AI in Charge: Large-Scale Experimental Evidence on Electric Vehicle Charging Demand*

Robert D. Metcalfe
Columbia University
Centre for Net Zero & NBER

Andrew R. Schein
Centre for Net Zero

Cohen R. Simpson
University College London

Yixin Sun
Centre for Net Zero

September 2025

DRAFT – CONFIDENTIAL – PLEASE DO NOT CIRCULATE

Abstract

One of the promising opportunities offered by AI to support the decarbonization of electricity grids is to align demand with low-carbon supply. We evaluate the effects of one of the world's largest AI managed EV charging tariff using a large-scale natural field experiment. The AI dynamically controlled vehicle charging to follow real-time wholesale electricity prices and coordinate and optimize charging for the grid and the consumer. We randomized financial incentives to encourage enrollment onto the tariff. Over more than a year, we found that the tariff led to a 42% reduction in household electricity demand during peak hours, with 100% of this demand shifted to low-cost, low-emission off-peak periods. The tariff generated substantial consumer and producer savings, while demonstrating potential to lower energy system costs and carbon emissions through significant load-shifting. Overrides of the AI algorithm were low suggesting that this tariff is likely more efficient than a real-time-pricing tariff without AI. Our findings highlight the potential for scalable AI managed charging and its substantial welfare gains for the electricity system and society. We also show that experimental estimates differ meaningfully from those obtained via non-randomized difference-in-differences analysis, due to differences in the samples in the two evaluation strategies, although we can reconcile the estimates with observables.

*We thank Lucy Bourke and Joe Grainger for their critical support with logistics and background research and Charlotte Avery and Jaye Cribb for devising the strategy used to detect possible ownership of an electric vehicle based on electricity usage. We are also grateful to Francesca Corby, Geraldine De Boisse, and Katrin Weingartner for all of their support and partnership for this project, to George Doucy, Ben Ferris, Lilla Hoddosy, Ryan Jenkinson, Will Reeves, and Henry Tagg for their guidance around operational aspects of the EV charging algorithms and associated data, and to Louise Bernard, Daniel Björkegren, David Byrne, Gus Chadney, Sheng Chai, Daniel Lopez Garcia, Andy Hackett, Charlie Hutchinson, Maria Jacob, Alex Schoch, Phil Steele, Chad Syverson, Lucy Yu, and participants at the EV-Flex Kickoff for their valuable insights and comments. This work would not have been possible without financial support by the Abdul Latif Jameel Poverty Action Lab's (J-PAL) King Climate Action Initiative (K-CAI). We thank Navya Kumar for excellent research assistance. All mistakes herein are ours alone. Conflicts of Interest: AS and YS are paid employees of Centre for Net Zero (CNZ), a part of the Octopus Energy Group Limited (OEG), RM is a consultant, and CS was a paid employee of CNZ at the time of the trial but is no longer employed by CNZ. AEARCTR-0013037 and USC IRB: UP-23-00733.

1 Introduction

Artificial intelligence (AI) presents both significant threats and promising opportunities for the future of the power grid (International Energy Agency, 2025*a*). On the threat side, a growing body of work highlights the increasing electricity demand, and resulting grid strain, driven by AI itself, particularly through the rapid expansion of data centers (De Vries, 2023; Aljbour et al., 2024; Bogmans et al., 2025; Chen, 2025; Pilz et al., 2025). This concern dominates the current public and policy discourse (Erdenesanaa, 2023; Kolbert, 2024). Yet AI also offers tools to improve the efficiency and flexibility of energy systems. In particular, it could be used to forecast, manage, coordinate, personalize, and optimize energy demand in the face of rising loads and variable renewable supply (Antonopoulos et al., 2020; Biswas et al., 2024; Boopathy et al., 2024; Sandalow et al., 2024).

AI has the potential to optimize energy demand management for both residential and industrial customers. The first-best solution is real-time pricing tariffs. However, a significant body of evidence suggests that consumers are reluctant to engage with time-varying prices, due to the cognitive burden and effort required to optimize their usage, along with the price uncertainty they must constantly monitor. In this context, AI can play a useful role by working on behalf of consumers to optimize their energy demand, significantly alleviating these cognitive costs. Furthermore, AI can simultaneously manage demand across millions of devices, helping to smooth out fluctuations in demand and maintain stability in wholesale prices.

Electric vehicles (EVs) are a compelling test case for these opportunities. As mobile, flexible loads with storage capacity, EVs can draw and inject power across the grid, potentially helping to smooth fluctuations from variable renewables and shift load across time and space. Realizing this potential, however, poses complex challenges: how to schedule charging to minimize costs without risking range anxiety, how to coordinate large numbers of vehicles to avoid local grid congestion, and how to adapt in real time to changing current and future electricity prices and user behavior. These are exactly the kinds of high-dimensional, dynamic problems where AI excels. As a result, EVs, and the systems designed to manage them, have become central to the computer science and engineering literatures exploring how AI can make energy systems more adaptive, resilient, and efficient.

The keen interest in managing EV charging through AI is because EVs will likely

become one the largest users of electricity in the future. Global sales of EVs have increased year-on-year over the past decade (50-fold increase), comprising more than 20% of new cars in 2024 (International Energy Agency, 2025*b*). With falling battery costs and mandates in many markets to phase out internal combustion engine (ICE) vehicles, the number EVs on the road is set to rise in the coming decade.[1] However, if left unmanaged, this new load could strain grid infrastructure, particularly during peak periods (Li and Jenn, 2024; Bailey, Brown, Myers, Shaffer and Wolak, 2025), threatening grid reliability, raising system costs, and exposing consumers to higher electricity prices, especially in areas with congestion or high marginal generation costs. Unmanaged charging also misses the opportunity for consumers to benefit from off-peak or low-emissions periods, which are likely to become more pronounced as renewable penetration increases.

Managed charging with AI, by shifting the timing and rate of EV charging to align with grid conditions, potentially offers a powerful tool for integrating EVs into the energy system. The core promise of managed charging is to shift electricity use to periods of low demand or high renewable output, rather than letting individual charging behavior drive aggregate load. At scale, this flexibility could reduce or defer the need for costly grid investments, such as peaking generation or local distribution upgrades (Anwar et al., 2022; Bennett et al., 2022; Barman et al., 2023). But realizing this potential requires more than technical feasibility; it depends on whether AI managed charging can be made routine, effective, and unobtrusive for consumers.

Despite increasing policy and commercial interest in AI managed charging (Hildermeier et al., 2022), there remains limited causal evidence on how consumers actually engage with AI or AI-based tariffs (Black et al., 2024). While utilities and regulators across the globe have produced forecasts of how managed charging might shape electricity demand and prices (Lowell et al., 2017; Bradley et al., 2018; Seamonds and Lowell, 2018; NYSERDA, 2021; Anwar et al., 2022; Jones et al., 2022), these projections often lack a credible counterfactual, limiting their empirical validity. In the United States alone, over $10 million has been invested across at least ten recent or ongoing managed EV charging pilots (SEPA, 2024), yet none constitutes a proper field experiment with a credible counterfactual. Moreover, none deploy AI-based charging management at meaningful scale, and all suffer from limited sample sizes, severely undermining statistical power and the reliability of their conclusions. These shortcomings underscore the need for rigorous em-

[1] The UK plans to end the sale of new cars powered solely by ICE by 2030, and hybrids and vans by 2035 (Department of Transportation, 2025). The European Union has legislated to effectively ban new ICE cars by 2035 as part of its "Fit for 55" climate package (European Commission, 2023). The US previously aimed for 50% of new vehicle sales to be electric by 2030, but this executive order has since been revoked (The White House, 2025).

pirical experiments.

Our paper addresses this gap by presenting what we believe is the only natural field experiment to date on the adoption and impact of a managed EV charging tariff – and notably, one driven by AI. In partnership with Octopus Energy, the UK's largest electricity supplier, we implemented a large-scale randomized encouragement design targeting 13,233 suspected EV owners across Great Britain.[2] The intervention promoted the Intelligent Octopus (IO) Go tariff, which combines dynamic time-of-use (ToU) pricing with AI managed charging.[3] The ToU component was designed as having an "off-peak" rate overnight approximately 50% lower than the standard rate, whereas all other times of the day, prices were set at or slightly above the standard rate (see Figure A1a). If the AI decided to charge the EV during the day, it would charge the customer the lower off-peak rate.

IO Go is now, to the best of our knowledge, the world's largest managed EV charging tariff, serving approximately 300,000 customers across the UK and expanding rapidly in the US, Germany, France, and Italy.[4] The AI algorithm behind IO Go coordinates and optimizes large fleets of EV charging. It combines linear programming with machine learning models to minimize expected energy costs, reduce grid congestion, and support grid balancing, all subject to customer-set constraints (a "ready-by" deadline and target state of charge). Updating as new information arrives, it delivers adaptive, price-responsive plans that scale from thousands to hundreds of thousands of vehicles and shift demand away from system peaks.

The field experiment allows us to estimate the causal effect of financial incentives on adoption of the AI managed charging tariff, and to use this variation to instrument for the impact of managed charging on electricity consumption. Treated trial participants were emailed and randomly assigned to one of four groups: (1) email only; (2) £5/month; (3) £50/month; or (4) £50/month conditional on not overriding the AI-managed charg-

[2]To determine whether a customer might own an EV, we used data on household electricity consumption to look for evidence of power draws consistent with ownership of a wall charger. This was done in conjunction with domain experts at Octopus Energy.

[3]The AI management is carried out by Kraken Technologies, a technology company owned by Octopus Energy Group. In practice, the AI generates charging schedules by combining wholesale price signals with forecasts of vehicle availability and user "ready-by" times. These schedules are updated in real time, and can be further adjusted to provide ancillary services or system balancing. The literature variously refers to such arrangements as "managed charging," "smart charging," or "utility-controlled charging." For example, Lagomarsino et al. (2022) define smart charging as coordinated systems that optimize charging for grid balancing and/or user preferences, while Axsen et al. (2017) describe utility-controlled charging as timing decisions delegated to a third party to reduce costs or support renewables (see also Kubli 2022; Bailey and Axsen 2015). By contrast, time-of-use tariffs are a form of user-managed charging in which customers shift demand directly in response to prices (Sovacool et al., 2017; Delmonte et al., 2020). Here we use "managed" and "smart" interchangeably to refer specifically to retailer-controlled charging.

[4]Currently, the tariff manages roughly 2GW of power in Great Britain (Green, 2025).

ing, with all payments lasting three months. A pure control group received no contact. We examined two primary outcomes: tariff take-up and electricity use, with particular focus on peak demand periods, over the 12 months following the encouragement, and we follow our pre-analysis plan (unless otherwise stated).

In addition to the experimental analysis, we conducted a supplementary difference-in-differences (DiD) analysis using observational data to estimate the impact of managed charging on electricity consumption among customers who voluntarily adopted the tariff outside of the trial. We compared electricity use before and after adoption for these customers, relative to a control group of similar customers who had not yet adopted, exploiting variation in adoption timing. This approach provides complementary evidence on the effects of managed charging in a real-world, self-selected yet scaled-up setting.

1.1 Primary findings

We report six main sets of findings.

First, we found that email-based encouragements increased take-up of the managed charging tariff. Across all incentive groups, assignment to treatment led to higher enrollment in the IO Go tariff relative to the control group. While the absolute increases in take-up were modest in absolute terms, even the email-only group, with no monetary incentive, raised enrollment by 3.4 percentage points, while the two £50/month arms nearly doubled that effect. Importantly, we suspect that these take-up rates represent a lower bound, as enrollment was constrained by technical compatibility: trial participants needed a supported charger or EV to join. Lacking data on compatibility, we cannot identify the true eligible population. Thus, the measured effects likely understate the potential impact in a fully compatible setting.[5] We find a price elasticity of the take up of the managed tariff to be -0.143.

Second, we observed strong retention and widespread acceptance of AI managed charging among trial participants. Trial participants who enrolled in IO Go largely remained on the tariff, with post-incentive take-up rates statistically indistinguishable from those during the first three months. Among adopters, we observe high adherence to the AI managed charging schedule: over half never overrode the supplier charging schedule, and only 2.3% of total electricity consumption occurred via overrides. These patterns

[5] According to informal estimates shared by Octopus Energy staff via personal communication, 60–70% of EV owners were likely compatible with IO Go in 2024, based on available EV and charger integrations.

suggest that, once adopted, AI managed charging integrates smoothly into daily routines with minimal disruption, underscoring its potential for long-term grid flexibility. Importantly, these override percentages are far below the threshold where RTP would dominate an AI-ToU tariff, according to a theoretical model we developed to quantify the conditions under which AI managed charging (with overrides) has welfare dominance over RTP. This is even true when the AI algorithm only controls 20% of the household's energy demand (roughly the amount of energy from EV charging in our sample). These empirical and theoretical findings suggest that AI may be the best feasible option in the presence of attention constraints and other real frictions.

Third, using treatment assignment as an instrument for tariff adoption, we estimated that AI managed charging reduced peak-period electricity consumption by 42%, with the entirety of that load shifted to overnight off-peak hours. There was no change in total electricity use, indicating that the program induced temporal load shifting rather than increased consumption. This shifting was similar across all of the randomized encouragement groups (email, low incentive, high incentive, and high incentive plus cost to override).

Fourth, the AI appeared to enhance responsiveness to wholesale electricity prices beyond the effects of the peak to off-peak shifting. We examined consumption patterns of trial participants who signed up to the AI tariff versus those who signed up to a similar time-of-use pricing regime but with no AI management of charging. Comparing the two groups, we found that participants who signed up to IO Go exhibited higher elasticity with respect to wholesale electricity prices during peak evening and off-peak overnight hours, whereas daytime responsiveness is similar across groups. While we cannot interpret these differences causally due to non-random assignment, the results suggest that the AI managed scheduling in IO Go helps shift consumption away from high system prices within periods of the day, i.e., above and beyond between-period shifting.

Fifth, DiD estimates were smaller than the experimental estimates. We believe this difference was due to differing samples. While both the experimental and DiD approaches reveal similar directional patterns of load shifting, the experimental instrumental variables estimates indicated a 42% reduction in peak-period consumption, compared to just 8% in the DiD estimates. An important contributor to this gap is that many customers in the DiD sample were already on time-of-use tariffs before adopting IO Go, leaving them with less potential for further consumption change. After reweighting the DiD sample to match the experimental group on prior tariff type, the estimated peak reduction rises to

22%: narrowing, but not eliminating, the gap. The remaining gap may be explained by slightly different automation patterns in 2023 and 2024, with more diffuse consumption reductions across the day in the DiD sample.

Sixth, we estimated the consumer, producer, grid, and social welfare benefits of this AI managed charging tariff. The reallocation of electricity use to periods with substantially lower rates reduced trial participant bills by £343 per trial participant per year (an estimated 18% reduction in cost per kWh).[6] The electricity retailer saw similar savings in procurement costs (procurement costs include the wholesale price of power and non-energy charges such as transmission and distribution fees), implying near 100% pass-through of savings to trial participants, at least based on the period of our analysis. Impacts on CO_2e emissions were large. In a well-functioning market where the benefits of flexibility are internalized by market participants, the case for subsidizing managed charging is weak; however, where there are market failures that cause grid balancing costs to be externalized, this "grid stabilization externality" may justify such a subsidy; the marginal value of public funds is infinite. The resource cost per tonne saving is extremely large (in the negative direction) for consumers. We estimate a lower bound resource cost per tonne of around -£2,445, which is an order of magnitude lower than the next-best technology (Gosnell et al., 2020; Hahn et al., 2024). This tariff and technology is a very cost-effective way to reduce CO_2e emissions.

1.2 Contribution to the existing literature

In our setting, the AI-managed charging tariff used real-time wholesale prices as a key input into scheduling charging, without exposing consumers to these granular prices. Instead, the utility centrally controlled EV charging when the vehicle was plugged in, scheduling it to minimize wholesale procurement costs and/or to maximize revenue from ancillary services, while satisfying consumer preferences. Consumers faced a higher rate during the day and receive an off-peak rate during a fixed six-hour window (23:30–05:30), as well as during any additional hours when the AI activates smart charging. The AI tariff we studied therefore represents a partial and centrally coordinated implementation of real-time pricing (RTP). It preserves key elements of allocative efficiency by shifting load

[6]The counterfactual used here is the average bill of the control group in the experimental period, which includes many tariffs. Using our estimated consumption treatment effects from Figure 6, we find that, relative to the bill under the standard flat tariff, the savings amount to approximately £650.

away from high-cost periods but bypasses price-based incentives at the individual level.[7]

Our research contributes to the literature on tariff switching and dynamic pricing. The closest empirical papers to ours in the tariff-switching literature are Fowlie et al. (2021) and Ito et al. (2023), who both studied consumer choices of time-varying tariffs. Both ran randomized experiments and studied adoption as well as consumption. Ito et al. (2023) documented selection on price-elasticity and consumption profiles in a framed field experiment and showed that providing consumers with take-up incentives encouraged more switching, and the more elastic consumers were more likely to adopt time-varying pricing. Fowlie et al. (2021) compared the adoption rates and aggregate demand response of ToU (and critical peak pricing) under opt-in and opt-out set-ups in a natural field experiment. They found that demand response decreases over time among always-takers and increases over time among complacents. However, both papers considered time-varying tariffs where rates are set ex-ante and therefore cannot address issues related to spot price uncertainty (except via critical peak pricing), which is a key element in the case of real-time pricing and/or AI managed tariffs that optimize around such prices. In their set-ups, the only uncertainty consumers face relates to their preferences and in particular how costly it is to change their consumption habits.[8] These papers also do not have an AI component.

In our setting, once consumers enroll, intra-day energy demand adjustments are fully managed by AI. This suggests little selection on levels and slopes, since the behavioral effort required from consumers is minimal, in contrast with Einav et al. (2013) and Ito et al. (2023).[9] These results raise broader welfare questions about the role of AI in managing consumption of intertemporal goods with time-varying prices. Three related studies speak to automation in similar contexts. Blonz et al. (2025) examined the impact of automated temperature set-points through a smart thermostat in an opt-in field experiment with ToU pricing. While their system had no AI component (the automation was deterministic, and did not update over time or use machine learning), they convincingly demonstrate the benefits of the automation (raising temperature set-points in the sum-

[7] While one way to assess the efficiency of this AI managed charging tariff is by estimating its correlation with wholesale prices (Hinchberger et al., 2024), this proves challenging in our case. The AI is performing a complex real-time optimization to maximize overall welfare, not just in wholesale markets, but also in ancillary services, while simultaneously optimizing consumer welfare. This includes aggregating and coordinating user preferences to ensure EVs are charged when needed. Additionally, the sufficient statistics approach to evaluating tariff efficiency overlooks total net benefits, particularly the role of tariff take-up and the own-price and cross-price elasticity (over time) of those who take up the tariff. AI has the promise of making RTP tariffs palatable (i.e., reducing the need for consumer attention and reducing high price uncertainty) in the future for consumers.

[8] This within-day variation and the associated welfare benefits of such change is also related to the work of Hahn et al. (2023).

[9] Other costs affecting selection on tariff levels may remain, such as lack of control (Bailey and Axsen, 2015) or privacy concerns (Moser, 2017).

mer during the peak price part of the day) and low override rates. Our findings of "set and forget" (via automation) having similarly low overrides support their results. Bailey et al. (2024) and Khanna et al. (2024) studied utility-managed demand response in opt-in experiments, again without AI. We extend these important contributions by leveraging a natural field experiment with AI deployed in a scaled-up market product.

Beyond these related studies, a broader literature on tariff switching and design documents significant inertia and inattention in energy markets, often attributed to high switching costs (Hortaçsu et al., 2017; Byrne et al., 2022; Gravert, 2024; Garcia-Osipenko et al., 2025). Our study is the first, to the best of our knowledge, to estimate switching costs in a natural field experiment by randomizing financial incentives for tariff switching. We found a price elasticity of switching to an AI tariff to be -0.143, suggesting the importance of switching costs.[10]

We also leveraged our empirical findings to create a theoretical model building on the conceptual foundations of Joskow and Tirole (2006*b*), who studied optimal tariff menus in competitive retail electricity markets. In their two-stage framework, consumers first choose among contracts and then respond to real-time wholesale prices under the terms of their chosen tariff. They showed that with sufficiently low transaction and metering costs, high-granularity tariffs such as RTP dominate coarser options like flat or ToU pricing. We retained the two-stage structure but add an AI managed option in which price response is automated. Critically, we also allowed for partial non-compliance in the form of overrides, which generate deviations from the optimal automated schedule. This extension reveals conditions, captured by a closed-form "crossover override rate" threshold, under which automation can outperform RTP even when RTP's transaction costs are small, and conversely when frequent overrides erode the automation advantage.

We also connect to Borenstein (2005*b*), who formally demonstrated that RTP can improve both allocative and investment efficiency relative to static retail tariffs. Borenstein's model, however, assumes full compliance under RTP and abstracts from behavioral frictions such as attention costs, risk aversion, or non-compliance with automated control. Our contribution is to embed these frictions explicitly in the welfare ranking exercise and to operationalize compliance as a stochastic decision margin with a measurable threshold. By doing so, we bridge the gap between the "full compliance" efficiency case for RTP and the more friction-aware tariff choice models in the retail competition literature.[11] In-

[10] Prior work in health, telecommunications, and energy has typically estimated switching elasticities using structural models or by experimentally varying the information provided about costs and benefits across options.

[11] For quantitative simulations of RTP's welfare gains, see Holland and Mansur (2006) for PJM, Poletti and Wright

deed, many RTP tariffs have struggled with low adoption and limited demand response, largely due to rational and/or irrational inattention and aversion to price uncertainty.[12]

Our experimental design allows us to measure override rates, estimate the welfare gap absent overrides, and compare the realized override frequency to the theoretically derived crossover rate. In this way, the model delivers not just comparative statics but concrete experimental benchmarks for when automation should, and should not, replace direct price exposure. The scalability, efficacy, and consumer acceptance of an AI hybrid tariff like this suggest they may become the dominant model for EV (and possibly heat pump) households, which will be the modal household in the next ten to fifteen years.

Our paper also contributes to the growing literature on EV charging, particularly work that explores how consumers respond to incentives to shift charging behavior (Garg et al., 2024; La Nauze et al., 2024; Bailey, Brown, Shaffer and Wolak, 2025; Bernard et al., 2025). While these studies provided important insights, they primarily examine manual charging decisions and typically involved selected or opt-in samples. A close and innovative study by Bailey, Brown, Myers, Shaffer and Wolak (2025) examined managed charging versus ToU in a small framed field experiment in Canada, but the managed charging component had no AI built in and the study lacked sufficient power to demonstrate the impact of managed charging for energy demand against the null. In contrast, our study is the first to evaluate a large-scale, AI managed charging product embedded within a real market tariff, with no selection into the experimental sample. We examined the world's largest managed EV charging program, implemented through a natural field experiment in partnership with a major electricity supplier, which allowed us to capture consumer demand in a real-world setting, free from selection biases into the sample. Moreover, our focus on AI-driven automation offers novel insights into how intelligent systems, not just prices, can coordinate household electricity demand at scale. We also randomized prices onto the AI tariff, allowing us to model the generalizability of the tariff.

There is a large body of work suggesting the importance of EV charging infrastructure and policies on long-term EV adoption (Zhou and Li, 2018; Archsmith et al., 2022; Powell et al., 2022; Rapson and Muehlegger, 2023*b*,*a*; Heid et al., 2024; Turk et al., 2024; Asensio et al., 2025; Dorsey et al., 2025; Gillingham et al., 2025). Our research contributes to this literature by causally estimating how an AI managed charging tariff can reduce the costs of adopting an EV, increase consumer surplus, optimize charging in line with grid condi-

(2020) for New Zealand, and Imelda et al. (2024) for renewable-heavy systems, such as the case of Oahu, Hawaii.

[12] See Borenstein (2007); Allcott (2011); Joskow and Wolfram (2012); Wolak (2013); Fabra et al. (2021); Leslie et al. (2021); Cahana et al. (2022); Pébereau and Remmy (2023); Enrich et al. (2024) for evidence.

tions so as to reduce costs building and operating electricity systems. These components are important for any welfare calculations of any EV policy. In our welfare framework, there are large cost reductions, but uncertainty on whether such reductions are captured all by firms and consumers or whether the Government will also spend less on the grid in the future (which would generate a Pareto improvement if so).

Finally, we contribute to the research on the economics of AI by demonstrating how AI can causally affect environmental outcomes.[13] While prior natural and field experiments have examined AI's impact on labor productivity (Agarwal et al., 2023; Cui et al., 2024; Björkegren et al., 2025; Brynjolfsson et al., 2025), we believe that our study is the first to provide causal evidence indicative of reduced energy consumption under AI-driven optimization, where we do not observe differences across broad household types.. Although scholars have highlighted the challenges of identifying AI's effects in labor markets (Brynjolfsson et al., 2019; Frank et al., 2019), these identification barriers are less pronounced in energy and environmental contexts. AI often lowers the cost of effort, making it difficult to disentangle AI from price effects and isolate its unique impact. This matters because reducing effort, through tools, automation, or process improvements, has long influenced labor markets, yet existing studies rarely separate this effect from AI itself. In our setting, the reduction in effort occurs on the retailer side in responding to wholesale prices or ancillary markets, while consumer prices remain unaffected by the AI's optimization.

The paper is structured as follows: Section 2 provides an overview of the natural field experiment, including the sampling, the intervention (i.e., the AI-EV tariff), the randomized encouragement, the randomization procedure, and available data. Section 3 develops the demand estimates from the field experiment. Section 4 complements Section 3 with a difference-in-difference analysis of the intervention, and attempts to reconcile with the experimental estimates. Section 5 presents the welfare estimates of the intervention and the randomized encouragement, and finally, Section 6 concludes.

2 Experimental design

This section describes the design and implementation of our field experiment. We begin by detailing the eligibility criteria and sampling strategy used to construct the study

[13] This complements the approach that AI data centers could impact on energy flexibility and thus emissions (Knittel et al., 2025).

population. We then describe the AI managed charging tariff and its underlying automation, before outlining the randomized encouragement design used to induce uptake. We conclude by summarizing our randomization procedure and the data used in the analysis.

2.1 Eligibility into the field experiment

We implemented the field trial in partnership with Octopus Energy, the United Kingdom's largest electricity supplier. We defined the target population as residential customers satisfying three criteria: (i) they were likely to own an EV; (ii) as of December 2023, they had only ever subscribed to conventional flat-rate or variable-rate import tariffs – that is, they had no prior engagement with time-of-use pricing or AI managed charging; and (iii) they resided in houses (rather than apartments/flats or mobile homes), which are likely suitable for at-home charging.

As we lacked administrative records on EV ownership, we inferred it using household electricity consumption data from smart meters. Following internal technical guidance, we defined "suspected charging events" as four to twelve consecutive half-hourly intervals with total usage exceeding 3.5 kWh, consistent with Level 2 (7 kW) home charging. We averaged the number of such events per week over a ten-week window in summer 2023 (July–August), and classified as "suspected EV owners" those customers with between 0.5 and 4 events per week (these thresholds were chosen in consultation with Octopus Energy colleagues with subject matter expertise on typical EV charging frequency among customers).

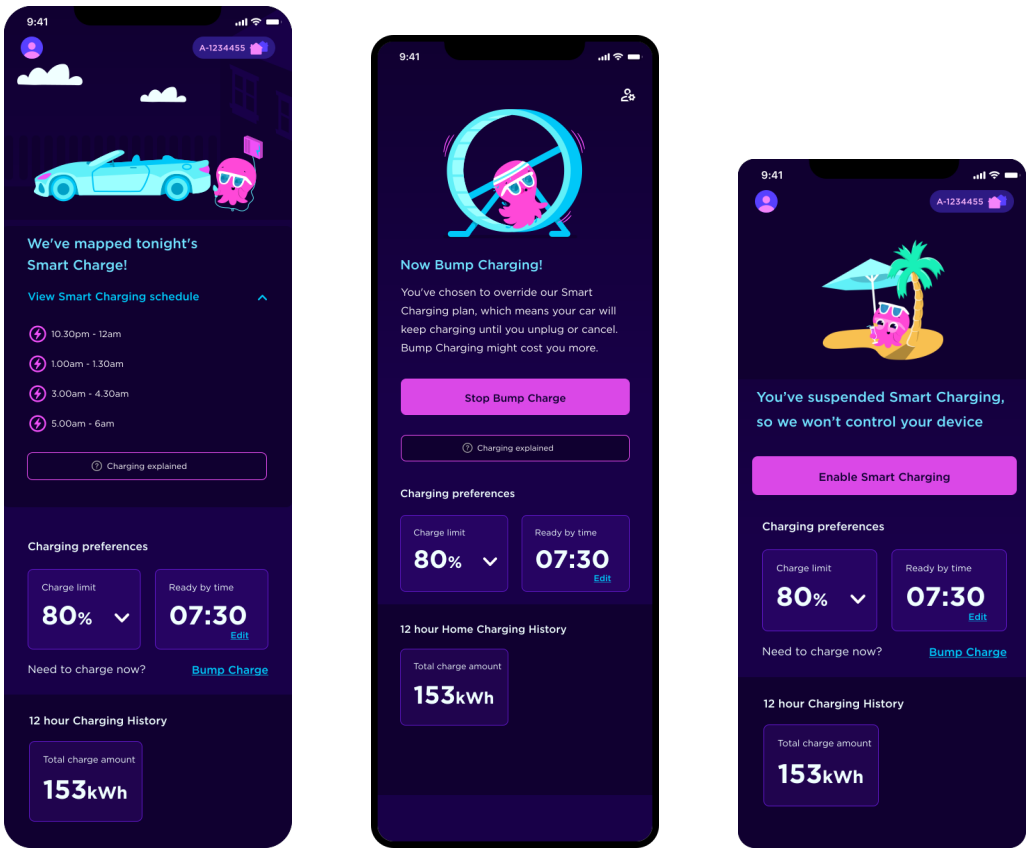
This sampling strategy yielded 13,233 trial participants who are more plausibly representative of mainstream British EV owners. Unlike early adopters already enrolled in smart tariffs, these trial participants are likely to be more price-sensitive, less motivated by environmental ideology, and more similar to the broader population of potential adopters of managed charging technologies.

2.2 The intervention: Intelligent Octopus (IO) Go

IO Go is a residential electricity tariff that combines time-of-use pricing with AI management of automated EV charging. Through the Octopus app, customers set a target battery level and departure time; Octopus then schedules charging to meet those targets (Figure 1a). In exchange, IO Go provides a favorable electricity rate of £0.07/kWh during

a fixed six-hour off-peak window (23:30–05:30).[14] This rate applies not only to EV charging but to all household consumption during that period. Charging that occurs outside the scheduled off-peak window is still billed at the off-peak rate if it is initiated by the AI automation. Customers retain the ability to manually override this schedule via a mobile app ("bump charging", shown in Figure 1b). When customers charge outside of the schedule, it is billed at the higher rate. The applicable tariff schedule for 2024 is shown in Figure 2.

Figure 1: IO Go Charging Controls via Mobile App



(a) Daily Charging Schedule

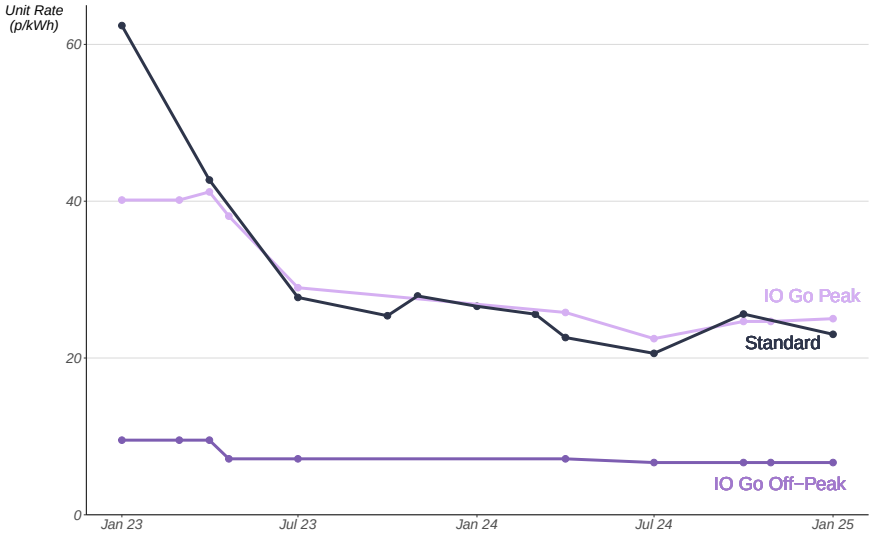
(b) Temporary Override

(c) Sustained Override

The IO Go tariff uses AI to generate charging schedules based on wholesale electricity prices, adapted to the market structure of each region where it is offered. In Great Britain, where wholesale prices are national, schedules are optimized using a blend of day-ahead and intraday wholesale prices, with intraday data available on a rolling-hourly basis. In

[14] Octopus Energy's standard tariff Flexible Octopus charges a flat rate of approximately £0.22–£0.27 per kWh throughout the day, depending on region, as shown in Figure 2.

Figure 2: Tariff Rates



Notes: This figure shows the tariff rates for Intelligent Octopus Go customers during the off-peak overnight period (23:30–05:30, dark purple) and the peak daytime period (05:30–23:30, light purple). For comparison, we also include the Flexible Octopus tariff from Octopus Energy, which maintains a flat rate throughout the day.

most European markets, the optimization is updated daily based on day-ahead wholesale prices, while in the United States, the algorithm relies on forecasts of real-time prices. The optimization window can extend up to 24 hours into the future, accommodating user-specific "ready-by" times. These may require intraday charging plans – such as afternoon readiness – or more commonly, overnight charging plans, particularly in Great Britain.

Following the baseline price-driven optimization, the system supports a secondary adjustment layer for participation in ancillary service markets and real-time grid-balancing operations. In this mode, the EV portfolio is managed as a coordinated fleet with portfolio-level volume targets, allowing for dynamic adjustment of charging schedules in response to market trades or direct requests from system operators. This enables dual operational modes: price-only optimization, in which vehicles charge during the lowest-cost periods (often aligning with high renewable output), and price-plus-volume optimization, in which charging is redistributed to meet specific aggregate energy delivery constraints while maintaining cost efficiency.

The optimization engine integrates two complementary computational approaches: classical linear programming and machine learning–based forecasting. The linear programming component solves for the cost-minimizing charging schedule subject to technical and operational constraints, such as maximum charging rates, user readiness requirements, and fleet-wide volume limits. The machine learning component enhances

this process by providing predictive inputs to the optimization, including forecasts of EV availability (based on historical connection patterns), expected charging requirements for vehicles not yet connected, and local load predictions, including home consumption and on-site solar generation when applicable. These forecasts are updated in real time, enabling the optimization to adapt to changing conditions.

From an efficiency point of view, a key advantage of IO Go's automation is to mitigate inefficiencies due to human error, bounded rationality, or rational inattention. In environments where consumers are either rationally inattentive or lack the sophistication to respond optimally to high-frequency price signals, automated optimization under IO Go may yield strictly higher social welfare (from a net benefits or resource cost-effectiveness point of view) than consumer-managed RTP (see Appendix A.5 for some predictions).

2.3 The randomized encouragement design

The trial participants in our sample were randomized into one of five arms:

1. Control group (n = 2,205): no outreach
2. Email (n = 7,720): encouragement email with no financial incentives for signing up to IO Go
3. Email + £5/month incentive to sign up (n = 1,101): encouragement email with offer of £5/month for three months for signing up to IO Go
4. Email + £50/month incentive to sign up (n = 1,102): encouragement email with offer of £50/month for three months for signing up to IO Go
5. Email + £50/month incentive to sign up (n = 1,105), no bump charging (overriding): encouragement email with offer of £50/month for three months for signing up to IO Go. Additionally, they pay £2 of their incentive for each day they "bump charged" – that is, overrode the AI control at least once per day. This was designed to probe trial participants' willingness to tolerate even less control over their charging schedule.

Each treatment arm received a single encouragement email from Octopus Energy's marketing team, reproduced in Appendix A.4.1. These messages included a prominent call to action, a theoretical £700/year savings estimate (based on historical usage modeling), and tariff-specific details, with minor variations in content to reflect the assigned

incentive level. The "no bump" group was explicitly informed that their monthly payment would be reduced by £2 for each day they initiated a manual override of managed charging.

Encouragement emails were dispatched in two waves: the first on February 15, 2024, and the second on March 20, 2024. Incentive payments were credited to trial participants' Octopus Energy account balances, accruing daily over a 90-day period conditional on maintaining an active IO Go contract. In effect, this structure functioned as a discount on the customer's electricity bill, lowering the effective cost of household energy use during the incentive window.

Note that it was required by our implementing partner, Octopus Energy, that our encouragement emails for IO Go inform customers that Octopus Energy offers other tariffs that may have better met their needs. Importantly, IO Go is not compatible with all chargers or vehicles; for these circumstances, Octopus Energy offers an alternative EV tariff: Octopus Go, a time-of-use EV tariff offering a fixed off-peak rate for electricity. The major differences from Intelligent Octopus Go are: (1) it offers one fewer hour of cheap overnight rate (00:30-5:30, instead of IO Go which is 23:30-5:30) (2) its off-peak rate is higher than IO Go's (£0.085/kWh, as compared to £0.07/kWh for IO Go; for exact rates, see Figure A1b), and (3) Octopus Go does not incorporate managed charging and thus has no bonus off-peak windows. This creates the possibility that encouragements influenced uptake of tariffs other than IO Go, posing a potential channel for exclusion restriction violations. We test and discuss this in Section 3.1.

2.3.1 Hypotheses

The design of the field experiment allowed us to test two main pre-specified hypotheses:

1. Increasing the incentive payment for adopting IO Go will increase actual adoption.
2. Adoption of IO Go tariff will shift electricity consumption from peak (16:00–20:00) to off-peak (23:30-5:30) hours.

These two hypotheses are not from the same family and thus we did not correct for multiple hypothesis testing across them. However, we will go into some of the testable predictions from the design below.

2.4 Tying the Experiment to a Theoretical Framework

From a welfare analysis perspective, we expect the IO Go tariff to generate distinct welfare outcomes relative to real-time pricing (RTP), depending on the nature of optimization opportunities and consumer behavior. IO Go can outperform RTP when it leverages price signals from multiple markets, such as ancillary markets that are typically excluded from the RTP signal.[15]

We consider EV charging under four tariff regimes: flat (the baseline counterfactual our sample is initially on), time-of-use (ToU) with no AI, real-time pricing without AI (RTP), and the AI tariff "IO Go" with possible household overrides ("bump charging"). Households require a certain amount of energy by a deadline, have specific plug-in hours, and face inconvenience costs when charging at less-preferred times.

The retailer's total cost depends on total load and wholesale prices, net of ancillary service benefits. Automated IO Go schedules charging centrally to minimize these costs, while RTP households schedule themselves but face *attention costs* and *risk from bill volatility*. ToU households respond to fixed peak/off-peak price differences.

We present the full model and predictions in Appendix A.5, which builds upon Borenstein (2005*a*,*b*); Joskow and Tirole (2006*a*,*b*), but here in this section we provide the general predictions. Automation under IO Go will have higher welfare than RTP with some attention and bill variance costs. Under IO Go, households may *override* the schedule when their immediate value from charging exceeds the benefit of deferring plus a hassle cost. Overrides can raise user utility but may increase system costs if they shift charging into expensive or high-emission hours. This will reduce the social welfare of the AI tariff and reduce the gap between it and RTP.

The two key parameters to estimate how many overrides mean that IO Go has lower welfare than RTP are: (i) ΔW_0: baseline welfare advantage of IO Go over RTP without overrides; and (ii) λ^{B}: per-override welfare loss bound. These allow us to estimate the crossover override rate:

$$\bar{\beta}^{\star} = \frac{\Delta W_0}{\lambda^{\mathrm{B}}}$$

which is the threshold above which IO Go with overrides yields lower welfare than RTP.

[15] In principle, RTP could incorporate such signals, but doing so would require a richer, possibly non-transparent, price vector and greater computational and behavioral demands on the consumer. However, IO Go may yield lower welfare than RTP in settings where intertemporal optimization across days offers higher gains compared to within-day load shifting, or when consumers systematically fail to plug in their EVs at times that maximize procurement cost savings, though there is little evidence from RTP papers that people switch over days, and no evidence from our data that people fail to plug in their EV.

In a stylized calibration with price elasticity of demand in RTP being -0.5, $\bar{\beta}^{\star} \approx 0.06$ to 0.2 overrides per EV–day at moderate attention costs; both higher attention costs and lower price elasticities under RTP raise this threshold. This assumes the amount of energy demand controlled by the AI algorithm is the same as under RTP. That is not true in practice, but we show that if EV demand is 20%, then $\bar{\beta}^{\star} \approx 0.012$ to 0.042.

Our model builds on strands of the energy economics and operations literature that model demand under dynamic electricity pricing. Threshold-style decision rules, in which users act only when net private benefits exceed a frictional cost, appear in online EV charging optimization under RTP (Yi et al., 2019), where a dissatisfaction penalty plays a role similar to our hassle cost ϕ_i. Menu-based contract designs for EV charging (Ghosh and Aggarwal, 2017) and joint welfare-maximizing scheduling algorithms (Huang et al., 2023) adopt multi-stage decision structures that parallel our two-stage extension, in which households first adopt a tariff and then choose whether to override automation.

Our model also relates to the generalized Roy model framework of Ito et al. (2023), who study voluntary take-up of dynamic pricing plans using marginal treatment effects to characterize heterogeneity in welfare gains from adoption. In their setting, the key policy-relevant object is the optimal adoption cutoff; in ours, it is the *crossover override rate* at which an AI schedule ceases to dominate RTP in welfare. Both identify a threshold along a behavioral margin, adoption or overrides, at which the welfare ranking of competing regimes changes.[16] Relative to this literature, our contribution is to embed such a threshold in an AI model for EV charging, explicitly linking override behavior to aggregate welfare outcomes and making sharp, testable predictions for our experimental design. Based on our setting and with a generous price elasticity of demand in the RTP set-up (-0.5), our simulations suggest that an override rate of less than 10% rate would make the AI tariff dominant over an RTP.

[16] Joskow and Tirole (2006*b*) present a general model of retail electricity competition in which tariff menus are designed to elicit efficient real-time demand response, subject to transaction and metering frictions. In their framework, sufficiently low frictions imply that high-granularity pricing schemes such as RTP dominate coarser alternatives. Borenstein (2005*b*) similarly show that RTP improves allocative and investment efficiency relative to flat or ToU rates in competitive markets, assuming full compliance and ignoring behavioral frictions. Our model extends these frameworks in two ways: (i) we introduce an AI managed regime (IO) that automates price response, removing household attention costs but allowing partial non-compliance through overrides, and (ii) we model overrides as a distinct behavioral margin with its own welfare implications. This richer friction structure implies that IO Go can outperform RTP even when RTP's transaction costs are small, provided override rates are below a crossover threshold $\bar{\beta}^{\star}$, and conversely that high override rates can reverse the ranking – an effect absent from the original Joskow & Tirole and Borenstein formulations.

2.5 Randomization

Random assignment was implemented using a block randomization procedure to improve covariate balance across trial arms (Moore and Schnakenberg, 2023). Prior to assignment, trial participants were grouped into 1,109 blocks of twelve customers based on Mahalanobis distance calculated over a set of pre-encouragement variables predictive of EV ownership and electricity consumption. These included historical electricity usage, tenure as a customer with Octopus Energy, and past engagement with smart tariff onboarding. Within each block, two trial participants were assigned to the pure control group, seven to the £0/month encouragement group, and one each to the £5/month, £50/month, and £50/month (no bump) groups – reflecting both budget constraints and the firm's commercial interest in maximizing exposure to IO Go. Randomization was performed separately within each of the United Kingdom's 14 electricity distribution regions to ensure geographic stratification. This procedure yielded excellent covariate balance across trial arms (Table A1, and Figure A2).

2.6 Data

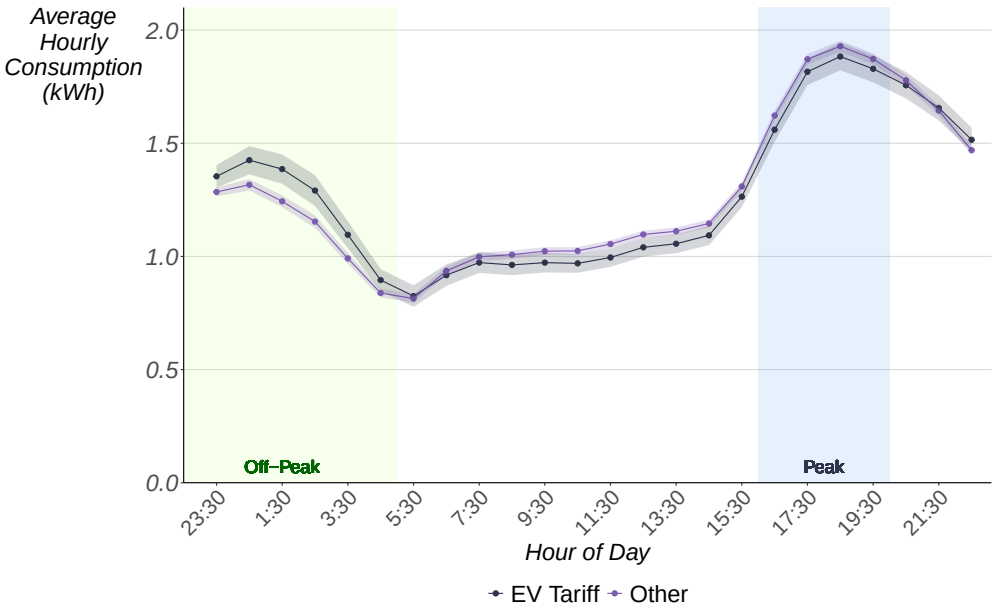
Our analysis drew on high-frequency administrative data from Octopus Energy, Great Britain's largest electricity supplier, with personal identifiers removed. We focused on two primary outcomes for our two hypotheses: (1) take-up of the IO Go tariff, measured as a binary indicator for whether a trial participant held an active IO Go contract during week t; and (2) electricity demand, observed at half-hourly resolution and expressed in kilowatt-hours (kWh). We aggregated consumption to the week × hour-of-day level. Electricity demand includes both EV-related and non-EV household load; concretely, we summed smart-meter readings (which measure consumption from all appliances, not just the EV charger) across all meter point administration numbers linked to each trial participant account. These outcomes are observed from January 1, 2024 through March 31, 2025.[17]

Figure 3 shows average hourly electricity consumption in the pre-encouragement period (i.e., January 2024), showing an expected evening peak beginning around 16:30. Notably, we also observe substantial overnight consumption, consistent with EV charging

[17] A meter point administration number (MPAN) is a unique ID for an electricity supply point (e.g, a house) in relation to a specific area of the UK's national electricity grid. Tariff-contracts and half-hourly measurements of electricity use are tied to account identifiers via MPANs. A single account can have multiple MPANs with different tariff agreements that are simultaneously active.

behavior. We believe this is partly driven by plugging the EV after work in the evening and the default scheduling behavior in common EV chargers, which sometimes pre-sets charging to off-peak times.[18] Our intervention thus tests the additional impact of managed charging in a setting where at least some users are already defaulted into off-peak charging schedules.

Figure 3: Pre-Trial (January 2024) Hourly Consumption



Notes: This figure shows average hourly electricity consumption across the sample in January 2024, prior to the start of the trial. Trial participants are grouped by (1) those who later spent more than 50% of the trial period on Intelligent Octopus Go or Octopus Go, and (2) those who remained on other tariffs. Shaded regions show the 95% confidence interval. Octopus Go is an alternative time-of-use EV tariff designed for customers who either had hardware that was incompatible with IO Go or did not wish to enroll in automation. Green shaded box indicates IO Go off-peak hours (23:30–05:30), when electricity is charged at £0.07/kWh; all other hours are billed at the standard variable rate. Blue shaded box indicates typical system peak hours (16:30–20:30), which are highlighted to show times of heightened grid stress.

For trial participants who adopted IO Go, we also collected customer settings and high-frequency telemetry on charging behavior. The settings include: (1) the ready-by time, defined as the user-specified time by which the vehicle should be fully or partially charged; and (2) the desired state of charge by that time. The telemetry data include: (1) plug-in and unplug timestamps; (2) charging start and charging end timestamps; and (3) an indicator for whether the charge was automatically dispatched by Octopus or manually overridden by the user. These data allowed us to reconstruct intended charging preferences, actual charging behavior, and deviations from automated control.

[18] Since June 2022, the UK's *Electric Vehicles (Smart Charge Points) Regulations 2021* have required that all new private EV chargers include a default charging mode set outside of peak hours (8–11am and 4–10pm), along with a randomized delay function to reduce grid strain.

We used additional administrative data at the daily level to estimate consumer and supplier benefits. First, to measure benefits to consumers, we used administrative data from Octopus Energy detailing each customer's unit rate per kWh. Second, to estimate benefits to the electricity supplier (in procurement cost savings), we used administrative data on each customer's wholesale energy costs, as well as non-energy costs, per settlement period. Non-energy costs include Transmission Network Use of System (TNUoS) and Distribution Use of System (DUoS) charges, capacity market payments, and policy costs (such as charges supporting Contracts for Difference); some of these charges vary by period of day (e.g., DUoS). This yields an imputed per-kWh cost that combines the energy and non-energy costs of supplying electricity.

For an exploration of sub-group heterogeneity, we used area-level deprivation data from the UK-wide composite Index of Multiple Deprivation (IMD), as constructed by Parsons and mySociety (2021) using methods from Abel et al. (2016). This version harmonizes the constituent country-specific IMDs to a common England-anchored scale, enabling deprivation comparisons across the UK's statistical reporting areas.

To quantify CO_2e impacts, we integrated data from WattTime, a U.S.-based nonprofit that produces historical estimates of the Marginal Operating Emissions Rate, or the emissions associated with the marginal change in load on the grid (WattTime, 2022). WattTime data is available at five-minute intervals.

3 Experimental results

This section presents the main empirical findings from our field experiment. We begin by documenting the impact of the encouragements on trial participant enrollment in a managed EV-charging tariff. We then examine how these changes in enrollment influenced electricity consumption patterns, using both reduced-form and instrumental variable approaches. Finally, we explore heterogeneity in treatment effects, user behavior under the managed charging regime, and the added value of automation beyond conventional time-of-use pricing. We followed our pre-analysis plan (PAP), but state where we added or deviated away from it, why we did, and their implications for interpretation in Appendix A.3.

3.1 Impact of encouragement on take-up of managed charging

We begin by estimating the effect of encouragement on trial participants' likelihood of adopting a managed EV-charging tariff. Specifically, we estimate the following linear probability model for whether trial participant i is enrolled in the IO Go tariff in week t:

$$D_{it}^{\text{IO}} = \pi_0 + \pi_1 \mathbf{Z}_i + \pi_2 P_{it} + \pi_3 (\mathbf{Z}_i \times P_{it}) + \mu_b + \mu_t + \epsilon_{iht} \tag{1}$$

Here, D_{it}^{IO} is a binary indicator equal to one if trial participant i is on an IO Go contract in week t. Z_i is a set of four binary indicators capturing the treatment assignment, with the control group omitted as the reference category. P_{it} denotes whether participant i is in the incentive period in week t, where the incentive period is the three months after the start of the trial. We include fixed effects for randomization block μ_b and calendar week μ_t. Standard errors are clustered at the level of participant and week.

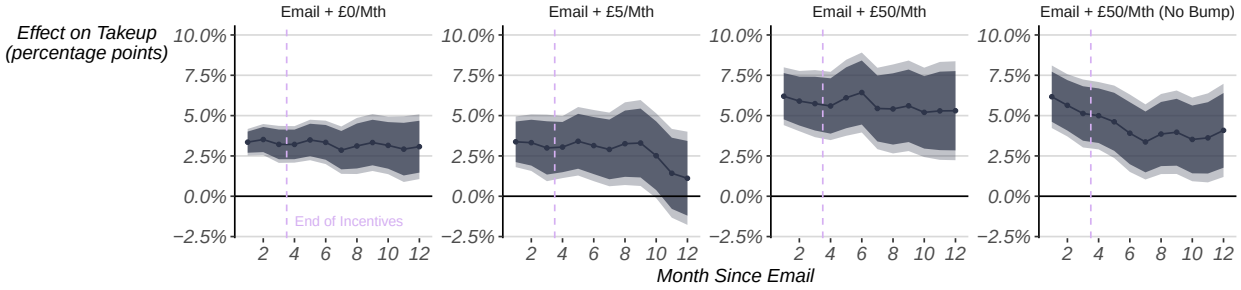
We observe an increase in take-up across all treatment groups, as shown in Table A2. During the 90-day incentive window, the £0/Month and £5/Month groups each increased the probability of take-up by approximately 3.4 percentage points, while the £50/Month and £50/Month (No Bump) groups nearly doubled that effect, reaching 5.9 and 5.7 percentage points, respectively. Take-up in the control group is also rising over time, but remains consistently lower than that observed in any encouragement arm (Figure A3). At the end of the three-month incentivization period, take-up in the control group stood at 2.7%, compared to 7.0% and 6.8% in the £0/Month and £5/Month groups, respectively. The higher incentives resulted in the greatest adoption, with the £50/Month and £50/Month (No Bump) groups reaching 9.3% and 9.2%.[19]. We calculated a price elasticity of 0.143 between the £0/Month and £50/Month groups using the arc elasticity formula, which compares the change in take-up rates relative to the midpoint of both the take-up and total incentive levels (£0 vs £150).

Across most encouragement arms, post-incentive enrollment remained stable. We formally test this in Table A2, interacting the encouragement indicators and the post-incentive period indicator in Table A2. The persistent retention after the incentive period suggests that the initial encouragements had durable effects even in the absence of continued subsidy payments. Only the £50/Month (No Bump) experienced a statistically significant drop of 1.6 percentage points after incentives ended. These results can be fur-

[19]Tariff contracts can change at the level of the day, and thus there could be some worry that take-up analysis at the weekly level could be biased. We run a robustness check of take-up at the daily level in Table A3, and find extremely similar results to Table A2

ther visualized in Figure 4, which shows estimates of Equation (1), split by month since treatment. Retention dropped only in the group facing restrictions on manual overrides, suggesting that managed charging is more likely to succeed when framed as a convenient default rather than a rigid mandate. Encouragements that showcase consumer benefits and flexibility can foster lasting adoption.

Figure 4: Impact of Encouragements on Take-up of Managed Charging Tariff (IO Go)



Notes: This figure plots the intention-to-treat effects of four email-based encouragements to adopt a AI managed charging tariff (Intelligent Octopus Go), split by month since receiving the email. The outcome is a binary indicator for weekly use of the tariff. Each panel corresponds to a different encouragement group, varying in the level of offered financial incentive. Shaded areas represent 90% (dark) and 95% (light) confidence intervals, with standard errors clustered at the participant and week level. The dashed vertical line indicates the end of the 90-day incentive period.

Importantly, take-up is mechanically constrained by compatibility: trial participants need either their charger or EV to be supported by IO Go in order to enroll. For those who do not enroll and have an EV, we do not observe their vehicle or charger details, and there is no comprehensive national data on EV and charger ownership to fill this gap. Therefore, we are unable to estimate what proportion of trial participants actually have compatible equipment.[20] While IO Go is compatible with some of the most popular home chargers (notably Ohme, MyEnergi/Zappi, and Hypervolt) this still represents only a subset of the overall charger market. Compatibility via direct vehicle integration covers several major EV brands, including Tesla, BMW, and VW, but again excludes some others. Thus, the measured take-up rate represents a conservative estimate, limited by the extent to which trial participants' vehicles or chargers were compatible with IO Go.[21][22]

As IO Go is not appropriate for all customers, Octopus Energy offers an alternative EV tariff, Octopus Go, described in Section 2.3. As a result of our email-based encourage-

[20] In February 2024, we sent a survey to 305 trial participants who we had emailed as part of a pre-trial pilot (where pilot trial participants are not in the sample of our main analyses). We received 68 responses. Among the 56 respondents who expressed interest in signing up for IO Go, 17 indicated that they had not done so due to device incompatibility.

[21] See this link for compatibility with charge points. IO Go is also compatible with several major EV brands.

[22] To further contextualize these take-up rates, the email open rate was 78% across all treatment groups, with most opens occurring on the day following delivery and no qualitative variation across treatments (Figure A4). Noncompliance with the intended IO Go take-up can thus be a result of not receiving or opening the encouragement email.

ments, we observed a nontrivial increase in take-up of Octopus Go; our speculation is that this was driven by customers who owned devices incompatible with IO Go. Compared to the effects on take-up of Intelligent Octopus Go, the impacts on Octopus Go adoption were smaller – roughly 30% as large as those for the £0/Month and £5/Month groups, and 8–11% as large as those for the £50/Month and £50/Month (No Bump) groups (Table A5). The results suggest that, although uptake was concentrated among adopters of the managed-charging product, the encouragements induced a more general shift toward tariffs designed for EV needs more generally. This affects how we interpret the impacts of IO Go on electricity consumption, which we will discuss more in Section 3.3.[23]

3.2 Impact of encouragements on electricity consumption

We next assess whether encouragement-induced take-up translated into changes in electricity consumption over the course of the day. Specifically, we estimated intention-to-treat (ITT) effects of each encouragement arm on hourly electricity use (kWh) using the following specification:

$$Y_{iht} = \alpha + \beta \mathbf{Z}_i + \gamma X_{ih} + \psi_b + \psi_t + \varepsilon_{it} \tag{2}$$

where Y_{iht} denotes mean electricity consumption for user i in hour h during week t, $\mathbf{Z}_i$ is a vector of binary indicators for each encouragement assignment, X_{ih} is user i's average pre-encouragement January 2024 consumption during hour h[24], and ψ_b and ψ_t are fixed effects for block and week, respectively. Standard errors are clustered by user and by week (Colin Cameron and Miller, 2015).

Figure 5 presents ITT estimates by encouragement arm, split by hour-of-day. All groups showed increased consumption during the overnight off-peak period (23:30–05:30) and modest reductions during peak hours (16:30–20:30), consistent with the tariff's incentive to shift usage. Patterns are broadly similar across arms, although the £0/Month and £50/Month (No Bump) treatments exhibit the largest declines in peak consumption.

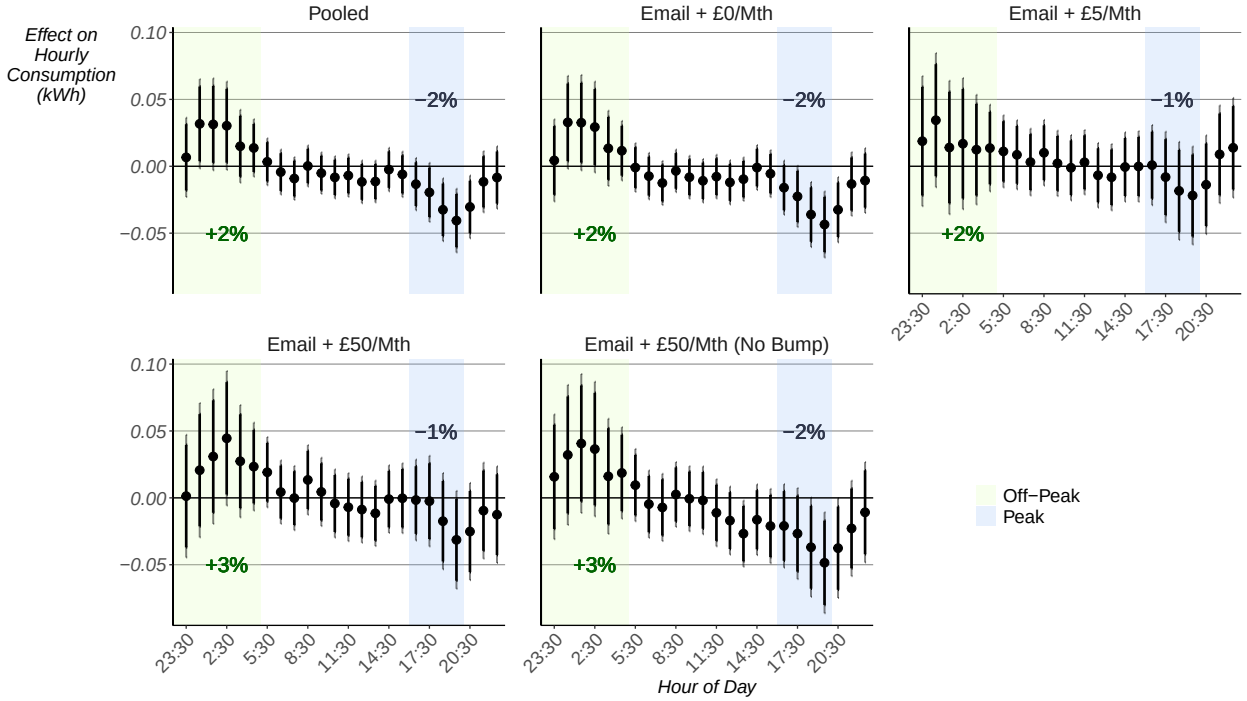
When pooling all treatments into a single encouragement indicator and estimating a regression over all hours within the specified peak and off-peak windows, we find that

[23] IO Go has another rival tariff, Agile Octopus, which is a variable-rate tariff with prices linked to the day-ahead wholesale cost of electricity. Take-up of Agile was quite low in our sample, and we found no effect of our encouragement design on take-up of Agile Octopus.

[24] This control variable was not specified in our pre-analysis plan, but we include it to enhance the precision of our estimates.

receiving any encouragement increased off-peak consumption by 2% (0.019 kWh) and reduced peak consumption by 2% (–0.026 kWh), with no overall effect on total usage (Table A4).

Figure 5: Impact of Encouragements on Hourly Electricity Consumption



Notes: This figure shows intention-to-treat effects of four email-based encouragements on hourly electricity use (kWh), using data from the 12 months after sending out the email (estimated using Equation (2), split by hour-of-day). The first panel defines encouragement as a binary indicator for whether the user received any encouragement. Each subsequent panel represents a separate encouragement arm with varying incentive levels. Lines depict 90% (dark) and 95% (light) confidence intervals. Standard errors are clustered by participant and week. Green shaded box indicates IO Go off-peak hours (23:30–05:30), when electricity is charged at £0.07/kWh; all other hours are billed at the standard variable rate. Blue shaded box indicates typical system peak hours (16:30–20:30), which are highlighted to show times of heightened grid stress. Our outcome measure is hourly consumption, with hours defined as starting on the half-hour to align with IO Go's pricing structure. Percentages represent treatment effects as a share of the control group trial participants who are not on an EV tariff, for off-peak (23:30–05:30, green) and peak (16:30–20:30, blue) periods. Estimates come from regressions pooling all hours in each window, as reported in Table A4 and defined in Equation (2).

3.3 Impact of managed charging on electricity consumption

To obtain the causal impact of adoption of the managed charging tariff on electricity consumption, we used an instrumental variables estimation. In the first stage, we instrumented tariff adoption with random assignment to any of the four email-based encouragements:

$$D_{iht} = \pi_0 + \pi_1 Z_i + \gamma X_{ih} + \mu_b + \mu_t + \epsilon_{iht} \tag{3}$$

where Z_i is a binary indicator for whether the user received any encouragement. X_{ih} denotes average baseline consumption for user i in hour h. The specification includes fixed effects for randomization block (μ_b) and calendar week (μ_t).

As shown in Section 3.1, our encouragements increased adoption of both the managed-charging tariff (IO Go) and the EV-oriented time-of-use alternative (Octopus Go). To preserve the exclusion restriction, given that our instrument affects both products, we define D_{it} as an indicator for take-up of either IO Go or Octopus Go.[25]

In the second stage, we regressed hourly electricity use on predicted tariff status:

$$Y_{iht} = \alpha + \beta \hat{D}_{iht} + \gamma X_{ih} + \psi_b + \psi_t + \varepsilon_{iht} \tag{4}$$

where Y_{iht} is the mean daily consumption for user i in hour h of week t (i.e., the mean daily consumption at each hour, averaged across the week). Standard errors are clustered at both the participant and week levels. We estimated this regression over the 12 months after encouragement emails were sent out.

This approach is valid under standard IV assumptions: (i) relevance — encouragement must increase adoption, which we confirm with strong first-stage effects already shown in Table A2; (ii) independence — random assignment ensures encouragement is uncorrelated with unobserved determinants of outcomes, which holds given our block randomization strategy (Appendix A.4.2); (iii) exclusion — encouragement should affect electricity use only through tariff adoption, which we preserve by pooling IO Go and Octopus Go as the treatment; and (iv) monotonicity — no customers should be less likely to adopt when encouraged, which is plausible given the nature of the intervention. Under these conditions, the IV estimates identify the Local Average Treatment Effect (LATE): the causal effect of adoption for customers who take up a tariff if they receive the encouragement.[26]

We found that adoption of EV tariff significantly shifted consumption from peak to off-peak hours. When we estimate Equation (4) over all hours within the specified peak and off-peak windows, our main specification shows that peak-period usage fell by 42% (0.581 kWh average hourly reduction), while off-peak usage rose by 50% (0.481 kWh

[25] This deviates from our pre-analysis plan, reflecting our uncertainty at the outset of the trial about whether we would face this exclusion restriction violation.

[26] The LATE identifies the effect of IO Go for compliers, not the population average treatment effect. Our randomized encouragement generates exogenous variation in assignment, but treatment can only be estimated for those who comply with encouragement. Compliance itself involves multiple stages—such as opening and reading the email—and these decisions may correlate with unobserved characteristics.

average hourly increase). This is a substantial reallocation of demand rather than an increase in overall consumption. Consistent with this interpretation, Table A9 reports no change in total electricity use. This load shifting pattern is further illustrated in Figure 6, which shows consumption estimates, split by hour-of-day.

Figure 6: Impact of EV Tariff on Electricity Consumption



Notes: This figure reports IV estimates of the effect of adopting an EV tariff on electricity consumption (kWh), split by hour-of-day. The instrument is an indicator for assignment to any email-based encouragement. Each panel reports a different specification: (1) is our main specification (Equation (4)); (2) defines a separate instrument for each encouragement group; (3) restricts to just the £0/Month group and the control group; (4) restricts to just the £50/Month group and the control group; (5) just the £50/Month (No Bump) group and the control group. All specifications control for baseline consumption, and fixed effects for randomization block and week. Lines depict 90% (dark) and 95% (light) confidence intervals. Standard errors are clustered by participant and week. Green shaded box indicates IO Go off-peak hours (23:30–05:30), when electricity is charged at £0.07/kWh; all other hours are billed at the standard variable rate. Blue shaded box indicates typical system peak hours (16:30–20:30), which are highlighted to show times of heightened grid stress. Our outcome measure is hourly consumption, with hours defined as starting on the half-hour to align with IO Go's pricing structure. Percentages represent average treatment effects as a share of the control group trial participants who are not on an EV tariff, for off-peak (23:30–05:30, green) and peak (16:30–20:30, blue) periods. Estimates come from regressions pooling all hours in each window, as reported Tables A7 and A8.

Our preferred specification uses a single binary instrument for assignment to any of the four encouragements. Combining the four encouragement groups into a single binary instrument mitigates concerns raised in Mogstad et al. (2021), particularly the potential for negative weights being assigned to individual instruments when multiple instruments are used. To assess robustness, Figure 6 also shows (i) a specification that includes the four encouragement indicators as separate instruments and (ii) 4 additional specifica-

tions that use each encouragement indicator as a stand-alone instrument. The estimated effects are similar across specifications, indicating that our findings are not sensitive to the instrument definition.[27]

As a comparison, the consumption profile of the baseline group – control households who did not adopt an EV tariff – exhibits a pronounced peak beginning around 16:30, consistent with typical residential demand patterns and uncoordinated EV charging behavior (see Figure 7). Using the estimated treatment effects from the main specification in Figure 6, we overlay the causal impact of EV tariff adoption onto the baseline profile. This constructed profile illustrates how the EV tariff shifts electricity demand away from the evening peak and toward the designated overnight off-peak period. The resulting pattern flattens the peak-hour hump and concentrates usage during hours when electricity is less expensive and there is less grid stress, underscoring the potential of managed charging to reshape intraday load without increasing total consumption.

We find some evidence of selection into IO Go takeup, but no evidence that it operates through selection on levels, i.e., expected winnings. In a logit specification of IO Go takeup on encouragement assignment, expected structural winnings, and baseline covariates, the coefficient on expected winnings is negative (i.e., in the opposite direction to what one might expect if expecting selection on winnings), but small. We provide the results as the marginal effects at the means of the covariates in Table A6. Take-up is more strongly associated with socioeconomic status: households in the second and third terciles of the Index of Multiple Deprivation are more likely to enroll than those in the most deprived tercile.

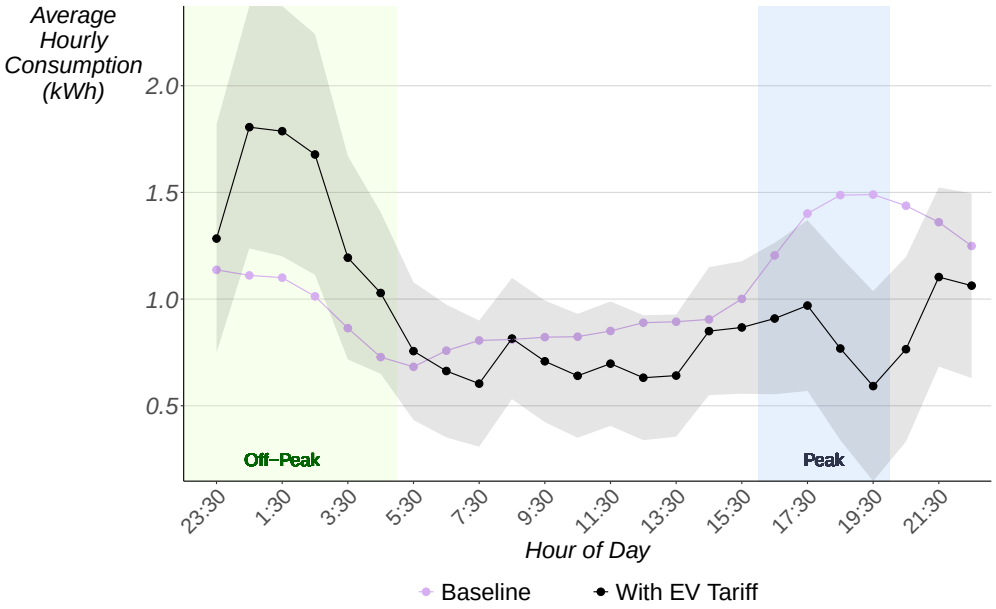
3.4 Heterogeneity

We explored heterogeneity in treatment effects across consumers by interacting the encouragement treatment with two key variables: (1) the Index of Multiple Deprivation (IMD) and (2) baseline electricity consumption. In these specifications, we treat interaction terms (e.g., EV tariff × baseline covariate) as endogenous and instrument them with the corresponding interaction of the randomized encouragement and the covariate, consistent with standard IV practice.

The IMD is a composite measure that captures multiple dimensions of deprivation

[27] Note that this preferred specification deviates from our pre-analysis plan of using the encouragement arms as four separate instruments. However, results from the pre-specified analysis, which includes all four encouragement indicators in the regression, are presented in the second panel of Figure 6.

Figure 7: Electricity Consumption With and Without EV Tariff Adoption



Notes: This figure displays mean hourly electricity consumption (kWh) for trial participants in the control group who were not enrolled in the EV tariff. The purple line ("Baseline") plots their observed consumption. The black line ("With EV Tariff") adds the estimated hourly treatment effects of EV tariff adoption, recovered from the IV analysis displayed in the first panel of Figure 6. The shaded area denotes the 95% confidence interval for the treatment effect estimates.

(e.g., crime, housing barriers, health) for small geographic areas, weighted to produce an overall deprivation score. For this analysis, we constructed IMD terciles by linking trial participants' meter-point-level postcodes to area-level IMD ranks.[28] Socioeconomic status shapes Socioeconomic status shapes both the ability to adopt EVs and the potential financial gains from managed charging. Households facing financial hardship could in theory benefit most from cheaper charging, but structural barriers may prevent uptake. Understanding these dynamics requires analyzing effects across the socioeconomic gradient.

For baseline electricity consumption, we computed the total kWh used per customer by aggregating all available half-hourly smart meter readings per day over the period from February 15, 2023, to August 31, 2023. Households with high baseline electricity consumption may have greater flexibility to shift charging, since larger batteries or multiple EVs provide more scope to delay without running short of range. At the same time, their greater and more time-sensitive energy needs can reduce flexibility, leaving less room to adjust without disrupting routines. This dimension highlights whether managed charging works chiefly for high-demand users or more broadly across households,

[28] We discretize the IMD into terciles to avoid very small cell sizes in the most deprived group.

shaping expectations about scalability and future generalizability.

We observed higher rates of tariff uptake in IMD terciles 2 and 3, with the greatest reduction in peak electricity consumption concentrated in the middle (second) tercile (Figure A5). These patterns, while suggestive of differential responsiveness by socioeconomic status, should be interpreted with caution given the limited precision of our subgroup estimates. Notably, only 9% of trial participants in our sample reside in the most deprived (bottom) IMD tercile, reflecting the well-established association between EV ownership and higher socioeconomic status (Figure A5c). As a result, our ability to detect precise effects among more deprived trial participants is limited by smaller sample sizes.

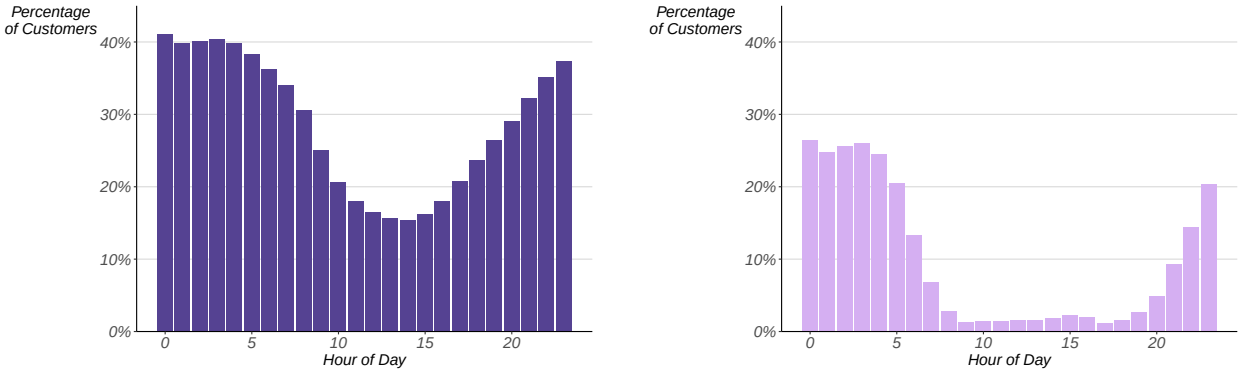
We observed additional heterogeneity in take-up by baseline electricity use. Households with higher pre-trial consumption are less likely to adopt the IO Go tariff, particularly under the £50/Month (No Bump) condition (Figure A6). This pattern suggests that higher-consuming households may be more reluctant to accept constraints on their charging behavior, potentially due to greater perceived disruption to their routines. However, conditional on adoption, we found little evidence that baseline consumption moderates the impact of IO Go on electricity usage; consumption effects are broadly similar across consumption terciles.

3.5 Intelligent Octopus Go user behavior

Using administrative data on trial participant settings and behaviors under the IO Go tariff, we explored how users from our sample engaged with automated EV charging. Two key patterns emerged: strong adherence to the automation schedule and a high degree of behavioral consistency across users. Together, these findings suggest that managed charging was generally well integrated into users' routines with minimal disruption.

We begin with evidence of adherence to the automation schedule, which provides a useful proxy for satisfaction with the tariff. We analyze 2,359 IO Go participants' use of the "bump" function, which allows them to override the default schedule and initiate immediate charging. We find that 55% of users never used the bump feature at all (Figure A10), and bump events accounted for only 2.3% of total IO Go electricity consumption. These infrequent overrides, consistent across treatment groups, indicate that the AI managed charging schedule is widely accepted and rarely disrupted. This is also consistent with what we have seen in Figure 4, where the sustained uptake of IO Go

Figure 8: Hourly Patterns of EV Plug-In and Charging Behavior



(a) Proportion Plugged In

(b) Proportion Charging

Notes: Panel (a) plots, by hour of day, the percentage of active users who had their EVs plugged in. Panel (b) plots, by hour of day, the percentage of active users who were charging. A user is considered active during a given hour if that hour falls between their first-ever and last-ever recorded plug-in event. All panels are averaged across all days in the study year. The analysis is based on data from 2,359 IO Go customers.

suggests that trial participants generally accepted the automated approach.

Preferences for charging settings further reinforce this uniformity. 61% of trial participants preferred their vehicle to finish charging between 07:00 and 09:00, and 93% set their desired state-of-charge (SOC) at 80% or higher (Figure A7a). Plug-in patterns also followed a predictable rhythm: more than half of plug-in events occur within 24 hours of the previous one, typically following a post-work return home and preceding the next morning's commute (Figure A7b).

If customers were to begin charging as soon as they plugged in, typically between 5:00 and 7:00pm, it would place significant strain on peak demand. Managed charging avoids this issue by decoupling plug-in time from charging time; customers still enjoy the convenience of plugging in when they arrive home, while AI managed scheduling shifts the actual charging to off-peak hours. This is illustrated in Figure 8. Between 20:00 and 08:00, about 30% of trial participants have their vehicles plugged in at any given time (Figure 8a), yet only around half are actively charging (Figure 8b). This idle plug-in time creates a reliable opportunity to implement demand-side strategies that serve both grid balancing and customer convenience.

Importantly, these behaviors are all consistent across our encouragement groups, suggesting that once trial participants opt into the tariff, they tend to use it in similar ways (see Figures A8, A9 and A11). This behavioral consistency is mirrored in the relatively homogeneous consumption impacts shown in Figure 6.

3.6 Value of AI managed charging beyond time-of-use tariff structure

The IO Go and Octopus Go tariffs both feature time-of-use pricing, but, importantly, IO Go includes AI managed charging, whereas Octopus Go relies on customers manually adjusting behavior in response to the tariff's day and off-peak rates or setting up their own automated charger schedules to align with the tariff (but note that we do not observe settings from these trial participants). To understand the added value of AI managed charging, we estimated how consumption under IO Go responds to the real-time system price, relative to consumption under Octopus Go.[29] While neither IO Go nor Octopus Go participants were directly exposed to system prices, the AI management of IO Go charging schedules responds to those prices on the retailer's behalf. Thus, the observed difference in elasticity reflects the supplier's algorithmic responsiveness, not household-level reactions to system costs.

Our analysis examined responsiveness to real-time system prices across three periods of the day: daytime (5:30–16:30), evening (16:30–23:30), and overnight off-peak (23:30–5:30). We examined responsiveness using a Poisson regression, which allows for a direct interpretation of the coefficients as elasticities of electricity consumption with respect to system prices, and allows for instances of zero consumption. Table 3 reports results from the following specification:

$$\log(\mathbb{E}[Y_{it}]) = \alpha + \beta_1 D_{it} + \beta_2 \log(P_t) + \beta_3 [D_{it} \times \log(P_t)] + \mu_d \tag{5}$$

where Y_{it} is electricity consumption for trial participant i at date-hour t, D_{it} is an indicator for IO Go participation, and P_t is the system price. The coefficient β_2 captures the price elasticity of demand among the baseline group (Octopus Go users), while β_3 captures the differential elasticity for IO Go users. We also include fixed effects for day d, so that coefficients are identified from within-day variation across participants. The sample in this analysis comprises the subset of our original trial sample (n = 13,233) who signed up for either IO Go or Octopus Go (n = 2,963).[30]

Given that tariff assignment is not random, this analysis should be interpreted as suggestive rather than causal. Many trial participants on Octopus Go are unable to enroll in IO Go due to compatibility constraints rather than personal preference, limiting selection

[29] In Great Britain, the system price is the market-wide wholesale price of electricity settled every half hour. It reflects the cost of balancing supply and demand on the grid and is published by the National Energy System Operator (NESO).

[30] This analysis was not pre-specified. We have included it as an exploratory analysis to help to isolate the automation-related mechanisms.

into IO Go to some extent; this can be seen in Table A13, which shows that IO Go users have made more attempts to enroll in Octopus smart tariffs prior to the encouragement emails. That said, some Octopus Go customers may actively choose not to enroll in IO Go to retain autonomy over their charging schedule. Even where this is not the case, tariff choice remains mechanically correlated with vehicle and charger type.

Overall we find within-period differences in the price elasticity of demand from IO Go and Octopus Go customers – where, again, note that the price is the system price that consumers themselves never see, but rather is a key input the AI uses when scheduling charging. During the evening period that encompasses peak hours (16:30–23:30), trial participants on IO Go exhibit significantly greater price responsiveness than Go customers : a -0.044 additional price elasticity of demand. During the overnight off-peak window (23:30–05:30), the price elasticity of demand is again significantly greater for IO Go customers than Octopus Go customers (by -0.024). By contrast, during daytime hours (05:30–16:30), the interaction term is statistically indistinguishable from zero, suggesting no meaningful difference in price responsiveness between the two groups, possibly due to less vehicles plugged in during these hours (and thus less available charge to shift).

When pooling hours across the full day, there is no evidence that IO Go participants systematically shift consumption away from periods of high system prices toward lower-priced periods (Column (4) of Table 1). If anything, Octopus Go customers appear to consume relatively more when prices are low, a counterintuitive pattern that is difficult to reconcile. One explanation is selection: Octopus Go adopters are those who received the IO Go encouragement but selected into Octopus Go, without any monetary incentive. These participants may be more motivated or attentive than the incentivized IO Go participants. It is also important to emphasize that these estimates reflect total household consumption, not just EV charging. Thus, the apparent responsiveness of Octopus Go users could reflect shifts in non-EV household load to overnight hours, rather than differences in automated charging. Unfortunately, we lack telemetry data for Octopus Go customers to directly test this mechanism.

Within the defined periods, the AI managed scheduling feature embedded in IO Go appears to dynamically shift consumption away from periods of high system prices more aggressively than trial participants relying on manual. Here, the enhanced responsiveness of IO Go customers likely reflects the supplier's ability to algorithmically optimize charging schedules when trial participants are most likely to be plugged in. In contrast, Octopus Go customers receive a static overnight rate and must manually determine their

charging behavior or set their charger settings accordingly without any AI optimizing on price in the background, leading to flatter responsiveness within-periods.

Table 1: Responsiveness to Systems Prices

period	5:30-16:30	16:30-23:30	23:30-5:30	All Hours
Model:	(1)	(2)	(3)	(4)
Variables				
D_IO_GO_onTRUE	0.075**	0.077**	-0.119***	0.089***
	(0.030)	(0.031)	(0.036)	(0.025)
log(Price)	0.047***	-0.003	0.010	-0.125***
	(0.008)	(0.010)	(0.013)	(0.010)
D_IO_GO_onTRUE × log(Price)	-0.004	-0.041***	-0.021*	0.017**
	(0.009)	(0.011)	(0.011)	(0.008)
Fixed-effects				
date	Yes	Yes	Yes	Yes
Fit statistics				
Octopus Go Mean	1.08	0.656	0.968	1.79
Observations	6,178,064	4,489,911	4,257,014	14,924,989

Clustered (User & date-hour) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: This table reports estimates from Equation (5), examining how average hourly electricity consumption under IO Go and Octopus Go responds to system prices. Column (1) covers daytime hours (05:30–16:30), column (2) evening (16:30–23:30), column (3) overnight off-peak (23:30–05:30), and column (4) pools together hours across the whole day. All regressions include day fixed effects, so coefficients are identified from within-day variation across participants; standard errors are clustered by user and hour.

Taken together, these results indicate that managed charging under IO Go not only shifted demand away from the peak, but also redistributed it within non-peak periods in a way that was responsive to real-time system costs. The automation embedded in IO Go appeared to improve response signals. These findings suggest that managed charging can outperform conventional time-of-use tariffs in aligning household electricity consumption with the dynamic needs of the grid.

4 Alternative estimation strategy: difference-in-differences

In addition to our experimental design, we leveraged a differences-in-differences (DiD) approach using observational data to estimate the impact of adopting IO Go on household electricity consumption. This strategy exploited the staggered, voluntary adoption of IO Go across Octopus Energy customers in 2023.

Comparing estimates from the DiD to RCT allows us to contrast the behavior of voluntary early adopters of managed charging (captured by the DiD) with that of harder-to-recruit individuals who required external encouragement to adopt (captured by the RCT). The RCT sample may better reflect the future mainstream population, who are not

proactively engaged, but require more pricing and marketing interventions to adopt managed charging. Some of the observed differences may also reflect other factors: changes to the Intelligent Octopus algorithm over time, evolving day-ahead price profiles (particularly as the volatility of the energy crisis subsided), or methodological differences, such as the potential for selection bias in the DiD estimates, as highlighted in critiques of observational methods (LaLonde, 1986; Imbens and Xu, 2024).

4.1 Empirical strategy

We began with a sample of 100,986 customers who adopted Intelligent Octopus Go (IO Go) at some point in 2023. To isolate the effect of IO Go from other contemporaneous changes, we further restricted the sample to customers who likely already owned an EV by December 2022, using the methodology outlined in Section 2.1. This ensured that observed changes in consumption patterns are due to changes in charging behavior, rather than the initial uptake of EVs. These restrictions result in a sample of 18,192 customers.[31]

We implemented a standard event-study difference-in-differences estimator, allowing for staggered adoption and dynamic treatment effects, following Callaway and Sant'Anna (2021). We make a parallel trends assumption based on "not-yet-treated" units. For each group of units first treated in week g, and for each week $t \geq g$, we defined the group-time average treatment effect on the treated (ATT) as:

$$ATT(g,t) = \mathbb{E}[Y_t - Y_{g-1} \mid G = g] - \mathbb{E}[Y_t - Y_{g-1} \mid D_t = 0,\ G \neq g] \tag{6}$$

where $G = g$ denotes the cohort of units first treated at time g. We estimate both (1) aggregate group-time effects, and (2) a single post-treatment estimate, constructed as the weighted average of all group-time ATT estimates, with weights proportional to group size. For both pre-treatment and post-treatment effects, we use a baseline period of -1, the week immediately before treatment (Roth, 2024).[32]

To enhance comparability, we limited control cohorts to those scheduled to adopt IO

[31] This restriction was not specified in our pre-analysis plan. However, preliminary analysis revealed that failing to condition on EV ownership by December 2022 would conflate the effects of tariff adoption with those of initial EV uptake, thereby biasing our estimates of charging behavior.

[32] We assume that once a customer first adopts IO Go, they remain "treated" in the sense that their experience with the tariff continues to shape their behavior, even if they subsequently switch to another Octopus tariff (Callaway and Sant'Anna, 2021). In practice, some customers do have more complex tariff histories. Under the irreversibility assumption, their electricity consumption patterns are considered to remain influenced by IO Go from the point of initial adoption. We view this as reasonable for two reasons: (1) 90% of customers who adopt IO Go subsequently remain on it for the entire research period, and (2) it is plausible that IO Go induces some degree of habit formation, both in EV charging routines and in household electricity use more broadly.

Go no later than eight weeks after the treated group's anticipation period ends. This ensures treated units are compared only to future adopters with similar adoption timing. We chose an eight-week window to balance comparability of treated and control groups against the length of the post-adoption estimation horizon. Comparability was assessed by examining pre-treatment trends, and we selected the longest horizon that yielded satisfactory pre-trend balance. Our treatment assessment therefore relies on the following parallel trends assumption: absent adoption, treated and not-yet-treated households would have experienced similar trends in electricity use. We assess the plausibility of this assumption by examining pre-treatment event-study coefficients, which show no systematic differences across groups.

In addition, to improve comparability with the RCT estimates and probe underlying mechanisms, we estimated weighted versions of Equation 6, reweighting the DiD sample match the RCT sample based on pre-treatment tariff type. The RCT explicitly tried to exclude customers with any prior smart tariff usage.[33] Since smart tariffs incorporate time-of-use pricing structures, this exclusion disproportionately removed time-of-use tariff users from the RCT sample. As a result, only 14% of RCT participants were on a time-of-use tariff at baseline, compared to 75% in the DiD sample.[34] To align the tariff composition across the two groups, we calculated the baseline shares of standard versus time-of-use tariff users in each sample. We then reweighted the DiD observations by the ratio of RCT to DiD shares, ensuring that the reweighted DiD sample better reflected the RCT's pre-treatment tariff distribution. [35]

4.2 Difference-in-differences results

Figure 9 shows that the difference-in-differences estimates are notably smaller than those from the RCT. In the unweighted specification (Column 1), IO Go adoption is associated with a 0.05 kWh (8%) average hourly reduction in peak-period consumption, and

[33] This exclusion was not perfect; a small number of customers who previously had smart tariffs *were* part of our trial sample. Smart tariffs are tariffs that require smart meters because their half-hourly unit rate changes.

[34] For historical reasons, there are a handful of time-of-use tariffs that are not "smart" tariffs; the most well-known of which is called "Economy 7", a tariff introduced in the 1970s to incentivize overnight electricity use, particularly for storage heaters by offering cheaper rates during a fixed seven-hour off-peak window. In recent years, some EV owners also adopted Economy 7 as a way to charge their vehicles at lower cost.

[35] We also implemented propensity-score reweighting. Specifically, we estimated the probability of being in the RCT (vs. DiD) sample using a logit regression with the following covariates: (1) tariff type prior to treatment, (2) total electricity consumption in December before the study period (December 2022 for DiD, December 2023 for RCT), (3) the share of consumption occurring during peak hours, (4) Octopus tenure, (5) IMD rank, and (6) property value. The resulting propensity scores were then used as weights in the DiD regression. However, we found that only the pre-treatment time-of-use tariff indicator had a substantive effect on the results. Given this, we opted to show the results only of the simpler tariff-based reweighting approach described in the main text.

a 0.202 kWh (12%) increase during off-peak hours. In contrast, the RCT estimates imply much larger shifts: a 0.581 kWh (42%) decrease during peak periods and a 0.481 kWh (50%) increase off-peak.

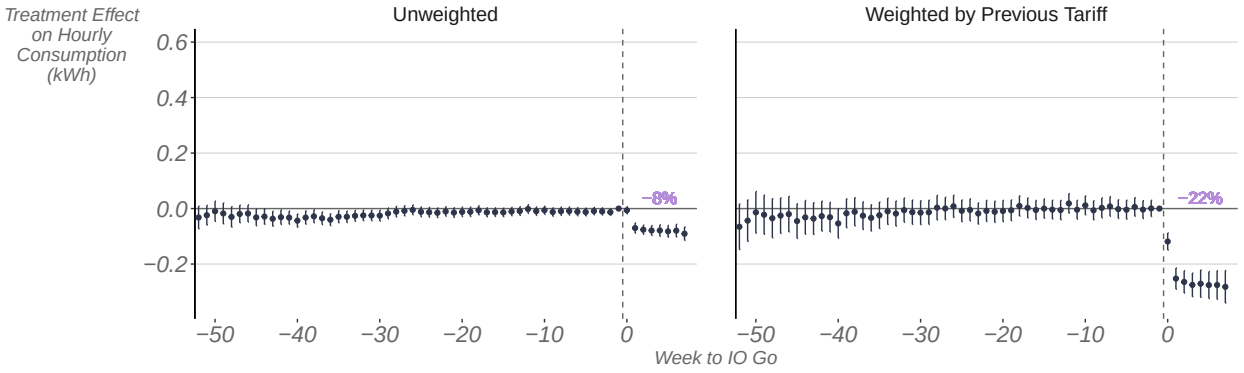
This discrepancy appears to be largely explained by differences in baseline time-of-use tariff usage. Column 2 presents the results after reweighting the DiD sample to match the RCT sample's pre-treatment tariff distribution. After reweighting, the estimated increase in off-peak consumption in the DiD analysis (0.485 kWh) closely matches the RCT analysis (0.481 kWh). The reduction in peak consumption is considerably smaller in the DiD analysis: 0.242 kWh compared to 0.581 kWh in the RCT. This weaker peak effect is compensated for by a decline in daytime, non-peak consumption in the DiD sample, as shown in Figure A14. We hypothesize that the observed differences in the timing of impacts throughout the day stem from variation in contemporaneous electricity prices and the IO Go algorithm's response to them. The DiD analysis covers data from 2023, while the RCT analysis reflects 2024 conditions. In 2023, afternoon prices were higher relative to overnight prices than in 2024, implying more need in 2023 to reduce afternoon (and not just evening peak) consumption. Relevant day ahead price trends can be seen in Figure A16.

Thus the RCT targeted a sample whose characteristics made their consumption more responsive to IO Go's optimization – i.e., because their charging behavior was not previously aligned with dynamic or off-peak pricing, leaving more scope for managed charging to change consumption in both peak and off-peak hours. We also estimate cohort-specific treatment effects, defining cohorts by the week of adoption as shown in Figure A15; we find that treatment effects are relatively homogeneous across cohorts.

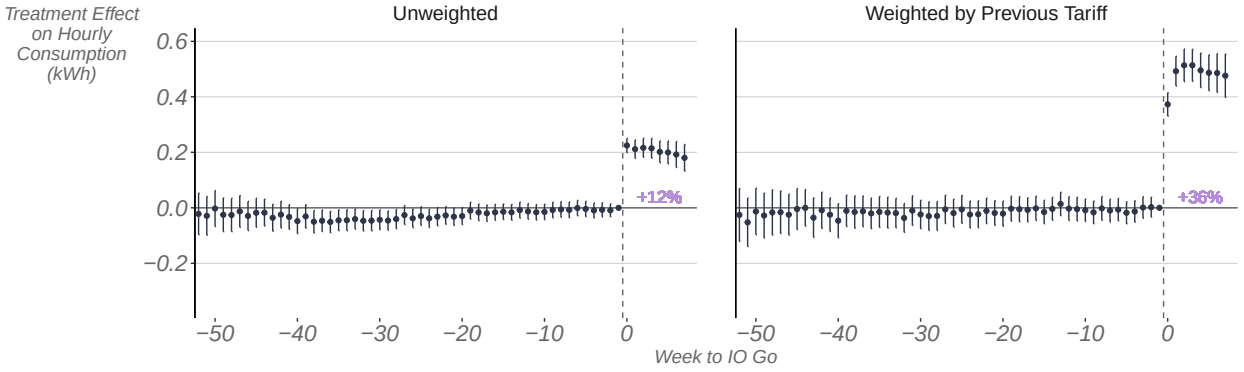
Taken together, this homogeneity in treatment effects across cohorts, combined with the close alignment of RCT and reweighted DiD estimates, suggests that the impact of IO Go is relatively stable across adopters. The primary source of variation appears to be baseline charging behavior, particularly whether customers were already on time-of-use tariffs prior to adoption, rather than any inherent heterogeneity in responsiveness to managed charging.

Figure 9: Difference-in-Differences Estimate of IO Go

(a) Peak



(b) Off-Peak



Notes: This figure reports staggered difference-in-differences estimates of the effect of adopting the IO Go tariff on hourly electricity consumption (in kWh) during (a) peak hours (16:30–20:30) and (b) off-peak hours (23:30–05:30), using a sample of 18,192 customers who first-ever enrolled in IO Go in 2023. Each panel plots treatment effects relative to the week before adoption. Estimates are reported under two specifications: (i) unweighted; (ii) and weighted by whether the trial participant was previously on a time-of-use tariff. Estimates are computed using the Callaway and Sant'Anna (2021) estimator. Percentages represent post-treatment effects as share of the pre-IO Go consumption levels. Post-treatment effects are estimated using average of all group-time average treatment effects, with weights proportional to the group size.

5 Welfare impacts

The large shifts of consumption from peak to off-peak hours, as seen in Section 3.3, have four potential benefits. We have (1) benefits to consumers from lower electricity bills, (2) lower electricity procurement costs to the electricity supplier, (3) climate change mitigation benefits via CO_2e abatement, and (4) reduced grid operation and stabilization costs. The CO_2e abatement arises via two pathways: directly, through reduced consumption during high-emissions hours (a shift we identified using our randomized encourage-

ment design); and indirectly, through potential substitution from internal combustion engine (ICE) vehicles to EVs caused by lower electricity bills (also identified through our randomized encouragement) that reduce the lifetime cost of EV ownership. We believe this latter behavioral response may be important, but we acknowledge that it is uncertain and speculative.

In this section, we first outline our methods for estimating the various benefits and costs of IO Go adoption. We then present the direct benefits accrued in 2024. Finally, we apply the Marginal Value of Public Funds (MVPF) framework developed by Hendren and Sprung-Keyser (2022) to assess the welfare implications of subsidizing managed charging.

5.1 Estimation

To estimate the magnitude of these benefits, we combined administrative data from Octopus Energy with external inputs and applied outcome-specific methods. For consumer benefits, we used administrative data on daily average electricity bills per kWh. We analyzed the impact on trial participants' bills using an instrumental variables specification analogous to Equation (4), but where Y_{iht} is the mean daily bill (in £) for user i on day t. Our outcome is amount customers paid per kwh of electricity, and we weight the regression by the consumption each day.

For supplier procurement costs, we applied an analogous IV. Octopus Energy provided administrative cost data, comprising of both the wholesale price of energy as well as non-energy charges such as transmission fees, distribution fees, and capacity payments. Here, the outcome is Octopus Energy's total daily procurement cost.

To assess direct emissions impacts from shifting consumption to lower-CO_2e-intensity hours, we multiplied trial participant electricity consumption in each half-hour interval by the corresponding average Marginal Operating Emissions Rate from WattTime, and aggregated these values to the daily level.[36] We analyzed the impact on CO_2e from electricity consumption using an instrumental variables specification analogous to Equation (4), but where Y_{iht} is CO_2e from electricity consumption (in grams of CO_2e) for trial participant i on day t. To estimate the potential CO_2e abatement benefits from induced substitution of internal combustion engine (ICE) vehicles to EVs, we combined IO Go bill

[36] As described in (as described in Section 2.6, the Marginal Operating Emissions Rate is the emissions associated with the marginal change in load on the grid (WattTime, 2022).

savings with existing estimates of EV price sensitivity.[37]

Finally, managed EV charging may contribute to avoiding grid costs, though the magnitude of these benefits is uncertain. Some top-down modeling studies suggest potentially large system-wide gains. We discuss the estimates from one such study in the contexts of our results in the next section. [38]

5.2 Benefits/savings from adoption of managed charging

AI managed charging presented large consumer bill benefits. Figure 10 presents the estimated benefits. We found that managed charging reduces consumer electricity bills by £0.0397 per kWh, an approximately 18% decline relative to the control group. Applied to average consumption over the analysis period, this implies a total bill reduction of approximately £343 during our study.[39]

The electricity retailer saw similar savings in procurement costs (procurement costs include the wholesale price of power and non-energy charges such as transmission and distribution fees), implying near 100% pass-through of savings to trial participants, at least based on the period of our analysis (2024).

Managed charging also yielded positive environmental benefits. Direct CO_2e abatement from shifting electricity consumption to cleaner off-peak hours generated a decrease of 124kg CO_2e per trial participant in 2024, which translates to estimated benefits of £35 based on a carbon value of £287 per tonne of CO_2e emitted (Department for Energy Security and Net Zero, 2023), though with wide uncertainty (95% CI: £-253 to £287). Indirect CO_2e abatement, arising from substitution of ICE vehicles with EVs due to lower operating costs of EVs,[40] generated a further decrease of 143kg CO_2e, or £40.9 (95% CI: £20.6

[37] We began by estimating the average lifetime electricity bill savings from IO Go over a ten-year vehicle lifespan, discounted to present value. We treated this as a reduction in the total cost of EV ownership. Applying a price elasticity of EV demand of -2.547 from Hahn et al. (2024), we inferred the corresponding proportional increase in EV adoption. To translate this into absolute uptake, we used estimates from Department for Transport (2024) that approximately 6.7% of UK households (1.95m of 28.8m households in the UK) purchase a new car each year. The resulting increase in EV uptake was multiplied by the difference in lifecycle CO_2e emissions between ICE vehicles (£8003.89) and EVs (£3259.66) (Hahn et al., 2024), yielding £4744.23 in CO_2e benefits per induced switch. We scaled these annual impacts using government projections of ICE vehicle sales, which decline over time (Department for Transport, 2023), and discounted future abatement to present value using a 3.5% rate recommended by HM Treasury (2020).

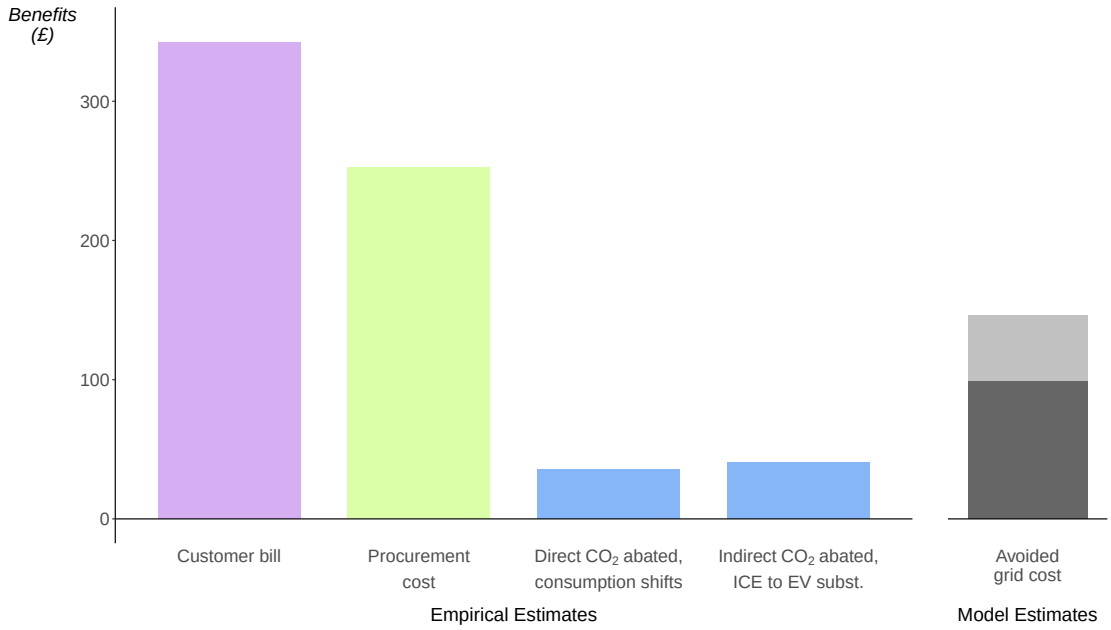
[38] In our pre-analysis plan, we pre-specified estimating CO_2e impacts using the ITT framework. We did not pre-specify the use of IV estimation for bills or CO_2e savings. We adopt the IV approach here because it more directly captures the causal effect of IO Go adoption—the quantity of substantive interest. While our pre-analysis plan focused on MVPF calculations rather than consumer bills, we now report bill savings as well, as they provide an important and policy-relevant measure of consumer benefits.

[39] Based on the average annual electricity consumption of control group trial participants not enrolled in an EV tariff (9,063 kWh).

[40] We are assuming that consumers value fuel efficiency in their decisions to buy a vehicle, which has some support Grigolon et al. (2018); Forsythe et al. (2023).

to £61.2), in estimated benefits.[41] From a resource cost per tonne of CO_2e abated, we find a lower bound resource cost per tonne of around -£2,445 for switching to the AI managed EV tariff, which is an order of magnitude lower than the next best technology (Gosnell et al., 2020; Hahn et al., 2024). If we only use the standard tariff as the counterfactual, the resource cost per CO_2e tonne abated is -£3253.

Figure 10: Benefits of Adopting EV Tariff per Household in 2024



Notes: This figure presents estimated benefits in 2024 of adopting an EV tariff. Trial participant bill savings are derived from causal estimates using administrative data on daily average electricity bills per kWh; the estimated savings are £343 per vehicle (95% CI: £173 to £512). Procurement cost savings are estimated at £237 (95% CI: £-261 to £766). Direct CO_2e abatement reflects emissions reductions from shifting electricity consumption to off-peak periods; these are valued at £35 (95% CI: £-253 to £324), using half-hourly marginal emissions intensity data matched to observed load-shifting. Indirect CO_2 abatement from ICE to EV substitution is valued at £40.9 (95% CI: £20.6 to £61.2), estimated by combining observed IO Go bill savings with price elasticity-based projections of EV adoption and associated emissions reductions. Grid benefits are sourced from Franken et al. (2025), showing per-vehicle value under a scenario of 100% smart charging adoption. The shaded area illustrates the range of estimates across multiple modeled scenarios.

For savings from avoided grid costs, we looked in particular at Franken et al. (2025), who estimated that fully flexible EV-related electricity demand, relative to a baseline with no flexible load, could reduce annual system costs by up to £0.25 billion in 2025 and as much as £4 billion by 2035.[42] These savings arose from both operational efficiencies, such as increased use of lower-marginal-cost renewable generation, and capital

[41] The confidence intervals for indirect CO_2e abatement given in Figure 10 reflect only the uncertainty from the regression of consumer bill savings on EV tariff adoption. This likely understates the true uncertainty, which is outside the scope of this research. Additional sources of uncertainty include the price elasticity of EV adoption, the future cost trajectory of EVs, and the relationship between EV adoption and net CO_2e damages.

[42] Their estimate was based on a whole-system linear cost optimization model, calibrated to assumptions used by the UK system operator. To the best of our knowledge, Franken et al. (2025) is the study most closely aligned with our context on two dimensions. First, it matches our outcome of interest: assessing the grid impacts of managed charging in monetary terms. By contrast, many other studies focus on EV deployment without more detailed modeling of managed charging (Heuberger et al., 2020), or in terms of electricity consumption, without translating to monetary benefits

savings, including deferred investment in firm capacity and grid infrastructure. Notably, the study found diminishing marginal returns: the first 25% of EV users adopting managed charging account for 50% (£2 billion in 2035) of the projected savings.

Our trial suggests that nearly all peak EV load can be shifted through managed charging, and that this behavior can be sustained for over a year, consistent with the assumptions in Franken et al. (2025). This implies that system-level benefits in the range of £2-4 billion could be feasible, conditional on widespread tariff uptake. When expressed on a per-vehicle basis, assuming 27 million EVs on the road by 2035 (consistent with Future Energy Scenarios from the UK's National Energy System Operator (2025)), this equates to approximately £146 in system savings per vehicle in 2035 under universal managed charging. Under "constrained" managed charging[43], savings fall modestly to £120 per vehicle. When adding flexible heat demand to the optimization, which cannibalizes some of the benefits from EV flexibility, the per-vehicle EV contribution falls to £99.

5.3 Welfare impacts of subsidizing managed charging

We next evaluated the welfare implications of *subsidizing* managed charging, using the £50/month incentive (total £150 across three months) offered to trial participants as a proxy subsidy. We considered impacts over one, five, and 10 years (2024–2033), discounting future values at the 3.5% rate recommended by HM Treasury (2020). Ten years is a reasonable lower-bound estimate of vehicle lifetime (Bento et al., 2018; Held et al., 2021; Kolli, 2011). However, the duration of benefits attributable to the subsidy also depends on how long subsidized customers remain more likely than non-subsidized customers to adopt IO Go. This compliance advantage may decay more quickly than the vehicle lifetime, although in our data the encouragement effect remains relatively stable over the full year of observation.

The relevant welfare benefits in this case are (1) partial transfer of the subsidy to consumers and (2) CO_2e abatement. To estimate the share of the subsidy that constitutes a transfer, we inferred the proportion of marginal versus inframarginal adopters. Adoption was 6.98% in the email-only group and 9.31% in the £50/month group, implying that

(Crozier et al., 2020). Second, it is aligned in the geographic focus on Great Britain. There are several relevant studies examining grid benefits of managed charging in California (Li and Jenn, 2024) or the United States more broadly (Powell et al., 2022), but to the best of our knowledge, Franken et al. (2025) provides the most directly applicable evidence for our setting in Great Britain. For an overview of studies relevant to our setting, see Thornhill and Deasley (2018).

[43] In Franken et al. (2025), the constraints are: automation only between 12 am and 4 am and 12% of consumers opt out of flexible charging each day.

25% of adopters were marginal.[44] We assumed inframarginal adopters (75%) valued the £150 incentive at its full face value (i.e., a 100% transfer of the total £150 to consumers). For marginal adopters (25%), we assumed an average valuation of 50% of the subsidy.[45] Combining these assumptions yields an average transfer of £0.875 per £1 of subsidy.

For welfare benefits from CO_2e abatement, we applied the estimated per-adoption benefit discussed above, but scaled it down to reflect that only 25% of adopters were induced by the subsidy.[46] We excluded reductions in trial participant bills from our welfare calculation, invoking the envelope theorem: these trial participants could have adopted IO Go without the subsidy but chose not to. We also assumed no producer surplus changes (e.g., procurement cost savings for Octopus Energy), under the assumption of a competitive retail electricity market. Finally, we excluded the indirect CO_2e benefits from increased EV adoption, since, again by the envelope theorem, a subsidy for managed charging should not affect EV uptake among trial participants who were already considering adoption.[47]

Changes in costs to the government include: (1) the subsidy itself, (2) lost VAT revenue from reduced electricity bills, (3) with greater uncertainty, a climate-related fiscal externality (increased tax revenue from higher economic growth due to the climate mitigation from CO_2e abatement), and, most speculatively, (4) avoided costs associated with electricity grid balancing. The VAT loss is calculated as 5% of the £343 annual reduction in electricity bills, totaling £12.42 per customer. This translates to a fiscal cost of approximately £3.10 for each marginal IO Go adopter. The climate-related fiscal externality is 1.07%[48] of the monetized value of CO_2e abatement attributable to IO Go adoption. This amounts to £3.22 per marginal IO Go adopter.

In total, these costs and benefits imply an MVPF of 0.887 over 1 year, 0.933 over 5 years, and 0.982 over 10 years.[49] For the remainder of the discussion, we focus on the

[44] Calculated as (9.31 – 6.98)/9.31 = 25%.

[45] For marginal adopters, we do not observe whether the first or last £1 of the subsidy induced adoption. If it were the first, the entire subsidy would be valued; if the last, the valuation would approach zero. Following the classic Harberger triangle approximation to deadweight loss (Harberger, 1964), and the approach in Hendren and Sprung-Keyser (2020) and Hahn et al. (2024), we assume a uniform distribution of latent subsidy valuations, consistent with a linear demand curve.

[46] We used the point estimate of the CO_2e impact, despite the statistical imprecision around that estimate.

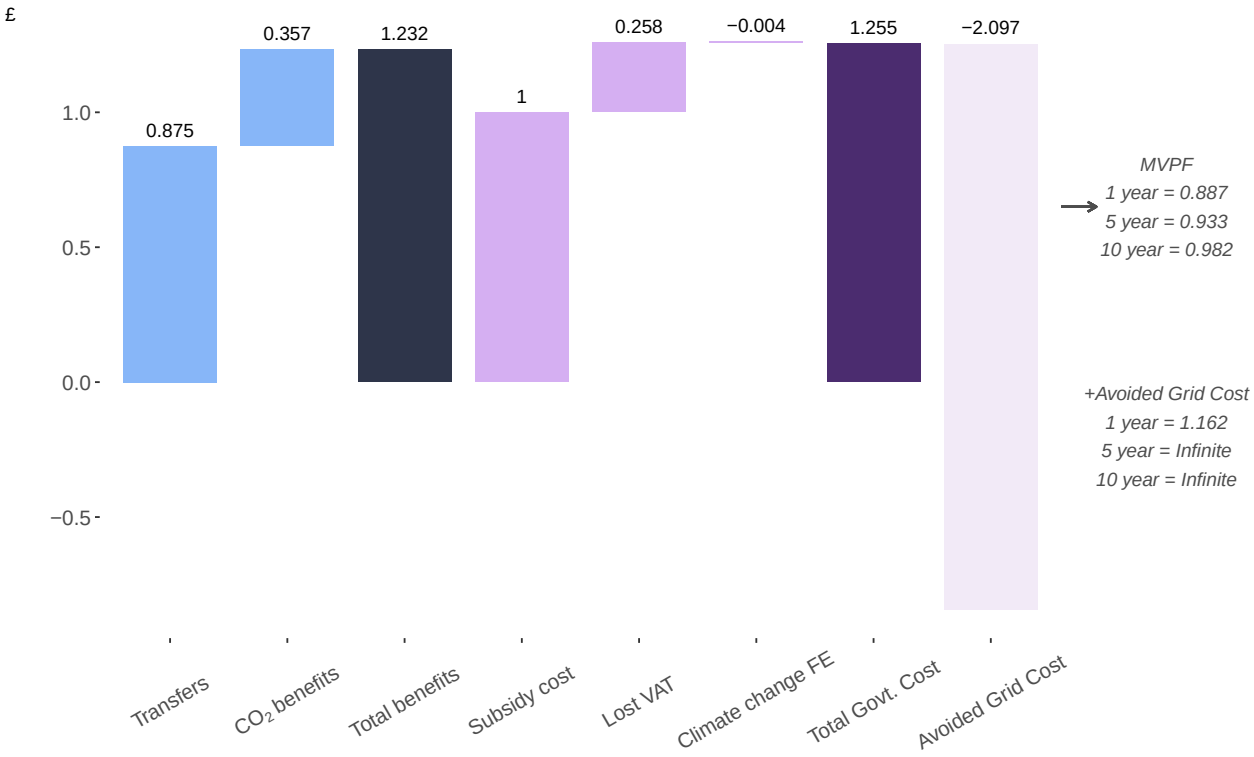
[47] This welfare calculation deviates from our pre-analysis plan, but we believe these deviations more accurately reflect the full set of social returns that would accrue under real-world implementation. For more details on exact deviations, please see Appendix A.3.

[48] We assume that the UK accounts for 3.2% of global GDP (PwC, 2024), and that 33.5% of UK GDP accrues to the government as tax revenue (Office for Budget Responsibility, 2024). The PwC estimates are a weighted average of projections from national statistical authorities, EIKON from Refinitiv, IMF, Consensus Economics, the OECD, and Fitch Solutions. We therefore use the PwC estimate as it consolidates projections from these institutions into a single figure. The product of these shares yields 1.07%.

[49] As noted above, the duration of benefits depends not just on vehicle lifetime, but also on how long subsidized

10-year estimate, while estimates for 1 and 5 years are presented in A17. This MVPF implies that for every £1 of fiscal cost to the government, the program generated £1.232 in participant benefits. The fiscal cost includes not only the direct subsidy but also the loss of VAT revenue due to lower energy bills, bringing the total cost per £1 of subsidy to £1.175. (i.e., $\frac{1.232}{1.255} = 1.173$).

Figure 11: Marginal value of public funds of subsidizing managed charging over 10 years, 2024-2033



Notes: This figure presents the estimated costs and benefits of subsidizing adoption of the IO Go managed EV charging tariff over 10 years, from 2024-2033. Customer surplus is based on a decomposition of marginal and inframarginal adoption under the £50/month offer, following the approach of Hahn et al. (2024). Direct CO_2e benefits reflect emissions reductions from shifting electricity use to cleaner hours, scaled to marginal adopters. Indirect CO_2e benefits are excluded under the assumption that AI managed charging subsidies do not affect EV uptake among inframarginal adopters. Estimated costs to government include the subsidy, lost VAT revenue, and increased tax receipts from climate-related fiscal externalities. Grid balancing benefits are shown separately, based on Franken et al. (2025) estimates of per-vehicle system savings under three scenarios. Only a share of these may accrue to government.

The change in government costs from avoided grid balancing are uncertain. They may be zero. In a well-functioning electricity market, the benefits of shifting to lower-marginal-cost generation and deferring investment in generation, transmission, and distribution infrastructure should be internalized by market participants. Thus, while the overall system benefits of IO Go may be large, only a small share of those costs would fall

customers remain more likely than non-subsidized customers to adopt IO Go.

to the government. However, if there are market failures such that the value of EV flexibility is not fully internalized, there may be a rationale for subsidizing managed charging. We find that if £70 (48%) of the £146 in per-customer benefits in 2035 identified by Franken et al. (2025) were borne by the government, the MVPF of the £150 subsidy would be infinite.[50]

6 Discussion

Our findings make three key contributions to the literature. First, using a field trial unique in its scale, we provide evidence on how real-time managed charging can reshape EV charging behavior at scale in the UK. Specifically, we found that managed charging led to a 42% reduction in household electricity demand during peak hours, with all of this demand shifted to low-cost, low-emission off-peak periods overnight. In contrast to prior observational studies or simulations, our randomized incentive design enabled causal inference about the elasticity of managed EV charging consumption to energy system price signals. This bridges a critical gap between the theoretical benefits of demand-side flexibility and the practical realities of consumer responsiveness.

These consumption impacts occurred without requiring manual input or sustained behavior change from trial participants: over half of adopters never overrode the automated schedule, and override events comprised just 2.3% of electricity use. This highlights the potential of automation through AI to unlock demand-side flexibility while respecting EV owners' preferences and constraints.

The automation embedded in IO Go appears to enhance responsiveness to grid signals, outperforming static time-of-use tariffs, especially during off-peak periods. This reinforces the argument that, as EVs move from early to mass adoption, managed charging responsive to real-time grid conditions is needed to avoid "herding" behavior and new peaks that may result from static time-of-use schedules. To reinforce this market price, network operators may consider more dynamic charges to manage congestion on low-voltage networks contending with high EV penetration.

Second, by comparing the outcomes of our randomized experiment with those from a standard difference-in-differences design, our results suggest that the impact of man-

[50] In this context, "infinite" does not mean literally unlimited welfare gains; rather, it is a formal term indicating that the government's net fiscal cost is negative, so, following Hendren and Sprung-Keyser (2020), the policy is treated as having an infinite MVPF.

aged charging is relatively stable across different types of adopters. Impacts were homogeneous across cohorts of adopters in our difference-in-differences sample, and – after reweighting our difference-in-differences sample to match the RCT sample on pre-adoption tariff, we found similar results between the two evaluations. Thus the primary source of variation in impact appears to be baseline charging behavior, particularly whether trial participants were already on smart or off-peak tariffs prior to adoption, rather than any inherent heterogeneity in responsiveness to managed charging.

Third, we quantified the welfare impacts across four dimensions – consumer costs, producer profits, environmental outcomes, and avoided costs associated with electricity grid balancing. We found that the managed charging reduced trial participants bills. For CO_2e emissions and retailer procurement costs, the estimates were imprecise and should be interpreted with caution. In a well-functioning market where the benefits of flexibility are internalized by market participants, the case for subsidizing managed charging is weak; however, where the are market failures that cause grid balancing costs to be externalized, this "grid stabilization externality" may justify such a subsidy.

These findings carry significant policy relevance. As electricity systems transition toward variable renewable generation, flexible demand from EVs presents a major opportunity for system balancing. Policymakers should ensure wholesale markets reflect the temporal and locational value of electricity and lower non-commodity electricity costs to ensure suppliers, aggregators, and other market participants are exposed to price as close as possible to the real-time marginal cost of electricity.

All in all, we have provided causal evidence that AI managed charging – when paired with real-time pricing on the retailer procurement side without direct pass-through to customers – can substantially reshape electricity consumption patterns at scale. By causally identifying behavioral responses, we have shown that, in Britain, managed charging of EVs can reduce peak load without contravening consumer preference. Our findings improve scientific understanding of the economics of energy consumption and market design, highlighting how well-structured incentives and dynamic pricing might align private behavior with policy objectives. Ultimately our study suggests that, as electrification expands in Britain and other advanced economies, managed charging can serve as a tool for aiding grid reliability and realizing environmental outcomes via a market-based framework.

References

Abel, G. A., Barclay, M. E. and Payne, R. A. (2016), 'Adjusted Indices of Multiple Deprivation to Enable Comparisons Within and Between Constituent Countries of the UK Including an Illustration Using Mortality Rates', *BMJ Open* **6**(11), e012750.

Agarwal, N., Moehring, A., Rajpurkar, P. and Salz, T. (2023), Combining human expertise with artificial intelligence: Experimental evidence from radiology, Technical report, National Bureau of Economic Research.

Aljbour, J., Wilson, T. and Patel, P. (2024), 'Powering intelligence: Analyzing artificial intelligence and data center energy consumption', *EPRI White Paper no. 3002028905* .

Allcott, H. (2011), 'Rethinking real-time electricity pricing', *Resource and energy economics* **33**(4), 820–842.

Antonopoulos, I., Robu, V., Couraud, B., Kirli, D., Norbu, S., Kiprakis, A., Flynn, D., Elizondo-Gonzalez, S. and Wattam, S. (2020), 'Artificial intelligence and machine learning approaches to energy demand-side response: A systematic review', *Renewable and Sustainable Energy Reviews* **130**, 109899.

Anwar, M. B., Muratori, M., Jadun, P., Hale, E., Bush, B., Denholm, P., Ma, O. and Podkaminer, K. (2022), 'Assessing the value of electric vehicle managed charging: a review of methodologies and results', *Energy and Environmental Science* **15**, 466–498.

Archsmith, J., Muehlegger, E. and Rapson, D. S. (2022), 'Future paths of electric vehicle adoption in the united states: Predictable determinants, obstacles, and opportunities', *Environmental and Energy Policy and the Economy* **3**, 71–110.

Asensio, O. I., Buckberg, E., Cole, C., Heeney, L., Knittel, C. R. and Stock, J. H. (2025), Charging uncertainty: Real-time charging data and electric vehicle adoption, Technical report, National Bureau of Economic Research.

Axsen, J., Langman, B. and Goldberg, S. (2017), 'Confusion of Innovations: Mainstream Consumer Perceptions and Misperceptions of Electric-Drive Vehicles and Charging Programs in Canada', *Energy Research & Social Science* **27**, 163–173.
URL: *https://linkinghub.elsevier.com/retrieve/pii/S2214629617300786*

Bailey, J. and Axsen, J. (2015), 'Anticipating pev buyers' acceptance of utility controlled charging', *Transportation Research Part A: Policy and Practice* **82**, 29–46.

Bailey, M., Brown, D. P., Shaffer, B. and Wolak, F. A. (2024), 'Centralized vs decentralized demand response: Evidence from a field experiment'.

Bailey, M. R., Brown, D., Myers, E., Shaffer, B. and Wolak, F. A. (2025), 'Electric vehicles

and the energy transition: Unintended consequences of time-of-use pricing', *American Economic Review: Insights* .

Bailey, M. R., Brown, D. P., Shaffer, B. and Wolak, F. A. (2025), 'Show Me the Money! A Field Experiment on Electric Vehicle Charge Timing', *American Economic Journal: Economic Policy* **17**(2), 259–284.

Barman, P., Dutta, L., Bordoloi, S., Kalita, A., Buragohain, P., Bharali, S. and Azzopardi, B. (2023), 'Renewable energy integration with electric vehicle technology: A review of the existing smart charging approaches', *Renewable and Sustainable Energy Reviews* **183**, 113518.

Bennett, J., Haines, T., Jones, B. C., Lave, M., Neuman, C., Scoffield, D. R. and Vining, W. (2022), 'Impact of electric vehicle customer response to time-of-use rates on distribution power grids', *Energy Reports* **8**, 8225–8235.

Bento, A. M., Roth, K. and Zuo, Y. (2018), 'Vehicle lifetime and scrappage behavior: Evidence from the used car market', *The Energy Journal* **39**(1).

Bernard, L., Hackett, A., Metcalfe, R. D., Panzone, L. A. and Schein, A. R. (2025), The impact of dynamic prices on electric vehicle public charging demand: Evidence from a nationwide natural field experiment, Technical report, Centre for Net Zero.

Biswas, P., Rashid, A., Biswas, A., Nasim, M. A. A., Chakraborty, S., Gupta, K. D. and George, R. (2024), 'Ai-driven approaches for optimizing power consumption: a comprehensive survey', *Discover Artificial Intelligence* **4**(1), 116.

Björkegren, D., Choi, J. H., Budihal, D., Sobhani, D., Garrod, O. and Atherton, P. (2025), 'Could ai leapfrog the web? evidence from teachers in sierra leone', *arXiv preprint arXiv:2502.12397* .

Black, D., Panossian, N., Jingjing, L., Nordman, B., Farrell, J., Hodge, C., Meintz, A., Abdullah, M., Kisacikoglu, M., Bennett, J. et al. (2024), 'Survey and gap prioritization of us electric vehicle charge management deployments'.

Blonz, J., Palmer, K., Wichman, C. J. and Wietelman, D. C. (2025), 'Smart thermostats, automation, and time-varying prices', *American Economic Journal: Applied Economics* **17**(1), 90–125.

Bogmans, C., Gomez-Gonzalez, P., Melina, G., Miranda-Pinto, J., Pescatori, A. and Thube, S. (2025), Power hungry: How ai will drive energy demand, Technical report, International Monetary Fund.

Boopathy, P., Liyanage, M., Deepa, N., Velavali, M., Reddy, S., Maddikunta, P. K. R., Khare, N., Gadekallu, T. R., Hwang, W.-J. and Pham, Q.-V. (2024), 'Deep learning for intelli-

gent demand response and smart grids: A comprehensive survey', *Computer science review* **51**, 100617.

Borenstein, S. (2005*a*), 'The long-run efficiency of real-time electricity pricing', *The Energy Journal* **26**(3), 93–116.

Borenstein, S. (2005*b*), 'Time-varying retail electricity prices: Theory and practice', *Electricity deregulation: choices and challenges* **4**, 317–356.

Borenstein, S. (2007), 'Customer risk from real-time retail electricity pricing: Bill volatility and hedgability', *The Energy Journal* **28**(2), 111–130.

Boyd, S. and Vandenberghe, L. (2004), *Convex Optimization*, Cambridge University Press.

Bradley, M. et al. (2018), 'Electric vehicle costbenefit analysis–plug-in electric vehicle cost-benefit analysis: Indiana'.

Brynjolfsson, E., Li, D. and Raymond, L. (2025), 'Generative ai at work', *The Quarterly Journal of Economics* p. qjae044.

Brynjolfsson, E., Rock, D. and Syverson, C. (2019), 'Artificial intelligence and the modern productivity paradox', *The economics of artificial intelligence: An agenda* **23**(2019), 23–57.

Byrne, D. P., Martin, L. A. and Nah, J. S. (2022), 'Price discrimination by negotiation: A field experiment in retail electricity', *The Quarterly Journal of Economics* **137**(4), 2499–2537.

Cahana, M., Fabra, N., Reguant-Rido, M. and Wang, J. (2022), 'This distributional impacts of real-time pricing'.

Callaway, B. and Sant'Anna, P. H. (2021), 'Difference-in-Differences with Multiple Time Periods', *Journal of Econometrics* **225**(2), 200–230.

Chen, S. (2025), 'How much energy will ai really consume? the good, the bad and the unknown', *Nature* **639**(8053), 22–24.

Colin Cameron, A. and Miller, D. L. (2015), 'A Practitioner's Guide to Cluster-Robust Inference', *Journal of Human Resources* **50**(2), 317–372.
URL: *http://jhr.uwpress.org/lookup/doi/10.3368/jhr.50.2.317*

Crozier, C., Morstyn, T. and McCulloch, M. (2020), 'The opportunity for smart charging to mitigate the impact of electric vehicles on transmission and distribution systems', *Applied Energy* **268**, 114973.
URL: *https://linkinghub.elsevier.com/retrieve/pii/S0306261920304852*

Cui, Z. K., Demirer, M., Jaffe, S., Musolff, L., Peng, S. and Salz, T. (2024), 'The effects of generative ai on high skilled work: Evidence from three field experiments with software developers', *Available at SSRN 4945566* .

De Vries, A. (2023), 'The growing energy footprint of artificial intelligence', *Joule* **7**(10), 2191–2194.

Delmonte, E., Kinnear, N., Jenkins, B. and Skippon, S. (2020), 'What Do Consumers Think of Smart Charging? Perceptions Among Actual and Potential Plug-In Electric Vehicle Adopters in the United Kingdom', *Energy Research & Social Science* **60**, 101318.
URL: *https://linkinghub.elsevier.com/retrieve/pii/S2214629619301422*

Department for Energy Security and Net Zero (2023), 'Green Book Supplementary Guidance: Valuation of Energy Use and Greenhouse Gas Emissions for Appraisal'. Published by the UK Government.
URL: *https://www.gov.uk/government/publications/valuation-of-energy-use-and-greenhouse-gas-emissions-for-appraisal*

Department for Transport (2023), Phasing out the sale of new petrol and diesel cars from 2030 and support for the zero-emission transition, Policy document, Department for Transport. Accessed 18 July 2025.

Department for Transport (2024), 'Vehicle licensing statistics, united kingdom: 2024', Official release via gov.uk. Accessed 18 July 2025.

Department of Transportation (2025), 'Phasing out sales of new petrol and diesel cars from 2030 and supporting the ZEV transition: summary of responses and joint government response - GOV.UK'.
URL: *https://www.gov.uk/government/consultations/phasing-out-sales-of-new-petrol-and-diesel-cars-from-2030-and-supporting-the-zev-transition/outcome/phasing-out-sales-of-new-petrol-and-diesel-cars-from-2030-and-supporting-the-zev-transition-summary-of-responses-and-joint-government-response*

Dorsey, J., Langer, A. and McRae, S. (2025), 'Fueling alternatives: Gas station choice and the implications for electric charging', *American Economic Journal: Economic Policy* **17**(1), 362–400.

Einav, L., Finkelstein, A., Ryan, S. P., Schrimpf, P. and Cullen, M. R. (2013), 'Selection on moral hazard in health insurance', *American Economic Review* **103**(1), 178–219.

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

Erdenesanaa, D. (2023), 'Ai could soon need as much electricity as an entire country', *The New York Times* .

European Commission (2023), 'Commission welcomes agreement on new rules to accel-

erate the deployment of zero-emission vehicles'. Press release. Available at: https://ec.europa.eu/commission/presscorner/detail/en/ip_23_48.

Fabra, N., Rapson, D., Reguant, M. and Wang, J. (2021), Estimating the elasticity to real-time pricing: evidence from the spanish electricity market, *in* 'AEA Papers and Proceedings', Vol. 111, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 425–429.

Forsythe, C. R., Gillingham, K. T., Michalek, J. J. and Whitefoot, K. S. (2023), 'Technology advancement is driving electric vehicle adoption', *Proceedings of the National Academy of Sciences* **120**(23), e2219396120.

Fowlie, M., Wolfram, C., Baylis, P., Spurlock, C. A., Todd-Blick, A. and Cappers, P. (2021), 'Default effects and follow-on behaviour: Evidence from an electricity pricing program', *The Review of Economic Studies* **88**(6), 2886–2934.

Frank, M. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., Feldman, M., Groh, M., Lobo, J., Moro, E. et al. (2019), 'Toward understanding the impact of artificial intelligence on labor', *Proceedings of the National Academy of Sciences* **116**(14), 6531–6539.

Franken, L., Hackett, A., Lizana, J., Riepin, I., Jenkinson, R., Lyden, A., Yu, L. and Friedrich, D. (2025), 'Power system benefits of simultaneous domestic transport and heating demand flexibility in great britain's energy transition', *Applied Energy* **377**, 124522.

Garcia-Osipenko, M., Kuminoff, N. V., Perry, S. and Vreugdenhil, N. (2025), Optimal second-best menu design: Evidence from residential electricity plans, Technical report, National Bureau of Economic Research.

Garg, T., Hanna, R., Myers, J., Tebbe, S. and Victor, D. G. (2024), Electric vehicle charging at the workplace: Experimental evidence on incentives and environmental nudges, Technical report, CESifo Working Paper.

Ghosh, A. and Aggarwal, V. (2017), 'Control of charging of electric vehicles through menu-based pricing', *IEEE Transactions on Smart Grid* **9**(6), 5918–5929.

Gillingham, K. T., Ovaere, M. and Weber, S. M. (2025), 'Carbon policy and the emissions implications of electric vehicles', *Journal of the Association of Environmental and Resource Economists* **12**(2), 313–352.

Gosnell, G. K., List, J. A. and Metcalfe, R. D. (2020), 'The impact of management practices on employee productivity: A field experiment with airline captains', *Journal of Political Economy* **128**(4), 1195–1233.

Gravert, C. (2024), From intent to inertia: Experimental evidence from the retail electricity market, Technical report, CEBI Working Paper Series.

Green, M. (2025), 'Ai-powered vpp kraken reaches 2gw managed domestic assets', Current News. Online; accessed 17 July 2025.

Grigolon, L., Reynaert, M. and Verboven, F. (2018), 'Consumer valuation of fuel costs and tax policy: Evidence from the european car market', *American Economic Journal: Economic Policy* **10**(3), 193–225.

Hahn, R., Hendren, N., Metcalfe, R. and Sprung-Keyser, B. (2024), A Welfare Analysis of Policies Impacting Climate Change, Technical Report w32728, National Bureau of Economic Research, Cambridge, MA.

Hahn, R. W., Metcalfe, R. D. and Tam, E. (2023), Welfare estimates of shifting peak travel, Technical report, National Bureau of Economic Research.

Harberger, A. C. (1964), 'The measurement of waste', *The American Economic Review* **54**(3), 58–76.

Heid, P., Remmy, K. and Reynaert, M. (2024), 'Equilibrium effects in complementary markets: Electric vehicle adoption and electricity pricing'.

Held, M., Ilgmann, C., Kuckshinrichs, W., Reuter, M. and Schneider, C. (2021), 'A new car fleet: From fossil-fuelled to fossil-free within 15 years?', *European Transport Research Review* **13**(4), 1–15.

Hendren, N. and Sprung-Keyser, B. (2020), 'A unified welfare analysis of government policies', *The Quarterly Journal of Economics* **135**(3), 1209–1318.

Hendren, N. and Sprung-Keyser, B. (2022), The Case for Using the MVPF in Empirical Welfare Analysis, Technical Report w30029, National Bureau of Economic Research, Cambridge, MA.

Heuberger, C. F., Bains, P. K. and Mac Dowell, N. (2020), 'The EV-olution of the power system: A spatio-temporal optimisation model to investigate the impact of electric vehicle deployment', *Applied Energy* **257**, 113715.
URL: *https://linkinghub.elsevier.com/retrieve/pii/S0306261919314023*

Hildermeier, J., Burger, J., Jahn, A. and Rosenow, J. (2022), 'A review of tariffs and services for smart charging of electric vehicles in europe', *Energies* **16**(1), 88.

Hinchberger, A. J., Jacobsen, M. R., Knittel, C. R., Sallee, J. M. and van Benthem, A. A. (2024), The efficiency of dynamic electricity prices, Technical report, National Bureau of Economic Research.

HM Treasury (2020), 'The Green Book: Appraisal and Evaluation in Central Government'. Published by HM Treasury.

URL: *https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-government/the-green-book-2020a6-discounting*

Holland, S. P. and Mansur, E. T. (2006), 'The short-run effects of time-varying prices in competitive electricity markets', *The Energy Journal* **27**(4), 127–156.

Hortaçsu, A., Madanizadeh, S. A. and Puller, S. L. (2017), 'Power to choose? an analysis of consumer inertia in the residential electricity market', *American Economic Journal: Economic Policy* **9**(4), 192–226.

Huang, Q., Jia, Q.-S., Wu, X. and Guan, X. (2023), 'Two-phase on-line joint scheduling for welfare maximization of charging station', *IEEE Transactions on Automation Science and Engineering* .

Imbens, G. and Xu, Y. (2024), 'LaLonde (1986) after Nearly Four Decades: Lessons Learned'. Issue: arXiv:2406.00827 _eprint: 2406.00827.
URL: *http://arxiv.org/abs/2406.00827*

Imelda, Fripp, M. and Roberts, M. J. (2024), 'Real-time pricing and the cost of clean power', *American Economic Journal: Economic Policy* **16**(4), 100–141.

International Energy Agency (2025*a*), 'Energy and AI'. Available at: https://www.iea.org/reports/global-ev-outlook-2025.
URL: *https://www.iea.org/reports/energy-and-ai*

International Energy Agency (2025*b*), 'Global ev outlook 2025'. Available at: https://www.iea.org/reports/global-ev-outlook-2025.

Ito, K., Ida, T. and Tanaka, M. (2023), 'Selection on welfare gains: Experimental evidence from electricity plan choice', *American Economic Review* **113**(11), 2937–2973.

Jones, C. B., Vining, W., Lave, M., Haines, T., Neuman, C., Bennett, J. and Scoffield, D. R. (2022), 'Impact of electric vehicle customer response to time-of-use rates on distribution power grids', *Energy Reports* **8**, 8225–8235.

Joskow, P. L. and Tirole, J. (2006*a*), 'Merchant transmission investment', *The Journal of Industrial Economics* **54**(2), 233–264.

Joskow, P. L. and Wolfram, C. D. (2012), 'Dynamic pricing of electricity', *American Economic Review* **102**(3), 381–385.

Joskow, P. and Tirole, J. (2006*b*), 'Retail electricity competition', *The RAND Journal of Economics* **37**(4), 799–815.

Khanna, S., Martin, R. and Muûls, M. (2024), 'Shifting household power demand across time: incentives and automation', *Work. Pap., Imp. Coll. London* .

Knittel, C. R., Senga, J. R. L. and Wang, S. (2025), Flexible data centers and the grid: Lower costs, higher emissions?, Technical report, National Bureau of Economic Research.

Kolbert, E. (2024), 'The obscene energy demands of ai', *The New Yorker* **9**.

Kolli, K. (2011), 'Car longevity: a biometric approach', *Population* **66**(2), 307–323.

Kubli, M. (2022), 'EV Drivers' Willingness to Accept Smart Charging: Measuring Preferences of Potential Adopters', *Transportation Research Part D: Transport and Environment* **109**, 103396.
URL: *https://linkinghub.elsevier.com/retrieve/pii/S1361920922002231*

La Nauze, A., Friesen, L., Lim, K. L., Menezes, F., Page, L., Philip, T. and Whitehead, J. (2024), Can electric vehicles aid the renewable transition? evidence from a field experiment incentivising midday charging, Technical report, CESifo Working Paper.

Lagomarsino, M., Van Der Kam, M., Parra, D. and Hahnel, U. J. (2022), 'Do I Need To Charge Right Now? Tailored Choice Architecture Design Can Increase Preferences for Electric Vehicle Smart Charging', *Energy Policy* **162**, 112818.
URL: *https://linkinghub.elsevier.com/retrieve/pii/S030142152200043X*

LaLonde, R. J. (1986), 'Evaluating the Econometric Evaluations of Training Programs with Experimental Data', *American Economic Review* **76**(4), 604–620. JSTOR: 1806062.

Leslie, G., Pourkhanali, A. and Roger, G. (2021), 'Can real-time pricing be progressive? identifying cross-subsidies under fixed-rate electricity tariffs', *Identifying Cross-Subsidies under Fixed-Rate Electricity Tariffs (January 27, 2021)* .

Li, Y. and Jenn, A. (2024), 'Impact of electric vehicle charging demand on power distribution grid congestion', *Proceedings of the National Academy of Sciences* **121**(18), e2317599121.

Lowell, D., Jones, B. and Seamonds, D. (2017), 'Mjb&a analyzes state-wide costs and benefits of plug-in vehicles in five northeast and mid-atlantic states', *MJ Bradley & Associates* .

Mogstad, M., Torgovitsky, A. and Walters, C. R. (2021), 'The Causal Interpretation of Two-Stage Least Squares with Multiple Instrumental Variables', *American Economic Review* **111**(11), 3663–3698.

Moore, R. T. and Schnakenberg, K. (2023), 'blockTools: Blocking, Assignment, and Diagnosing Interference in Randomized Experiments v.0.6.4'.
URL: *https://cran.r-project.org/web/packages/blockTools/index.html*

Moser, C. (2017), 'The role of perceived control over appliances in the acceptance of electricity load-shifting programmes', *Energy Efficiency* **10**(5), 1115–1127.

National Energy System Operator (2025), 'Electricity capacity report data workbook, tab ec.11'. Accessed: 2025-07-21.
URL: *https://www.neso.energy/document/283061/download*

NYSERDA (2021), Benefit-cost analysis of electric vehicle deployment in new york state, Technical report, NYSERDA.

Office for Budget Responsibility (2024), 'The UK's Tax Burden in Historical and International Context'. Accessed: 25 July 2024.
URL: *https://obr.uk/box/the-uks-tax-burden-in-historical-and-international-context*

Parsons, A. and mySociety (2021), 'Composite 2020 UK Index of Multiple Deprivation (England-Anchored) v.3.3.0'. Place: https://github.com/mysociety/composite_uk_imd.
URL: *https://pages.mysociety.org/composite$_u k_i$md/datasets/uk$_i$ndex/latest*

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

Pilz, K. F., Mahmood, Y. and Heim, L. (2025), *AI's Power Requirements Under Exponential Growth Extrapolating AI Data Center Power Demand and Assessing Its Potential Impact on US Competitiveness*, RAND.

Poletti, S. and Wright, J. (2020), 'Real-time pricing and imperfect competition in electricity markets', *The Journal of Industrial Economics* **68**(1), 93–135.

Powell, S., Cezar, G. V., Min, L., Azevedo, I. M. and Rajagopal, R. (2022), 'Charging infrastructure access and operation to reduce the grid impacts of deep electric vehicle adoption', *Nature Energy* **7**(10), 932–945.

PwC (2024), 'Global Economy Watch: Projections'. Accessed: 25 July 2024.
URL: *https://www.pwc.com/gx/en/research-insights/global-economy-watch-projections.html*

Rapson, D. and Muehlegger, E. (2023*a*), 'Global transportation decarbonization', *Journal of Economic Perspectives* **37**(3), 163–188.

Rapson, D. S. and Muehlegger, E. (2023*b*), 'The economics of electric vehicles', *Review of Environmental Economics and Policy* **17**(2), 274–294.

Roth, J. (2024), 'Interpreting Event-Studies from Recent Difference-in-Differences Methods'. Place: arXiv.
URL: *https://doi.org/10.48550/arXiv.2401.12309*

Sandalow, D. B., McCormick, C., Kucukelbir, A., Friedmann, J., Nachmany, M., Lee, H., Hill, A., Loehr, D., Wald, M., Halff, A. M. et al. (2024), 'Artificial intelligence for climate change mitigation roadmap'.

Seamonds, D. and Lowell, D. (2018), 'Electric vehicle cost-benefit analysis'.

SEPA (2024), The state of managed charging in 2024, Technical report, Smart Electic Power Alliance.

Sovacool, B. K., Axsen, J. and Kempton, W. (2017), 'The Future Promise of Vehicle-to-Grid (V2G) Integration: A Sociotechnical Review and Research Agenda', *Annual Review of Environment and Resources* **42**(1), 377–406.
URL: *https://www.annualreviews.org/doi/10.1146/annurev-environ-030117-020220*

The White House (2025), 'Initial rescissions of harmful executive orders and actions', https://www.whitehouse.gov/presidential-actions/2025/01/initial-rescissions-of-harmful-executive-orders-and-actions/.

Thornhill, C. and Deasley, S. (2018), 'My Electric Benefits Case: Peer Review by Frontier Economics'.

Turk, G., Schittekatte, T., Martínez, P. D., Joskow, P. L. and Schmalensee, R. (2024), *Designing Distribution Network Tariffs Under Increased Residential End-user Electrification: Can the US Learn Something from Europe?*, MIT Center for Energy and Environmental Policy Research.

WattTime (2022), Marginal Emissions Modeling: WattTime's Approach to Modeling and Validation, Technical report.
URL: *https://watttime.org/wp-content/uploads/2023/11/WattTime-MOER-modeling-20221004.pdf*

Wolak, F. A. (2013), 'Economic and political constraints on the demand-side of electricity industry re-structuring processes', *Review of Economics and Institutions* **4**(1), 42.

Yi, H., Lin, Q. and Chen, M. (2019), Balancing cost and dissatisfaction in online ev charging under real-time pricing, *in* 'IEEE INFOCOM 2019-IEEE Conference on Computer Communications', IEEE, pp. 1801–1809.

Zhou, Y. and Li, S. (2018), 'Technology adoption and critical mass: the case of the us electric vehicle market', *The Journal of Industrial Economics* **66**(2), 423–480.

A Appendix

A.1 Tables

Table A1: Experimental Balance

Variable	Control Mean	Email + £0/Mth	Email + £5/Mth	Email + £50/Mth	Email + £50/Mth (No Bump)
Not Always Octopus Customer	0.26	0.00	0.00	0.00	0.00
		(0.01)	(0.02)	(0.02)	(0.02)
Octopus Tenure	2.72	-0.03	-0.01	-0.03	-0.01
	[1.63]	(0.04)	(0.06)	(0.06)	(0.06)
Proportion Total kWh Peak	0.23	0.00	0.00	0.00	0.00
	[0.08]	(0.00)	(0.00)	(0.00)	(0.00)
Smart Tariff Onboarding Processes	0.22	0.00	0.00	-0.01	0.00
	[0.47]	(0.01)	(0.02)	(0.02)	(0.02)
Structural Winnings (GBP/kWh)	752.64	-9.72	-17.76	-26.70	-18.12
	[684.94]	(16.50)	(25.08)	(23.84)	(25.02)
Total kWh	4,082.53	-42.42	-33.33	-65.43	-38.96
	[2,599.11]	(62.80)	(97.97)	(91.35)	(94.82)
Total kWh Stdev	0.77	-0.01	-0.01	0.00	0.00
	[0.25]	(0.01)	(0.01)	(0.01)	(0.01)
N	2,205	7,720	1,101	1,102	1,105

Note: The first column shows the mean and [standard deviation] for the control group. Each row and each subsequent encouragement column represents an individual regression of the row variable on an indicator for receiving the encouragement in the column. The encouragements appear to be balanced on baseline characteristics. The standard errors are in parentheses. Density plots of these covariates are shown in Figure A2. Detailed definitions of these covariates can be found in Appendix A.4.2.

Table A2: Impact of Encouragements on Take-up of IO Go

Dependent Variable:	IO GO
Model:	(1)
Variables	
Email + £0/Month	0.034***
	(0.004)
Email + £5/Month	0.033***
	(0.007)
Email + £50/Month	0.059***
	(0.008)
Email + £50/Month (No Bump)	0.057***
	(0.008)
Post-Incentive Period	0.004
	(0.004)
Email + £0/Month × Post-Incentive Period	-0.003
	(0.004)
Email + £5/Month × Post-Incentive Period	-0.006
	(0.007)
Email + £50/Month × Post-Incentive Period	-0.003
	(0.006)
Email + £50/Month (No Bump) × Post-Incentive Period	-0.016**
	(0.006)
Fixed-effects	
Block	Yes
Week of Year	Yes
Fit statistics	
Control Mean (Incentive Period)	0.024
Test £0 = £50	0.002
Test £0 = £50 (No Bump)	0.003
Observations	661,891

Clustered (Participant & Week of Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table shows the effects of four email-based encouragements on adopting IO Go in the 12 months since emails were sent out. The outcome is a binary indicator for weekly use of the Octopus Go tariff. Encouragement indicators are interacted an indicator for whether the week is during the incentive period, three months after the start of the trial. The specification controls for fixed effects for randomization block and week-of-year. Standard errors, clustered by participant and week, are reported in the parentheses. Mean IO Go take-up rate during the incentive period is reported for the control group. "Test £0 = £50" is the p-value on the test of equality between the first and third coefficient; "Test £0 = £50 (No bump)" is the p-value on the test of equality between the first and fourth coefficient.

Table A3: Robustness, Impact of Encouragements on Take-up of IO Go, Daily Data

Dependent Variable:	IO GO
Model:	(1)
Variables	
Email + £0/Month	0.034***
	(0.004)
Email + £5/Month	0.034***
	(0.007)
Email + £50/Month	0.059***
	(0.008)
Email + £50/Month (No Bump)	0.057***
	(0.008)
Post-Incentive Period	0.004
	(0.005)
Email + £0/Month × Post-Incentive Period	-0.003
	(0.005)
Email + £5/Month × Post-Incentive Period	-0.006
	(0.007)
Email + £50/Month × Post-Incentive Period	-0.004
	(0.007)
Email + £50/Month (No Bump) × Post-Incentive Period	-0.017**
	(0.007)
Fixed-effects	
Block	Yes
Day	Yes
Fit statistics	
Control Mean (Incentive Period)	0.023
Test £0 = £50	0.002
Test £0 = £50 (No Bump)	0.003
Observations	4,580,025

Clustered (Participant & Day) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table shows the effects of four email-based encouragements on adopting IO Go in the 12 months since emails were sent out. The outcome is a binary indicator for daily use of the Octopus Go tariff. Encouragement indicators are interacted an indicator for whether the date is during the incentive period, which is the three months after the start of the trial. The specification controls for fixed effects for randomization block and week-of-year. Standard errors, clustered by participant and day, are reported in the parentheses. Mean IO Go take-up rate during the incentive period is reported for the control group. "Test £0 = £50" is the p-value on the test of equality between the first and third coefficient; "Test £0 = £50 (No bump)" is the p-value on the test of equality between the first and fourth coefficient.

Table A4: Impact of Encouragements on Electricity Consumption

Model:	Peak (1)	Off-Peak (2)	Overall (3)	Peak (4)	Off-Peak (5)	Overall (6)
Variables						
Email + £0/Month	-0.029***	0.018	-0.005			
	(0.011)	(0.012)	(0.007)			
Email + £5/Month	-0.011	0.018	0.004			
	(0.016)	(0.017)	(0.010)			
Email + £50/Month	-0.013	0.025	0.003			
	(0.016)	(0.017)	(0.010)			
Email + £50/Month (No Bump)	-0.033**	0.026	-0.006			
	(0.016)	(0.018)	(0.010)			
Any Encouragement				-0.026**	0.019*	-0.004
				(0.010)	(0.012)	(0.007)
Fixed-effects						
Week of Year	Yes	Yes	Yes	Yes	Yes	Yes
Block	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Control Mean	1.4	1	1	1.4	1	1
Observations	2,646,712	4,631,639	15,879,842	2,646,712	4,631,639	15,879,842

Clustered (Participant & Week of Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Each column reports coefficients from separate regressions of hourly electricity consumption (in kWh) on indicators for encouragement assignment. Columns (1), (2), and (3) estimate intention-to-treat effects for each arm; columns (4), (5), and (6) pool all treatment arms into a single binary indicator. The dependent variable is consumption during either the peak (16:30–20:30), off-peak (23:30–05:30), or overall hours. All regressions control for baseline consumption, and include fixed effects for week-of-year and randomization block. Standard errors, clustered by participant and week, are reported in parentheses. Means consumption for the control group are reported at the bottom of each panel.

Table A5: Impact of Encouragements on Octopus Go Take-up

Dependent Variable:	Octopus Go
Model:	(1)
Variables	
Email + £0/Month	0.0101***
	(0.0031)
Email + £5/Month	0.0099*
	(0.0052)
Email + £50/Month	0.0045
	(0.0049)
Email + £50/Month (No Bump)	0.0066
	(0.0051)
Fixed-effects	
Block	Yes
Week of Year	Yes
Fit statistics	
Control Mean	0.025
Observations	661,891

Clustered (Participant & Week of Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table shows the effects of four email-based encouragements on adopting Octopus Go in the 12 months since emails were sent out. The outcome is a binary indicator of for weekly use of the Octopus Go tariff. The specification controls for fixed effects for randomization block and week. Standard errors, clustered by participant and week, are reported in parentheses. Mean Octopus Go take-up rate is reported for the control group.

Table A6: Selection into IO Go and Baseline Characteristics

Variable	estimate	std.error	p.value
Email + £0/Month	0.034	0.005	0.000
Email + £5/Month	0.037	0.009	0.000
Email + £50/Month	0.057	0.010	0.000
Email + £50/Month (No Bump)	0.059	0.010	0.000
IMD Tercile 2	0.030	0.007	0.000
IMD Tercile 2	0.033	0.006	0.000
Octopue Tenure (Years)	-0.003	0.002	0.032
Structural Winnings (Z-Score)	-0.007	0.003	0.008

Note: This table shows the estimation results of a logit regression of takeup of IO Go on encouragement group and participants' baseline characteristics. We show the marginal effects at the means of the covariates. The sample size is 13,227.

Table A7: Robustness, Impact of EV Tariff on Peak Consumption (kWh)

	Main	4 Instruments	£0/Mth	£5/Mth	£50/Mth	£50/Mth (No Bump)	6 Months	No Baseline
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
EV Tariff	-0.581**	-0.498**	-0.697***	-0.269	-0.305	-0.587**	-0.473**	-0.554
	(0.224)	(0.209)	(0.247)	(0.338)	(0.230)	(0.236)	(0.214)	(0.336)
Fixed-effects								
Week of Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Block	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics								
Control Mean	1.4	1.4	1.4	1.4	1.4	1.4	1.4	1.4
Observations	2,646,712	2,646,712	1,980,876	663,401	660,592	662,359	1,307,900	2,646,712
First Stage F-Stat	52.720	14.066	44.154	20.302	37.764	40.647	78.180	52.715

Clustered (Participant & Week of Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table reports IV estimates of the effect of adopting an EV tariff on hourly electricity consumption (kWh) during peak hours (16:30-20:30). The instrument is an indicator for assignment to any email-based encouragement. Columns (1) is our main specification, also reported in Figure 6; (2) defines a separate instrument for each encouragement group; (3) restricts to just the £0/Month group and the control group; (4) restricts to just the £5/Month group and the control group; (5) restricts to just the £50/Month group and the control group; (6) just the £50/Month (No Bump) group and the control group; (7) restricts to the 6 months after encouragement emails were sent out; (8) does not control for baseline consumption. All specifications control for baseline consumption (except column 6), and fixed effects for randomization block and week. Standard errors, clustered by participant and week, are reported in parentheses. The control mean is calculated over the same time periods for control group trial participants not enrolled in an EV tariff. The bottom row shows the first-stage Wald F-statistic.

Table A8: Robustness, Impact of EV Tariff on Off-Peak Consumption (kWh)

	Main	4 Instruments	£0/Mth	£5/Mth	£50/Mth	£50/Mth (No Bump)	6 Months	No Baseline
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
EV Tariff	0.481*	0.468*	0.501*	0.284	0.280	0.414	0.686***	0.118
	(0.261)	(0.244)	(0.282)	(0.400)	(0.269)	(0.270)	(0.236)	(0.392)
Fixed-effects								
Week of Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Block	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics								
Control Mean	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99
Observations	3,970,004	3,970,004	2,971,234	995,188	990,964	993,611	1,961,825	3,970,004
First Stage F-Stat	53.449	14.283	44.658	20.564	38.473	41.213	78.912	52.756

Clustered (Participant & Week of Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table reports IV estimates of the effect of adopting an EV tariff on hourly electricity consumption (kWh) during off-peak hours (23:30-05:30). The instrument is an indicator for assignment to any email-based encouragement. Columns (1) is our main specification, also reported in Figure 6; (2) defines a separate indicator for each encouragement group; (3) restricts to just the £0/Month group and the control group; (4) restricts to just the £50/Month group and the control group; (5) just the £50/Month (No Bump) group and the control group; (6) restricts to the 6 months after encouragement emails were sent out; (7) we do not control for baseline consumption. All specifications control for baseline consumption (except column 6), and fixed effects for randomization block and week. Standard errors, clustered by participant and week, are reported in parentheses. The control mean is calculated over the same time periods for control group trial participants not enrolled in an EV tariff. The bottom row shows the first-stage Wald F-statistic.

Table A9: Robustness, Impact of EV Tariff on Overall Consumption (kWh)

	Main	4 Instruments	£0/Mth	£5/Mth	£50/Mth	£50/Mth (No Bump)	6 Months	No Baseline
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
EV Tariff	-0.079	-0.060	-0.125	0.032	-0.038	-0.132	0.002	-0.203
	(0.150)	(0.140)	(0.163)	(0.226)	(0.150)	(0.159)	(0.154)	(0.247)
Fixed-effects								
Week of Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Block	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics								
Control Mean	1	1	1	1	1	1	1	1
Observations	15,879,842	15,879,842	11,884,810	3,980,493	3,963,619	3,974,217	7,847,164	15,879,842
First Stage F-Stat	52.809	14.090	44.257	20.465	37.912	40.763	78.281	52.747

Clustered (Participant & Week of Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table reports IV estimates of the effect of adopting an EV tariff on electricity consumption over the whole day. The instrument is an indicator for assignment to any email-based encouragement. Columns (1) is our main specification; (2) defines a separate indicator for each encouragement group; (3) restricts to just the £0/Month group and the control group; (4) restricts to just the £50/Month group and the control group; (5) just the £50/Month (No Bump) group and the control group; (6) restricts to the 6 months after encouragement emails were sent out; (7) we do not control for baseline consumption. All specifications control for baseline consumption (except column 6), and fixed effects for randomization block and week. Standard errors, clustered by participant and week, are reported in parentheses. The control mean is calculated over the same time periods for control group trial participants not enrolled in an EV tariff. The bottom row shows the first-stage Wald F-statistic.

Table A10: Heterogeneity of IO Go Take-up, by IMD and Baseline Electricity Consumption

Model:	x IMD (1)	x Baseline kWh (2)
Variables		
Email + £0/Mth × Tercile 1	0.021	0.029***
	(0.018)	(0.010)
Email + £0/Mth × Tercile 2	0.045***	0.050***
	(0.009)	(0.009)
Email + £0/Mth × Tercile 3	0.026***	0.016*
	(0.007)	(0.009)
Email + £5/Mth × Tercile 1	0.001	0.035**
	(0.026)	(0.016)
Email + £5/Mth × Tercile 2	0.052***	0.033**
	(0.017)	(0.015)
Email + £5/Mth × Tercile 3	0.022*	0.019
	(0.011)	(0.015)
Email + £50/Mth × Tercile 1	0.009	0.069***
	(0.031)	(0.018)
Email + £50/Mth × Tercile 2	0.062***	0.064***
	(0.018)	(0.017)
Email + £50/Mth × Tercile 3	0.059***	0.036**
	(0.013)	(0.016)
Email + £50/Mth (No Bump) × Tercile 1	0.050	0.080***
	(0.033)	(0.019)
Email + £50/Mth (No Bump) × Tercile 2	0.062***	0.045***
	(0.017)	(0.016)
Email + £50/Mth (No Bump) × Tercile 3	0.035***	0.009
	(0.012)	(0.015)
Tercile 2	-0.009	-0.013
	(0.017)	(0.012)
Tercile 3	0.010	-0.0005
	(0.017)	(0.013)
Fixed-effects		
Block	Yes	Yes
Week of Year	Yes	Yes
Fit statistics		
Observations	15,879,842	15,879,842

Clustered (Participant & Week of Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table presents heterogeneity in impact of encouragements on take-up of IO Go, interacting encouragement group with (1) Index of Multiple Deprivation (IMD) terciles and (2) baseline consumption terciles. The outcome is a binary indicator for weekly use of the Octopus Go tariff. The specification controls for baseline consumption, and fixed effects of randomization block and week of year. Standard errors, clustered by participant week, are reported in parentheses.

Table A11: Heterogeneity of Impact of EV Tariffs, by IMD and Baseline Electricity Consumption

	x IMD		x Baseline kWh	
	Off-Peak	Peak	Off-Peak	Peak
Model:	(1)	(2)	(3)	(4)
Variables				
(EV Tariff) x (Tercile 1)	0.924	1.16	0.295	-0.300
	(0.960)	(1.08)	(0.291)	(0.221)
(EV Tariff) x (Tercile 2)	0.604	-0.882**	0.504	-0.399
	(0.388)	(0.383)	(0.316)	(0.277)
(EV Tariff) x (Tercile 3)	0.508*	-0.207	0.998	-0.759
	(0.297)	(0.269)	(0.778)	(0.725)
(Other Tariff) x (Tercile 2)	0.078	0.327**	0.059	0.182***
	(0.127)	(0.140)	(0.072)	(0.060)
(Other Tariff) x (Tercile 3)	0.098	0.218	0.131	0.441***
	(0.127)	(0.138)	(0.134)	(0.125)
Fixed-effects				
Block	Yes	Yes	Yes	Yes
Week of Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	3,970,004	2,646,712	3,970,004	2,646,712

Clustered (Participant & Week of Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table presents heterogeneity in IV effects of EV tariff adoption on hourly electricity consumption during off-peak (23:30-5:30) and peak (16:30-20:30) periods, interacting EV tariff adoption with (1) Index of Multiple Deprivation (IMD) terciles and (2) baseline consumption terciles. All specifications control for baseline consumption, and fixed effects for randomization block and week. Standard errors, clustered by participant and week, are reported in parentheses.

Table A12: Bump Charging and Baseline Characteristics

Dependent Variable:	I(Ever Bumped)
Model:	(1)
Variables	
Constant	0.385***
	(0.050)
ZEmail+£0/Mth	0.004
	(0.032)
ZEmail+£5/Mth	-0.063
	(0.047)
ZEmail+£50/Mth	-0.002
	(0.045)
ZEmail+£50/Mth(NoBump)	-0.002
	(0.046)
Frac. of kWh During Peak (Z-Score)	0.039***
	(0.012)
Total Consumption (Z-Score)	-0.006
	(0.015)
Structural Winnings (Z-Score)	0.005
	(0.016)
IMD Tercile 2	0.029
	(0.043)
IMD Tercile 3	0.052
	(0.041)
Octopus Tenure (Years)	0.007
	(0.007)
I(Baseline TOU)	0.078**
	(0.033)
Fit statistics	
Observations	2,146
Dependent variable mean	0.45573

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table reports regression results where the dependent variable is an indicator for whether a participant has ever engaged in bump charging. Explanatory variables include the encouragement group assignment and baseline characteristics. Detailed definitions of these baseline characteristics can be found in Appendix A.4.2.

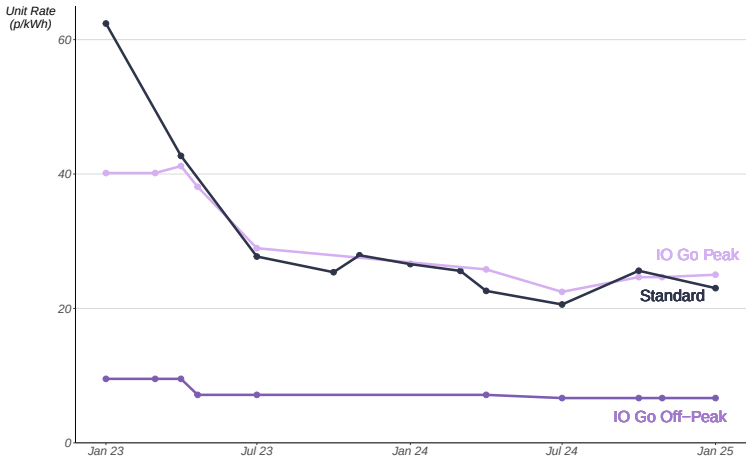
Table A13: Baseline Differences Between IO Go and Octopus Go Trial Participants

Variable	IO Go Mean	Octopus Go Mean	IO Go - Octopus Go	p-value
Octopus Tenure	2.51	2.58	-0.06	0.34
	[1.65]	[1.64]		
Total kWh	3912.32	3753.20	159.12	0.076*
	[2343.05]	[2174.77]		
Total kWh Stdev	0.77	0.77	0.00	0.74
	[0.24]	[0.22]		
Not Always Octopus Customer	0.24	0.24	0.00	0.95
	[0.43]	[0.43]		
Smart Tariff Onboarding Processes	0.34	0.26	0.08	0.00025***
	[0.58]	[0.51]		
Structural Winnings (GBP/kWh)	687.70	684.59	3.11	0.9
	[630.75]	[592.95]		
Proportion Total kWh Peak	0.22	0.23	-0.01	0.057*
	[0.08]	[0.08]		

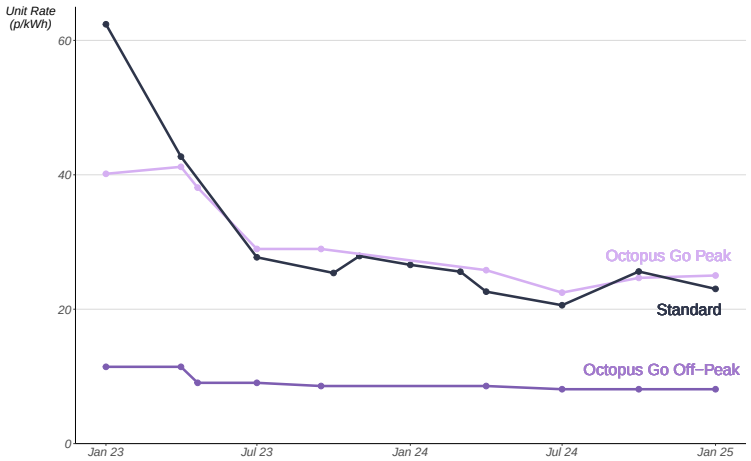
Note: Column (1) shows the mean and [standard deviation] for IO Go users; column (2) reports the same for Octopus Go participants; column (3) shows the difference in means between columns (1) and (2); column (4) reports the p-value from a two-sided t-test of equality of means. Detailed definitions of these covariates can be found in Appendix A.4.2.

A.2 Figures

Figure A1: Tariff Rates in 2024



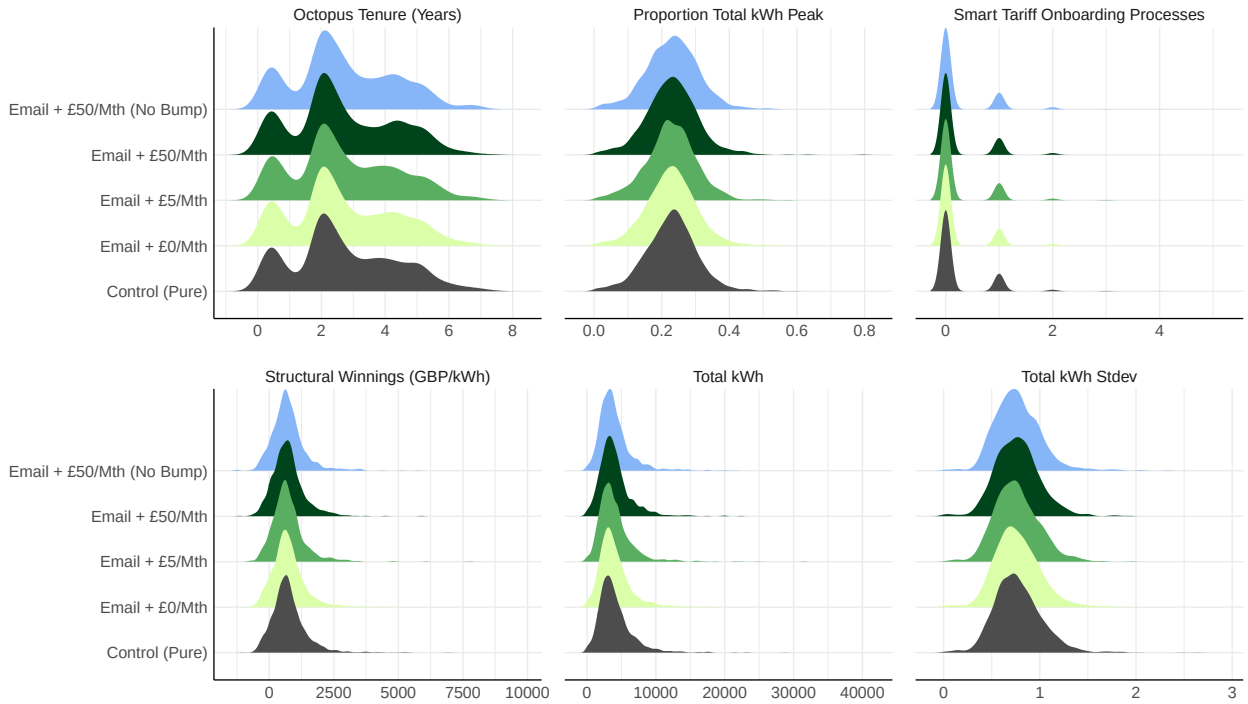
(a) Intelligent Octopus Go



(b) Octopus Go

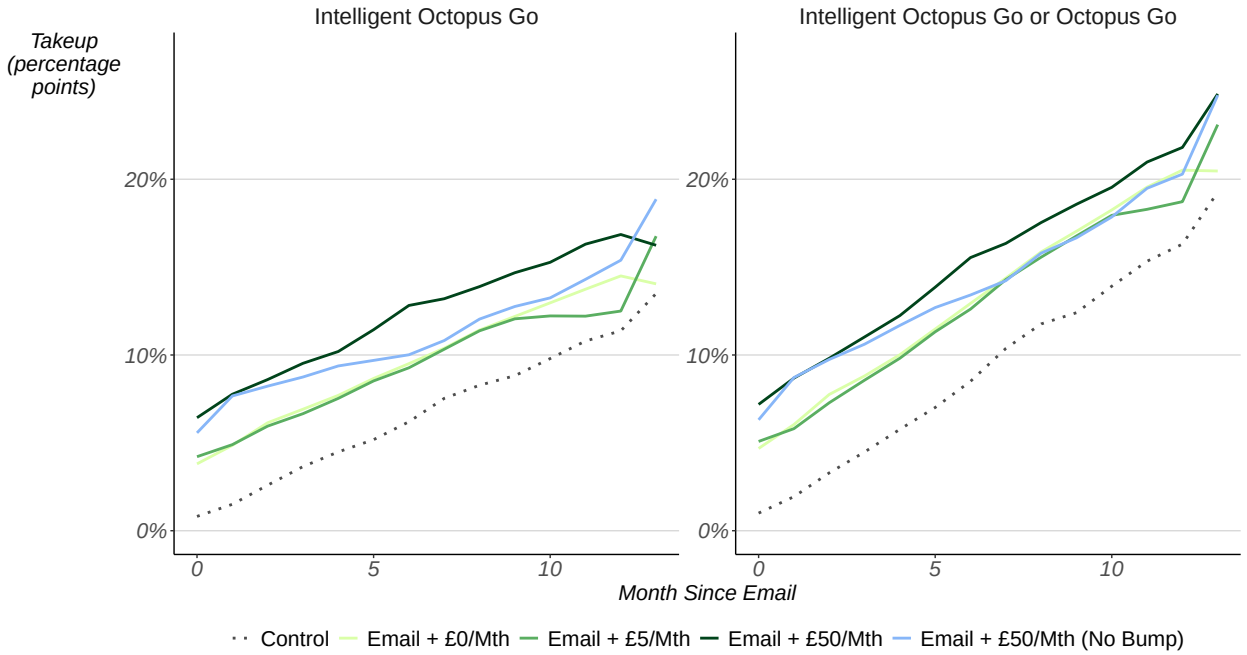
Notes: Panel (a) shows the tariff rates for Intelligent Octopus Go customers during the off-peak overnight period (23:30–05:30, dark purple) and the peak daytime period (05:30–23:30, light purple). For comparison, we also include the Flexible Octopus tariff from Octopus Energy, which maintains a flat rate throughout the day. Panel (b) shows analogous tariff rates for Octopus Go's off-peak overnight period (00:30-5:30) and peak daytime period (5:30-00:30).

Figure A2: Distribution of Baseline Covariates Across Encouragement Arms



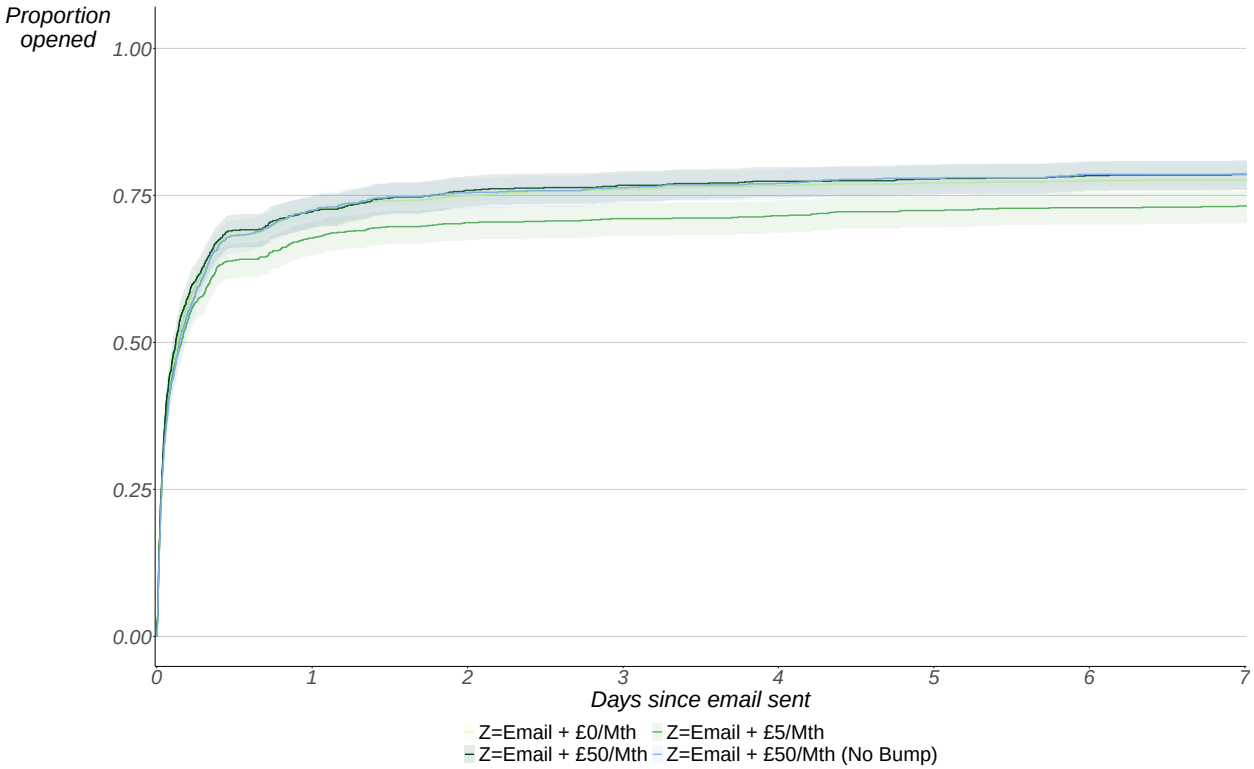
Notes: This figure plots kernel density estimates of baseline covariates across encouragement arms. The similarity of these distributions illustrates that block randomization achieved good covariate balance across arms, consistent with the regression-based balance tests reported in Table A1. Note that the variable 'Not Always Octopus Customer' has been excluded, as it is a binary variable that is unsuitable for a density plot. Detailed definitions of these covariates can be found in Appendix A.4.2.

Figure A3: Take-up of EV Tariffs Over Time by Trial Arms



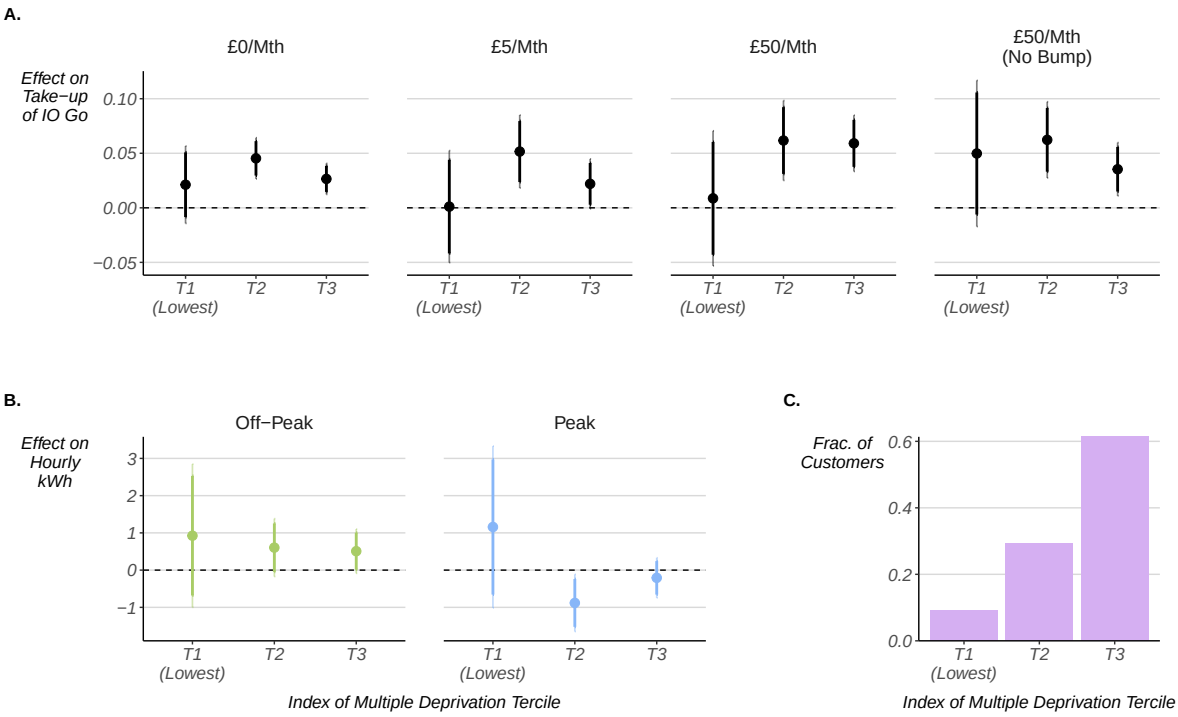
Notes: This figure plots the fraction of trial participants in each trial group that has taken up an EV tariff. The left panel shows take-up of the AI managed charging tariff, IO Go, which combines static time-of-use pricing with remote control of EV charging. The right panel shows take-up of either IO Go or Octopus Go, the latter being a static time-of-use tariff without supplier control.

Figure A4: Probability of First-Ever Opening Emailed Encouragement



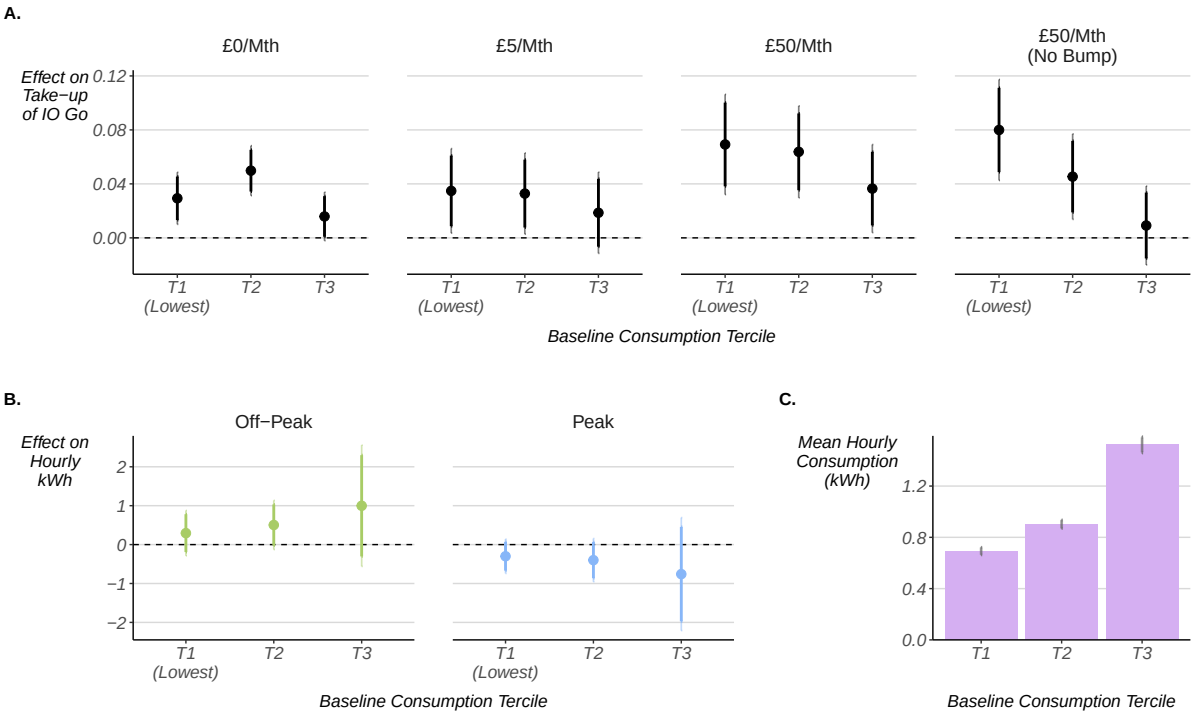
Notes: This figure plots the cumulative proportion of customers who first opened an encouragement email over time. Solid lines show Kaplan–Meier survival estimates of the probability of not yet opening, with shaded areas denoting 95% confidence intervals. Time is measured from the date of the encouragement email until the first observed open.

Figure A5: Heterogeneity by Index of Multiple Deprivation



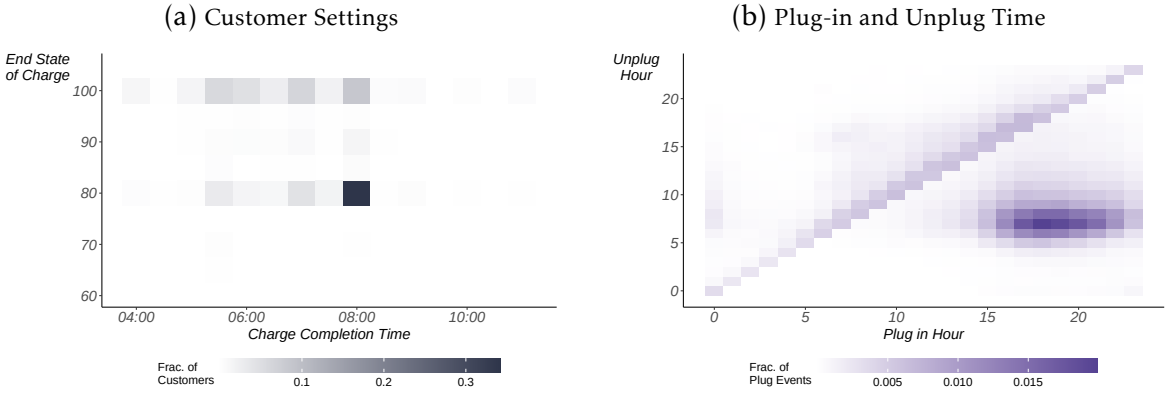
Notes: This figure presents heterogeneity in treatment effects by Index of Multiple Deprivation (IMD) terciles. Panel A shows the effect of each encouragement group on take-up of IO Go, interacting encouragement with deprivation tercile (T1: most deprived); regression results are also presented in Table A10. Panel B displays the estimated IV effects of EV tariff adoption on hourly electricity consumption during off-peak (23:30-5:30, green) and peak (16:30-20:30, blue) periods, again interacting EV tariff adoption with deprivation tercile; regression results are also presented in Table A11. Panel C shows the proportion of trial participants in the experimental sample in each tercile. Confidence intervals are shown at the 95% level. All specifications control for baseline consumption, and fixed effects for randomization block and week. Lines depict 90% (dark) and 95% (light) confidence intervals. Standard errors are clustered by participant and week.

Figure A6: Heterogeneity by Baseline Electricity Consumption



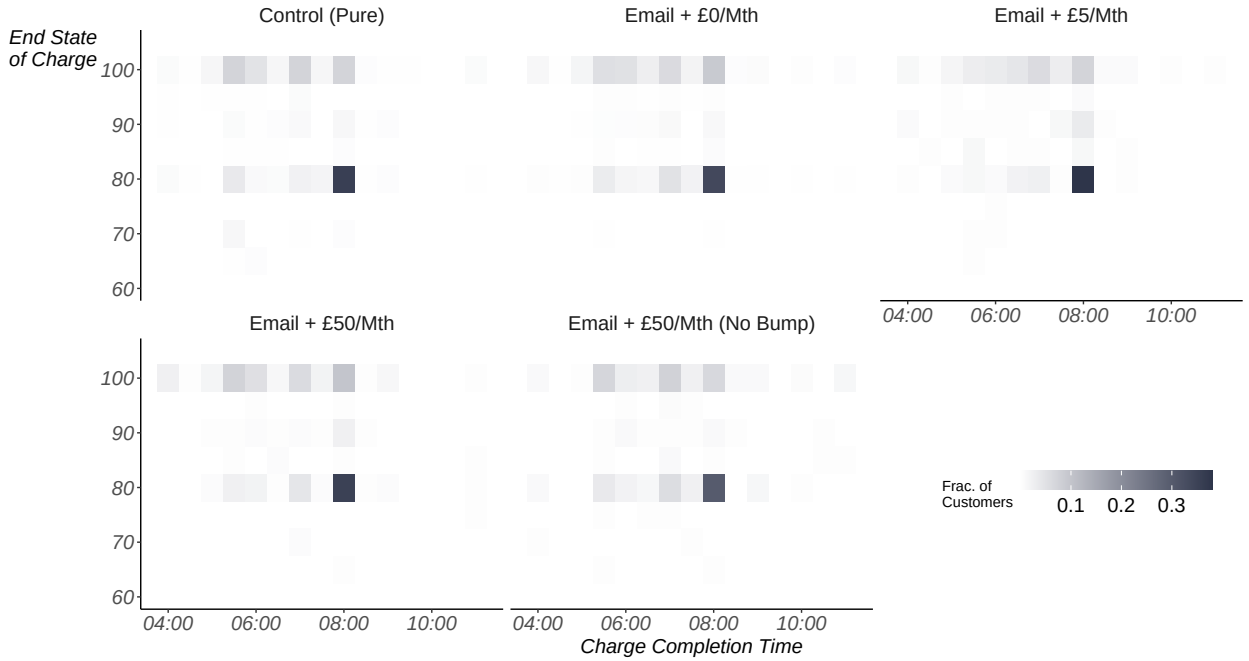
Notes: This figure presents heterogeneity in treatment effects by baseline total consumption (kWh). Panel A shows the effect of each encouragement group on take-up of IO Go, interacting encouragement with deprivation tercile (T1: lowest consumption); regression results are also presented in Table A10. Panel B displays the estimated IV effects of EV tariff adoption on hourly electricity consumption during off-peak (23:30-5:30, green) and peak (16:30–20:30, blue) periods, again interacting EV tariff adoption with baseline consumption tercile; regression results are also presented in Table A11. Panel C shows mean hourly consumption for each tercile, with 95% standard error bars. Confidence intervals are shown at the 95% level. All specifications control for baseline consumption, and fixed effects for randomization block and week. Lines depict 90% (dark) and 95% (light) confidence intervals. Standard errors are clustered by participant and week.

Figure A7: Participant Preferences and Behaviors



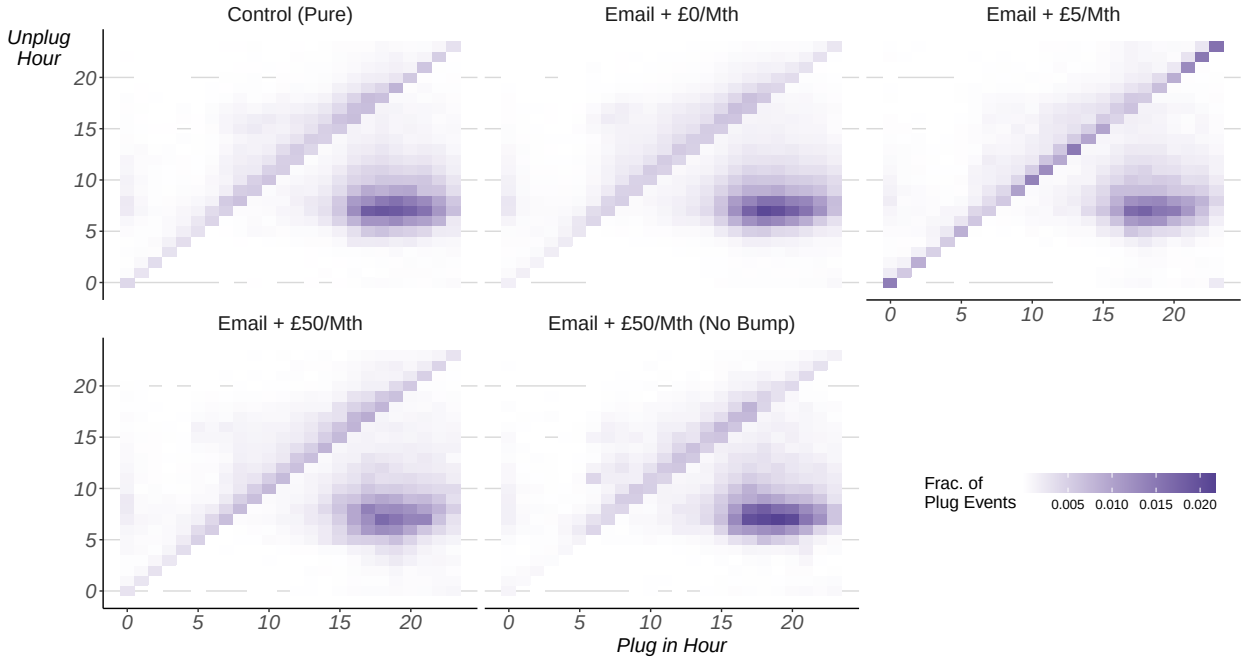
Notes: This figure summarizes IO Go customer preferences and behavioral patterns. Panel (a) shows the joint distribution of customer-selected end state of charge and charge completion time, as specified via the Octopus app. Panel (b) displays the frequency of plug-in and unplug times.

Figure A8: Preferences for End State of Charge and Completion Time - By Encouragement Group



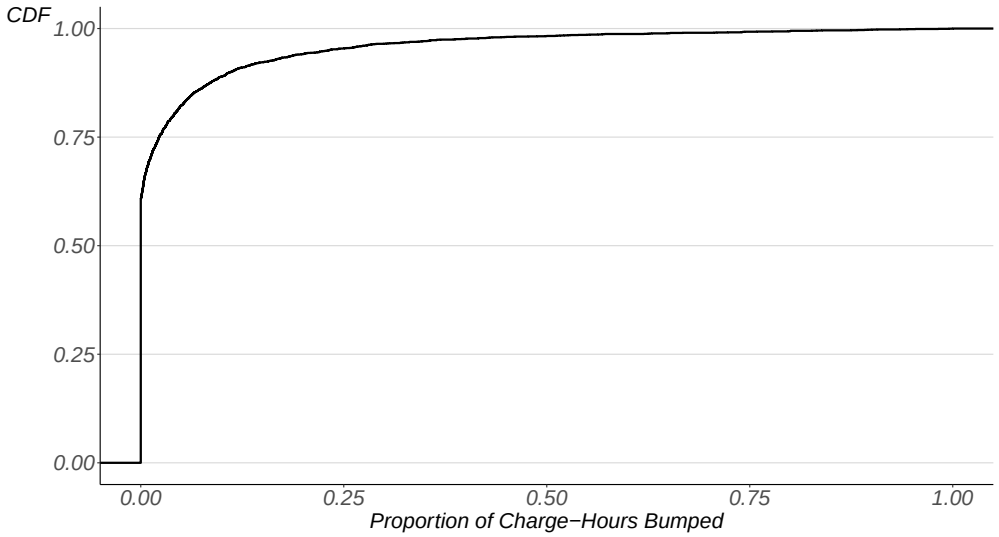
Notes: This figure shows the joint distribution of customer-selected end state of charge and charge completion time, as specified via the Octopus app. The trial participants are split by the encouragement group they were randomly assigned.

Figure A9: Preferences for End State of Charge and Completion Time - By Encouragement Group



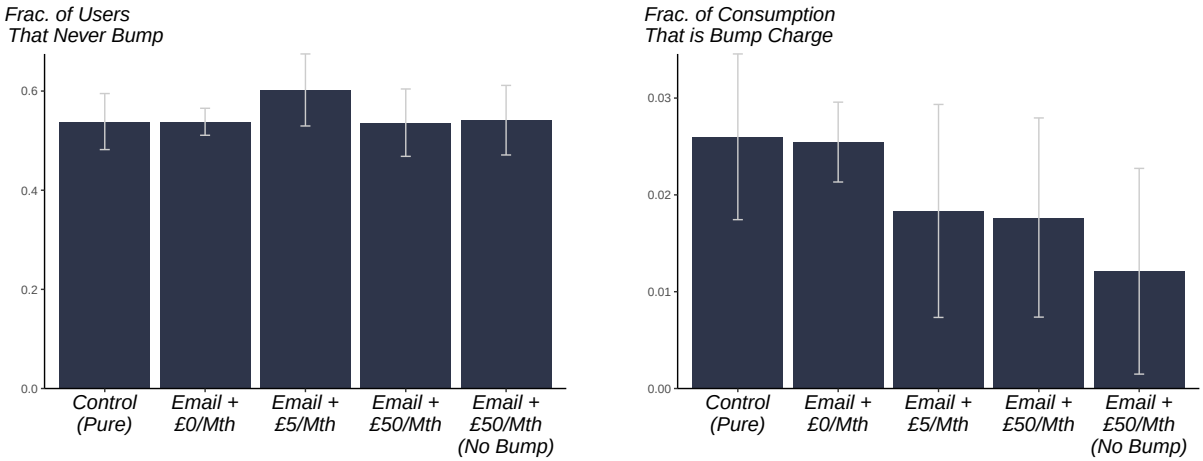
Notes: This figure shows the frequency of plug-in and unplug times. The results are shown separately by randomized encouragement group.

Figure A10: Distribution of Proportion of Charge-Hours Bumped Per Customer



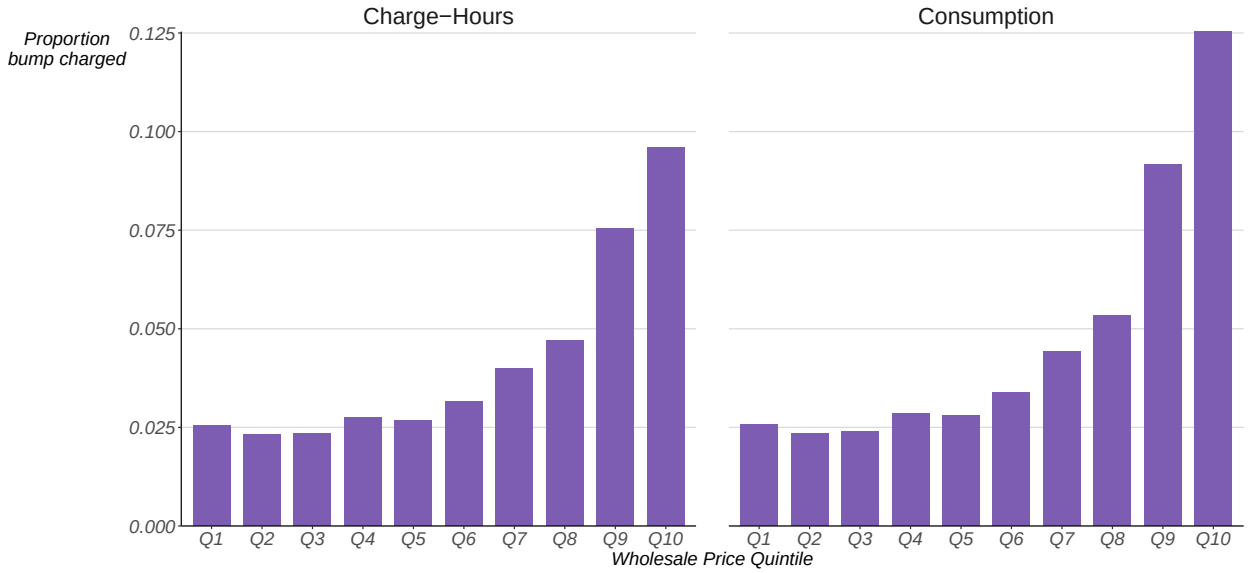
Notes: This figure illustrates bump charging behavior, where trial participants manually overrode the AI managed charging schedule. The horizontal axis shows, for each customer, what share of their total charge-hours was "bumped" (overridden). The vertical axis shows the cumulative proportion of customers. Charge-hours here are hours where any charging occurs.

Figure A11: Bump Charge Behaviors - By Encouragement Group



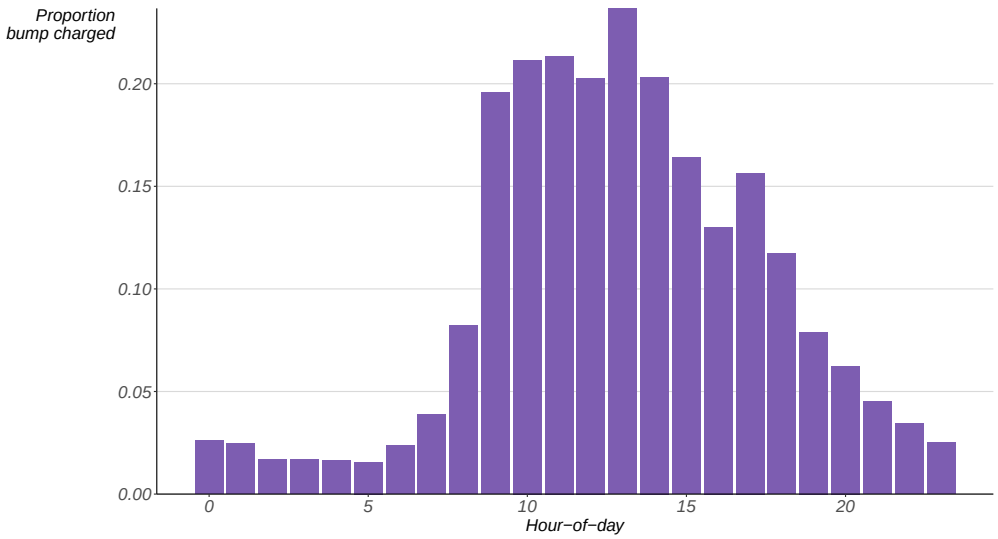
Notes: This figure illustrates bump charging behavior, where trial participants manually overrode the AI managed charging schedule. Results are shown separately by randomized encouragement group. The left panel shows the proportion of trial participants who never used bump charging. The right panel shows the share of total electricity consumption that came from bump charging.

Figure A12: Bump Behavior, by Quintile of System Price



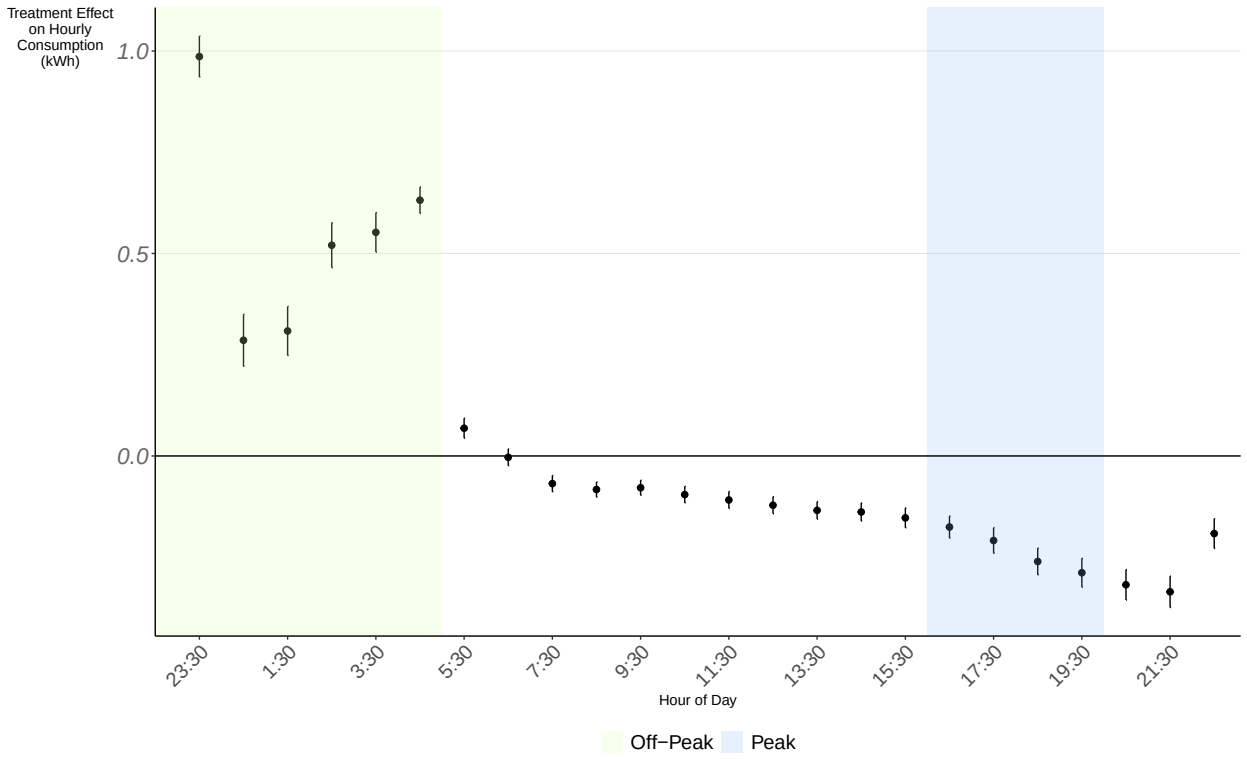
Notes: This figure shows bump charging behavior, where trial participants manually overrode the AI-managed charging schedule. The left panel reports the proportion of charge-hours that were bumped, within each quintile of wholesale system prices. A charge-hour is defined as any hour in which charging occurs. The right panel reports the proportion of total consumption that came from bump charging, again by system price quintile.

Figure A13: Bump Behavior, by Hour-of-day



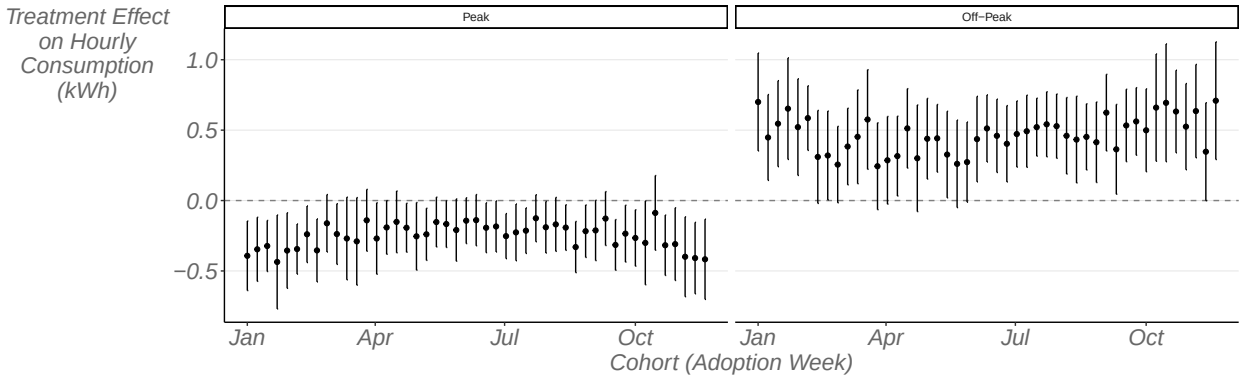
Notes: This figure shows bump charging behavior, where trial participants manually overrode the AI-managed charging schedule. It reports the share of charge-hours that were bumped at each hour of the day. A charging hour means any hour during which charging took place.

Figure A14: Difference-in-Differences Estimate of IO Go, by Hour-of-day



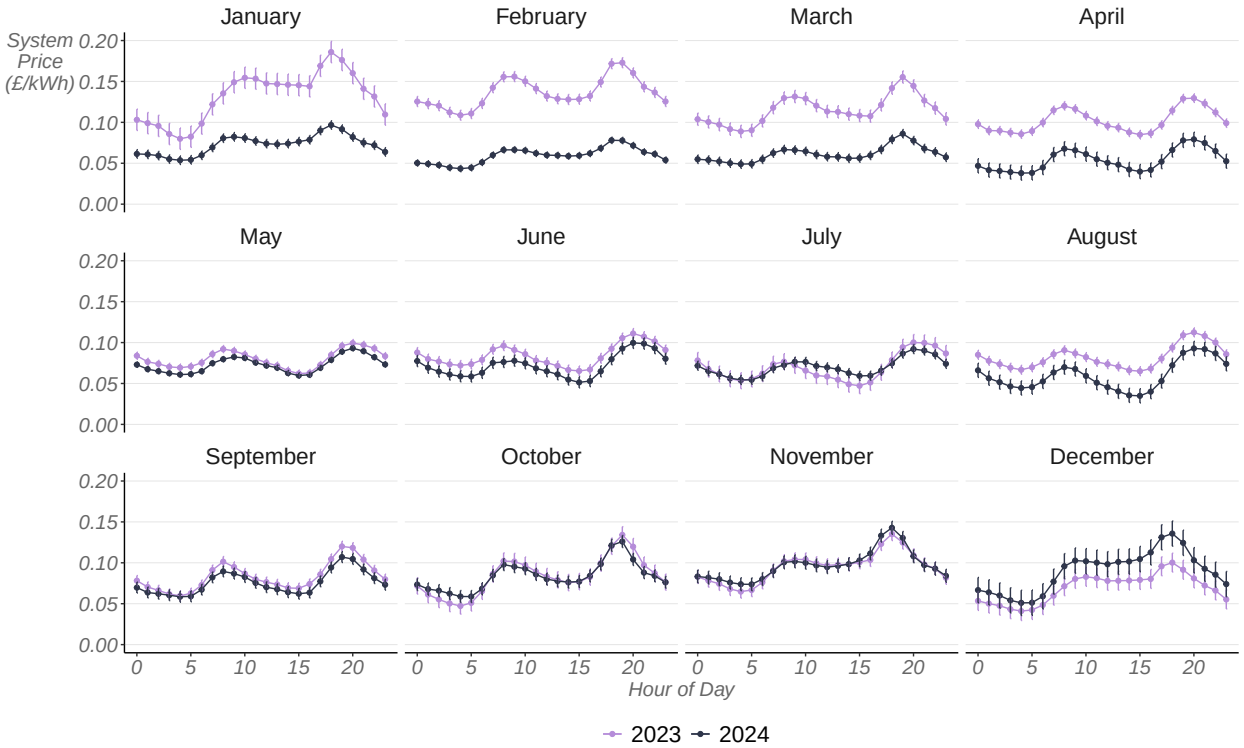
Notes: This figure reports effect of adopting the IO Go tariff on hourly electricity consumption (in kWh), split by hour of the day. These are estimated using a sample of 18,192 customers who first-ever enrolled in IO Go in 2023, weighted by whether the customer was previously on a time-of-use tariff. Estimates are computed using the Callaway and Sant'Anna (2021) estimator. Confidence intervals are shown at the 95% level.

Figure A15: Cohort Specific Difference-in-Differences Estimate of IO Go



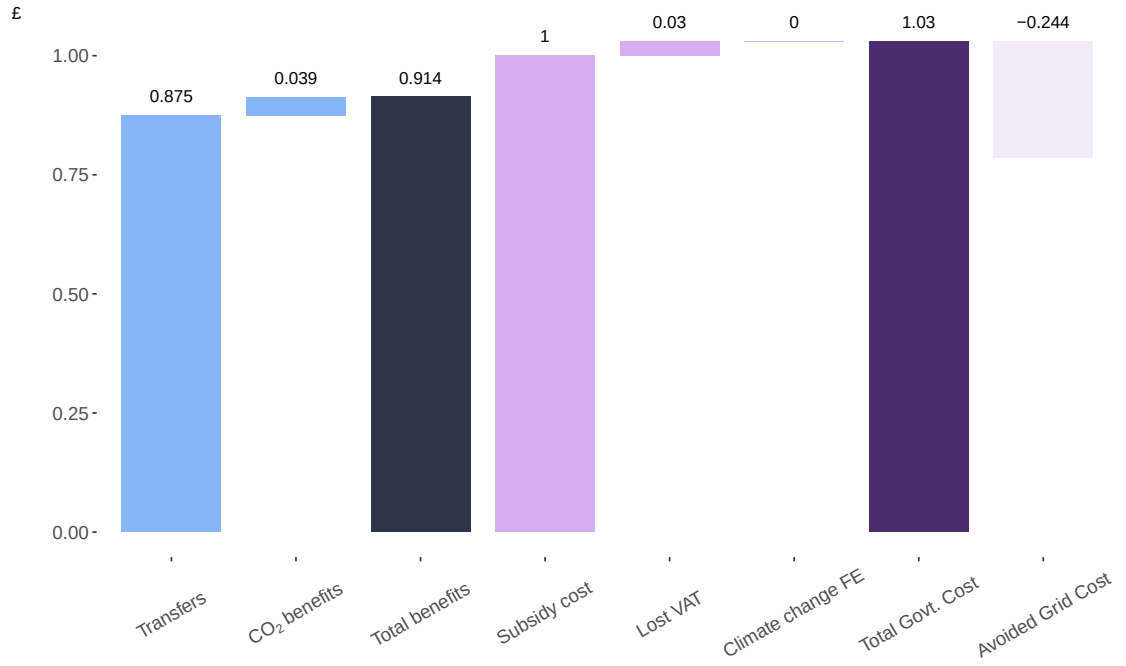
Notes: This figure reports cohort-specific estimates of the effect of adopting the IO Go tariff on hourly electricity consumption (in kWh) during (a) peak hours (16:30-20:30) and (b) off-peak hours (23:30-5:30). These are estimated using a sample of 18,192 customers who first-ever enrolled in IO Go in 2023, weighted by whether the customer was previously on a time-of-use tariff. Estimates are computed using the Callaway and Sant'Anna (2021) estimator. Confidence intervals are shown at the 95% level.

Figure A16: Day Ahead Price: 2023 vs 2024

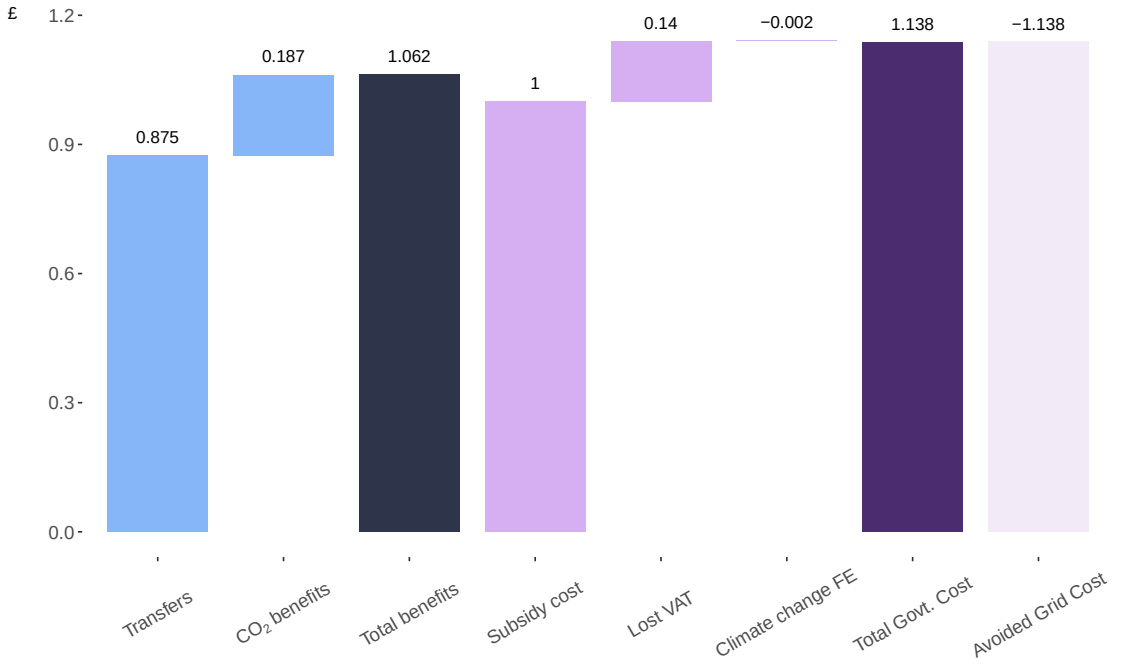


Notes: This figure reports mean day ahead prices in the UK for January 2023 to December 2024. Confidence intervals are shown at the 95% level.

Figure A17: Marginal value of public funds of subsidizing AI managed charging



(a) 1 year, 2024



(b) 5 years, 2024-2028

Notes: This figure presents the estimated costs and benefits of subsidizing adoption of the IO Go managed EV charging tariff over the course of (a) 1 year, and (b) 5 years. Customer surplus is based on a decomposition of marginal and inframarginal adoption under the £50/month offer, following the approach of Hendren and Sprung-Keyser (2020). Direct CO_2e benefits reflect emissions reductions from shifting electricity use to cleaner hours, scaled to marginal adopters. Indirect CO_2e benefits are excluded under the assumption that managed charging subsidies do not affect EV uptake among inframarginal adopters. Estimated costs to government include the subsidy, lost VAT revenue, and increased tax receipts from climate-related fiscal externalities. Grid stabilization benefits are shown separately, based on Franken et al. (2025) estimates of per-vehicle system savings under three scenarios. Only a share of these may accrue to government.

A.3 Deviations from the Pre-Analysis Plan

This experiment was pre-registered with the American Economic Association (AEA) under Trial No. 0013037. While we aimed to follow the pre-analysis plan (PAP) as closely as possible, a number of deviations became necessary in the course of implementation and analysis. Below, we detail the key departures from the PAP. Where applicable, corresponding changes to regression specifications are noted via footnotes in the main text.

- **Data:** Originally, our consumption data was to be aggregated on the hour (e.g., 00:00, 01:00, etc.). However, because the Intelligent Octopus Go (IO Go) tariff rates begin on the half-hour (e.g., 23:30 rather than 00:00), we adjusted our data aggregation accordingly. All half-hourly data was therefore aligned to begin on the half-hour mark.

- **Data:** We had pre-specified that electricity consumption data would be aggregated to the hour-day level. However, due to computational constraints in processing and analyzing high-frequency data, we instead aggregated consumption to the week × hour-of-day level.

- **Data:** We included 12 months of post-encouragement data, rather than the 6 months originally proposed, as we initiated our analysis later than anticipated and took advantage of the longer available data window. This change in data window does not substantively change our results, as shown in column (7) of Table A7, Table A8, and Table A9.

- **Analysis:** To improve precision and account for pre-existing consumption patterns, we included a control for baseline average hourly electricity consumption in our regression models. This adjustment was not pre-specified, but we show the version of the regression that does not include baseline consumption in column (8) of Table A7, Table A8, and Table A9. Coefficients for the peak consumption and overall consumption are similar, while the coefficient for the off-peak consumption changes from 0.515 (main specification) to 0.164 (no baseline consumption). We view this more as a loss of precision than of a substantive change in the underlying effect. As can be seen in Column (8), without baseline consumption as a control, the estimates become extremely noisy, and are statistically indistinguishable from the main specification.

- **Analysis:** Our pre-specified IV analysis planned to use the randomized encouragements as instruments for just IO Go adoption. However, given that the randomized encouragements increased enrollment in both Intelligent Octopus Go and Octopus Go, we defined treatment as adoption of either EV tariff, rather than IO Go alone. This adjustment was necessary to preserve the exclusion restriction in our instrumental variables analysis, a requirement we did not know would be necessary at the outset of the trial.

- **Analysis:** Following the concerns raised in Mogstad et al. (2021) regarding the interpretation of IV estimates with multiple instruments, we use a single binary instrument combining all encouragement groups. This approach helps avoid complications such as negative weights and improves interpretability under a common first-stage assumption. We report results from our pre-specified analysis using the four instruments separately in Figure 6, Table A7, Table A8, and Table A9. These include both the joint specification with all four instruments and separate regressions where each encouragement serves as an instrument individually.

- **Analysis:** Our pre-specified DiD sample included all customers who adopted IO Go in 2023. In the final analysis, our DiD analysis restricted the sample to customers who likely owned an EV by December 2022. This was to ensure that observed changes in consumption patterns are due to changes in charging behavior, rather than the initial uptake of EVs. We also re-weighted this restricted DiD sample by pre-adoption tariff to better understand the extent to which differences between RCT and DiD were due to pre-adoption tariff (this re-weighting was not prespecified).

- **Analysis:** For our DID analysis, we prespecified restricting control cohorts to customers who adopt IO GO within 30 days. However, our final analysis used an eight-week window. This was to balance comparability of treated and control groups against the length of the post-adoption estimation horizon. Comparability was assessed by examining pre-treatment trends, and we selected the longest horizon that yielded satisfactory pre-trend balance.

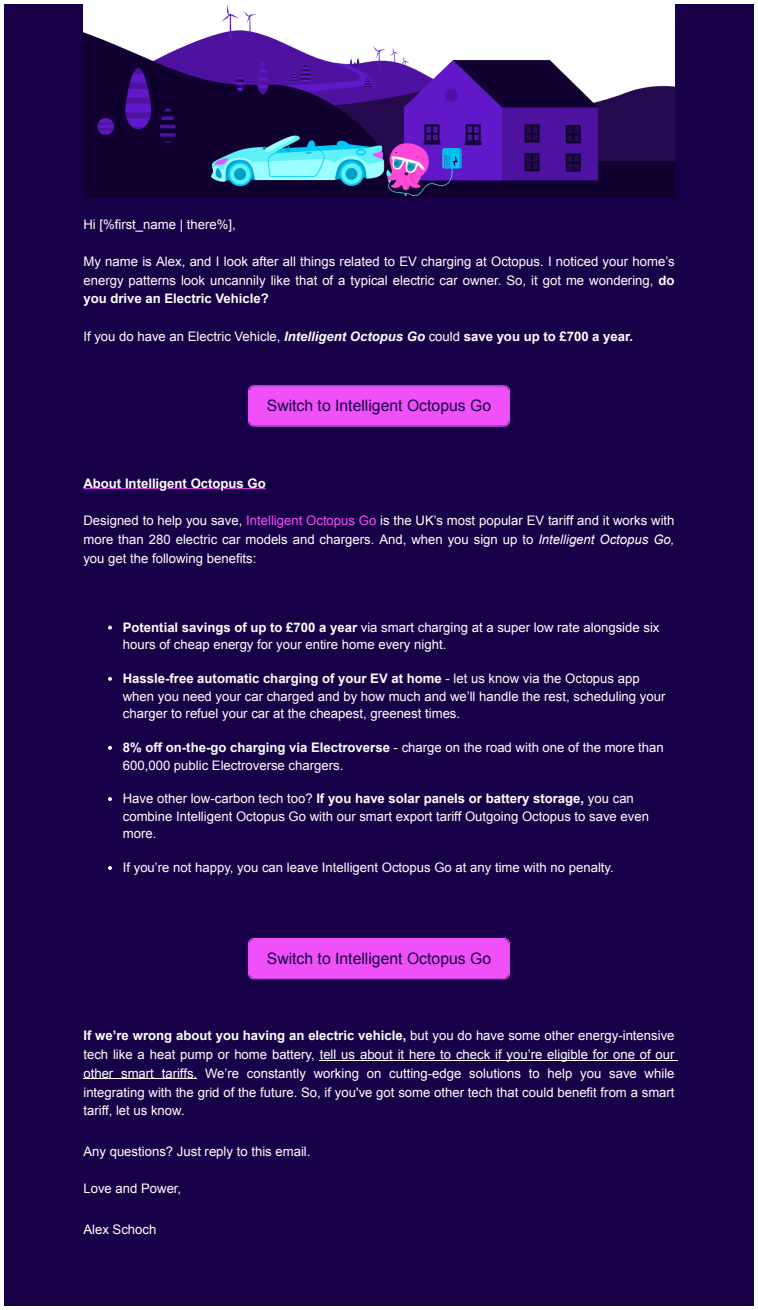
- **Welfare:** In keeping with the pre-analysis plan, we used the £150 incentive as a proxy subsidy, with J-PAL as proxy Government. However, we refined our welfare analysis to better capture the long-term and system-wide implications of the intervention. To this end, we looked at CO_2e impacts over a 10-year time period;

examined how the MVPF would change when we included avoided grid balancing costs; and included lost VAT as an extra cost to the Government of the subsidy. While these additions increase the measured benefits, we believe they more accurately reflect the full set of social returns that would accrue under real-world implementation. In keeping with the pre-analysis plan, we assumed that trial participants who enrolled in response to the £0/month email reflected inframarginal participants, while the incremental take-up in the £50/month group represented marginal adopters with an average willingness to pay equal to 50% of the subsidy, consistent with standard MVPF assumptions.

- **Welfare**: We originally pre-specified estimating CO_2e impacts using the ITT framework. We did not pre-specify the use of IV estimation for bills or CO_2e savings. We adopt the IV approach here because it more directly captures the causal effect of IO Go adoption—the quantity of substantive interest. While our pre-analysis plan focused on MVPF calculations rather than consumer bills, we now report bill savings as well, as they provide an important and policy-relevant measure of consumer benefits.

A.4 Additional details on design of field trial

A.4.1 Reproduction of email-based encouragements



Hi [%first_name | there%],

My name is Alex, and I look after all things related to EV charging at Octopus. I noticed your home's energy patterns look uncannily like that of a typical electric car owner. So, it got me wondering, **do you drive an Electric Vehicle?**

If you do have an Electric Vehicle, ***Intelligent Octopus Go*** could **save you up to £700 a year.**

Switch to Intelligent Octopus Go

About Intelligent Octopus Go

Designed to help you save, Intelligent Octopus Go is the UK's most popular EV tariff and it works with more than 280 electric car models and chargers. And, when you sign up to *Intelligent Octopus Go*, you get the following benefits:

- **Potential savings of up to £700 a year** via smart charging at a super low rate alongside six hours of cheap energy for your entire home every night.
- **Hassle-free automatic charging of your EV at home** - let us know via the Octopus app when you need your car charged and by how much and we'll handle the rest, scheduling your charger to refuel your car at the cheapest, greenest times.
- **8% off on-the-go charging via Electroverse** - charge on the road with one of the more than 600,000 public Electroverse chargers.
- Have other low-carbon tech too? **If you have solar panels or battery storage,** you can combine Intelligent Octopus Go with our smart export tariff Outgoing Octopus to save even more.
- If you're not happy, you can leave Intelligent Octopus Go at any time with no penalty.

Switch to Intelligent Octopus Go

If we're wrong about you having an electric vehicle, but you do have some other energy-intensive tech like a heat pump or home battery, tell us about it here to check if you're eligible for one of our other smart tariffs. We're constantly working on cutting-edge solutions to help you save while integrating with the grid of the future. So, if you've got some other tech that could benefit from a smart tariff, let us know.

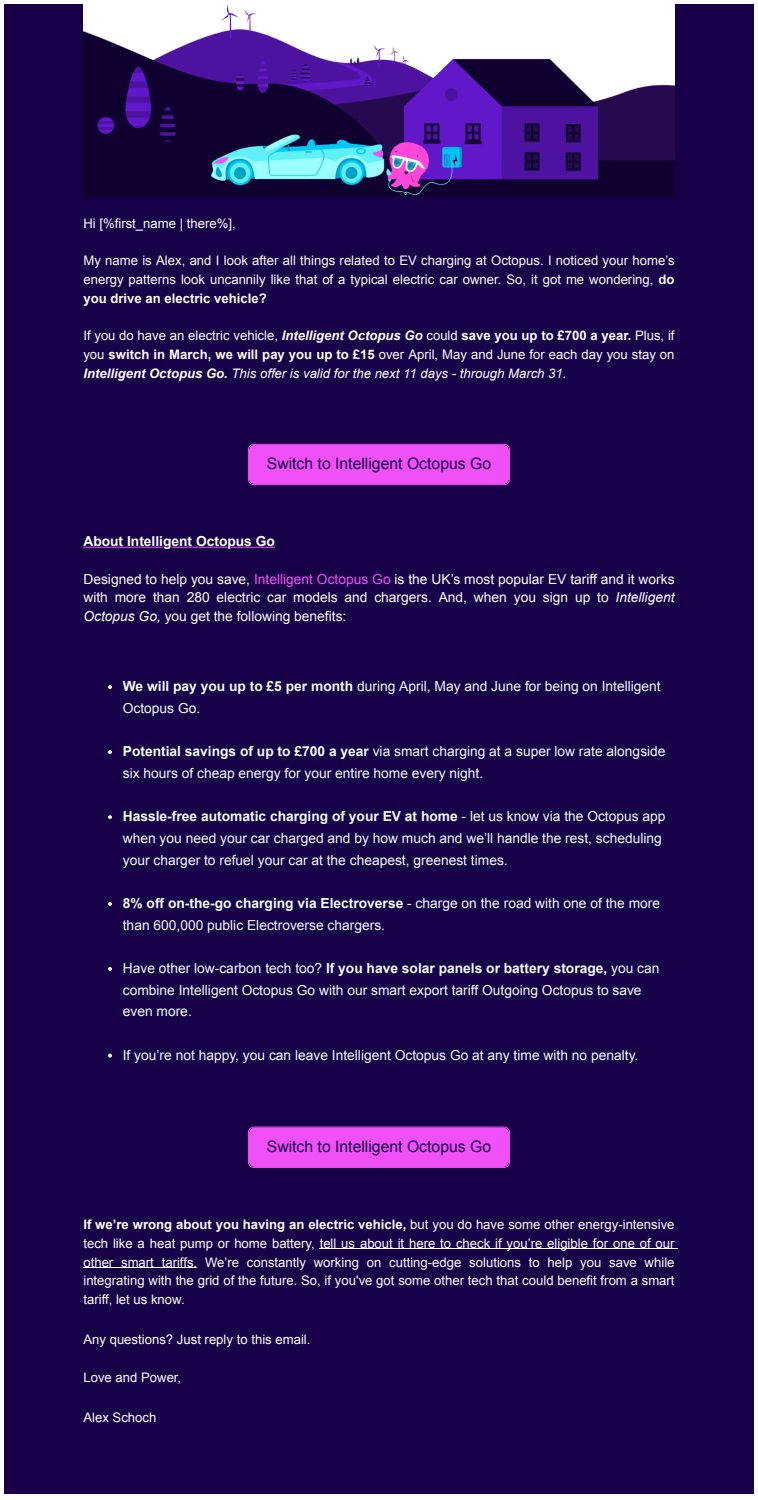
Any questions? Just reply to this email.

Love and Power,

Alex Schoch

Figure A18: Randomized Encouragement Group 1 (Email + £0/Mth)

Note: The subject line read FYI: Do you drive an EV? You could save hundreds on your energy bills.



Hi [%first_name | there%],

My name is Alex, and I look after all things related to EV charging at Octopus. I noticed your home's energy patterns look uncannily like that of a typical electric car owner. So, it got me wondering, **do you drive an electric vehicle?**

If you do have an electric vehicle, ***Intelligent Octopus Go*** could **save you up to £700 a year.** Plus, if you **switch in March, we will pay you up to £15** over April, May and June for each day you stay on ***Intelligent Octopus Go.*** *This offer is valid for the next 11 days - through March 31.*

Switch to Intelligent Octopus Go

About Intelligent Octopus Go

Designed to help you save, Intelligent Octopus Go is the UK's most popular EV tariff and it works with more than 280 electric car models and chargers. And, when you sign up to *Intelligent Octopus Go,* you get the following benefits:

- **We will pay you up to £5 per month** during April, May and June for being on Intelligent Octopus Go.
- **Potential savings of up to £700 a year** via smart charging at a super low rate alongside six hours of cheap energy for your entire home every night.
- **Hassle-free automatic charging of your EV at home** - let us know via the Octopus app when you need your car charged and by how much and we'll handle the rest, scheduling your charger to refuel your car at the cheapest, greenest times.
- **8% off on-the-go charging via Electroverse** - charge on the road with one of the more than 600,000 public Electroverse chargers.
- Have other low-carbon tech too? **If you have solar panels or battery storage,** you can combine Intelligent Octopus Go with our smart export tariff Outgoing Octopus to save even more.
- If you're not happy, you can leave Intelligent Octopus Go at any time with no penalty.

Switch to Intelligent Octopus Go

If we're wrong about you having an electric vehicle, but you do have some other energy-intensive tech like a heat pump or home battery, tell us about it here to check if you're eligible for one of our other smart tariffs. We're constantly working on cutting-edge solutions to help you save while integrating with the grid of the future. So, if you've got some other tech that could benefit from a smart tariff, let us know.

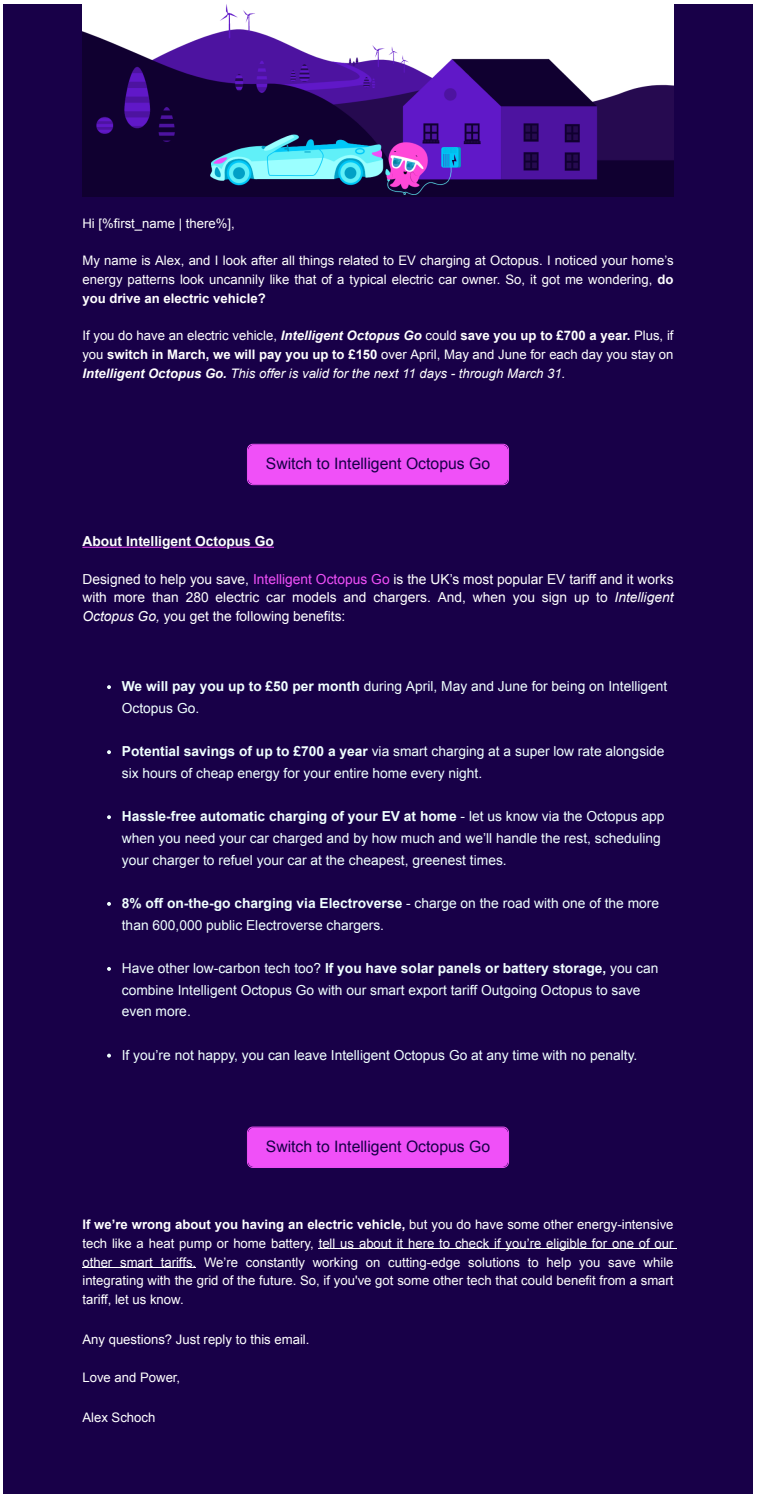
Any questions? Just reply to this email.

Love and Power,

Alex Schoch

Figure A19: Randomized Encouragement Group 2 (Email + £5/Mth)

Note: The subject line read FYI: Do you drive an EV? You could save hundreds on your energy bills and get paid £15.



Hi [%first_name | there%],

My name is Alex, and I look after all things related to EV charging at Octopus. I noticed your home's energy patterns look uncannily like that of a typical electric car owner. So, it got me wondering, **do you drive an electric vehicle?**

If you do have an electric vehicle, ***Intelligent Octopus Go*** could **save you up to £700 a year.** Plus, if you **switch in March, we will pay you up to £150** over April, May and June for each day you stay on ***Intelligent Octopus Go.*** *This offer is valid for the next 11 days - through March 31.*

Switch to Intelligent Octopus Go

About Intelligent Octopus Go

Designed to help you save, Intelligent Octopus Go is the UK's most popular EV tariff and it works with more than 280 electric car models and chargers. And, when you sign up to *Intelligent Octopus Go,* you get the following benefits:

- **We will pay you up to £50 per month** during April, May and June for being on Intelligent Octopus Go.
- **Potential savings of up to £700 a year** via smart charging at a super low rate alongside six hours of cheap energy for your entire home every night.
- **Hassle-free automatic charging of your EV at home** - let us know via the Octopus app when you need your car charged and by how much and we'll handle the rest, scheduling your charger to refuel your car at the cheapest, greenest times.
- **8% off on-the-go charging via Electroverse** - charge on the road with one of the more than 600,000 public Electroverse chargers.
- Have other low-carbon tech too? **If you have solar panels or battery storage,** you can combine Intelligent Octopus Go with our smart export tariff Outgoing Octopus to save even more.
- If you're not happy, you can leave Intelligent Octopus Go at any time with no penalty.

Switch to Intelligent Octopus Go

If we're wrong about you having an electric vehicle, but you do have some other energy-intensive tech like a heat pump or home battery, tell us about it here to check if you're eligible for one of our other smart tariffs. We're constantly working on cutting-edge solutions to help you save while integrating with the grid of the future. So, if you've got some other tech that could benefit from a smart tariff, let us know.

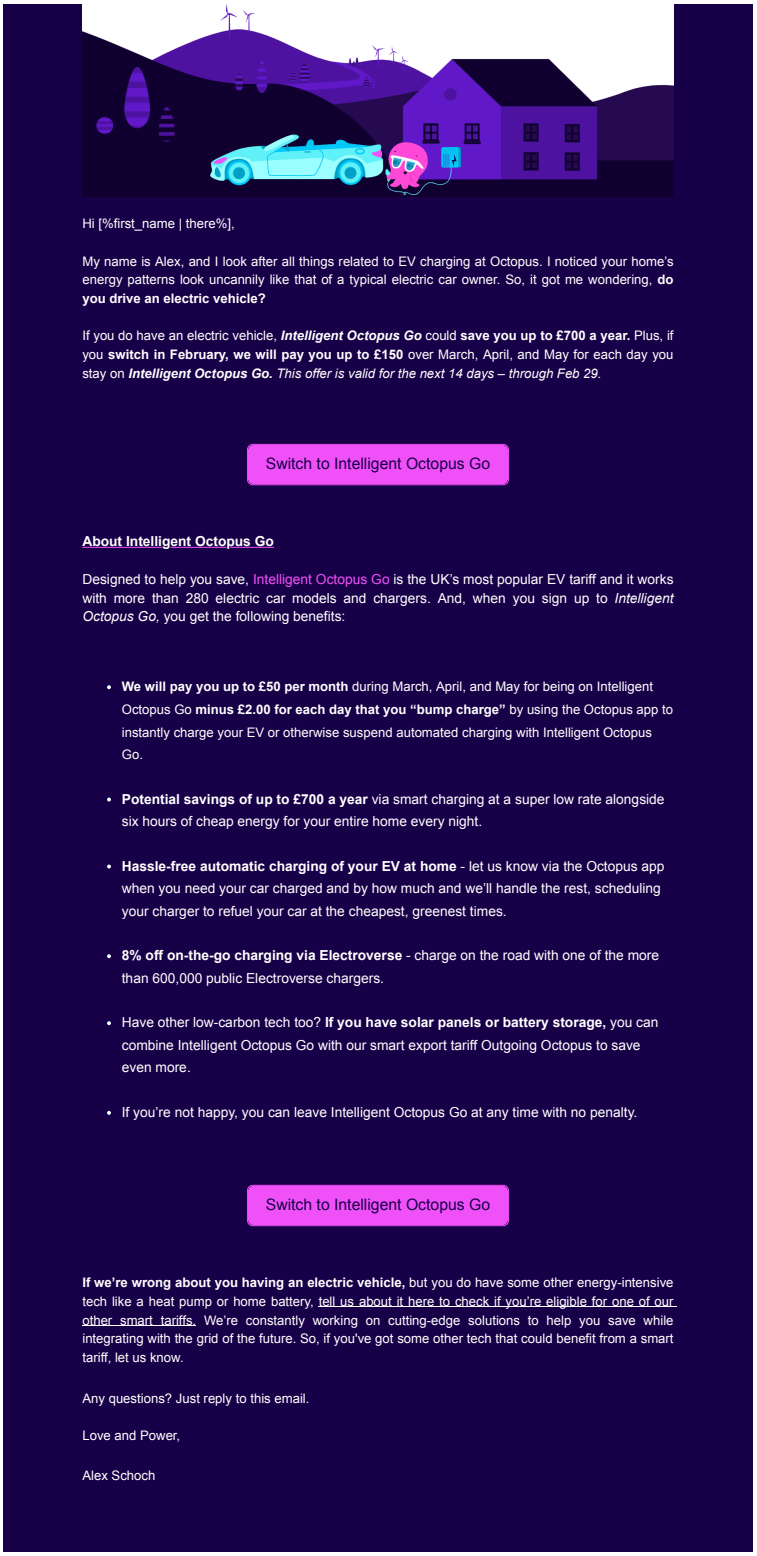
Any questions? Just reply to this email.

Love and Power,

Alex Schoch

Figure A20: Randomized Encouragement Group 3 (Email + £50/Mth)

Note: The subject line read FYI: Do you drive an EV? You could save hundreds on your energy bills and get paid £150.



Hi [%first_name | there%],

My name is Alex, and I look after all things related to EV charging at Octopus. I noticed your home's energy patterns look uncannily like that of a typical electric car owner. So, it got me wondering, **do you drive an electric vehicle?**

If you do have an electric vehicle, ***Intelligent Octopus Go*** could **save you up to £700 a year.** Plus, if you **switch in February, we will pay you up to £150** over March, April, and May for each day you stay on ***Intelligent Octopus Go.*** *This offer is valid for the next 14 days – through Feb 29.*

Switch to Intelligent Octopus Go

About Intelligent Octopus Go

Designed to help you save, Intelligent Octopus Go is the UK's most popular EV tariff and it works with more than 280 electric car models and chargers. And, when you sign up to *Intelligent Octopus Go,* you get the following benefits:

- **We will pay you up to £50 per month** during March, April, and May for being on Intelligent Octopus Go **minus £2.00 for each day that you "bump charge"** by using the Octopus app to instantly charge your EV or otherwise suspend automated charging with Intelligent Octopus Go.
- **Potential savings of up to £700 a year** via smart charging at a super low rate alongside six hours of cheap energy for your entire home every night.
- **Hassle-free automatic charging of your EV at home** - let us know via the Octopus app when you need your car charged and by how much and we'll handle the rest, scheduling your charger to refuel your car at the cheapest, greenest times.
- **8% off on-the-go charging via Electroverse** - charge on the road with one of the more than 600,000 public Electroverse chargers.
- Have other low-carbon tech too? **If you have solar panels or battery storage,** you can combine Intelligent Octopus Go with our smart export tariff Outgoing Octopus to save even more.
- If you're not happy, you can leave Intelligent Octopus Go at any time with no penalty.

Switch to Intelligent Octopus Go

If we're wrong about you having an electric vehicle, but you do have some other energy-intensive tech like a heat pump or home battery, tell us about it here to check if you're eligible for one of our other smart tariffs. We're constantly working on cutting-edge solutions to help you save while integrating with the grid of the future. So, if you've got some other tech that could benefit from a smart tariff, let us know.

Any questions? Just reply to this email.

Love and Power,

Alex Schoch

Figure A21: Randomized Encouragement Group 4 (Email + £50/Mth, No Bump)

Note: The subject line read FYI: Do you drive an EV? You could save hundreds on your energy bills and get paid £150.

A.4.2 Block randomization implementation

We implemented block randomization on trial participant account identifiers using Mahalanobis distance calculated over the following pre-encouragement variables:

1. Tenure with Octopus: Years since a customer's earliest import tariff contract with Octopus Energy (as of August 31, 2023).
2. Total Consumption (kWh): Total electricity use (kWh) from February 15 to August 31, 2023, aggregated across all half-hourly smart meter readings.
3. Consumption Variability (kWh): Standard deviation of half-hourly consumption over the same period.
4. Not Always Octopus Energy Customer: Binary flag for whether a trial participant originally joined Octopus via acquisition (e.g., from Bulb, Co-op) or Supplier of Last Resort procedures.
5. Smart Tariff Onboarding Attempts: Count of historical attempts to enroll in Octopus smart tariffs (e.g., Intelligent Octopus Go) prior to February 15, 2024. Includes cases where customers initiated but did not complete the process.
6. DNO Region: Categorical variable for the customer's Distribution Network Operator region. For accounts with multiple active meter points, we used the region linked to the most recent tariff as of August 31, 2023.
7. Expected Structural Winnings: Estimated monetary difference between a customer's actual electricity cost (under observed tariff contracts) and a counterfactual IO Go contract from February 15 to August 31, 2023. We assumed an off-peak rate of £0.075/kWh and a peak rate of £0.30/kWh for IO Go, and ignored taxes, standing charges, and regional price variation.
8. Peak-Hour Consumption Share: Proportion of total electricity usage (Feb–Aug 2023) occurring during 16:00–20:00, a period of high grid constraint.

Full details of structural savings calculations and data preparation are available upon request.

A.5 Theoretical Model

A.5.1 Environment and Notation

Time is divided into settlement periods $h = 1,\dots,H$. Each EV-owning household i requires E_i kWh by ready-by time T_i, has plug-in availability $\mathcal{A}_i \subseteq \{1,\dots,H\}$, maximum charge rate $\bar{x}_i$, and charging efficiency η_{ih}. Baseline (non-EV) load is b_{ih}, with aggregate $b_h \equiv \sum_i b_{ih}$. Total system load is

$$Q_h \;=\; b_h \;+\; \sum_i x_{ih}.$$

Retailer costs are $c_h(Q_h)$ (convex) and ancillary/avoided benefits $r_h(Q_h)$ (concave). Define the system shadow price

$$\tilde{c}_h \;=\; c_h'(Q_h) - r_h'(Q_h), \tag{7}$$

as in convex scheduling (Boyd and Vandenberghe, 2004; Joskow and Tirole, 2006*b*). Households derive mobility utility $U_i(E_i)$ (increasing, concave) and timing disutility $\psi_i(h) \geq 0$. Under RTP, they face attention/optimization cost $a_i \geq 0$ and risk penalty $\gamma_i \geq 0$ on bill variance (Borenstein, 2007). Under IO (AI scheduling), they can override at a hassle cost $\phi_i \geq 0$.

We compare four tariffs:[51]

1. **Flat**: price p^{flat}.
2. **ToU**: $p^{\text{peak}} > p^{\text{off}}$ across fixed windows.
3. **RTP**: hourly p_h^{RTP}; households self-schedule.
4. **IO (AI managed)**: centralized schedule $\{x_{ih}\}$ against $\{\tilde{c}_h\}$; users may override.

[51] Our framework builds on Joskow and Tirole (2006*b*), who study the welfare properties of RTP under convex scheduling. Their analysis contrasts RTP with the absence of RTP. Where we depart is in the introduction of (i) algorithmic intermediation (IO) as a distinct pricing/coordination regime, and (ii) frictions such as override behavior, execution costs (m_{ih}, ϕ_i), and aggregate override thresholds ($\bar{\beta}^\star$). These extensions allow us to bring the theoretical framework into closer alignment with experimental data, in particular by capturing behavioral deviations from the frictionless convex scheduling benchmark.

The feasible EV-charging schedules satisfy[52]

$$\sum_{h \le T_i,\, h \in \mathcal{A}_i} \eta_{ih} x_{ih} \ge E_i, \qquad 0 \le x_{ih} \le \bar{x}_i. \tag{8}$$

Elasticity by operative signal. Let π_h^k be the operative signal under regime $k \in \{\text{IO}, \text{RTP}, \text{ToU}\}$ with $\pi_h^{\text{IO}} = \tilde{c}_h$, $\pi_h^{\text{RTP}} = p_h^{\text{RTP}}$, $\pi_h^{\text{ToU}} = p_h^{\text{ToU}}$. Define

$$\epsilon_h^k \;=\; \frac{\partial \ln Q_h^k}{\partial \ln \pi_h^k}, \tag{9}$$

evaluated in high plug-in hours ($Q_h^k > 0$). Let e_h denote marginal CO_2 emissions (kg/kWh).

A.5.2 Household and Aggregator Problems

RTP household problem. A household chooses $\{x_{ih}\}$ to minimize expected cost plus frictions:

$$\min_{\{x_{ih}\}} \quad \mathbb{E}\Big[\sum_h p_h^{\text{RTP}} x_{ih}\Big] + \gamma_i \,\text{Var}\Big(\sum_h p_h^{\text{RTP}} x_{ih}\Big) + \sum_h \psi_i(h)\, x_{ih} + a_i \cdot \mathbf{1}\{\text{actively scheduling}\} \tag{10}$$

$$\text{s.t.} \quad \sum_{h \le T_i,\, h \in \mathcal{A}_i} \eta_{ih} x_{ih} \ge E_i, \quad 0 \le x_{ih} \le \bar{x}_i.$$

At interior hours, first-order conditions (FOC) take the (mean–variance) form

$$\mathbb{E}[p_h^{\text{RTP}}] + 2\gamma_i \Big(p_h^{\text{RTP}},\, \sum_\ell p_\ell^{\text{RTP}} x_{i\ell}\Big) + \psi_i(h) \;=\; \mu_i \eta_{ih}, \tag{11}$$

with $\mu_i \ge 0$ the KKT multiplier on the energy requirement. The attention cost $a_i > 0$ drives corner solutions by discouraging small reallocations.[53]

[52] We model EV-only RTP for comparability to IO (which optimizes EV load). If non-EV end uses are price-responsive under whole-household RTP, the welfare gap ΔW_0 generalizes accordingly; our empirical mapping focuses on the EV subproblem.

[53] Derivation of (11): Rewrite the energy requirement as a KKT-friendly inequality $g_i(x) \equiv E_i - \sum_{h \le T_i,\, h \in \mathcal{A}_i} \eta_{ih} x_{ih} \le 0$ with multiplier $\mu_i \ge 0$. Let $S_i(x,p) \equiv \sum_h p_h^{\text{RTP}} x_{ih}$ be the (random) daily bill. Ignoring the fixed attention cost a_i (which

IO aggregator problem. The IO scheduler minimizes social cost plus timing disutility:

$$\min_{\{x_{ih}\}} \quad \sum_h \Big[c_h(Q_h) - r_h(Q_h)\Big] + \sum_{i,h} \psi_i(h)\, x_{ih} \tag{12}$$
$$\text{s.t.} \quad \sum_{h \le T_i,\, h \in \mathcal{A}_i} \eta_{ih} x_{ih} \ge E_i, \quad 0 \le x_{ih} \le \bar{x}_i.$$

The objective is convex and the constraints are affine; under Slater's condition (capacity slack), KKT are necessary and sufficient (Boyd and Vandenberghe, 2004).

A.5.3 Core Lemmas and Propositions

Standing assumptions for this subsection. (i) $c_h(\cdot)$ convex, $r_h(\cdot)$ concave, so $\tilde{c}_h$ is non-decreasing in Q_h; (ii) price-taking (a single household does not affect $\tilde{c}_h$); (iii) Slater's condition holds for Eq. (12)–Eq. (8) (there exists feasible slack).

Lemma A.5.1 (Peak shaving under IO (merit-order via KKT))**.** *Under assumptions (i)–(iii), define the adjusted hourly cost*

$$\kappa_{ih} \equiv \frac{\tilde{c}_h + \psi_i(h)}{\eta_{ih}}.$$

(a) For each i, an IO optimum allocates charging to hours with the lowest available κ_{ih} subject to Eq. (8). (b) If there exist feasible peak and off-peak hours p, o with $\kappa_{ip} > \kappa_{io}$, then $x_{ip} > 0$

---

does not depend on $\{x_{ih}\}$ and thus only affects the extensive/corner decision), the Lagrangian is

$$\mathcal{L} = \mathbb{E}[S_i] + \gamma_i \operatorname{Var}(S_i) + \sum_h \psi_i(h)\, x_{ih} + \mu_i \left(E_i - \sum_{\substack{h \le T_i \\ h \in \mathcal{A}_i}} \eta_{ih} x_{ih} \right) + \sum_h \alpha_{ih}(x_{ih} - \bar{x}_i) - \sum_h \beta_{ih} x_{ih},$$

with box-constraint multipliers $\alpha_{ih}, \beta_{ih} \ge 0$. Since prices are exogenous to the household, $\frac{\partial}{\partial x_{ih}} \mathbb{E}[S_i] = \mathbb{E}[p_h^{\mathrm{RTP}}]$ and $\frac{\partial}{\partial x_{ih}} \operatorname{Var}(S_i) = 2\big(p_h^{\mathrm{RTP}}, S_i\big)$ because $\operatorname{Var}(S_i) = \mathbb{E}\Big[(S_i - \mathbb{E}[S_i])^2\Big] \Rightarrow \partial \operatorname{Var}(S_i)/\partial x_{ih} = 2\,\mathbb{E}\Big[(S_i - \mathbb{E}[S_i])(p_h^{\mathrm{RTP}} - \mathbb{E}[p_h^{\mathrm{RTP}}])\Big]$. Stationarity w.r.t. x_{ih} gives

$$\mathbb{E}[p_h^{\mathrm{RTP}}] + 2\gamma_i \Big(p_h^{\mathrm{RTP}}, \sum_\ell p_\ell^{\mathrm{RTP}} x_{i\ell}\Big) + \psi_i(h) - \mu_i \eta_{ih} + \alpha_{ih} - \beta_{ih} = 0.$$

At interior hours ($0 < x_{ih} < \bar{x}_i$), $\alpha_{ih} = \beta_{ih} = 0$, yielding

$$\mathbb{E}[p_h^{\mathrm{RTP}}] + 2\gamma_i \Big(p_h^{\mathrm{RTP}}, \sum_\ell p_\ell^{\mathrm{RTP}} x_{i\ell}\Big) + \psi_i(h) = \mu_i \eta_{ih},$$

which is (11). Complementary slackness for α_{ih}, β_{ih} covers corner hours.

implies all lower-cost hours $\{o : \kappa_{io} < \kappa_{ip}\}$ are saturated or infeasible. Aggregating over i, IO (weakly) lowers peak load and (weakly) raises off-peak load while daily kWh per EV is unchanged.

Proof. Let $Q_h = b_h + \sum_i x_{ih}$. The Lagrangian is

$$\mathcal{L}(\{x_{ih}\},\{\mu_i\},\{\mu^{\pm}_{ih}\}) = \sum_h \Big[c_h(Q_h) - r_h(Q_h)\Big] + \sum_{i,h} \psi_i(h) x_{ih} + \sum_i \mu_i \left(E_i - \sum_{\substack{h \le T_i \\ h \in \mathcal{A}_i}} \eta_{ih} x_{ih} \right) + \sum_{i,h} \mu^{+}_{ih}(x_{ih} - \bar{x}_i) - \sum_{i,h} \mu^{-}_{ih} x_{ih}. \quad (13)$$

Stationarity w.r.t. x_{ih} gives

$$\frac{\partial \mathcal{L}}{\partial x_{ih}} = \underbrace{c'_h(Q_h) - r'_h(Q_h)}_{\tilde{c}_h} + \psi_i(h) - \mu_i \eta_{ih} + \mu^{+}_{ih} - \mu^{-}_{ih} = 0.$$

If $0 < x_{ih} < \bar{x}_i$ then $\mu^{\pm}_{ih} = 0$ and $\mu_i = \frac{\tilde{c}_h + \psi_i(h)}{\eta_{ih}} \equiv \kappa_{ih}$. If $x_{ih} = 0$ then $\mu^{-}_{ih} \ge 0$ and $\kappa_{ih} \ge \mu_i$; if $x_{ih} = \bar{x}_i$ then $\mu^{+}_{ih} \ge 0$ and $\kappa_{ih} \le \mu_i$. Hence x_{ih} is nonincreasing in κ_{ih}: the IO scheduler fills the lowest adjusted-cost hours first, up to feasibility. If $x_{ip} > 0$ while some feasible o has $\kappa_{io} < \kappa_{ip}$ and $x_{io} < \bar{x}_i$, then $\kappa_{ip} \le \mu_i \le \kappa_{io}$, a contradiction. Summing over i yields load shifted from high-$\tilde{c}_h$ to low-$\tilde{c}_h$ hours, conserving daily energy by Eq. (8).

Lemma A.5.2 (Elasticity ordering)**.** *In high plug-in hours,*

$$|\epsilon^{\mathrm{IO}}_h| \;\ge\; |\epsilon^{\mathrm{RTP}}_h| \;\ge\; |\epsilon^{\mathrm{ToU}}_h|.$$

Proof. Step 1 (IO upper-bound response). Consider a small perturbation $d\tilde{c}$ concentrated in hour h. Differentiating the KKT system in Lemma A.5.1 yields a linear system in $\{dx_{ih}, d\mu_i, d\mu^{\pm}_{ih}\}$ whose solution preserves the merit order: $dx_{ih} \le 0$ at the perturbed hour and $dx_{i\ell} \ge 0$ at some lower-$\kappa_{i\ell}$ hours, with $\sum_{\ell \le T_i, \ell \in \mathcal{A}_i} \eta_{i\ell} dx_{i\ell} = 0$ (intra-day reallocation). Aggregating over i,

$$\frac{\partial Q^{\mathrm{IO}}_h}{\partial \tilde{c}_h} \;\le\; 0, \qquad \left|\frac{\partial Q^{\mathrm{IO}}_h}{\partial \tilde{c}_h}\right| \text{ is maximal subject to Eq. (8).}$$

Thus $|\epsilon^{\mathrm{IO}}_h|$ is an upper bound among feasible reallocations.

Step 2 (RTP attenuation by risk/attention). From the household FOC Eq. (11), totally differentiate across hours. Stacking in vector form,

$$H_i\, dx_i \;=\; -\Big(I + 2\gamma_i\, \Sigma_p\, S_i\Big)\, dp,$$

where H_i is the Hicksian substitution matrix (negative semidefinite), Σ_p is the covariance matrix of prices, and S_i maps p to $\sum_\ell p_\ell x_{i\ell}$. The matrix $(I + 2\gamma_i \Sigma_p S_i)$ is positive semidefinite for $\gamma_i \geq 0$; multiplying a negative semidefinite H_i by such a factor weakly *shrinks* the response (Loewner order). Furthermore, $a_i > 0$ creates inactive hours (corners), reducing the response support. Aggregating over i preserves attenuation:

$$\left\| \frac{\partial Q^{\mathrm{RTP}}}{\partial p} \right\| \;\leq\; \left\| \frac{\partial Q^{\mathrm{IO}}}{\partial \tilde{c}} \right\|.$$

Step 3 (ToU as projection). Let P be the block-averaging operator mapping hourly prices to ToU blocks ($P^2 = P$, $\|P\|_2 \leq 1$). For small perturbations, $dQ^k \approx H^k\, d\pi^k$ with H^k negative semidefinite. Under ToU, $d\pi^{\mathrm{ToU}} = P\, d\pi^{\mathrm{RTP}}$, hence

$$\|dQ^{\mathrm{ToU}}\|_2 \;=\; \|H^{\mathrm{ToU}}\, P\, d\pi^{\mathrm{RTP}}\|_2 \;\leq\; \|H^{\mathrm{RTP}}\, d\pi^{\mathrm{RTP}}\|_2,$$

so $|\partial Q_h^{\mathrm{ToU}}/\partial p_h| \leq |\partial Q_h^{\mathrm{RTP}}/\partial p_h|$. Combining Steps 1–3 yields the stated ordering.

54

Lemma A.5.3 (Welfare ranking with frictions)**.** *With $a_i, \gamma_i > 0$,*

$$W^{\mathrm{IO}} \;\geq\; W^{\mathrm{RTP}} \;\geq\; W^{\mathrm{ToU}} \;\geq\; W^{\mathrm{Flat}}.$$

Proof. Let W^k be per-EV social welfare (net of pure transfers). IO maximizes W for the

[54] We do not analyze a bill-variance ordering in our empirical set-up, but with $\mathrm{Bill}^k = \sum_h p_h^k X_h^k$, the law of total variance gives

$$\mathrm{Var}(\mathrm{Bill}^k) = \mathbb{E}\left[\mathrm{Var}\left(\sum_h p_h^k X_h^k \,\middle|\, \mathbf{p}^k\right)\right] + \mathrm{Var}\left(\mathbb{E}\left[\sum_h p_h^k X_h^k \,\middle|\, \mathbf{p}^k\right]\right).$$

Moving from RTP to ToU replaces $\mathbf{p}$ by $P\mathbf{p}$, where P is a contraction ($\|P\|_2 \leq 1$) and a Blackwell coarsening. Both the within-state term and the between-state term weakly fall. Under IO (EV sub-load), consumers face a flat off-peak retail rate; wholesale volatility is internalized by the aggregator, yielding

$$\mathrm{Var}(\mathrm{Bill}^{\mathrm{RTP}}) \;\geq\; \mathrm{Var}(\mathrm{Bill}^{\mathrm{ToU}}) \;\geq\; \mathrm{Var}(\mathrm{Bill}^{\mathrm{IO}}).$$

EV sub-load given $\{\tilde{c}_h\}$ (Lemma A.5.1 and KKT optimality). Relative to IO,

$$W^{\mathrm{RTP}} = W^{\mathrm{IO}} - \underbrace{\Delta W_{\mathrm{coord}}}_{\geq 0} - \underbrace{\mathbb{E}[a_i]}_{>0} - \underbrace{\mathbb{E}[\gamma_i \operatorname{Var}(\mathrm{Bill}^{\mathrm{RTP}})]}_{>0}.$$

ToU coarsens the signal (projection loss $\Delta W_{\mathrm{gran}} \geq 0$) and Flat removes intertemporal incentives entirely. Hence the stated ordering.

Proposition A.5.1 (Emissions ordering)**.** *If e_h is (weakly) lower off-peak, then relative to baseline: IO achieves the largest emissions reduction, followed by RTP, then ToU, then Flat.*

Proof. By Lemma A.5.1, IO shifts the most load into low-$\tilde{c}_h$ hours, which coincide with low-e_h by assumption; RTP and ToU shift less; Flat does not reallocate. Summing $e_h x_{ih}$ across hours yields the ordering.

A.5.4 Overrides ("Bump Charging"): Behavior and Welfare

Override decision and probability. A household overrides in hour h iff

$$v_{ih} > m_{ih} + \phi_i, \tag{14}$$

where v_{ih} is the immediate utility from charging now and m_{ih} the marginal benefit from deferring to the IO-planned hour. If F_{ih} is the CDF of v_{ih}, the override probability is

$$\beta_{ih} \;=\; \Pr\left[v_{ih} > m_{ih} + \phi_i\right] \;=\; 1 - F_{ih}(m_{ih} + \phi_i). \tag{15}$$

Lemma A.5.4 (Override monotonicity)**.** *If F_{ih} is nondecreasing, then β_{ih} is weakly decreasing in ϕ_i and in m_{ih}.*

Proof. Differentiate Eq. (15) (where densities exist): $\partial\beta_{ih}/\partial\phi_i = -f_{ih}(m_{ih} + \phi_i) \leq 0$ and similarly for m_{ih}. Without densities, monotonicity follows from the CDF order.

Lemma A.5.5 (Welfare effect of an override)**.** *Let h' be the IO-planned hour and h the override hour for EV energy $q_{ih} > 0$. Private net benefit is $v_{ih} - m_{ih} - \phi_i > 0$ when overriding. Social cost changes by $(\tilde{c}_h - \tilde{c}_{h'})\, q_{ih}$; if $\tilde{c}_h > \tilde{c}_{h'}$ the override raises system cost.*

Proof. Private part is by Eq. (14). Social part: the IO plan equalizes marginal costs across used hours; deviating to higher-$\tilde{c}_h$ increases procurement net of r_h by $(\tilde{c}_h - \tilde{c}_{h'})\, q_{ih}$.

Proposition A.5.2 (Aggregate override rate)**.** *If v_{ih} are i.i.d. with CDF F_v and $m_{ih} \in \{m^{\text{chg}}, m^{\text{def}}\}$ depending on whether IO planned charging in h, and if ρ is the share of hours planned to charge, then*

$$\bar{\beta} \;=\; (1-\rho)\Big[1 - F_v(m^{\text{def}} + \bar{\phi})\Big] \;+\; \rho\Big[1 - F_v(m^{\text{chg}} + \bar{\phi})\Big],$$

where $\bar{\phi}$ is the (mean) hassle cost.

Proof. Law of total probability conditioning on planned status; apply Eq. (15) in each state and average.

Welfare with overrides and crossover. Let W^{IO} and W^{RTP} be per-EV welfare absent overrides and define the baseline gap

$$\Delta W_0 \;=\; W^{\text{IO}} - W^{\text{RTP}}. \tag{16}$$

If overrides occur at rate $\bar{\beta}$ with per-override loss λ, IO welfare becomes

$$W^{\text{IO+O}} \;=\; W^{\text{IO}} - \bar{\beta}\,\lambda. \tag{17}$$

We calibrate λ by

$$\hat{\lambda} \;\approx\; p_{\text{peak}} \cdot \Delta\tilde{c}_{\text{eff}} \cdot q^{O}, \qquad \lambda^{\text{UB}} \;=\; (\max_h \tilde{c}_h - \min_h \tilde{c}_h) \cdot q^{O}_{\max}, \tag{18}$$

where p_{peak} is the probability an override lands in peak, $\Delta\tilde{c}_{\text{eff}}$ the peak–off-peak spread (£/kWh), and q^{O} kWh shifted per override.

Lemma A.5.6 (IO–RTP crossover threshold)**.** *If $\Delta W_0 > 0$ and $\lambda > 0$, the override rate at which IO with overrides equals RTP is*

$$\bar{\beta}^{\star} \;=\; \frac{\Delta W_0}{\lambda}. \tag{19}$$

Proof. Equate Eq. (17) to $W^{\text{RTP}} = W^{\text{IO}} - \Delta W_0$ and solve for $\bar{\beta}$.

Proposition A.5.3 (Welfare ordering with overrides)**.** *If $\bar{\beta} < \bar{\beta}^{\star}$, then $W^{\text{IO+O}} > W^{\text{RTP}}$; if $\bar{\beta} > \bar{\beta}^{\star}$, RTP can dominate.*

Proof. From Eq. (17) and Lemma A.5.6, the sign of $W^{\text{IO+O}} - W^{\text{RTP}} = \Delta W_0 - \bar{\beta}\lambda$ is determined by $\bar{\beta}$ relative to $\bar{\beta}^{\star}$.

A.5.5 Alignment with our field experiment

The experiment (IO vs. alternatives) identifies (i) peak-to-off-peak reallocation magnitudes (Lemma A.5.1); (ii) override frequencies $\bar{\beta}$ and their timing relative to peak/off-peak (for $\hat{\lambda}$ via Eq. (18)); and (iii) the baseline welfare gap ΔW_0 via cost/benefit accounting under no overrides. These map directly into the decision rule induced by Lemma A.5.6. Results on bill-variance are acknowledged but not tested (footnote above).

Decision rule (calibrated simulations). We simulate crossover thresholds under UK prices (off-peak £0.07/kWh, peak £0.27/kWh; spread $\Delta\tilde{c}_{\text{eff}}$=£0.20/kWh). Per-override loss uses a point estimate and bounds:

$$\hat{\lambda} \;=\; p_{\text{peak}}\Delta\tilde{c}_{\text{eff}}q^{O} \quad \text{with} \quad (p_{\text{peak}}, q^{O}) = (0.7, 6) \;\Rightarrow\; \hat{\lambda} = £0.84,$$

$$\lambda^{\text{LB}}:\ (0.5, 4) \Rightarrow £0.40, \qquad \lambda^{\text{UB}}:\ (0.9, 8) \Rightarrow £1.44.$$

The IO–RTP crossover is $\bar{\beta}^{\star} = \Delta W_0/\lambda$.

Table 14: Crossover override rate $\bar{\beta}^{\star} = \Delta W_0/\lambda$ with different levels of inattention

		λ (£/override)			$\bar{\beta}^{\star}$ (events/EV–day)		
a (£/EV–day)	ΔW_0 (£/EV–day)	λ^{LB}	$\hat{\lambda}$	λ^{UB}	$\bar{\beta}^{\star}(\lambda^{\text{LB}})$	$\bar{\beta}^{\star}(\hat{\lambda})$	$\bar{\beta}^{\star}(\lambda^{\text{UB}})$
0.00	0.074	0.40	0.84	1.44	0.185	0.088	0.051
0.10	0.079	0.40	0.84	1.44	0.198	0.094	0.055
0.20	**0.083**	**0.40**	**0.84**	**1.44**	**0.208**	**0.099**	**0.058**
0.30	0.086	0.40	0.84	1.44	0.215	0.102	0.060

Columns 1–2 vary the RTP attention cost a and the implied no-override welfare gap ΔW_0.
Columns 3–5 report the per-override loss calibrations: conservative lower bound λ^{LB}, point estimate $\hat{\lambda}$, and conservative upper bound λ^{UB}.
Columns 6–8 compute $\bar{\beta}^{\star} = \Delta W_0/\lambda$. Larger λ (more peak kWh moved per override) $\Rightarrow$ smaller crossover. Values rounded.

Decision rule (using the baseline row). Let $\bar{\beta}_{\text{exp}}$ be the experimental average override rate. Using Table 14, a=0.20 and the point estimate $\hat{\lambda}$ imply $\bar{\beta}^{\star} \approx 0.099$. Thus: if $\bar{\beta}_{\text{exp}} < 0.058$ (the λ^{UB} column), IO dominates RTP even under conservative losses; if $\bar{\beta}_{\text{exp}} < 0.099$, IO dominates at the point estimate; and if $\bar{\beta}_{\text{exp}} \in [0.099, 0.208)$, the ranking depends on whether the true λ is closer to the optimistic lower bound or the point estimate.

Table 15: Crossover with different price elasticities and with point-estimate $\lambda = £0.84$

Elasticity ε	ΔW_0 (£/EV–day)	λ (£/override)	$\bar{\beta}^\star$ (events/EV–day)
−0.20	0.095	0.84	0.113
−0.50	0.083	0.84	0.099
−0.80	0.070	0.84	0.083

$\Delta W_0(\varepsilon)$ is locally linearized around $|\varepsilon_0|$=0.5 with endpoints chosen to match the range $|\varepsilon| \in [0.2, 0.8]$.
Higher $|\varepsilon|$ (more RTP responsiveness) lowers ΔW_0 and hence lowers the crossover $\bar{\beta}^\star$.

Crossover thresholds with partial AI control. The preceding analysis assumed that the entire EV load was subject to AI scheduling. If only a fraction α of household load is controllable (e.g. $\alpha = 0.2$ so that just 20% of usage is managed by IO while the remaining 80% is on flat pricing), then the effective welfare gap shrinks to $\Delta W_0^{(\alpha)} = \alpha \Delta W_0$. Because the override loss λ is unchanged, the crossover threshold becomes

$$\bar{\beta}^\star(\alpha) = \frac{\Delta W_0^{(\alpha)}}{\lambda} = \alpha \cdot \frac{\Delta W_0}{\lambda}.$$

Thus, reducing α proportionally lowers the override cutoffs: with only 20% controllable load, crossovers fall to one-fifth of their baseline values.

Simulation results. Figure 22 plots the implied crossover rates $\bar{\beta}^\star$ for $\alpha = 0.2$ across a sweep of RTP elasticities ($\varepsilon \in [0, -0.5]$). Three curves show the lower-bound, midpoint, and upper-bound calibrations of per-override loss λ. The midline ($\hat{\lambda} = 0.84$) implies that with $\varepsilon = -0.2$, the crossover is only $\bar{\beta}^\star \approx 0.04$ overrides per EV–day (i.e. roughly one override every 25 days). When λ is large (upper bound), thresholds fall even further. Hence, with partial AI control, IO remains dominant only if override frequencies are kept extremely low.

Limitations. Our theoretical framework abstracts from several important considerations. First, the welfare thresholds depend on calibrated parameters for attention costs (a_i), hassle costs (ϕ_i), and risk-aversion penalties (γ_i) that are not directly observed in our experiment; while we draw on the literature and perform sensitivity analysis, these inputs remain assumption-driven.

Second, we model overrides as independent events with constant per-override wel-

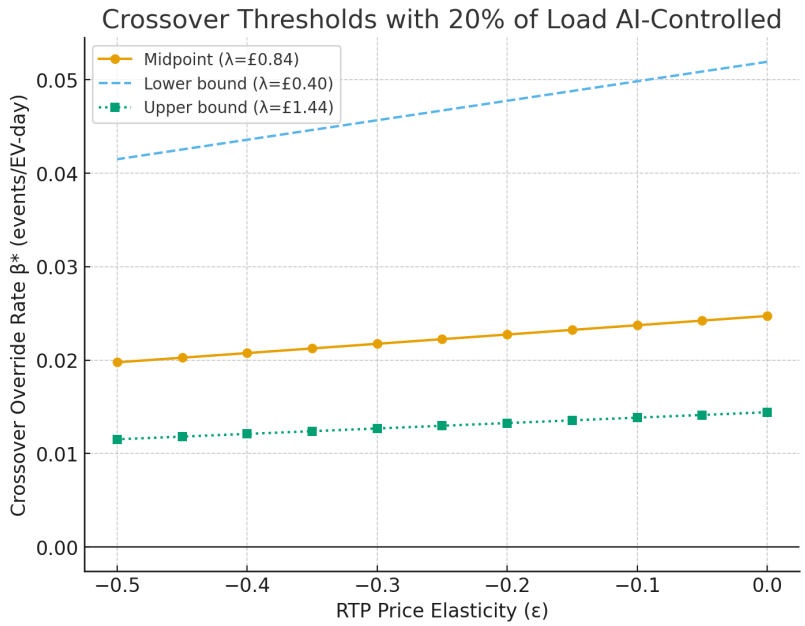


Figure 22: Crossover override thresholds

Notes: The figure plots the override crossover rate $\bar{\beta}^{\star} = \Delta W_0^{(\alpha)}/\lambda$ across a range of assumed RTP elasticities ($\varepsilon \in [0, -0.5]$) when only 20% of household load is managed by IO. Each line corresponds to a different calibration of the per-override loss λ: a conservative lower bound, a midpoint point estimate, and a conservative upper bound. Larger λ values shrink the crossover, meaning IO only dominates if override frequencies are very low. Note that even if RTP users are perfectly inelastic ($\varepsilon = 0$), IO with extremely high override rates (e.g. 99%) cannot outperform RTP under any of the calibrated λ values. At $\varepsilon = 0$, the crossover cutoff is only $\bar{\beta}^{\star} \approx 0.124$ for the full-load case ($\alpha = 1$) and $\bar{\beta}^{\star} \approx 0.025$ for the partial-load case ($\alpha = 0.2$).

fare loss λ, whereas in reality override behavior and costs may vary systematically across hours, days, and households.

Finally, our analysis is partial-equilibrium: prices $\{\tilde{c}_h\}$ are taken as given and households are price-takers, omitting possible feedback effects if IO Go or RTP adoption is widespread. These simplifications are deliberate, allowing a tractable link between model predictions and experimental data, but they should be borne in mind when interpreting the welfare rankings.