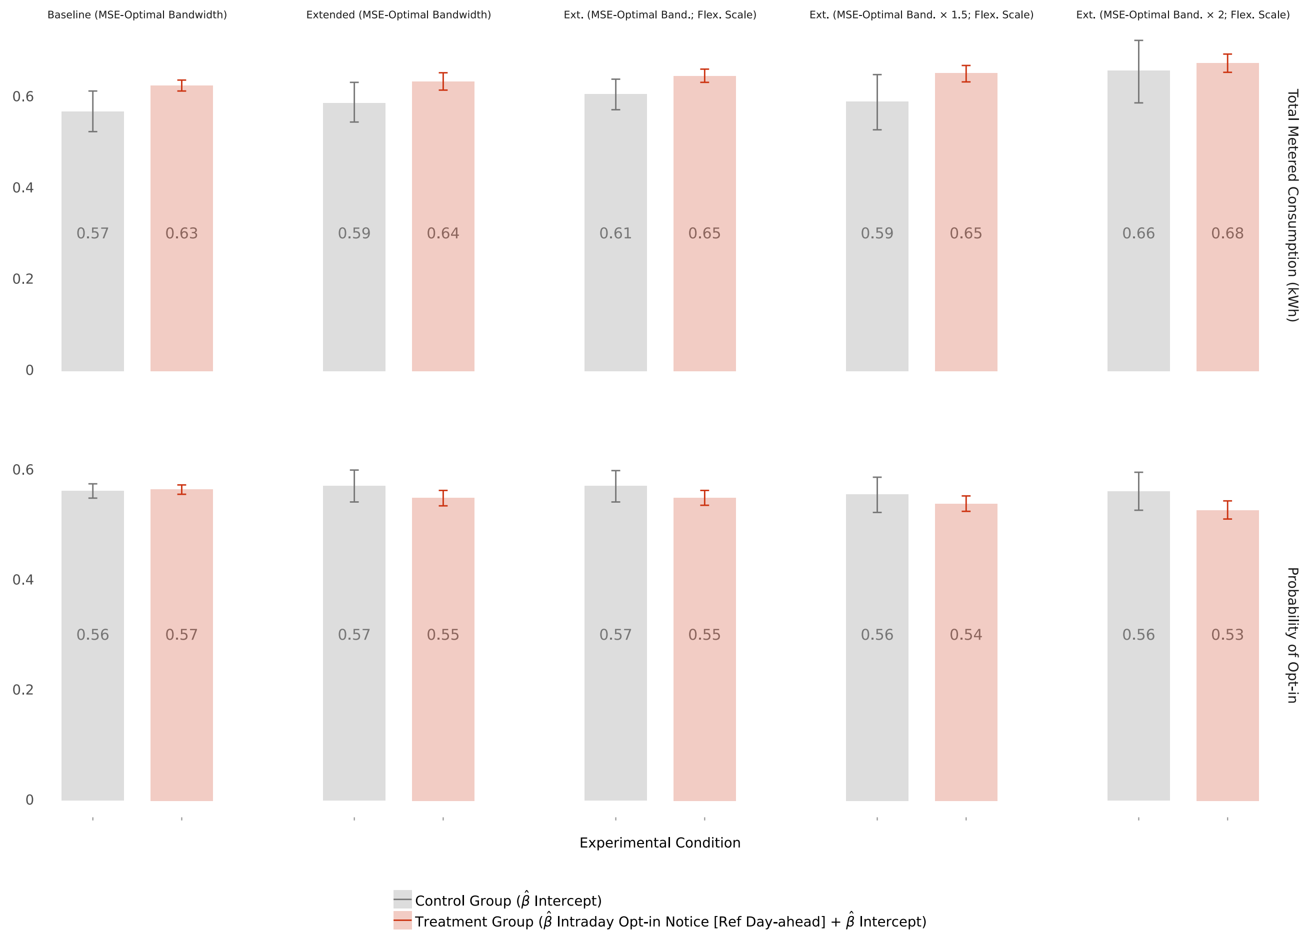

Figure 6: Expected value of the outcome y for the control and treatment condition, holding all other covariates constant. Outcome variable y given on right-hand side of plot. Posterior mean parameter values appear in Figure 10. However, summation for the experimental condition is performed using *posterior-sample-specific values* for each parameter. This yields a distribution of expected outcomes which we summarise using 95% HDIs.

Sharp Regression Discontinuity Design (RDD): Results (Exploratory Modelling)

Before advancing, we briefly consider treatment-effect heterogeneity in relation to the timing of intraday notice using a pair of ancillary models. Recall that our binary indicator for intraday notice covers all notices sent roughly between 08:00AM and 11:00AM on Feb. 13 (Figure 2). This leads to a comparison between customers sent notice within minutes of 11PM the day prior (i.e., the control group; Figure 2) and customers sent day-of notice over multiple hours across the morning of the Saving Session. Accordingly, we explore whether our results might be consistent throughout the morning-to-afternoon period by fitting two addition models wherein we swap on singular

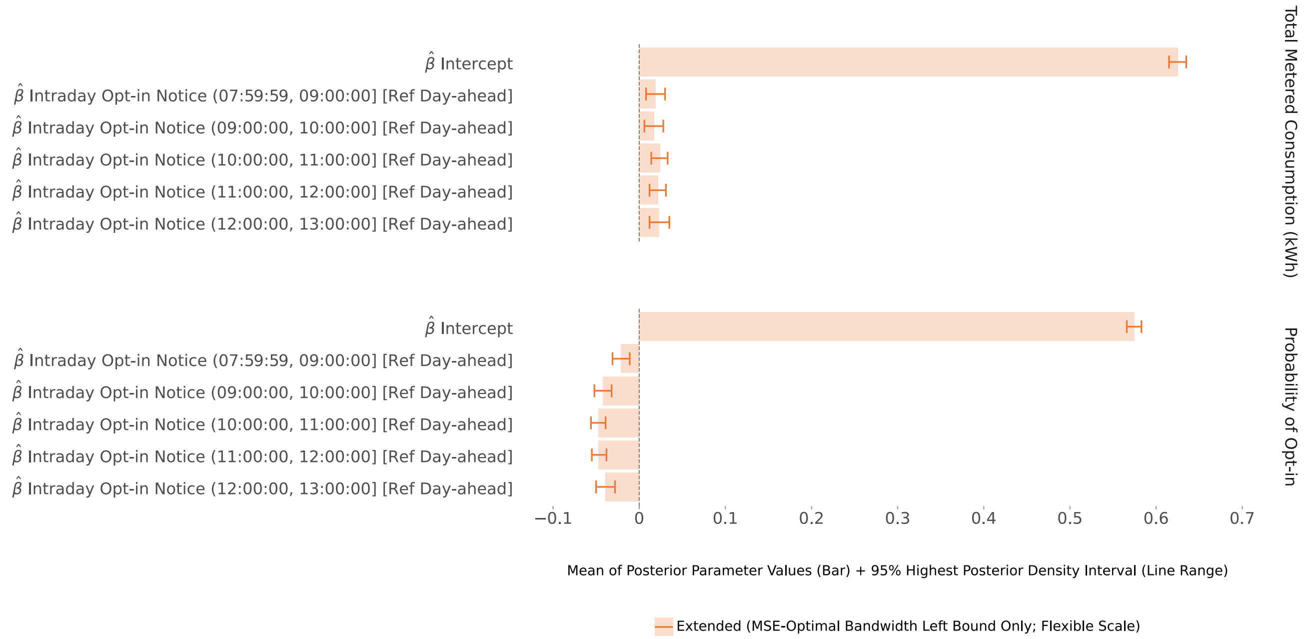

Figure 7: "Forrest plot" depicting posterior means for the hour-specific ATEs ($\hat{\beta}$ Intraday Opt-in Notice (Time Interval) [Ref Day-ahead Notice]) and the expected average outcome in the control group ($\hat{\beta}$ Intercept), holding all other covariates constant, from regression discontinuity models fit to data from the 10th Saving Session (Feb. 13, 2023) filtered only using our lower bandwidth h_{Left}. Posterior means are accompanied by a 95% HDI. Models fit using our "Extended (Flexible Scale)" specification and thus adjust for relevant pre-treatment variables such as a customer's historical energy usage and the degree to which their postcode is deprived. Outcome variable y given on right-hand side of plot. Full results given in Supplementary Tables [-tbl-results-RDD-consumption-flexible-extended-mse-leftonly] and [-tbl-results-RDD-optin-flexible-extended-mse-leftonly])

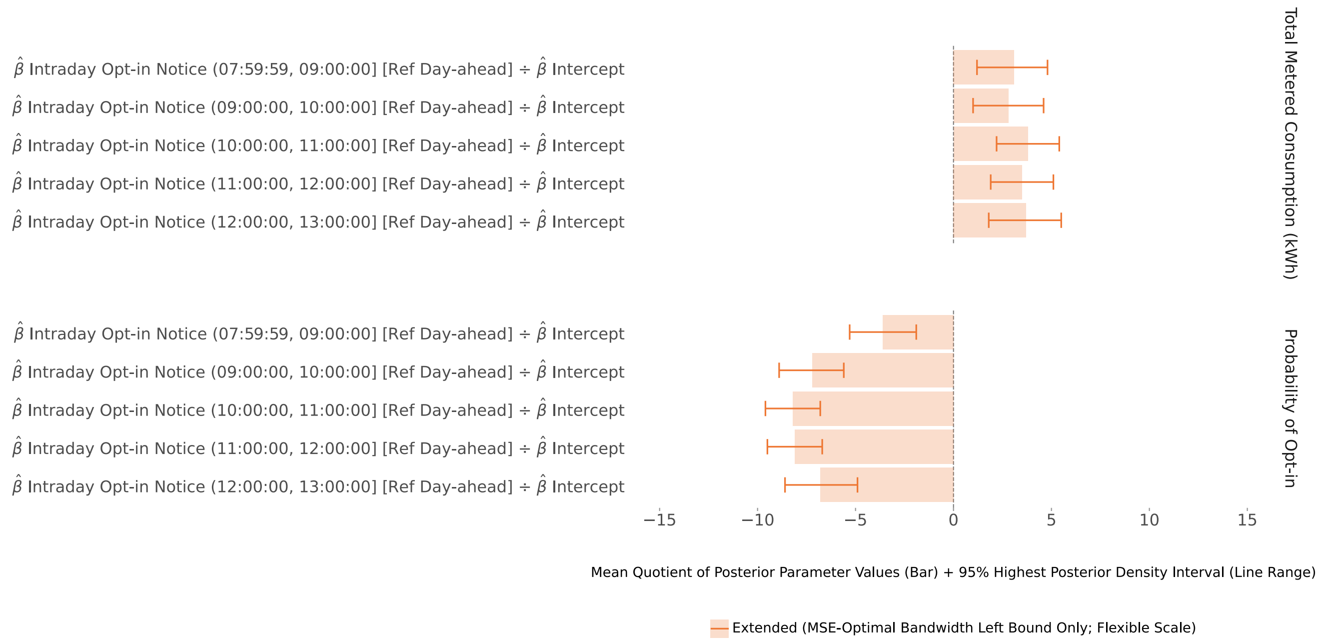

Figure 8: Percentage change in outcome y over the baseline for the control group (i.e., $\hat{\beta}$ Intercept) for the binary treatment given the hour-specific ATE ($\hat{\beta}$ Intraday Opt-in Notice (Time Interval) [Ref Day-ahead Notice]), holding all other covariates constant. Outcome variable y given on right-hand side of plot. Posterior mean parameter values appear in Figure 7. However, division is performed using *posterior-sample-specific values* for each parameter. This yields a distribution of percentage changes which we summarise using 95% HDIs.

binary indicator for intraday notice with a series of binary indictors for the 60-min period within which day-of notice was sent (i.e., every hour from 08:00AM and 13:00PM, the cut-off for the sending of opt-in notice).

In so doing, we necessarily approach our RDD through the lens of a standard observational study wherein we wish to estimate a causal effect by adjusting for all confounders (Gelman et al. 2020:437). Thus, we assume that, conditional on pre-treatment covariates — most importantly customer tenure — assignment to the control or to the hour-specific treatment groups is independent of the potential outcomes within a restricted range of our running variable (See Equation (1) as well as Cattaneo et al. on (Forthcoming) on "local randomisation"). Put simply, we assume ignorability of our notice-based treatments conditional on tenure for a limited range of account IDs, amongst other factors — recalling that the ID-based assignment variable is itself a function of tenure (Figure 3).[12]

We estimate our ancillary models simply by filtering our data using our lower bandwidth bound ($C - h_{Left}$) such that all customers with account IDs A_i greater than this bound are used for model fitting. Thus, we continue to use as a control group customers sent notice immediately before 11:00PM on Feb. 12th.

As this form of RDD focuses on the average difference between the control and treatment groups (i.e., ATEs) as opposed to the causal effect at the cutoff C (i.e., the LATE), it makes no functional form assumption (Cattaneo et al. Forthcoming) and standard techniques analytic techniques can be used (Cattaneo et al. Forthcoming). Accordingly, the two ancillary models are fit using our "Extended (Flexible Scale)" specification whilst dropping as predictors account ID, the quadratic transformation of account ID, and the multiplicative interactions between treatment and account ID.

Keeping all of this in mind, estimates (Figure 7, Figure 8) indicate that results from our main RDD models are consistent throughout the morning of the 13th in terms of their signs, where the magnitudes of the ATEs do not dramatically change over time.

Table 1: Number of observations N used to fit each model for our primary analyses of Total Metered Consumption (kWh) or the Probability of Opt-in for a given Saving Session in relation to our regression discontinuity design (RDD) or randomised controlled trial (RCT) and after dropping missing values in relation to the model specifications in Figure 4 and Figure 10

Design	y	Specification	N	Results
RDD	Total Consumption (kWh)	Baseline (MSE-Optimal Band.)	78,724	Table 3
RDD	Total Consumption (kWh)	Extended (MSE-Optimal Band.)	69,223	Table 4
RDD	Total Consumption (kWh)	Extended (MSE-Optimal Band.; Flex. Scale)	69,223	Table 5
RDD	Total Consumption (kWh)	Extended (MSE-Optimal Band. × 1.5; Flex. Scale)	96,558	Table 6
RDD	Total Consumption (kWh)	Extended (MSE-Optimal Band. × 2; Flex. Scale)	123,373	Table 7
RDD	Total Consumption (kWh)	Extended (MSE-Optimal Band. Left Only; Flex. Scale)	350,834	Table 8
RDD	Probability of Opt-in	Baseline (MSE-Optimal Band.)	99,678	Table 9
RDD	Probability of Opt-in	Extended (MSE-Optimal Band.)	88,501	Table 10
RDD	Probability of Opt-in	Extended (MSE-Optimal Band.; Flex. Scale)	88,501	Table 11
RDD	Probability of Opt-in	Extended (MSE-Optimal Band. × 1.5; Flex. Scale)	104,509	Table 12
RDD	Probability of Opt-in	Extended (MSE-Optimal Band. × 2; Flex. Scale)	125,236	Table 13
RDD	Probability of Opt-in	Extended (MSE-Optimal Band. Left Only; Flex. Scale)	378,091	Table 14
RCT	Total Consumption (kWh)	Baseline	638,242	Table 15
RCT	Total Consumption (kWh)	Extended	547,055	Table 16
RCT	Total Consumption (kWh)	ITT (Extended Variant)	547,055	Table 17
RCT	Total Consumption (kWh)	ITT (Extended Variant; Flex. Scale)	547,055	Table 18
RCT	Probability of Opt-in	Baseline	650,809	Table 19
RCT	Probability of Opt-in	Extended	558,290	Table 20
RCT	Probability of Opt-in	ITT (Extended Variant)	558,290	Table 21
RCT	Probability of Opt-in	ITT (Extended Variant; Flex. Scale)	558,290	Table 22

[12] Note, we drop P376 Baseline Consumption from our ancillary model for the probability of opt in as its inclusion led to divergent transitions (i.e., ill-behaved chains) during posterior sampling.

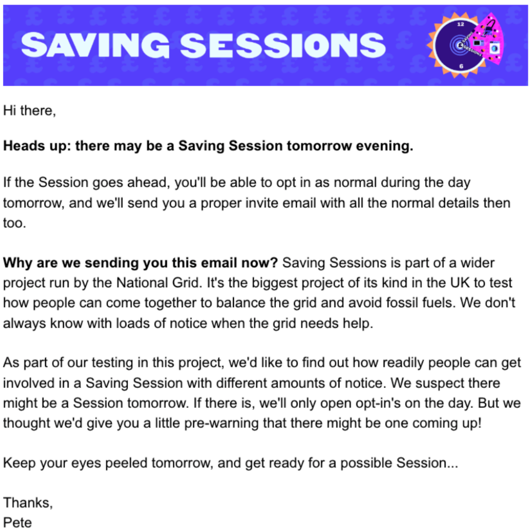

SAVING SESSIONS

Hi there,

Heads up: there may be a Saving Session tomorrow evening.

If the Session goes ahead, you'll be able to opt in as normal during the day tomorrow, and we'll send you a proper invite email with all the normal details then too.

Why are we sending you this email now? Saving Sessions is part of a wider project run by the National Grid. It's the biggest project of its kind in the UK to test how people can come together to balance the grid and avoid fossil fuels. We don't always know with loads of notice when the grid needs help.

As part of our testing in this project, we'd like to find out how readily people can get involved in a Saving Session with different amounts of notice. We suspect there might be a Session tomorrow. If there is, we'll only open opt-in's on the day. But we thought we'd give you a little pre-warning that there might be one coming up!

Keep your eyes peeled tomorrow, and get ready for a possible Session...

Thanks,
Pete

Figure 9: Day-ahead "heads-up" email sent to customers to raise awareness about an upcoming Saving Session.

Study 2: Random Assignment Mechanism With Non-Compliance

Randomised Encouragement Design (RED): Basic Setup

Like Study 1, our second study concerns the causal impact of opt-in notice timing. However, it benefits from a far larger sample for the purposes of model fitting and a randomised controlled trial based on an inversion of the customer-messaging set-up used for the Feb. 13 quasi-experiment.

Specifically, Study 2 concerns 650,809 OE customers who had signed-up to participate in Demand Flexibility Service events by Mar. 14 and whom OE tracked during the Mar. 14 Saving Session. Unlike the other Saving Sessions throughout the Winter of 2022-23, all 650,809 customers were sent *intraday* opt-in notices for the Mar. 15 Saving Session. Accordingly, for our RCT, we randomly selected a subset of customers to receive *supplementary messaging* and, in some cases, a *supplementary price incentive* on top of the standard intraday notice.

Of the 650,809 OE customers, 19,182 were randomly assigned to receive a *day-ahead* "heads-up" email on Mar. 14 (Figure 9). We refer to this notice as a "heads-up" email as those in receipt of this message *could not* use it to formally agree to participate in the Saving Session on Mar. 15 (c.f. opt-in notices). Indeed, the heads-up email only informed DFS-participating customers about the upcoming Saving Session and its general importance. Distribution of the day-ahead heads-up email was managed using a third-party platform (i.e., "SendGrid") that is operationally distinct from OE's internal customer-messaging platform used to send the intraday notices (discussed above in relation to Study 1).

Furthermore, 19,220 of the 650,809 OE customers were randomly assigned to receive an *intraday* "reminder" SMS text message on Mar. 15. And these 19,220 customers were made eligible for a performance-related bonus price incentive of 1,000 "OctoPoints" worth £1.25. Like the heads-up email, the SMS reminder text raised awareness of the upcoming Saving Session. However, customers in receipt of a text may have already received primary notice (hence, "reminder"). Furthermore, the SMS text disclosed the level of bonus on offer subject to positive session performance.[13] Owing to their abbreviated length, the SMS text did not reference the general importance of Saving Sessions. The exact text of the SMS was as follows:

> *"SAVING SESSION TODAY 1830-1930. SPOT PRIZE: Octobot has chosen you at random to win 1000 OctoPoints if you save energy in tonight's Session. Opt in before 1830!"*

Thus, our RCT has three experimental conditions:

1. Control Group ($N = 627{,}155$): Intraday Opt-in Notice Only
2. Treatment Group 1 ($N = 19{,}182$): Intraday Opt-in Notice *plus* Day-ahead "Heads-up" Email
3. Treatment Group 2 ($N = 19{,}220$): Intraday Opt-in Notice *plus* Intraday "Reminder" SMS Text *plus* Eligibility for £1.25 Bonus

Two factors make random assignment of our second treatment imperfect. First, some customers assigned to the SMS-plus-bonus condition had disallowed SMS communications from OE. Second, we were unable to send SMS texts to every customer assigned to the SMS-plus-bonus condition who *had* allowed SMS messages from OE.

[13] Session performance was deemed positive when a customer's session consumption was less than their historical average consumption.

Specifically, we were limited to sending intraday SMS texts to a *maximum* of 5,000 OE customers. In total, 4,731 of the 19,220 OE customers assigned to the SMS-plus-bonus condition *did not* allow SMS texts from OE. Of the remaining 14,489 customers who did allow SMS communications, 4,472 were randomly sub-sampled to receive an intraday SMS reminder. Thus, there were 14,748 customers (i.e., 4,731 + 10,017) assigned the second treatment who did not actually receive the second treatment. Nevertheless, all 19,220 customers assigned to the SMS-plus-bonus condition were made eligible to receive the bonus price incentive regardless of whether they disallowed SMS communications from OE and irrespective of whether they were a part of the random sub-sample. In these respects, our second treatment suffers from imperfect compliance.

We are interested in the causal effect of actually receiving an intraday SMS text — again assuming, similarly to Study 1, that *sent* notices are actually *received*. Thus, our binary indicator for the second treatment only reflects the 4,472 OE customers who had allowed SMS communication from OE and who were sent an intraday SMS reminder after random sub-sampling. Consequently, the 4,731 customers who had disallowed SMS communication and the 10,017 customers who had allowed SMS communication but who were not randomly sub-sampled are only included in our binary indicator for *eligibility* for our second treatment.

Note well that these 14,748 customers only received the intraday email. And they were not made aware of their eligibility for the bonus price incentive unless they met the bonus criterion by the end of the Saving Session — where winnings were disclosed after the Saving Session. If these individuals *had* been informed, it would represent a distinct form of treatment and it would have been prudent to create a fourth treatment group for "intraday notice plus bonus eligible". But this was not the case.

Keeping all of this in mind, we draw the following contrasts across the 650,809 OE customers who had signed-up to participate in DFS events by Mar. 14 and for whom OE tracked during the Mar. 15 Saving Session as a part of our RCT:

1. Intraday Opt-in Notice + Day-ahead "Heads-up" Email (Treatment Group 1) vs. Intraday Opt-in Notice Only (Control Group)
2. Intraday Opt-in Notice + Intraday "Reminder" SMS Text [Actually Received] + Eligibility for £1.25 Bonus (Treatment Group 2) vs. Intraday Opt-in Notice Only (Control Group)

Given non-compliance, we formulate an answer to RQ1 and RQ2 through the lens of a randomised encouragement design (RED) — i.e., a type of experimental setup wherein variation in some difficult-to-directly-manipulate treatment is induced using a source of random variation (i.e., the random "encouragement") that is related to the difficult-to-directly-manipulate treatment *and not* related to the outcome of interest. Here, our random "encouragement" is our *original*, explicit random assignment to the SMS-plus-bonus condition.

Note, due to non-compliance, we can only estimate the Complier Average Causal Effect (CACE) for the SMS-plus-bonus condition. The CACE — a kind of LATE — is the causal estimand for customers whose receipt of an intraday SMS and awareness of their eligibility for the bonus price incentive could be altered by our randomisation. This is distinct from the Average Treatment Effect (ATE) which we estimate for the day-ahead-email condition alongside the CACE using instrumental variable (IV) estimation following Gelman et al. (2020), Bhuller and Sigstad (2005), Rossi et al. (2005), Lopes and Polson (2014), and McElreath (2020). We further explain and justify our RED framing and our use of IV estimation in the *Supplementary Information*.

Randomised Encouragement Design (RED): Methods

We estimate the ATE and the CACE for our supplementary-notice conditions using a (recursive) simultaneous-equation Bayesian regression model. This is the standard recommend approach to IV estimation amongst Bayesian statisticians (Gelman et al. 2020; Lopes and Polson 2014; McElreath 2020; Rossi et al. 2005; see also, Greene 2019; Kleibergen and Zivot 2003; Wooldridge 2010). It involves using a two-dimensional multivariate Normal distribution to jointly fit two models. The first model is of our outcome y (i.e., session consumption or session participation) conditional on our treatment with imperfect compliance T (i.e., a binary indicator for the 4,472 customers who received the SMS-plus-bonus treatment). The second model is of T conditional on our binary random encouragement Z (i.e., an indicator for the 19,220 customers merely assigned to the SMS-plus-bonus condition).

Put succinctly, we fit models comprised of two, interrelated regression equations with correlated errors $\epsilon_{i,y}$ and $\epsilon_{i,T}$. This is broadly analogous to a frequentist two-stage least-squares procedure using OLS regression.

Formally, our combined model for the ATE for the day-ahead-email condition and the *CACE* for the SMS-plus-bonus condition is as follows:

$$\begin{bmatrix} y_i \\ T_{\text{SMS},i} \end{bmatrix} \sim MVNormal_2\left(\begin{bmatrix} \mu_{y,i} \\ \mu_{T_{\text{SMS},i}} \end{bmatrix}, \Sigma\right) \tag{3}$$

$$\mu_{y,i} = \beta_0 + \beta_1 T_{\text{Day-ahead Email},i} + \beta_2 T_{\text{SMS},i} + X_i\vec{\beta} + \epsilon_{i,y} \tag{a}$$

$$\mu_{T_{\text{SMS}},i} = \gamma_0 + \gamma_1 Z_{\text{SMS},i} + X_i\vec{\gamma} + \epsilon_{i,T} \tag{b}$$

$$\Sigma = \begin{bmatrix} \sigma_y^2 & \rho\sigma_y\sigma_{T_{\text{SMS}}} \\ \rho\sigma_y\sigma_{T_{\text{SMS}}} & \sigma_{T_{\text{SMS}}}^2 \end{bmatrix} = \underbrace{\begin{bmatrix} \sigma_y & 0 \\ 0 & \sigma_{T_{\text{SMS}}} \end{bmatrix}}_{\text{Diag}(\sigma)} \cdot \underbrace{\begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}}_{\Omega} \cdot \underbrace{\begin{bmatrix} \sigma_y & 0 \\ 0 & \sigma_{T_{\text{SMS}}} \end{bmatrix}}_{\text{Diag}(\sigma)}$$

$$\beta_1,\ \beta_2,\ \vec{\beta},\ \gamma_1,\ \vec{\gamma} \sim Normal(\mu = 0,\ \sigma = 0.5)$$

$$\beta_0,\ \gamma_0 \sim Normal(\mu = 0,\ \sigma = 0.5)$$

$$\sigma_y\ , \sigma_{T_{\text{SMS}}} \sim Gamma(\alpha = 2,\ \beta = 0.5)$$

$$\Omega \sim LKJ_2(\eta = 1),$$

where (3b) and (3a) are analogous to the first- and second-stage in two-stage least-squares frameworks, $T_{\text{Day-ahead Email},i}$ is the binary indicator for treatment one (i.e., day-ahead "heads-up" email), $T_{\text{SMS},i}$ is the binary indicator for treatment two *for OE customers who received it*, and $Z_{\text{SMS},i}$ is our random instrument/encouragement (i.e., random assignment to the SMS-plus-bonus condition).

Additionally, X is a $N \times p$ matrix containing p pre-treatment covariates and/or confounders x believed to jointly determine $Z_{\text{SMS},i}$ and the *response* potential outcomes (y^0, y^1) or $Z_{\text{SMS},i}$ and the *treatment* potential outcomes (T^0, T^1) for study units $i \in N$. Furthermore, $\vec{\beta}$ and $\vec{\gamma}$ are p-length vectors of coefficients relating the pre-treatment covariates/confounders to y and T, respectively. Moreover, σ_y and σ_T are the standard deviations of the *equation-specific* error terms $\epsilon_{i,y}$ and $\epsilon_{i,T}$ — where ρ is the correlation between the errors terms.

Thus, β_1 is the ATE of the day-ahead condition, β_2 is the CACE of the SMS-plus-bonus condition, and ρ captures the expected amount and direction of "endogeneity" (i.e., the correlation between unexplained aspects of Y and unexplained aspects of T) — although this is expected to be a lower, not upper, estimate (Lopes and Polson 2014:102). In the *Supplementary Information*, we discuss our choice of prior for each parameter in Equation (3).

Randomised Encouragement Design (RED): Results (Main)

Figure 10 depicts the posterior mean for three causal effects from models of total consumption (kWh) during, and actual participation in, the Mar. 15 Saving Session. The first and second causal effect come from our simultaneous-equation models. And they are, respectively, the ATE of receiving a day-ahead heads-up (vs. an intraday notice only) and the CACE of receiving an intraday SMS reminder plus being made bonus-eligible (vs. an intraday notice only). The third quantity comes from a single-equation, "reduced form" model (see *Supplementary Information*) and it is the intent-to-treat effect (ITT) for the SMS-plus-bonus condition — i.e., the effect of merely being made eligible to receive treatment (Gelman et al. 2020:426). Posterior means are accompanied by 95% HDIs. For both outcomes, posterior means and HDIs are obtained using two model specifications. As in Study 1, these specifications differ based on inclusion of potential confounders (i.e., "Baseline" vs. "Extended"). Supplementary Figure 16 depicts the distribution of posterior values for each causal effect in relation to our Region of Practical Equivalence (ROPE), the Probability of Direction (POD), the 95% HDI, and the posterior mode — i.e., *the* most-likely value for a parameter under the fitted model. 95% HDIs in Figure 10 are tiny due to posterior standard deviations being uniformly small.

The posterior mean ATE (i.e., $\hat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]) from our simplest model of consumption ("Baseline") indicates that OE customers sent a supplementary day-ahead email used, on average, 0.021 kWh *less* (95% HDI = $[-0.032, -0.01]$) during the Saving Session compared to OE customers only sent an intraday notice. Consumption is estimated to be 0.655 kWh (95% HDI = $[0.653, 0.657]$), on average, amongst customers who only received day-of notice. Thus the posterior mean ATE represents a 3.2% (95% HDI = $[-5, -1.5]$) reduction in consumption over baseline (Figure 11, Figure 12). However, the posterior mean ATE is attenuated to -0.013 (95% HDI = $[-0.022, -0.003]$) in the presence of pre-treatment covariates ("Extended" model) — where it virtually disappears when we explicitly model heteroscedasticity (i.e., "Flexibly Scale" model; PMEAN = -0.006; 95% HDI = $[-0.013, 0.001]$) Moreover, despite the ATE's POD for being less than zero being 100% under our simplest model (Supplementary Figure 16), $\approx$ 27% of its posterior distribution includes our ROPE when adjusting for pre-treatment variables (Supplementary Figure 16) and 85% of its posterior includes our ROPE when we explicitly model heteroscedasticity (Supplementary Figure 17).

Similar to the LATE in our models of consumption in Study 1, the ATE's POD for being less than zero is $\approx$ 95-100% regardless of covariates (Supplementary Figures 16 and 17). Thus, under our models, there is near certainty

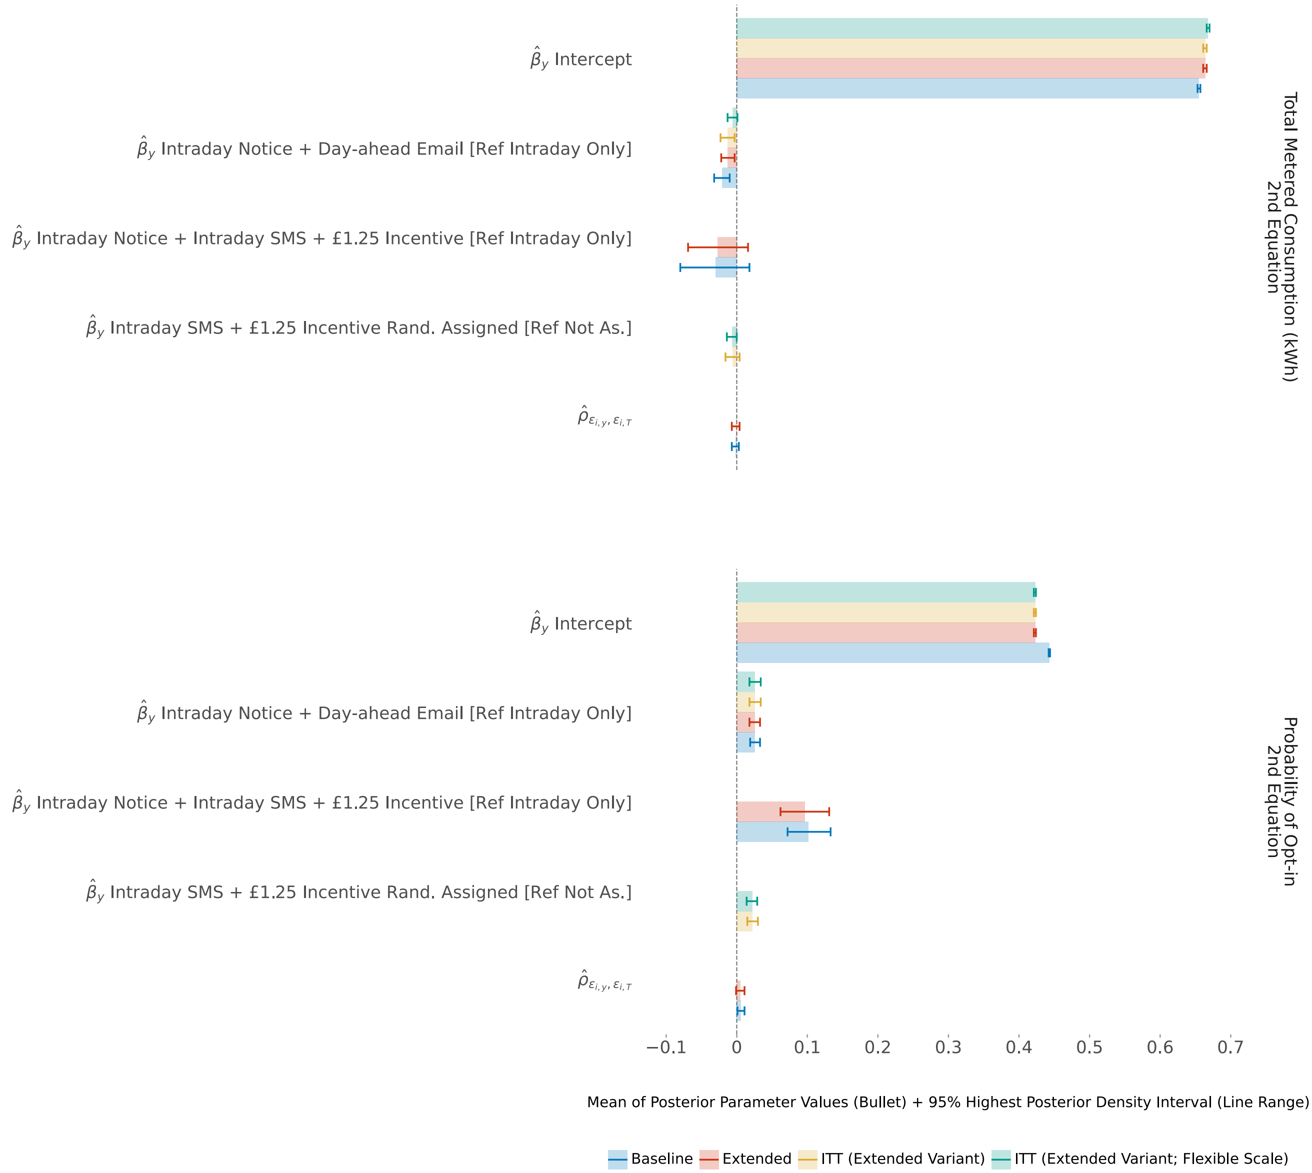

Figure 10: "Forrest plot" depicting posterior means for the ATE ($\hat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]), the CACE ($\hat{\beta}_y$ Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]), the intent-to-treat effect or "ITT" ($\hat{\beta}_y$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]), and the expected average outcome in the control group ($\hat{\beta}_y$ Intercept), holding all other covariates constant. Outcome variable y given on right-hand side of plot. Results obtained either from the second part of a simultaneous-equation regression model (i.e., the ATE and CACE) or a single-equation regression model (i.e., the ITT) fit to data from the randomised controlled trial for the 12th Saving Session (Mar. 15, 2023). Four model specifications are used. The first ("Baseline") only adjusts for our binary treatment and encouragement (Supplementary Tables 15, and 19). The second ("Extended") adds to the first relevant pre-treatment variables such as a customer's historical energy usage, their region of residence, and the degree to which their postcode is deprived (Supplementary Tables 16, and 20). The third ("ITT (Extended Variant)") is a version of the second wherein we replace the binary indicator for actually receiving the SMS-plus-bonus treatment with a binary indicator for randomised assignment to the SMS-plus-bonus condition. The fourth specification is a variant of the third wherein we account for heteroscedasticity using a linear predictor for the scale of the outcome σ (see parameters ϕ in Supplementary Table 18, and 22; see also *Supplementary Information*). The number of observation used to fit each of our models appear in Table 1.

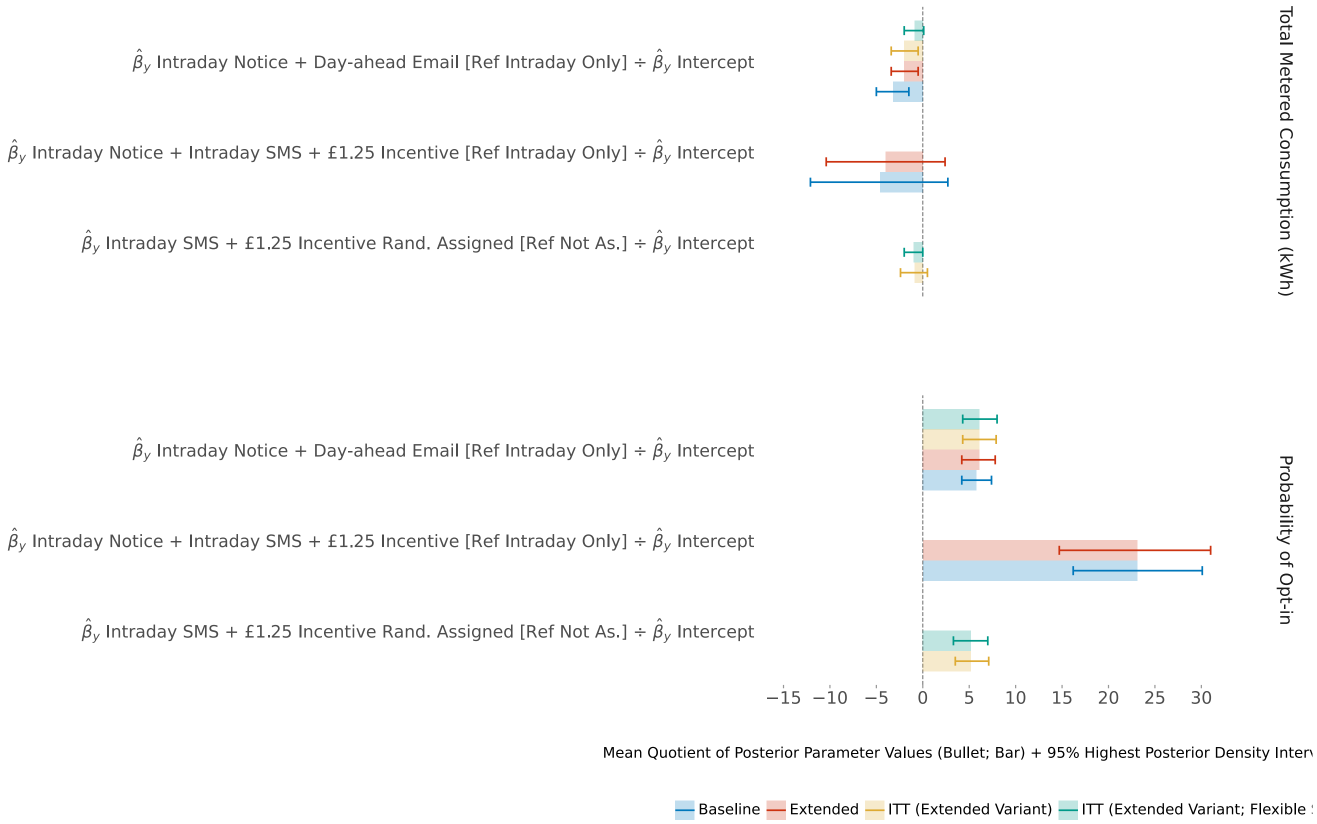

Figure 11: Percentage change in outcome y over the baseline for the control group ($\widehat{\beta}_y$ Intercept) for the binary treatment given the ATE ($\widehat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]), the CACE ($\widehat{\beta}_y$ Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]), or the ITT ($\widehat{\beta}_y$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]), holding all other covariates constant. Outcome variable y given on right-hand side of plot. Posterior mean parameter values appear in Figure 10. However, division is performed using *posterior-sample-specific values* for each parameter. This yields a distribution of percentage changes which we summarise using 95% HDIs.

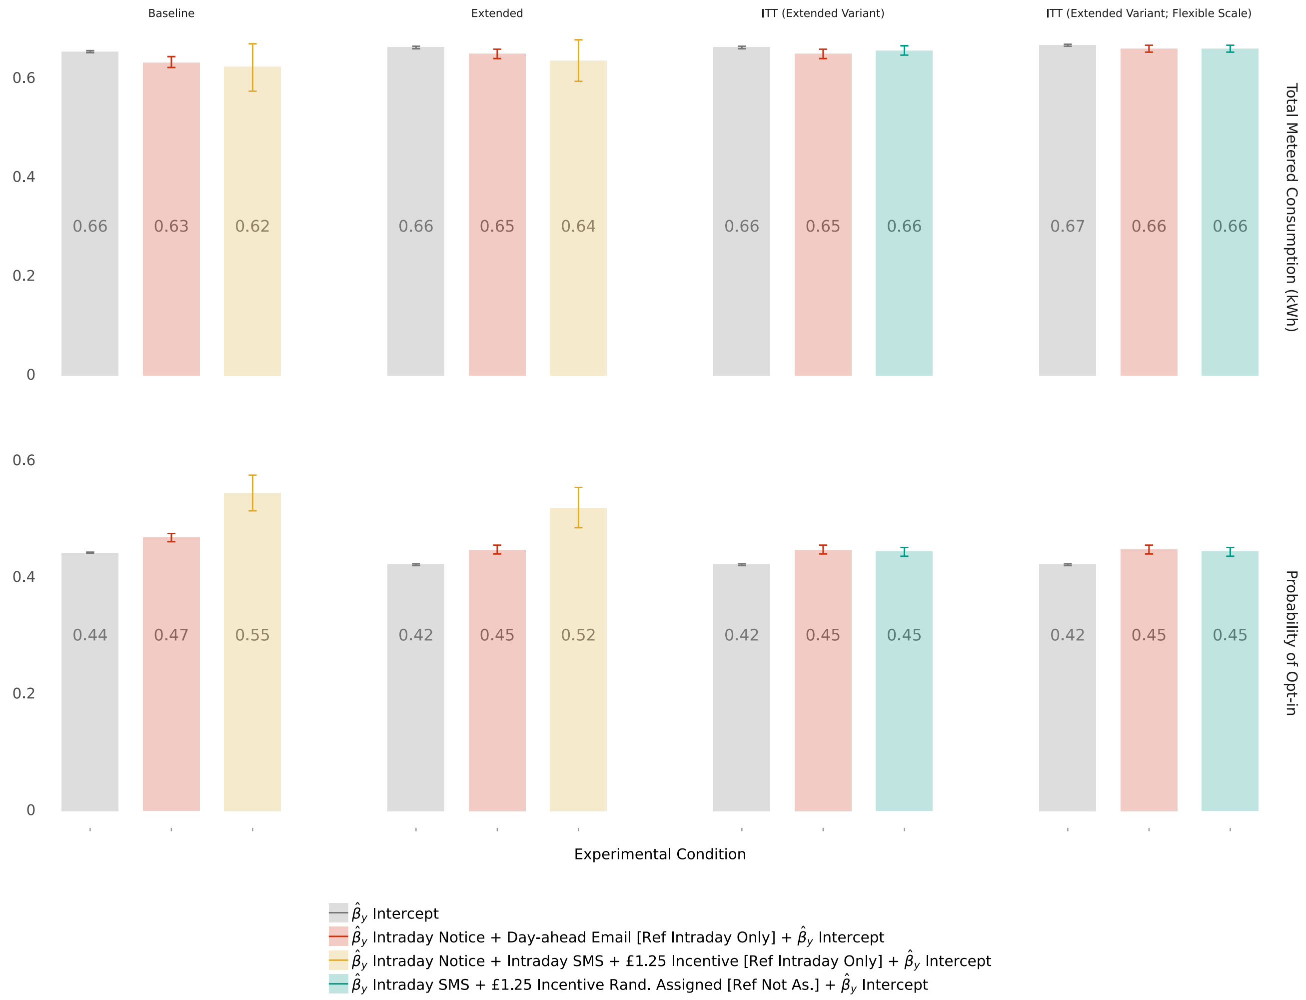

Figure 12: Expected value of the outcome y for the control and treatment conditions, holding all other covariates constant. Outcome variable y given on right-hand side of plot. Posterior mean parameter values appear in Figure 10. However, summation for the experimental conditions is performed using *posterior-sample-specific values* for each parameter. This yields a distribution of expected outcomes which we summarise using 95% HDIs.

that the ATE "exists" and that it is negative (i.e., decreased consumption under earlier notice). Accordingly, there is uncertainty around the ATE's "significance". And, as with Study 1, the clear POD leads us to conclude that there is compelling evidence to suggest that being sent supplementary day-ahead notice had a negative causal impact on energy consumption during the Mar. 15 Saving Session with the proviso that there is a trivial possibility that the ATE is so tiny as to be practically unimportant. Put simply, early email-based notice appears to have mattered in this particular scenario (RQ1), albeit possibly to a very small degree.

Turning to our treatment with non-compliance, the first equation in our Simple and Expanded model of consumption captures the expected association between our instrument and our SMS-based treatment. In the present scenario, this is the expected proportion of compliers ($\widehat{\beta}_y$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]) under the simple model (PMEAN = 0.233; 95% HDI = $[0.232, 0.234]$; Supplementary Table 15) and the expanded model (PMEAN = 0.225; 95% HDI = $[0.224, 0.227]$; Supplementary Table 16).

Nevertheless, evidence in relation to the CACE in both our simple and expanded models is inconclusive. That is to say, there is no compelling evidence to suggest that being sent an intraday SMS reminder whilst being made eligible for the bonus price incentive had a causal impact on consumption during the Mar. 15 Saving Session in either direction. Specifically, the posterior mean CACEs (i.e., $\widehat{\beta}_y$ Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]) in our Baseline and Extended models are, respectively -0.030 (95% HDI = $[-0.08, 0.018]$) and -0.027 (95% HDI = $[-0.069, 0.016]$). However, under both models, the CACE has a non-trivial probability of having the opposite sign (POD greater than zero $\approx$ 11%; Supplementary Figure 16). And, depending on model, $\approx$ 16-18% of the posterior for the CACE, alongside a portion of 95% of its most-probable values, include our ROPE (Supplementary Figure 16).

To the extent that $\widehat{\rho}_{\epsilon_{i,y},\epsilon_{i,T}}$ — i.e., the estimated correlation between errors for the first and second equations in our simultaneous-equation models — captures endogeneity (Lopes and Polson 2014), we do not expect the posterior mean CACE to be biased. Indeed $\widehat{\rho}_{\epsilon_{i,y},\epsilon_{i,T}}$ is minuscule in the Baseline model (PMEAN = -0.002; 95% HDI = $[-0.007, 0.003]$) and in the Extended model (PMEAN = -0.001; 95% HDI = $[-0.007, 0.004]$). We also stress our use of an instrument that was randomised alongside our adjustment for a wide-array of potential sources of residual imbalance.

All in all, the most conservative conclusion based on behaviour of the posterior for the CACE is that it is unclear whether and how the timing of supplementary SMS-based notice, alongside the availability of the bonus price incentive, mattered for energy consumption during the Mar. 15 Saving Session for the particular OE customers in our sample (RQ1). Evidence around ITT effect on consumption support this conclusion (Supplementary Figure 17).

As for participation, our Baseline and Extended model (both linear probability sub-models) unambiguously indicate that our treatments positively impacted the likelihood of opting into the Mar. 15 Saving Session (RQ2). Specifically, the posterior mean ATE for the day-ahead-heads-up-email condition from our simplest model (PMEAN = 0.026; 95% HDI = $[0.019, 0.033]$) indicates a roughly 6% increase in the the probability of opt-in over the control group for the event (95% HDI = $[4.2, 7.8]$, Figure 11), the latter of whom are estimated as opting-in with a probability of 0.443, on average (95% HDI = $[0.442, 0.444]$, Figure 12). Furthermore, under all of our model specifications, the ATE's POD for being greater than zero is 100% and the ATE's posterior distribution comfortably clears our ROPE (Supplementary Figures 16, and 17). Thus, there is compelling evidence for us to conclude that the supplementary, email-based heads-up caused and increase in the probability of participation.

Further still, the posterior mean CACE for the SMS-plus-bonus condition is 0.102 (95% HDI = $[0.072, 0.133]$) in our Baseline model of the probability of formally agreeing to participate in the Mar. 15 Saving Session. The POD for being greater than zero is 100% for the CACE. And the posterior for the CACE comfortably clears our ROPE in both our simple and expanded models (Supplementary Figure 16). Thus, there is also compelling evidence for us to conclude that being sent a supplementary intraday SMS-based reminder whilst being made eligible for an additional price incentive caused and increase in the probability of opt-in by $\approx$ 23% (95% HDI = $[14.7, 31]$, Figure 11) over the probability of event participation in the control group.

Results related to the CACE for opt-in are consistent with our models for the ITT for the SMS-plus-bonus condition (Supplementary Figure 17). The latter models also suggest that our results for the ITT/CACE are not vulnerable to heteroscedasticity. Accordingly, the broad timing and the general type of supplementary notice appears to have mattered for participation in the Mar. 15 Saving Session for the particular OE customers in our sample (RQ2).

Randomised Encouragement Design (RED): Results (Exploratory Modelling)

Last, we briefly consider treatment-effect heterogeneity in relation to customers' attributes using a set of ancillary, exploratory models. Given the large scale of our data, we narrowly focus on geography as this is a plausible source of variation in consumption-related behaviour in absence of a clear scientific theory of heterogeneity. Thus, ancillary models concern Great Britain's 14 Distribution Network Operator (DNO) regions — i.e., geographically-defined areas under the control of different commercial firms that are used to administer electricity from the National Grid to British households. Specifically, ancillary models are used to estimate region-specific ATEs for the day-ahead-email condition alongside region-specific ITTs for mere eligibility for the SMS-plus-bonus condition. Recall that the ITT is the effect of eligibility itself. And it is useful in scenarios wherein one expects a gap between the number of individuals targeted for some treatment and the number of individuals actually receiving the treatment. Indeed, there are likely to be challenges around compliance with SMS-based treatments in real-world scenarios such that policymakers and practitioners will benefit from knowing the effect of *possibly* receiving treatment.

Following Blackwell and Olson's (2021) warnings around omitted-interaction-effect bias, we eschew multiplicative interactions between region and the ATE and between region and the ITT to instead subset our data by DNO region and then use these subsets to fit 14 separate models. This is equivalent to including multiplicative interactions between DNO region and *all* other parameters in our model (Blackwell and Olson 2021), a practice we avoid here due to concerns around possible model instability. Mean parameter values appear in Figure 13 and mean percentage changes over baseline appear in Figure 14.

There is clear heterogeneity in posterior mean ATEs and posterior mean ITTs conditional on DNO region. In line with our pooled models (Figure 10), evidence of causal effects is generally more compelling in our region-specific models of the probability of event participation as opposed to our region-specific models of consumption, where the largest participation effects appear to be for Wales and the South of England. Still, 95% HDI are generally wide and span zero such that most regions have plausible values for their ATE and ITT that are similar to those in other regions, save for the ITT in South Wales and Merseyside and North Wales.[14]

[14] Note, we drop P376 Baseline Consumption from our ancillary model for the probability of opt in in the DNO region "South Eastern" as its inclusion led to divergent transitions (i.e., ill-behaved chains) during posterior sampling.

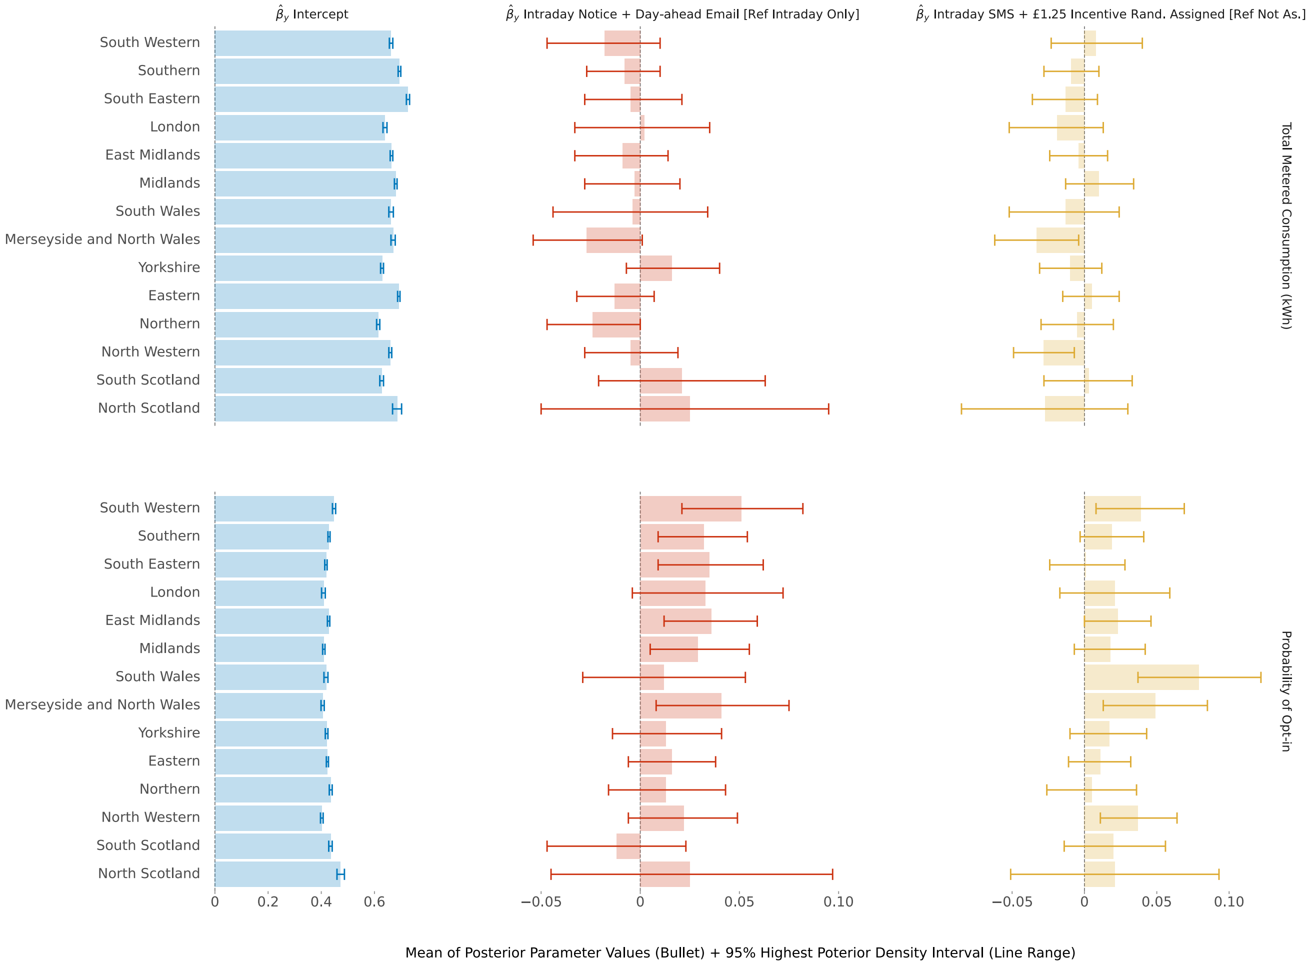

Figure 13: "Forrest plot" depicting posterior means for the ATE ($\hat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]), the ITT ($\hat{\beta}_y$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]), and the expected average outcome in the control group ($\hat{\beta}_y$ Intercept) in each of Great Britain's 14 DNO regions, holding all other covariates constant. Result obtained using single-equation regression models fit to data specific to each DNO region from the randomised controlled trial for the 12th Saving Session (Mar. 15, 2023). Region-specific models fit using our "Extended (Flexible Scale)" specification and thus we adjust for relevant pre-treatment variables such as a customer's historical energy usage and the degree to which their postcode is deprived whilst including a linear predictor for the scale of the outcome σ.

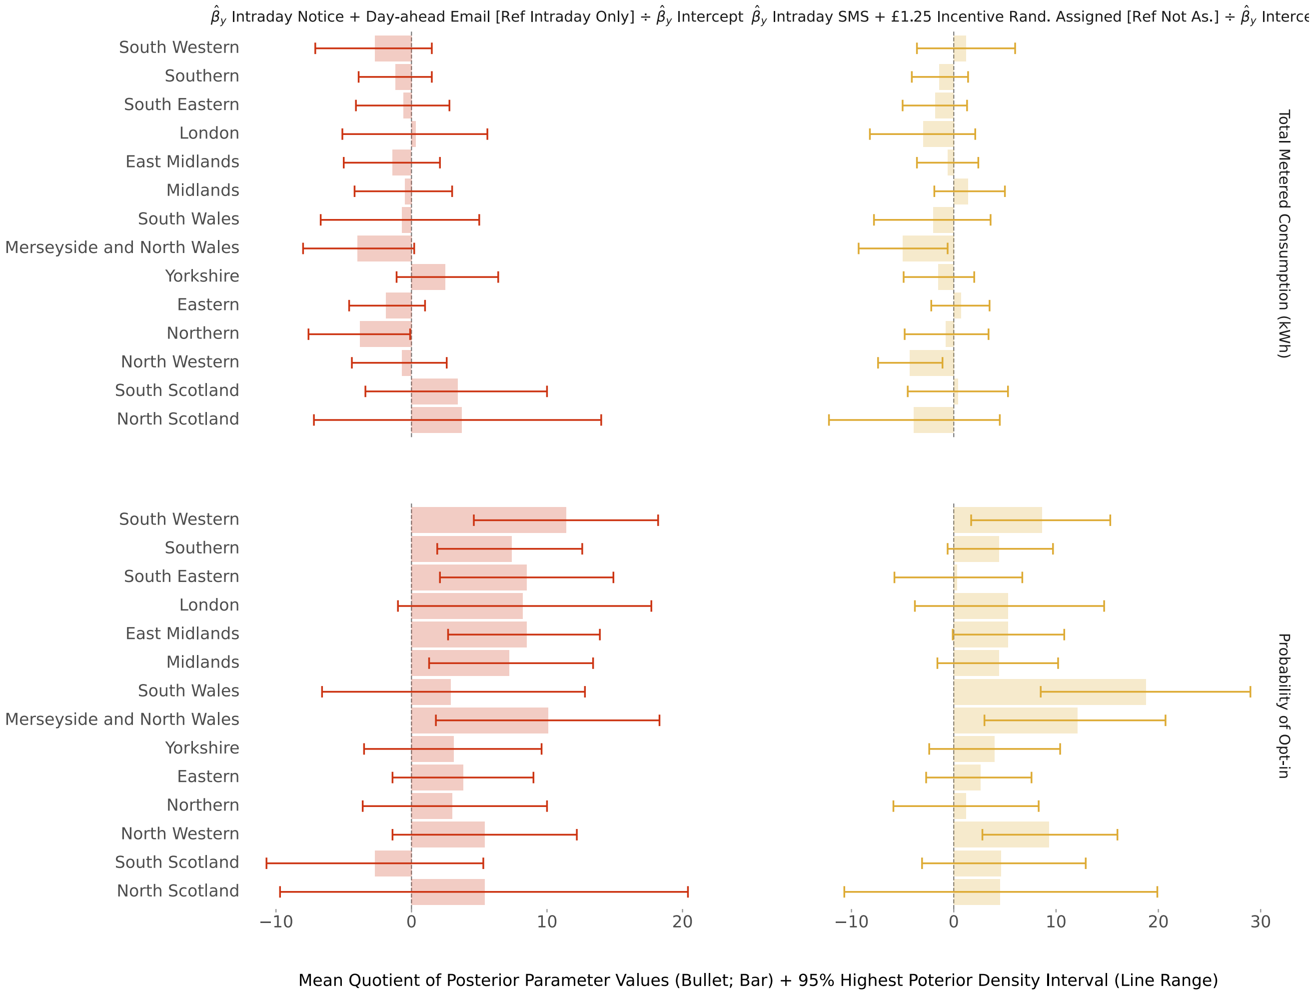

Figure 14: Percentage change in outcome y over the baseline for the control group (i.e., $\hat{\beta}_y$ Intercept) for the binary treatment given the ATE ($\hat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]) or the ITT ($\hat{\beta}_y$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]), holding all other covariates constant. Posterior mean parameter values appear in Figure 13. However, quotients are obtained using *posterior-sample-specific values* for each parameter. This yields a distribution of percentage-changes which we summarise using 95% HDIs.

Discussion

Changing the behaviour of private individuals is a vital component of the transformation of energy systems. And here we have investigated the potential utility of energy retailers' modulation of their appeals to consumers to engage in flexible domestic energy use. Specifically, we have presented two studies — a quasi-experiment and a randomised controlled trial — wherein we analysed data on total consumption (kWh) during — and formal agreement to participate in — two 60-minute energy-savings events in 2023 on the part of 666,441 customers of Octopus Energy alongside data on the time horizon (course-grained) and channel (general) of OE's requests to its customers to consume less electricity.

Our RCT provides clear evidence of a positive, causal association between customers' formal agreement to participate in the second 60-min event and OE's *supplementary* notice about that event — where day-of, SMS-based notice, alongside a bonus price incentive, was more effective than day-ahead email-based notice (Figure 11). More formally, we find that, conditional on pre-treatment variables (e.g., historical energy usage), being sent a supplementary "heads-up" email the day prior or a day-of "reminder" SMS text message respectively increased the probability of event participation by $\approx 6\%$ (95% $\text{HDI}_{\text{Extended Model}} = [4.2\%, 7.8\%]$) and $\approx 23\%$ (95% $\text{HDI}_{\text{Extended Model}} = [14.7\%, 31\%]$) over baseline (i.e., day-of primary notice [email and/or SMS]) — where we observe an increase of $\approx 5.2\%$ (95% $\text{HDI}_{\text{Extended (Flexible Scale) Model}} = [3.3\%, 7\%]$) in the probability of participation when considering mere eligibility for the SMS-based treatment (Figure 11). Nevertheless, we find no compelling evidence to suggest that ancillary contact via an intraday SMS text led to increased or decreased turn down at the designated time.

As for email and consumption, our quasi-experiment using a donut-hole RDD clearly indicates that being sent an intraday notice had a positive causal impact on energy consumption during the Feb. 13 Saving Session. Recall that our models place noteworthy probability density on particularly small values for this causal effect. Still, on a broad scale (e.g., aggregation across whole geographic regions) and conditional on pre-treatment variables, the posterior mean causal effect of short notice in Study 1 ($\text{PMEAN}_{\text{Extended (Flexible Scale) Model}} = 0.040$, 95% $\text{HDI}_{\text{Extended (Flexible Scale) Model}} = [0.004, 0.075]$) is considerable as it represents an increased consumption of $\approx 6.7\%$ (95% HDI $= [0.7\%, 12.9\%]$). Indeed, the general direction of our effects from Study 1 and Study 2 are consistent in that they indicate longer time horizons for email-based notice perhaps reduced energy consumption amongst the particular OE customers in our samples (see the *positive* posterior values for the LATE $\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead Notice] in Study 1 Figure 4 and the *negative* posterior values for the ATE $\hat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only] in Study 2 Figure 10). Yet, despite their PODs being near 100%, we note again that there is some question as to whether any of these effects are, in practice, not zero given our ROPE (Supplementary Figures 15, 16, 17) — especially for our RCT which only indicates a $\approx 1\%$ turndown (95% HDI $= [-2\%, 0.1\%]$) under early, supplementary email-based notice after we account for heteroscedasticity.

Practically speaking, what then do our results mean for institutional actors (e.g., energy suppliers, transmission system operators, governments) and applied researchers hoping to change electricity-consumption-related behaviours on the part of private individuals? Our most unambiguous results concern engagement with energy-saving events. And our second study strongly suggests the providing domestic consumers with supplementary notice about these events over and above primary messaging can play an important part in converting mere awareness of the need to conserve electricity into agreement to *try* to do so. Note well, however, that the design of our RCT prevents us from partitioning the CACE (or the ITT) for the second treatment condition into the effect of being sent an SMS reminder versus the effect of the bonus price incentive. Thus, the potential for more money — and/or the greater accessibility afforded through use of SMS as a communication channel — could undergird this particular result. Still, even if the expected sizeable impact of the second treatment is primarily driven by the price incentive, our two-equation regression models also indicate that day-ahead supplementary notice via email also causally increased participation.

Given all of this, we ultimately conclude that the broad timing and the general type of appeals matter in the sense that they stand to shape engagement with demand-flexibility initiatives (RQ2) yet may still fail to drive large change in levels of consumption (RQ1) — at least for the individuals in our sample.

When making sense of this conclusion, there are two limitations to our research that practitioners and applied researchers ought to keep in mind.

First, as we allude to above, our mix of experimental conditions is awkward. That is to say, we do not fit models using all possible combinations of distinct treatments with respect to timing, channel, and price incentive (e.g., intraday opt-in notice + day-ahead "heads-up" email vs. intraday opt-in notice + intraday "reminder" SMS text vs. intraday opt-in notice + £1.25 bonus vs. intraday opt-in notice + intraday "reminder" SMS text + £1.25 bonus, etc.). Unfortunately, we were restricted in our ability to assign an expansive set of treatments owing to limited budget and difficulties related to orchestrating an experiment involving customers at a large commercial organisation. Thus, we can not estimate the direct and synergistic effects of timing, channel, and price incentive, making aspects of our findings difficult to interpret. Along this line, and given our results, future research should make a special effort to measure the impact of notice-based treatments using more-granular time horizons (e.g., 3, 6, 12, 24, and 48 hours ahead of an energy-savings event). *Mutatis mutandis* for varying price incentives.

Finally, save for a series of ancillary models focused on geographic region, we have narrow focused on various flavours of unconditional treatment effects (i.e., the ATE, LATE, and CACE). However, other conditional average treatment effects (e.g., those related to variation in socio-economic status) may also be relevant. Broadly speaking, conditional ATEs may be recovered in a regression framework using multiplicative interactions (see Gelman et al. (2020)). Nevertheless, we generally opt for unconditional models owing to a lack of clear scientific theory and strong domain priors (c.f., DNO region) around mechanisms for treatment-effect heterogeneity — a recipe for bad research practice and poor reproducibility. We also eschew multiplicative interactions owing to the non-trivial changes to model interpretation induced by their inclusion alongside increased potential for model instability — particularly when incorporating interactions in the most-rigorous manner. Accordingly, future work should explore how best to sensibly investigate treatment-effect heterogeneity in a manner sensitive to the warnings of both Blackwell and Olson (2021) and Brambor et al. (2006) around multiplicative interactions.

References

Barreca, A. I., M. Guldi, J. M. Lindo, and G. R. Waddell. 2011. "Saving Babies? Revisiting the Effect of Very Low Birth Weight Classification." *The Quarterly Journal of Economics* 126(4):2117–23. doi: 10.1093/qje/qjr042.

Barreca, Alan I., Jason M. Lindo, and Glen R. Waddell. 2016. "Heaping-Induced Bias in Regression-Discontinuity Designs." *Economic Inquiry* 54(1):268–93. doi: 10.1111/ecin.12225.

Bergquist, Magnus, Maximilian Thiel, Matthew H. Goldberg, and Sander Van Der Linden. 2023. "Field Interventions for Climate Change Mitigation Behaviors: A Ssecond-Order Meta-Analysis." *Proceedings of the National Academy of Sciences* 120(13):e2214851120. doi: 10.1073/pnas.2214851120.

Blackwell, Matthew, and Michael P. Olson. 2021. "Reducing Model Misspecification and Bias in the Estimation of Interactions." *Political Analysis* 1–20. doi: 10.1017/pan.2021.19.

Bogdanov, Dmitrii, Manish Ram, Arman Aghahosseini, Ashish Gulagi, Ayobami Solomon Oyewo, Michael Child, Upeksha Caldera, Kristina Sadovskaia, Javier Farfan, Larissa De Souza Noel Simas Barbosa, Mahdi Fasihi, Siavash Khalili, Thure Traber, and Christian Breyer. 2021. "Low-Cost Renewable Electricity as the Key Driver of the Global EnergyTtransition Towards Sustainability." *Energy* 227:120467. doi: 10.1016/j.energy.2021.120467.

Brambor, Thomas, William Roberts Clark, and Matt Golder. 2006. "Understanding Interaction Models: Improving Empirical Analyses." *Political Analysis* 14(1):63–82. doi: 10.1093/pan/mpi014.

Cattaneo, Matias D., Nicolás Idrobo, and Rocío Titiunik. Forthcoming. *A Practical Introduction to Regression Discontinuity Designs: Extensions*. Cambridge University Press.

Fankhauser, Sam, Stephen M. Smith, Myles Allen, Kaya Axelsson, Thomas Hale, Cameron Hepburn, J. Michael Kendall, Radhika Khosla, Javier Lezaun, Eli Mitchell-Larson, Michael Obersteiner, Lavanya Rajamani, Rosalind Rickaby, Nathalie Seddon, and Thom Wetzer. 2022. "The Meaning of Net Zero and How to Get It Right." *Nature Climate Change* 12(1):15–21. doi: 10.1038/s41558-021-01245-w.

Gelman, Andrew, Jennifer Hill, and Aki Vehtari. 2020. *Regression and Other Stories*. 1st ed. Cambridge University Press.

Goldberg, Matthew H., Abel Gustafson, Seth A. Rosenthal, and Anthony Leiserowitz. 2021. "Shifting Republican Views on Climate Change Through Targeted Advertising." *Nature Climate Change* 11(7):573–77. doi: 10.1038/s41558-021-01070-1.

Greene, William. 2019. *Econometric Analysis*. 8th Edition. Pearson.

Gustafson, Abel, Matthew H. Goldberg, Parrish Bergquist, Karine Lacroix, Seth A. Rosenthal, and Anthony Leiserowitz. 2022. "The Durable, Bipartisan Effects of Emphasizing the Cost Savings of Renewable Energy." *Nature Energy* 7(11):1023–30. doi: 10.1038/s41560-022-01099-2.

Huntington-Klein, Nick. 2021. *The Effect: An Introduction to Research Design and Causality*. Boca Raton: CRC Press, Taylor & Francis Group.

Jost, John T., Delia S. Baldassarri, and James N. Druckman. 2022. "Cognitive–Motivational Mechanisms of Political Polarization in Social-Communicative Contexts." *Nature Reviews Psychology* 1(10):560–76. doi: 10.1038/s44159-022-00093-5.

Kleibergen, Frank, and Eric Zivot. 2003. "Bayesian and Classical Approaches to Instrumental Variable Regression." *Journal of Econometrics* 114(1):29–72. doi: 10.1016/S0304-4076(02)00219-1.

Kruschke, John K. 2015. *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan*. Edition 2. Boston: Academic Press.

Kruschke, John K. 2018. "Rejecting or Accepting Parameter Values in Bayesian Estimation." *Advances in Methods and Practices in Psychological Science* 1(2):270–80. doi: https://psycnet.apa.org/doi/10.1177/2515245918771304.

Lopes, Hedibert F., and Nicholas G. Polson. 2014. "Bayesian Instrumental Variables: Priors and Likelihoods." *Econometric Reviews* 33(1-4):100–121. doi: 10.1080/07474938.2013.807146.

Makowski, Dominique, Mattan S. Ben-Shachar, S. H. Annabel Chen, and Daniel Lüdecke. 2019. "Indices of Effect Existence and Significance in the Bayesian Framework." *Frontiers in Psychology* 10:2767. doi: 10.3389/fpsyg.2019.02767.

McElreath, Richard. 2020. *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. 2nd ed. Boca Raton: CRC Press, Taylor & Francis Group.

Przepiorka, Wojtek, and Christine Horne. 2020. "How Can Consumer Trust in Energy Utilities Be Increased? The Effectiveness of Prosocial, Proenvironmental, and Service-Oriented Investments as Signals of Trustworthiness." *Organization & Environment* 33(2):262–84. doi: 10.1177/1086026618803729.

Rossi, Peter E., Greg M. Allenby, and Robert E. McCulloch. 2005. *Bayesian Statistics and Marketing*. Hoboken, NJ: Wiley.

Stephenson, Janet, Barry Barton, Gerry Carrington, Daniel Gnoth, Rob Lawson, and Paul Thorsnes. 2010. "Energy Cultures: A Framework for Understanding Energy Behaviours." *Energy Policy* 38(10):6120–29. doi: 10.1016/j.enpol.2010.05.069.

Vehtari, Aki, Andrew Gelman, Daniel Simpson, Bob Carpenter, and Paul-Christian Bürkner. 2021. "Rank-Normalization, Folding, and Localization: An Improved R^ for Assessing Convergence of MCMC (with Discussion)." *Bayesian Analysis* 16(2). doi: 10.1214/20-BA1221.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data*. 2nd ed. Cambridge, Mass: MIT Press.

Zhang, Baobao, Sander Van Der Linden, Matto Mildenberger, Jennifer R. Marlon, Peter D. Howe, and Anthony Leiserowitz. 2018. "Experimental Effects of Climate Messages Vary Geographically." *Nature Climate Change* 8(5):370–74. doi: 10.1038/s41558-018-0122-0.

Supplementary Information

Adjudication of Evidence

Here we clarify how we marshal evidence to answer RQ1 and RQ2. Null-Hypothesis Significance Testing (NHST) is controversial amongst researchers working within a Bayesian framework as it is thought to promote binary thinking and overconfidence (Kruschke 2018; Kruschke and Liddell 2018; Makowski et al. 2019; McShane et al. 2019). And, in the wake of the generalised reproducibility crisis, NHST is increasingly viewed with scepticism — where NHST is outright avoided by some social, physical, and biomedical scientists (Benjamin et al. 2017; Demidenko 2016; Halsey 2019; Lakens et al. 2018; McShane et al. 2019; van de Schoot et al. 2021). Nevertheless, we wish to judge our results in a manner that allows us to make decisions and pose recommendations for practitioners. Bayesian inference yields *entire distributions of possible values for some parameter of interest* (e.g., the LATE), not just a single point estimate. These "posterior" distributions may be discussed in terms of the probability of particular values of a parameter. This is expected to be more useful for practitioners. Thus, we judge whether there is evidence of a treatment effect by using properties of posterior distributions in line with the existence-significance decision-making framework of Makowski et al. (2019) (see also Kruschke (2018; Kruschke 2021) and Kruschke and Liddell (2018).

As mentioned in the main text, by "existence" we mean "the consistency of an effect in one particular direction (i.e., positive or negative), without any assumptions or conclusions as to its size, importance, relevance, or meaning" (Makowski et al. 2019:9). "Existence" is conceptually distinct from "significance" — the latter of which indicates whether an effect is "worthy of attention" such that its magnitude is not "likely to be too small to be of high importance in real-world scenarios or applications" (Makowski et al. 2019:9)?

Recall that we judge existence of a treatment effect using the Probability of Direction (POD). The POD is the proportion of the posterior distribution on either side of zero (i.e., no effect). And recall that we judge significance of a treatment effect using a whole-posterior or "full" Region of Practical Equivalence (ROPE). Specifically, we report the proportion of the posterior distribution contained within a region of values around zero that are, for practical purposes, *all* considered to be null. Thus the POD and the ROPE yield *continuous metrics of evidence.*

We stress that the ROPE is subjective in the sense that it is chosen by the researcher (Kruschke 2018; Kruschke and Liddell 2018; Makowski et al. 2019). This is no different from the frequentist p-value which has arbitrary decision thresholds (e.g., $\alpha = 0.05$) that only appear objective due to long-standing norms (Benjamin et al. 2017; Demidenko 2016; Halsey 2019; Kruschke 2018; Kruschke and Liddell 2018; Lakens et al. 2018; Makowski et al. 2019; McShane et al. 2019).

For both Study 1 and Study 2, we set the ROPE for our treatment effects on consumption to $[-0.01 \text{ kWh},\ 0.01 \text{ kWh}]$. We do so as we found that median prompted turn down can be as low as 0.06 kWh (e.g., during overnight periods) in an unpublished analysis of requests for flexible energy use across thousands of British households in early 2022.[15] Moreover, small shifts in consumption can be meaningful to energy retailers and transmission system operators when aggregating across hundreds of thousands of consumers and whole geographic regions. Indeed, the ability of utility providers to facilitate even modest turn down — especially closer to real time under short notice — is vital to the transformation of energy systems such that limited energy reduction represents valuable balancing action.

Moreover, for both Study 1 and Study 2, we set the ROPE for our treatment effects on the probability of participation to $[-(\sigma_{y,\text{Opt-in}} \times 0.01),\ (\sigma_{y,\text{Opt-in}} \times 0.01)]$ — i.e., $\pm$ 1% of the observed, unconditional standard deviation of our binary indicator for session participation (i.e., $\sigma_{y,\text{Opt-in}} = \sqrt{\mu_{y,\text{Opt-in}} \times (1 - \mu_{y,\text{Opt-in}})}$). This ROPE is somewhat similar to that which is recommended by Kruschke (2018) for detecting a small effect using linear regression. However, we use a multiplicative factor of 0.01 instead of 0.1 owing to our assumption that even very small participation effects could be meaningful to energy retailers and transmission system operators due to the scale of the number of households under their remit.

Posterior distributions for our causal effects in relation to our ROPEs are given in Supplementary Figures 15, 16, and 17.

[15] See "The Big Dirty Turn Down: Deep Analysis of a Domestic Energy Flexibility Trial" by Centre for Net Zero (2021).

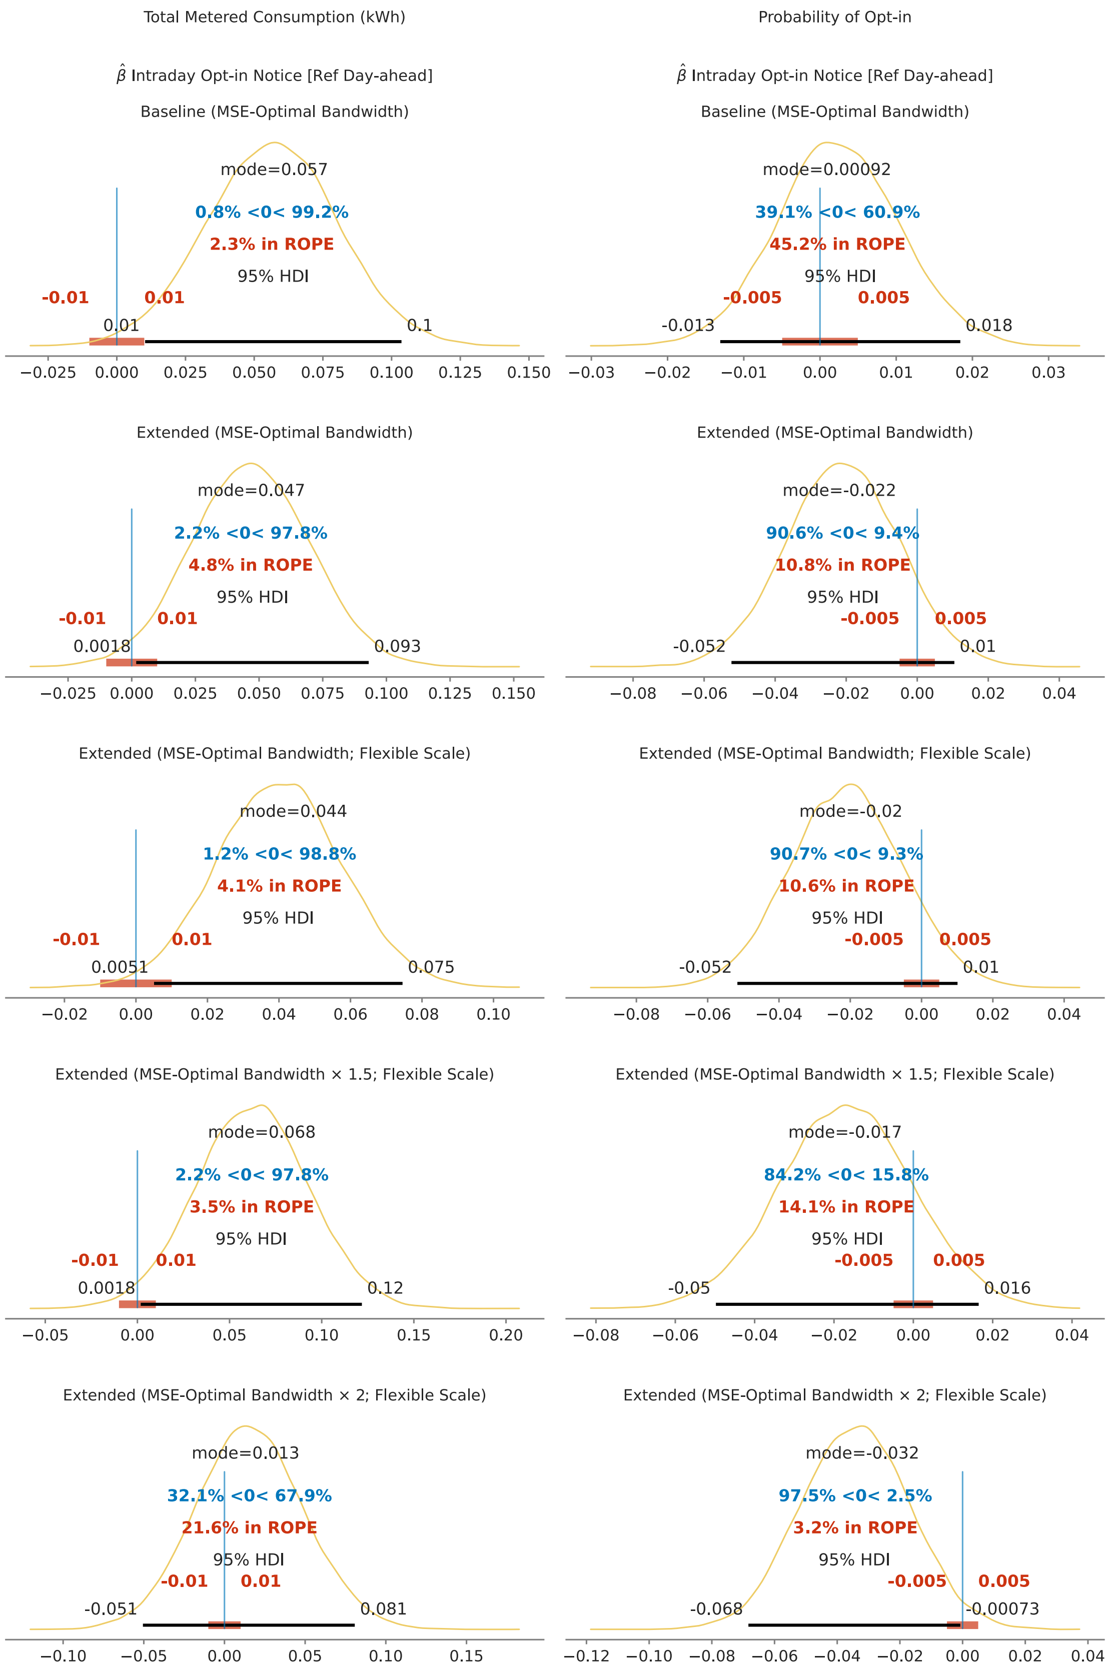

Figure 15: Kruschke-style (2018) posterior density plots for the LATE ($\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead Notice]) from regression discontinuity models fit to data from the 10th Saving Session (Feb. 13, 2023). Outcome variable y given on right-hand side of plot. Probability of Direction (POD; Blue) indicates the proportion of the posterior distribution (Yellow) that is greater than zero and less than zero. The Region of Practical Equivalence (ROPE) is the range of values around zero considered to be null for practical purposes. 95% HDI represented by the black horizontal bar.

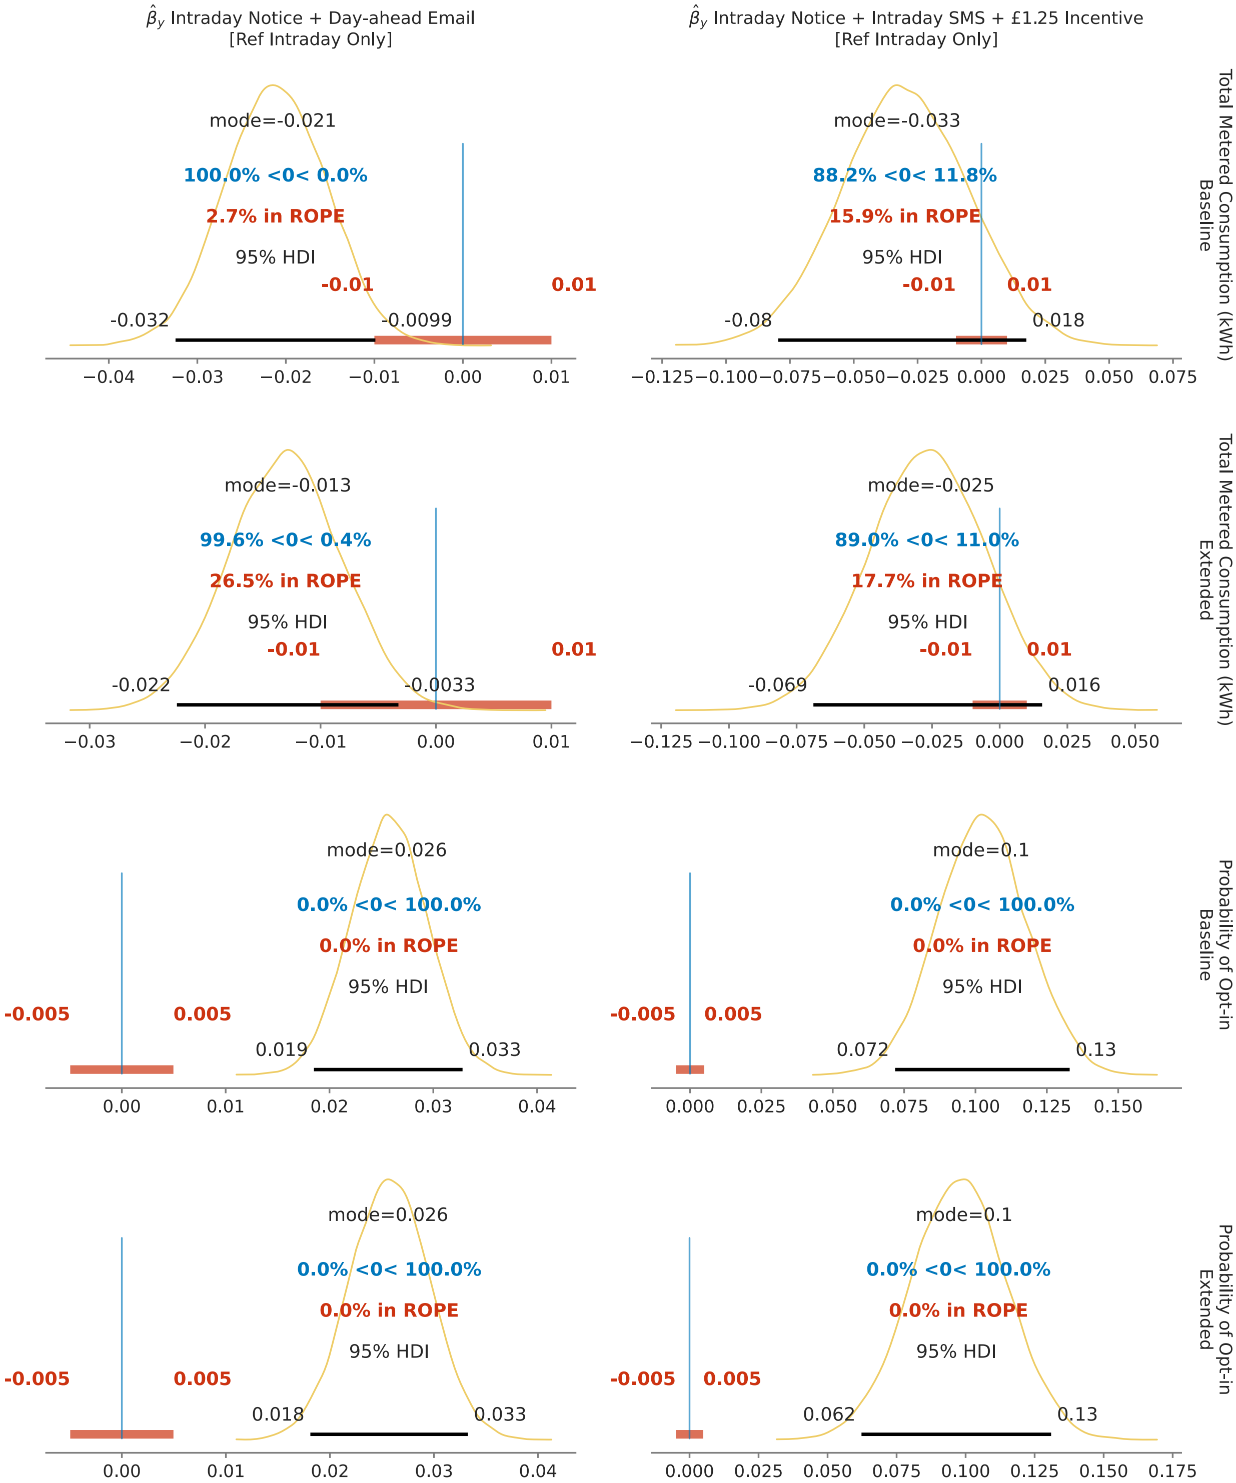

Figure 16: Kruschke-style (2018) posterior density plots for the ATE (left column; $\hat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]) and the CACE (right column; $\hat{\beta}_y$ Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]) from simultaneous-equation regression models fit to data from the randomised controlled trial for the 12th Saving Session (Mar. 15, 2023). Outcome variable y given on right-hand side of plot. Probability of Direction (POD; Blue) indicates the proportion of the posterior distribution (Yellow) greater than zero and less than zero. The Region of Practical Equivalence (ROPE) is the range of values around zero considered to be null for practical purposes. 95% HDI indicated by the black horizontal bar.

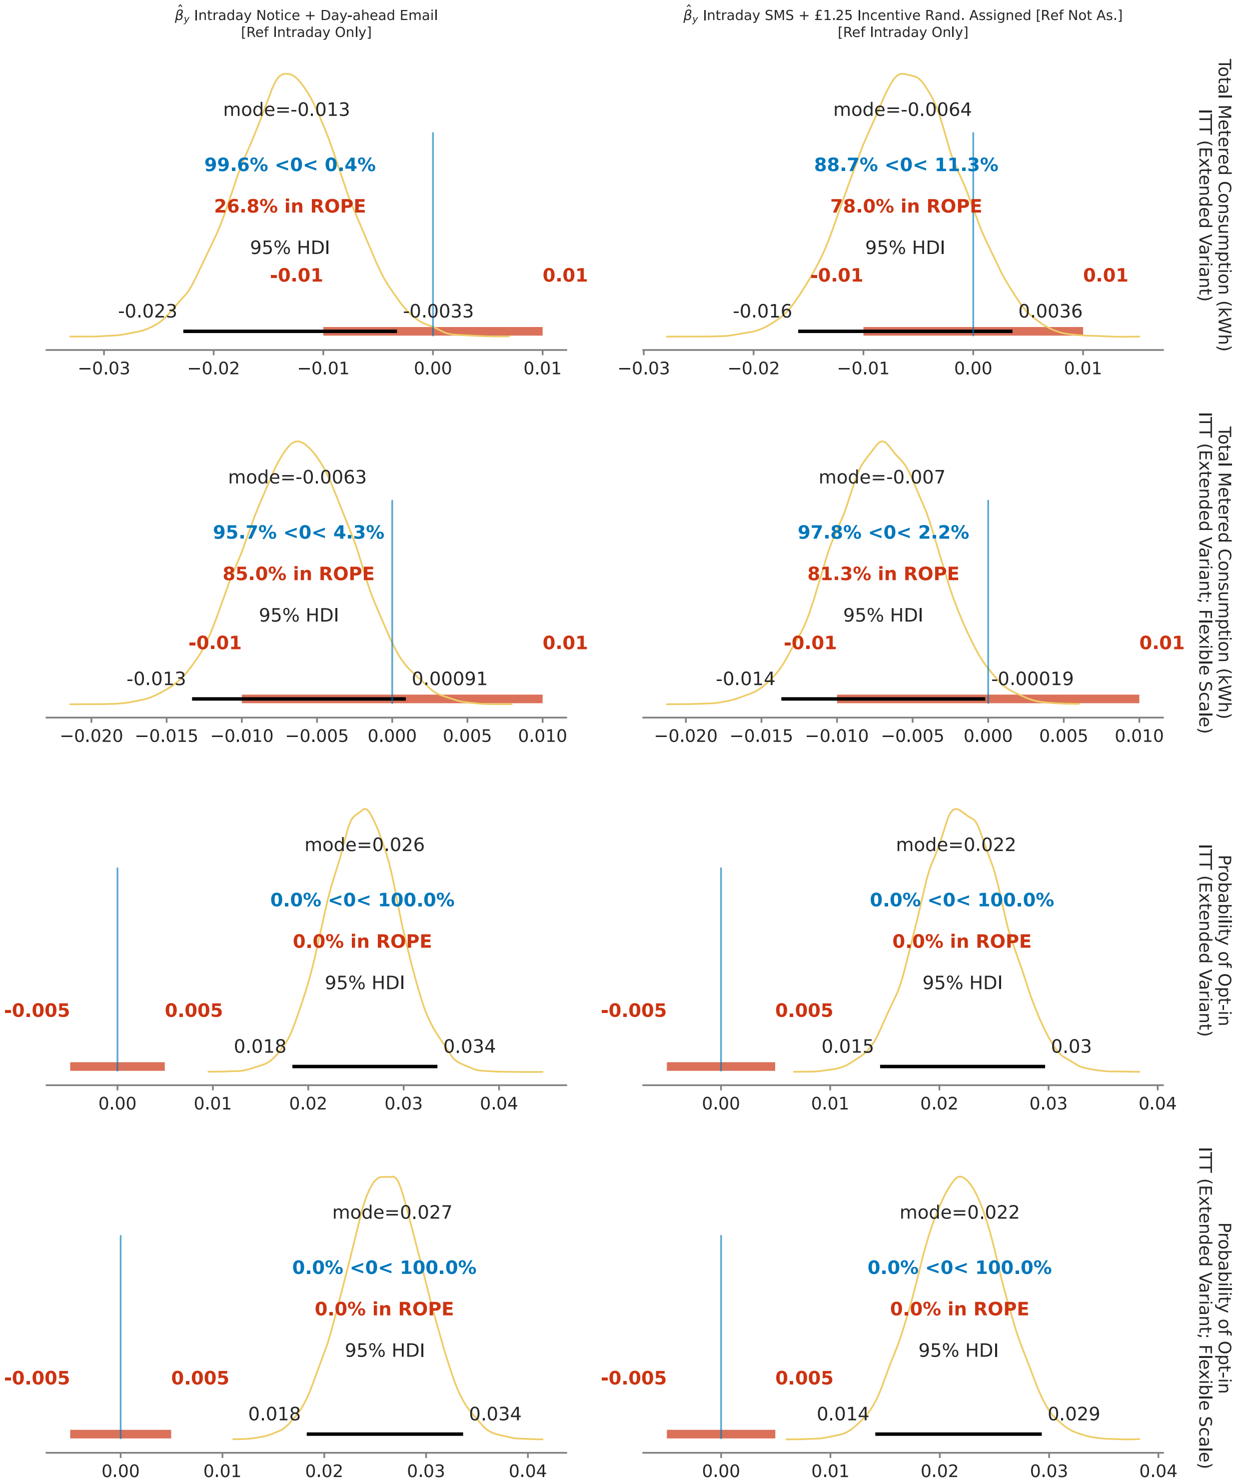

Figure 17: Kruschke-style (2018) posterior density plots for the ATE (left column; $\hat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]) and the ITT (right column; $\hat{\beta}_y$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]) from single-equation regression models fit to data from the randomised controlled trial for the 12th Saving Session (Mar. 15, 2023). Outcome variable y given on right-hand side of plot. Probability of Direction (POD; Blue) indicates the proportion of the posterior distribution (Yellow) greater than zero and less than zero. The Region of Practical Equivalence (ROPE) is the range of values around zero considered to be null for practical purposes. 95% HDI indicated by the black horizontal bar.

Study 1 Extended Details

Background — Known Assignment Mechanism, But No Randomisation

Our first study concerns the causal impact of receiving an intraday opt-in notice as opposed to a day-ahead notice. Like the other Demand Flexibility Service (DFS) events delivered by Octopus Energy (OE) throughout the Winter of 2022-23, OE customers who agreed to participate in the Saving Session on Feb. 13 should have received a notice the day prior. However, notices were delayed (Figure 2). And they were ultimately sent in accordance with the ordering of customers' account IDs (i.e., a string of integers ranging in length) which are, in turn, a function of each customer's *tenure* — where customers new to OE generally have account IDs that are larger in magnitude (Figure 3).

To clarify, opt-in notices for Saving Sessions were distributed to customers in scheduled batches using a roster of account IDs. Batching is a standard practice used to minimise error in the delivery of messages to a large number of customers. And, as a general rule, OE does not release batched communications to customers between 8PM and 8AM. Nevertheless, the process by which messages to customers are generated, batched, and sent is inexact. This is owing to idiosyncratic server delays — where opt-in notices for the Saving Session on Feb. 13 were deferred to an unusual degree into the evening of the 12th and through to the morning of the 13th (Figure 2). Thus, some DFS-participating OE customers received day-ahead opt-in notices whereas others received an opt-in notice for the Feb. 13 Saving Session sometime before 1PM on the day of the Saving Session itself (i.e., intraday).

Note well that account IDs were not manually batched. Instead, a batch size was first manually chosen. And then the internal platform OE uses for external communication was scheduled to dispense batches of notices in order of the magnitude of customers' account IDs. Still, owing to the batching process and standard server lag, account ID does not strictly (i.e., monotonically) increase with time (Figure 3).[16] Furthermore, owing to the above-mentioned imperfect nature of message delivery in relation to OE's customer-communication platform, batched messages began to be sent again around 7:45AM on Feb 13 (Figure 2). Nevertheless, we use 8AM as our temporal cutoff as this is the point at which OE formally allows customer contact.

Specification and Bandwidth Choice given "Donut-Hole" RDD

To fit our regression discontinuity (RD) models, we only use a 2nd-degree polynomial (i.e., a quadratic fit) for our centred assignment variable ($A_i - C$). This choice is made in light of Gelman and Imbens' (2019) and Gelman and Zelizer's (2015) warnings of overfitting and nonsensical conclusions when fitting RD models that include higher-order polynomials (see also Huntington-Klein (2021:516–18)).

Indeed, we prefer a *less*-flexible fit reflective of information further from the cutoff. This is because we utilise a "donut-hole" regression discontinuity design (RDD) (Barreca et al. 2011; Barreca, Lindo, and Waddell 2016) and must necessarily *extrapolate* the regression line forward across the region of our account-ID-based running variable to the immediate left of our cutoff for which we exclude all OE customers sent overnight notices (Figure 2). For this reason, we also eschew weighting observations with values for the assignment variable closer to the threshold during model fitting. Instead, all observations are weighted equally (i.e., a "uniform kernel") and we simply present models using multiple bandwidths (wide and narrow) following the recommendations of Lee and Lemieux (2010).

Given our use of donut-hole RDD, we estimate one optimal *asymmetric* bandwidth using the techniques of Cattaneo et al. (2017; Cattaneo, Idrobo, and Titiunik 2019; Cattaneo, Idrobo, and Titiunik Forthcoming; Cattaneo and Vazquez-Bare 2017) as implemented in the newer, Python-based version of their popular STATA function "rdbwselect" (Calonico et al. 2017). The bandwidth is referred to as "optimal" as it is automatically estimated given: (a) the data; (b) a polynomial order (here, 2nd-degree); (c) a kernel weighting function (here, uniform), and (d) a means of calculating variance (here, K-nearest-neighbours, where $K = 3$), amongst other factors (Cattaneo and Vazquez-Bare 2017). We limit our attention to an optimal bandwidth expected to minimise mean-squared error or "MSE" (i.e., the average of the squared deviations between predicted and observed values). However, we probe the sensitivity of our results by fitting ancillary models after expanding our MSE-optimal bandwidth by a factor of 1.5 and a factor of 2 (both arbitrarily chosen). We use "rdbwselect" to obtain our MSE-Optimal bandwidth without the use of pre-treatment covariates.

Note that an optimal bandwidth is specific to outcome variable (Cattaneo and Vazquez-Bare 2017:143). Thus, we use two *sets* of asymmetric MSE-optimal bandwidths. The first set is specific to our models of total consumption ($h_{Left,\,\text{MSE-Optimal, Consumption}}$ and $h_{Right,\,\text{MSE-Optimal, Consumption}}$) and the second is specific to our models of the probability of participation in the 10th saving session ($h_{Left,\,\text{MSE-Optimal, Opt-in}}$ and $h_{Right,\,\text{MSE-Optimal, Opt-in}}$).

[16] Details of the message-delivery process obtained during conversation between the authors of this research and technical experts at Octopus Energy Group.

To actually filter our data, we construct a range of valid account IDs by taking our constructed ID-based threshold $C = 2,454,839$ and subtracting h_{Left} and adding h_{Right}. The bandwidths obtained using "rdbwselect" are as follows:

- $h_{Left,\text{ MSE-Optimal, Consumption}} = 244,339.92$
- $h_{Right,\text{ MSE-Optimal, Consumption}} = 495,274.84$
- $h_{Left,\text{ MSE-Optimal, Opt-in}} = 489,367.98$
- $h_{Right,\text{ MSE-Optimal, Opt-in}} = 448,440.77$

We how our results might behave under different ranges of our running variable by expanding our bandwidths as follows:

- $h_{Left,\text{ MSE-Optimal, Consumption}} = (244,339.92 \div 1.5)$
- $h_{Right,\text{ MSE-Optimal, Consumption}} = (495,274.84 \times 1.5)$
- $h_{Left,\text{ MSE-Optimal, Consumption}} = (244,339.92 \div 2)$
- $h_{Right,\text{ MSE-Optimal, Consumption}} = (495,274.84 \times 2)$
- $h_{Left,\text{ MSE-Optimal, Opt-in}} = (489,367.98 \div 1.5)$
- $h_{Right,\text{ MSE-Optimal, Opt-in}} = (448,440.77 \times 1.5)$
- $h_{Left,\text{ MSE-Optimal, Opt-in}} = (489,367.98 \div 2)$
- $h_{Right,\text{ MSE-Optimal, Opt-in}} = (448,440.77 \times 2)$

Finally, we note that we observe balance on relevant pre-treatment covariates (e.g., historical energy usage) for our MSE-optimal bandwidth (see Supplementary Figure 18).

Study 2 Extended Details

Background — Random Assignment Mechanism, But Non-Compliance

Random assignment to our first experimental condition (i.e., day-ahead email) was uncomplicated. However, random assignment to our second condition (i.e., SMS plus bonus) suffers from non-compliance. Accordingly, threats to our ability to credibly estimate the causal effect of the SMS-plus-bonus condition include: (a) unobserved third factors determinant of both energy consumption and an OE customer's decision to allow SMS communications from Octopus (e.g., socio-economic status; environmental conscientiousness); and (b) residual imbalance of potential outcomes stemming from our random sub-sampling (Gelman, Hill, and Vehtari 2020).[17]

Recall from the main text that, given non-compliance, we formulate an answer to RQ1 and RQ2 through the lens of a randomised encouragement design (RED) — i.e., a type of experimental setup wherein variation in some difficult-to-directly-manipulate treatment is induced using a source of random variation (i.e., the random "encouragement") that is related to the difficult-to-directly-manipulate treatment *and not* related to the outcome of interest. For example, consider randomised assignment of advertisements across geographic markets and hours of watching a video-streaming service platform, like Netflix or Hulu, in an analysis relating the probability of service renewal to hours of television watched, the latter of which may be difficult to control.

As discussed by Gelman et al. (2020), REDs allow one to estimate causal effects for participants whose behaviour *could be altered* by the random encouragement. These individuals are known as *"compliers"* — i.e., a latent sub-population of entities and the only individuals for whom a RED provides counterfactual information. This is because the treatment-related behaviour of compliers could differ under heterogenous experimental conditions. Using the above streaming example, the group of people for whom differential exposure to advertisement could results in different amounts of television watched. Along this line, REDs allow one to recover a special kind of local average treatment effect (LATE) known as the complier average causal effect (CACE). Importantly, the CACE is distinct from and, possibly, unrepresentative of the average treatment effect or "ATE" in both the broader sample and the wider population from which that sample is drawn (Basu, Coe, and Chapman 2018; Gelman et al. 2020). For Study 2, the ATE is the overall effect of receiving an intraday SMS reminder and being made bonus-incentive eligible.

IV estimation is sometimes viewed with scepticism due to its multiple, peculiar assumptions that are distinct from the familiar (conditional) ignorability assumption. For example, see Gelman et al. (2020, ch. 21), Hernán and Robins (2023, ch. 16), and Cunningham (2021:315–86). Accordingly, here we consider the formal assumptions of IVs in relation to the design of Study 2.

Recall that all individuals assigned to receive our second treatment were made eligible for the bonus price incentive such that the only dimension along which compliance could vary is receipt of the SMS head-up. Accordingly, for our study there can be no *always takers* (i.e., OE customers who would have received an intraday SMS reminder

[17] We observe good balance on relevant pre-treatment covariates across the three experimental conditions (Figure 19).

even if they had not been randomly assigned to the SMS-plus-bonus condition). This is by design as we randomly selected which OE customers were sent SMS text messages. Customers were fully unaware of this randomisation, only learning of the SMS reminder at the time that they received it — if they ever received it. And it was impossible for customers not assigned to the SMS-plus-bonus condition to be inadvertently sent an intraday SMS reminder. Thus, our design also precludes *defiers* — i.e., individuals whose treatment status is always opposite of their randomised encouragement (e.g., customers who would manoeuvre to receive an intraday SMS when not assigned to the SMS group and customers who would manoeuvre to never receive a SMS when assigned to the SMS group). Nevertheless, our design does not rule out *never takers* as there are customers who would not receive an intraday SMS no matter their random assignment (e.g., those customers who disallow SMS communications from OE; customers who inadvertently agree to receive OE's SMS communications but who screen for "spam" to automatically block texts from unknown senders/phone numbers). As discussed by Gelman et al. (2020, p. 425), the presence of never-takers exacerbates violations of the exclusion restriction such that the CACE is expected to be biased by α (i.e., the associated between the instrument Z and the outcome y for the never-takers) multiplied by a factor stemming from the ratio of the proportion of never-takers and the proportion of compliers. Here, this ratio is approximately 1.05 — i.e., 4,731 ÷ 4,472 — where this first figure is the number of individuals randomly assigned to the SMS-plus-bonus condition but who had disallowed SMS messages from OE.

Beyond the assumptions related to the aforementioned sub-populations, there are five primary assumptions of IV estimation (Gelman et al. (2020:421–26, Section 21.1)) in addition to a sixth assumption (Bhuller and Sigstad (2022)) resulting from our decision to estimate the effect of two binary treatments in the same model.

1. Ignobility of the Instrument With Respect to Treatment Potential Outcomes (T^0, T^1) and With Respect to the Response Potential Outcomes (y^0, y^1): Recall that potential outcomes are one's value for a response variable (e.g., session consumption) when exposed to different experimental conditions at the same point in time. And this assumption implies that our instrument (i.e., assignment to the SMS-plus-bonus condition) is independent of the potential outcomes for actually receiving the SMS-plus-bonus treatment *and* the potential outcomes for our response variables (i.e., kWh of energy consumption; participation in the Mar. 15 savings session). This is clearly satisfied as our assignment to the SMS-plus-bonus condition was randomised. Thus we do not anticipate third factors determinant of both condition assignment and actual receipt of treatment or condition assignment and session consumption.

2. Monotonicity: This assumptions implies that there are no *defiers*. As discussed above, this is precluded by the design of our study as OE customers were unaware of our randomisation, OE customers in the control condition (i.e., intraday notice only) could not themselves join the SMS-plus bonus-condition, and OE customers who actually received the SMS-plus-bonus treatment could not somehow reject it to join the control group.

3. Nonzero Association Between the Instrument and the Treatment: This assumption implies that the covariance of the instrument and the treatment is not equal to zero. This assumption is testable. And our models (Figure 4, Figure 10) unsurprisingly indicate a positive association between random assignment to the SMS-plus-bonus condition and receiving the SMS-plus-bonus treatment.

4. Exclusion Restriction: This assumption implies that the randomised instrument (i.e., assignment to the SMS-plus-bonus condition) is not directly associated with the outcome (i.e.session consumption and session participation) for never-takers and always-takers (Gelman et al. 2020:423). This assumption is trivially satisfied in our case as it was impossible for OE customers to receive the SMS-based-treatment on their own (i.e., always-takers) or to independently learn about and ultimately reject the SMS-based treatment (i.e., never-takers).

5. No Crossed Effects (Bhuller and Sigstad 2022): This assumption implies that our randomised instrument is not associated with the other treatments in the model (here, the day-ahead-heads-up-email condition). This assumption is easily satisfied in our case as our two treatment groups are distinct by design and because the day-ahead-heads-up-email condition has (to our knowledge) perfect compliance.

Given all of this, we maintain that instrumental variables estimation can be defensibly used to recover the CACE of the SMS-plus-bonus condition on consumption during, and participation in, the Mar. 15 Saving Session.

Balance on Pre-Treatment Covariates

As discussed by Gelman et al. (2020), under a strict interpretation, balance relate to similarity in the *distributions* of potential-outcomes-relevant pre-treatment variables across levels of a treatment variable — not merely summary statistics (e.g., the mean). Moreover, hypothesis tests for "statistically significant" differences across experimental groups have been subject to repeated critique under the view that balance is a property of a give sample, not the population from which it is drawn (Harvey 2018; Imai, King, and Stuart 2008; Senn 1994).

Accordingly, we eschew comparisons like t-tests to instead qualitatively compare distributions of pre-treatment covariates across experimental groups using quantile-quantile (Q-Q) plots, following Imai et al. (2008). As our concern is energy-consumption related behaviour, we narrowly focus on historic energy usage — i.e., Total P376 (Unadjusted) Baseline (kWh) and Estimated Annual Consumption (kWh). as mentioned in the main text, the former is an unweighted average of consumption during the same half-hour of the day during the ten most-recent working days as governed by the the P376 amendment to Great Britain's electricity balancing and settlement code. And the latter is OE's predicted customer consumption based on meter reading over one year. Given the large size of our data, which are drawn from across Great Britain, we create Q-Q plots specific to region and we also consider balance on the degree to which an OE customer's *postcode* is deprived using an index of multiple deprivation (IMD) which combines, in a weighted manner, multiple facets of poverty (e.g., crime, barriers to housing, health, amongst other factors).[18] Last, for Study 1, we also include Q-Q plots of account ID and customer tenure to show the by-design discontinuity of the former and the sneaking overlap of the latter.

The Intent-to-Treat Effect and the Reduced Form Equation

Alongside results from our two-equation regression models (Figure 10), we also consider a single-equation regression model that replaces the binary indicator for actually receiving the SMS-plus-bonus treatment with a binary indicator for randomised assignment to the SMS-plus-bonus condition. This is a kind of "reduced form" version of our two-equation model (Angrist 2006:32–33; Wooldridge 2010:90–91) that provides us with the intent-to-treat (ITT) effect (Gelman et al. 2020:426) — i.e., the effect of eligibility itself.

Formally, our single-equation model for the ITT is defined as follows:

$$
\begin{gathered}
y_i \sim Normal(\mu_{y,i}, \sigma_y) \\
\mu_{y,i} = \theta_0 + \theta_1 T_{\text{Day-ahead Email},i} + \theta_2 Z_{\text{SMS Randomly Assigned},i} + X_i\vec{\theta} + \epsilon_i \\
\theta_1,\ \theta_2,\ \vec{\theta} \sim Normal(\mu = 0,\ \sigma = 0.5) \\
\theta_0 \sim Normal(\mu = 0,\ \sigma = 0.5) \\
\sigma_y \sim Gamma(\alpha = 2,\ \beta = 0.5),
\end{gathered} \tag{4}
$$

where θ_1 is the ATE of the day-ahead condition on the response y, θ_2 is the ITT, and X is a $N \times p$ matrix containing p pre-treatment covariates and/or confounders.

We also use this general specification for our models fit using data specific to each DNO region.

[18] For further details, see the UK Ministry of Housing, Communities & Local Government.

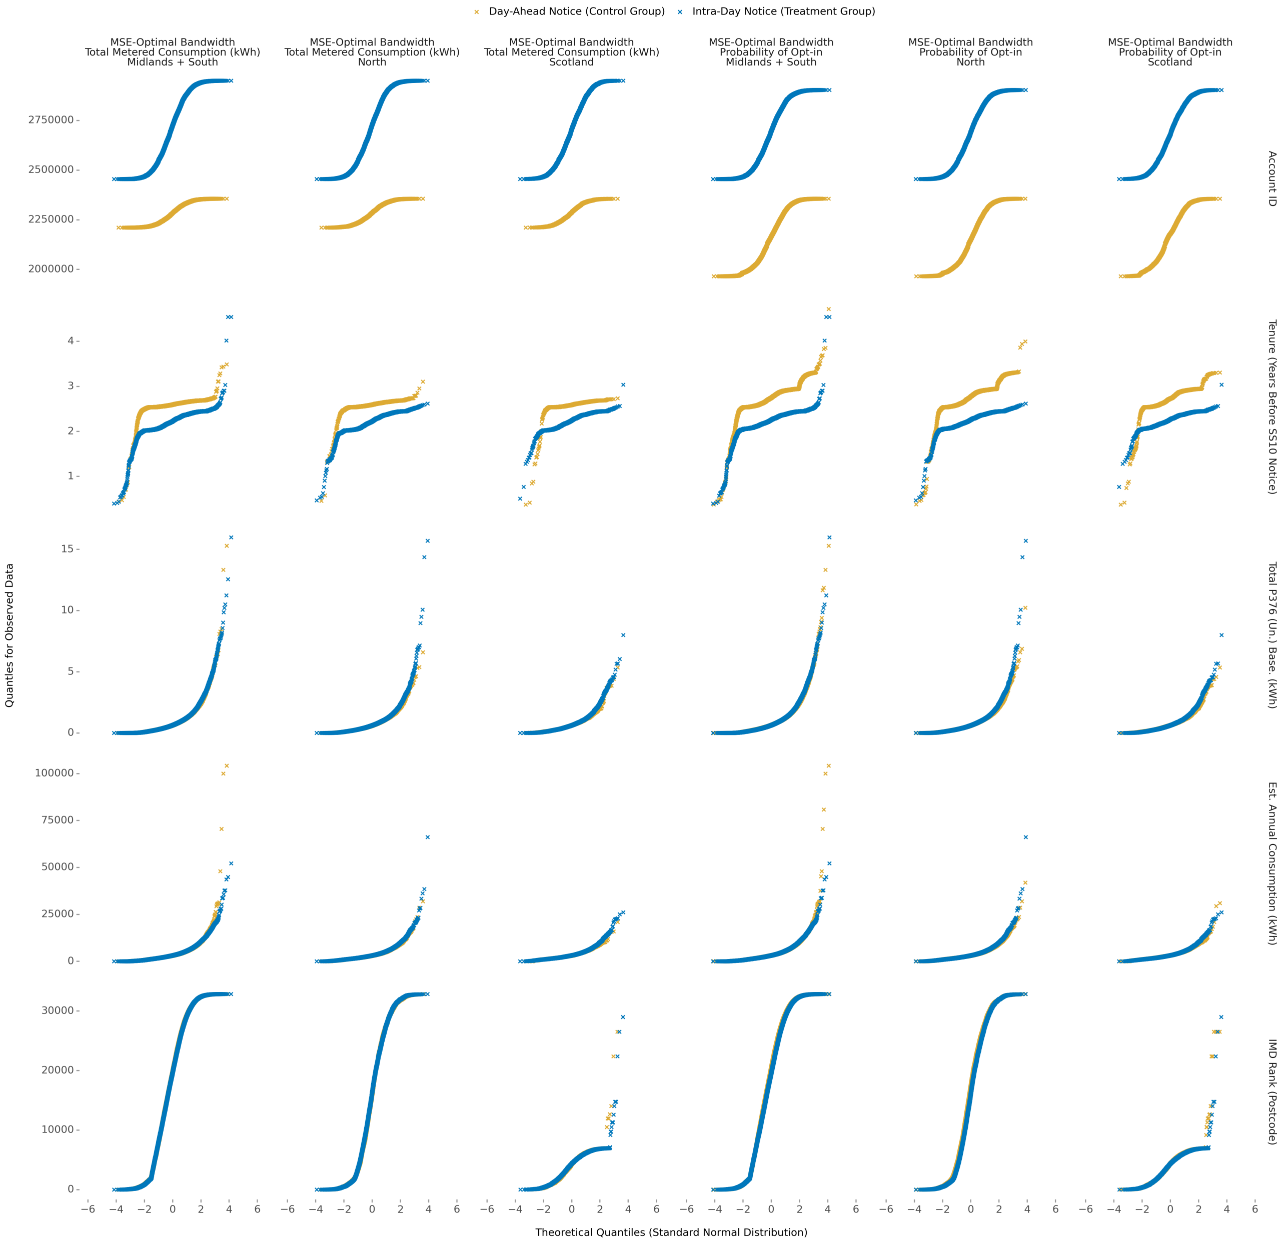

Figure 18: "Small-multiple" Q-Q Plot for balance of relevant pre-treatment covariates across non-randomised experimental groups depending on MSE-optimal bandwidth and outcome (i.e., session consumption or session participation). For the Index of Multiple Deprivation (IMD), more deprived areas have *lower* postcode ranks.

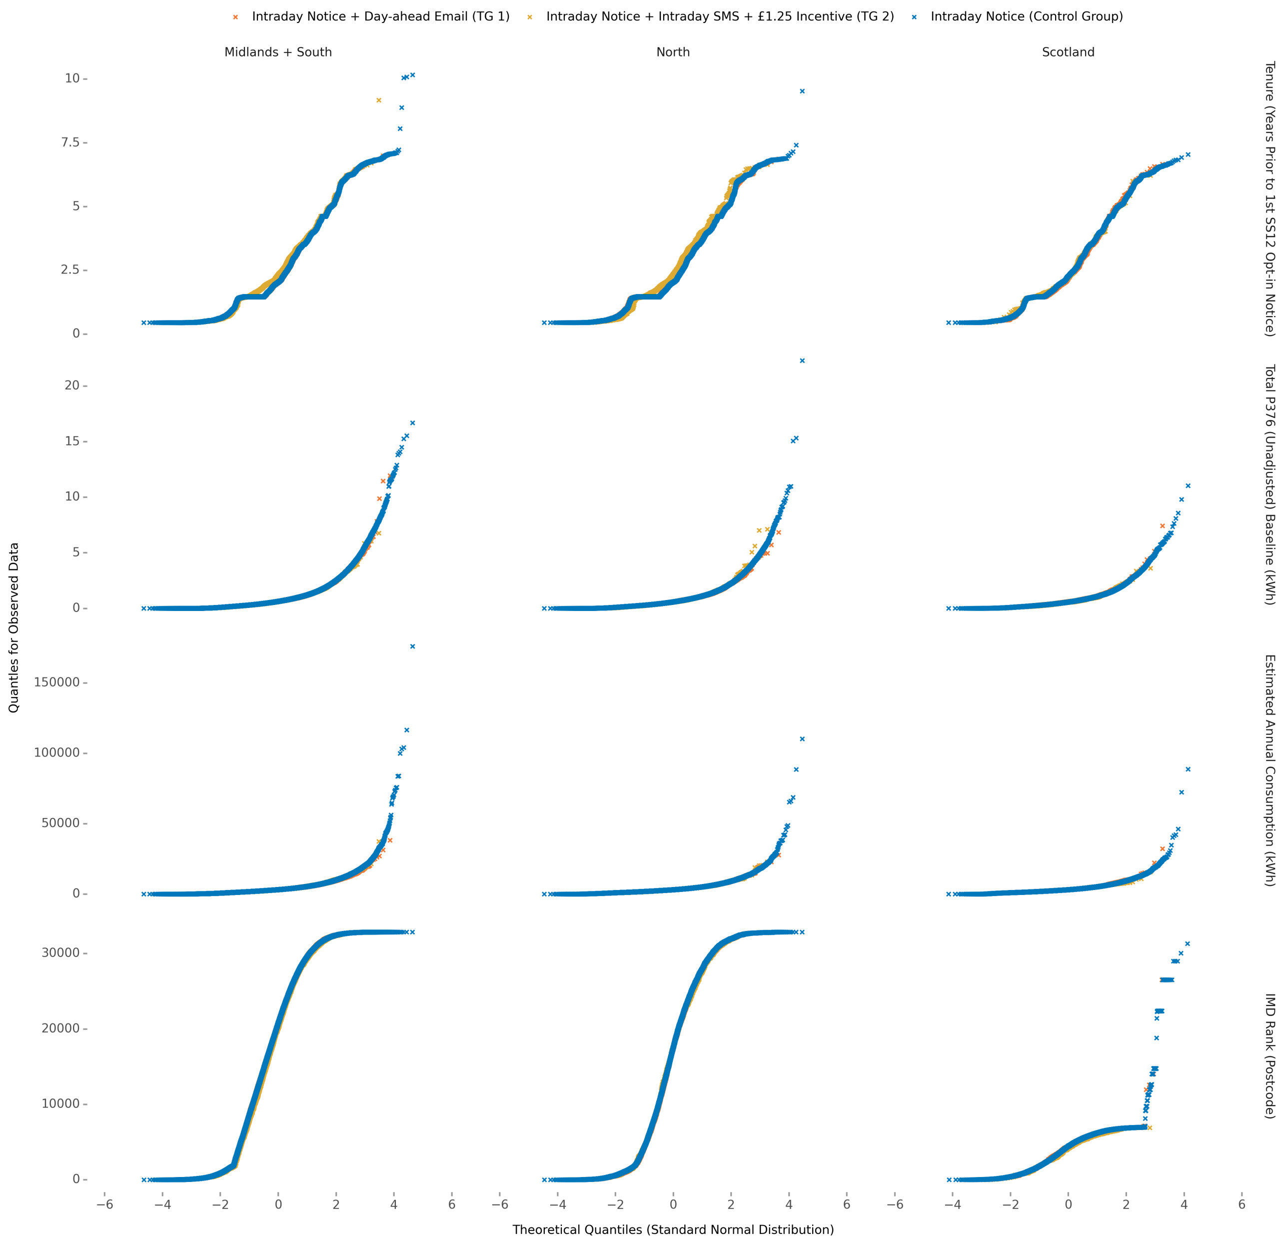

Figure 19: "Small-multiple" Q-Q Plot for balance of pre-treatment covariates across randomised experimental conditions. For the Index of Multiple Deprivation (IMD), more deprived areas have *lower* postcode ranks.

Simple Linear Regression and Limited Response Variables

Recall that, in the style of ordinary-least-squares (OLS) regression, we use and identity link function and Normal (i.e., Gaussian) distributions as the functional form for all models of our two response variables — i.e., (a) total kWh of electricity usage during a 60 min. session; and (b) actual participation in a given Saving Session (i.e., opt-in). The later is of course binary. And consumption is a non-negative, real-valued quantity with true (i.e., valid) zeros (n.b. some customers' may use stored energy captured via their solar panels and retained with a domestic battery).

Accordingly, it may seem odd that we use simple linear regression to model two limited (i.e., range-restricted) variables as opposed to, for instance, Gamma regression for our continuous positive outcome and logistic regression for our binary outcomes. This choice may seem especially odd to readers unfamiliar with econometric techniques or (Bayesian) readers who maintain that a statistical model ought to encapsulate theoretical mechanisms such that it is *generative* of the observed data (e.g., see Gabry et al. (2019)).

Nevertheless, simple linear regression, specifically OLS, is the standard approach to obtaining causal effects amongst econometricians. These individuals tend to narrowly focus on accurately estimating treatment effects in the form of conditional means as opposed to performing in-sample or out-of-sample prediction tasks. Angrist (2006), Angrist and Pischke (2009), Basu et al. (2018), and Wooldridge (2010) provide comprehensive discussions of this OLS-centric logic, which we adopt here.

That said, use of simple linear regression to model limited dependent variables is not without controversy — particularly with respect to instrumental variable (IV) estimation with a binary instrument, binary treatment, and a binary outcome (Study 2). Indeed, researchers from multiple scientific disciplines, including economics, have cautioned against using linear regression in this scenario (Bhattacharya, Goldman, and McCaffrey 2006; Dong and Lewbel 2015; Hollenbach, Montgomery, and Crespo-Tenorio 2019; Huntington-Klein 2021; Kleibergen and Zivot 2003; Li et al. 2022).

Still, we prefer to use Normal functional forms for both of our outcome variables and for both our regression discontinuity design (Study 1) and our IV estimation (Study 2) in light of the popularity of linear probability models (Angrist 2006; Angrist and Pischke 2009; Wooldridge 2010) and their interpretive benefits compared to *non*-linear probability models (see Battey et al. (2019), Breen et al. (2018), and Gomila (2021) for comparisons). Moreover, we anticipate that our results are robust to use of simple linear regression following discussions by Angrist (2006), Basu et al. (2018), Breen et al. (2018), and Wooldridge (2010), particularly given our sample size for Study 1 ($\approx$ 80K-350K, depending on model) and Study 2 ($\approx$ 650K).

Indeed, Chiburis et al. (2012) suggest that, in the binary-instrument-binary-treatment-binary-outcome scenario, results from IV estimation obtained using linear probability models or a probit model are only likely to differ for sample sizes less than 5,000 and when the probability of treatment and/or the probability of a positive outcome for the response (here, session opt-in) are close to 0.1 of 0.9. This latter warning is perhaps relevant to Study 2 as roughly 23% of individuals received the SMS-plus-bonus treatment. However, Basu et al. (2018) point to the favourable performance of two-stage OLS for estimating the LATE in a binary-instrument-binary-treatment-binary-outcome scenario across a wide range of treatment rarities, where we note again that our sample size for Study 2 is substantial.

Given all of this, we opt for simple linear regression. In so doing, we necessarily adopt an interpretation of our models as linear projections of our nonlinear outcomes (Gomila 2021; Wooldridge 2010). Furthermore, we limit our attention to point estimates as opposed to predicted quantities across a wide range of values for our covariates which could, in some instances, be nonsensical (e.g., negative consumption or negative probabilities). As discussed by Gelman et al. (2020:154–55; also Gelman et al. 2014), the outcome variable is not required to be normally-distributed when using simple linear regression and both the unequal variance of errors (i.e., heteroscedasticity) and the non-normality of errors are expected to be minor issues when using simple linear regression to estimate conditional means and to recover the regression line as opposed to the prediction of individual data points. Along this line, all continuous pre-treatment variables and confounders enter our regression equations as Z-scores created by subtracting the sample mean and dividing by the sample standard deviation.

That said, we still wish to consider the extent to which heteroscedasticity might underlie our results. Accordingly, we modify the equation used for our regression discontinuity models (i.e., Equation (2), Study 1) by using a linear predictor for both location (i.e, the conditional mean) *and* scale (i.e., the standard deviation of the response).

Specifically, our single-equation model for the LATE with flexible scale is defined as follows:

$$
\begin{aligned}
y_i &\sim Normal(\mu_{y,i}, \sigma_{y,i}) \\
\mu_{y,i} &= \beta_0 + \beta_1 z_i + \beta_2(A_i - C_{\text{ID}}) + \beta_3(z_i \times (A_i - C_{\text{ID}})) + \beta_4(A_i - C_{\text{ID}})^2 + \beta_5(z_i \times (A_i - C_{\text{ID}})^2) + X_i\vec{\beta} + \epsilon_i \\
\sigma_{y,i} &= \exp(\phi_0 + \phi_1 z_i + \phi_2(A_i - C_{\text{ID}}) + \phi_3(z_i \times (A_i - C_{\text{ID}})) + \phi_4(A_i - C_{\text{ID}})^2 + \phi_5(z_i \times (A_i - C_{\text{ID}})^2) + X_i\vec{\beta}) \\
&\forall \quad A_i \in (C_{\text{ID}} - h_{Left}, C_{\text{ID}} + h_{Right}) \\
&\beta_{1,\ldots,5},\ \vec{\beta},\ \phi_{1,\ldots,5},\ \vec{\phi} \sim Normal(\mu = 0,\ \sigma = 0.5) \\
&\beta_0,\ \phi_0 \sim Normal(\mu = 0,\ \sigma = 0.5).
\end{aligned} \tag{5}
$$

Note that the set of correlates for $\sigma_{y,i}$ may be arbitrarily chosen. However, we use an identical set of correlates for the location and scale sub-models owing to a lack of theoretical rationale for the exclusion of various predictors and due to the fundamental heteroscedasticity of linear probability models (see Rencher and Schaalje (2008:508–9)). As $\sigma_{y,i}$ must be strictly positive, we use the log link function. Thus the regression coefficients ϕ relate to changes in $\ln_{\sigma_{y,i}}$, where the expected value of $\sigma_{y,i}$ on its original scale may be obtained using the exponential function.

Unfortunately, we are unable to include a linear predictor for the scale parameters in our simultaneous-equation models (i.e., Equation (3), Study 2). This is owing to particularities of the implementation of the multivariate normal distribution in the probabilistic programming language that we rely on. Specifically, the multivariate normal distribution implemented in PyMC does not, at the time of writing, allow one to calculate the overall log-likelihood using multiple, observation-specific covariance matrices (i.e., "batched" operations). And the massive scale of the data from our RCT ($N \approx$ 600k) presents us from dealing with this issue through utilisation of block diagonal covariance matrices. This is because performing operations on $N \times N$ arrays of the size we consider for our research is computationally burdensome, easily overwhelming our available RAM. Thus, we instead use the reduced-form model (i.e., Equation (4), Study 2) to consider whether non-constant variance might drive our RCT-based results as it is a standard single-equation linear regression model amenable to batched computation with PyMC.

Accordingly, our single-equation model for the ITT with flexible scale is defined as follows:

$$
\begin{aligned}
y_i &\sim Normal(\mu_{y,i}, \sigma_{y,i}) \\
\mu_{y,i} &= \theta_0 + \theta_1 T_{\text{Day-ahead Email},i} + \theta_2 Z_{\text{SMS Randomly Assigned},i} + X_i\vec{\theta} + \epsilon_i \\
\sigma_{y,i} &= \exp(\phi_0 + \phi_1 T_{\text{Day-ahead Email},i} + \phi_2 Z_{\text{SMS Randomly Assigned},i} + X_i\vec{\theta}) \\
&\theta_{1,\ldots,5},\ \theta_\beta,\ \phi_{1,\ldots,5},\ \vec{\phi} \sim Normal(\mu = 0,\ \sigma = 0.5) \\
&\theta_0,\ \phi_0 \sim Normal(\mu = 0,\ \sigma = 0.5).
\end{aligned} \tag{6}
$$

Why are we Bayesian?

Bayesian inference is a widely-used approach to model estimation wherein available (i.e., "prior", "a priori") information about unknown parameters in a statistical model (e.g., the regression coefficients) is combined or "updated" with information on these parameters encoded in the observed data to which a model is to fit. This combination produces "posterior" (i.e., "a posteriori") information summarising the likelihood of each model parameter given the data and the priors (van de Schoot et al. 2021).

As mentioned in the main text, we estimate our regression models within a Bayesian framework using the powerful and flexible Python-based probabilistic programming language PyMC. We do so as Bayesian inference provides intuitive uncertainty intervals that describe the plausibility of parameter values in terms of probabilities. This is in contrast to frequentist confidence intervals which relate to the repeated execution of an analytic procedure using different samples. Furthermore, Bayesian inference does not rely on p-values which have been widely critiqued as metrics of evidence and decision-making as social and natural scientists have sought alternative research workflows in the name of reproducible research (Halsey 2019; Kruschke 2018; Kruschke and Liddell 2018; Makowski et al. 2019; McShane et al. 2019). Moreover, the large scale of our data makes p-values less-meaningful metrics of evidence. This is because p-values are generally expected to shrink in size as a sample grows (Demidenko 2016) inducing a phenomenon whereby many, possibly all, parameters are "statistically significant".

Readers unversed in Bayesianism should see van de Schoot et al. (2021) for an introduction and Kruschke (2021) and Makowski et al. (2019) on its general benefits, parameter "credibility", and the effective reporting of Bayesian analyses. Furthermore, Gelman (2008) and Krushke and Liddell (2018) provide rebuttals of common objections to Bayesian analyses by frequentists (see also Gelman et al. (Gelman et al. 2014, 2020; Gelman and Shalizi 2013)).

What are Our Priors?

We generally follow the prior-specification recommendations of Gelman and colleagues on their popular "Prior Choice Wiki". That said, we also tailor our prior specifications to aspects of our particular regression problems (e.g., IV estimation).

Specifically, priors for our regression coefficients for the location and scale sub-models in both our regression discontinuity (RD) and our instrumental variable (IV) models are normal distributions (i.e., bell-curve-like) with a mean (μ) of zero, and a scale parameter (σ) equal to the value of 0.5 for both our non-varying intercepts/constants and our non-varying slope coefficients. Given the scale of our outcome (i.e., kWh over 60 minutes with an average of ≈ 0.6; the probability of session participation) and the scaling of our covariates (generally binary or Z-Scores; n.b., predictions on the log scale for the location sub-models), these priors are expected to be "weakly informative" the narrower range of values with probability density around zero (i.e., no effect). Note, the "Prior Choice Wiki" proposes a "generic weakly informative prior" in the form of normal distribution with a mean of zero and standard deviation of one. However, we found that using tighter normal distributions mitigated issues during sampling, namely slow chain mixing and, for a few models, the occurrence of (divergent transitions)[https://www.martinmodrak.cz/2018/02/19/taming-divergences-in-stan-models/].

With regard to our main models for IV estimation, our principle concern is the prior for the covariance of the errors between the simultaneous regression equations (i.e., $\epsilon_{i,y}$ and $\epsilon_{i,T}$) as it governs the expected degree of "endogeneity" — i.e., the extent to which unobserved aspects of the treatment-with-imperfect-compliance ($\epsilon_{i,T}$) are associated with unobserved aspects of the outcome used for the second equation ($\epsilon_{i,y}$) (Rossi, Allenby, and McCulloch 2005).

Following methodological work on IV estimation in a Bayesian context (Lopes and Polson 2014; Rossi et al. 2005), in addition to standard practice in contemporary applications of Bayesian inference (McElreath 2020), we decompose the variance-covariance matrix (i.e., Σ) of the two-dimension multivariate normal distribution into matrix factors — namely: (a) a 2 x 2 matrix with standard deviations for the equation-specific errors (i.e., σ_y and σ_T) along the prime diagonal and zeros along the anti-diagonal; and (b) a 2 x 2 correlation matrix (Ω) with ones along the prime diagonal and the correlation between $\epsilon_{i,y}$ and $\epsilon_{i,T}$ — i.e., ρ — along the anti-diagonal.

Because σ_y and σ_T are strictly positive, for their priors, we uses Gamma distributions with shape parameters $\alpha = 2$ and $\beta = 1$. This results in a peak in the probability density around 2.5 with a decay in density from the value of 2.5 to roughly the value of 15. We use the same Gamma prior for our RD models.

Following Lopes and Polson (2014), we use a popular distribution developed by Lewandowski, Kurowicka and Joe (2009) or "LKJ" as the prior for the correlation matrix Ω. The LKJ distribution has one shape parameter η. When set to the value of one, the LKJ distribution is a uniform distribution over correlations (i.e., [-1, 1]). Here, we set η equal to the value of one. Note, one can manually recover the variance-covariance matrix Σ using the matrix dot product of the diagonal matrix of standard deviations (Diag(σ)) and the correlation matrix Ω as depicted in Equation (3).

Supplementary References

Angrist, Joshua D. 2006. "Instrumental Variables Methods in Experimental Criminological Research: What, Why and How?" *Journal of Experimental Criminology* 2(1):23–44. doi: 10.1007/s11292-005-5126-x.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Barreca, A. I., M. Guldi, J. M. Lindo, and G. R. Waddell. 2011. "Saving Babies? Revisiting the Effect of Very Low Birth Weight Classification." *The Quarterly Journal of Economics* 126(4):2117–23. doi: 10.1093/qje/qjr042.

Barreca, Alan I., Jason M. Lindo, and Glen R. Waddell. 2016. "Heaping-Induced Bias in Regression-Discontinuity Designs." *Economic Inquiry* 54(1):268–93. doi: 10.1111/ecin.12225.

Basu, Anirban, Norma B. Coe, and Cole G. Chapman. 2018. "2SLS Versus 2SRI: Appropriate Methods for Rare Outcomes and/or Rare Exposures." *Health Economics* 27(6):937–55. doi: 10.1002/hec.3647.

Battey, Heather S., David R. Cox, and Michelle V. Jackson. 2019. "On the Linear in Probability Model for Binary Data." *Royal Society Open Science* 6(5):190067. doi: 10.1098/rsos.190067.

Benjamin, Daniel J., James O. Berger, Magnus Johannesson, Brian A. Nosek, E.-J. Wagenmakers, Richard Berk, Kenneth A. Bollen, Björn Brembs, Lawrence Brown, Colin Camerer, David Cesarini, Christopher D. Chambers, Merlise Clyde, Thomas D. Cook, Paul De Boeck, Zoltan Dienes, Anna Dreber, Kenny Easwaran, Charles Efferson, Ernst Fehr, Fiona Fidler, Andy P. Field, Malcolm Forster, Edward I. George, Richard Gonzalez, Steven Goodman, Edwin Green, Donald P. Green, Anthony G. Greenwald, Jarrod D. Hadfield, Larry V. Hedges, Leonhard Held, Teck Hua Ho, Herbert Hoijtink, Daniel J. Hruschka, Kosuke Imai, Guido Imbens, John P. A. Ioannidis, Minjeong Jeon, James Holland Jones, Michael Kirchler, David Laibson, John List, Roderick Little, Arthur Lupia, Edouard Machery, Scott E. Maxwell, Michael McCarthy, Don A. Moore, Stephen L. Morgan, Marcus Munafó, Shinichi Nakagawa, Brendan Nyhan, Timothy H. Parker, Luis Pericchi, Marco Perugini, Jeff Rouder, Judith Rousseau, Victoria Savalei, Felix D. Schönbrodt, Thomas Sellke, Betsy Sinclair, Dustin Tingley, Trisha Van Zandt, Simine Vazire, Duncan J. Watts, Christopher Winship, Robert L. Wolpert, Yu Xie, Cristobal Young, Jonathan Zinman, and Valen E. Johnson. 2017. "Redefine Statistical Significance." *Nature Human Behaviour* 2(1):6–10. doi: 10.1038/s41562-017-0189-z.

Bhattacharya, Jay, Dana Goldman, and Daniel McCaffrey. 2006. "Estimating Probit Models with Self-Selected Treatments." *Statistics in Medicine* 25(3):389–413. doi: 10.1002/sim.2226.

Bhuller, Manudeep, and Henrik Sigstad. 2022. "2SLS with Multiple Treatments." Retrieved July 28, 2023 (https://arxiv.org/abs/2205.07836).

Breen, Richard, Kristian Bernt Karlson, and Anders Holm. 2018. "Interpreting and Understanding Logits, Probits, and Other Nonlinear Probability Models." *Annual Review of Sociology* 44(1):39–54. doi: 10.1146/annurev-soc-073117-041429.

Calonico, Sebastian, Matias D. Cattaneo, Max H. Farrell, and Rocío Titiunik. 2017. "Rdrobust: Software for Regression-Discontinuity Designs." *The Stata Journal: Promoting Communications on Statistics and Stata* 17(2):372–404. doi: 10.1177/1536867X1701700208.

Cattaneo, Matias D., Nicolas Idrobo, and Rocio Titiunik. 2019. *A Practical Introduction to Regression Discontinuity Designs: Foundations* [computer program]. Version 1. https://arxiv.org/abs/1911.09511.

Cattaneo, Matias D., Nicolás Idrobo, and Rocío Titiunik. Forthcoming. *A Practical Introduction to Regression Discontinuity Designs: Extensions*. Cambridge University Press.

Cattaneo, Matias D., and Gonzalo Vazquez-Bare. 2017. "The Choice of Neighborhood in Regression Discontinuity Designs." *Observational Studies* 3(2):134–46. doi: 10.1353/obs.2017.0002.

Chiburis, Richard C., Jishnu Das, and Michael Lokshin. 2012. "A Practical Comparison of the Bivariate Probit and Linear IV Estimators." *Economics Letters* 117(3):762–66. doi: 10.1016/j.econlet.2012.08.037.

Cunningham, Scott. 2021. *Causal Inference: The Mixtape*. Yale University Press.

Demidenko, Eugene. 2016. "The p -Value You Can't Buy." *The American Statistician* 70(1):33–38. doi: 10.1080/00031305.2015.1069760.

Dong, Yingying, and Arthur Lewbel. 2015. "A Simple Estimator for Binary Choice Models with Endogenous Regressors." *Econometric Reviews* 34(1-2):82–105. doi: 10.1080/07474938.2014.944470.

Gabry, Jonah, Daniel Simpson, Aki Vehtari, Michael Betancourt, and Andrew Gelman. 2019. "Visualization in Bayesian Workflow." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 182(2):389–402. doi: 10.1111/rssa.12378.

Gelman, Andrew. 2008. "Objections to Bayesian Statistics." *Bayesian Analysis* 3(3):445–49. doi: 10.1214/08-BA318.

Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. 2014. *Bayesian Data Analysis*. 3rd Edition. CRC Press, Taylor & Francis Group.

Gelman, Andrew, Jennifer Hill, and Aki Vehtari. 2020. *Regression and Other Stories*. 1st ed. Cambridge University Press.

Gelman, Andrew, and Guido Imbens. 2019. "Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs." *Journal of Business & Economic Statistics* 37(3):447–56. doi: 10.1080/07350015.2017.1366909.

Gelman, Andrew, and Cosma Rohilla Shalizi. 2013. "Philosophy and the Practice of Bayesian Statistics: *Philosophy and the Practice of Bayesian Statistics*." *British Journal of Mathematical and Statistical Psychology* 66(1):8–38. doi: 10.1111/j.2044-8317.2011.02037.x.

Gelman, Andrew, and Adam Zelizer. 2015. "Evidence on the Deleterious Impact of Sustained Use of Polynomial Regression on Causal Inference." *Research & Politics* 2(1):205316801556983. doi: 10.1177/2053168015569830.

Gomila, Robin. 2021. "Logistic or Linear? Estimating Causal Effects of Experimental Treatments on Binary Outcomes Using Regression Analysis." *Journal of Experimental Psychology: General* 150(4):700–709. doi: 10.1037/xge0000920.

Halsey, Lewis G. 2019. "The Reign of the p-Value Is Over: What Aalternative Analyses Could We Employ to Fill the Power Vacuum?" *Biology Letters* 15(5):20190174. doi: 10.1098/rsbl.2019.0174.

Harvey, Lisa A. 2018. "Statistical Testing for Baseline Differences Between Randomised Groups Is Not Meaningful." *Spinal Cord* 56(10):919–19. doi: 10.1038/s41393-018-0203-y.

Hernán, Miguel M., and James M. Robins. 2023. *Causal Inference: What If*. Chapman & Hall/CRC.

Hollenbach, Florian M., Jacob M. Montgomery, and Adriana Crespo-Tenorio. 2019. "Bayesian Versus Maximum Likelihood Estimation of Treatment Effects in Bivariate Probit Instrumental Variable Models." *Political Science Research and Methods* 7(3):651–59. doi: 10.1017/psrm.2018.15.

Huntington-Klein, Nick. 2021. *The Effect: An Introduction to Research Design and Causality*. Boca Raton: CRC Press, Taylor & Francis Group.

Imai, Kosuke, Gary King, and Elizabeth A. Stuart. 2008. "Misunderstandings Between Experimentalists and Observationalists about Causal Inference." *Journal of the Royal Statistical Society Series A: Statistics in Society* 171(2):481–502. doi: 10.1111/j.1467-985X.2007.00527.x.

Kleibergen, Frank, and Eric Zivot. 2003. "Bayesian and Classical Approaches to Instrumental Variable Regression." *Journal of Econometrics* 114(1):29–72. doi: 10.1016/S0304-4076(02)00219-1.

Kruschke, John K. 2018. "Rejecting or Accepting Parameter Values in Bayesian Estimation." *Advances in Methods and Practices in Psychological Science* 1(2):270–80. doi: https://psycnet.apa.org/doi/10.1177/2515245918771304.

Kruschke, John K. 2021. "Bayesian Analysis Reporting Guidelines." *Nature Human Behaviour* 5(10):1282–91. doi: 10.1038/s41562-021-01177-7.

Kruschke, John K., and Torrin M. Liddell. 2018. "The Bayesian New Statistics: Hypothesis Testing, Estimation, Meta-Analysis, and Power Analysis from a Bayesian Perspective." *Psychonomic Bulletin & Review* 25(1):178–206. doi: 10.3758/s13423-016-1221-4.

Lakens, Daniel, Federico G. Adolfi, Casper J. Albers, Farid Anvari, Matthew A. J. Apps, Shlomo E. Argamon, Thom Baguley, Raymond B. Becker, Stephen D. Benning, Daniel E. Bradford, Erin M. Buchanan, Aaron R. Caldwell, Ben Van Calster, Rickard Carlsson, Sau-Chin Chen, Bryan Chung, Lincoln J. Colling, Gary S. Collins, Zander Crook, Emily S. Cross, Sameera Daniels, Henrik Danielsson, Lisa DeBruine, Daniel J. Dunleavy, Brian D. Earp, Michele I. Feist, Jason D. Ferrell, James G. Field, Nicholas W. Fox, Amanda Friesen, Caio Gomes, Monica Gonzalez-Marquez, James A. Grange, Andrew P. Grieve, Robert Guggenberger, James Grist, Anne-Laura Van Harmelen, Fred Hasselman, Kevin D. Hochard, Mark R. Hoffarth, Nicholas P. Holmes, Michael Ingre, Peder M. Isager, Hanna K. Isotalus, Christer Johansson, Konrad Juszczyk, David A. Kenny, Ahmed A. Khalil, Barbara Konat, Junpeng Lao, Erik Gahner Larsen, Gerine M. A. Lodder, Jiří Lukavský, Christopher R. Madan, David Manheim, Stephen R. Martin, Andrea E. Martin, Deborah G. Mayo, Randy J. McCarthy, Kevin McConway, Colin McFarland, Amanda Q. X. Nio, Gustav Nilsonne, Cilene Lino De Oliveira, Jean-Jacques Orban De Xivry, Sam Parsons, Gerit Pfuhl, Kimberly A. Quinn, John J. Sakon, S. Adil Saribay, Iris K. Schneider, Manojkumar Selvaraju, Zsuzsika Sjoerds, Samuel G. Smith, Tim Smits, Jeffrey R. Spies, Vishnu Sreekumar, Crystal N. Steltenpohl, Neil Stenhouse, Wojciech Świątkowski, Miguel A. Vadillo, Marcel A. L. M. Van Assen, Matt N. Williams, Samantha E. Williams, Donald R. Williams, Tal Yarkoni, Ignazio Ziano, and Rolf A. Zwaan. 2018. "Justify Your Alpha." *Nature Human Behaviour* 2(3):168–71. doi: 10.1038/s41562-018-0311-x.

Lee, David S., and Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature* 48(2):281–355. doi: 10.1257/jel.48.2.281.

Lewandowski, Daniel, Dorota Kurowicka, and Harry Joe. 2009. "Generating Random Correlation Matrices Based on Vines and Extended Onion Method." *Journal of Multivariate Analysis* 100(9):1989–2001. doi: 10.1016/j.jmva.2009.04.008.

Li, Chuhui, Donald S. Poskitt, Frank Windmeijer, and Xueyan Zhao. 2022. "Binary Outcomes, OLS, 2SLS and IV Probit." *Econometric Reviews* 41(8):859–76. doi: 10.1080/07474938.2022.2072321.

Lopes, Hedibert F., and Nicholas G. Polson. 2014. "Bayesian Instrumental Variables: Priors and Likelihoods." *Econometric Reviews* 33(1-4):100–121. doi: 10.1080/07474938.2013.807146.

Makowski, Dominique, Mattan S. Ben-Shachar, S. H. Annabel Chen, and Daniel Lüdecke. 2019. "Indices of Effect Existence and Significance in the Bayesian Framework." *Frontiers in Psychology* 10:2767. doi: 10.3389/fpsyg.2019.02767.

McElreath, Richard. 2020. *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. 2nd ed. Boca Raton: CRC Press, Taylor & Francis Group.

McShane, Blakeley B., David Gal, Andrew Gelman, Christian Robert, and Jennifer L. Tackett. 2019. "Abandon Statistical Significance." *The American Statistician* 73:235–45. doi: 10.1080/00031305.2018.1527253.

Rencher, Alvin C., and G. Bruce Schaalje. 2008. *Linear Models in Statistics*. 2nd ed. Hoboken, N.J: Wiley-Interscience.

Rossi, Peter E., Greg M. Allenby, and Robert E. McCulloch. 2005. *Bayesian Statistics and Marketing*. Hoboken, NJ: Wiley.

Senn, Stephen. 1994. "Testing for Baseline Balance in Clinical Trials." *Statistics in Medicine* 13(17):1715–26. doi: 10.1002/sim.4780131703.

van de Schoot, Rens, Sarah Depaoli, Ruth King, Bianca Kramer, Kaspar Märtens, Mahlet G. Tadesse, Marina Vannucci, Andrew Gelman, Duco Veen, Joukje Willemsen, and Christopher Yau. 2021. "Bayesian Statistics and Modelling." *Nature Reviews Methods Primers* 1(1):1. doi: 10.1038/s43586-020-00001-2.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data*. 2nd ed. Cambridge, Mass: MIT Press.

Supplementary Tables

Octopus Energy "Saving Sessions"

Table 2: Dates of the 13 DFS-related events delivered by Octopus Energy (OE) throughout Winter 2022-23. Price incentives for OE customers administered via a points-based rewards scheme (i.e., "OctoPoints") as opposed to direct payment or cash credit.

Session Start Time	Session End Time	OctoPoints (per kW-hr)	Incentive (£/kW-hr)
2022-11-15 at 17:00	2022-11-15 at 18:00	1800	2.25
2022-11-22 at 17:30	2022-11-22 at 18:30	1800	2.25
2022-11-30 at 17:30	2022-11-30 at 18:30	1800	2.25
2022-12-01 at 17:00	2022-12-01 at 18:00	1800	2.25
2022-12-12 at 17:00	2022-12-12 at 19:00	1800	2.25
2023-01-19 at 09:00	2023-01-19 at 10:00	1800	2.25
2023-01-23 at 17:00	2023-01-23 at 18:00	2700	3.38
2023-01-24 at 16:30	2023-01-24 at 18:00	3200	4.00
2023-01-30 at 09:00	2023-01-30 at 10:00	1800	2.25
2023-02-13 at 17:30	2023-02-13 at 18:30	1800	2.25
2023-02-21 at 17:30	2023-02-21 at 18:30	1800	2.25
2023-03-15 at 16:30	2023-03-15 at 17:30	1800	2.25
2023-03-23 at 16:30	2023-03-23 at 17:30	1800	2.25

All Study 1 (RDD) Results in Tabular Format

Below, we include all numeric results from each model fitted for Study 1.

Table 3: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response Y = "Total Metered Consumption (kWh)") fitted using the "Baseline (MSE-Optimal Bandwidth)" specification ($N = 78,724$).

Parameter	Mode	Mean	Std.	95% HDI
$\hat{\beta}$ Intercept	0.567	0.569	0.023	[0.525, 0.614]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead]	0.057	0.056	0.024	[0.01, 0.104]
$\hat{\beta}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	-0.174	-0.181	0.130	[-0.435, 0.071]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)	0.179	0.190	0.131	[-0.068, 0.445]
$\hat{\sigma}_y$	0.755	0.756	0.002	[0.752, 0.759]

Table 4: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response $Y =$ "Total Metered Consumption (kWh)") fitted using the "Extended (MSE-Optimal Bandwidth)" specification ($N = 69,223$).

Parameter	Mode	Mean	Std.	95% HDI
$\hat{\beta}$ Intercept	0.586	0.588	0.022	[0.546, 0.633]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead]	0.047	0.047	0.023	[0.002, 0.093]
$\hat{\beta}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	-0.260	-0.259	0.160	[-0.569, 0.053]
$\hat{\beta}$ Account ID (Millions)^2	-0.010	0.029	0.343	[-0.64, 0.699]
$\hat{\beta}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	0.351	0.352	0.003	[0.346, 0.357]
$\hat{\beta}$ Estimated Annual Consumption (kWh) [Z-Score]	0.374	0.373	0.014	[0.347, 0.4]
$\hat{\beta}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.010	-0.010	0.003	[-0.015, -0.005]
$\hat{\beta}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.011	-0.011	0.006	[-0.024, 0.001]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)	0.550	0.553	0.113	[0.341, 0.783]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)^2	-0.046	-0.044	0.041	[-0.124, 0.034]
$\hat{\beta}$ Business Entity [Ref Non-Business]	-0.101	-0.101	0.009	[-0.119, -0.084]
$\hat{\beta}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.008	0.008	0.006	[-0.003, 0.019]
$\hat{\beta}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.004	0.003	0.010	[-0.016, 0.023]
$\hat{\beta}$ DNO Region — North [Ref Midlands & South]	0.192	0.209	0.180	[-0.133, 0.571]
$\hat{\beta}$ DNO Region — Scotland [Ref Midlands & South]	0.034	0.075	0.349	[-0.601, 0.762]
$\hat{\sigma}_y$	0.636	0.636	0.002	[0.633, 0.639]

Table 5: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response Y = "Total Metered Consumption (kWh)") fitted using the "Extended (MSE-Optimal Bandwidth; Flexible Scale)" specification ($N = 69,223$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}$ Intercept	0.607	0.607	0.017	[0.573, 0.64]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead]	0.044	0.040	0.018	[0.005, 0.075]
$\widehat{\beta}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	-0.175	-0.165	0.138	[-0.445, 0.097]
$\widehat{\beta}$ Account ID (Millions)^2	0.011	-0.016	0.334	[-0.676, 0.63]
$\widehat{\beta}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	0.389	0.389	0.004	[0.381, 0.396]
$\widehat{\beta}$ Estimated Annual Consumption (kWh) [Z-Score]	0.428	0.428	0.016	[0.396, 0.459]
$\widehat{\beta}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.007	-0.007	0.002	[-0.01, -0.003]
$\widehat{\beta}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.005	-0.005	0.005	[-0.014, 0.005]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)	0.133	0.140	0.111	[-0.078, 0.355]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)^2	-0.041	-0.042	0.028	[-0.098, 0.013]
$\widehat{\beta}$ Business Entity [Ref Non-Business]	-0.086	-0.086	0.007	[-0.1, -0.072]
$\widehat{\beta}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.002	0.001	0.004	[-0.007, 0.01]
$\widehat{\beta}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.005	0.005	0.007	[-0.009, 0.019]
$\widehat{\beta}$ DNO Region — North [Ref Midlands & South]	0.124	0.125	0.151	[-0.165, 0.425]
$\widehat{\beta}$ DNO Region — Scotland [Ref Midlands & South]	0.146	0.091	0.339	[-0.593, 0.727]
$\widehat{\phi}_{\sigma_y}$ Intercept	-0.709	-0.711	0.024	[-0.759, -0.663]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead]	0.100	0.097	0.026	[0.046, 0.147]
$\widehat{\phi}_{\sigma_y}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	-0.308	-0.337	0.168	[-0.674, -0.016]
$\widehat{\phi}_{\sigma_y}$ Account ID (Millions)^2	0.097	0.101	0.345	[-0.58, 0.778]
$\widehat{\phi}_{\sigma_y}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	0.348	0.348	0.004	[0.341, 0.356]
$\widehat{\phi}_{\sigma_y}$ Estimated Annual Consumption (kWh) [Z-Score]	0.748	0.749	0.016	[0.717, 0.781]
$\widehat{\phi}_{\sigma_y}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.001	-0.001	0.003	[-0.007, 0.005]
$\widehat{\phi}_{\sigma_y}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.011	-0.012	0.007	[-0.025, 0.002]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)	0.174	0.187	0.127	[-0.061, 0.433]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)^2	-0.009	-0.005	0.045	[-0.095, 0.082]
$\widehat{\phi}_{\sigma_y}$ Business Entity [Ref Non-Business]	-0.010	-0.010	0.010	[-0.028, 0.009]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.009	0.010	0.006	[-0.002, 0.022]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	-0.006	-0.008	0.011	[-0.029, 0.014]
$\widehat{\phi}_{\sigma_y}$ DNO Region — North [Ref Midlands & South]	0.159	0.169	0.190	[-0.199, 0.538]
$\widehat{\phi}_{\sigma_y}$ DNO Region — Scotland [Ref Midlands & South]	0.199	0.233	0.353	[-0.472, 0.919]

Table 6: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response Y = "Total Metered Consumption (kWh)") fitted using the "Extended (MSE-Optimal Bandwidth $\times$ 1.5; Flexible Scale)" specification ($N = 96,558$).

Parameter	Mode	Mean	Std.	95% HDI
$\hat{\beta}$ Intercept	0.588	0.591	0.031	[0.529, 0.65]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead]	0.068	0.062	0.031	[0.002, 0.122]
$\hat{\beta}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	-0.353	-0.337	0.235	[-0.793, 0.129]
$\hat{\beta}$ Account ID (Millions)^2	0.132	0.072	0.348	[-0.605, 0.75]
$\hat{\beta}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	0.383	0.383	0.003	[0.376, 0.389]
$\hat{\beta}$ Estimated Annual Consumption (kWh) [Z-Score]	0.349	0.349	0.011	[0.328, 0.371]
$\hat{\beta}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.006	-0.006	0.002	[-0.009, -0.002]
$\hat{\beta}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.006	-0.006	0.005	[-0.016, 0.003]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)	0.123	0.124	0.105	[-0.075, 0.335]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)^2	-0.005	-0.006	0.030	[-0.064, 0.052]
$\hat{\beta}$ Business Entity [Ref Non-Business]	-0.088	-0.088	0.006	[-0.101, -0.076]
$\hat{\beta}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.004	0.003	0.003	[-0.003, 0.01]
$\hat{\beta}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.003	0.003	0.006	[-0.01, 0.015]
$\hat{\beta}$ DNO Region — North [Ref Midlands & South]	0.342	0.330	0.236	[-0.159, 0.768]
$\hat{\beta}$ DNO Region — Scotland [Ref Midlands & South]	-0.134	-0.103	0.349	[-0.782, 0.576]
$\hat{\phi}_{\sigma_y}$ Intercept	-0.751	-0.754	0.039	[-0.831, -0.679]
$\hat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead]	0.149	0.146	0.039	[0.073, 0.225]
$\hat{\phi}_{\sigma_y}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	-0.637	-0.643	0.283	[-1.207, -0.097]
$\hat{\phi}_{\sigma_y}$ Account ID (Millions)^2	0.198	0.144	0.351	[-0.539, 0.843]
$\hat{\phi}_{\sigma_y}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	0.356	0.355	0.003	[0.349, 0.362]
$\hat{\phi}_{\sigma_y}$ Estimated Annual Consumption (kWh) [Z-Score]	0.549	0.549	0.012	[0.525, 0.572]
$\hat{\phi}_{\sigma_y}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	0.000	0.000	0.003	[-0.005, 0.005]
$\hat{\phi}_{\sigma_y}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.012	-0.012	0.006	[-0.025, 0.0]
$\hat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)	0.123	0.123	0.118	[-0.107, 0.351]
$\hat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)^2	0.124	0.124	0.042	[0.042, 0.208]
$\hat{\phi}_{\sigma_y}$ Business Entity [Ref Non-Business]	-0.007	-0.006	0.008	[-0.022, 0.011]
$\hat{\phi}_{\sigma_y}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.015	0.015	0.005	[0.005, 0.025]
$\hat{\phi}_{\sigma_y}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.007	0.008	0.010	[-0.012, 0.027]
$\hat{\phi}_{\sigma_y}$ DNO Region — North [Ref Midlands & South]	0.612	0.590	0.285	[0.037, 1.145]
$\hat{\phi}_{\sigma_y}$ DNO Region — Scotland [Ref Midlands & South]	-0.260	-0.199	0.351	[-0.888, 0.488]

Table 7: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response Y = "Total Metered Consumption (kWh)") fitted using the "Extended (MSE-Optimal Bandwidth $\times$ 2; Flexible Scale)" specification ($N = 123,373$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}$ Intercept	0.661	0.659	0.035	[0.588, 0.725]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead]	0.013	0.016	0.034	[-0.051, 0.081]
$\widehat{\beta}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	0.150	0.132	0.278	[-0.421, 0.671]
$\widehat{\beta}$ Account ID (Millions)^2	0.181	0.141	0.316	[-0.472, 0.756]
$\widehat{\beta}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	2.896	2.896	0.021	[2.852, 2.937]
$\widehat{\beta}$ Estimated Annual Consumption (kWh) [Z-Score]	0.298	0.298	0.009	[0.281, 0.316]
$\widehat{\beta}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.007	-0.007	0.001	[-0.01, -0.004]
$\widehat{\beta}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.009	-0.010	0.005	[-0.019, 0.0]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)	0.070	0.080	0.087	[-0.085, 0.257]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)^2	-0.002	-0.006	0.027	[-0.061, 0.045]
$\widehat{\beta}$ Business Entity [Ref Non-Business]	-0.086	-0.087	0.005	[-0.097, -0.076]
$\widehat{\beta}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.002	0.002	0.003	[-0.004, 0.009]
$\widehat{\beta}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.004	0.004	0.006	[-0.008, 0.015]
$\widehat{\beta}$ DNO Region — North [Ref Midlands & South]	-0.206	-0.189	0.278	[-0.737, 0.357]
$\widehat{\beta}$ DNO Region — Scotland [Ref Midlands & South]	-0.140	-0.113	0.316	[-0.731, 0.496]
$\widehat{\phi}_{\sigma_y}$ Intercept	-0.583	-0.585	0.043	[-0.668, -0.501]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead]	0.106	0.103	0.042	[0.021, 0.184]
$\widehat{\phi}_{\sigma_y}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	-0.149	-0.123	0.346	[-0.806, 0.55]
$\widehat{\phi}_{\sigma_y}$ Account ID (Millions)^2	0.092	0.062	0.348	[-0.603, 0.76]
$\widehat{\phi}_{\sigma_y}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	2.451	2.453	0.021	[2.413, 2.494]
$\widehat{\phi}_{\sigma_y}$ Estimated Annual Consumption (kWh) [Z-Score]	0.504	0.504	0.009	[0.485, 0.522]
$\widehat{\phi}_{\sigma_y}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	0.001	0.001	0.002	[-0.003, 0.006]
$\widehat{\phi}_{\sigma_y}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.062	-0.062	0.007	[-0.076, -0.049]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)	0.028	0.057	0.108	[-0.146, 0.276]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)^2	0.128	0.125	0.037	[0.053, 0.198]
$\widehat{\phi}_{\sigma_y}$ Business Entity [Ref Non-Business]	-0.007	-0.007	0.007	[-0.022, 0.006]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.005	0.006	0.005	[-0.004, 0.015]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.002	0.002	0.009	[-0.014, 0.019]
$\widehat{\phi}_{\sigma_y}$ DNO Region — North [Ref Midlands & South]	-0.278	-0.289	0.346	[-0.988, 0.37]
$\widehat{\phi}_{\sigma_y}$ DNO Region — Scotland [Ref Midlands & South]	0.115	0.145	0.348	[-0.527, 0.836]

Table 8: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response Y = "Total Metered Consumption (kWh)") fitted using the "Extended (MSE-Optimal Bandwidth Left Bound Only; Flexible Scale)" specification ($N = 350,834$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}$ Intercept	0.626	0.625	0.005	[0.615, 0.635]
$\widehat{\beta}$ Intraday Opt-in Notice (07:59:59, 09:00:00] [Ref Day-ahead]	0.019	0.019	0.006	[0.008, 0.03]
$\widehat{\beta}$ Intraday Opt-in Notice (09:00:00, 10:00:00] [Ref Day-ahead]	0.017	0.017	0.006	[0.006, 0.028]
$\widehat{\beta}$ Intraday Opt-in Notice (10:00:00, 11:00:00] [Ref Day-ahead]	0.024	0.024	0.005	[0.014, 0.033]
$\widehat{\beta}$ Intraday Opt-in Notice (11:00:00, 12:00:00] [Ref Day-ahead]	0.022	0.022	0.005	[0.012, 0.031]
$\widehat{\beta}$ Intraday Opt-in Notice (12:00:00, 13:00:00] [Ref Day-ahead]	0.023	0.023	0.006	[0.012, 0.035]
$\widehat{\beta}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	5.181	5.182	0.022	[5.139, 5.224]
$\widehat{\beta}$ Estimated Annual Consumption (kWh) [Z-Score]	0.170	0.170	0.003	[0.164, 0.177]
$\widehat{\beta}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.007	-0.007	0.001	[-0.008, -0.005]
$\widehat{\beta}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	0.006	0.006	0.001	[0.003, 0.009]
$\widehat{\beta}$ Business Entity [Ref Non-Business]	-0.039	-0.039	0.026	[-0.09, 0.012]
$\widehat{\beta}$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.008	-0.009	0.012	[-0.034, 0.014]
$\widehat{\beta}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	-0.076	-0.076	0.003	[-0.083, -0.07]
$\widehat{\beta}$ DNO Region — North [Ref Midlands & South]	0.006	0.006	0.002	[0.003, 0.01]
$\widehat{\beta}$ DNO Region — Scotland [Ref Midlands & South]	0.002	0.002	0.004	[-0.005, 0.01]
$\widehat{\phi}_{\sigma_y}$ Intercept	-0.628	-0.629	0.008	[-0.644, -0.615]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice (07:59:59, 09:00:00] [Ref Day-ahead]	0.045	0.043	0.009	[0.026, 0.06]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice (09:00:00, 10:00:00] [Ref Day-ahead]	0.029	0.029	0.009	[0.012, 0.046]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice (10:00:00, 11:00:00] [Ref Day-ahead]	0.028	0.029	0.008	[0.014, 0.044]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice (11:00:00, 12:00:00] [Ref Day-ahead]	-0.015	-0.015	0.008	[-0.029, 0.0]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice (12:00:00, 13:00:00] [Ref Day-ahead]	-0.049	-0.049	0.009	[-0.067, -0.031]
$\widehat{\phi}_{\sigma_y}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	4.913	4.917	0.023	[4.872, 4.961]
$\widehat{\phi}_{\sigma_y}$ Estimated Annual Consumption (kWh) [Z-Score]	0.258	0.258	0.004	[0.251, 0.265]
$\widehat{\phi}_{\sigma_y}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.004	-0.004	0.001	[-0.006, -0.001]
$\widehat{\phi}_{\sigma_y}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.017	-0.016	0.002	[-0.021, -0.012]
$\widehat{\phi}_{\sigma_y}$ Business Entity [Ref Non-Business]	-0.187	-0.176	0.045	[-0.265, -0.089]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.081	0.080	0.018	[0.045, 0.116]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	-0.013	-0.013	0.005	[-0.022, -0.004]
$\widehat{\phi}_{\sigma_y}$ DNO Region — North [Ref Midlands & South]	0.011	0.011	0.003	[0.005, 0.016]
$\widehat{\phi}_{\sigma_y}$ DNO Region — Scotland [Ref Midlands & South]	0.037	0.036	0.006	[0.026, 0.047]

Table 9: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response Y = "Probability of Opt-in") fitted using the "Baseline (MSE-Optimal Bandwidth)" specification ($N = 99,678$).

Parameter	Mode	Mean	Std.	95% HDI
$\hat{\beta}$ Intercept	0.563	0.563	0.007	[0.55, 0.576]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead]	0.001	0.002	0.008	[-0.013, 0.018]
$\hat{\beta}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	-0.050	-0.053	0.021	[-0.095, -0.011]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)	-0.027	-0.027	0.027	[-0.079, 0.026]
$\hat{\sigma}_y$	0.496	0.496	0.001	[0.494, 0.498]

Table 10: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response Y = "Probability of Opt-in") fitted using the "Extended (MSE-Optimal Bandwidth)" specification ($N = 88,501$).

Parameter	Mode	Mean	Std.	95% HDI
$\hat{\beta}$ Intercept	0.572	0.572	0.015	[0.543, 0.601]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead]	-0.022	-0.021	0.016	[-0.052, 0.01]
$\hat{\beta}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	0.133	0.144	0.109	[-0.067, 0.356]
$\hat{\beta}$ Account ID (Millions)^2	0.201	0.221	0.180	[-0.136, 0.568]
$\hat{\beta}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	0.002	0.002	0.002	[-0.002, 0.006]
$\hat{\beta}$ Estimated Annual Consumption (kWh) [Z-Score]	-0.050	-0.051	0.008	[-0.067, -0.034]
$\hat{\beta}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	0.017	0.017	0.002	[0.014, 0.021]
$\hat{\beta}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	0.016	0.016	0.005	[0.005, 0.026]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)	0.072	0.077	0.072	[-0.066, 0.218]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)^2	0.046	0.046	0.028	[-0.008, 0.103]
$\hat{\beta}$ Business Entity [Ref Non-Business]	0.108	0.108	0.006	[0.096, 0.12]
$\hat{\beta}$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.005	-0.005	0.004	[-0.013, 0.002]
$\hat{\beta}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.023	0.023	0.007	[0.009, 0.037]
$\hat{\beta}$ DNO Region — North [Ref Midlands & South]	-0.180	-0.173	0.130	[-0.429, 0.079]
$\hat{\beta}$ DNO Region — Scotland [Ref Midlands & South]	-0.233	-0.242	0.220	[-0.677, 0.179]
$\hat{\sigma}_y$	0.495	0.495	0.001	[0.493, 0.497]

Table 11: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response Y = "Probability of Opt-in") fitted using the "Extended (MSE-Optimal Bandwidth; Flexible Scale)" specification ($N = 88,501$).

Parameter	Mode	Mean	Std.	95% HDI
$\hat{\beta}$ Intercept	0.571	0.572	0.015	[0.543, 0.6]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead]	-0.020	-0.021	0.016	[-0.052, 0.01]
$\hat{\beta}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	0.147	0.143	0.108	[-0.064, 0.356]
$\hat{\beta}$ Account ID (Millions)^2	0.240	0.219	0.179	[-0.127, 0.571]
$\hat{\beta}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	0.001	0.002	0.002	[-0.002, 0.006]
$\hat{\beta}$ Estimated Annual Consumption (kWh) [Z-Score]	-0.051	-0.052	0.009	[-0.069, -0.035]
$\hat{\beta}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	0.017	0.017	0.002	[0.014, 0.021]
$\hat{\beta}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	0.016	0.016	0.005	[0.005, 0.026]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)	0.063	0.074	0.076	[-0.074, 0.226]
$\hat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)^2	0.046	0.046	0.028	[-0.01, 0.1]
$\hat{\beta}$ Business Entity [Ref Non-Business]	0.107	0.108	0.006	[0.096, 0.119]
$\hat{\beta}$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.005	-0.005	0.004	[-0.012, 0.003]
$\hat{\beta}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.024	0.023	0.007	[0.01, 0.037]
$\hat{\beta}$ DNO Region — North [Ref Midlands & South]	-0.181	-0.171	0.129	[-0.426, 0.082]
$\hat{\beta}$ DNO Region — Scotland [Ref Midlands & South]	-0.222	-0.242	0.220	[-0.679, 0.188]
$\hat{\phi}_{\sigma_y}$ Intercept	-0.703	-0.702	0.019	[-0.737, -0.664]
$\hat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead]	0.004	0.003	0.020	[-0.038, 0.042]
$\hat{\phi}_{\sigma_y}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	0.004	-0.014	0.134	[-0.272, 0.254]
$\hat{\phi}_{\sigma_y}$ Account ID (Millions)^2	-0.016	-0.036	0.222	[-0.472, 0.399]
$\hat{\phi}_{\sigma_y}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	-0.000	0.000	0.003	[-0.005, 0.006]
$\hat{\phi}_{\sigma_y}$ Estimated Annual Consumption (kWh) [Z-Score]	0.018	0.017	0.012	[-0.007, 0.04]
$\hat{\phi}_{\sigma_y}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.003	-0.003	0.003	[-0.008, 0.002]
$\hat{\phi}_{\sigma_y}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.002	-0.002	0.007	[-0.016, 0.012]
$\hat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)	-0.010	0.015	0.106	[-0.189, 0.226]
$\hat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)^2	-0.008	-0.007	0.040	[-0.086, 0.07]
$\hat{\phi}_{\sigma_y}$ Business Entity [Ref Non-Business]	-0.042	-0.043	0.009	[-0.059, -0.025]
$\hat{\phi}_{\sigma_y}$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.001	-0.000	0.006	[-0.011, 0.01]
$\hat{\phi}_{\sigma_y}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	-0.006	-0.004	0.010	[-0.024, 0.015]
$\hat{\phi}_{\sigma_y}$ DNO Region — North [Ref Midlands & South]	0.019	0.023	0.169	[-0.301, 0.36]
$\hat{\phi}_{\sigma_y}$ DNO Region — Scotland [Ref Midlands & South]	0.007	0.031	0.278	[-0.504, 0.579]

Table 12: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response Y = "Probability of Opt-in") fitted using the "Extended (MSE-Optimal Bandwidth × 1.5; Flexible Scale)" specification ($N = 104,509$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}$ Intercept	0.557	0.557	0.016	[0.524, 0.588]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead]	-0.017	-0.017	0.017	[-0.05, 0.016]
$\widehat{\beta}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	0.081	0.088	0.136	[-0.176, 0.359]
$\widehat{\beta}$ Account ID (Millions)^2	0.189	0.143	0.302	[-0.43, 0.747]
$\widehat{\beta}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	0.002	0.002	0.002	[-0.002, 0.006]
$\widehat{\beta}$ Estimated Annual Consumption (kWh) [Z-Score]	-0.050	-0.049	0.007	[-0.063, -0.035]
$\widehat{\beta}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	0.019	0.019	0.002	[0.016, 0.022]
$\widehat{\beta}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	0.011	0.011	0.006	[-0.0, 0.022]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)	0.068	0.059	0.077	[-0.086, 0.216]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)^2	0.041	0.040	0.027	[-0.012, 0.092]
$\widehat{\beta}$ Business Entity [Ref Non-Business]	0.114	0.114	0.006	[0.103, 0.124]
$\widehat{\beta}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.000	0.000	0.004	[-0.007, 0.007]
$\widehat{\beta}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.024	0.025	0.006	[0.012, 0.037]
$\widehat{\beta}$ DNO Region — North [Ref Midlands & South]	-0.163	-0.155	0.141	[-0.431, 0.118]
$\widehat{\beta}$ DNO Region — Scotland [Ref Midlands & South]	-0.080	-0.035	0.305	[-0.65, 0.539]
$\widehat{\phi}_{\sigma_y}$ Intercept	-0.699	-0.700	0.020	[-0.738, -0.661]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead]	0.002	0.002	0.021	[-0.038, 0.043]
$\widehat{\phi}_{\sigma_y}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	-0.011	-0.003	0.148	[-0.292, 0.29]
$\widehat{\phi}_{\sigma_y}$ Account ID (Millions)^2	-0.023	-0.011	0.314	[-0.645, 0.586]
$\widehat{\phi}_{\sigma_y}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	0.000	-0.000	0.003	[-0.005, 0.005]
$\widehat{\phi}_{\sigma_y}$ Estimated Annual Consumption (kWh) [Z-Score]	0.015	0.015	0.010	[-0.005, 0.035]
$\widehat{\phi}_{\sigma_y}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.002	-0.002	0.002	[-0.007, 0.002]
$\widehat{\phi}_{\sigma_y}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.001	-0.002	0.008	[-0.017, 0.013]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)	0.012	0.029	0.105	[-0.172, 0.237]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] × Account ID (Millions)^2	-0.008	-0.004	0.038	[-0.08, 0.07]
$\widehat{\phi}_{\sigma_y}$ Business Entity [Ref Non-Business]	-0.044	-0.043	0.008	[-0.059, -0.027]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.001	-0.000	0.005	[-0.01, 0.01]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	-0.003	-0.003	0.009	[-0.021, 0.015]
$\widehat{\phi}_{\sigma_y}$ DNO Region — North [Ref Midlands & South]	0.012	0.017	0.157	[-0.281, 0.33]
$\widehat{\phi}_{\sigma_y}$ DNO Region — Scotland [Ref Midlands & South]	0.001	-0.013	0.318	[-0.645, 0.605]

Table 13: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response Y = "Probability of Opt-in") fitted using the "Extended (MSE-Optimal Bandwidth $\times$ 2; Flexible Scale)" specification ($N = 125,236$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}$ Intercept	0.559	0.562	0.018	[0.528, 0.597]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead]	-0.032	-0.034	0.017	[-0.068, -0.001]
$\widehat{\beta}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	0.171	0.149	0.122	[-0.084, 0.393]
$\widehat{\beta}$ Account ID (Millions)^2	0.182	0.167	0.278	[-0.395, 0.692]
$\widehat{\beta}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	0.001	0.001	0.002	[-0.002, 0.005]
$\widehat{\beta}$ Estimated Annual Consumption (kWh) [Z-Score]	-0.044	-0.044	0.006	[-0.055, -0.031]
$\widehat{\beta}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	0.019	0.019	0.002	[0.016, 0.022]
$\widehat{\beta}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	0.009	0.008	0.005	[-0.002, 0.018]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)	0.047	0.051	0.073	[-0.09, 0.196]
$\widehat{\beta}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)^2	0.050	0.049	0.025	[-0.0, 0.098]
$\widehat{\beta}$ Business Entity [Ref Non-Business]	0.115	0.115	0.005	[0.106, 0.125]
$\widehat{\beta}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.004	0.004	0.003	[-0.002, 0.01]
$\widehat{\beta}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.028	0.029	0.006	[0.018, 0.041]
$\widehat{\beta}$ DNO Region — North [Ref Midlands & South]	-0.126	-0.147	0.123	[-0.388, 0.094]
$\widehat{\beta}$ DNO Region — Scotland [Ref Midlands & South]	-0.195	-0.181	0.279	[-0.712, 0.377]
$\widehat{\phi}_{\sigma_y}$ Intercept	-0.705	-0.703	0.025	[-0.75, -0.653]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead]	0.011	0.007	0.024	[-0.041, 0.054]
$\widehat{\phi}_{\sigma_y}$ Account ID (Millions) [Relative to Cutoff 2,454,839]	-0.050	-0.025	0.160	[-0.336, 0.292]
$\widehat{\phi}_{\sigma_y}$ Account ID (Millions)^2	0.008	-0.052	0.323	[-0.688, 0.571]
$\widehat{\phi}_{\sigma_y}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	-0.000	-0.000	0.002	[-0.005, 0.005]
$\widehat{\phi}_{\sigma_y}$ Estimated Annual Consumption (kWh) [Z-Score]	0.014	0.013	0.009	[-0.003, 0.03]
$\widehat{\phi}_{\sigma_y}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.002	-0.002	0.002	[-0.006, 0.002]
$\widehat{\phi}_{\sigma_y}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.000	-0.001	0.007	[-0.015, 0.014]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)	0.004	0.023	0.102	[-0.177, 0.221]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice [Ref Day-ahead] $\times$ Account ID (Millions)^2	-0.009	-0.009	0.037	[-0.079, 0.065]
$\widehat{\phi}_{\sigma_y}$ Business Entity [Ref Non-Business]	-0.040	-0.040	0.007	[-0.054, -0.025]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.001	-0.001	0.005	[-0.01, 0.008]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	-0.004	-0.003	0.009	[-0.019, 0.014]
$\widehat{\phi}_{\sigma_y}$ DNO Region — North [Ref Midlands & South]	0.036	0.029	0.162	[-0.29, 0.344]
$\widehat{\phi}_{\sigma_y}$ DNO Region — Scotland [Ref Midlands & South]	0.092	0.050	0.325	[-0.577, 0.69]

Table 14: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a regression discontinuity model (Response Y = "Probability of Opt-in") fitted using the "Extended (MSE-Optimal Bandwidth Left Bound Only; Flexible Scale)" specification ($N = 378,091$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}$ Intercept	0.574	0.575	0.004	[0.566, 0.583]
$\widehat{\beta}$ Intraday Opt-in Notice (07:59:59, 09:00:00] [Ref Day-ahead]	-0.020	-0.021	0.005	[-0.031, -0.011]
$\widehat{\beta}$ Intraday Opt-in Notice (09:00:00, 10:00:00] [Ref Day-ahead]	-0.041	-0.042	0.005	[-0.052, -0.032]
$\widehat{\beta}$ Intraday Opt-in Notice (10:00:00, 11:00:00] [Ref Day-ahead]	-0.047	-0.047	0.004	[-0.056, -0.039]
$\widehat{\beta}$ Intraday Opt-in Notice (11:00:00, 12:00:00] [Ref Day-ahead]	-0.046	-0.047	0.004	[-0.055, -0.038]
$\widehat{\beta}$ Intraday Opt-in Notice (12:00:00, 13:00:00] [Ref Day-ahead]	-0.040	-0.039	0.006	[-0.05, -0.028]
$\widehat{\beta}$ Estimated Annual Consumption (kWh) [Z-Score]	-0.021	-0.021	0.002	[-0.025, -0.018]
$\widehat{\beta}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	0.018	0.018	0.001	[0.016, 0.02]
$\widehat{\beta}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	-0.005	-0.005	0.002	[-0.008, -0.001]
$\widehat{\beta}$ Business Entity [Ref Non-Business]	0.013	0.009	0.031	[-0.05, 0.071]
$\widehat{\beta}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.052	0.051	0.013	[0.025, 0.075]
$\widehat{\beta}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.105	0.105	0.003	[0.099, 0.111]
$\widehat{\beta}$ DNO Region — North [Ref Midlands & South]	0.002	0.002	0.002	[-0.002, 0.005]
$\widehat{\beta}$ DNO Region — Scotland [Ref Midlands & South]	0.026	0.025	0.004	[0.018, 0.033]
$\widehat{\phi}_{\sigma_y}$ Intercept	-0.706	-0.706	0.006	[-0.718, -0.694]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice (07:59:59, 09:00:00] [Ref Day-ahead]	0.007	0.006	0.007	[-0.008, 0.02]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice (09:00:00, 10:00:00] [Ref Day-ahead]	0.010	0.009	0.007	[-0.005, 0.024]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice (10:00:00, 11:00:00] [Ref Day-ahead]	0.010	0.010	0.006	[-0.002, 0.023]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice (11:00:00, 12:00:00] [Ref Day-ahead]	0.011	0.011	0.006	[-0.002, 0.023]
$\widehat{\phi}_{\sigma_y}$ Intraday Opt-in Notice (12:00:00, 13:00:00] [Ref Day-ahead]	0.010	0.010	0.008	[-0.006, 0.026]
$\widehat{\phi}_{\sigma_y}$ Estimated Annual Consumption (kWh) [Z-Score]	0.006	0.006	0.003	[0.0, 0.011]
$\widehat{\phi}_{\sigma_y}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.003	-0.003	0.001	[-0.005, -0.0]
$\widehat{\phi}_{\sigma_y}$ Tenure (Years Prior to 1st SS10 Opt-in Notice) [Z-Score]	0.001	0.001	0.003	[-0.004, 0.006]
$\widehat{\phi}_{\sigma_y}$ Business Entity [Ref Non-Business]	-0.002	-0.001	0.044	[-0.087, 0.083]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.012	-0.011	0.018	[-0.046, 0.023]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	-0.039	-0.038	0.004	[-0.047, -0.03]
$\widehat{\phi}_{\sigma_y}$ DNO Region — North [Ref Midlands & South]	0.000	-0.000	0.003	[-0.005, 0.005]
$\widehat{\phi}_{\sigma_y}$ DNO Region — Scotland [Ref Midlands & South]	-0.004	-0.004	0.005	[-0.014, 0.007]

All Study 2 (RCT) Results in Tabular Format

Below, we include all numeric results from each model fitted for Study 2.

Table 15: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a simultaneous-equation regression model (2nd Equation Response Y = "Total Metered Consumption (kWh)"; 1st Equation Response T = "Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]") fitted using the "Baseline" specification ($N = 638,242$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}_y$ Intercept	0.655	0.655	0.001	[0.653, 0.657]
$\widehat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]	-0.021	-0.021	0.006	[-0.032, -0.01]
$\widehat{\beta}_y$ Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]	-0.033	-0.030	0.025	[-0.08, 0.018]
$\widehat{\beta}_T$ Intercept	-0.000	0.000	0.000	[-0.0, 0.0]
$\widehat{\beta}_T$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]	0.233	0.233	0.001	[0.232, 0.234]
$\widehat{\sigma}_y$	0.784	0.784	0.001	[0.783, 0.785]
$\widehat{\sigma}_T$	0.073	0.073	0.000	[0.073, 0.073]
$\widehat{\rho}_{\epsilon_{i,y},\epsilon_{i,T}}$	-0.002	-0.002	0.003	[-0.007, 0.003]

Table 16: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a simultaneous-equation regression model (2nd Equation Response Y = "Total Metered Consumption (kWh)"; 1st Equation Response T = "Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]") fitted using the "Extended" specification ($N = 547,055$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}_y$ Intercept	0.664	0.664	0.001	[0.661, 0.666]
$\widehat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]	-0.013	-0.013	0.005	[-0.022, -0.003]
$\widehat{\beta}_y$ Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]	-0.025	-0.027	0.022	[-0.069, 0.016]
$\widehat{\beta}_y$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	1.452	1.452	0.004	[1.445, 1.458]
$\widehat{\beta}_y$ Estimated Annual Consumption (kWh) [Z-Score]	0.137	0.137	0.002	[0.133, 0.141]
$\widehat{\beta}_y$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.002	-0.002	0.001	[-0.004, -0.001]
$\widehat{\beta}_y$ Tenure (Years Prior to 1st SS12 Opt-in Notice) [Z-Score]	-0.009	-0.009	0.001	[-0.011, -0.007]
$\widehat{\beta}_y$ Business Entity [Ref Non-Business]	0.141	0.142	0.035	[0.073, 0.208]
$\widehat{\beta}_y$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.044	-0.044	0.004	[-0.051, -0.036]
$\widehat{\beta}_y$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	-0.075	-0.075	0.003	[-0.082, -0.069]
$\widehat{\beta}_y$ DNO Region — North [Ref Midlands & South]	0.008	0.008	0.002	[0.005, 0.012]
$\widehat{\beta}_y$ DNO Region — Scotland [Ref Midlands & South]	0.010	0.011	0.004	[0.003, 0.018]
$\widehat{\beta}_T$ Intercept	-0.000	-0.000	0.000	[-0.0, 0.0]
$\widehat{\beta}_T$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]	0.226	0.225	0.001	[0.224, 0.227]
$\widehat{\beta}_T$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	-0.000	-0.000	0.000	[-0.001, 0.001]
$\widehat{\beta}_T$ Estimated Annual Consumption (kWh) [Z-Score]	-0.000	-0.000	0.000	[-0.001, 0.0]
$\widehat{\beta}_T$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.000	-0.000	0.000	[-0.0, -0.0]
$\widehat{\beta}_T$ Tenure (Years Prior to 1st SS12 Opt-in Notice) [Z-Score]	0.001	0.001	0.000	[0.001, 0.001]
$\widehat{\beta}_T$ Business Entity [Ref Non-Business]	0.002	0.002	0.004	[-0.006, 0.01]
$\widehat{\beta}_T$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.005	-0.005	0.000	[-0.006, -0.004]
$\widehat{\beta}_T$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.001	0.001	0.000	[0.001, 0.002]
$\widehat{\beta}_T$ DNO Region — North [Ref Midlands & South]	0.000	0.000	0.000	[-0.0, 0.001]
$\widehat{\beta}_T$ DNO Region — Scotland [Ref Midlands & South]	0.001	0.001	0.000	[-0.0, 0.002]
$\widehat{\sigma}_y$	0.617	0.616	0.001	[0.615, 0.618]
$\widehat{\sigma}_T$	0.072	0.072	0.000	[0.072, 0.072]
$\widehat{\rho}_{\epsilon_{i,y},\epsilon_{i,T}}$	-0.001	-0.001	0.003	[-0.007, 0.004]

Table 17: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a single-equation regression model (Response Y = "Total Metered Consumption (kWh)") fitted using the "ITT (Extended Variant)" specification ($N = 547,055$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}_y$ Intercept	0.663	0.664	0.001	[0.661, 0.666]
$\widehat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]	-0.013	-0.013	0.005	[-0.023, -0.003]
$\widehat{\beta}_y$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]	-0.006	-0.006	0.005	[-0.016, 0.004]
$\widehat{\beta}_y$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	1.452	1.452	0.004	[1.445, 1.459]
$\widehat{\beta}_y$ Estimated Annual Consumption (kWh) [Z-Score]	0.137	0.137	0.002	[0.133, 0.141]
$\widehat{\beta}_y$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.002	-0.002	0.001	[-0.004, -0.001]
$\widehat{\beta}_y$ Tenure (Years Prior to 1st SS12 Opt-in Notice) [Z-Score]	-0.009	-0.009	0.001	[-0.011, -0.008]
$\widehat{\beta}_y$ Business Entity [Ref Non-Business]	0.140	0.142	0.035	[0.073, 0.211]
$\widehat{\beta}_y$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.044	-0.044	0.004	[-0.051, -0.036]
$\widehat{\beta}_y$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	-0.075	-0.075	0.003	[-0.082, -0.069]
$\widehat{\beta}_y$ DNO Region — North [Ref Midlands & South]	0.008	0.008	0.002	[0.005, 0.012]
$\widehat{\beta}_y$ DNO Region — Scotland [Ref Midlands & South]	0.011	0.010	0.004	[0.003, 0.018]
$\widehat{\sigma}_y$	0.617	0.617	0.001	[0.615, 0.618]

Table 18: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a single-equation regression model (Response $Y =$ "Total Metered Consumption (kWh)") fitted using the "ITT (Extended Variant; Flexible Scale)" specification ($N = 547,055$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}_y$ Intercept	0.668	0.668	0.001	[0.666, 0.67]
$\widehat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]	-0.006	-0.006	0.004	[-0.013, 0.001]
$\widehat{\beta}_y$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]	-0.007	-0.007	0.003	[-0.014, -0.0]
$\widehat{\beta}_y$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	1.513	1.513	0.005	[1.504, 1.523]
$\widehat{\beta}_y$ Estimated Annual Consumption (kWh) [Z-Score]	0.140	0.140	0.002	[0.136, 0.145]
$\widehat{\beta}_y$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.004	-0.004	0.001	[-0.005, -0.002]
$\widehat{\beta}_y$ Tenure (Years Prior to 1st SS12 Opt-in Notice) [Z-Score]	-0.007	-0.007	0.001	[-0.009, -0.006]
$\widehat{\beta}_y$ Business Entity [Ref Non-Business]	-0.010	-0.009	0.035	[-0.076, 0.059]
$\widehat{\beta}_y$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.033	-0.033	0.003	[-0.038, -0.028]
$\widehat{\beta}_y$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	-0.061	-0.061	0.003	[-0.066, -0.055]
$\widehat{\beta}_y$ DNO Region — North [Ref Midlands & South]	0.003	0.002	0.001	[-0.0, 0.005]
$\widehat{\beta}_y$ DNO Region — Scotland [Ref Midlands & South]	0.004	0.004	0.003	[-0.002, 0.009]
$\widehat{\phi}_{\sigma_y}$ Intercept	-0.715	-0.715	0.001	[-0.717, -0.713]
$\widehat{\phi}_{\sigma_y}$ Intraday Notice + Day-ahead Email [Ref Intraday Only]	0.038	0.039	0.006	[0.028, 0.05]
$\widehat{\phi}_{\sigma_y}$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]	-0.011	-0.011	0.006	[-0.022, -0.0]
$\widehat{\phi}_{\sigma_y}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	1.424	1.423	0.005	[1.414, 1.432]
$\widehat{\phi}_{\sigma_y}$ Estimated Annual Consumption (kWh) [Z-Score]	0.197	0.197	0.002	[0.192, 0.202]
$\widehat{\phi}_{\sigma_y}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.009	-0.009	0.001	[-0.011, -0.007]
$\widehat{\phi}_{\sigma_y}$ Tenure (Years Prior to 1st SS12 Opt-in Notice) [Z-Score]	-0.025	-0.025	0.001	[-0.027, -0.023]
$\widehat{\phi}_{\sigma_y}$ Business Entity [Ref Non-Business]	0.307	0.306	0.040	[0.228, 0.385]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.014	-0.014	0.004	[-0.022, -0.005]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.127	0.127	0.004	[0.12, 0.135]
$\widehat{\phi}_{\sigma_y}$ DNO Region — North [Ref Midlands & South]	0.016	0.016	0.002	[0.011, 0.02]
$\widehat{\phi}_{\sigma_y}$ DNO Region — Scotland [Ref Midlands & South]	0.043	0.043	0.004	[0.034, 0.051]

Table 19: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a simultaneous-equation regression model (2nd Equation Response Y = "Probability of Opt-in"; 1st Equation Response T = "Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]") fitted using the "Baseline" specification ($N = 650,809$).

Parameter	Mode	Mean	Std.	95% HDI
$\hat{\beta}_y$ Intercept	0.443	0.443	0.001	[0.442, 0.444]
$\hat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]	0.026	0.026	0.004	[0.019, 0.033]
$\hat{\beta}_y$ Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]	0.102	0.102	0.016	[0.072, 0.133]
$\hat{\beta}_T$ Intercept	0.000	-0.000	0.000	[-0.0, 0.0]
$\hat{\beta}_T$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]	0.233	0.233	0.001	[0.232, 0.234]
$\hat{\sigma}_y$	0.497	0.497	0.000	[0.496, 0.498]
$\hat{\sigma}_T$	0.073	0.073	0.000	[0.072, 0.073]
$\hat{\rho}_{\epsilon_{i,y},\epsilon_{i,T}}$	0.006	0.006	0.003	[0.001, 0.011]

Table 20: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a simultaneous-equation regression model (2nd Equation Response $Y =$ "Probability of Opt-in"; 1st Equation Response $T =$ "Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]") fitted using the "Extended" specification ($N = 558,290$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}_y$ Intercept	0.423	0.423	0.001	[0.421, 0.424]
$\widehat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]	0.026	0.026	0.004	[0.018, 0.033]
$\widehat{\beta}_y$ Intraday Notice + Intraday SMS + £1.25 Incentive [Ref Intraday Only]	0.100	0.097	0.017	[0.062, 0.131]
$\widehat{\beta}_y$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	-0.000	-0.000	0.001	[-0.001, 0.001]
$\widehat{\beta}_y$ Estimated Annual Consumption (kWh) [Z-Score]	-0.023	-0.024	0.001	[-0.026, -0.021]
$\widehat{\beta}_y$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	0.008	0.008	0.001	[0.007, 0.009]
$\widehat{\beta}_y$ Tenure (Years Prior to 1st SS12 Opt-in Notice) [Z-Score]	0.041	0.041	0.001	[0.04, 0.043]
$\widehat{\beta}_y$ Business Entity [Ref Non-Business]	0.106	0.108	0.028	[0.054, 0.164]
$\widehat{\beta}_y$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.053	0.053	0.003	[0.047, 0.059]
$\widehat{\beta}_y$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.139	0.139	0.003	[0.134, 0.144]
$\widehat{\beta}_y$ DNO Region — North [Ref Midlands & South]	-0.006	-0.006	0.001	[-0.009, -0.003]
$\widehat{\beta}_y$ DNO Region — Scotland [Ref Midlands & South]	0.024	0.024	0.003	[0.018, 0.03]
$\widehat{\beta}_T$ Intercept	0.000	0.000	0.000	[-0.0, 0.0]
$\widehat{\beta}_T$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]	0.225	0.225	0.001	[0.224, 0.227]
$\widehat{\beta}_T$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	-0.000	-0.000	0.000	[-0.0, 0.0]
$\widehat{\beta}_T$ Estimated Annual Consumption (kWh) [Z-Score]	-0.000	-0.000	0.000	[-0.001, 0.0]
$\widehat{\beta}_T$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	-0.000	-0.000	0.000	[-0.0, -0.0]
$\widehat{\beta}_T$ Tenure (Years Prior to 1st SS12 Opt-in Notice) [Z-Score]	0.001	0.001	0.000	[0.001, 0.001]
$\widehat{\beta}_T$ Business Entity [Ref Non-Business]	0.002	0.002	0.004	[-0.006, 0.011]
$\widehat{\beta}_T$ Does Not Have Octopus Product [Ref Has Octopus Product]	-0.005	-0.005	0.000	[-0.006, -0.004]
$\widehat{\beta}_T$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.001	0.001	0.000	[0.001, 0.002]
$\widehat{\beta}_T$ DNO Region — North [Ref Midlands & South]	0.000	0.000	0.000	[-0.0, 0.001]
$\widehat{\beta}_T$ DNO Region — Scotland [Ref Midlands & South]	0.001	0.001	0.000	[-0.0, 0.002]
$\widehat{\sigma}_y$	0.493	0.493	0.000	[0.492, 0.494]
$\widehat{\sigma}_T$	0.072	0.072	0.000	[0.072, 0.072]
$\widehat{\rho}_{\epsilon_{i,y},\epsilon_{i,T}}$	0.005	0.005	0.003	[-0.001, 0.011]

Table 21: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a single-equation regression model (Response Y = "Probability of Opt-in") fitted using the "ITT (Extended Variant)" specification ($N = 558,290$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}_y$ Intercept	0.423	0.423	0.001	[0.421, 0.424]
$\widehat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]	0.026	0.026	0.004	[0.018, 0.034]
$\widehat{\beta}_y$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]	0.022	0.022	0.004	[0.015, 0.03]
$\widehat{\beta}_y$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	-0.000	-0.000	0.001	[-0.001, 0.001]
$\widehat{\beta}_y$ Estimated Annual Consumption (kWh) [Z-Score]	-0.024	-0.024	0.001	[-0.026, -0.021]
$\widehat{\beta}_y$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	0.008	0.008	0.001	[0.007, 0.009]
$\widehat{\beta}_y$ Tenure (Years Prior to 1st SS12 Opt-in Notice) [Z-Score]	0.041	0.042	0.001	[0.04, 0.043]
$\widehat{\beta}_y$ Business Entity [Ref Non-Business]	0.106	0.108	0.028	[0.054, 0.164]
$\widehat{\beta}_y$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.053	0.053	0.003	[0.047, 0.058]
$\widehat{\beta}_y$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.139	0.139	0.003	[0.134, 0.145]
$\widehat{\beta}_y$ DNO Region — North [Ref Midlands & South]	-0.006	-0.006	0.001	[-0.009, -0.003]
$\widehat{\beta}_y$ DNO Region — Scotland [Ref Midlands & South]	0.025	0.024	0.003	[0.018, 0.03]
$\widehat{\sigma}_y$	0.493	0.493	0.000	[0.492, 0.494]

Table 22: Posterior mean, posterior standard deviation, and 95% highest density interval (HDI) for a single-equation regression model (Response Y = "Probability of Opt-in") fitted using the "ITT (Extended Variant; Flexible Scale)" specification ($N = 558,290$).

Parameter	Mode	Mean	Std.	95% HDI
$\widehat{\beta}_y$ Intercept	0.423	0.423	0.001	[0.421, 0.424]
$\widehat{\beta}_y$ Intraday Notice + Day-ahead Email [Ref Intraday Only]	0.027	0.026	0.004	[0.018, 0.034]
$\widehat{\beta}_y$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]	0.022	0.022	0.004	[0.014, 0.029]
$\widehat{\beta}_y$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	-0.014	-0.014	0.003	[-0.02, -0.008]
$\widehat{\beta}_y$ Estimated Annual Consumption (kWh) [Z-Score]	-0.019	-0.019	0.002	[-0.022, -0.016]
$\widehat{\beta}_y$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	0.008	0.008	0.001	[0.007, 0.009]
$\widehat{\beta}_y$ Tenure (Years Prior to 1st SS12 Opt-in Notice) [Z-Score]	0.042	0.042	0.001	[0.041, 0.043]
$\widehat{\beta}_y$ Business Entity [Ref Non-Business]	0.109	0.110	0.029	[0.053, 0.164]
$\widehat{\beta}_y$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.053	0.053	0.003	[0.048, 0.059]
$\widehat{\beta}_y$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.139	0.139	0.003	[0.133, 0.144]
$\widehat{\beta}_y$ DNO Region — North [Ref Midlands & South]	-0.006	-0.007	0.002	[-0.01, -0.004]
$\widehat{\beta}_y$ DNO Region — Scotland [Ref Midlands & South]	0.024	0.024	0.003	[0.018, 0.03]
$\widehat{\phi}_{\sigma_y}$ Intercept	-0.710	-0.710	0.001	[-0.712, -0.707]
$\widehat{\phi}_{\sigma_y}$ Intraday Notice + Day-ahead Email [Ref Intraday Only]	0.007	0.007	0.006	[-0.004, 0.018]
$\widehat{\phi}_{\sigma_y}$ Intraday SMS + £1.25 Incentive Rand. Assigned [Ref Not As.]	0.004	0.003	0.006	[-0.008, 0.014]
$\widehat{\phi}_{\sigma_y}$ Total P376 (Unadjusted) Baseline (kWh) [Z-Score]	0.004	0.005	0.001	[0.003, 0.007]
$\widehat{\phi}_{\sigma_y}$ Estimated Annual Consumption (kWh) [Z-Score]	-0.006	-0.006	0.002	[-0.01, -0.002]
$\widehat{\phi}_{\sigma_y}$ Index of Multiple Deprivation Rank (Postcode) [Z-Score]	0.002	0.002	0.001	[-0.0, 0.004]
$\widehat{\phi}_{\sigma_y}$ Tenure (Years Prior to 1st SS12 Opt-in Notice) [Z-Score]	0.012	0.012	0.001	[0.01, 0.013]
$\widehat{\phi}_{\sigma_y}$ Business Entity [Ref Non-Business]	0.033	0.031	0.040	[-0.051, 0.107]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Octopus Product [Ref Has Octopus Product]	0.010	0.010	0.004	[0.002, 0.019]
$\widehat{\phi}_{\sigma_y}$ Does Not Have Smart Tariff [Ref Has Smart Tariff]	0.015	0.014	0.004	[0.007, 0.022]
$\widehat{\phi}_{\sigma_y}$ DNO Region — North [Ref Midlands & South]	-0.003	-0.003	0.002	[-0.007, 0.001]
$\widehat{\phi}_{\sigma_y}$ DNO Region — Scotland [Ref Midlands & South]	0.009	0.008	0.004	[-0.001, 0.016]