



COLUMBIA | SIPA

Center for Environmental Economics and Policy

CEEP Working Paper Series
Working Paper Number 43

October 2026

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

Robert D. Metcalfe, Andrew R. Schein, Cohen R. Simpson, and Yixin Sun

<https://ceep.columbia.edu/sites/default/files/content/papers/n43.pdf>

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

Robert D. Metcalfe

Andrew R. Schein

Cohen R. Simpson

Yixin Sun

Columbia University

Centre for Net Zero

University College London

Centre for Net Zero

Centre for Net Zero & NBER

October 2026

Abstract

One of the promising opportunities offered by AI to support the decarbonization of electricity grids is to align demand with low-carbon supply, yet there is little causal evidence on household acceptance of AI managed control, and its impact on their consumption. We evaluated the effects of the world's largest AI managed EV charging tariff (a retail electricity pricing plan) using a large-scale natural field experiment. This tariff is designed to optimize EV charging by dynamically controlling charging in response to real-time wholesale electricity prices. We randomized financial incentives to encourage enrollment onto the tariff and observed consumption for 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 lower-cost and lower-carbon-intensity periods. The tariff generated substantial consumer savings, while demonstrating potential to lower producer costs, energy system costs, and carbon emissions through significant load shifting. Overrides of the AI algorithm were low, suggesting that this tariff was likely more efficient than a real-time-pricing tariff without AI, given our theoretical framework. We also showed that experimental estimates differed meaningfully from those obtained via non-randomized difference-in-differences analysis, due to differences in the potential outcomes of the samples in the two evaluation strategies, although we can reconcile the estimates somewhat with observables. Our findings highlight the potential for scalable managed charging and its substantial welfare gains for the electricity system and society.

*We thank Lucy Bourke and Joe Grainger, from Centre for Net Zero (CNZ), for their critical support with logistics and background research. We thank Jaye Cribb, from Octopus Energy Limited (OEL), for devising the strategy used to detect ownership of an electric vehicle based on electricity usage, and Charlotte Avery, from CNZ, for further refining and implementing that strategy. We thank Francesca Corby, Geraldine De Boisse, and Katrin Weingartner, from OEL, for their comments and collaboration regarding study design and implementation, in particular construction of the email-based treatments used for customer experimentation. We thank George Doucy, Ben Ferris, Lilla Hoddosy, Ryan Jenkinson, Will Reeves, Richard Sephton, and Henry Tagg, all from OEL, for their guidance around operational aspects of the EV charging algorithms and associated data collection. We thank Andy Hackett, from CNZ, for specific comments on the text of the paper itself. We thank Navya Kumar for excellent research assistance. We thank Omar Asensio, Susanna Berkouwer, Louise Bernard, Daniel Björkegren, David Byrne, Gus Chadney, Sheng Chai, Sylvain Chassang, Daniel Lopez Garcia, Charlie Hutchinson, Alex Imas, Maria Jacob, Seung Min Kim, Justin Kirkpatrick, Alex Schoch, Ben Sprung-Keyser, Phil Steele, Chad Syverson, Arthur van Benthem, Casey Wichman, Izzy Woolgar, Lucy Yu, and seminar participants at ASU, TREE, Wharton, Gothenburg, Copenhagen Business School, Manchester Workshop on the Economics of Energy Transition, and Harvard's Workshop on the Economics and Policy of Electric Transportation Charging Infrastructure for their valuable insights and comments. This work would not have been possible without financial support from the Abdul Latif Jameel Poverty Action Lab's (J-PAL) King Climate Action Initiative (K-CAI). 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. The four authors made final decisions regarding experimental design, data analysis, and writing; all mistakes herein are ours alone. AEARCTR-0013037 and USC IRB: UP-23-00733.

Game theory has been justly criticized for its "hyper-rational" view of human behavior. However, such hyper-rationality may actually be an appropriate model for software agents: presumably software agents have much better computational powers than human beings. The whole framework of game theory and mechanism design may well find its most exciting and practical application with computerized agents rather than human agents...

-Hal Varian, 1995

1 Introduction

The textbook homo economicus is an optimizer who responds instantaneously to prices and never suffers from inattention or cognitive overload. For decades, that vision sat uneasily beside the empirical record that humans suffer from a host of cognitive biases. Thirty years ago, Varian (1995) suggested a resolution: hyper-rationality may be the wrong model for humans but the right one for AI. The software agents he imagined are now becoming real, and with them the potential to dissolve this tension.

Electricity markets are a case in point. In theory, the most efficient approach to energy demand management is real-time pricing (RTP), which directly exposes consumers to the marginal cost of electricity (Borenstein, 2005b; Hinchberger et al., 2026). However, a large body of evidence suggests that consumers are reluctant to engage with day-ahead or real-time prices, due to the cognitive burden and effort required to optimize their usage, along with the price uncertainty they must constantly monitor (Harding and Sexton, 2017; Fabra et al., 2021). What changes this calculus is AI: software that optimize on behalf of consumers who will not, or cannot, do it themselves.

As more households and businesses are electrifying heating, transport, and storage, there is greater societal value for managed electricity demand as opposed to unmanaged demand. EVs are a clear example; flexible, mobile demand loads with storage capacity, EVs can absorb or supply power to help balance renewables in real time. Yet optimal managed charging requires solving complex, real-time problems of cost minimization across dozens of markets, congestion avoidance, and user preferences, precisely the kinds of tasks AI excels at (Rigas et al., 2014; Kaack et al., 2022; Yaghoubi et al., 2024).¹ With EV sales now exceeding 20% of new cars globally (International Energy Agency, 2025) and

¹There are several wholesale markets (day-ahead, intra-day, and last hour—all across three international markets), auxiliary markets, capacity markets, and balancing mechanism markets, in addition to local distribution network markets. For the former markets, our own analysis suggests that the R-squared between day ahead and settlement prices is around 0.6, suggesting that there is a lot of potential value to be had with an algorithm trading in real time to shift demand around to line up with these markets.

projected to drive 15% of future energy demand growth (International Energy Agency, 2024), unmanaged charging will cause two problems. The first is the large increase in daily peak electricity demand straining national infrastructure, and the second is that low-voltage distribution networks may struggle to cope with an increase in simultaneous EV charging (Li and Jenn, 2024; Bailey, Brown, Shaffer and Wolak, 2025; Bernard et al., 2025). Managed charging could instead align electricity demand with low-cost, low-emission periods, reducing national and local-level system stresses.

Partnering with Octopus Energy Limited, at the time Britain’s largest electricity supplier, we investigated the adoption and impact of a popular AI managed EV charging tariff, named "Intelligent Octopus Go" (IO Go).² This tariff is, to the best of our knowledge, the world’s largest managed EV charging tariff, currently serving approximately 300,000 customers across the UK and expanding into the US, Germany, France, Italy, and Spain. It manages over 2GW of power in Great Britain alone (Green, 2025). We estimate that over 17% of all British EV owners used the tariff at the time of our analysis.

The IO Go tariff represents a partial, centrally coordinated implementation of RTP. The tariff combines time-of-use (ToU) pricing with managed charging: customers face a lower off-peak rate (approximately 60% discount) overnight. In return, the retailer controls when the vehicle charges, shifting load away from high-cost periods using real-time wholesale prices as a key input, and many other ancillary market prices, all subject to customer-set constraints (a “ready-by” deadline and target state of charge).³

We used a large-scale randomized encouragement design to estimate the causal effect of financial incentives on adoption of the managed charging tariff, and to use this variation to instrument for the impact of managed charging on electricity consumption. 13,233 suspected EV owners⁴ across Great Britain were randomly assigned to one of five groups. The first treatment group received an email only. Three other treatment groups received an email with a sign-up incentive of £5/month, £50/month, or £50/month con-

²Note that some researchers use "smart charging" (Lagomarsino et al., 2022; Kubli, 2022), and others "utility controlled charging" — which Axsen et al. (2017) describe as a "utility controlling the timing or direction of electricity flowing from the grid to a [PEV's] battery," possibly implemented by a third party and oriented toward cost savings rather than environmental benefits (see also Bailey and Axsen, 2015). By contrast, time-of-use tariffs ("stepwise tariffs" from Daneshzand et al. (2023)) are managed charging controlled on-site by the end-user in response to price, not by a supplier (Delmonte et al., 2020). Unless qualified, we use "AI managed", "managed" and "smart" charging to refer only to remote, supplier-controlled arrangements implemented via Wi-Fi. Within this family, we call the intervention AI managed because the supplier's schedules are not fixed windows or deterministic rules but are generated by machine-learning forecasting and real-time optimization over prices, loads, weather, and grid conditions.

³AI optimization is performed by Kraken Technologies, a company at the time of our experiment owned by Octopus Energy Group. The system ingests real-time data on EVs, household loads, weather forecasts, wholesale prices, and grid constraints, applying machine learning-based forecasting and linear programming to determine optimal charging schedules across batches of devices in real time.

⁴Ownership was modeled from load profiles consistent with home chargers, in conjunction with domain experts at Octopus Energy Limited.

ditional on not overriding the managed charging, with all payments lasting three months. A pure control group received no encouragement or contact. We define our trial as a natural field experiment as it meets the major external validity assumptions in List (2026).

In addition to the experimental analysis, we conducted a supplementary pre-specified 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, exploiting variation in adoption timing to use not-yet-adopted customers as the comparison group. This approach provides complementary evidence on the effects of managed charging in a real-world, self-selected yet scaled-up setting.

The intervention we studied (i.e., the tariff) is a bundled treatment that combines ToU pricing, automated scheduling, and real-time algorithmic optimization across dozens of markets. We provide several empirical strategies to separately identify the contribution of each component. As electricity systems grow more complex, with greater renewables-driven variability in wholesale prices, the proliferation of distributed assets like EVs, and increased participation of those assets across multiple markets, purely deterministic or rule-based scheduling (like classic ToU rates/times) may become less effective. In such settings, more advanced AI architectures for coordination and adaptive control will become increasingly valuable (Si et al., 2025; Wang et al., 2025). At the same time, understanding consumer acceptance of managed systems remains critical, as there is far less empirical evidence on how households respond to AI-driven automation, as opposed to more familiar ToU or deterministic scheduling systems.

1.1 Primary findings

First, using randomized encouragement treatment assignment as an instrument for tariff adoption, we estimated that managed charging reduced household peak-period electricity consumption by 0.581 kWh per hour, a 42% decline relative to the control group mean, with all of that electricity load shifted to overnight off-peak hours. There was no overall change in total electricity use, indicating that the program induced temporal load shifting rather than increased consumption. This 42% peak reduction was large relative to most estimates from the dynamic-pricing literature (Wolak, 2011; Fowle et al., 2021; Ito et al., 2023).⁵ We did not find treatment effects varied by the household area

⁵Our estimate is close to the upper end of recent EV-specific framed experimental estimates: Bailey, Brown, Shaffer and Wolak (2025) reported a 0.20 kWh (49%) reduction in peak EV charging, while La Nauze et al. (2024) reported

income, the value of their property, or the baseline electricity use of the home, suggesting some stability in treatment effects.

Second, we observed strong retention and widespread acceptance of managed charging among trial participants. Trial participants who enrolled in the managed charging tariff 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 managed charging schedule: over half never overrode the managed charging schedule, on any given day there is a 1% likelihood of overriding, and only 2.3% of total electricity consumption occurred via overrides. These patterns suggest that, once adopted, managed charging integrates smoothly into daily routines with minimal disruption, underscoring its potential for long-term grid flexibility. Additionally, these overrides tended to be customer-specific rather than correlated across space or time, indicating that they are unlikely to generate coincident demand spikes that would strain the grid. We found similar charging patterns from our experimental sample and the whole population of EV households on the tariff outside of the experiment, both in the UK and in other countries where this tariff operates, including Spain, Germany, and the U.S. (Texas).

Third, beyond the ToU price structure, managed charging by algorithms delivered value through several channels. It smoothed the shadow peaks induced by ToU pricing, spreading charging across the off-peak window rather than concentrating it at the price change. When EVs were plugged in and under automation, charging was effectively governed by the wholesale price rather than the ToU rate, so the tariff effective price signal tracked wholesale prices substantially more closely than the static ToU rate did. This translated into consumption: within each ToU band, customers' consumption was more responsive to wholesale price variation than was the consumption of customers on "Octopus Go", a static time-of-use (ToU) tariff, consistent with the managed charging tariff's algorithm shifting load in response to price information households never observed. Managed charging also enabled participation in real-time grid-balancing markets — the Capacity Market, Balancing Mechanism, and local flexibility services — creating value a static ToU tariff without AI cannot reach.

Fourth, a DiD analysis of the tariff adopters outside the trial (i.e., among the 17%

reductions of 0.05–0.09 kWh (17%–27%) during peak hours. Those studies measured EV-specific charging demand, whereas our outcome is whole-household electricity consumption, so these comparisons are informative but not exact apples-to-apples benchmarks. Table A1 places our estimates in the context of the broader literature on dynamic pricing and EV charging. All studies save ours and Burlig et al. (2026) (see below) are framed field experiments, whereas ours is a natural field experiment.

of EV drivers in the United Kingdom who adopted the tariff before our trial) yielded two findings on selection and cohort variation. First, baseline tariff shapes the treatment effect: voluntary adopters were disproportionately already on time-of-use tariffs before switching, and reweighting the DiD sample to match the experimental group on prior tariff type substantially increased the estimated treatment effect. Second, cohort-specific effects were homogeneous across 2023, indicating that treatment effects did not degrade as the adopter pool broadened. This is consistent with the algorithm, rather than consumer behavior, driving the response. The experimental IV remains our primary causal estimate, as voluntary tariff adoption is endogenous to other behavioral changes we cannot observe.

Fifth, we estimated the consumer, producer, grid, and social welfare benefits of this managed charging tariff. The reallocation of electricity use to periods with substantially lower rates reduced trial participants' bills by £343 per year, an estimated 18% reduction in cost per kWh. When benchmarked against the retailer's standard flat tariff, the estimated savings rise to roughly £650 per year.⁶ The electricity retailer experienced similar savings in procurement costs (which include wholesale power prices and non-energy charges such as transmission and distribution fees), implying near-complete pass-through of savings to participants, at least during our study period. Using 2050 grid projections, we further estimated that if all EV households adopted this tariff, the resulting load shifting would flatten the load-duration curve, lowering peaks and raising off-peak demand, thereby improving efficiency and reducing the need to build at least five to six large nuclear power stations to meet future electricity demand.

In terms of societal welfare, the impacts of the tariff on CO₂e emissions were substantial.⁷ The resource cost per tonne of CO₂e abated was extremely low (negative) for consumers. We estimated a resource cost of -£946 per tonne, an order of magnitude lower than that of the next-best technology (Gillingham and Stock, 2018; Gosnell et al., 2020; Hahn et al., 2024). We also estimated whether governments should be subsidizing such tariffs, using the marginal value of public funds (MVPF) framework (Hahn et al., 2024). When the numerator includes only the subsidy transfer and carbon benefits, the 10-year MVPF is 0.991. We also considered avoided grid costs, since the tariff shifts consumption to off-peak periods with lower marginal generation costs and thereby defers investment

⁶The counterfactual for our £343 estimate is the average bill of the control group during the experimental period, which includes many tariffs. The £650 figure uses our estimated consumption treatment effects relative to the bill under the standard flat tariff, which most customers in our trial sample were on prior to the experiment.

⁷Electricity-related emissions are generally lowest overnight, when wind and nuclear generation dominate, and highest during the evening peak, when falling solar output and rising demand bring more gas generation online.

in generation, transmission, and distribution infrastructure. If these benefits are not fully captured by private market participants, they may constitute a positive externality. Assuming that 100% of these grid benefits are passed through to consumers (in the UK, as in many electricity systems, consumers ultimately bear capacity costs through their electricity bills), the 10-year MVPF rises to 2.8.

1.2 Contribution to the existing literature

Our study contributes to three strands of literature: (1) EV charging, (2) tariff switching and dynamic pricing, and (3) the economics of AI.

EV Charging: The closest to our work are two framed field experiments of managed and flexible EV charging. In particular, Bailey, Brown, Shaffer and Wolak (2025) and La Nauze et al. (2024) showed that financial incentives and time-of-use pricing alone can induce substantial flexibility in EV charging demand. At the same time, Bailey, Brown, Myers, Shaffer and Wolak (2025) highlighted an important limit of time-of-use pricing: by rewarding the same low-price hours, it can generate shadow peaks from simultaneous charging. Their results point to a natural role for centrally coordinated managed charging, which can preserve flexibility while avoiding synchronized demand.⁸

The existing evidence already demonstrates the promise of managed charging, but it comes from smaller scale samples of customers who select into the experiment.⁹ We contribute the first natural field experiment of a managed EV tariff at scale, studying the world’s most popular AI managed EV charging tariff. The AI component is of independent interest because it speaks to a broader question: whether households accept automation that responds to grid conditions on their behalf. That question matters beyond

⁸We also relate to several other empirical studies of EV charging at home, at work, and in public charging contexts (Burkhardt et al., 2021; Garg et al., 2024; Bernard et al., 2025; Garnache et al., 2025). Methodologically, Burlig et al. (2026) is closest to ours: they also evaluated managed charging in a natural field experiment, and likewise found low take-up of a marginal offer. The settings differed in some respects that are important for interpretation. The tariff we studied had already reached a majority of EV drivers with the retailer, so our randomized encouragement design estimated take-up among a residual, and potentially harder-to-convert, population rather than from the near-zero base in Burlig et al. (2026). Design differences also plausibly explain that gap – their offer required a separate app and session-by-session opt-in and left the per-kWh price unchanged, whereas ours sat within the retailer’s main app and lowered the off-peak price. Perhaps most important, both our study and Burlig et al. (2026) examined the impact of encouragement in the form of a single email with a switching incentive, whereas in reality managed charging tariffs such as the one we studied in this work are actively and repeatedly marketed across multiple channels. With all that in mind, we caution that weak responsiveness to a single marginal offer need not imply that managed charging cannot reach high adoption.

⁹There are also dozens of small managed-charging pilots in the field (SEPA, 2024), but most lack credible counterfactuals or sufficient scale. In the United States alone, over \$100 million has been allocated 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 empirical experiments.

EV charging. As electricity systems rely more heavily on variable renewable generation, they will increasingly require demand that can respond to prices and system conditions without requiring constant consumer attention.¹⁰

Tariff switching and dynamic pricing: Our paper relates closely to the tariff design and real-time pricing (RTP) literature. The tariff we study is a hybrid: a time-of-use retail tariff combined with real-time-price-based automation for EV charging. Much of the RTP literature emphasizes the potential efficiency gains from granular pricing (Borenstein, 2005a; Nicolson et al., 2018; Hinchberger et al., 2026), but it also highlights the practical limits imposed by consumer attention, price uncertainty, and low adoption (Harding and Sexton, 2017; Fabra et al., 2021). More broadly, there is substantial inertia and inattention in retail energy markets (Hortaçsu et al., 2017; Byrne et al., 2022; Gravert, 2024; Garcia-Osipenko et al., 2025). Our findings show that financial incentives and set-and-forget automation can relax behavioral frictions, while still achieving substantial adoption, strong retention, and limited overrides.

This result is encouraging for the broader feasibility of real-time pricing. The closest paper on this margin is Blonz et al. (2025), who showed that automation tied to a time-of-use price schedule materially increased demand response by adjusting thermostat set-points, and, as in our setting, found low override rates. We build on this insight by studying a tariff that goes beyond responding to a fixed TOU schedule; the algorithm we study adjusts charging in response to underlying wholesale market conditions and participates in ancillary service markets and real-time grid balancing operations.¹¹ As such, it is a highly efficient tariff.

Economics of AI: More broadly, our paper relates to emerging work on AI agents as intermediaries. Generative search can lower the cost of query formation by allowing users to express needs in natural language (Zheng et al., 2025), and evidence from Sierra Leone shows that AI can also reduce the practical cost of information access relative to conventional web search (Björkegren et al., 2025). On the theoretical side, Liang (2026) modeled AI "clones" as representations that expand search capacity but may imperfectly encode users' preferences. Rusak et al. (2025) further showed that LLM-based preference

¹⁰Our findings also inform the broader EV policy literature on infrastructure, adoption, and system costs (Zhou and Li, 2018; Archsmith et al., 2022; Powell et al., 2022; Rapson and Muehlegger, 2023b,a; Heid et al., n.d.; Turk et al., 2024; Asensio et al., 2025; Dorsey et al., 2025; Gillingham et al., 2025). By showing that AI scheduling lowered consumer bills, reduced grid system peaks, and shifted electricity load to low-emission periods, we demonstrated a key pathway through which managed charging can cut grid investment needs and raise welfare.

¹¹Another related paper in a very different setting is Khanna et al. (2024), who study automated appliance control in India. That setting differs from ours in two important respects: take-up was low, and the flexible load is quite different from EV charging, which makes direct comparison difficult. Among those who did participate, however, they also found that automation combined with incentives reduced household electricity demand.

elicitation can make richer market designs feasible when human preference reporting is otherwise too costly.

In our setting, AI does not primarily improve search across products or assist with one-off choices. Instead, it intermediates households’ exposure to high-frequency wholesale electricity conditions. This matters because the consequences are not only private. By reducing the attention and effort required to respond to dynamic conditions, AI managed charging can also generate system-level benefits through load shifting, improved alignment of consumption with grid needs, and emissions reductions.

Although some 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 effort cost remains unaffected by the AI’s optimization.

The paper is structured as follows: Section 2 provides background on the intervention (i.e., the AI managed EV tariff). Section 3 provides an overview of the natural field experiment, including the sampling, the randomized encouragement, the randomization procedure, and available data. Section 4 develops the demand estimates from the field experiment. Section 5 presents descriptive evidence on the value added by managed charging beyond the time-of-use price signal. Section 6 complements Section 4 with a difference-in-difference analysis of the intervention, and attempts to reconcile with the experimental estimates. Section 7 presents the welfare estimates of the intervention and the randomized encouragement. Finally, Section 8 concludes.

2 Background

This research was conducted in partnership with Octopus Energy Limited (OEL). At the time of our experiment, OEL was the largest residential electricity retailer in the United Kingdom, supplying approximately 7.7 million households, or roughly one quarter of all domestic energy accounts. The experiment evaluates Intelligent Octopus Go (IO

Go), a residential electricity tariff that combines time-of-use pricing with AI managed EV charging, and which had accumulated approximately 286,000 enrolled customers as of the time of writing. We describe the tariff design in detail below.

2.1 The intervention: a managed charging tariff


IO Go is a residential electricity tariff that combines time-of-use pricing with AI management of automated EV charging. Through the retailer’s app, customers set a target battery level (i.e., state of charge) and departure time (i.e., ready by time); the retailer then schedules charging to meet those targets (Figure 1a). In exchange, the tariff provides a favorable electricity rate of £0.07/kWh during a fixed six-hour off-peak window (23:30–05:30).¹² 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, or sustained override shown in Figure 1c). The overriding is easy and costless – customers open the app, go to the tab in Figure 1a, and press the “Bump Charge” button. When customers charge outside of the schedule through overriding, it is billed at the higher rate, so there is a financial cost to overriding. The applicable tariff schedule for 2024 is shown in Figure 2. Generally, the higher rate was very similar to the utility’s standard flat rate.

Beyond its engineering implementation, the tariff design embodies a core principle from mechanism design. It functions as a large-scale restricted-revelation mechanism. Each customer specifies two primitives — a ready-by time and a desired state of charge — which together summarize individual flexibility (Figure 1a). With these constraints, the algorithm allocates charging across time to minimize system costs while meeting user needs. No detailed reporting of preferences, utilities, or price sensitivities is required. In mechanism-design terms, it is a simple direct mechanism defined over a coarse message space, yet it achieves near-efficient outcomes because the AI optimizes using observed demand and wholesale price data. The design echoes the notion of approximate competitive equilibrium from Budish (2011), although it is possible that customers may exaggerate their preferences (i.e., setting an earlier ready-by-time than necessary, or higher state of

¹²During the time of our natural field experiment, the retailer’s standard tariff [Flexible Octopus](#) charged a flat rate of approximately £0.22–£0.27 per kWh throughout the day, depending on the month, as shown in Figure 2, with small variations by region of country.

Figure 1: Managed Charging Tariff Controls via Mobile App




Notes: Example screens from the managed charging tariff mobile application. Panel (a) shows the AI-generated charging schedule, while panels (b) and (c) depict temporary and sustained user overrides of managed charging.

charge) given there are no penalties for doing so.

The tariff uses AI to generate charging schedules based on prices in many markets. The first is 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 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 consumer-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-

Figure 2: Tariff Rates



Notes: This figure shows the tariff rates for customers on the managed charging tariff, 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 standard non-time-varying tariff (black) from the retailer, the main counterfactual tariff.


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. A diagram outlining the structure of the AI scheduling framework is shown in Figure 3.¹³

2.2 Redistributing the frictions of dynamic pricing

Most existing literature on time-varying electricity pricing concerns retail products offering either time-of-use (ToU) pricing or real-time pricing (RTP). The tariff we study straddles the two: the household faces a simple ToU bill with fixed peak and off-peak rates, while the supplier, during plug-in windows, schedules charging in response to wholesale prices, ancillary service opportunities, and grid balancing signals. The con-

¹³The optimization engine integrates two complementary computational approaches: classical linear programming and machine learning-based forecasting (the AI currently uses no large language models to predict or optimize). 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.

Figure 3: AI Managed Scheduling Framework



Notes: In contrast to a static time-of-use tariff or a deterministic supplier-controlled schedule (e.g., a simple “charge only overnight” rule), the scheduling depicted here is produced by machine-learning forecasting of wholesale prices, household and EV loads, weather, and grid conditions, combined with real-time linear-programming optimization that resolves charging schedules across batches of devices as those forecasts update. The optimization is implemented by Kraken Technologies, which ingests both demand-side inputs (customer requirements, plug-in behavior, and local load predictions) and supply-side inputs provided by Octopus Energy. The supply side inputs can be complex. Distribution Network Operators procure local flexibility to manage constraints and reliability. Sustain services address long-term capacity needs, Secure services manage predictable network peaks, Dynamic services respond to real-time constraints, and Restore services support recovery after faults. At the national level, NESO procures ancillary and balancing services to maintain system stability. Frequency response maintains 50 Hz operation through Dynamic Containment, Dynamic Moderation, and Dynamic Regulation. Reserve services provide backup capacity, split into Quick and Slow Reserve, while the Balancing Mechanism matches supply and demand in real time. Technical services include Black Start, inertia, reactive power, short-circuit level, and constraint management. The Capacity Market provides availability payments for system stress events via T-1 (one year ahead) and T-4 (four years ahead) auctions. Wholesale markets provide energy price signals across various timeframes. Day-ahead markets allow trading for the following day. Intraday markets enable near-real-time adjustments, while imbalance pricing penalizes deviations from contracted positions. Peak-related levies – including the Capacity Market levy, Distribution Use of System (DUoS), and Transmission Network Use of System (TNUoS) charges – further incentivize load shifting and peak reduction. Not all of these markets and signals are directly acted on; they are noted to illustrate the range of supply-side inputs considered in optimizing demand. The retailer is responsible for forecasting wholesale and ancillary market prices and other market conditions, while Kraken integrates these inputs to generate charging schedules subject to user constraints and technical limits.

sumer sees ToU; the supplier acts on RTP. The household can always override at a hassle cost, reverting to immediate charging under the ToU schedule.

This hybrid redistributes the frictions of dynamic pricing. Pure RTP is allocatively efficient in principle but loads attention costs and bill-volatility risk onto the household — frictions that, under rational inattention or bounded rationality, prevent consumers from responding optimally to high-frequency price signals. The tariff shifts both burdens to the supplier: the household sees only a ToU bill, while the algorithm absorbs the wholesale-price response. In return, the supplier inherits a new behavioral margin: the household’s option to override.

Each override recovers the household’s immediate preference but displaces load toward hours the algorithm had intentionally avoided, partially undoing the welfare gain. Whether the managed charging tariff welfare-dominates RTP therefore relies on the override rate. Rare overrides preserve the automation gains; frequent overrides dissipate them. Appendix H formalizes this intuition and characterizes the welfare ordering. And as we discuss in our results, we are able to directly observe the override rate, and find it

to be very low (Section 4.3).

In addition, the British electricity system, like many others, includes multiple ancillary and flexibility markets, with overlapping mechanisms for valuing flexibility. Market participants respond not only to wholesale price signals but also to a range of network, balancing, and capacity markets, including the Capacity Market, Balancing Mechanism, and various frequency response and distribution network services. This complexity means that optimizing against a single real-time or day-ahead wholesale price would fail to capture value available in other markets. The tariff’s automation enables the retailer to allocate charging efficiently across these interacting price signals.

2.3 Adoption of the managed charging tariff

Intelligent Octopus Go has achieved substantial enrollment since its launch in January 2022, with approximately 286,000 customers enrolled as of writing. According to accredited official statistics published by the [Department for Transport](#), there were approximately 1.63 million licensed battery electric cars in the United Kingdom as of September 2025. This implies that over 17% of British EV owners were signed up to the tariff, and over 20% were signed up to one of the retailer’s EV tariffs (i.e., including its non-managed-charging EV tariff option). Among the retailer’s own EV customers, take-up is necessarily higher, but we lack data on their total EV customer base to quantify take-up.

3 Experimental design

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

3.1 Eligibility into the field experiment

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 tariffs – that is, they had no prior engagement with managed tariffs; 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 half-hourly 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),¹⁴ 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 Limited colleagues with subject matter expertise on typical EV charging frequency among customers).¹⁵

In summary, our sampling strategy yielded 13,233 trial participants. In this sample, 86% were on flat-price tariffs. Given that most customers in the UK are on flat tariffs, our trial population is likely to resemble the broader set of prospective adopters of managed charging technologies along this dimension. This contrasts with early adopters of the tariff, who were typically more engaged, more technologically inclined, and potentially more environmentally motivated than the mainstream EV-owning population.¹⁶ In other words, unlike early adopters already enrolled in smart tariffs, we believe our trial participants were more similar to the broader population of potential adopters of managed charging technologies.¹⁷

¹⁴We used the summer for this approach as there was less likely to be heat pumps or any other electric heating blocking the EV charging signal.

¹⁵We are unable to provide formal performance statistics for the EV detection algorithm because we do not have ground-truth data on customers’ low-carbon technology ownership. However, after the first round of encouragements in February 2024, we surveyed 305 customers who received an encouragement but had not signed up for the managed charging tariff. We received 73 responses. Of these respondents, only four reported not owning an EV (5%). This high proportion of true EV owners among those we predicted to have an EV suggests that the detection algorithm likely had a low Type I error rate (i.e., few false positives).

¹⁶According to a recent [report](#), 91% of customers are on flat-rate tariffs in the UK.

¹⁷Given that they had adopted their EV for a few months before we contacted them, we expect response rates would be slightly lower than if we contacted them immediately after buying their EV (Nicolson et al., 2017). However, that would raise difficult identification issues since we would not have a clear baseline time period.

3.2 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 managed charging
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 managed charging
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 managed charging
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 managed charging. Additionally, they paid £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 make overriding even more costly and thus understand trial participants’ willingness to tolerate even less control over their charging schedule.

Each treatment arm received a single encouragement email from the retailer’s marketing team, reproduced in Appendix C.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’ account balances, accruing daily over a 90-day period conditional on maintaining an active managed charging tariff 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. This setup was very natural and in-line with how the retailer administers incentives in other programs.

Note that it was required by our implementing partner that our encouragement emails for the managed charging tariff inform customers that the retailer offered other tariffs that

may have better met their needs. Importantly, the managed charging tariff is not compatible with all chargers or vehicles; for these circumstances, the retailer offers an alternative EV tariff: a ToU tariff offering a fixed off-peak rate for electricity. The major differences from the managed charging tariff are: (1) it offers one fewer hour of cheap overnight rate (00:30–05:30, compared to 23:30–05:30 for the managed charging tariff) (2) its off-peak rate is higher than the managed charging tariff (£0.085/kWh, as compared to £0.07/kWh for the managed charging tariff; for exact rates, see Figure A1b), and (3) the ToU tariff 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 the managed charging tariff, posing a potential channel for exclusion restriction violations. We test and discuss this in Section D.1.

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

1. Increasing the incentive payment for adopting the managed charging tariff will increase actual adoption.
2. Adoption of the tariff will shift electricity consumption from peak (16:30–20:30) to off-peak (23:30–05:30) hours.

These two hypotheses are not from the same family and thus we did not correct for multiple hypothesis testing across them.

3.3 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.¹⁸ 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. This distribution across treatment arms reflected a balance between budget constraints and the desire to offer realistic incentive levels, on the one hand, with the goal of maximizing power (by increasing sign-up to the tariff), on the other. Randomization was performed

¹⁸These included historical electricity usage, tenure as a customer with the retailer, and past engagement with smart tariff onboarding.

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 A2, and Figure A2).¹⁹

We used the randomized encouragement to estimate the intent-to-treat effect (ITT) of receiving the encouragement on take-up of managed charging, and then used the encouragement as an instrument to identify the local average treatment effect (LATE) of managed charging on outcomes. To identify the ITT effect, we satisfied the three standard requirements for a valid randomized assignment mechanism: probabilistic assignment, individualistic assignment, and unconfoundedness.²⁰

In addition to having used an assignment mechanism that ensures unbiased estimation, the study benefited from high observability. We tracked each participant’s electricity consumption before, during, and after tariff adoption. Data loss only occurred when customers left the retailer or stopped using smart meters. Twelve months after encouragement, 97% of participants still provided usable smart-meter readings, with attrition rates balanced between the treatment and control groups. Finally, there was no selection into the sample and no notification of randomization or an experiment to the sample, limiting any Hawthorne or John Henry effects. Thus, we can call our experiment a natural field experiment (Harrison and List, 2004).

3.4 Data

Our analysis drew on high-frequency administrative data from the retailer, with personal identifiers removed. We focused on two primary outcomes for our two hypotheses: (1) take-up of the tariff, measured as a binary indicator for whether a trial participant held an active managed charging tariff 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 $\times$ 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

¹⁹In our pre-analysis plan, we calculated the minimum detectable LATE when using the encouragements as instruments for take-up of the tariff, and performed power calculations separately for each encouragement-specific instrument. We found that to achieve 80% power required detecting a 30–40% reduction in peak-period consumption.

²⁰First, every participant had a positive probability of being assigned to any of the encouragement groups or control (probabilistic assignment). Second, assignment depended only on pre-encouragement covariates and is therefore independent of other participants’ potential outcomes (individualism). Finally, encouragement assignment was independent of potential outcomes (unconfoundedness).

outcomes are observed from January 1, 2024 through March 31, 2025.²¹ All of these decisions were in our pre-analysis plan (AEARCTR-0013037).

In Figure 4, we show the average hourly electricity consumption in the pre-encouragement period (i.e., January 2024), showing an expected evening peak beginning around 16:30, coinciding with the system peak in Great Britain when wholesale prices, network congestion, and emissions intensity are typically highest. Baseline electricity consumption for our trial participants was much higher than that of a random sample of 5,000 other customers with smart meters, but without EVs. Notably, we also observe substantial overnight consumption, consistent with EV charging demand. We believe this is partly driven by plugging in the EV after work in the evening and the default scheduling behavior in common EV chargers, which sometimes pre-set charging to off-peak times.²² 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.

We focus our analysis on off-peak periods (23:30–05:30) and peak periods (16:30–20:30), following our pre-registration. Off-peak periods correspond to the tariff’s hours of cheap overnight rates. Peak hours capture the period of highest intensive domestic electricity consumption (Few et al., 2022). That said, it is worth noting that the definition of “peak” and “off-peak” may change and themselves become more variable by day and season in the coming years.


For experiment participants who adopted the tariff, 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 the retailer’s managed charging algorithm or manually overridden by the user. These data allowed us to reconstruct intended charging preferences, actual charging demand, and deviations from automated control.

We used additional administrative data at the daily level to estimate consumer and

²¹A 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.

²²Since 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.

Figure 4: 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 the managed charging tariff or ToU tariff, (2) those who remained on other tariffs, and (3) a random 20,000 sample of other customers. Shaded regions show the 95% confidence interval. The ToU tariff is an alternative time-of-use EV tariff designed for customers who either had hardware that was incompatible with the managed charging tariff or did not wish to enroll in automation. Green shaded box indicates the managed charging tariff's 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.

supplier benefits. First, to measure benefits to consumers, we used administrative data from the retailer 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), small area measures of relative deprivation across each of the United Kingdom. Areas were ranked from the most deprived area (rank 1) to the least deprived area, based on income, employment, education, health, crime, barriers to housing and services, and the living environment. This index was 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 CO₂e 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.

4 Experimental results

This section presents the main empirical findings from our field experiment. We begin by documenting the impact of enrollment in managed charging on electricity consumption patterns, using our randomized encouragement as an instrumental variable for enrollment. We then examined treatment effect heterogeneity by socioeconomic status and baseline consumption. These analyses followed our pre-analysis plan, but we state where we added or deviated from it, why we did, and their implications for interpretation in Appendix F.²³

We conducted several additional exploratory analyses. First, we used telemetry data to characterize managed charging tariff behavior along three margins: plug-ins, charging, and overrides. Second, we compared responsiveness to prices between managed charging tariff and ToU tariff customers to isolate the value added by automation. Third, we brought in telemetry data from Germany, Spain, and the United States to compare UK behavior against other markets where the managed charging tariff operates.

4.1 Empirical strategy

Adoption of the managed charging tariff is likely endogenous to household characteristics that also drive electricity consumption. Tech-savvy households, heavy EV users, and price-sensitive customers are likely to both self-select into the tariff and consume differently from non-adopters, confounding direct comparisons. We addressed this with randomized encouragement as an instrument for tariff adoption. The resulting estimates identify a Local Average Treatment Effect: the causal impact of adoption on compliers —

²³Our pre-analysis plan also specified take-up rates and intent-to-treat effects of the encouragement on consumption; we report these in Appendix D.

customers who take up a tariff when encouraged.

The first stage instruments tariff adoption with random assignment to any of the four email-based encouragements:

$$D_{iht} = \pi_0 + \pi_1 Z_i + \gamma X_{ih} + \mu_b + \mu_t + \epsilon_{iht} \quad (1)$$

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).

Even simple email contact, without incentives, raised take-up by 3.4 percentage points; £50/month incentives roughly doubled that (Appendix D.1). Importantly, these take-up rates represented a lower bound for two reasons. First, eligibility was inferred from an EV detection algorithm rather than observed directly. Second, enrollment required a supported charger or EV, and we lack data on compatibility in the customer base. The measured effects therefore understate the potential response in a fully compatible, correctly identified sample. We treat these take-up rates as evidence of statistical power for the IV analysis, not as a measure of managed charging’s underlying appeal.

As shown in Appendix D.1, our encouragements increased adoption of both the managed-charging tariff and an alternative EV-oriented static time-of-use (ToU) tariff. To preserve the exclusion restriction, given that our instrument affected both products, we define D_{it} as an indicator for take-up of either the managed charging or ToU tariff.²⁴

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

$$Y_{iht} = \alpha + \beta \hat{D}_{iht} + \gamma X_{ih} + \psi_b + \psi_t + \epsilon_{iht} \quad (2)$$

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 instrumental variables (IV) assumptions: (i) relevance — encouragement must increase adoption, which we confirm with strong first-stage effects already shown in Table A3; (ii) independence — random assignment ensures

²⁴This deviates from our pre-analysis plan, reflecting our uncertainty at the outset of the trial about whether we would face this exclusion restriction violation. However, we find the coefficients are similar whether we include or not the ToU customers into the D_{it} indicator.

encouragement is uncorrelated with unobserved determinants of outcomes, which holds given our assignment mechanism, as discussed in Section 3.3; (iii) exclusion — encouragement should affect electricity use only through tariff adoption, which we preserve by pooling managed charging and ToU 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.²⁵

4.2 Impact of the managed charging tariff on consumption

We found that adoption of an EV tariff significantly shifted consumption from peak to off-peak hours. When we estimate Equation (2) over all hours within the specified peak and off-peak windows, our main specification shows that households’ peak-period usage fell by 42% (0.581 kWh average hourly reduction), while their off-peak usage rose by 50% (0.481 kWh average hourly increase). This is a substantial reallocation of demand rather than an increase in overall consumption. Consistent with this interpretation, Table A10 reports no change in total electricity use. This load shifting pattern is further illustrated in Figure 5, which shows consumption estimates, split by hour-of-day.


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 5 also shows (i) a specification that includes the four encouragement indicators as separate instruments and (ii) 4 additional specifications 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.²⁶

We show the consumption profile of the baseline group – control households who did not adopt either the managed charging tariff or its similar non-managed ToU version – to

²⁵The LATE identifies the effect of tariff adoption 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.

²⁶Note 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 5.

Figure 5: 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 (2)); (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 the tariff's 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 the tariff'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 A8 and A9.

contextualize these treatment effects. These customers' electricity consumption exhibits a pronounced peak beginning around 16:30, consistent with typical residential demand patterns and uncoordinated EV charging behavior (see Figure A6). Using the estimated treatment effects from the main specification in Figure 5, 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 also explored heterogeneity in treatment effects across consumers by interacting the encouragement treatment with three key variables: (1) the Index of Multiple Deprivation (IMD), (2) property value of their home, and (3) baseline electricity consumption. In these specifications, we treat interaction terms (e.g., EV tariff $\times$ baseline covariate) as endogenous and instrument them with the corresponding interaction of the randomized encouragement and the covariate, consistent with standard IV practice. Across all three dimensions, we found limited evidence of heterogeneity in the impact of managed charging on consumption patterns. We discuss these results in Appendix D.3.

Our estimated peak reduction is large relative to most of the dynamic-pricing literature, though that contrast is consistent with EV charging being a particularly flexible form of electricity demand (Wolak, 2011; Fowlie et al., 2021; Ito et al., 2023). Bailey, Brown, Shaffer and Wolak (2025) report that financial incentives reduced peak EV charging by 0.20 kWh, or 49%, while La Nauze et al. (2024) report peak-hour charging reductions of 0.05–0.09 kWh, or 17%–27%. Those studies measure EV-specific charging demand, whereas our outcome is whole-household electricity consumption, so the comparison is informative but not exact. Table A1 places our estimate in the context of the broader literature.

Burlig et al. (2026) is methodologically closest to our setting in studying managed charging through broad customer outreach rather than opt-in recruitment. They show that when managed charging is offered at population scale, enrollment is low and intent-to-treat effects on electricity consumption are correspondingly small. We also found modest take-up in our experimental sample. However, as discussed in Section 2, take-up of this managed charging tariff is substantially higher in the retailer’s broader customer base. By design, our experimental sample targeted customers with especially high inertia.

4.3 Override, plug-in, and charging behavior using telemetry data

For participants who adopted the managed charging tariff, we had telemetry data on their plug-in events and charging sessions, allowing us to examine their behavior more closely. On average, adopting households consumed 1.14 kWh per hour (for their home overall, not just their EV), equivalent to 9,997 kWh per year. Of this total, 23.0% was attributable to EV charging, or around 2,295 kWh annually. If all charging occurred at home and the vehicle averaged about 3–4 miles per kWh, this implies 7,000 – 9,000 miles

driven per year. These figures are consistent with the average annual car mileage in the UK (approximately 7,000 miles; [Department for Transport, 2024](#)), suggesting that the mileage behavior in our sample is comparable to that of primary household vehicles in the broader UK population.


We further 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 consistency in treatment effects across users. Together, these findings suggested that managed charging was generally well integrated into users' routines with minimal disruption.

We begin with evidence of adherence to the automation schedule. By adherence, we mean their likelihood of overriding the AI charging algorithm, a useful proxy for satisfaction with the tariff. We analyzed 2,152 participants' use of the "bump" function, which allowed them to override the default schedule and initiate immediate charging. We found that 54% of users never used the bump feature at all (Figure [A13](#)), and bump events accounted for only 2.3% of total electricity consumption among managed charging tariff users. These infrequent overrides, consistent across treatment groups, suggested that the managed charging schedule was generally accepted and rarely disrupted. This was also consistent with what we saw in Figure [A31](#), where the sustained uptake of the tariff suggests that trial participants generally accepted the automated approach.

Although bump charging was infrequent, its implications for grid planning depend on whether charging is correlated across locations or time. To assess this, we used a fixed-effects regression to decompose the variance in bump charging across customer, temporal, spatial, and temperature components. As shown in Appendix [D.5](#), most of the variation was driven by customer-specific differences, with little contribution from time, location, or temperature. This dominance of customer fixed effects indicated that bump charging was largely idiosyncratic rather than correlated, making it unlikely to generate coincident demand spikes that would strain the grid.

Preferences for charging settings further reinforce this uniformity. 67% of trial participants preferred their vehicle to finish charging between 07:00 and 09:00, and 94% set their desired state-of-charge (SOC) at 80% or higher (Figure [A10a](#)). 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 [A10b](#)).

Figure 6: Hourly Patterns of EV Plug-In and Charging Behavior



Notes: Panel (a) shows, by hour of day, the percentage of electricity consumption that occurred during that hour. This consumption is further divided into charging triggered by bump (charging initiated by users overriding the schedule) versus charging scheduled by the managed charging tariff. Panel (b) plots, by hour of day, the percentage of active users who had their EVs plugged in. A user is considered active during a given hour if that hour falls between their first-ever and last-ever recorded plug-in event. The analysis is based on data from 2,359 managed charging tariff customers.

We also document that the mechanism operates effectively with sparse and infrequently updated private information. The median customer only changes their settings twice after the initial onboarding. The bulk of preference discovery occurs in the first week: one week after adopting the tariff, 54% of customers never updated their settings again (Figure A9). This stability underscores the mechanism's minimal information requirements. Once users report a small number of primitives, the algorithm can implement near-efficient charging schedules without ongoing preference elicitation or strategic interaction.

If customers were to begin charging as soon as they plugged in, typically between 5:00pm 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 managed scheduling shifts the actual charging to off-peak hours, as illustrated in Figure 6. 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 A11, A12 and A14). This behavioral consistency is mirrored in the relatively homogeneous consumption impacts shown in Figure 5.

4.4 Managed charging tariff customers in other countries

A natural external validity question is whether the UK effects generalize to other markets. Octopus Energy Group retailers in seven countries currently offer the managed charging tariff. Four of these — Germany, Spain, the United States (Texas), and the United Kingdom — have a substantial customer base. Appendix D.4 gives the full cross-country analysis; here we summarize the implications for external validity. In addition, we also show the descriptive statistics for early adopters in the UK, as defined by those who took up IO Go in the first six months since it was rolled out.

Cross-country descriptive statistics on telemetry data showed similar behavior across the other countries to the UK. Importantly, Figure A38 plots average hourly EV consumption for managed charging tariff users alongside average household consumption for a random sample of non-EV households in each country. In all four markets, consumption under the managed charging tariff showed a pronounced overnight spike of approximately 0.4 kWh per hour, with the rest of the day flat. Charging concentrated overnight in Germany, Spain, Texas, and the UK alike.

The magnitude lined up with the experimental estimate. In the UK trial, adoption of the tariff shifted approximately 0.5 kWh per hour of consumption from peak to off-peak periods, the 42% peak reduction in Figure 5. The 0.4 kWh overnight spike observed in each of the four countries was within rounding distance of that figure, which implied a comparable volume of flexible load shifted by EV charging in each market. The consumption response identified experimentally in the UK is therefore consistent with the descriptive cross-country evidence.

5 Value of managed charging beyond time-of-use tariff structure

The tariff we studied bundled ToU prices with automated scheduling and real-time AI optimization. As a result, the experiment identifies the effect of the bundle, not the separate contribution of each component. To probe the mechanisms, we present four pieces of descriptive evidence on the added value of AI managed charging, beyond the ToU price incentive. First, managed charging avoids the shadow peaks that ToU pricing induces (Section 5.1). Second, since the algorithm responds to real-time prices when the

EVs are plugged in, it produces a price signal that tracks the wholesale costs more closely than a static TOU rate (Section 5.2). Third, we show that this price signal was effective in shifting consumption towards lower-priced half-hours within ToU periods (Section 5.3). Finally, managed charging enables participation in real-time grid-balancing markets that a static tariff cannot participate in (Section 5.4).

5.1 ToU pricing creates shadow peaks

As documented by Bailey, Brown, Myers, Shaffer and Wolak (2025), ToU pricing shifts EV charging into off-peak hours, but it induces “shadow peaks” from simultaneous charging at the off-peak boundary. We observed this in our sample; restricting the sample to customers on the ToU tariff - the EV tariff with ToU pricing but *without* managed charging, with cheaper overnight pricing from 00:30 to 05:30 - consumption jumped from 1.01 kWh at 23:30 to 2.68 kWh at 00:30, a 165% increase in a single hour (Figure 7).


Customers on the managed charging tariff also saw a jump when their off-peak window opened, rising from 1.37 kWh at 22:30 to 2.09 kWh at 23:30. The jump was smaller, and the profile that followed was flatter. The algorithm distributed charging across the off-peak window and shifted some of it into the afternoon and evening hours, avoiding the sharp coordinated spike observed under ToU-only pricing.²⁷

Note that this was not a clean comparison. The managed charging tariff off-peak window began one hour earlier than the ToU tariff window, mechanically widening the interval over which charging could be spread. In addition, selection into the two tariffs was not random, and our randomized encouragement design is not powered to recover the causal contrast between managed charging tariff and ToU tariff consumption profiles. We therefore treat the comparison as descriptive rather than causal.

The descriptive evidence does suggest that managed charging addresses a known failure of relying on ToU pricing at scale. ToU pricing coordinates households around cheap hours, not around grid conditions, and produces new peaks as a result. Managed scheduling coordinates them around grid conditions directly, attenuating coordination problems, a distinction that grows in importance as EVs become a dominant share of residential load.

²⁷Managed charging tariff customers had slightly higher average consumption overall. Figure A15 shows the same pattern when each hour is plotted as a fraction of total daily consumption - elevated consumption from 19:30 onward and a flatter off-peak profile.

Figure 7: Electricity consumption by hour-of-day on EV tariff



Notes: The figure plots average hourly electricity consumption for customers on the managed charging tariff (purple) and ToU tariff (blue) over the trial period, estimated from 2,131 managed charging tariff customers and 889 ToU tariff customers. Shaded regions denote 95% confidence intervals. The ToU tariff is a time-of-use EV tariff without managed charging, available to customers whose hardware is incompatible with the managed charging tariff or who did not wish to enroll in automation. The managed charging tariff off-peak window runs 23:30–05:30; the ToU tariff off-peak window runs 00:30–05:30.

5.2 Managed charging produces a more efficient price signal

It is well known that ToU pricing leaves substantial deadweight loss on the table relative to real-time pricing, because flat within-band rates fail to communicate the high-frequency variation in wholesale costs (Hinchberger et al., 2026; Borenstein, 2005a). Managed charging offers a way to recover some of this lost value without exposing households to volatile prices: the algorithm can act on wholesale-price variation that the posted tariff does not. We assess this channel by comparing how closely each tariff’s effective price signal tracks wholesale prices.

Following Hinchberger et al. (2026), the squared gap between wholesale prices and retail price signals approximates the deadweight loss from mispricing, and the R^2 from a regression of wholesale prices on a retail price signal is proportional to the efficiency gain achieved relative to real-time pricing.

We first applied this benchmark to the ToU tariff rate alone. Regressing the half-hourly ToU tariff unit rate on Great Britain electricity costs yielded an R^2 of 0.061 against full procurement cost (wholesale price plus non-energy charges; Column 2, Table A14)

and 0.005 against wholesale price alone (Column 2, Table A15).

We then constructed the analogous signal for the managed charging tariff. Because the tariff bundles a ToU schedule with automated charging, the price that effectively governs consumption in a given half-hour depends on charging state. In each active device-half-hour, customers were in one of three states: not plugged in, plugged in and automated, or plugged in and bump charging. We constructed the managed charging tariff effective price as

$$\hat{p}_h^{MC} = s_h^{noplug} p_h^{ToU} + s_h^{auto} y_h + s_h^{bump} p_h^{peak},$$

where y_h is the wholesale price and s_h are the population shares in each state. This decomposed the managed tariff into the part governed by automation (which inherits the wholesale price) and the part still governed by the ToU schedule.

The managed charging price signal tracked full procurement cost more closely than the static ToU tariff, with R^2 equal to 0.115, against 0.061 for the static ToU rate (Table A14). The managed charging price was also substantially better at tracking wholesale prices alone than the static ToU tariff, with R^2 equal to 0.096, against 0.005 for the ToU tariff (Table A15). This pattern is consistent with automation adding price information beyond the two-rate schedule. The magnitude was bounded by plug-in behavior: when vehicles were not plugged in or were bump charging, the effective price reverted to the ToU signal, so the ceiling on R^2 improvement is set by the share of half-hours in which the algorithm is in control.

The managed charging tariff effective price is not a posted tariff schedule of the kind Hinchberger et al. (2026) evaluate; it is an ex-post object that weights the ToU rate and the wholesale price by realized charging states. We apply their framework in a simplified form, and the comparison is best read as descriptive of how much wholesale-price information reaches the effective margin of consumption under each tariff.

5.3 Managed charging responds to real-time prices within ToU bands

The R^2 comparison establishes that the managed charging tariff price tracks wholesale costs more closely than the pure ToU tariff. Whether consumption actually responded to that variation is a separate question. If automation adds value beyond the ToU schedule, it should show up in how consumption responds to wholesale price variation *within* a ToU band, where retail prices are flat and any scheduling response must come from

the algorithm rather than the household. Neither managed charging tariff nor ToU tariff participants were directly exposed to wholesale prices, but the AI management of managed charging tariff schedules responds to those prices on the retailer’s behalf, such that when an EV is plugged in, it is responding to system prices.²⁸ Any observed difference in elasticity therefore reflects the supplier’s algorithmic responsiveness, not household-level reactions to system costs.

We estimated price responsiveness within the two ToU bands of the managed charging tariff: daytime peak (05:30–23:30) and overnight off-peak (23:30–05:30).²⁹ We used a Poisson regression, which accommodates zero consumption and yields coefficients that can be interpreted directly as elasticities of electricity consumption with respect to system prices (Mullahy and Norton, 2024). Table A16 reports results from the following specification:

$$\log(\mathbb{E}[Y_{it}]) = \alpha + \beta_1 D_{it} + \beta_2 \log(P_t) + \beta_3 [D_{it} \times \log(P_t)] + \mu_d \quad (3)$$

where Y_{it} is electricity consumption for trial participant i at date-hour t , D_{it} is an indicator for managed charging tariff participation, and P_t is the system price. The coefficient β_2 captures the price elasticity of demand among the baseline group (ToU tariff users), while β_3 captures the differential elasticity for managed charging tariff users. We also included fixed effects for day d , so that coefficients are identified from within-day variation across participants. The sample comprised the subset of our original trial sample ($n = 13,233$) who signed up for either the managed charging tariff or ToU tariff ($n = 2,894$).³⁰

The within-band estimates were consistent with an automation effect. Within the daytime peak band, managed charging tariff consumption was more price-elastic than ToU tariff consumption by approximately -0.016 (Table A16, Column 2). Within the overnight off-peak band, the differential was approximately -0.021 (Table A16, Column 3). In both bands, managed charging tariff consumption was more responsive to within-band wholesale price variation than ToU tariff consumption was. Because households faced flat retail prices inside each band, this differential is attributable to the algorithm rather than to household response to a price signal.³¹

²⁸In 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).

²⁹As a robustness check, we also estimate price responsiveness within the time bands of the static ToU tariff: daytime peak (05:30–00:30) and overnight off-peak (00:30–05:30).

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

³¹When pooling hours across the full day (Table A16, Column 1), ToU tariff customers appeared to consume relatively more when prices were low, a counterintuitive pattern. We cannot fully disentangle three contributors. First, selection:

5.4 Managed charging as a virtual power plant

Beyond improving the price signal that individual customers face, managed charging creates value by turning household EVs into a virtual power plant: an aggregation of distributed batteries that can be dispatched, verified, and settled in other electricity markets. In Great Britain, these include the Capacity Market, the Balancing Mechanism, and local flexibility services procured by Distribution Network Operators. The Capacity Market pays resources for being available during system stress events. The Balancing Mechanism is the real-time market used by the system operator to correct supply-demand imbalances close to delivery. Distribution Network Operators manage local distribution grids and can procure flexibility to relieve neighborhood-level congestion.

These markets require more than a static response to a fixed ToU schedule. Participation requires the aggregator to make commitments that the system operator can verify and settle against. This creates a precision requirement: the aggregator must forecast vehicle availability, infer charging needs, respect customer ready-by constraints, coordinate thousands of small batteries, and deliver a specified quantity of flexibility at a specified time. The optimization problem is therefore both high-frequency and high-dimensional. A fixed off-peak price can shift load across broad windows, but it cannot reliably produce accountable, dispatchable flexibility for real-time system operations.

Since 2022, Octopus Energy Limited has participated in the British Capacity Market, offering flexible demand from managed-charging EV customers as a substitute for keeping a fossil generator on standby. Its contracted role has steadily grown. Across delivery years 2022–2029, the retailer has been awarded an average of about 226 MW of capacity obligations per year. In the most recent auction, for delivery year 2029/30, it was awarded 507 MW, corresponding to a gross agreement value of £13.7 million for that delivery year alone.³² Figure A16 reports estimated gross agreement values by delivery year.

The retailer has also used the tariff as a virtual power plant in the UK Balancing Mechanism since 2024. It has bid 13 small EV fleets of approximately 5 MWh into the market,

ToU tariff adopters who declined the managed charging tariff incentive may have had more demand flexibility on other margins, which would have allowed them to time non-EV load to overnight hours. Second, the outcome captured total household consumption, and we lacked telemetry data for ToU tariff customers to isolate EV charging. Third, residual compositional differences between the two groups may have remained despite covariate balance (Table A17). Tariff assignment was not random. Many trial participants on the ToU tariff were unable to enroll in the managed charging tariff for compatibility reasons rather than preference, which bounded the scope for active selection. Tariff choice nonetheless remained mechanically correlated with vehicle and charger type, and a subset of ToU tariff customers may have declined the managed charging tariff to retain manual control over their charging schedule.

³²These values should not be interpreted as profits, since they are gross agreement values before penalties and delivery costs.

resulting in an average of 32 MWh of accepted volume per month (Figure A17). This volume is small relative to the full managed charging tariff portfolio, but it provides evidence that household EVs can be aggregated into a dispatchable resource. With approximately 300,000 vehicles on the tariff, the potential scale is much larger.³³

The retailer has also begun selling flexibility into local distribution markets, typically on a contract-by-contract or trial-by-trial basis.³⁴ Beyond formal market participation, the retailer also captures flexibility value through its own procurement, by shifting managed charging tariff charging in response to intraday wholesale and imbalance price movements rather than contracting with the system operator. The AI algorithm is what turns a fleet of household EVs into a resource that can respond across these channels.

6 Alternative estimation strategy: difference-in-differences of a larger and earlier-adopting sample

In addition to our experimental design, we leveraged a difference-in-differences (DiD) approach using observational data to estimate the impact of adopting the managed charging tariff on household electricity consumption. This strategy exploited the staggered, voluntary adoption of the tariff across the retailer’s customers in 2023 (one year prior to our experiment). Everything that follows was pre-specified in our pre-analysis plan, except where deviations are explicitly noted.

The natural field experiment from Section 4.2 remains our primary causal estimate. Customers often switch tariffs in conjunction with other behavior changes, including purchasing an EV, moving, or income shocks. As we discuss below, this threatens the parallel-trends assumption underlying DiD (LaLonde, 1986; Imbens and Xu, 2024). Randomized encouragement isolates the tariff effect from these co-occurring changes. Nonetheless, the DiD estimates contribute two pieces of evidence the experiment cannot provide. First, they reveal the impact of managed charging among customers who selected into the tariff without exogenous encouragement. Second, the staggered 2023 adoption window allows us to estimate effects separately by adoption cohort, showing how the tariff’s impact shifted over time.

³³A simple benchmark illustrates the order of magnitude: if each vehicle has a 60 kWh battery, the portfolio contains approximately 18 GWh of gross storage capacity. Only a fraction of that capacity is available at any point in time, but even a 10% available share would imply 1.8 GWh of flexible energy.

³⁴Examples include [Scottish and Southern Electricity Networks](#), [Northern Power Grid](#), and [UK Power Networks](#).

6.1 Empirical strategy

Switching to the tariff is often nearly coincident with other behavioral changes, most notably EV purchases. Figure A18 plots daily consumption for a random sample of 2023 managed charging tariff adopters without pre-existing EV load signatures. Consumption rose sharply around the tariff switch, increasing by 30% in the week after enrollment relative to the week before, consistent with EV acquisition coinciding with tariff adoption.

To isolate the effect of the managed charging tariff from EV adoption, we restricted the sample to customers who already owned an EV by December 2022, using the methodology in Section 3.1. The final sample contains 20,248 customers, drawn from an initial pool of 100,986 managed charging tariff adopters in 2023.³⁵ Other coincident behavioral changes remain a threat to identification. Customers who switch tariffs may also purchase additional low-carbon technologies, move, or change employment, none of which we observe. We cannot sign the direction of bias from these confounders, and the DiD estimates should be interpreted with this in mind.

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 made 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:

$$ATT(g, t) = \mathbb{E}[Y_t - Y_{g-1} \mid G = g] - \mathbb{E}[Y_t - Y_{g-5} \mid D_t = 0, G \neq g] \quad (4)$$

where $G = g$ denotes the cohort of units first treated at time g . We estimated 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. To mitigate potential bias from anticipation effects, we excluded the four weeks prior to adoption from our estimation. Accordingly, our reference period, Y_{g-5} , was hourly consumption measured five weeks prior to the treatment (Roth, 2024).³⁶

³⁵This 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.

³⁶We assumed that once a customer first adopts the managed charging tariff, they remain “treated” in the sense that their experience with the tariff continues to shape their behavior, even if they subsequently switched to another electricity tariff (Callaway and Sant’Anna, 2021). In practice, some customers did have more complex tariff histories. Under the irreversibility assumption, their electricity consumption patterns are considered to remain influenced by the tariff from the point of initial adoption. We view this as reasonable for two reasons: (1) 78% of customers who adopt the tariff subsequently remain on it for at least 12 months afterwards, and (2) it is plausible that the tariff induces some degree of habit formation, both in EV charging routines and in household electricity use more broadly.

To enhance comparability, we limited control cohorts to those scheduled to adopt the managed charging tariff no later than twelve weeks after the treated group’s anticipation period ends. This ensured treated units were compared only to future adopters with similar adoption timing. We chose a twelve-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 relied on the following parallel trends assumption: absent adoption, treated and not-yet-treated households would have experienced similar trends in electricity use.

In addition, to improve comparability with the field experiment estimates and probe underlying mechanisms, we estimated weighted versions of Equation 4, reweighting the DiD sample match the field experiment sample based on pre-treatment tariff type. The field experiment explicitly tried to exclude customers with any prior smart tariff usage.³⁷ Since smart tariffs incorporate time-of-use pricing structures, this exclusion disproportionately removed time-of-use tariff users from the field experiment sample. As a result, only 14% of the field experiment participants were on a time-of-use tariff at baseline, compared to 75% in the DiD sample.³⁸ 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 field experiment to DiD shares, ensuring that the reweighted DiD sample better reflected the field experiment’s pre-treatment tariff distribution.³⁹

Reweightings the DiD customers was important because their baseline consumption profiles differed substantially based on prior tariff. Figure A19 shows that DiD customers on flat tariffs, DiD customers on ToU tariffs, and the field experimental control group had distinct consumption patterns before adopting the managed charging tariff, even though overall consumption was similar across the three groups. Customers who had been on ToU tariffs already displayed a consumption pattern closely aligned with the

³⁷This 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.

³⁸For 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.

³⁹We also implemented propensity-score reweighting. Specifically, we estimated the probability of being in the field experiment (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 the field experiment), (3) the share of consumption occurring during peak hours, (4) tenure with the retailer, (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.

managed charging tariff’s incentives: minimal afternoon peaking and high overnight usage, consistent with many having been on the ToU tariff, which offered cheaper overnight rates from 00:30 to 05:30. In contrast, DiD customers who had been on flat tariffs still exhibited an afternoon peak, and while they also showed an overnight spike in consumption, its magnitude was only about one-third that of customers on ToU tariffs. Together, these patterns indicated that many DiD customers were already partially aligned with the managed charging tariff’s incentivized charging profile prior to adoption.


6.2 Difference-in-differences results

Figure 8 shows that the difference-in-differences estimates are notably smaller than those from the field experiment. In the unweighted specification (Column 1), managed charging tariff adoption is associated with a 0.06 kWh (7%) average hourly reduction in peak-period consumption, and a 0.1350 kWh (8%) increase during off-peak hours. In contrast, the field experiment estimates imply much larger shifts: a 0.581 kWh (42%) decrease during peak periods and a 0.481 kWh (50%) increase off-peak.

Comparing the DiD and experimental results: 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 field experiment sample’s pre-treatment tariff distribution. After reweighting, the estimated increase in off-peak consumption in the DiD analysis (0.451 kWh) matches the field experiment analysis (0.481 kWh). However, the reduction in peak consumption is still smaller in the DiD analysis: 0.24 kWh compared to 0.581 kWh in the field experiment. This weaker peak effect is compensated for by a decline in daytime, non-peak consumption in the DiD sample, as shown in Figure A20, resulting in no overall increase in consumption.

Two mechanisms could explain this difference in daytime treatment effects. First, DiD customers may have different plug-in behavior, giving the algorithm more flexibility to schedule daytime charging. We tested the first mechanism using telemetry data on plug-in rates. Figure A24a plots the difference in hourly plug-in rates between DiD and RCT managed charging tariff users, with positive values indicating greater availability among DiD customers. DiD adopters left their vehicles plugged in for longer morning windows, and Figure A24b shows this translated into more daytime charging. The differences are small, however, and cannot account for the full gap in daytime impacts between the two

Figure 8: Difference-in-Differences Estimate of Managed Charging Tariff



Notes: This figure reports staggered difference-in-differences estimates of the effect of adopting the managed charging 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 20,248 customers who first-ever enrolled in the managed charging tariff 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-adoption consumption levels. Post-treatment effects are estimated using average of all group-time average treatment effects, with weights proportional to the group size.

samples.

Second, DiD customers may have greater non-EV daytime flexibility, since the consumption impacts we measure are at the household level and reflect all appliances. We cannot test this directly, as we lack appliance and technology data linked to individual customers in our sample. Managed charging tariff customers do often have multiple low-carbon technologies. Internal surveys from the retailer indicate that of customers on the managed charging tariff, approximately 7% have a heat pump, 25.4% have solar panels, and 15% have a battery. To the extent that voluntary adoption selects on flexibility-related characteristics, this is consistent with this second mechanism. Both mechanisms are also difficult to separate from the broader identification concern in Section 6.1: voluntary managed charging tariff adoption may coincide with other behavioral changes that we cannot observe.

Treatment effects by adoption cohort: A benefit of a staggered DiD is that one can compare treatment effects depending on when an individual was treated (i.e., in this case, adopted the managed charging tariff). This is of policy relevance insofar as treatment effects are affected by whether an individual is an earlier or later adopter. It could be that earlier adopters are especially keen, tech savvy, or environmentally conscious. It could also be that later adopters are more unaware and thus have more still-shiftable consumption. How treatment effects change as a program scales is thus key to future forecasts. To get at this question, we estimated cohort-specific treatment effects, defining cohorts by the week of adoption. We found that treatment effects were homogeneous across cohorts, as shown in Figure A21.

This stability across cohorts supports the conclusion that the treatment effect of the managed charging tariff is stable across adopters. This conclusion is partially supported by two other pieces of evidence from our experimental and quasi-experimental analyses: first, that the reweighted DiD estimates align with the RCT estimates; and second, that the RCT itself shows no meaningful heterogeneity in treatment effects (Section 4.2).

We believe this stability is driven by the fact that the primary source of variation is baseline charging behavior, or whether customers were already on a time-of-use tariff before adoption, rather than heterogeneity in responsiveness to the tariff itself. This pattern reflects the mechanics of managed charging: once a customer enrolls, the algorithm schedules charging, and the consumer's role is limited to setting a target state of charge and a departure time. Treatment effects depend on the algorithm's optimization, not on

consumer behavior, which leaves little room for the kind of heterogeneity typically observed in behavioral price-response studies.⁴⁰

7 Welfare impacts

The large causal shifts in consumption from peak to off-peak hours from the managed EV tariff, as seen in Section 4.2, have four potential benefits: (1) benefits to consumers from lower electricity bills, (2) lower electricity procurement costs to the electricity supplier, (3) climate change mitigation benefits via CO₂e abatement, and (4) reduced grid operation and stabilization costs. In this section, we first present the benefits accrued in 2024. Then, we apply the Marginal Value of Public Funds (MVPF) framework developed by Hendren (2016) and Hendren and Sprung-Keyser (2020) to assess the welfare implications of subsidizing managed charging.

7.1 Benefits/savings from adoption of managed charging

By shifting consumption during lower-priced hours, the tariff generated several benefits, which we quantified using the methodology described in Appendix G.

Managed charging presented large consumer bill benefits. Using administrative data on customer bills, we found that managed charging reduced consumer electricity bills by £0.04 per kWh, an approximately 18% decline relative to the control group (Figure 9). Applied to average consumption over the analysis period, this implies a total bill reduction of approximately £343 during our study.^{41,42}

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; 95% CI: £-261 to £766). Procurement savings should not be interpreted as additional to consumer bill savings; the similarity of their magnitudes suggests near

⁴⁰We also conducted a DiD analysis focusing on customers who adopted the managed charging tariff in 2024, to make the comparison period more directly aligned with the field experiment timeframe. However, identifying which customers already owned an EV at the start of 2024 proved difficult. We thus use the 2023 as our main DiD specification. Full results and further discussion are provided in Appendix E

⁴¹Based on the average annual electricity consumption of control group trial participants not enrolled in an EV tariff (8,635 kWh).

⁴²The counterfactual for our £343 annual savings estimate is the average electricity bill of the control group during the experimental period, which includes customers starting on and adopting a mix of tariffs. If instead we apply a counterfactual based on a standard flat tariff, the estimated annual savings increase to £650. This figure is calculated using our estimated consumption treatment effects from Figure 5, applied to the bill under the flat tariff. The £650 estimate represents a 34% bill reduction.

100% pass-through of savings to trial participants, at least based on the period of our analysis (2024).⁴³

Managed charging also yielded positive environmental benefits. Direct CO₂e abatement from shifting electricity consumption to cleaner off-peak hours generated a decrease of 124 kg CO₂e per trial participant in 2024, which translates to estimated benefits of £35 based on a carbon value of £287 per tonne of CO₂e emitted (Department for Energy Security and Net Zero, 2023), though with wide uncertainty (95% CI: £-253 to £287).

Indirect CO₂e abatement, arising from substitution from ICE vehicles to EVs, due to lower operating costs of EVs,⁴⁴ generated a further decrease of 143 kg CO₂e, or £40.9 (95% CI: £20.6 to £61.2), in estimated benefits. The confidence intervals for indirect CO₂e abatement given in Figure 9 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 CO₂e damages.


We estimated a resource cost per tonne of approximately –£1,283, based on annual bill savings of £343 and total emissions reductions of 0.267 tonnes CO₂e per customer (this combines 0.124 tonnes CO₂e of direct and 0.143 tonnes CO₂e of indirect savings). Note that this calculation uses the annual savings as the cost, rather than taking the net present value of present and future savings. The £343 figure corresponds to average bill savings relative to the experimental control group; savings are £650 relative to the retailer’s standard flat tariff. The latter implies a resource cost of –£2,434 per tonne ($-650 \div 0.267$). Using a more conservative measure based on supplier procurement savings (£253 per customer per year), the implied cost is –£946 per tonne. In summary, the tariff and technology combination was a very cost-effective way to reduce CO₂e emissions, much lower than the next best technology (Gosnell et al., 2020; Gillingham and Stock, 2018; Hahn et al., 2024).

For savings from avoided system costs, we used the model in 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 in Great Britain by up to £0.25

⁴³The point estimate on procurement cost savings was smaller than the point estimate on bill savings, but the estimates’ confidence intervals overlap each other’s point estimates; also note that we have not included revenue from ancillary markets in the estimation of procurement cost savings.

⁴⁴We 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).

Figure 9: 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 £253 (95% CI: £-261 to £766). (These costs include both the wholesale price of power and non-energy charges such as transmission and distribution fees, but exclude revenue from grid services, such as participation in ancillary markets.) The similarity in the magnitudes of bill and procurement savings suggests that the supplier passed through nearly all cost reductions to customers over the 2024 analysis period. Direct CO₂e 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 CO₂e abatement from ICE to EV substitution is valued at £40.9 (95% CI: £20.6 to £61.2), estimated by combining observed managed charging tariff 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, where benefits range from £99 to £146 per year in 2035.

billion in 2025 and as much as £4 billion by 2035.⁴⁵ These savings arose from both operational efficiencies, such as increased use of lower-marginal-cost renewable generation, and capital 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.⁴⁶

⁴⁵Their 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 (Crozier et al., 2020). Second, it is aligned in the geographic focus on Great Britain. There are several relevant studies examining system 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.

⁴⁶As Franken et al. (2025) notes, “The first units of flexible EV charging tap into uncontested renewable generation, unlocking large benefits with a relatively modest flexibility rollout. However, beyond the 25% mark, excess renewable generation becomes more scarce, slowing down further gains.”

Our trial finds 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 (from 2035 onward, in £₂₀₂₄) could be feasible, conditional on widespread EV and tariff uptake, relative to a baseline with no flexible demand. When expressed on a per-vehicle basis, assuming 27 million EVs on the road in the UK by 2035 (National Energy System Operator, 2023), this equates to approximately £146 in system savings per vehicle in 2035 under universal managed charging.⁴⁷ Under “constrained” managed charging,⁴⁸ 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 £99.⁴⁹

To better understand how these reduced grid system costs might materialize, we estimated that if all EVs had used the managed charging tariff in 2024, 1.87 GW of load would have been shifted away from the highest demand hour in the country. Using demand data from the National Energy System Operator’s Future Energy Scenarios, we added our estimated hourly consumption effects from Figure 5 to hourly load values, and constructed what the load duration curve would have looked like for 2024 with 100% adoption of managed charging (Panel A, Figure A7). Panel B zooms in on the top 100 consumption hours, illustrating how managed charging could significantly reduce extreme peaks.

Looking ahead to 2050, the precise hourly impacts of managed charging are uncertain, given expected changes in electricity pricing, renewable generation patterns, and climate conditions. To approximate potential effects, we modified the 2050 load duration curve by applying the estimated impacts from the four peak hours in Figure 5 to the top four hours of each day in the 2050 load profile, assuming that the displaced consumption would be evenly redistributed across the remaining hours. Under this scenario, the highest load hour in 2050 would experience approximately 5 GW less consumption (Panel C and D in Figure A7).⁵⁰

⁴⁷In 2025, Franken et al. (2025) estimates the value is £94 per EV; the value is lower than in 2035 due to greater grid constraints in 2035 potentially solved by EV flexibility.

⁴⁸In 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.

⁴⁹These system benefits estimates are consistent with existing industry and research findings, and sit at the lower end of reported ranges. For instance, *The Utility Playbook: Turning EV Grid Risk into a \$30 Billion Opportunity* (2025) projects system savings of between \$145 and \$575 per actively managed EV by 2035.

⁵⁰The consumption profiles were obtained from the National Energy System Operator’s [Future Energy Scenarios](#), which provides baseline profiles for 2024 and projections for 2050 under the assumption of 0% unmanaged charging. To estimate the total load shift, our IV estimates were scaled by the total number of EVs on the road. Data on the number of battery electric vehicles (BEVs) and plug-in hybrid electric vehicles (PHEVs) in 2024 were obtained from

7.2 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 the managed charging tariff. 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) CO₂e 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 25% of adopters were marginal.⁵¹ 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.⁵² Combining these assumptions yields an average transfer of £0.875 per £1 of subsidy.

For welfare benefits from CO₂e 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.⁵³ We excluded reductions in trial participant bills from our welfare calculation, invoking the envelope theorem: these trial participants could have adopted the managed charging tariff without the subsidy but chose not to. We also assumed no producer surplus changes (e.g., procurement cost savings for the retailer), under the assumption of a competitive retail electricity market. Finally, we excluded the indirect CO₂e benefits from increased EV adoption, since, again by the envelope theorem, a sub-

the [UK's Vehicle Licensing Statistics](#), indicating a total of 2,107,897 vehicles. Projections for 2050 were also taken from the [Future Energy Scenarios](#), which forecast approximately 37.4 million EVs on UK roads by 2050.

⁵¹Calculated as $(9.31 - 6.98)/9.31 = 25\%$. We used this formulation of the marginals as opposed to using the elasticity approach as in Hahn et al. (2024) since we do not have a cost of adopting the tariff.

⁵²For 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.

⁵³We used the point estimate of the CO₂e impact, despite the statistical imprecision around that estimate.

sidy for managed charging should not affect EV uptake among trial participants who were already considering adoption.⁵⁴

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 CO₂e 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 less government revenue per customer. This translates to a fiscal cost of approximately £3.10 for each marginal managed charging tariff adopter. The climate-related fiscal externality is 1.07%⁵⁵ of the monetized value of CO₂e abatement attributable to managed charging tariff adoption. This amounts to £3.22 in additional government revenue per marginal managed charging tariff 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.⁵⁶ For the remainder of the discussion, we focus on the 10-year estimate, while estimates for 1 and 5 years are presented in A22.⁵⁷ This MVPF implies that for every £1 of fiscal cost to the government, the program generated £0.982 in societal 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 £1.255. (i.e., $\frac{1.232}{1.255} = 0.982$). The ratio of these benefits to costs is then 0.982.⁵⁸

We also think about the MVPF inclusive of avoided grid costs, though these benefits are uncertain, and may be zero. In a well-functioning electricity market, the benefits of shifting to lower-marginal-cost generation and deferring investment in generation, trans-

⁵⁴This 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 F.

⁵⁵We 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%.

⁵⁶As noted above, the duration of benefits depends not just on vehicle lifetime, but also on how long subsidized customers remain more likely than non-subsidized customers to adopt the managed charging tariff.


⁵⁷We also present MVPFs using two alternative SCC values: (a) Bilal and Känzig (2024), which estimates an SCC of \$1,367 for 2024, and (b) [Interagency Working Group on Social Cost of Greenhouse Gases, United States Government](#), which estimates an SCC of \$55.3 for 2024.

⁵⁸We calculated this using:

$$MVPF = \frac{xds + Edx}{xds + Vdx + Cdx + Gdx} \quad (5)$$

where x is quantity and s the subsidy. E represents CO₂ benefits to individuals; V , C and G represent VAT, climate change fiscal externalities, and avoided grid balancing costs respectively.

Figure 10: 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 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 CO₂e benefits reflect emissions reductions from shifting electricity use to cleaner hours, scaled to marginal adopters. Indirect CO₂e 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 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.

mission, and distribution infrastructure should be reflected in price signals and internalized by market participants. Thus, while the overall system benefits of the managed charging tariff may be large, only a small share of those would accrue to the government. However, market failures may prevent the value of EV flexibility from being fully internalized. In the presence of such failures, there may be rationale for a second-best policy instrument to incentivize managed charging, such as subsidies.⁵⁹

In the UK, as in many electricity systems, grid costs are recovered through customers' electricity bills. When customers bear these costs, we log it as a benefit in the numerator, using the £146 in per-customer benefits in 2035 estimated by Franken et al. (2025).

⁵⁹This issue has some symmetry to the MVPF of healthcare subsidies, as the MVPF varies greatly depending on whether low-income individuals pay for their own healthcare, hospitals do, or the government does (Finkelstein et al., 2019).

Under this accounting, the 10-year MVPF is 2.8. If the government were instead to fund the additional capacity, the saving would reduce its net fiscal cost rather than benefit customers. If £70 (48%) of the £146 were borne by the government, the MVPF of the £150 subsidy would be infinite.⁶⁰

8 Discussion

Our findings make three key contributions to the literature. First, we provided 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 clear causal inference about the elasticity of managed EV charging consumption to energy system price signals.

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. We showed that this override metric is the sufficient statistic to compare such AI tariff versus RTP, and thus highlights the potential of automation through AI to provide demand-side flexibility while respecting EV owners' preferences and constraints. The automation embedded in the tariff appeared to enhance responsiveness to grid signals, especially during the evening and overnight periods of the day.

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 managed charging was 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 field experiment 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.

⁶⁰In 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 (a Pareto improvement).

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 participants bills substantially. It also caused reductions in CO₂e emissions and retailer procurement costs.

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.
- 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.
- Bailey, M. R., Brown, D. P., 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* **7**(4), 550–566.
- 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.
- 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.
- Bernard, L., Hackett, A., Metcalfe, R. and Schein, A. (2024), Decarbonizing Heat: The Impact of Heat Pumps and a Time-of-Use Heat Pump Tariff on Energy Demand, Technical Report w33036, National Bureau of Economic Research, Cambridge, MA.

- Bilal, A. and Känzig, D. R. (2024), ‘The Macroeconomic Impact of Climate Change: Global vs. Local Temperature’.
- 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:2502.12397 [cs].
- 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.
- Bollinger, B. K. and Hartmann, W. R. (2020), ‘Information vs. Automation and Implications for Dynamic Pricing’, *Management Science* **66**(1), 290–314.
- Borenstein, S. (2005a), ‘The long-run efficiency of real-time electricity pricing’, *Energy Journal* **26**(3).
- Borenstein, S. (2005b), ‘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. Publisher: SAGE Publications Sage CA: Los Angeles, CA.
- Boyd, S. and Vandenberghe, L. (2004), *Convex Optimization*, Cambridge University Press.
- 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.
- Budish, E. (2011), ‘The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes’, *Journal of Political Economy* **119**(6), 1061–1103.
- Burkhardt, J., Gillingham, K. and Kopalle, P. (2021), Field Experimental Evidence on the Effects of Information and Pricing on Residential Electricity Conservation, Technical Report 25576, National Bureau of Economic Research, Cambridge, MA.
- Burlig, F., Bushnell, J. and Rapson, D. (2026), If You Build It, They May Not Come: Willingness to Participate in Managed EV Charging, Technical Report w35086, National Bureau of Economic Research, Cambridge, MA.
- 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.
- Callaway, B. and Sant’Anna, P. H. (2021), ‘Difference-in-Differences with Multiple Time Periods’, *Journal of Econometrics* **225**(2), 200–230.

- Colin Cameron, A. and Miller, D. L. (2015), 'A Practitioner's Guide to Cluster-Robust Inference', *Journal of Human Resources* **50**(2), 317–372.
- 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.
- Daneshzand, F., Coker, P. J., Potter, B. and Smith, S. T. (2023), 'EV Smart Charging: How Tariff Selection Influences Grid Stress and Carbon Reduction', *Applied Energy* **348**, 121482.
- 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.
- Department for Energy Security and Net Zero (2023), 'Green Book Supplementary Guidance: 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. Published: PDF report via gov.uk.
- Department for Transport (2024), 'Vehicle Licensing Statistics, United Kingdom: 2024'. Published: Official release via gov.uk.
- 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.
- Fabra, N., Rapson, D., Reguant, M. and Wang, J. (2021), 'Estimating the Elasticity to Real-Time Pricing: Evidence from the Spanish Electricity Market', *AEA Papers and Proceedings* **111**, 425–429.
- Few, J., McKenna, E., Pullinger, M., Elam, S., Webborn, E. and Oreszczyn, T. (2022), Smart Energy Research Lab: Energy Use in GB Domestic Buildings 2021, Technical Report 1, University College London.
- Finkelstein, A., Hendren, N. and Shepard, M. (2019), 'Subsidizing health insurance for low-income adults: Evidence from Massachusetts', *American Economic Review* **109**(4), 1530–1567.
- 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. and others (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.
- Gabaix, X. (2019), Behavioral inattention, in ‘Handbook of behavioral economics: Applications and foundations 1’, Vol. 2, Elsevier, pp. 261–343.
- 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.
- Garnache, C., Hernæs, O. and Imenes, A. G. (2025), ‘Demand-Side Management in Fully Electrified Homes’, *Journal of the Association of Environmental and Resource Economists* **12**(2), 257–283.
- Gillan, J. (2017), ‘Dynamic Pricing, Attention, and Automation: Evidence from a Field Experiment in Electricity Consumption’.
- Gillingham, K. and Stock, J. H. (2018), ‘The cost of reducing greenhouse gas emissions’, *Journal of Economic Perspectives* **32**(4), 53–72.
- 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’. Published: Current News.

- 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. W., Hendren, N., Metcalfe, R. D. and Sprung-Keyser, B. (2024), 'A Welfare Analysis of Policies Impacting Climate Change'.
- Harberger, A. C. (1964), 'The measurement of waste', *The American Economic Review* **54**(3), 58–76.
- Harding, M. and Sexton, S. (2017), 'Household Response to Time-Varying Electricity Prices', *Annual Review of Resource Economics* **9**(1), 337–359.
- Harrison, G. W. and List, J. A. (2004), 'Field experiments', *Journal of Economic literature* **42**(4), 1009–1055.
- Heid, P., Remmy, K. and Reynaert, M. (n.d.), '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. (2016), 'The policy elasticity', *Tax Policy and the Economy* **30**(1), 51–89.
- Hendren, N. and Sprung-Keyser, B. (2020), 'A unified welfare analysis of government policies', *The Quarterly Journal of Economics* **135**(3), 1209–1318.
- 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.
- Hinchberger, A. J., Jacobsen, M. R., Knittel, C. R., Sallee, J. M. and van Benthem, A. A. (2026), 'The Efficiency of Dynamic Electricity Prices', *NBER Working Paper*.
- HM Treasury (2020), 'The Green Book: Appraisal and Evaluation in Central Government'.
- 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.
- Imbens, G. and Xu, Y. (2024), 'LaLonde (1986) after Nearly Four Decades: Lessons Learned'.
- International Energy Agency (2024), 'World Energy Outlook 2024'.
- International Energy Agency (2025), '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.

- Jessoe, K. and Rapson, D. (2014), 'Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use', *American Economic Review* **104**(4), 1417–1438.
- Joskow, P. and Tirole, J. (2006), 'Retail electricity competition', *The RAND Journal of Economics* **37**(4), 799–815.
- Kaack, L. H., Donti, P. L., Strubell, E., Kamiya, G., Creutzig, F. and Rolnick, D. (2022), 'Aligning artificial intelligence with climate change mitigation', *Nature Climate Change* **12**(6), 518–527.
- Khanna, S., Martin, R. and Muûls, M. (2024), 'Shifting household power demand across time: incentives and automation', *Work. Pap., Imp. Coll. London*.
- 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.
- 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.
- LaLonde, R. J. (1986), 'Evaluating the Econometric Evaluations of Training Programs with Experimental Data', *American Economic Review* **76**(4), 604–620. JSTOR: 1806062.
- 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.
- Liang, A. (2026), 'Artificial Intelligence Clones'. arXiv:2501.16996 [econ].
- List, J. A. (2026), *Experimental economics: Theory and practice*, University of Chicago Press.
- 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'.
- Mullahy, J. and Norton, E. C. (2024), 'Why Transform Y? The Pitfalls of Transformed Regressions with a Mass at Zero', *Oxford Bulletin of Economics and Statistics* **86**(2), 417–447. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/obes.12583>.

- National Energy System Operator (2023), Future Energy Scenarios (2023), Technical report, National Energy System Operator.
- Nicolson, M., Huebner, G. M., Shipworth, D. and Elam, S. (2017), ‘Tailored Emails Prompt Electric Vehicle Owners To Engage With Tariff Switching Information’, *Nature Energy* 2(6), 17073.
- Nicolson, M. L., Fell, M. J. and Huebner, G. M. (2018), ‘Consumer demand for time of use electricity tariffs: A systematized review of the empirical evidence’, *Renewable and Sustainable Energy Reviews* 97, 276–289.
- Office for Budget Responsibility (2024), ‘The UK’s 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’.
- Powell, S., Cezar, G. V., Min, L., Azevedo, I. M. L. 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’.
- Rapson, D. and Muehlegger, E. (2023a), ‘Global transportation decarbonization’, *Journal of Economic Perspectives* 37(3), 163–188.
- Rapson, D. S. and Muehlegger, E. (2023b), ‘The Economics of Electric Vehicles’, *Review of Environmental Economics and Policy* 17(2), 274–294.
- Rigas, E. S., Ramchurn, S. D. and Bassiliades, N. (2014), ‘Managing electric vehicles in the smart grid using artificial intelligence: A survey’, *IEEE Transactions on Intelligent Transportation Systems* 16(4), 1619–1635.
- Roth, J. (2024), ‘Interpreting Event-Studies from Recent Difference-in-Differences Methods’.
- Rusak, G., Manning, B. S. and Horton, J. J. (2025), ‘AI Agents Can Enable Superior Market Designs’.
- SEPA (2024), The State of Managed Charging in 2024, Technical report, Smart Electric Power Alliance.
- Si, C., Chau, K., Liu, W. and Hou, Y. (2025), ‘Optimal scheduling for electric vehicle charging: A review of methods, technologies, and uncertainty management’, *Journal of Energy Storage* 131, 117500.
- The Utility Playbook: Turning EV Grid Risk into a \$30 Billion Opportunity* (2025), Technical report, ev.energy.

- 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.
- Varian, H. R. (1995), ‘Economic Mechanism Design for Computerized Agents’.
- Wang, Z., Zheng, F. and Liu, M. (2025), ‘Charging Scheduling of Electric Vehicles Considering Uncertain Arrival Times and Time-of-Use Price’, *Sustainability* **17**(3), 1100.
- WattTime (2022), Marginal Emissions Modeling: WattTime’s Approach to Modeling and Validation, Technical report.
- Wolak, F. A. (2011), ‘Do Residential Customers Respond to Hourly Prices? Evidence from a Dynamic Pricing Experiment’, *American Economic Review* **101**(3), 83–87.
- Yaghoubi, E., Yaghoubi, E., Khamees, A., Razmi, D. and Lu, T. (2024), ‘A systematic review and meta-analysis of machine learning, deep learning, and ensemble learning approaches in predicting EV charging behavior’, *Engineering Applications of Artificial Intelligence* **135**, 108789.
- Zheng, S., Zhu, Y., Ye, X. and Shen, L. (2025), ‘Generative Search: Evidence from a Large-Scale Field Experiment’.
- 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.

Appendix

Table of Contents

A	Tables	56
B	Figures	71
C	Additional details on design of field trial	89
C.1	Reproduction of email-based encouragements	89
C.2	Block randomization implementation	93
D	Empirical results from natural field experiment	94
D.1	Impact of encouragement on take-up of managed charging	94
D.2	Impact of encouragements on electricity consumption	97
D.3	Heterogeneity	98
D.4	Generalizability to other countries	104
D.5	Variance decomposition of bump behavior	108
E	2024 difference-in-differences	109
F	Deviations from the Pre-Analysis Plan	112
G	Methodology for calculating welfare impacts	115
H	Theoretical Model	117
H.1	Environment and Notation	118
H.2	Household and Aggregator Problems	119
H.3	Core Lemmas and Propositions	120
H.4	Overrides (“Bump Charging”): Behavior and Welfare	123
H.5	Cross-country calibration and crossover analysis	125

A Tables

Table A1: Literature Review: Residential Electricity Demand Flexibility

Paper	Load	Flexibility Mechanism	Research design	Main Peak Effect	Study length
<i>Panel A. Pricing only (no automation)</i>					
Wolak (2011); US	Whole-home	Hourly RTP, CPP, or critical-peak rebate	Framed field experiment; N = 1,245 (857 T / 388 C)	Hourly RTP: -2.6% summer demand; CPP: -8.9%	12 months + 3 CPP events
Jessoe and Rapson (2014); US	Whole-home	Short-run CPP events; in-home display for subset	Framed field experiment; N = 437 households	Price-only: -0 to -7%; Prices + information: -8 to -22%	Summer 2011
Bailey, Brown, Shaffer and Wolak (2025); Canada	Home EV charging	Off-peak payment (3.5¢/kWh) and/or prosocial nudge	Framed field experiment; N = 150 EVs	-49% peak-hour EV charging; off-peak share 59%→77%	12 months
La Nauze et al. (2024); Australia	EV charging (all locations)	Peak penalty + midday reward (for non-solar)	Framed field experiment; N = 390 EVs	-27% (non-solar) / -17% (solar) peak EV; +34% midday for non-solar	Several months
<i>Panel B. Pricing + deterministic automation or enabling technology</i>					
Bollinger and Hartmann (2020); US	Whole-home	ToU or variable peak pricing + info/IHD/PCT	Framed field experiment; N = 2,210 households	PCT automation: -1.15 to -1.17 kW in critical period; price response ~3× larger with automation	Summer 2011
Gillan (2017); US	Whole-home	Dynamic prices + messaging + up to \$240 device rebate	Framed field experiment; N = 5,531 households	31% price rise: -11% consumption; automation effects ~5× larger than average	3 months
Khanna et al. (2024); India	Connected appliances (mostly A/C); household electricity	30-min automated smart-switch shutoffs + per-kWh rewards + notifications + free hardware	Framed field experiment; N = 584 smart-switch users (>1,000 overall)	Device-level: -59 to -60%; household-level: -8.5% during event	3 months
Blonz et al. (2025); Canada	HVAC / A/C	Smart thermostat (Ecobee Eco+) auto-adjusts set-points to existing ToU	Framed field experiment; N = 2,133 households	Compliers: A/C compressor run-time -7.28 min/peak hr (~88% vs control)	3 months
Bailey, Brown, Myers, Shaffer and Wolak (2025); Canada	Home EV charging + transformer-level congestion	Utility-managed charging vs. ToU	Framed field experiment; N = 202 EVs	ToU: -20.3% peak EV kWh; Managed: -4.8%	6 months
<i>Panel C. Pricing + AI-managed charging</i>					
Burlig et al. (2026); US	Home EV charging (+ whole-home)	Managed-charging app + monthly incentives (\$0/\$5/\$20/\$40); ToU-steepening arms	Natural field experiment; N = >12,000 eligible EV households	Low enrollment, and no statistically meaningful ITT effect on electricity consumption	—
This paper (2026); UK	EV charging + whole-household	AI-managed charging on ToU tariff (IO Go); real-time wholesale + ancillary market optimisation	Natural field experiment; N = 13,233	-42% whole-household peak	12 months

Note: This table contextualises our research findings against the literature on residential electricity demand flexibility. Studies are grouped by the nature of the intervention: (A) pricing only, (B) pricing combined with deterministic automation or enabling technology (e.g., smart thermostats, programmable control), and (C) pricing combined with AI-managed charging. “T” / “C” denote treatment / control. “CPP” is critical-peak pricing; “RTP” is real-time pricing; “ToU” is time-of-use pricing; “PCT” is programmable communicating thermostat; “IHD” is in-home display; “DiD” is difference-in-differences. Our paper is the final row.

Table A2: 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 (0.01)	0.00 (0.02)	0.00 (0.02)	0.00 (0.02)
Octopus Tenure	2.72 [1.63]	-0.03 (0.04)	-0.01 (0.06)	-0.03 (0.06)	-0.01 (0.06)
Proportion Total kWh Peak	0.23 [0.08]	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Smart Tariff Onboarding Processes	0.22 [0.47]	0.00 (0.01)	0.00 (0.02)	-0.01 (0.02)	0.00 (0.02)
Structural Winnings (GBP)	752.64 [684.94]	-9.72 (16.50)	-17.76 (25.08)	-26.70 (23.84)	-18.12 (25.02)
Total kWh	4,082.53 [2,599.11]	-42.42 (62.80)	-33.33 (97.97)	-65.43 (91.35)	-38.96 (94.82)
Total kWh Stdev	0.77 [0.25]	-0.01 (0.01)	-0.01 (0.01)	0.00 (0.01)	0.00 (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 C.2.

Table A3: Impact of Encouragements on Take-up of Managed Charging Tariff

Dependent Variable: Model:	managed charging tariff (1)
<i>Variables</i>	
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)
<i>Fixed-effects</i>	
Block	Yes
Week of Year	Yes
<i>Fit statistics</i>	
Control Mean (Incentive Period)	0.0365
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 the managed charging tariff in the 12 months since emails were sent out. The outcome is a binary indicator for weekly use of the managed charging 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 managed charging tariff 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: Robustness, Impact of Encouragements on Take-up of Managed Charging Tariff, Daily Data

Dependent Variable: Model:	managed charging tariff (1)
<i>Variables</i>	
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)
<i>Fixed-effects</i>	
Block	Yes
Day	Yes
<i>Fit statistics</i>	
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 the managed charging tariff in the 12 months since emails were sent out. The outcome is a binary indicator for daily use of the managed charging 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 managed charging tariff 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 A5: Impact of Encouragements on Electricity Consumption

Model:	Peak (1)	Off-Peak (2)	Overall (3)	Peak (4)	Off-Peak (5)	Overall (6)
<i>Variables</i>						
Email + £0/Month	-0.029*** (0.011)	0.018 (0.012)	-0.005 (0.007)			
Email + £5/Month	-0.011 (0.016)	0.018 (0.017)	0.004 (0.010)			
Email + £50/Month	-0.013 (0.016)	0.025 (0.017)	0.003 (0.010)			
Email + £50/Month (No Bump)	-0.033** (0.016)	0.026 (0.018)	-0.006 (0.010)			
Any Encouragement				-0.026** (0.010)	0.019* (0.012)	-0.004 (0.007)
<i>Fixed-effects</i>						
Week of Year	Yes	Yes	Yes	Yes	Yes	Yes
Block	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
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 A6: Impact of Encouragements on ToU Tariff Take-up

Dependent Variable: Model:	ToU tariff (1)
<i>Variables</i>	
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)
<i>Fixed-effects</i>	
Block	Yes
Week of Year	Yes
<i>Fit statistics</i>	
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, the ToU tariff in the 12 months since emails were sent out. The outcome is a binary indicator for weekly use of the ToU tariff. The specification controls for fixed effects for randomization block and week. Standard errors, clustered by participant and week, are reported in parentheses. Mean ToU tariff take-up rate is reported for the control group.

Table A7: Selection into Managed Charging Tariff and Baseline Characteristics

Variable	(1)	(2)	(3)	(4)
Email + £0/Month	0.044	0.043	0.044	0.048
Email + £5/Month	0.043	0.043	0.044	0.048
Email + £50/Month	0.064	0.063	0.064	0.070
Email + £50/Month (No Bump)	0.060	0.059	0.061	0.066
Structural Winnings (Z-Score)	-0.010	-0.001	0.004	
Baseline TOU Tariff		0.022	0.022	0.018
Baseline kWh		-0.010	-0.012	-0.010
IMD Tercile 2		0.032	0.034	0.036
IMD Tercile 3		0.029	0.035	0.037
Octopus Tenure (Years)		-0.006	-0.006	-0.007
Covariates			+ Covariates interacted	+ Nonparametric struc. winnings
Sq. Corr. Coef	0.0047	0.0078	0.0093	0.0103
p-score range	[0.027, 0.17]	[0.018, 0.21]	[0.0000082, 0.28]	[0.0011, 0.24]
p-score R2 with (1)		0.60	0.51	0.46

Note: This table shows the estimation results of a logit regression of take-up of the managed charging tariff on encouragement group and participants' baseline characteristics. We show the marginal effects at the means of the covariates.

Table A8: 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	No Week FE
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>Variables</i>									
EV Tariff	-0.581** (0.224)	-0.498** (0.209)	-0.697*** (0.247)	-0.269 (0.338)	-0.305 (0.230)	-0.587** (0.236)	-0.473** (0.214)	-0.554 (0.336)	-0.568** (0.224)
<i>Fixed-effects</i>									
Week of Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Block	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>									
Control Mean	1.4	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	2,646,712
First Stage F-Stat	52.720	14.066	44.154	20.302	37.764	40.647	78.180	52.715	52.910

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 5; (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 A9: Robustness, Impact of EV Tariff on Off-Peak Consumption (kWh)

Model:	Main	4 Instruments	£0/Mth	£5/Mth	£50/Mth	£50/Mth (No Bump)	6 Months	No Baseline	No Week FE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>Variables</i>									
EV Tariff	0.481*	0.468*	0.501*	0.284	0.280	0.414	0.686***	0.118	0.488*
	(0.261)	(0.244)	(0.282)	(0.400)	(0.269)	(0.270)	(0.236)	(0.392)	(0.260)
<i>Fixed-effects</i>									
Week of Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Block	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>									
Control Mean	0.99	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	3,970,004
First Stage F-Stat	53.449	14.283	44.658	20.564	38.473	41.213	78.912	52.756	53.647

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 5; (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: Robustness, Impact of EV Tariff on Overall Consumption (kWh)

Model:	Main	4 Instruments	£0/Mth	£5/Mth	£50/Mth	£50/Mth (No Bump)	6 Months	No Baseline	No Week FE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>Variables</i>									
EV Tariff	-0.079	-0.060	-0.125	0.032	-0.038	-0.132	0.002	-0.203	-0.071
	(0.150)	(0.140)	(0.163)	(0.226)	(0.150)	(0.159)	(0.154)	(0.247)	(0.150)
<i>Fixed-effects</i>									
Week of Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Block	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>									
Control Mean	1	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	15,879,842
First Stage F-Stat	52.809	14.090	44.257	20.465	37.912	40.763	78.281	52.747	53.001

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 A11: Heterogeneity of Managed Charging Tariff Take-up, by IMD and Baseline Electricity Consumption

Model:	x IMD (1)	x Baseline kWh (2)
<i>Variables</i>		
Email + £0/Mth × Tercile 1	0.021 (0.018)	0.029*** (0.010)
Email + £0/Mth × Tercile 2	0.045*** (0.009)	0.050*** (0.009)
Email + £0/Mth × Tercile 3	0.026*** (0.007)	0.016* (0.009)
Email + £5/Mth × Tercile 1	0.001 (0.026)	0.035** (0.016)
Email + £5/Mth × Tercile 2	0.052*** (0.017)	0.033** (0.015)
Email + £5/Mth × Tercile 3	0.022* (0.011)	0.019 (0.015)
Email + £50/Mth × Tercile 1	0.009 (0.031)	0.069*** (0.018)
Email + £50/Mth × Tercile 2	0.062*** (0.018)	0.064*** (0.017)
Email + £50/Mth × Tercile 3	0.059*** (0.013)	0.036** (0.016)
Email + £50/Mth (No Bump) × Tercile 1	0.050 (0.033)	0.080*** (0.019)
Email + £50/Mth (No Bump) × Tercile 2	0.062*** (0.017)	0.045*** (0.016)
Email + £50/Mth (No Bump) × Tercile 3	0.035*** (0.012)	0.009 (0.015)
Tercile 2	-0.009 (0.017)	-0.013 (0.012)
Tercile 3	0.010 (0.017)	-0.0005 (0.013)
<i>Fixed-effects</i>		
Block	Yes	Yes
Week of Year	Yes	Yes
<i>Fit statistics</i>		
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 the managed charging tariff, 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 managed charging 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 A12: Heterogeneity of Impact of EV Tariffs, by IMD and Baseline Electricity Consumption

Model:	x IMD		x Baseline kWh	
	Off-Peak (1)	Peak (2)	Off-Peak (3)	Peak (4)
<i>Variables</i>				
(EV Tariff) x (Tercile 1)	0.924 (0.960)	1.16 (1.08)	0.295 (0.291)	-0.300 (0.221)
(EV Tariff) x (Tercile 2)	0.604 (0.388)	-0.882** (0.383)	0.504 (0.316)	-0.399 (0.277)
(EV Tariff) x (Tercile 3)	0.508* (0.297)	-0.207 (0.269)	0.998 (0.778)	-0.759 (0.725)
(Other Tariff) x (Tercile 2)	0.078 (0.127)	0.327** (0.140)	0.059 (0.072)	0.182*** (0.060)
(Other Tariff) x (Tercile 3)	0.098 (0.127)	0.218 (0.138)	0.131 (0.134)	0.441*** (0.125)
<i>Fixed-effects</i>				
Block	Yes	Yes	Yes	Yes
Week of Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
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 (23:30–05:30) 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 A13: Bump Charging and Baseline Characteristics

Dependent Variable: Model:	I(Ever Bumped) (1)
<i>Variables</i>	
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)
<i>Fit statistics</i>	
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 C.2.

Table A14: R^2 from regressions of managed charging and ToU tariffs on full electricity procurement cost

Model:	Managed tariff (1)	ToU tariff (2)	ToU, Off-peak Dummy (3)
<i>Variables</i>			
Constant	11.03*** (0.13)	11.64*** (0.16)	17.70*** (0.05)
Managed charging, avg rate	0.32*** (0.01)		
ToU rate		0.24*** (0.01)	
I(Off-peak)			-4.55*** (0.11)
<i>Fit statistics</i>			
R^2	0.115	0.061	0.083

IID standard-errors in parentheses

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Notes: This table reports OLS regressions of half-hourly full electricity procurement costs on tariff-implied retail price signals. Prices are measured in pence per kWh, and include both wholesale energy and non-energy costs. For wholesale energy costs, we used Elexon system buy prices, colloquially known as the "system" or "imbalance" price. Tariff unit rates came from retailer half-hourly tariff rate tables. Column 1 used the constructed managed charging tariff price signal: in each active device-half-hour, the price equaled the managed charging tariff ToU rate when the EV was not plugged in, the system price when the EV was plugged in and automated, and the managed charging tariff peak rate when the customer used bump charging. These customer states were aggregated to the half-hour level before estimation. Column 2 used the observed ToU tariff half-hourly unit rate. Column 3 replaced the unit rate with an indicator for the ToU tariff off-peak window, defined as 00:30–05:30, so that the reported R^2 captures the explanatory power of the tariff's timing if the tariff rates were set optimally within the bands. Plug-in and bump-charging states came from retailer telemetry for the randomized trial sample. Standard errors are reported in parentheses.

Table A15: R^2 from regressions of managed charging and ToU tariff rates on wholesale price

Model:	Managed tariff (1)	ToU tariff (2)	ToU, Off-peak Dummy (3)
<i>Variables</i>			
Constant	2.12*** (0.13)	6.17*** (0.14)	8.02*** (0.05)
Managed charging, avg rate	0.32*** (0.01)		
ToU rate		0.07*** (0.01)	
I(Off-peak)			-1.68*** (0.10)
<i>Fit statistics</i>			
R^2	0.096	0.005	0.011

IID standard-errors in parentheses

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Notes: This table reports OLS regressions of half-hourly electricity system prices on tariff-implied retail price signals. As in Table A14, prices are measured in pence per kWh, and we use Elexon system buy prices as the wholesale energy cost. Here, however, we ignore non-energy costs. Column 1 used the constructed managed charging tariff price signal: in each active device-half-hour, the price equaled the managed charging tariff ToU rate when the EV was not plugged in, the system price when the EV was plugged in and automated, and the managed charging tariff peak rate when the customer used bump charging. These customer states were aggregated to the half-hour level before estimation. Column 2 used the observed ToU tariff half-hourly unit rate. Column 3 replaced the unit rate with an indicator for the ToU tariff off-peak window, defined as 00:30–05:30, so that the reported R^2 captures the explanatory power of the tariff's timing if the tariff rates were set optimally within the bands. Tariff unit rates came from retailer half-hourly tariff rate tables. Plug-in and bump-charging states came from retailer telemetry for the randomized trial sample. Standard errors are reported in parentheses.

Table A16: Responsiveness to systems prices

Model:	All Hours (1)	05:30-23:30 (2)	23:30-05:30 (3)	05:30-00:30 (4)	00:30-5:30 (5)
<i>Variables</i>					
Managed charging tariff	0.089*** (0.025)	0.091*** (0.027)	-0.118*** (0.036)	0.052* (0.030)	-0.241*** (0.027)
log(Price)	-0.132*** (0.011)	0.047*** (0.008)	0.004 (0.015)	0.009 (0.009)	0.011 (0.015)
Managed charging tariff $\times$ log(Price)	0.017** (0.008)	-0.016* (0.008)	-0.021* (0.011)	-0.056*** (0.010)	-0.021*** (0.006)
<i>Fixed-effects</i>					
Date	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>					
ToU tariff Mean	1.08	0.78	1.79	0.80	1.91
Observations	14,924,989	10,667,975	4,257,014	11,300,042	3,624,947

Clustered (User & date-hour) standard-errors in parentheses

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Notes: This table reports estimates from Equation (3), examining how average hourly electricity consumption under the managed charging tariff and ToU tariff responds to system prices. Column (1) pools together hours across the whole day, column (2) covers daytime, peak hours (05:30–23:30), and column (3) overnight off-peak hours (23:30–05:30). Columns (4) and (5) are analogous to (2) and (3), but using the time bands according to the static ToU tariff schedule. All regressions included day fixed effects, so coefficients are identified from within-day variation across participants; standard errors were clustered by user and hour.

Table A17: Baseline Differences Between Managed Charging Tariff and ToU Tariff Trial Participants

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

Note: Column (1) shows the mean and [standard deviation] for managed charging tariff users; column (2) reports the same for ToU tariff 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 C.2.

Table A18: Overview of Managed Charging Tariff Features by Country

Country	Cheap charging	Price during charging window	Additional Benefits	Launch Date
UK	6 hour window: 23:30–05:30	7p/kWh	7p/kWh for the whole home between 23:30 and 05:30	May 2022
US	Dynamic based on forecasts	As per regular tariff price	Lower per kWh rate when connected to IO Go	December 2022
Germany (NB: changed as of June 2025)	5 hour window: 00:00–05:00	maximum of €0.20/kWh	Super-cheap electricity for the entire home between 00:00 and 05:00	August 2023
Spain	Dynamic based on forecasts	€0.07/kWh	n/a	September 2024


Table A19: Bump Charging per Charge Day

Country	Bumps/day	Peak bumps/day	Overnight bumps/day	Daytime bumps/day
United Kingdom	0.019	0.003	0.003	0.013
United Kingdom - Early Adopters	0.022	0.005	0.003	0.015
Germany	0.107	0.019	0.008	0.080
Spain	0.105	0.007	0.027	0.071
United States	0.035	0.011	0.004	0.020


Note: This table presents the frequency of bumping each day there is a charge event, for each of the four major countries where the managed charging tariff is active. This uses data from a random sample of 4,442 users across the four countries. Peak times are defined as 16:30–20:30 for the UK, 17:00–21:00 for Germany, 20:00–00:00 for Spain, and 15:00–19:00 for the United States. Overnight times are defined as 23:00–5:00am, except in Germany, where the managed charging tariff defines overnight off-peak as 00:00–05:00. Daytime hours are all other non-peak and non-evening hours.

B Figures

Figure A1: Tariff Rates in 2024




(a) Managed Charging Tariff



(b) ToU Tariff


Notes: Panel (a) shows the tariff rates for managed charging tariff 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 retailer's standard flat tariff, which maintains a flat rate throughout the day. Panel (b) shows analogous tariff rates for the ToU tariff's off-peak overnight period (00:30–05:30) and peak daytime period (05:30–00:30).

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 A2. 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 C.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 managed charging tariff, which combines static time-of-use pricing with remote control of EV charging. The right panel shows take-up of either the managed charging tariff or ToU tariff, 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: Completion Rate for Participants Signing Up for Managed Charging Tariff




Notes: This figure shows the percentage of participants who, after attempting to sign up for the managed charging tariff during the incentive period, successfully completed the onboarding process. Completion of onboarding includes testing the compatibility of their vehicle or charger with the managed charging tariff app.

Figure A6: 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 5. The shaded area denotes the 95% confidence interval for the treatment effect estimates.

Figure A7: Load Duration Curves




Notes: This figure shows load duration curves for 2024 and the forecasted load duration curve for 2050. The unmanaged charging consumption profile is based on demand modelling data from the UK National Energy System Operator's [Future Energy Scenarios](#). To estimate the impact of managed charging in 2024, we applied the hourly coefficients from our IV estimates (Figure 5) to the corresponding hourly load values. For 2050, we assumed that managed charging would primarily affect the four daily peak hours. We therefore applied the estimated effects from the four peak hours in Figure 5 to the corresponding top four hours in each day of the 2050 load profile. The resulting reductions in peak consumption were assumed to be evenly redistributed across the remaining hours of the day.

Figure A8: Difference Between Unmanaged and Managed Load Duration Curves at Selected Hours




Notes: The bars show the difference in load (unmanaged minus managed) at several positions along the load-duration curve in Fig. A7. Values are taken at load hours 1, 25, 50, 75, and 100 for the years 2024 and 2050.

Figure A9: Customers' Preferences Over Time




Notes: This figure summarizes how frequently managed charging tariff customers change their preferences. Panel (a) shows what proportion of customers changed their preferences in the weeks after adopting the managed charging tariff; (b) displays the CDF of the number of times customers change their preferences, split by week since adoption of the managed charging tariff

Figure A10: Participant Preferences and Behaviors




Notes: This figure summarizes managed charging tariff 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 retailer app. Panel (b) displays the frequency of plug-in and unplug times.

Figure A11: 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 A12: 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 A13: Distribution of Proportion of Charge-Hours Bumped Per Customer




Notes: This figure illustrates bump charging behavior, where trial participants manually overrode the 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 A14: Bump Charge Behaviors - By Encouragement Group




Notes: This figure illustrates bump charging behavior, where trial participants manually overrode the 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 A15: Electricity consumption by hour-of-day on EV tariff




Notes: The figure plots hourly electricity consumption as a fraction of total daily consumption for customers on the managed charging tariff (purple) and ToU tariff (blue) over the trial period, estimated from 2,131 managed charging tariff customers and 889 ToU tariff customers. Shaded regions denote 95% confidence intervals. The ToU tariff is a time-of-use EV tariff without managed charging, available to customers whose hardware is incompatible with the managed charging tariff or who did not wish to enroll in automation. The managed charging tariff off-peak window runs 23:30–05:30; the ToU tariff off-peak window runs 00:30–05:30.

Figure A16: Payments to Octopus from Capacity Market Participation, by Delivery Year




Notes: The figure shows estimated gross Capacity Market agreement values for Octopus-linked CMUs by delivery year. Octopus-linked CMUs are identified from the [NESO CMU register](#) using applicant, registered-holder, and parent-company fields. Agreement values equal awarded capacity obligations (MW) times 1,000 times the relevant auction [clearing price](#) (£/kW-year), and are assigned to the delivery year of the obligation. Values are gross of penalties, secondary trading, indexation, and pass-through arrangements, and should not be interpreted as Octopus-retained profit.

Figure A17: Monthly accepted MWh




Notes: Monthly accepted Balancing Mechanism volumes from Octopus EV-fleet Balancing Mechanism Units (BMUs). The figure aggregates [bid-offer acceptances](#) for 13 EV-fleet BMUs by settlement month. The average monthly accepted volume over October 2024 to March 2026 was 34 MWh, £2,650/month in cashflow.

Figure A18: Daily electricity consumption, random sample of managed charging tariff adopters




Notes: This figure plots average daily electricity consumption by week relative to first adoption of the managed charging tariff. The sample consists of 10,000 randomly selected accounts that first adopted the managed charging tariff in 2023. Shaded bands report 95 percent confidence intervals clustered at the account level. The dashed vertical line marks the week of managed charging tariff adoption.

Figure A19: Comparison of Baseline Consumption Profiles




Notes: This figure shows baseline hourly electricity consumption patterns for customers included in the DiD and RCT analyses. For the DiD sample, baseline refers to consumption prior to adopting the managed charging tariff. For the RCT sample, baseline refers to control group participants who were not on an EV tariff. "Flat" indicates customers who were on a flat tariff before switching to the managed charging tariff, while "ToU" refers to those previously on a time-of-use tariff. The average daily consumption during the baseline period was 1.1 kWh for DiD–Flat customers, 1.02 kWh for DiD–ToU customers, and 1.0 kWh for the RCT control group. The shaded areas represent 95% confidence intervals, with errors clustered at the account user level.

Figure A20: Difference-in-Differences Estimate of Managed Charging Tariff, by Hour-of-day




Notes: This figure reports effect of adopting the managed charging tariff on hourly electricity consumption (in kWh), split by hour of the day. These are estimated using a sample of 20,248 customers who first-ever enrolled in the managed charging tariff 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 A21: Cohort Specific Difference-in-Differences Estimate of Managed Charging Tariff




Notes: This figure reports cohort-specific estimates of the effect of adopting the managed charging tariff on hourly electricity consumption (in kWh) during (a) peak hours (16:30-20:30) and (b) off-peak hours (23:30-05:30). These are estimated using a sample of 9,303 customers who first-ever enrolled in the managed charging tariff in 2024, 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 A22: Marginal value of public funds of subsidizing 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 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 CO₂e benefits reflect emissions reductions from shifting electricity use to cleaner hours, scaled to marginal adopters. Indirect CO₂e 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.

Figure A23: Marginal value of public funds of subsidizing managed charging, Alternative SCC




(a) 10 years, SCC from Bilal and Känzig (2024)



(b) 10 years, SCC from Interagency Working Group

Notes: This figure presents the estimated costs and benefits of subsidizing adoption of the managed EV charging tariff over the course of 10 years. Panel (a) uses the SCC estimate from Bilal and Känzig (2024); panel (b) uses the SCC estimate from [Interagency Working Group on Social Cost of Greenhouse Gases, United States Government](#). 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 CO₂e benefits reflect emissions reductions from shifting electricity use to cleaner hours, scaled to marginal adopters. Indirect CO₂e 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.

Figure A24: Differences in Plug-in and Charging Behavior Between DiD and RCT Managed Charging Tariff Customers



Notes: Panel (a) plots the difference in plug-in rates between DiD and RCT managed charging tariff customers by hour of day. Panel (b) plots the difference in the share of charging kWh occurring in each hour. Positive values indicate higher plug-in availability or charging shares among DiD customers relative to RCT customers. Shaded bands report 95 percent confidence intervals clustered at the account level.




Figure A25: Estimated Impact of Managed Charging Tariff, by "Bump" (Override) Behavior

Notes: This figure reports staggered difference-in-differences estimates of the effect of adopting the managed charging tariff on hourly electricity consumption, for participants who were a part of our experiment. We report separate estimates for customers who (i) have "bumped", or overridden the supplier managed schedule, and (2) who have never bumped. Estimates are also separately estimated by peak (16:30-20:30) and off-peak hours (23:30-05:30). Estimates are computed using the Callaway and Sant'Anna (2021) estimator. Error bars represent the 95% confidence interval.




Figure A26: Estimated Impact of Managed Charging Tariff, by User Preferences

Notes: This figure reports staggered difference-in-differences estimates of the effect of adopting the managed charging tariff on hourly electricity consumption, for participants who were a part of our experiment. We report separate estimates for customers by their settings for when they need the car ready and how much charge is needed: (1) default or less ambitious - 8:00 a.m. ready-by time and 80% charge, or later/lower, (2) ambitious — either an earlier ready-by time or higher charge, and (3) most ambitious — both earlier and higher charge. Estimates are also separately estimated by peak (16:30-20:30) and off-peak hours (23:30–05:30). Estimates are computed using the Callaway and Sant’Anna (2021) estimator. Error bars represent the 95% confidence interval.

C Additional details on design of field trial

C.1 Reproduction of email-based encouragements




Figure A27: 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.”




Figure A28: 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.”




Figure A29: 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.”




Figure A30: 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.”

C.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., the managed charging tariff) 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 managed charging tariff 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 the tariff, 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.

D Empirical results from natural field experiment

This section presents empirical findings from our field experiment. We begin by documenting the impact of the encouragements on trial participant enrollment into the AI managed EV charging tariff. We then examine how these changes in enrollment influenced electricity consumption patterns, using both reduced-form and instrumental variable approaches. We followed our pre-analysis plan, but state where we added or deviated from it, why we did, and their implications for interpretation in Appendix F.

D.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 managed charging tariff in week t :

$$D_{it}^{MC} = \pi_0 + \pi_1 \mathbf{Z}_i + \pi_2 P_{it} + \pi_3 (\mathbf{Z}_i \times P_{it}) + \mu_b + \mu_t + \epsilon_{it} \quad (6)$$

Here, D_{it}^{MC} is a binary indicator equal to one if trial participant i is on a managed charging tariff contract in week t . $\mathbf{Z}_i$ is a set of four binary indicators capturing the encouragement 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 A3. During the 90-day incentive window, the £0/Month and £5/Month groups increased the probability of take-up by approximately 3.4 percentage points and 3.3 percentage points, while the £50/Month and £50/Month (No Bump) groups nearly doubled that effect, reaching 5.9 percentage points 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 3.6%, compared to 6.9% and 6.7% 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

9.5% and 8.7%.⁶¹ We calculated a price elasticity of -0.159 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. 83% of customers who adopted the managed charging tariff during the incentive period remained on the tariff 12 months later. We formally test this in Table A3, interacting the encouragement indicators and the post-incentive period indicator in Table A3. 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) group experienced a statistically significant drop of 1.6 percentage points after incentives ended. These results can be further visualized in Figure A31, which shows estimates of Equation (6), 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.

Importantly, take-up is mechanically constrained by compatibility: trial participants need either their charger or EV to be supported by the managed charging tariff 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.⁶² While the tariff is compatible with some of the most popular home chargers (notably Ohme, MyEnergi/Zappi, and Hyper-volt) 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

⁶¹Tariff 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 A4, and find extremely similar results to Table A3.

⁶²In 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 the managed charging tariff, 17 indicated that they had not done so due to device incompatibility. We also show in Figure A5 the completion rate for participants who started signing up for the tariff but did not complete onboarding. 23% did not complete onboarding, and this appeared balanced across encouragement groups. These individuals were not included in our take-up rate. This does not represent the overall incompatibility rate, for two reasons. First, many customers likely checked whether their device was compatible before starting sign-up, since Octopus provides a compatibility survey on the sign-up page. Second, there may also be other reasons beyond compatibility for not completing onboarding, such as hassle factors associated with onboarding.

Figure A31: Impact of Encouragements on Take-up of Managed Charging Tariff



Notes: This figure plots the intention-to-treat effects of four email-based encouragements to adopt a managed charging tariff, 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.

were compatible with the managed charging tariff.^{63,64}

As the managed charging tariff is not yet compatible for all customers, Octopus Energy offers an alternative EV tariff, the ToU tariff described in Section 3.2. As a result of our email-based encouragements, we observed a nontrivial increase in take-up of the ToU tariff; our speculation is that this was driven by customers who owned devices incompatible with the managed charging tariff. Compared to the effects on take-up of the managed charging tariff, the impacts on ToU tariff 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, although the effects in the latter two groups are not statistically significant (Table A6). 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 the managed charging tariff on electricity

⁶³See this link for compatibility with [charge points](#). The tariff is also compatible with several major [EV brands](#).

⁶⁴To 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 managed charging tariff take-up can thus be a result of not receiving or opening the encouragement email.

consumption, which we will discuss more in Section 4.2.⁶⁵

Additionally, we tested for selection on levels – that is, whether adoption of the managed charging tariff could be explained by observable baseline characteristics, and in particular whether structural winners (those with higher expected bill savings) were more likely to enroll. We estimated a logit regression of tariff take-up on encouragement assignment, expected structural winnings, and baseline covariates. We estimated four specifications that progressively increased flexibility: (1) a baseline model including only the incentive and expected winnings; (2) an expanded model adding additional covariates; (3) a model allowing for interactions among covariates; and (4) a final specification incorporating nonparametric controls for expected savings.

We found no evidence of selection on levels: the coefficient on expected winnings was small and statistically insignificant across all specifications. The results, presented as marginal effects at the means of the covariates in Table A7, suggest that customers with greater expected financial savings were not systematically more likely to adopt. There is some evidence of selection on socioeconomic status – households in the second and third terciles of the Index of Multiple Deprivation were more likely to enroll than those in the most deprived tercile. In addition, customers who were already on a time-of-use tariff prior to the trial were more likely to take up the managed charging tariff.

The explanatory power of these models is extremely low: the squared correlation coefficients from the propensity-score estimates are near zero, and the estimated propensity scores themselves occupy a narrow range. This limited variation implies that most of the heterogeneity in managed charging tariff take-up arose from unobserved factors rather than observable characteristics. Moreover, the propensity scores produced by the four models are not highly correlated with one another, underscoring the instability of the selection equations. As a result, we were unable to estimate marginal treatment effects to further understand selection on slope and heterogeneous treatment effects.

D.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

⁶⁵The managed charging tariff 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.

the following specification:

$$Y_{iht} = \alpha + \beta \mathbf{Z}_i + \gamma X_{ih} + \psi_b + \psi_t + \varepsilon_{iht} \quad (7)$$

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 ,⁶⁶ and ψ_b and ψ_t are fixed effects for block and week, respectively. Standard errors were clustered by user and by week (Colin Cameron and Miller, 2015). To obtain hour-specific effects, we estimated the model separately for each hour of the day. To analyze broader periods (i.e., peak vs. off-peak), we grouped hours into the relevant period and ran the regression on those aggregates.

In Figure A32, we present the 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 receiving any encouragement increased off-peak consumption by 2% (0.020 kWh) and reduced peak consumption by 2% (-0.026 kWh), with no overall effect on total usage (Table A5).


D.3 Heterogeneity

We explored heterogeneity in treatment effects across consumers by interacting the encouragement treatment with three key variables: (1) the Index of Multiple Deprivation (IMD), (2) property value of their home, and (3) baseline electricity consumption. In these specifications, we treat interaction terms (e.g., EV tariff $\times$ 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 (e.g., crime, housing barriers, health) for small geographic areas, weighted to produce

⁶⁶This control variable was not specified in our pre-analysis plan, but we included it to enhance the precision of our estimates.

Figure A32: 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 (7), 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 managed charging tariff 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 the managed charging tariff’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 A5 and defined in Equation (7).

an overall deprivation score. For this analysis, we constructed IMD terciles by linking trial participants’ meter-point-level postcodes to area-level IMD ranks.⁶⁷ 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, such as lack of off-street parking or neighborhood safety concerns, may prevent uptake. Understanding these dynamics requires analyzing effects across the socioeconomic gradient.

We used property value as an additional indicator of socioeconomic status to complement our location-based Index of Multiple Deprivation (IMD) measure. Property value data were obtained from WhenFresh, a provider that aggregated daily listings from

⁶⁷We discretize the IMD into terciles to avoid very small cell sizes in the most deprived group.

Zoopla, a major UK property portal. Each property was associated with a Unique Property Reference Number (UPRN), a persistent identifier that remained constant throughout the property’s lifetime. Using customers’ UPRNs, we matched each record to the most recent property value available in the WhenFresh dataset. This matching procedure yielded property value information for 98% of the sample. We used the full WhenFresh dataset (29,783,591 properties) to establish national cutoffs for property-value terciles. Thus, the heterogeneity bins in our analysis reflected thresholds derived from the national property distribution rather than from our sample alone.


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, shaping expectations about scalability and future generalizability.

We observed mixed evidence of heterogeneity in take-up. Adoption of the managed charging tariff was higher in IMD terciles 2 and 3 (Figure A33a), but we did not observe the same pattern in looking at heterogeneity by property values (Figure A34). Take-up also declined with baseline electricity use: households with higher pre-trial consumption were less likely to adopt, particularly under the £50/Month (No Bump) condition (Figure A35a). This pattern suggests that higher-consuming households may be more reluctant to accept restrictions on charging, perhaps due to greater perceived disruption to their routines.

Turning to impacts on electricity consumption, we found limited evidence of heterogeneous impacts. Across the three dimensions we considered, there was no evidence of differences in off-peak effects. For peak hours, the largest reductions occurred in the middle IMD tercile (Figure A33b) and the highest property value tercile (Figure A34b). However, these results should be interpreted with caution given the limited precision of subgroup estimates: only 9% of our sample resided in the most deprived tercile and 7% occupied properties in the lowest tercile of the national property-value distribution. This reflects the strong association between EV ownership and higher socioeconomic status (Figure A33c, Figure A34c). The small sample size in this group restricted our ability

to detect precise effects among more deprived households. By contrast, baseline electricity use showed little role in shaping treatment impacts. Conditional on adoption, consumption effects are broadly similar across all terciles of pre-trial consumption (Figure A35b).⁶⁸


Figure A33: 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 the managed charging tariff, interacting encouragement with deprivation tercile (T1: most deprived); regression results are also presented in Table A11. Panel B displays the estimated IV effects of EV tariff adoption on hourly electricity consumption during off-peak (23:30–05:30, green) and peak (16:30–20:30, blue) periods, again interacting EV tariff adoption with deprivation tercile; regression results are also presented in Table A12. 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.


⁶⁸We also looked at heterogeneity by weekday versus weekend, but found demand impacts to be similar regardless of the day of the week.

Figure A34: Heterogeneity by Property Value



Notes: This figure presents heterogeneity in treatment effects by property value terciles. Panel A shows the effect of each encouragement group on take-up of the managed charging tariff, interacting encouragement with property value tercile (T1: most deprived); regression results are also presented in Table A11. Panel B displays the estimated IV effects of EV tariff adoption on hourly electricity consumption during off-peak (23:30–05:30, green) and peak (16:30–20:30, blue) periods, again interacting EV tariff adoption with property value tercile; regression results are also presented in Table A12. 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 A35: 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 the managed charging tariff, interacting encouragement with deprivation tercile (T1: lowest consumption); regression results are also presented in Table A11. Panel B displays the estimated IV effects of EV tariff adoption on hourly electricity consumption during off-peak (23:30–05:30, 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 A12. 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.

D.4 Generalizability to other countries

Octopus currently offers the managed charging tariff in seven countries, of which four — Germany, Spain, the United States (Texas), and the United Kingdom — have a substantial number of customers. Exact customer counts are commercially sensitive and cannot be disclosed, but all four markets have sufficient adoption to allow for meaningful comparisons of consumer behavior. Details on how the tariff is structured in each market are presented in Table [A18](#).

To assess whether treatment effects are likely to hold beyond the UK setting, we present descriptive statistics on customer behavior and charging demand patterns under the managed charging tariff in other countries. We drew random samples of 1,000 customers in 2024 and 1,000 customers in 2022 (early-adopters) in the UK. Germany, Spain, and the US had substantially smaller customer bases; for these markets, we used all customers who were on the tariff by the end of 2024.⁶⁹ We then recalibrated our theoretical model to incorporate the behavioral patterns observed in the data, to assess the welfare margins of RTP versus managed charging.

We examined three key EV charging dimensions in order to compare managed charging tariff users across countries. First, plug-in behavior, referring to how often and when vehicles are connected to be ready to charge. Second, override (“bump”) behavior, which measures how frequently users intervene in supplier managed schedules. Third, consumption profile throughout the day. Together, they show how the behavioral foundations of managed charging vary across markets.

We present these three dimensions for the four countries: UK, Germany, Spain, and United States. In addition, we also show the descriptive statistics for early adopters in the UK, as defined by those who took up the managed charging tariff in the first six months since it was rolled out. The UK has by far the most widespread usage of the tariff, with about two orders of magnitude more customers than the other countries we analyze; we therefore present statistics of early adopters, which could be more comparable to the customers from other countries.


Across countries, UK customers were plugged in the least overall, and exhibited the strongest variation in plug-in rates over the day (Figure [A36](#)). In contrast, customers in other countries kept their vehicles connected for a greater share of the day. This may be a pattern associated with earlier adopters rather than with the country itself; early adopters

⁶⁹Per-country sample sizes are withheld as commercially sensitive.

in the UK had higher plug-in rates than later UK adopters. The lower daytime plug-in rates among all but the earliest adopters in the UK suggest that customers there provided the AI with fewer opportunities to optimize charging, while more constant plug-in availability likely enabled greater algorithmic flexibility in Germany, Spain, and the United States.


Overrides (“bump charging”) were more common outside the UK, although still infrequent overall. Many customers across all countries never used the bump function at all: 55% in the UK, 43% in Germany, 43% in Spain, and 56% in the United States (Figure A37). The highest share of bump-charged electricity occurred in Germany, where 3.3% of EV electricity consumption resulted from overrides, and there was an 10.7% probability of bumping for each charge-day (Table A19). The patterns suggest a behavioral substitution: customers who plug in less frequently also tend to intervene less often, whereas those who leave vehicles connected more consistently appear more likely to occasionally override the AI schedule, though note that we do not have causal evidence to support these hypotheses.

Figure A36: Plug-in Rate By Hour-of-Day



Notes: This figure plots, by hour of day, the percentage of active users who had their EVs plugged in. A user is considered active during a given hour if that hour falls between their first-ever and last-ever recorded plug-in event. We use a sample of 4,442 managed charging tariff users across the four countries.

Figure A37: Proportion of Consumption Bump Charged 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 consumption was through overriding the schedule. The vertical axis shows the cumulative proportion of customers. We use a sample of 4,442 managed charging tariff users across the four countries.

Despite higher rates of overriding outside the UK, the vast majority of charging still occurred during overnight periods, as shown in Figure A38. Using EV charging data, the figure plots average hourly EV consumption alongside average household electricity use for a random sample of non-EV households in each country. Across all settings, managed charging tariff users displayed a pronounced overnight spike in charging — around 0.4 kWh per hour — indicating that most charging remains concentrated in off-peak night-time hours.

In the UK field experiment, adoption of the managed charging tariff shifted roughly 0.5 kWh per hour of consumption from peak to off-peak periods (i.e., our 42% reduction in peak consumption observed in Figure 5). The magnitude of the overnight 0.4 kWh spike observed across all four countries is broadly consistent with this shift, implying a similar volume of flexible load from EV charging. This consistency suggests that the consumption response identified experimentally in the UK is likely to be similar across other markets.

In this section, we observed higher and more consistent plug-in rates in non-UK markets (??). As shown in Section 4.3, plug-in behavior is an important aspect of customer flexibility that influences when load shifting can occur. Thus, the observed plug-in patterns abroad align with the potential for similar flexibility conditions outside the UK. Taken together with the charging activity shown in Figure A38, these descriptive results

Figure A38: Cross-Country Consumption Profiles



Notes: The figure plots hourly electricity consumption for a random sample of customers across managed charging tariff countries. Blue lines show EV charging consumption, based on telemetry data from managed charging tariff users, while light purple lines show household consumption among non-EV households. All panels use 2024 data, except Panel 2, which presents 2023 consumption data for customers who adopted the managed charging tariff during its initial rollout period (September 2022 – March 2023).

indicate that flexibility gains associated with AI managed charging could plausibly extend across countries.

D.5 Variance decomposition of bump behavior

We conducted a variance decomposition to quantify the relative contributions of different factors to the overall variability in bump charging behavior. Using hourly charging data, we regressed an indicator for whether charging occurs via bump charging on a set of fixed effects and temperature controls. Specifically, we estimated the following specification:

$$I(\text{Bump}_{iht}) = \alpha_i + \mu_t + \gamma_h + \psi_r + \sum_k \beta_k \text{Temp}_{i,t+1}^k + \varepsilon_{iht}, \quad (8)$$

where $I(\text{Bump}_{iht})$ is an indicator equal to one if customer i engaged in bump charging during hour h on date t . The term α_i denotes customer fixed effects, μ_t date fixed effects, γ_h hour-of-day fixed effects, and ψ_r region fixed effects corresponding to the 14 UK Grid Supply Point (GSP) regions. The temperature controls, $\text{Temp}_{i,t+1}^k$, are indicators for ten bins of next-day temperature, included to capture anticipatory charging responses to expected weather conditions. For each customer, temperature was assigned using observations from the geographically closest UK Met Office weather station to the customer's postcode. The UK Met Office is the national meteorological service of the United Kingdom. We run this regression for our experimental sample, as well as the random sample of UK managed charging tariff users, described in Section 4.4.

The variance decomposition can be expressed as:

$$\text{Var}(Y_{it}) = \underbrace{\text{Var}(\alpha_i)}_{\text{Household FE}} + \underbrace{\text{Var}(\mu_t)}_{\text{Date FE}} + \underbrace{\text{Var}(\gamma_h)}_{\text{Hour FE}} + \underbrace{\text{Var}(\psi_r)}_{\text{Region FE}} \quad (9)$$

$$+ \underbrace{\text{Var}(\beta \cdot \text{Temp}_{i,t+1})}_{\text{Temperature effect}} + \underbrace{2 \cdot \text{Cov}(\cdot, \cdot)}_{\text{Covariance terms}} + \underbrace{\text{Var}(\varepsilon_{it})}_{\text{Residual}} \quad (10)$$

As shown in Table A20, customer fixed effects accounted for the vast majority of the variation in bump charging behavior, explaining 0.274 of the total variance in the RCT sample. Hour-of-day and region fixed effects contributed an additional 0.088. This dominance of customer fixed effects indicates that bump charging is largely idiosyncratic rather than correlated, making it unlikely to generate coincident demand spikes that would strain the grid. The UK managed charging tariff adopters sample show similar results, with customer fixed effects as the main source of variation.

Table A20: Decomposition of Variance in Bump Behavior

Component	Share (RCT Sample)	Share (UK Sample)
Customer FE	0.274	0.296
Date FE	0.003	0.001
Hour FE	0.060	0.008
Region (GSP)	0.031	0.001
Next day temperature	0.000	0.000

Note: This table reports the variance decomposition of bump charging behavior based on the regression specification in Equation (9). Each entry shows the share of total variance in the bump charging indicator explained by the corresponding component.

E 2024 difference-in-differences

As discussed in Section 6.1, we also conducted a DiD analysis focusing on customers who adopted the managed charging tariff in 2024. However, identifying which customers already owned an EV at the start of 2024 was difficult. This challenge likely reflected increased uptake of low-carbon technologies (LCTs). Of particular note, heat pump installations rose sharply at the end of 2023, following the UK Government’s expansion of heat pump subsidies in October 2023.⁷⁰ This increase in LCT ownership likely produced additional consumption spikes resembling EV charging, but originating from other devices. Lacking direct observation of LCT ownership, we cannot be sure whether changes in consumption among customers adopting the managed charging tariff in 2024 were driven by tariff adoption, adoption of EVs, or adoption of other LCT devices that encouraged more attention paid to which tariff they were on.⁷¹

To mitigate this issue, we restricted the sample to customers who appeared to own an EV as of August 2023. We chose August 2023 because it was both (1) before the increase in heat pump subsidies and (2) during summer months when heat pump use was minimal. Nonetheless, adoption of other LCTs may still have confounded our estimates by increasing customers’ engagement with tariff selection. This confounder was documented by Bernard et al. (2024), who examined households that received heat pump installations from Octopus Energy. They found that following the installation, two-thirds of these households adopted a smart tariff, with the managed charging tariff being the most popular choice, possibly due to adoption of an EV at a similar time.⁷² We presented

⁷⁰Coverage on heat pump takeup can be found in the [BBC](#) and [Guardian](#). Academic research on the electricity consumption impacts of heat pumps is documented in Bernard et al. (2024)

⁷¹This tariff-switching behavior is documented by Bernard et al. (2024), who examined households that received heat pump installations from Octopus Energy. They found that following the installation, two-thirds of these households adopted a smart tariff, with the managed charging tariff being the most popular choice, possibly due to adoption of an EV at a similar time.

⁷²Additionally, in an internal survey, 25% of customers with a heat pump reported being on the managed charging

the results of the 2024 DiD analysis but advised caution in their interpretation.

We began with a sample of 146,143 customers who adopted the managed charging tariff at some point in 2024. To be included in the final sample, customers had to have been with Octopus Energy by August 2023, had smart-meter data available at that time, and likely owned an EV as of August 2023. We excluded 2,039 customers who were already part of our randomized controlled trial. After these restrictions, the final analysis sample consisted of 9,303 customers.


Our empirical strategy mirrored the 2023 DiD analysis. We used "not-yet-treated" customers as controls and defined an anticipation period of four weeks, using the five weeks prior to managed charging tariff adoption as the reference period. We limited control cohorts to those scheduled to adopt the tariff no later than twelve weeks after the treated group's anticipation period ended. Finally, we estimated a weighted version of the model to match the 2024 DiD sample to the RCT sample based on pre-treatment tariff type.

Similar to the 2023 results, Figure [A39](#) showed that the unweighted difference-in-differences estimates were smaller than those from the RCT. After reweighting the 2024 DiD sample to align with the RCT's pre-treatment tariff distribution, the estimated increase in off-peak consumption (0.548 kWh) closely matched the RCT estimate (0.481 kWh). The reduction in peak consumption was smaller in the DiD analysis: 0.269 kWh compared to 0.581 kWh in the RCT.

Diverging from the 2023 analysis, the decline in daytime, non-peak consumption did not fully offset the increase in off-peak consumption, resulting in an overall rise of 1.53 kWh per day, an 8% increase in total consumption. We believe this likely reflected concurrent adoption of other LCTs, such as heat pumps, during the same period. There was tentative evidence of this in the off-peak anticipation period in Figure [A39b](#), although without direct data on LCT ownership, we could not confirm this hypothesis.


tariff; it is therefore plausible that heat pump adoption itself encouraged switching to the tariff.

Figure A39: 2024 Difference-in-Differences Estimate of Managed Charging Tariff



Notes: This figure reports staggered difference-in-differences estimates of the effect of adopting the managed charging 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 9,303 customers who first-ever enrolled in the managed charging tariff in 2024. 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-adoption consumption levels. Post-treatment effects are estimated using average of all group-time average treatment effects, with weights proportional to the group size.

Figure A40: 2024 Difference-in-Differences Estimate of Managed Charging Tariff, by Hour-of-day



Notes: This figure reports effect of adopting the managed charging tariff on hourly electricity consumption (in kWh), split by hour of the day. These are estimated using a sample of 9,303 customers who first-ever enrolled in the managed charging tariff in 2024, 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.

F 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, managed charging 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.

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 $\times$ hour-of-day level.

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 A8, Table A9, and Table A10.

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 A8, Table A9, and Table A10. 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.

Our pre-specified IV analysis planned to use the randomized encouragements as instruments for just managed charging tariff adoption. However, given that the randomized encouragements increased enrollment in both the managed charging tariff and the ToU tariff, we defined treatment as adoption of either EV tariff, rather than the managed charging tariff 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.

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 5, Table A8, Table A9, and Table A10. These include both the joint specification with all four instruments and separate regressions where each encouragement serves as an instrument individually.

Our pre-specified DiD sample included all customers who adopted the managed charging tariff in 2023. In the final analysis, our DiD analysis restricted the sample to cus-

tomers 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 pre-specified). We also added a separate DiD analysis of customers who adopted DiD in 2024, which was not pre-specified. However, this analysis is not a part of our main findings, as detailed in Appendix E.

For our DiD analysis, we prespecified restricting control cohorts to customers who adopt the managed charging tariff within 30 days. However, our final analysis used a twelve-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. Additionally, we added an anticipation period, which was not part of the original specification. This decision followed a review of pre-trends, especially in the 2024 analysis, where we observed a rise in off-peak consumption in the four weeks before managed charging tariff adoption (Figure A39b).

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 CO₂e 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.

We originally pre-specified estimating CO₂e impacts using the ITT framework. We did not pre-specify the use of IV estimation for bills or CO₂e savings. We adopted the IV approach here because it more directly captures the causal effect of managed charging tariff 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.

G Methodology for calculating welfare impacts

This appendix describes how we constructed the welfare measures reported in Section 7. The objective was to measure the private and social surplus generated by adoption of the managed charging tariff, using the randomized encouragement as an instrument for enrollment. The analysis used the daily experimental billing panel, which combines smart-meter consumption, observed tariff charges, retailer procurement costs, and marginal-emissions measures. Let i index customers and d index dates. For each welfare component Y_{id} , we estimated

$$Y_{id} = \alpha_d + \gamma_b + \delta D_{id}^{MC} + \rho X_i + \varepsilon_{id},$$

where α_d are date fixed effects, γ_b are randomization-block fixed effects, D_{id}^{MC} indicates that customer i was enrolled on the managed charging tariff on date d , and X_i is pre-period electricity consumption. Enrollment was instrumented with assignment to any encouragement arm, Z_i . Standard errors were clustered by customer and date. This specification estimated the local average effect of enrollment for customers induced to adopt by the randomized encouragement.

Consumer benefits: We used administrative data on daily average electricity bills per kWh. The outcome was the amount customers paid per kWh of electricity, and we weighted the regression by daily consumption.⁷³ We converted the estimated reduction in unit cost into an annual bill saving by multiplying by annual baseline consumption among pure-control customers who were not enrolled on the managed charging tariff. Let $\widehat{\delta}^{bill}$ denote the estimated effect on the daily unit bill and $\overline{kWh}^C$ denote mean daily consumption in the non-enrolled pure-control group.

Supplier procurement costs: We applied the same IV design to Octopus Energy’s total daily procurement cost. Octopus Energy provided administrative data on half-hourly electricity costs. These costs included wholesale power and non-energy charges such as

⁷³Because managed charging tariff customers typically charged only on a subset of days, an unweighted comparison across all days would place equal weight on charging and non-charging days. The consumption-weighted specification estimated the bill effect for the electricity actually consumed.

transmission and distribution fees, but excluded revenues or costs from grid services, including ancillary-market participation. The wholesale component combined hedged prices, day-ahead prices, intraday prices, and the final system price for the half-hour. The exact weighting of these components may vary over time, but this measure more closely approximated the supplier’s realized procurement cost than the system price alone. Comparing bill savings and procurement savings provided a pass-through calculation from supplier cost savings to customer bill savings. As discussed in 5.4, there are other revenue streams that the managed charging tariff provided for Octopus from participating in other real-time grid-balancing operations and providing grid services; these were not included as the per-vehicle value is proprietary information.

CO₂e impacts : For direct emissions impacts, we multiplied each trial participant’s half-hourly electricity consumption by the corresponding average Marginal Operating Emissions Rate from WattTime, then aggregated to the customer-day level. As described in Section 3.4, the Marginal Operating Emissions Rate measures the emissions associated with a marginal change in electricity load on the grid (WattTime, 2022). Using this, we estimated the IV effect of managed charging tariff enrollment on daily CO₂e emissions from electricity consumption. Let $\widehat{\delta}^{CO_2}$ denote the estimated daily effect, measured in kg CO₂e.

We monetized annual emissions reductions using the UK government’s social cost of carbon, which is approximately £250 per tonne of CO₂e. This is calculated by estimating the marginal abatement cost (i.e., resource costs) per tonne of CO₂. The government estimates the amount of CO₂e that is needed to meet the UK’s future CO₂e targets and walks up the marginal abatement cost curve until it hits that CO₂e target, which hits costs at £250 per tonne of CO₂e. Our methodology was:

$$\widehat{B}_t^{CO_2} = -365 \cdot \widehat{\delta}^{CO_2} \cdot \frac{SCC_t}{1000} \cdot (1 - s_{GDP}^{UK} s_{tax}^{UK}).$$

The division by 1000 converted kg to tonnes. The term $(1 - s_{GDP}^{UK} s_{tax}^{UK})$ removed the fiscal component of global climate damages already internalized through UK tax revenue. We converted monetary values to 2023 prices using the UK GDP deflator. The UK government SCC is based on the marginal abatement cost required to meet domestic emissions targets, rather than the integrated-assessment-model damage estimates used in some other jurisdictions, like the US.

We also estimated a potential indirect emissions channel through EV adoption. Lower charging costs reduce the total cost of EV ownership. We first computed the present value of managed charging tariff electricity bill savings over a ten-year vehicle lifespan, using a 3.5% discount rate recommended by HM Treasury (2020). We then treated this saving as a reduction in the ownership cost of an EV and applied an EV own-price elasticity of -2.547 from Hahn et al. (2024). To translate the proportional adoption response into annual vehicle purchases, we used Department for Transport (2024), which reported that approximately 6.7% of UK households purchased a new car each year.

We valued each induced substitution from an internal-combustion vehicle to an EV using the difference in lifecycle CO₂e damages between ICE vehicles (£8,003.89) and EVs (£3,259.66), based on Hahn et al. (2024). This yielded £4,744.23 in CO₂e benefits per induced switch. We scaled the annual impacts by government projections of the non-EV share of new vehicle sales, which declined over time under the phase-out path in Department for Transport (2023). Because the pool of potential ICE-to-EV substitutions declines as EVs become a larger share of new car sales, we scaled the annual effect by the projected non-EV share of new vehicle sales. Specifically, we used non-EV shares of 78% in 2024, 72% in 2025, 67% in 2026, 62% in 2027, 48% in 2028, 34% in 2029, 20% in 2030, 16% in 2031, 12% in 2032, and 8% in 2033, based on the phase-out path in Department for Transport (2023). This scaling assumed that induced EV purchases displaced ICE purchases only among households that would otherwise have bought a non-EV vehicle.

System benefits: Finally, managed EV charging may avoid broader electricity-system costs. These benefits were not estimated from the randomized experiment. Instead, we reported model-based estimates from Franken et al. (2025), who quantified per-vehicle system savings from smart charging under alternative grid-flexibility scenarios. We presented these grid benefits separately from the experimentally estimated bill, procurement, and emissions components because they came from an external system model rather than from trial variation.

H Theoretical Model

From a welfare analysis perspective, we expect the managed charging tariff to generate distinct welfare outcomes relative to real-time pricing (RTP), depending on the scope for optimization and the behavioral frictions consumers face. The tariff can outperform RTP

when it leverages price signals from multiple markets, such as ancillary service opportunities that are typically excluded from the RTP signal.⁷⁴ Our theoretical model focuses on behavioral frictions rather than multi-market optimization; however, we acknowledge that both are important drivers of the extent to which automation might or might not outperform RTP.

H.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

$$\tilde{c}_h = c'_h(Q_h) - r'_h(Q_h), \quad (11)$$

as in convex scheduling (Boyd and Vandenberghe, 2004; Joskow and Tirole, 2006). 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 managed charging, they can override at a hassle cost $\phi_i \geq 0$.

We compare four tariffs:⁷⁵

1. **Flat:** price p^{flat} .
2. **ToU:** $p^{\text{peak}} > p^{\text{off}}$ across fixed windows.

⁷⁴In principle, RTP could incorporate such signals, but doing so would require a richer and less transparent price vector, as well as greater computational and behavioral demands on consumers. The managed charging tariff may instead yield lower welfare than RTP when intertemporal optimization across days offers larger gains than within-day shifting, or if users systematically fail to plug in their EVs when charging would be most cost-effective. However, evidence from RTP studies suggests that cross-day shifting is rare, and our data show no systematic plug-in failures.

⁷⁵Our framework builds on Joskow and Tirole (2006), 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 as a distinct pricing/coordination regime, and (ii) frictions such as override behavior, execution costs (m_{ih}, ϕ_i), and aggregate override thresholds ($\tilde{\beta}^*$). 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.

3. **RTP**: hourly p_h^{RTP} ; households self-schedule.

4. **Managed charging**: centralized schedule $\{x_{ih}\}$ against $\{\tilde{c}_h\}$; users may override.

The feasible EV-charging schedules satisfy⁷⁶

$$\sum_{h \leq T_i, h \in \mathcal{A}_i} \eta_{ih} x_{ih} \geq E_i, \quad 0 \leq x_{ih} \leq \bar{x}_i. \quad (12)$$

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

$$\epsilon_h^k = \frac{\partial \ln Q_h^k}{\partial \ln \pi_h^k}, \quad (13)$$

evaluated in high plug-in hours ($Q_h^k > 0$). Let e_h denote marginal CO₂ emissions (kg/kWh).

H.2 Household and Aggregator Problems

RTP household problem. A household chooses $\{x_{ih}\}$ to minimize expected cost plus frictions:⁷⁷

$$\begin{aligned} \min_{\{x_{ih}\}} \quad & \mathbb{E} \left[\sum_h p_h^{\text{RTP}} x_{ih} \right] + \gamma_i \text{Var} \left(\sum_h p_h^{\text{RTP}} x_{ih} \right) + \sum_h \psi_i(h) x_{ih} + a_i \cdot \mathbf{1}\{\text{actively scheduling}\} \\ \text{s.t.} \quad & \sum_{h \leq T_i, h \in \mathcal{A}_i} \eta_{ih} x_{ih} \geq E_i, \quad 0 \leq x_{ih} \leq \bar{x}_i. \end{aligned} \quad (14)$$

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

$$\mathbb{E}[p_h^{\text{RTP}}] + 2\gamma_i \text{Cov} \left(p_h^{\text{RTP}}, \sum_{\ell} p_{\ell}^{\text{RTP}} x_{i\ell} \right) + \psi_i(h) = \mu_i \eta_{ih}, \quad (15)$$

⁷⁶We model EV-only RTP for comparability to managed charging, 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.

⁷⁷This functional form is inspired by Borenstein (2007); Gabaix (2019) who modeled variations of an expected expenditure minimization augmented with cognitive and risk terms.

with $\mu_i \geq 0$ the KKT multiplier on the energy requirement. The attention cost $a_i > 0$ drives corner solutions by discouraging small reallocations.⁷⁸

Managed charging aggregator problem. The managed charging scheduler minimizes social cost plus timing disutility:

$$\begin{aligned} \min_{\{x_{ih}\}} \quad & \sum_h \left[c_h(Q_h) - r_h(Q_h) \right] + \sum_{i,h} \psi_i(h) x_{ih} \\ \text{s.t.} \quad & \sum_{h \leq T_i, h \in \mathcal{A}_i} \eta_{ih} x_{ih} \geq E_i, \quad 0 \leq x_{ih} \leq \bar{x}_i. \end{aligned} \quad (16)$$

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

H.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. (16)–Eq. (12) (there exists feasible slack).

⁷⁸Derivation of (15): Rewrite the energy requirement as a KKT-friendly inequality $g_i(x) \equiv E_i - \sum_{h \leq T_i, h \in \mathcal{A}_i} \eta_{ih} x_{ih} \leq 0$ with multiplier $\mu_i \geq 0$. 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 does not depend on $\{x_{ih}\}$ and thus only affects the extensive/corner decision), the Lagrangian is

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

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

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

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

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

which is (15). Complementary slackness for α_{ih}, β_{ih} covers corner hours.

Lemma H.3.1 (Peak shaving under managed charging (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 , the managed charging optimum allocates charging to hours with the lowest available κ_{ih} subject to Eq. (12).* (b) *If there exist feasible peak and off-peak hours p, o with $\kappa_{ip} > \kappa_{io}$, then $x_{ip} > 0$ implies all lower-cost hours $\{o : \kappa_{io} < \kappa_{ip}\}$ are saturated or infeasible. Aggregating over i , managed charging (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

$$\begin{aligned} \mathcal{L}(\{x_{ih}\}, \{\mu_i\}, \{\mu_{ih}^\pm\}) = & \sum_h [c_h(Q_h) - r_h(Q_h)] + \sum_{i,h} \psi_i(h) x_{ih} \\ & + \sum_i \mu_i \left(E_i - \sum_{\substack{h \leq T_i \\ h \in A_i}} \eta_{ih} x_{ih} \right) + \sum_{i,h} \mu_{ih}^+ (x_{ih} - \bar{x}_i) - \sum_{i,h} \mu_{ih}^- x_{ih}. \end{aligned} \quad (17)$$

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 $\mu_{ih}^\pm = 0$ and $\mu_i = \frac{\tilde{c}_h + \psi_i(h)}{\eta_{ih}} \equiv \kappa_{ih}$. If $x_{ih} = 0$ then $\mu_{ih}^- \geq 0$ and $\kappa_{ih} \geq \mu_i$; if $x_{ih} = \bar{x}_i$ then $\mu_{ih}^+ \geq 0$ and $\kappa_{ih} \leq \mu_i$. Hence x_{ih} is nonincreasing in κ_{ih} : the managed charging 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 $\kappa_{ip} \leq \mu_i \leq \kappa_{io}$, a contradiction. Summing over i yields load shifted from high- $\tilde{c}_h$ to low- $\tilde{c}_h$ hours, conserving daily energy by Eq. (12).

Lemma H.3.2 (Elasticity ordering). *In high plug-in hours,*

$$|\epsilon_h^{\text{MC}}| \geq |\epsilon_h^{\text{RTP}}| \geq |\epsilon_h^{\text{ToU}}|.$$

Proof. Step 1 (managed charging upper-bound response). Consider a small perturbation $d\tilde{c}$ concentrated in hour h . Differentiating the KKT system in Lemma H.3.1 yields a linear

system in $\{dx_{ih}, d\mu_i, d\mu_{ih}^\pm\}$ whose solution preserves the merit order: $dx_{ih} \leq 0$ at the perturbed hour and $dx_{i\ell} \geq 0$ at some lower- $\kappa_{i\ell}$ hours, with $\sum_{\ell \leq T_i, \ell \in A_i} \eta_{i\ell} dx_{i\ell} = 0$ (intra-day reallocation). Aggregating over i ,

$$\frac{\partial Q_h^{\text{MC}}}{\partial \tilde{c}_h} \leq 0, \quad \left| \frac{\partial Q_h^{\text{MC}}}{\partial \tilde{c}_h} \right| \text{ is maximal subject to Eq. (12).}$$

Thus $|\epsilon_h^{\text{MC}}|$ is an upper bound among feasible reallocations.

Step 2 (RTP attenuation by risk/attention). From the household FOC Eq. (15), totally differentiate across hours. Stacking in vector form,

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

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:

$$\left\| \frac{\partial Q^{\text{RTP}}}{\partial p} \right\| \leq \left\| \frac{\partial Q^{\text{MC}}}{\partial \tilde{c}} \right\|.$$

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, $d\pi^{\text{ToU}} = P d\pi^{\text{RTP}}$, hence

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

so $|\partial Q_h^{\text{ToU}} / \partial p_h| \leq |\partial Q_h^{\text{RTP}} / \partial p_h|$. Combining Steps 1–3 yields the stated ordering.⁷⁹

⁷⁹We do not analyze a bill-variance ordering in our empirical set-up, but with $\text{Bill}^k = \sum_h p_h^k X_h^k$, the law of total variance gives

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

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 managed charging (EV sub-load), consumers face a flat off-peak retail rate; wholesale volatility is internalized by the aggregator, yielding

$$\text{Var}(\text{Bill}^{\text{RTP}}) \geq \text{Var}(\text{Bill}^{\text{ToU}}) \geq \text{Var}(\text{Bill}^{\text{MC}}).$$

Lemma H.3.3 (Welfare ranking with frictions). *With $a_i, \gamma_i > 0$,*

$$W^{\text{MC}} \geq W^{\text{RTP}} \geq W^{\text{ToU}} \geq W^{\text{Flat}}.$$

Proof. Let W^k be per-EV social welfare (net of pure transfers). Managed charging maximizes W for the EV sub-load given $\{\tilde{c}_h\}$ (Lemma H.3.1 and KKT optimality). Relative to managed charging,

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

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

Proposition H.3.1 (Emissions ordering). *If e_h is (weakly) lower off-peak, then relative to baseline: managed charging achieves the largest emissions reduction, followed by RTP, then ToU, then Flat.*

Proof. By Lemma H.3.1, managed charging 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.

H.4 Overrides (“Bump Charging”): Behavior and Welfare

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

$$v_{ih} > m_{ih} + \phi_i, \tag{18}$$

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

$$\beta_{ih} = \Pr[v_{ih} > m_{ih} + \phi_i] = 1 - F_{ih}(m_{ih} + \phi_i). \tag{19}$$

Lemma H.4.1 (Override monotonicity). *If F_{ih} is nondecreasing, then β_{ih} is weakly decreasing in ϕ_i and in m_{ih} .*

Proof. Differentiate Eq. (19) (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 H.4.2 (Welfare effect of an override). *Let h' be the managed charging 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. (18). Social part: the managed charging 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 H.4.1 (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 managed charging planned charging in h , and if ρ is the share of hours planned to charge, then*

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

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

Proof. Law of total probability conditioning on planned status; apply Eq. (19) in each state and average.

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

$$\Delta W_0 = W^{\text{MC}} - W^{\text{RTP}}. \quad (20)$$

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

$$W^{\text{MC+O}} = W^{\text{MC}} - \bar{\beta} \lambda. \quad (21)$$

We calibrate λ by

$$\hat{\lambda} \approx p_{\text{peak}} \cdot \Delta \tilde{c}_{\text{eff}} \cdot q^O, \quad \lambda^{\text{UB}} = (\max_h \tilde{c}_h - \min_h \tilde{c}_h) \cdot q_{\text{max}}^O, \quad (22)$$

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 H.4.3 (Managed charging-RTP crossover threshold). *If $\Delta W_0 > 0$ and $\lambda > 0$, the*

override rate at which managed charging with overrides equals RTP is

$$\bar{\beta}^* = \frac{\Delta W_0}{\lambda}. \quad (23)$$

Proof. Equate Eq. (21) to $W^{\text{RTP}} = W^{\text{MC}} - \Delta W_0$ and solve for $\bar{\beta}$.

Proposition H.4.2 (Welfare ordering with overrides). *If $\bar{\beta} < \bar{\beta}^*$, then $W^{\text{MC+O}} > W^{\text{RTP}}$; if $\bar{\beta} > \bar{\beta}^*$, RTP can dominate.*

Proof. From Eq. (21) and Lemma H.4.3, the sign of $W^{\text{MC+O}} - W^{\text{RTP}} = \Delta W_0 - \bar{\beta}\lambda$ is determined by $\bar{\beta}$ relative to $\bar{\beta}^*$.

H.5 Cross-country calibration and crossover analysis

The experiment (managed charging vs. alternatives) identifies (i) peak-to-off-peak reallocation magnitudes (Lemma H.3.1); (ii) override frequencies $\bar{\beta}$ and their timing relative to peak/off-peak (for $\hat{\lambda}$ via Eq. (22)); 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 H.4.3.

H.5.1 Objective and overview

This section develops a cross-country calibration of the *override–welfare* relationship in managed charging systems. The goal is to identify, for each major managed charging market, the *crossover override rate* – that is, the frequency of user overrides that would eliminate managed charging’s welfare advantage over real-time pricing (RTP).

The analysis combines three empirical components: (1) the observed probability of an override per charge day, from a random sample of users across the United Kingdom, Germany, Spain, and the United States; (2) the conditional peak-landing probability $p_{\text{peak}} = \Pr(\text{peak bump/day})/\Pr(\text{bump/day})$; and (3) country-level parameters describing tariff spreads and average energy per override q^O (kWh per bump event), drawn from managed charging tariff pricing data.

H.5.2 Estimating ΔW_0 for the United Kingdom

This appendix documents how we estimate the baseline welfare gap $\Delta W_0^{(\text{UK})}$ – the per-EV daily welfare advantage of managed charging relative to RTP in the absence of overrides. The welfare gap is defined as the expected daily difference in total surplus between managed charging and RTP regimes:

$$\Delta W_0 = [U^{\text{MC}} - C^{\text{MC}}] - [U^{\text{RTP}} - C^{\text{RTP}}],$$

where U^k is the household mobility utility net of timing disutility and C^k the expected system procurement cost. Under RTP, households face optimization and attention frictions that reduce responsiveness; under managed charging, scheduling is automated, internalizing system prices without these frictions.

For the UK, we set the representative elasticity of RTP demand at $\varepsilon = -0.2$ and assume a daily attention cost of $a = 0.20$ per day. Using observed load and tariff data (off-peak £0.07/kWh, peak £0.27/kWh; spread $\Delta \tilde{c}_{\text{eff}} = 0.20/\text{kWh}$), we simulate hourly charging schedules under both managed charging and RTP. Automated coordination real-locates approximately 0.375 kWh of energy per day from peak to off-peak hours, yielding procurement-cost savings of about $0.20/\text{kWh} \times 0.375 = 0.075$ per day. In addition, managed charging reduces timing disutility and bill-variance penalties by a further £0.020, as households experience fewer inconvenient charging hours and lower price volatility exposure. Together these components yield: $\Delta W_0^{(\text{UK})} = 0.075 + 0.020 = 0.095$ per day.⁸⁰

H.5.3 Empirical calibration of override behavior

Table A19 reports observed override behavior across countries. The United Kingdom exhibits a low overall bump frequency (1.7% per charge day) and correspondingly low peak-landing probability (0.3%), yielding $p_{\text{peak}} = 0.176$. Early adopters show slightly higher values ($p_{\text{peak}} = 0.23$). In Germany and Spain, users override much more frequently (10% per day), while in the United States roughly one-third of overrides fall in peak windows. Average energy per override q^O ranges from 1.25 kWh in Spain to around 2.4 kWh in the UK and US (see Table A21 below). Tariff spreads $\Delta \tilde{c}_{\text{eff}}$ vary widely, from

⁸⁰For the remaining countries, we scale ΔW_0 proportionally to the observed effective price spread $\Delta \tilde{c}_{\text{eff}}^{(c)}$, following $\Delta W_0^{(c)} = \Delta W_0^{(\text{UK})} \times (\Delta \tilde{c}_{\text{eff}}^{(c)} / \Delta \tilde{c}_{\text{eff}}^{(\text{UK})})$. This assumes similar elasticities and flexible load shares across markets, attributing cross-country differences in welfare to the strength of the intertemporal tariff gradient. Empirically, this yields $\Delta W_0^{(\text{DE})} \approx 0.067$, $\Delta W_0^{(\text{ES})} \approx 0.032$, and $\Delta W_0^{(\text{US})} \approx 0.021$.

£0.17/kWh in the UK to only £0.037/kWh in the US. These differences strongly affect the expected welfare loss per event.

Table A21: Tariff and override characteristics by country

Country	Off-peak rate	Standard rate	q^O (kWh)	$\Delta\tilde{c}_{\text{eff}}$ (£/kWh)	% bumped in peak
UK	£0.24	£0.07	2.41	0.17	0.3
Germany	€0.27	€0.39	2.06	0.12	1.9
Spain	€0.07	€0.128	1.25	0.058	0.7
US	\$0.11	\$0.147	2.45	0.037	1.1

H.5.4 Per-override welfare loss

Using the expression $\hat{\lambda} = p_{\text{peak}} \Delta\tilde{c}_{\text{eff}} q^O$, we obtain the average loss per override (in £). Table A22 presents results under $\varepsilon = -0.2$ and $\Delta W_0 = 0.095$. Price spreads are expressed in pounds for comparability.

Table A22: Cross-country crossover analysis (RTP $\varepsilon = -0.2$)

Country	p_{peak}	$\Delta\tilde{c}_{\text{eff}}$ (£/kWh)	q^O (kWh)	$\hat{\lambda}$ (£/override)	Full-load $\tilde{\beta}^*$	Partial ($\alpha = 0.25$)	Observed rate
UK	0.1	0.17	2.41	0.072	1.32	0.33	0.0186
Germany	0.205	0.12	2.06	0.051	1.87	0.47	0.107
Spain	0.083	0.058	1.25	0.006	15.83	3.96	0.105
US	0.367	0.037	2.45	0.033	2.88	0.72	0.035

Notes: The expected welfare loss per override is $\hat{\lambda} = p_{\text{peak}} \Delta\tilde{c}_{\text{eff}} q^O$, 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. Crossover rates $\tilde{\beta}^*$ (events/EV-day) are computed as $\Delta W_0 / \hat{\lambda}$, with $\Delta W_0 = 0.095$ and $\alpha = 0.25$ for partial-load automation. Observed rates reflect empirical daily override frequencies. The US is from Texas.

We can see that for the four countries, the observed override rate is well below the crossover override rate. In the UK, there is a low override frequency (1.9% per day) and a moderate price spread yield $\hat{\lambda} \approx 0.072$. Managed charging remains welfare-superior up to $\tilde{\beta}^*(0.25) = 0.33$ per day, an order of magnitude above observed behavior. For Germany, more frequent overrides and a nontrivial tariff gradient narrow managed charging's margin, but the observed rate (0.107) remains below $\tilde{\beta}^*(0.25) = 0.47$. Germany is the tightest case, where welfare parity could emerge if overrides doubled. For Spain, flat tariffs imply tiny welfare penalties ($\hat{\lambda} \approx 0.006$). Even frequent overrides barely affect welfare: the crossover is 15.8/day (full-load) and 3.96/day ($\alpha=0.25$), far above observed 0.105. In

the US (Texas), high p_{peak} but very small $\Delta\tilde{c}_{\text{eff}}$ yield $\hat{\lambda} = 0.033$ and a partial crossover of 0.72/day. Observed 0.035/day lies safely below this threshold, though steeper future tariffs could tighten the margin. To note, the $\bar{\beta}^*$ are expressed as override events per EV-day rather than probabilities. A value above 1 therefore indicates the rate at which managed charging would lose its welfare advantage if users could, on average, override more than once per day. In practice, each vehicle can generate at most one override per plug-in episode, so observed rates (0.02–0.11 per day) are far below this threshold. Thus, $\bar{\beta}^* > 1$ should be interpreted as meaning that the tariff remains welfare-superior even if every vehicle overrode on every charging day.

Across all countries, observed rates (0.019–0.11 per day) are an order of magnitude below the corresponding $\alpha=0.25$ crossovers, confirming that managed charging remains welfare-superior even with only 25% of load automated.⁸¹

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

Third, we treat plug-in behavior as exogenous to the tariff, assuming it remains fixed across pricing regimes. In practice, customers may adapt when and how they charge in response to tariff signals. For instance, RTP could encourage shifting plug-in times toward periods of lower prices and across days, whereas the managed charging tariff’s ToU pricing is fixed, and cannot flexibly incentivize plug-in. We lack data on plug-in behavior for customers who are not on the managed charging tariff, but future work could incorporate these behavioral responses to better capture the interaction between tariff design and charging behavior.

Finally, our analysis is partial-equilibrium: prices $\{\tilde{c}_h\}$ are taken as given and households are price-takers, omitting possible feedback effects if managed charging tariff or RTP adoption is widespread. These simplifications are deliberate, allowing a tractable

⁸¹Higher elasticity (e.g. $\varepsilon = -0.5$) reduces ΔW_0 and proportionally lowers $\bar{\beta}^*$, narrowing managed charging’s advantage by roughly 20%. Yet even under $\varepsilon = -0.5$, managed charging dominates in all countries except possibly Germany.

link between model predictions and experimental data, but they should be borne in mind when interpreting the welfare rankings.