Managing Risk in Electricity Markets

I n a u g u r a l d i s s e r t a t i o n

zur

Erlangung des Doktorgrades

der

Wirtschafts- und Sozialwissenschaftlichen Fakultät

der

Universität zu Köln

Vorgelegt von

Simon Michael Brandkamp, Master of Science

aus Isselburg

2025

Referent:
Prof. Dr. Axel Ockenfels

Erster Korreferent:
Prof. Peter Cramton, Ph.D.

Zweiter Korreferent:
Jun.-Prof. Dr. Oliver Ruhnau

Vorsitzender der Prüfungskommission:
Prof. Dr. Christoph Schottmüller

Datum der Disputation:
07.04.2025

Acknowledgments

I am deeply grateful to my supervisors, Peter Cramton and Axel Ockenfels, for their guidance and support over the last years. Peter's and Axel's passion for applying market design to solve pressing societal problems has inspired me and has profoundly influenced my perspective on the world. Our frequent discussions and their valuable feedback have been essential to shaping this thesis. I also thank Emmanuele Bobbio and David Malec for countless stimulating debates on the "nitty-gritty" details of modeling electricity systems and for never hesitating to engage with my seemingly trivial questions – especially in the early stages of my Ph.D. journey.

My sincere thanks go to my co-authors, Stephanie Chan, Jason Dark, Darell Hoy, Chris Wilkens, and Lucy Yu for the fruitful collaborations and their invaluable feedback on my work. I am also grateful to my fellow Ph.D. students and colleagues for all the enjoyable mensa and coffee breaks, including Rebekka Cordes, Marius Gramb, Max Grossmann, Melisa Kurtis, Aenne Läufer, Mark Marner-Hausen, Jennifer Mayer, Yero Ndiaye, Eva Oeß, Saskia Opitz, Rastislav Rehak, Marten Ritterrath, Nicolas Röver, Lennart Struth, Marina Talantceva, Max Thon, and Luisa Santiago Wolf.

I could not have completed this thesis without the constant emotional support of my friends, who always felt when I needed a welcome distraction. In particular, I owe a special debt of gratitude to my flatmates, Jule Bay, Martti Lehmann, and Pato Tensing, who were always there whenever I feared I might not make it through this challenge.

Finally, I am very grateful to my parents Silvia and Hubert and to my sisters Lena, Vera, and Wiebke for sparking my curiosity about political questions in early years, for always believing in me, and for their incomparable emotional support.

Contents

List of Figures

List of Tables

Introduction

Many societies strive to build electricity systems that supply low-cost, reliable, and green power. A key challenge is that these systems must also be resilient to extreme weather events, which often simultaneously reduce power supply and increase demand. In the coming decades, electricity systems may become more vulnerable because intermittent renewables and the electrification of heating will make the system even more sensitive to extreme weather (Cramton, 2022).

A particularly extreme weather event sparked the motivation for writing this thesis: The winter storm Uri in Texas in February 2021. Uri brought sustained cold weather for multiple days. The cold snap reduced electricity supply because 30 GW of thermal plants became unavailable when gas pipelines froze and technical instruments failed. At the same time, demand for electric heating soared to record levels. With demand far outstripping supply, system operator ERCOT imposed controlled rolling outages on roughly 20% of the system. These outages had catastrophic consequences: 246 deaths and over $100 billion in property damage (Cramton, 2022).

The Texas winter storm illustrates how extreme weather can create severe scarcity of electricity. Many economists argue that high short-term electricity prices are critical for managing scarcity and avoiding blackouts. High prices send scarcity information to market participants, incentivizing them to reduce demand and increase supply (Cramton, 2017).

Yet, scarcity prices are unpopular because they yield seemingly unjustified profits for some generators and expose market participants to enormous financial risk. Consequently, many regulators weaken scarcity price signals by setting price caps or by promoting fixed-price tariffs as default for consumers. These measures erode incentives to cut demand or boost supply during scarcity. When price signals are diluted, regulators use alternative tools to manage scarcity, such as rationing. However, rationing creates severe problems during extreme weather events, as the rolling outages in Texas reveal.

Motivated by the Texas winter storm, this thesis examines how prices can be used to manage scarcity during extreme weather. It investigates two overarching research questions: First, can dynamic electricity prices incentivize consumers to play a more active role in mitigating scarcity events? If consumers' response to dynamic prices was weak, it would not be effective to expose them to short-term scarcity prices. Second, how can consumers and generators manage the financial risks created by high scarcity prices? Enabling market participants to effectively manage these risks is likely essential for increasing the popularity of scarcity pricing.

Chapter 1, "Resilient electricity requires consumer engagement"[1], co-authored with Emmanuele Bobbio, Stephanie Chan, Peter Cramton, David Malec, Axel Ockenfels, and Lucy Yu, addresses the first of the above research questions. We analyze a large sample of UK households on dynamic tariffs linked to day-ahead prices. Using a fixed-effects regression model, we estimate households' price elasticities and investigate whether certain technologies (e.g., electric vehicles) boost price responsiveness.

We find a relatively large and significant average own-price elasticity of -0.26. Cross-price elasticities are mostly small and insignificant, indicating that consumers do not substantially shift usage across time. Households with at least one low-carbon technology exhibit higher price responsiveness, particularly electric vehicle owners who also have batteries and solar PV, with elasticities exceeding -0.5. In addition, we conduct a thought experiment to derive the share of price-responsive demand in Texas that would have been sufficient to avoid outages during the winter storm if price-responsive consumers in Texas were as price-responsive as the average UK household in our sample. We find that the winter storm outages could have been avoided if only 44% of Texan consumers had been as price-responsive as the UK households that we study.

Chapter 1 points out that residential consumers on dynamic electricity tariffs are price-sensitive and can, therefore, play a vital role in lowering demand during scarcity events. However, scarcity prices create risks for households on dynamic tariffs because they can lead to very high electricity bills. Chapter 2, therefore, studies how to help households manage these risks by combining dynamic electricity tariffs with forward contracts. Such a hedged tariff is supposed to protect households from high scarcity prices while preserving the incentive to reduce electricity consumption during scarcity events.

[1]This chapter contains the current version of the working paper Cramton et al. (2025a). Earlier versions of this working paper were published under the same title by Bobbio et al. (2022b) and Bobbio et al. (2024). An earlier version of this working paper was also published under the preliminary title "Price responsive demand in Britain's electricity market" by Bobbio et al. (2022a).

Chapter 2 is titled "Hedging households against extreme electricity prices"[2] and was single-authored. It employs the same large dataset of UK households on a dynamic electricity tariff as Chapter 1. In Chapter 2, I use a utility maximization model to derive each household's optimal hedge share, i.e., the optimal share of its typical electricity consumption that the household should buy forward. I also investigate how effectively the optimal hedge protects households from volatility of their monthly electricity bills.

I find that the average optimal hedge share is 59%, varying strongly across households. Parts of this between-household variation in hedge shares relate to technology ownership: Homes with electric heating and electric vehicles opt for higher hedge shares, while those with solar PV and battery storage choose lower ones. In addition, I reveal that households can better manage risks if they choose different hedge shares for different times of day. Time-dependent hedge shares take into account how a household's daily demand pattern is correlated with aggregate load and prices.

My novel theoretical contribution in Chapter 2 is that an increase in price elasticity of demand raises households' optimal hedge shares if they experience a positive correlation between electricity prices and their weather-related desire to consume electricity. The more price-elastic the household is, the more it uses the forward hedge to mitigate its exposure to spot prices.

Moreover, Chapter 2 highlights that the optimal forward hedge is effective in reducing the volatility of monthly electricity bills, on average, by 18%. On the other hand, I show that the optimal hedge only leads to small welfare gains for the households in the UK. The reason is that, during the sample period, the UK households are only exposed to relatively small price spikes compared to the day-ahead price spikes observable in Texas. The welfare gains from optimal hedging strongly increase when exposing UK households to a simulated Texas-style extreme weather event with high prices for multiple days in a row.

Chapter 2 points out that hedging with forward contracts can be an effective risk management strategy for households. In Chapter 3, we study hedging strategies for electricity generation companies and load-serving entities (LSEs) in the ERCOT day-ahead electricity market in Texas. Chapter 3 is titled "Hedging electricity price spikes with forwards and options"[3] and is co-authored with Peter Cramton, Jason Dark, Darrel Hoy, and David Malec. In our model, generators and LSEs choose the optimal mix of forward contracts and European call options. We are particularly interested in how this

[2]This chapter contains the current version of the working paper Brandkamp (2025). An abstract of this working paper was also published in the conference proceedings of IAEE (2024) and EAERE (2024).

[3]This chapter contains the current version of the working paper Brandkamp et al. (2025).

optimal mix is affected by the frequency and size of day-ahead price spikes in a delivery period.

To model price spikes, we estimate a regime-switching model that allows simulating joint distributions of hourly day-ahead prices, net load, renewable generation, and daily gas prices in the ERCOT market between 2011 and 2022 (Coulon et al., 2013). We also run a merit order dispatch model for a large sample of power plants to obtain distributions of hourly profits for different power plant technologies.

We show that frequent and large price spikes in a delivery period can cause significant worst-case losses for generators and LSEs. Depending on their level of risk aversion, the agents trade off minimizing worst-case losses versus minimizing profit variance when selecting their optimal hedge strategies. In delivery periods with large price spikes, generators and LSEs choose larger option holdings and smaller forward positions to be protected against worst-case losses. The LSE relies more on options than the generator, as price spikes result in more extreme negative profit tails for the LSE.

Yet, when forwards and options are combined, the resulting reduction in profit volatility and worst-case losses is only marginal compared to a hedging strategy that only uses forward contracts. Hedging with forwards-only is almost as effective as combining forwards and options because day-ahead profits are roughly linear in day-ahead prices. Finally, we show that raising the option strike price of the option drives agents to choose more forwards and fewer options. Overall, our findings in Chapter 3 suggest that hedging with forwards and options is effective in lowering profit variability and worst-case losses, even under intense price spikes.

Chapter 3 characterizes optimal hedging strategies for arbitrage-free forward and option prices. In Chapter 4, we build on these optimal hedging strategies to simulate market equilibria in a forward energy market with a novel market design. In equilibrium, forward and option prices for electricity typically deviate from arbitrage-free levels (Bessembinder & Lemmon, 2002, Redl et al., 2009, Botterud et al., 2010). Therefore, Chapter 4 extends the analysis in Chapter 3 and simulates hedging strategies for forwards and option prices that deviate from arbitrage-free levels.

Chapter 4 is titled “A Forward Energy Market to Improve Reliability and Resiliency”[4] and is co-authored with Peter Cramton, Jason Dark, Darrel Hoy, David Malec, Axel Ockenfels, and Chris Wilkens. The main contribution of this chapter is to propose a novel market design for forward electricity contracts and options. Participants in this market

[4]This chapter contains the current version of the working paper Cramton et al. (2025b). Parts of the analysis in the paper were also published as related policy white papers under Cramton (2023) and Cramton et al. (2024b).

can trade thousands of granular hourly delivery periods for up to four years ahead. To facilitate trading so many granular products simultaneously, we apply the flow trading technology by Budish et al. (2023) to electricity markets.

As a proof of concept, Chapter 4 develops a large-scale simulation of ERCOT's day-ahead and forward energy market between 2011 and 2022, building on the day-ahead market simulations in Chapter 3. We use these simulations to analyze generators' and LSEs' demand curves for forwards and options.

Our main finding is that the slope of net demand curves for forward and option products depends strongly on agents' risk preferences. Highly risk-averse agents display almost vertical demand curves. High risk aversion would, therefore, likely lead to market clearing prices for forwards and options that are well above arbitrage-free levels. By contrast, when risk aversion is more moderate, market participants are willing to exploit arbitrage opportunities, likely causing market clearing prices to align closer to arbitrage-free benchmarks. The study also reveals that demand for both forwards and options is more sensitive to changes in forward prices than option prices. Moreover, agents often take especially large arbitrage positions during peak periods prone to extreme price spikes. Finally, lowering the strike price of the option encourages heavier use of options and smaller forward positions. A lower strike price also induces agents to take larger arbitrage positions in both forwards and options.

In the future, we plan to extend the research in Chapter 4 to analyze how equilibrium prices and quantities in our proposed forward energy market evolve over time as they get closer to their physical delivery periods. We also aim to investigate how the forward energy market impacts generators' and LSEs' expected profit, profit volatility, and downside tail risks.

Overall, the chapters in this thesis highlight how scarcity prices can play a crucial role in improving the resilience of electricity systems to extreme weather events. At the same time, it is essential to combine scarcity pricing with granular hedging products to allow consumers and electricity producers to manage the financial risks involved.

From a policy perspective, regulators should foster the development of accessible hedging products, which may require substantial efforts to educate consumers about the benefits of hedged dynamic electricity tariffs. In wholesale markets, policymakers should establish centralized platforms offering granular and liquid hedging instruments.

Taken together, these measures would enable consumers and producers to manage the risks posed by scarcity pricing effectively. At the same time, combining scarcity prices and hedging preserves consumers' and producers' incentives to align their behavior with the needs of the electricity system during extreme weather events.

Chapter overview

Chapter 1

Cramton, P., Bobbio, E., Brandkamp, S., Chan, S., Malec, D., Ockenfels, A., & Yu, L. (2025a). Resilient electricity requires consumer engagement. *Working Paper*. https://doi.org/10.2139/ssrn.5097987

This chapter is co-authored with Emmanuele Bobbio, Stephanie Chan, Peter Cramton, David Malec, Axel Ockenfels, and Lucy Yu. All authors contributed equally to the project. The research idea was jointly developed by all authors. Stephanie Chan and David Malec provided and cleaned the dataset and conducted preliminary analyses. Emmanuele Bobbio and I carried out the econometric and statistical analysis. Peter Cramton, Emmanuele Bobbio, and I wrote parts of the policy sections. Axel Ockenfels and Lucy Yu also contributed to writing the policy sections and provided feedback on the overall draft and analysis.

Chapter 2

Brandkamp, S. (2025). Hedging households against extreme electricity prices. *University of Cologne, Working Paper Series in Economics*, (106). https://ideas.repec.org/p/kls/series/0106.html

This chapter was single-authored.

Chapter 3

Brandkamp, S., Cramton, P., Dark, J., Hoy, D., & Malec, D. (2025). Hedging electricity price spikes with forwards and options. *University of Maryland Working Paper*. https://cramton.umd.edu/papers2020-2024/brandkamp-et-al-hedging-with-forwards-and-options.pdf

This chapter is co-authored with Peter Cramton, Jason Dark, Darrell Hoy, and David Malec. I developed the research idea with substantial support from all other co-authors, in particular Peter Cramton. Jason Dark and I estimated the regime-switching model and implemented the numerical optimization for finding optimal forwards and option quantities. Darrell Hoy, David Malec, and I developed and simulated the merit order dispatch model. I conducted the analysis of the optimal hedging strategies and wrote the draft. Peter Cramton, Jason Dark, Darrell Hoy, and David Malec gave feedback on the draft.

Chapter 4

Cramton, P., Brandkamp, S., Dark, J., Hoy, D., Malec, D., Ockenfels, A., & Wilkens, C. (2025b). A forward energy market to improve reliability and resiliency. *University of Maryland Working Paper.* https://cramton.umd.edu/papers2020-2024/cramton-et-al-forward-energy-market.pdf

This chapter is co-authored with Peter Cramton, Jason Dark, Darell Hoy, David Malec, Axel Ockenfels, and Chris Wilkens. All authors contributed equally to this project. Peter Cramton developed the research idea with support from all other team members. Jason Dark and I recalibrated the regime-switching model and set up the numerical optimization for deriving the net demand curves. Darrell Hoy, David Malec, and I developed and simulated the merit order dispatch model. David Malec run the simulation for finding market clearing prices and quantities. David Malec and I analyzed the results of the simulations. Peter Cramton and I wrote the draft. All other authors gave feedback on the draft.

Chapter 1

Resilient electricity requires consumer engagement

This chapter is co-authored with Emmanuele Bobbio, senior lead economist at PJM Interconnection, Stephanie Chan, a data scientist at the Centre for Net Zero, Peter Cramton, Professor of Economics at University of Maryland, David Malec a research associate at the University of Maryland, Axel Ockenfels, Professor of Economics at the University of Cologne and Director at the Max Planck Institute for Research on Collective Goods in Bonn, and Lucy Yu, chief executive officer at the Centre for Net Zero.

This chapter contains the current version of the working paper Cramton et al. (2025a). Earlier versions of this working paper were published under the same title by Bobbio et al. (2022b) and Bobbio et al. (2024). An earlier version of this working paper was also published under the preliminary title "Price responsive demand in Britain's electricity market" by Bobbio et al. (2022a).

Abstract

Active consumers are essential to transitioning to a flexible, resilient energy system. Electricity markets balance supply and demand with price. Historically, this price response has come almost entirely from supply. However, when much of supply is intermittent or in-flexible, price-responsive demand becomes essential for energy reliability. It is also key to building resiliency into a system facing more extreme events in a changing climate. We measure how price-responsive consumers are in Britain from August 2020 to August 2021 with half-hourly individual household data. Our sample includes customers with a dynamic rate that tracks wholesale cost and flat-rate customers used to control for

weather and other factors. A one percent increase in price reduces demand by 0.26 percent. This elasticity is larger for consumers owning low-carbon technologies. This price response is sufficient to maintain system balance in extreme events even when most consumers are unresponsive. Regulators can encourage price-responsive demand through retail choice and subsidizing enabling technologies. Regulators can also protect consumers with man-dated hedging in dynamic plans. Low-income households benefit most from such policies.

Declarations

This research is funded by the Deutsche Forschungsgemeinschaft (German Research Foundation) under Germany's Excellence Strategy–EXC 2126/1–390838866 and the European Research Council under the European Union's Horizon 2020 research and innovation program, grant 741409.

1.1 Introduction

In February 2021, winter storm Uri blanketed Texas with extreme cold for several days. Thermal generators failed, and heating demand surged, creating a 37% shortfall between electricity demand and supply. The system operator, ERCOT, had to order controlled outages to keep the system balanced and avoid a cascading blackout. Over four million Texans were without power for several days. Many were without water as the interconnected critical infrastructures failed - first gas, then electricity, then water. Hundreds of people died. Dollar damages totaled many tens of billions. Cramton (2022) provides details.

Climate change makes extreme events more salient, frequent, and severe (Cohen et al., 2021). Research on how to make critical infrastructure more resilient to extreme weather is needed. So far, the burden of balancing the power system during extreme events has fallen almost exclusively on the supply side. Demand hardly reacts to improve resiliency since households and businesses generally pay a flat rate per kilowatt-hour of electricity, irrespective of market conditions. Flat rates are simple and limit variations in electricity bills. Still, flat rates are inefficient in today's electricity markets, where the social cost of generating electricity varies from -$0.50 to $9.00/kWh, depending on time and location.

Here, we provide new and strong evidence that demand-side policies promoting price-responsive demand and energy efficiency offer significant opportunities to build system resiliency. They are an essential element of any least-cost approach to improving resiliency - a lesson relevant to electricity markets worldwide.

Price-responsive demand is enabled by smart meters - measuring use at each time and location - and smart markets that enable consumers' price response to be expressed, either individually for wholesale consumers or in aggregate for retail consumers. The increase in bid-in demand has many efficiency, reliability, and resiliency benefits.

We divide our analysis into three parts. Firstly, we show that domestic consumers significantly respond to electricity prices using a large sample of 4,148 households in the UK. The households receive dynamic rates that follow wholesale electricity prices on a half-hourly basis. We estimate households' price elasticity of electricity demand using a fixed effects regression model while exploiting the price-consumption patterns of households on flat rates as controls for other sources of variation.

Our results indicate that a one percent price increase induces domestic consumers to reduce consumption by, on average, 0.26 percent. The elasticity is larger for households owning low-carbon technologies, particularly electric vehicles, which also tend to be more flexible and have stronger financial incentives to respond to prices. For future work, it will be crucial to estimate how consumers' price response changes in a dynamic framework in which a rising share of households adopts these low-carbon technologies.

Importantly, we do not aim to estimate the price elasticity of demand for the general population of electricity consumers in the UK. The price elasticity of the UK's general population is likely low, as other studies suggest (Fabra et al., 2021). In our study, we estimate the price elasticity of first movers who self-selected into a dynamic tariff and who are very different from today's average consumers. First movers likely have a larger price elasticity as they are potentially more aware of their electricity prices and are more likely to have low-carbon technologies. These characteristics make first movers more representative than today's average consumers for an increasingly large share of households in the future. In the upcoming decade, more households will adopt low-carbon technologies and will likely have larger price elasticities than today's consumers. Therefore, it is important to study the price elasticity of first movers while emphasizing that we do not generalize our results to the population of today's consumers.

In the second part of the analysis, we conduct a thought experiment to study how price-responsive consumers can make power systems more resilient against extreme events like the Texas winter storm. In this event, the electricity price surged to $9.00/kWh for multiple days due to the electricity shortage. About 999/1000 Texan households did not respond to the high shortage price because they had a conventional fixed-price tariff. In contrast, the 1/1000 set of consumers on dynamic rates had a strong incentive to reduce consumption in response to extreme shortage prices. Price-responsive consumers on dynamic rates help restore the balance between supply and demand and prevent outages. We estimate the share of price-responsive consumers on dynamic rates that would have been necessary to prevent outages during the Texas electricity crisis if these consumers were as price-elastic as the UK consumers in our sample.

We find that power outages during the Texas winter storm could have been avoided if 44% of Texan consumers had a price-responsive rate and responded in line with UK consumers. The thought experiment is not meant to provide a precise estimate of the share of price-responsive demand required to avoid outages but rather illustrates that an achievable share of price-responsive consumers can help make power systems resilient to extreme weather events.

That said, dynamic rates without proper risk protection are unacceptable for domestic consumers. Designed incorrectly, dynamic rates can make households vulnerable to extremely high prices during shortage events, leading to exorbitant electricity bills. In the third part of our analysis, we subsequently argue that it is essential and possible to accompany dynamic electricity rates with additional regulatory measures to protect households from price risk.

Our results suggest regulators should promote dynamic electricity pricing with forward hedging to protect consumers from high prices while encouraging price responsiveness. We conjecture that low-income households would benefit most from reduced consumption incentives and energy efficiency subsidies. Additionally, policies should accelerate smart meter adoption and enhance competition to offer innovative pricing plans. These measures will help prevent power outages, drive investments in low-carbon technologies, and support a faster and more resilient green energy transition. We will get back to policy implications in our concluding section.

Regulators should also combine dynamic pricing with low-income subsidies for energy efficiency investments. Dynamic pricing makes investments in energy efficiency measures like insulation more appealing since energy efficiency protects households from high prices. These subsidies ensure that low-income households also benefit from dynamic pricing.

Moreover, regulators need to implement enabling policies for rapid adoption of demand response. For instance, they should support an accelerated smart meter roll-out and intensify retail competition to encourage electricity suppliers to offer innovative, dynamic plans.

Implementing the above regulatory measures is essential to make dynamic electricity rates with a forward hedge appealing to consumers. Attractive dynamic rate plans encourage domestic consumers to actively participate in the energy transition. They induce consumers to mitigate demand peaks and help prevent power outages during extreme weather events. Demand response also motivates households to invest in low-carbon technologies like electric vehicles, battery storage, and energy efficiency - which are essential to decarbonization. Dynamic electricity prices are vital for a fast, efficient, and resilient green transition, and should be rapidly adopted at scale.

We structure the paper as follows: In Section 1.2, we connect our research to the literature on dynamic electricity prices and price-responsive demand. Section 1.3 describes the UK household-level data on electricity prices and consumption. Section 1.4 presents the regression model we employ to estimate households' price elasticity of electricity demand. Section 1.5 discusses our estimation results. In section 1.6, we conduct a thought experiment to illustrate how price-responsive consumers can strengthen resiliency during extreme weather events. Section 1.7 evaluates additional policy measures regulators should implement alongside dynamic electricity rates. Section 1.8 concludes.

1.2 Literature

A large literature stresses the benefits of exposing domestic consumers to dynamic electricity prices (Allcott, 2011, Borenstein, 2005, Houthakker, 1951). Dynamic prices increase the efficiency of the power system as they help align electricity consumption with the short-run marginal cost of power generation. Thereby, they create incentives for efficient short-run generation and long-term investment (Borenstein & Holland, 2005), mitigate market power (Poletti & Wright, 2020), and potentially reduce carbon emissions (Cahana et al., 2022, Holland & Mansur, 2008).

Adding to the above literature, we highlight an additional benefit: dynamic prices for domestic consumers make power systems more resilient to extreme weather events. Previous research has recognized that price-responsive consumers can improve system resiliency by lowering peak demand during extreme weather. However, this research mainly focused on demand response from large industrial consumers (Wang et al., 2017). We analyze the essential role of price-responsive households for resiliency in the future.

Previous estimates for short-term price elasticities differ widely, given that existing studies vary in sample sizes and granularity of price changes. Most research builds on experiments and small-scale field studies (Allcott, 2011, Faruqui & Sergici, 2010). Moreover, rather than analyzing fully dynamic real-time prices, most studies examine less granular time-of-use (Braithwait, 2000, Harding & Lamarche, 2016, Train & Mehrez, 1994) or critical peak prices (Bollinger & Hartmann, 2020, Faruqui & George, 2005, Faruqui et al., 2014, Jessoe & Rapson, 2014). Only a few papers investigate large datasets containing individual customers exposed to dynamic real-time prices (Fabra et al., 2021, Stumpe, 2022).

Overall, the above literature suggests that the price response of domestic electricity consumers is small, with own- and cross-price elasticities mostly below −0.2 (Andruszkiewicz et al., 2019, Harding & Sexton, 2017). The literature also discusses three types of enabling technologies that raise consumers' price elasticity. The first type contains devices like in-home displays and phone apps that give consumers feedback about their consumption and prices. Multiple studies find that these information technologies increase demand response, especially during peak-price hours (Allcott, 2011, Faruqui et al., 2014). For instance, Jessoe and Rapson (2014) find that in-home displays increase households' price elasticity by roughly three standard deviations. The second technology type consists of devices that automize consumers' demand response, e.g., smart thermostats and smart EV charging (Faruqui & Sergici, 2011, Wolak, 2010). Bollinger and Hartmann (2020) show that these automation technologies increase price elasticities more than information technologies. The third type includes flexibility technologies like electric vehicles, electric heating, and battery storage that make households' electric load more flexible.

There is limited household-level research on the effect of flexibility technologies on consumers' price elasticities. Reiss and White (2005) and Wolak (2010) reveal that electric heating and cooling ownership strongly increases price response for small experimental samples. Ruan et al. (2022) simulate price-responsive demand with various electric appliances and estimate time-varying price elasticities. Their simulations show that a smart, dynamic pricing mechanism can reduce the peak-to-average demand ratio when low-carbon technologies are applied.

We contribute to this literature by estimating the effect of several flexibility technologies on price elasticities for a large observational sample of households. Our sample also allows studying the impact of the combination of multiple flexibility technologies (e.g., electric vehicle plus battery storage) on demand response.

We also add to a literature addressing concerns that might explain why regulators rarely implement dynamic prices. One concern is that dynamic pricing exposes domestic consumers to unacceptable risk by making their electricity bills more volatile (Burger et al., 2020, Faruqui, 2012). To tackle this concern, we discuss combining dynamic electricity tariffs with a forward hedge to shield residential consumers from high prices. Several authors propose forward hedges for industrial electricity consumers (Borenstein, 2007b, Schlecht et al., 2024, Wolak & Hardman, 2022). Some US electricity suppliers have already offered forward hedges to industrial consumers (Barbose et al., 2005, Braithwait & Eakin, 2002). However, few studies look at households. One exception is a pilot study by Stavrogiannis (2010) who finds that forward hedging effectively reduces household bill volatility.

Yet, several authors raise concerns about dynamic pricing from a social justice perspective. Cahana et al. (2022) and Horowitz and Lave (2014) find that low-income households are likely to be worse off on dynamic rates since they often do not have the necessary appliances to respond to prices. Our contention is that consumer flexibility will be a central part of the future energy system, and therefore developing our understanding of dynamic pricing will allow us to address any distributional impacts. We argue that forward hedging can actually make dynamic pricing more attractive for low-income households. At the same time, we emphasize that regulators should address social justice concerns by accompanying dynamic pricing with low-income subsidies for energy efficiency measures and for other technologies (e.g., smart thermostats) that help low-income households to be price-responsive.

1.3 Data Sources

Global electricity supplier Octopus Energy provided anonymized, half-hour electricity smart meter customer readings on three plans: fixed, dynamic (wholesale-linked), and electric vehicle (EV) rates. Our sample of 15,000 British customers consists of approximately 5,000 consumers randomly sampled from each plan from August 2020 to August 2021. All had smart meters since at least August 2020.

Fixed-rate

Fixed-rate customers pay a constant price for electricity for all half hours of the day. The fixed rate differs between households depending on their grid supply area and the specific fixed-rate plan they chose. Eight percent of fixed-rate customers experience a minor adjustment of their fixed rate once during the sample period.

Dynamic rate

The dynamic rate reflects day-ahead auction prices for electricity. Figure 1.1 depicts the half-hourly day-ahead wholesale prices for the Great Britain price zone obtained from

EPEX SPOT (2023). The graph reveals that wholesale electricity prices are relatively low and stable during the sample period and rarely exceed 10 p/kWh. However, while the overall volatility of wholesale prices is low, there are a few price spike hours in which wholesale prices are very high.

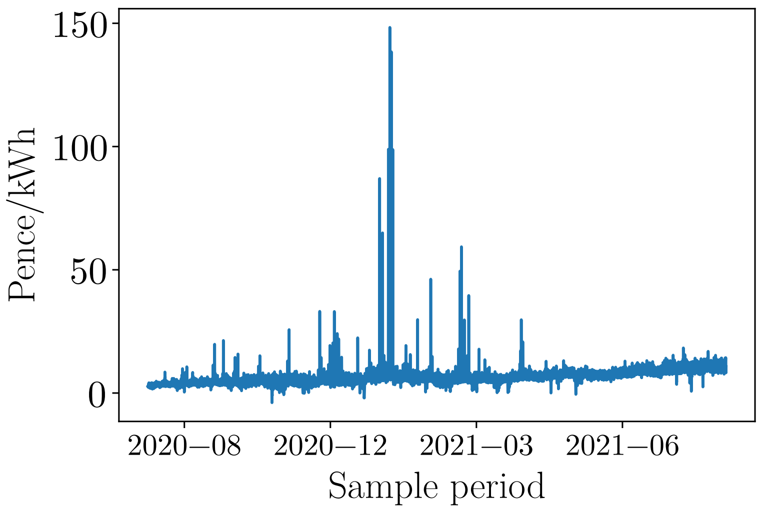

Figure 1.1: Half-hourly day-ahead wholesale electricity prices in the Great Britain price zone obtained from EPEX SPOT (2023)

The final dynamic prices that households pay add distribution costs and a peak time premium to the wholesale prices.[1] Octopus designed the plan to encourage consumers to shift their consumption outside the 4-7 pm peak. Customers have forward notice of these half-hourly prices, made available every evening between 4-8 pm for the next day. A negative wholesale price can result in a negative customer price, known as plunge pricing. However, a cap at 35p/kWh – roughly double a fixed rate – protects customers from surge pricing.

EV rate

EV plans offer electricity at two rates: an off-peak price during fixed charging hours, such as 00:30 - 04:30, and a peak price approximately three times higher. Pricing depends slightly on geographic location. Octopus designed the plan to incentivize consumers to charge their EVs in the off-peak window.

[1] For every half-hourly interval, Octopus Energy multiplies the day-ahead auction price with a distribution charge multiplier that ranges from 2 to 2.4, depending on the grid supply area the household lives in. Between 4 pm and 7 pm, Octopus Energy also adds a peak-time premium that ranges from 11 to 14 p/kWh, depending on the grid supply area. Afterward, VAT is added. The resulting price is the dynamic price for the half-hourly interval unless it exceeds the price cap of 35 p/kWh. If the price exceeds the price cap, the price is set to 35 p/kWh Octopus Energy (2019).

Our analysis below only focuses on customers with dynamic and fixed rates. Each consumer is associated with up to one year of smart meter readings from August 2020 to August 2021. As users are free to switch plans across this period, they may belong to multiple groups across the whole period. For instance, a user on a dynamic rate may choose to migrate to an EV rate at any point.[2]

Figure 1.2 shows the number of customers on the fixed and dynamic tariffs during the sample period. It highlights that surprisingly many households switch to another tariff or supplier during the sample period. Switching tariffs or suppliers is very common in Britain. 20 percent of British electricity consumers switched to a different supplier in 2020 (DESNZ, 2023). The share of customers who switch internally to another tariff offered by the same supplier was even 29 percent in Britain in 2019 (ACER, 2021). The frequent switching behavior explains why many consumers leave their tariff during the sample period.

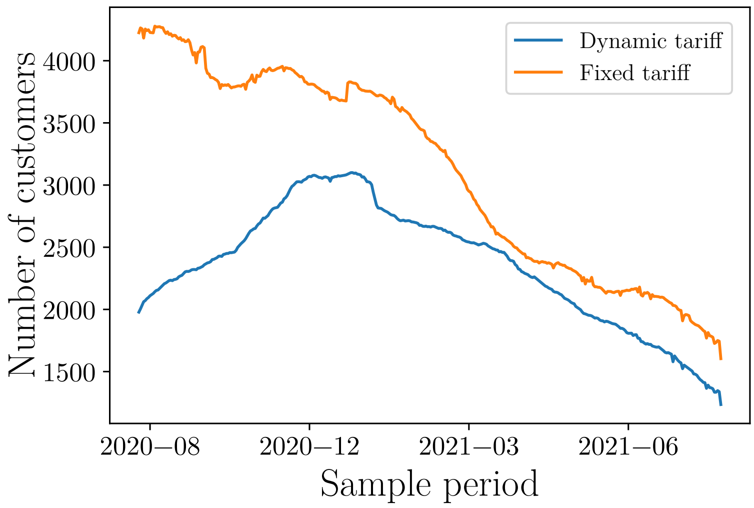

Figure 1.2: Number of customers by tariff type

Figure 1.2 also highlights that the consumers in our sample self-selected in and out of the dynamic tariff. Since dynamic tariffs are new and uncommon in the UK, the households on dynamic rates are likely first movers and not representative of the UK's general population. As discussed in the introduction, we do not aim to estimate the price elasticity for today's general population, but only for the first movers because the first movers might be more representative of a growing share of tomorrow's consumers. The

[2]A small proportion of users participate in export plans, which are structured as separate plans. A user may be in multiple plans at the same point in time in this instance.

estimation results in section 5 should, therefore not be generalized to the population of today's consumers.

Table 1.1 provides summary statistics for households' electricity consumption on dynamic and fixed tariffs respectively, broken down by season and by time of day. Electricity consumption refers to consumption per half-hourly interval. Households on dynamic tariffs have a higher mean consumption and a higher average standard deviation of their consumption (Avg SD) than households on fixed tariffs. These statistics reinforce the hypothesis that households who self-selected into dynamic tariffs are systematically different from households on fixed tariffs. The higher consumption level of dynamic households suggests that these households likely have higher income levels.

	Electricity consumption in kWh				Price in pence/kWh			
Tariff type	**Mean**	**Avg SD**	**Min**	**Max**	**Mean**	**Avg SD**	**Min**	**Max**
				All observations				
Dynamic	0.37	0.49	0	5.29	14.78	7.12	-9.88	35.00
Fixed	0.24	0.27	0	5.29	14.90	0.19	5.78	33.86
				By season				
Dynamic fall	0.43	0.52	0	5.29	11.99	6.89	-9.88	35.00
Dynamic spring	0.31	0.42	0	5.29	17.55	6.14	-1.33	35.00
Dynamic summer	0.28	0.40	0	5.29	16.05	6.59	0.80	35.00
Dynamic winter	0.43	0.48	0	5.29	14.71	7.12	0.29	35.00
Fixed fall	0.26	0.27	0	5.29				
Fixed spring	0.21	0.24	0	5.29				
Fixed summer	0.20	0.23	0	5.28				
Fixed winter	0.28	0.27	0	5.29				
				By time of day				
Dynamic afternoon	0.34	0.36	0	5.29	19.73	9.06	0.42	35.00
Dynamic evening	0.39	0.41	0	5.29	15.02	6.40	-2.52	35.00
Dynamic morning	0.33	0.36	0	5.29	13.95	4.26	-5.04	35.00
Dynamic night	0.43	0.59	0	5.29	10.40	3.76	-9.88	29.74
Fixed afternoon	0.31	0.30	0	5.28				
Fixed evening	0.28	0.26	0	5.29				
Fixed morning	0.25	0.25	0	5.29				
Fixed night	0.13	0.13	0	5.29				

Table 1.1: Number of customers by customer groups and plans

Consumption is higher in fall and winter than in summer and spring for both customer groups. Interestingly, consumption is highest at night for households on dynamic tariffs, while it is lowest at night for households on fixed tariffs. This suggests that dynamic households might shift significant load from electric heating or vehicles to night hours. Across almost all seasons and times of day, the maximum consumption is 5.29 kWh, which is likely an administrative upper measurement limit set by the retailer. This maximum limit is only reached for 60 half-hourly observations.

On average, retail prices for households on dynamic and fixed tariffs are almost equal. The average standard deviation of half-hourly prices per household on dynamic tariffs is quite large (the non-zero standard deviation for fixed tariff stems from period changes in fixed electricity prices). Dynamic prices can go significantly negative but do so only in 0.3% of the half-hourly periods we observe. Dynamic prices are largest in spring and summer and peak in the afternoon.

Low-carbon technology ownership survey data

When users sign up for a dynamic rate, Octopus Energy asks them to complete a survey to indicate their ownership of various low-carbon technologies (LCTs). The survey queries information on four LCTs: 1) smart thermostats, 2) electric vehicles, 3) residential solar, and 4) battery storage. Boolean flags indicate the stated ownership of these technologies. In addition, we add information on electric heating ownership that is not survey-based but inferred from the lack of a gas contract with Octopus Energy—households generally purchase electricity and gas from the same provider in Britain.[3]

Table 1.2 shows the number of customers on dynamic and fixed rates by ownership of low-carbon technologies (LCTs). We only have information on LCT ownership for a subsample of households since many households did not participate in the survey.

1.4 Methodology

We employ a time segment fixed effects regression model to estimate price elasticities of electricity demand. A time segment is defined as the combination of year, month, day of the week, and half-hour interval.

The fixed effects control for periodicity and trends. They capture how consumption and prices move relative to one another after accounting for their characteristic values in a particular time segment. Electricity consumption and prices are positively correlated and vary systematically across years, seasons, and throughout the day. Consumption and prices are higher in the early morning than at night, and the afternoon peak occurs later in the summer than in the winter. However, their joint distribution is approximately

[3]Our indicator of electric heating ownership is relatively imprecise. Some customers who do not have a gas contract with Octopus Energy might purchase gas from another supplier or might use oil heating. Therefore, we might falsely assume some gas heating owners use electric heating. This imprecision will likely lead to a downward bias of our estimates for electric heating ownership because gas heating owners are arguably less responsive to electricity prices than electric heating owners.

	Customer groups	**Dynamic rate**	**Fixed-rate**
1	All customers	4148	5904
3	At least one LCT	2593	647
2	No LCT	99	12
4	Inferred electric heating only	414	162
5	Inferred electric heating + smart thermostat	411	94
6	EV only	1025	352
7	Solar only	138	27
8	Battery only	49	12
9	EV + solar	280	45
10	EV + battery	66	16
11	Solar + battery	124	21
12	EV + solar + battery	177	20

Table 1.2: Number of customers by low-carbon technology ownership

stable for a particular time segment. For example, consumption and prices on Monday at 7:30 behave similarly, regardless of whether it is the first or second Monday of March 2021. We can view each time segment as a repeated experiment where we draw four or five consumption-price pairs—depending on the number of weeks in a month.

Using fixed effects for these time segments, we estimate the following model by OLS:

$$\ln(C_t^d) = c^{s(t)} + \sum_{j=-16}^{16} \gamma_j \ln(P_{t+j}) + \theta \ln(C_t^f) + \varepsilon_t$$

t denotes time in half-hour increments. C is the average consumption of households on dynamic plans (superscript d) and fixed rates (superscript f); P is the average retail price paid by households on dynamic rates. $c^{s(t)}$ denotes the fixed effect for time segment $s(t)$. ε is the error term, which we assume is independently and identically distributed and uncorrelated with regressors P and C^f. Since we are interested in capturing the average price elasticity of households on dynamic tariffs, we do not estimate the above equation on a household level. We regress the log average consumption of households on dynamic tariffs on the log dynamic price and the log average consumption of households on fixed tariffs.

The specification assumes that the demand curve is isoelastic – if the price increases by $x\%$, consumption varies by $\gamma \times x\%$ regardless of the consumption and price level. Households may respond to a higher price by shifting consumption to other time intervals when the price is lower, e.g., running the dishwasher or charging the electric vehicle earlier or later. At 16:00, households on dynamic plans learn prices for the following day. Thus, the model includes contemporaneous and earlier and later prices, spanning a ±8-hour window for 33 elasticity coefficients γ_{j} , $j = -16, \dots , +16$.

The model includes the average consumption of households on a fixed rate as control for demand shocks. A positive demand shock increases the price, affecting households' consumption on dynamic plans both directly and indirectly – via the price. The demand shock impacts households on fixed rates only directly since they do not face price changes. Controlling for fixed-rate consumption, therefore, allows isolating the effect of a price change on the consumption of households on dynamic plans.

Failure to control for demand shocks positively biases the contemporaneous elasticity coefficient. Interestingly, the response to prices turns out to be sufficiently strong; estimating the model without control delivers a statistically and economically significant negative coefficient. As expected, introducing the control increases the coefficient's magnitude in absolute value. The estimated coefficient on the control is not statistically different from one, indicating that the direct effect of demand shocks is similar for customers on dynamic and fixed rates.

On the other hand, we assume that the price response of consumers on dynamic rates does not have a measurable impact on wholesale electricity prices. Only a tiny share of households receives dynamic rates in the UK. Moreover, households' electricity consumption only accounts for 38% of aggregate electricity demand in the UK (DESNZ, 2023). Aggregate demand hardly changes if only a tiny share of households adjusts demand in response to prices, even if this price response is large. Hence, the responsive demand has a negligible effect on aggregate demand and wholesale electricity prices. In contrast, the price response of households likely impacts aggregate demand and prices in countries like Spain, where most households receive dynamic prices. In such a setting, using an instrument like wind generation forecasts is necessary to avoid a simultaneous equation bias (Fabra et al., 2021).

Our analysis first examines how households on dynamic plans respond to prices. Then, we stratify the analysis by low-carbon technology (LCT) ownership status. For example, we estimate the model by restricting attention to households with electric vehicles.

One drawback of our analysis is that most households did not participate in the LCT survey. Moreover, only customers who switched to a dynamic rate were asked to complete the survey. Therefore, we only have LCT information for households on fixed rates if they switch to a dynamic rate at some point. These customers on fixed rates might differ from fixed-rate customers who never switched to a dynamic rate.

We drop all households without LCT ownership information for all regressions that analyze the effect of owning an LCT. This filtering leaves only a few observations for some

LCTs, especially for fixed-rate customers (see Table 1.2). For example, we only have twelve customers on fixed rates with no LCTs. Thus, we estimate the model using a generic control group with all 5,904 households on fixed rates, irrespective of LCT ownership status. As a robustness check, we also run the model using smaller LCT-group-specific control groups. Our results are robust to using these narrower control groups compared to using the generic group. Results for the LCT-group-specific controls are shown in Figure A.1 in Appendix A.

1.5 Empirical findings on households' consumption response to prices

Figure 1.3 displays the price elasticity coefficients of our baseline regression model. The coefficients describe how households react when electricity prices marginally increase in period 0h. The figure's x-axis shows the seventeen half-hourly periods before and after the price increase. The time lags and leads capture whether households shift consumption to adjacent periods in response to a price increase in 0h. The y-axis depicts the price elasticity of electricity consumption.

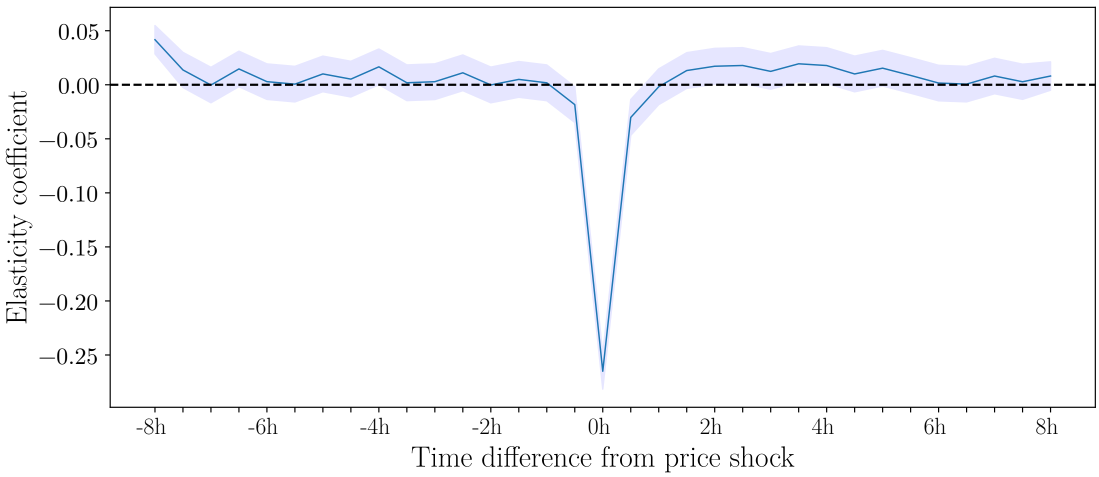

Figure 1.3: Consumption response due to price shock at period 0h

Figure 1.3 reveals that households react substantially to a price increase in period 0h with an average own-price elasticity of -0.26. The figure also suggests that customers modestly increase demand in adjacent periods when the price increases in period 0h. However, these cross-price elasticities are generally not statistically significant – see Table

A.1 in Appendix A. Thus, customers' willingness or ability to shift consumption over time appears limited.

Figure 1.4 highlights that customers' response to prices only moderately differs across seasons. Average own-price elasticities are slightly larger in summer (-0.297) and winter (-0.273) than in spring (-0.246) and fall (-0.223). As Table A.1 in Appendix A reveals, the differences in own-price elasticity between summer and winter and fall or spring are statistically significant. Households might be more able to adjust their consumption in seasons with more extreme temperatures since they can adjust their use of air conditioning and electric heating. However, differences across seasons are overall minor.

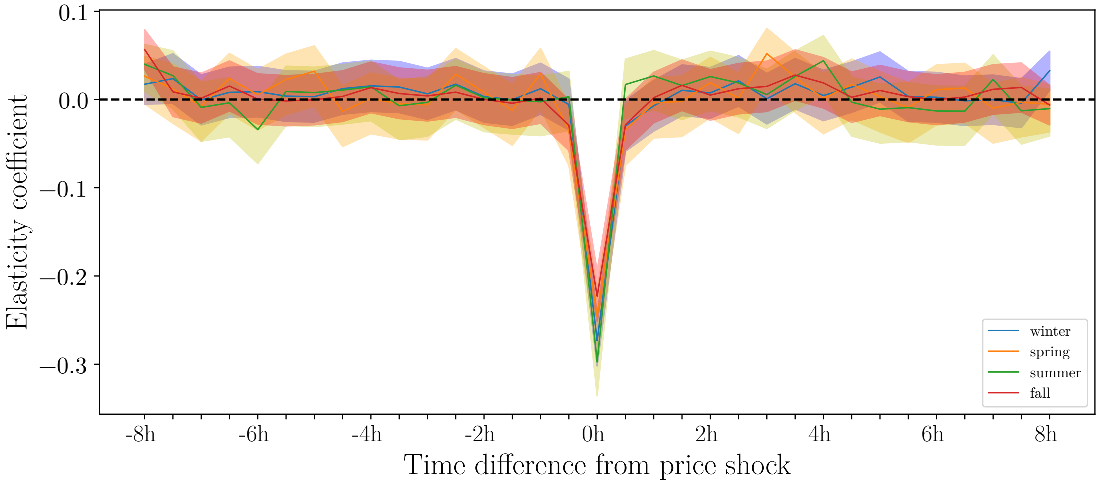

Figure 1.4: Consumption response due to price shock at period 0h by season

In contrast, the consumption response varies significantly over the day, as Figure 1.5 illustrates. Customers' response to price changes is significantly stronger at night (the own-price elasticity is -0.302 from midnight to 5:30) compared to all other times of day (see Table A.1 for the precise estimation results). The price response in the evening (-0.240 from 18:00 to 23:30) is also significantly larger than the response in the afternoon and morning. The own-price elasticities in the afternoon (-0.157 from noon to 17:30) and morning (-0.154 from 6:00 to 11:30) are smaller and statistically indistinguishable from each other. Customers seem less willing to adjust their electricity usage in the morning and afternoon. Postponing electricity-consuming activities during these times of day might be impractical and too costly due to fixed working hours.

Low-carbon technologies: Octopus Energy surveys customers who sign up for a dynamic rate and collects ownership information of low-carbon technologies (LCT). We use this data to analyze whether LCT ownership impacts customers' reactions to price

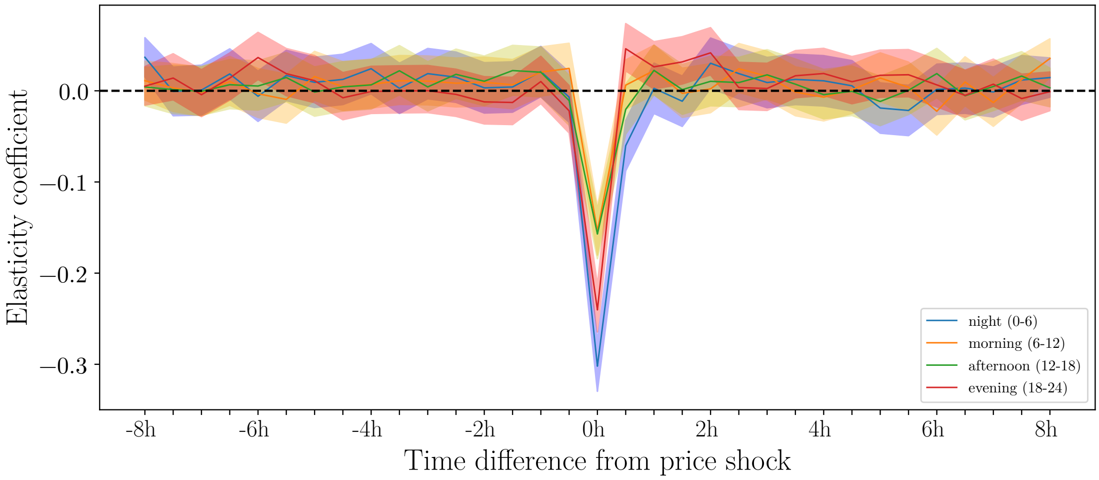

Figure 1.5: Consumption response due to price shock at period 0h by time of day

changes. We consider five LCTs: 1) electric heating, 2) smart thermostats, 3) electric vehicles, 3) residential solar, and 4) battery storage.

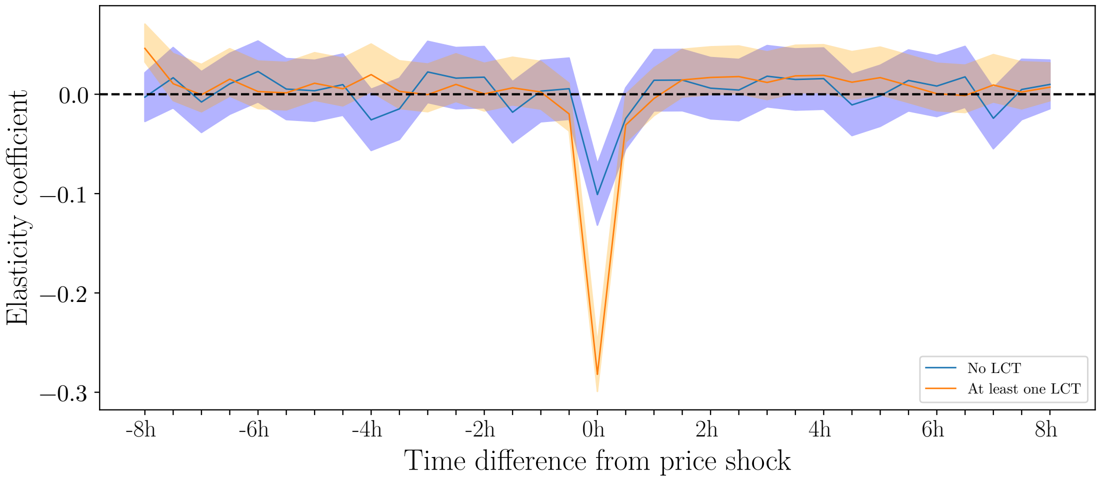

Figure 1.6: Consumption response due to price shock at period 0h by low-carbon technology status

Figure 1.6 shows that customers owning at least one LCT are almost three times as price responsive as customers who do not own any LCT. No-LCT customers have an own-price elasticity of -0.101 compared to -0.282 for customers with at least one LCT.

Next, we consider specific combinations of LCTs. For conciseness, Figure 1.7 only shows the own-elasticity coefficients. Table A.1 in Appendix A displays the complete estimation results. The left side of Figure 1.7 indicates that low-carbon heating technologies do not increase customers' price responsiveness. The own-price elasticities are nearly the same for customers with electric heating or electric heating and smart thermostats (and no other LCT) as for customers without LCTs.

On the right side of Figure 1.7, we focus on the effect of the remaining LCTs, namely electric vehicles, solar PV, and batteries.[4] The own-price elasticity of customers who only have an electric vehicle (EV only) is more than three times larger than that of customers without any LCT. Electric cars require substantial energy, motivating consumers to charge when electricity is cheaper. Pairing electric vehicles with solar and batteries increases price responsiveness even further.

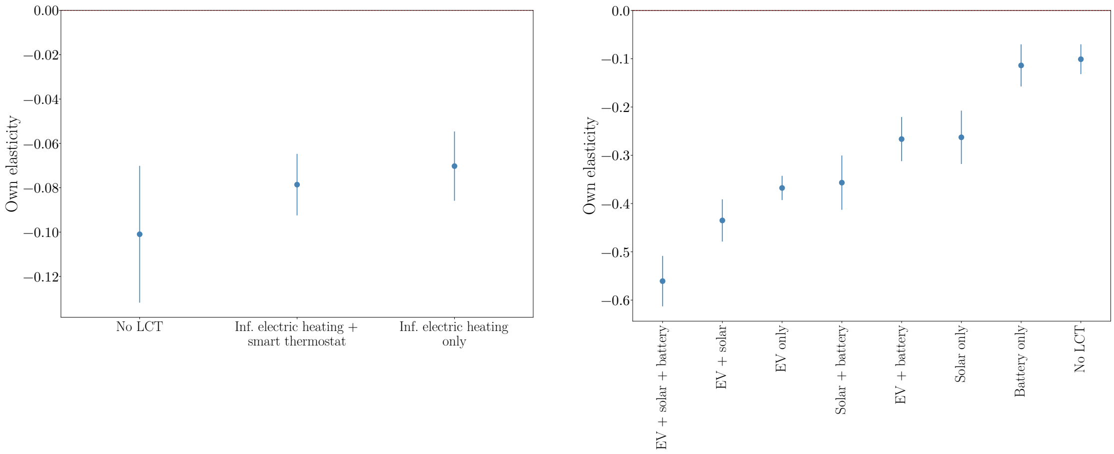

Figure 1.7: Consumption response by low-carbon technology ownership status

Solar PV ownership also has a substantial effect on price responsiveness. Surprisingly, batteries alone do not make customers more price-responsive than customers without LCT. Moreover, customers who pair batteries with electric vehicles or solar PV are not significantly more price responsive than customers who only use an EV or a solar panel. Batteries are not yet endowed with software to take advantage of price changes because dynamic rates remain rare (Green & Staffell, 2017).

1.6 Demand response and resiliency

In this section, we simulate a thought experiment to analyze if price-responsive consumers can help make the power system resilient to extreme weather events like the Texas winter

[4]To increase the size of our sub-samples, we do not control for ownership of the low-carbon heating technologies. For instance, "EV only" customers have an electric vehicle, no solar PV, and no battery, but they may have electric heating or own a smart thermostat. Low-carbon heating technologies do not have a significant effect on price responsiveness, as discussed above.

storm in February 2021. The left panel of Figure 1.8 illustrates a typical winter peak in Texas. There is ample supply to meet demand. The market clears at a typical price under \$40 per megawatt-hour. The right panel in Figure 1.8 shows the supply and demand picture during the height of the winter storm crisis. Generation outages caused supply to fall about 35 gigawatts less than expected. Electric heating caused demand (the red dashed line) to surge about 20 gigawatts higher than prior winter peaks. The gap between demand and supply during the height of the storm was this difference between the red dashed line and the maximum supply: $76 - 48 = 28$ gigawatts or 37% of demand.

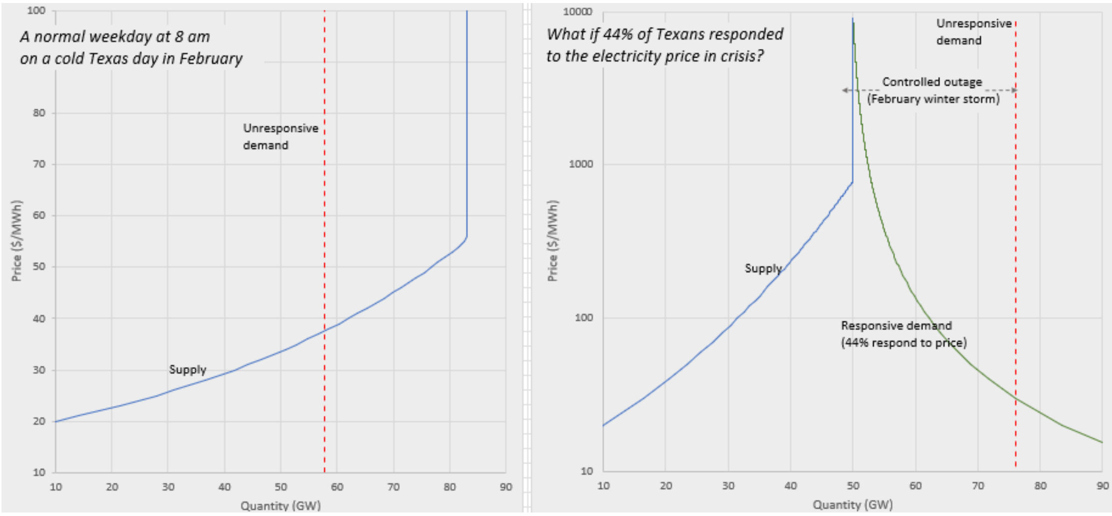

Figure 1.8: Price-responsive demand improves resiliency

When supply and demand curves do not intersect, the Texan system operator, ERCOT, sets a high administrative shortage price of \$9000 per megawatt-hour. However, domestic demand was unresponsive to this shortage price during the winter storm because 99.75% of Texas households had a fixed-rate plan.[5] These households paid a fixed price of about \$110 per megawatt-hour, only 1.2% of electricity's value in a crisis.

In our thought experiment, we investigate whether the outage could have been avoided if a larger share of consumers had been exposed to the shortage price during the winter storm. We calculate the share of price-responsive consumers necessary to prevent outages if the price-responsive consumers in Texas were as price-elastic as the UK consumers in our sample, with an average price elasticity of -0.26. For this calculation, we assume demand to be isoelastic, i.e., consumers' price elasticity is always -0.26 for all price levels.

[5]Personal communications with the two service providers offering dynamic rates revealed that fewer than 50 thousand out of 26 million households had wholesale-passthrough rates at the time of the winter storm.

We find that if 44% of Texan consumers had a price-responsive rate and responded as the UK consumers in our sample with a price-elasticity of -0.26, the need for controlled outages is eliminated even at the height of the storm. As the green line in the right panel of Figure 1.8 illustrates, a price increase causes electricity demand to fall if 44% of consumers are price-responsive. Falling demand shrinks the gap between supply and demand. The gap vanishes at the clearing price of $9000 per megawatt-hour. Thus, if a sizeable minority of Texans had been price-responsive, Texas would have survived the 2021 storm without shortage.

Figure 1.9 illustrates the share of price-elastic consumers required to prevent power outages during the Texas Winter storm based on varying levels of average price elasticity among price-responsive households in our thought experiment. The results indicate that if 79% of Texan households exhibited price elasticity like the UK households in our sample without low-carbon technology (an average elasticity of 0.1 in absolute terms), the outages could have been averted. Similarly, if 39% of Texan households had displayed an average price elasticity of 0.37, as seen among EV-only owners in our sample, the outages could have been avoided. Furthermore, if 36% of households had been as price-elastic as those owning EV, solar, and battery (with an elasticity of 0.56 in our sample), the outages could also have been prevented.

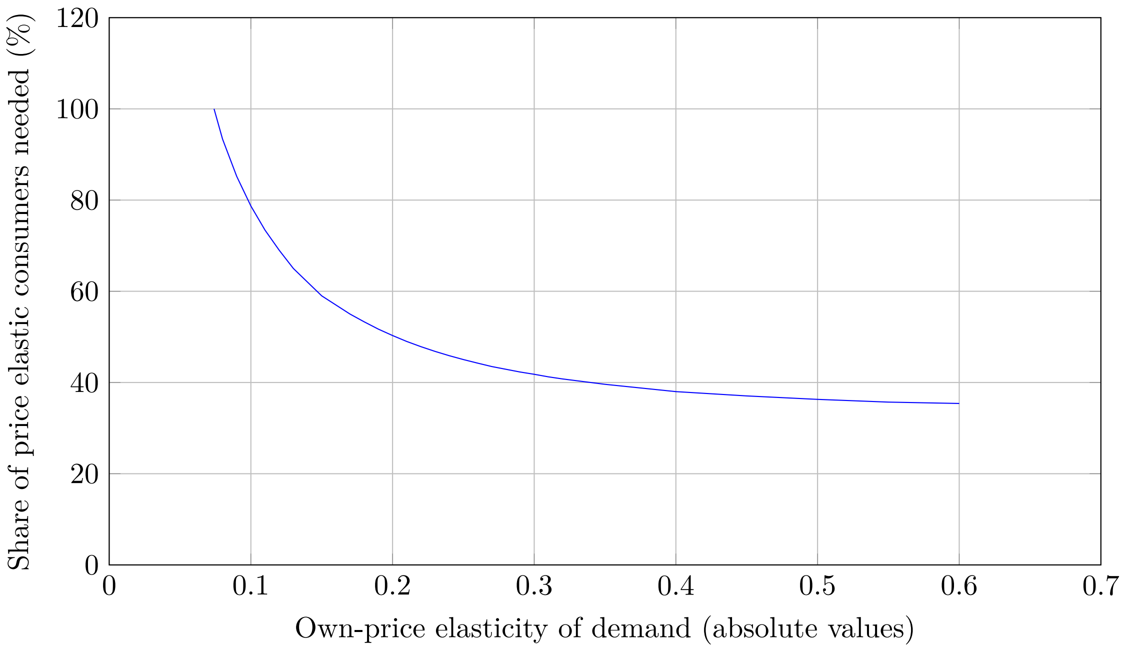

Figure 1.9: Share of price elastic consumers needed to avoid outages during the Texas winter storm (%)

This thought experiment rests on a strong assumption: the price-responsive customers in Texas have the same average price elasticity as the British customers in our sample. However, prices in Texas varied by a factor of eighty, while they only varied by a factor of three for the British sample. During our study, the British retail price was capped at 35 pence per kilowatt-hour (£350/MWh). Whether Texan consumers would react more or less strongly to the far higher price increases during the winter storm remains an open question. It also rests on the strong assumption that supply is unresponsive to demand elasticity, which appears unlikely in the long run.

We emphasize that the point of this thought experiment is not to precisely estimate the price-responsive demand needed to survive the winter storm Uri. Instead, it demonstrates the vital role engaged and price-responsive customers can play in making power systems resilient during extreme weather events.

1.7 Policy levers and consumer protection

Regulators are concerned that dynamic rates make households vulnerable to extremely high prices, and rightfully so. The 0.25% of Texas consumers with dynamic rates generated widespread news coverage and public interest. Sensationalist reports of multi-thousand-dollar bills were frequent, although these bills were rare. While these dynamic rates save consumers money in the long run, they are problematic during a crisis if they are not introduced with safeguards in place. The Texas Legislature's first law addressing the crisis was to ban dynamic rates (Ferman, 2021).

Regulators in other jurisdictions still allow dynamic rates but implement or plan to implement measures to protect consumers from high electricity bills by limiting their exposure to dynamic prices. For example, in response to the EU energy prices following Russia's invasion of Ukraine, the European Commission recently stated that a key objective of the envisioned electricity market design reform is to "enhance the protection of consumers from volatile prices and to empower them with greater contract choice" (European Commission, 2023) and that "this wider choice will allow consumers, if they wish, to lock in secure, long-term prices to be shielded from sudden price shocks." (European Commission, 2023). Similarly, in the context of our study, the dynamic retail price in Britain was capped at £350 per megawatt-hour in 2021. The cap is effective in protecting consumers from high electricity bills. However, it is too low to ensure demand response during a crisis. A price of £350 per megawatt-hour might not be high enough to induce households to lower their heating in a winter storm.

Moreover, with retail prices capped at £350 and wholesale prices uncapped, sustained scarcity can cause provider bankruptcies. Britain's 2021 energy crisis demonstrated this vulnerability. A sustained high gas price has led to high wholesale electricity prices, bankrupting inadequately hedged providers (Thomas, 2021). The California 2000-2001 energy crisis had the same root cause (Cramton, 2017).

Instead of banning or capping dynamic rates, regulators should mandate forward hedging in plans that expose consumers to the wholesale spot price. A hedge is commonplace among industrial consumers and can be implemented in understandable ways for retail consumers. The service provider buys forward the consumer's expected consumption and then only exposes the consumer on the margin to settle deviations from expected consumption. The real-time price rewards the consumer for consuming less during a crisis. In this way, a penalty (a large bill) becomes a reward (a large rebate).

For example, consider an electric-heated home on a fixed rate. Suppose the home is near a hospital, so the electricity stayed on throughout the storm and consumed 0.4 megawatt-hours during the 4-day storm, double its typical usage. The household pays $0.4 \times \$110 = \$44$ for electricity during the event.

By contrast, suppose the household had a dynamic rate with hedging. The consumer pays the $\$9,000$/MWh price only for deviations from expected consumption. The service provider bought the consumer's expected demand of 0.2 megawatt-hours at the \$110 forward price, a cost of $0.2 \times \$110 = \$22$. The high marginal price motivates the consumer to put jackets on, turn down the thermostat, and consume only one-half of the typical amount. The responsive consumer's bill for the crisis is $\$22 - 0.1 \times \$9000 = -\$878$.

Instead of exposing the consumer to downside risk, the price-responsive consumer enjoys the opportunity to be rewarded for being flexible and making a socially beneficial decision – consuming less so that others can warm their houses, too. Hedging transforms downside risk into upside opportunity. Poorer consumers are more price-sensitive because they spend more of their income on electricity and are apt to benefit the most from the flexibility option. Furthermore, hedging brings an additional resiliency benefit by reducing the chances of service provider default during extended periods of high wholesale prices.

The Texas Legislature's ban was a short-termist political response to the crisis. Instead, the Legislature should instruct the regulator to prohibit unhedged plans but welcome plans with hedging as an essential innovation. Indeed, the two Texas service providers offering dynamic plans were about to introduce improved plans with hedging when the storm struck (Cramton, 2022).

Rather than limit marginal exposure to spot prices, a necessary condition of economic efficiency, regulators should foster resiliency and social justice with low-income subsidies that promote energy efficiency. The low-hanging fruits are improved insulation, caulking, and other energy efficiency programs. The levelized cost of energy ($/MWh) of energy efficiency programs in the US ranges from $12-49, with an average of $24 (Cohn, 2021). By contrast, the levelized costs of energy for the most efficient production technologies are $29-42 for solar photovoltaic, $26-54 for wind, and $44-73 for gas combined cycle (Lazard, 2020). Mandates and low-income subsidies for high-efficiency heat pumps and other appliances are also desirable. These steps create the most significant savings during extreme weather events – and help achieve intelligent demand in the long term. Peak demand falls, reducing the need for additional electricity infrastructure.

Many markets, especially in Europe, have taken great strides in promoting energy efficiency through mandates and subsidies. Germany is a good example. By contrast, much of the US still has poor energy efficiency, as seen in Texas. The reason is simple. Traditional energy providers, the loudest policy voice, do not benefit from energy efficiency. Consumers benefit, but the gains are often hidden and distant.

Apart from energy efficiency, a thoughtful acceleration of enabling policies is needed to foster the rapid adoption of dynamic pricing. To allow price-responsive demand, we must measure real-time electricity use. Smart meters need to be installed in each home. This essential step has been progressed in many US markets and has been completed in countries including Denmark, Sweden, and Finland, leaving the UK, Germany, and others lagging.

Regulators can also promote price-responsive demand by supporting retail choice so consumers can select a competitive rate that allows price response or by requiring utilities to offer a dynamic rate with hedging to manage risk. As more consumers choose the dynamic rate, resiliency improves. Regulators could even nudge consumers by making the dynamic rate with hedging the default rate for those with supporting devices, such as smart thermostats or electric vehicles. The dynamic rate is the default if the consumer fails to select a rate (Berger et al., 2023).

Texas illustrates the enormous cost of a multi-day outage in our modern world. Regulators can mitigate this cost with policies that foster electricity resiliency and the benefits that flexibility affords as we power ourselves with an increasing share of renewables. Price-responsive demand and energy efficiency are low-hanging fruit.

1.8 Conclusions

In this paper, we stress the vital role of engaged and price-responsive electricity consumers for a successful green transition. This is a valuable, low-cost tool to help keep supply and demand in balance in a system powered increasingly by intermittent renewables. Using domestic customer-level data from the UK, we first show that households are able and willing to adjust their consumption in response to dynamic electricity prices. On average, the price elasticity of demand is -0.26 for the households in our sample.

While this effect may partly reflect self-selection, rapid innovation in demand-response technology and digitalization of energy systems will increase adoption and enable greater price-responsive demand among all households. We provide suggestive evidence that low-carbon technologies like electric vehicles significantly strengthen households' price response. As more consumers adopt these technologies in the coming decades, demand response from domestic consumers is expected to change and increase further. Therefore, it should be a key focus of future research to study in a dynamic model how households' response to electricity changes over time when an increasing share of consumers starts using these technologies.

Several policy implications emerge. Most importantly, regulators should mandate electricity suppliers to combine dynamic rates with a forward hedge. The hedge shields households from high prices while preserving an incentive to be price-responsive, especially during shortage events. The hedge makes dynamic pricing attractive for low-income households since it rewards them for reducing consumption during shortages.

Regulators should also combine dynamic pricing with low-income subsidies for energy efficiency investments. Dynamic pricing makes investments in energy efficiency measures like insulation more appealing since energy efficiency protects households from high prices. These subsidies ensure that low-income households also benefit from dynamic pricing.

Regulators need to implement enabling policies for rapid adoption of demand response. For instance, they should support an accelerated smart meter roll-out and intensify retail competition to encourage electricity suppliers to offer innovative, dynamic plans.

Implementing the above regulatory measures is essential to make dynamic electricity rates with a forward hedge appealing to consumers. Attractive dynamic rate plans encourage domestic consumers to actively participate in the energy transition. They induce consumers to mitigate demand peaks and help prevent blackouts during extreme weather events. Demand response also motivates households to invest in low-carbon technologies like electric vehicles, battery storage, and energy efficiency - which are essential to decarbonization. Dynamic electricity prices are vital for a fast, efficient, and resilient green transition and should be rapidly adopted at scale.

We also highlight that price-responsive consumers can make power systems resilient to extreme weather events like the Texas winter storm. In a thought experiment, we estimate that power outages could have been avoided during the winter storm if 44% of

Texan electricity consumers had been as price-responsive as the consumers in our UK sample. Dynamic prices give households a strong incentive to reduce their consumption when needed most, namely during extreme weather events. The broader adoption of electric heating and cooling will make price-responsive households even more crucial for system resiliency since domestic electricity consumption will likely correlate more positively with extreme weather (Staffell et al., 2015).

However, the Texas winter storm also revealed that households cannot be fully exposed to extremely high dynamic prices without proper risk protection. Hence, we argue that regulators should mandate suppliers to combine dynamic rates with a forward hedge. The forward hedge protects households from soaring prices and preserves an incentive to be price-responsive, especially during extreme events. Moreover, regulators should combine dynamic electricity rates with additional policy levers to address social justice concerns like low-income subsidies for energy efficiency measures. Regulators should also implement policies for a faster smart-meter rollout and intensified retail competition to promote demand response.

Overall, price-responsive demand is fundamental for a reliable and resilient energy system. Electricity, heat, and transport account for over 60% of global emissions (EPA, 2021). Dynamic rates strengthen incentives to adopt low-carbon technologies in all these sectors. For instance, dynamic pricing supports investments in renewable generation assets since it incentivizes consumers to shift consumption to periods when energy from renewables is abundant. Similarly, price-responsive demand increases incentives to invest in slow-ramping resources, like nuclear, because part of the adjustment to balance supply and demand shifts to the demand side (Roques et al., 2005).

In the transport sector, dynamic electricity rates and electric vehicles exhibit strong complementarity. Electric vehicles motivate consumers to be price-responsive since they benefit from charging at night when the price is low. On the other hand, dynamic pricing is a precondition for smart charging and discharging of electric vehicles. Electric vehicles can generate additional value via vehicle-to-grid solutions. In combination with dynamic pricing and battery storage, electric vehicles also have enormous potential for smoothing the intermittent generation from renewables. By 2050, electric vehicles are projected to increase electricity demand by 25% (Steinberg et al., 2017). Dynamic pricing can transform electric vehicles from a challenge to an opportunity for the power grid.

Dynamic electricity prices thereby help decarbonize not only the electricity sector but also other sectors of the economy. As such, dynamic pricing should be a central component of any least-cost strategy aimed at accelerating the green transition.

Chapter 2

Hedging households against extreme electricity prices

This chapter contains the current version of the working paper Brandkamp (2025). An abstract of this working paper was also published in the conference proceedings of IAEE (2024) and EAERE (2024).

Abstract

Dynamic electricity prices expose households to the risk of extremely high electricity bills during scarcity events. To protect households from high scarcity prices, I explore how to combine dynamic electricity prices with forward hedging. I derive household-specific optimal forward hedge shares by applying a utility maximization model to 2,159 UK households exposed to dynamic prices. The average optimal hedge share is 59% of households' baseline consumption. Hedge shares are higher for electric heating and electric vehicle owners and lower for solar PV and battery storage owners. My key theoretical finding is that an increase in households' price elasticity of demand raises optimal hedge shares if households face positive correlation between electricity prices and their weather-related desire to consume electricity. Forward hedging effectively reduces electricity bill volatility by 18% for price-inelastic households. When exposing households to scarcity events, hedging achieves sizable welfare gains equivalent to 19% reduction in average electricity prices.

Acknowledgments

I thank Emmanuele Bobbio, Peter Cramton, David Malec, Axel Ockenfels, Oliver Ruhnau, and seminar participants of the EWI seminar at the University of Cologne, the YEP ECONtribute seminar at the universities of Bonn and Cologne, and the YEEES32 workshop at the FAU in Nuremberg for helpful comments.

Declarations

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. This research is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1-390838866.

2.1 Introduction

Many economists emphasize the benefits of exposing households to dynamic electricity prices. Dynamic prices are linked to day-ahead electricity prices and fluctuate on a short-term basis, e.g., every half hour. They incentivize households to reduce consumption during periods of high aggregate demand or scarce electricity supply. Thereby, dynamic prices help mitigate demand peaks and improve the integration of intermittent renewable generation into the energy system. This makes electricity markets more efficient (Houthakker, 1951, Borenstein & Holland, 2005, Borenstein, 2005, Allcott, 2011).[1]

A drawback of dynamic pricing is that it makes consumers vulnerable to extremely high prices when electricity is scarce. When power demand exceeds supply, regulators

[1]Borenstein (2005) finds that broad adoption of dynamic prices would lead to substantial efficiency gains in the long run, reducing consumers' total electricity bill by 3-11% in California. In the short run, Holland and Mansur (2006) obtain only small welfare gains for consumers that translate to a 0.24-2.5% reduction in total electricity bills. Holland and Mansur (2006) also reveal that dynamic prices reduce the volatility of wholesale electricity prices and load.

set high scarcity prices for electricity in many power markets (Cramton, 2017).[2] For example, during the winter storm in Texas in February 2021, a scarcity price of 9$/kWh was enforced for more than 64 hours over multiple days in response to a severe electricity shortage (ERCOT, 2022). As a result, the small minority of households exposed to dynamic prices received electricity bills exceeding $100 per day (EIA, 2022). These extremely high bills sparked a public debate about whether dynamic pricing is suitable for households (McDonnel et al., 2021).

To address these concerns, economists proposed to complement dynamic price tariffs with a forward hedge (Borenstein, 2007a, Wolak & Hardman, 2022, Winzer et al., 2024, Schlecht et al., 2024). The objective of the forward hedge is to protect households from high electricity prices while preserving incentives to reduce demand when needed most, namely during scarcity events (Cramton et al., 2025a). The German government recently introduced subsidized forward contracts known as gas and electricity price brakes to reduce consumers' exposure to soaring energy prices while incentivizing them to save scarce energy (Dertwinkel-Kalt & Wey, 2023).

To illustrate how a forward hedge works, consider a household with a typical daily consumption of 20 kWh. The household is offered to buy 100% of its typical consumption forward at a price of 0.2$/kWh. It pays 4$ (=20kWh*0.2$/kWh) per day for this hedge. During a Texas-style scarcity event, when the electricity price is 9$/kWh all day, the household pays the scarcity price only for the deviation of its realized consumption from the hedged quantity of 20 kWh. If the realized consumption is 20kWh, the household only pays 4$ for the hedge and is fully protected from the high scarcity price.

If the household reduces its consumption below the hedged quantity, e.g., to 15 kWh, it is remunerated for the negative deviation of 5 kWh at the scarcity price. This results in a daily bill of $-41\$ \ (= 4\$ + (-5\text{kWh}) * 9\$/\text{kWh})$. The hedge rewards households for doing what is socially desirable: reducing consumption when electricity is scarce. Conversely, if the household consumes more than 20 kWh, say 25kWh, the positive deviation of 5kWh leads to a high daily bill of 49$ ($=4\$ + 5\ \text{kWh} * 9\$/\text{kWh}$). This high bill creates a strong incentive to reduce consumption during scarcity. At the same time, the hedge protects the household from the far higher 225$ bill it would face without the hedge. Moreover, the household can choose to hedge more than 100% of its typical consumption to be better protected against high prices (Cramton et al., 2025a).

In this paper, I calculate household-specific *optimal hedge shares*, i.e., the optimal share of typical consumption that a household should buy forward. I simulate these

[2]The objective of scarcity prices is to induce all available power plants to operate at maximum level and to provide incentives for investments in additional generation capacity (Cramton, 2017).

optimal shares for a dataset of 2,159 UK households that received half-hourly dynamic prices but no forward hedge.

A hedge share is defined as *optimal* if it maximizes households' expected utility. Optimal hedge shares may vary by time of day to account for intraday load patterns. I study how optimal hedge shares are influenced by households' price elasticity of demand, risk aversion, and ownership of technologies like electric heating and vehicles, and battery storage. To estimate welfare gains, I calculate the price premia households should be willing to pay for being optimally hedged.

To derive the optimal hedge shares, I solve a two-stage utility maximization model via backward induction. In the second stage, households do not face any uncertainty. They maximize their utility from electricity consumption and aggregate consumption of other goods. I assume a constant-elasticity-of-substitution utility function and exogenously choose households' elasticity of substitution and their coefficient of relative risk aversion. Moreover, I suppose that each household spends, on average, 2% of its income on electricity since I observe household expenditure only for electricity and not for other goods. This assumption is consistent with electricity's average expenditure share in the UK (UK ONS, 2021).

In the first stage of the model, households choose optimal hedge shares subject to stochastic electricity prices and quantity shocks. Quantity shocks capture all factors except prices that impact agents' desire to consume electricity. For instance, the weather or a national holiday influence households' desire to consume. The model allows deriving the unobservable quantity shocks as a residual of households' observable unhedged electricity demand.

My main contribution is to examine the interaction between price elasticity of demand and optimal hedging when agents face two volatile factors: Prices and quantity shocks. Previous literature analyzed how price elasticity affects optimal hedge shares when prices are the unique volatile factor (Moschini & Lapan, 1992, Dionne & Santugini, 2015). They argue that optimal hedge shares decline if agents' price elasticity increases. Price-elastic agents are less vulnerable to high prices because they can reduce demand when prices are high. Hence, price-elastic agents need to hedge less when prices are the unique volatile factor.

However, electricity consumers also face volatile quantity shocks. Their desire to consume electricity depends heavily on external shocks like extreme weather. Only a few authors consider how quantity shocks influence hedge shares (Losq, 1982, Cowan, 2004). They find that optimal hedge shares increase if prices and quantity shocks are positively correlated.

My key result is that this correlation also determines how an increase in price elasticity affects hedge shares. When prices and quantity shocks are positively correlated, an increase in price elasticity raises hedge shares. Assume price elasticity increases for a household with electric heating. During a winter storm, the increased price elasticity induces this household to reduce consumption more strongly in response to high scarcity prices. However, reducing consumption during a winter storm causes large disutility for an electric heating owner since she needs electricity to heat her home. As a response, the more price-elastic electric heating owner increases her hedge share to reduce her exposure to spot prices. Lower spot price exposure allows her to maintain an acceptable level of consumption even during a winter storm when both prices and her desire to consume electricity are extremely high.

In contrast, when prices and quantity shocks are negatively correlated, an increase in price elasticity decreases hedge shares. In this case, responding more strongly to high prices causes small disutility because households' desire to consume is low when prices are high. Hence, increasing price elasticity makes households less vulnerable to high prices since they respond by reducing consumption without suffering large disutility. Therefore, they need to hedge less.

From these theoretical observations, I calculate optimal hedge shares for each household in my sample. The average optimal hedge share is 59% of households' typical consumption. Optimal hedge shares differ widely both between customers and for a specific customer by time of day. Ownership of low-carbon technologies partly explains the heterogeneity in hedge shares. Electric heating or electric vehicle owners have high average hedge shares of 77% and 69%, respectively. These technologies likely increase the positive correlation between prices and quantity shocks. On the other hand, solar PV owners hedge on average only 22% of their typical consumption, indicating that their electricity consumption from the grid is negatively correlated with prices. Battery storage ownership is also associated with a slightly lower mean hedge share of 52%.

An exogenous increase in price elasticity of demand further amplifies the dispersion in hedge shares. Some households respond to higher price elasticity by reducing their hedge shares to exploit the negative correlation between prices and quantity shocks. Other households with a positive correlation increase hedge shares. On average, optimal hedge shares decline when price elasticity rises. In contrast, increasing risk aversion has a positive effect on hedge shares. This positive effect rapidly diminishes with the level of risk aversion.

The optimal hedge is effective in reducing the volatility of monthly electricity bills. In my main specification, the optimally hedged tariff reduces households' coefficient of variation of monthly bills, on average, by 18% compared to an unhedged tariff. Hedging is particularly effective for households with large average hedge shares beyond 100% of their typical consumption. For these households, the hedge reduces their coefficient of bill volatility by 45%. Such households even experience lower bill volatility when optimally hedged than under a fixed-price tariff. The fixed tariff protects households only from price volatility, while the optimally hedged tariff protects from volatility in both prices and consumption (Borenstein, 2007a).

An increase in price elasticity makes the optimal hedge less effective in reducing bill volatility. Price-elastic households consume more when prices are high on the optimally hedged tariff than on the unhedged tariff because the hedge weakens households' exposure to spot prices and lowers their price response. Consuming more when prices are high raises bill volatility. When assuming high price elasticity, some households even face higher bill volatility on the optimally hedged tariff than on the unhedged one.

Similarly, if households become more risk-averse, the optimally hedged tariff achieves smaller reductions in bill volatility compared to the unhedged tariff. High risk aversion induces households to choose higher hedge shares and thereby lowers households' exposure to spot prices. As explained above, low spot price exposure can increase bill volatility for price-elastic households.

Optimal forward hedging results in tiny welfare gains for households during my sample period. Augmenting a dynamic tariff with a forward hedge leads to a welfare gain compared to an unhedged dynamic tariff that translates, on average, to a 0.3% reduction in mean electricity prices. In contrast, dynamic pricing itself achieves more significant average welfare gains equivalent to a 0.8-4% reduction in mean electricity prices.

The welfare gain from hedging increases to a more substantial 19% reduction in mean electricity prices when simulating a scarcity event like the Texas winter storm. Hedging is more valuable to consumers if they face worst-case events with extremely high prices. I simulate such a scarcity event by artificially exposing households' to extremely high electricity prices over multiple days.

The small welfare gain from hedging without scarcity events can be partly explained by electricity's small share in household expenditure. UK households spend only 2% of their income on electricity (UK ONS, 2021). Electricity's expenditure share will likely increase substantially in the upcoming decades due to electrification of mobility and heating. Therefore, the relevance of hedging electricity prices might increase. However, for the given sample, the average simulated welfare gain from hedging only increases to 1.6%, even when assuming that households spend a much larger income share of 10% on electricity.

My paper presents a partial equilibrium analysis. Households' hedging decisions do not affect electricity prices. In general equilibrium, where all consumers are on dynamic and optimally hedged tariffs, the welfare effects of hedging are likely lower. Dynamic pricing mitigates price peaks and, thereby, reduces price volatility. When price volatility is lower, hedging creates smaller welfare benefits.

I structure my paper as follows: Section 2.2 relates this paper to the optimal hedge literature. Section 2.3 describes a utility maximization model to derive optimal hedge shares and price premia for the optimal forward hedge. In section 2.4, I parameterize the model to simulate optimal hedge shares for real-world domestic consumers. Section 2.5 presents the data set, and section 2.6 discusses the simulation results. Section 2.7 concludes.

2.2 Literature review

This paper contributes to the literature on optimal forward hedge shares. Danthine (1978), Holthausen (1979), and Feder et al. (1980) identify optimal hedge shares and production decisions for price-inelastic firms when selling prices are uncertain. Their main insight is the "separation result": Production decisions only depend on the forward price and not on risk attitudes. In contrast, hedge shares depend on risk attitudes and expectations about future prices. Moreover, they find that price-inelastic firms should hedge 100% of their production when forward hedge prices equal expected spot prices.

McKinnon (1967), Rolfo (1980), Losq (1982), and Lapan and Moschini (1994) derive optimal hedge shares for firms in industries like agriculture where both prices and production are uncertain. McKinnon (1967) highlights that a positive correlation between prices and production is a major motive for firms to increase hedge shares beyond 100%.

The literature also explores how optimal hedge shares change when alternative risk management tools are available. For example, if firms are price-elastic, their optimal hedge share is lower since they can flexibly adjust their production after price uncertainty resolves (Moschini & Lapan, 1992, Dionne & Santugini, 2015). Similarly, storage and buffer stocks provide flexibility and reduce the need for high hedge shares (McKinnon, 1967, Newbery & Stiglitz, 1981, Gemmill, 1985, Gilbert, 1985). In the context of this paper, battery storage will likely gain importance as a risk management tool for electricity consumers.

Moreover, multiple authors derive optimal hedge shares when firms combine forward contracts with alternative derivatives like options (Moschini & Lapan, 1992, Brown & Toft, 2002, Gay et al., 2002, Mnasri et al., 2017). They argue that when the uncertainty concerning quantity is high, forward contracts are less effective than options to protect firms from profit fluctuations. The key difference is that forward contracts have a linear payoff structure, while options have a nonlinear one. When both prices and quantities are uncertain, firms' profits are nonlinear in prices. Therefore, nonlinear options provide better risk protection (Moschini & Lapan, 1992, Sakong et al., 1993). Brown and Toft (2002) customize the optimal payoff structure for a value-maximizing firm using a portfolio of forward contracts and exotic nonlinear derivatives.

Studies on hedging for small retail consumers are rare since they typically do not have access to forward markets (Newbery, 1989). Various researchers examine hedge decisions for large consumers like load-serving entities in electricity markets. Load-serving entities buy electricity on wholesale markets with volatile prices and sell it to retail consumers at fixed rates. They face high price and demand uncertainty caused by unique characteristics of electricity. Electricity demand fluctuates and is mostly price-inelastic. Large-scale storage is often unavailable or expensive. Moreover, electricity prices and demand are positively correlated as they both depend on weather conditions. Oum and Oren (2010), Zhou et al. (2017), and Azevedo et al. (2007) characterize the optimal hedge position for load-serving entities. They stress the importance of augmenting forward contracts with nonlinear options, given the large quantity uncertainty that load-serving entities face.

Another tool that load-serving entities can use to mitigate risk is time-varying tariffs like time-of-use, critical peak pricing, peak time rebate, or dynamic tariffs (Faruqui et al., 2012). Time-varying tariffs reduce risk for load-serving entities as they mitigate the positive correlation between wholesale prices and demand (Zhou et al., 2017). They do so by charging end customers higher prices when wholesale prices are (expected to be) high. Numerous studies reveal that end customers lower their electricity consumption in response to high short-term prices. In a literature survey of 36 studies, Espey and Espey (2004) report an average short-term price elasticity of electricity demand of -0.35.

On the other hand, time-varying tariffs increase the risk exposure for retail consumers, especially on fully dynamic tariffs (Borenstein, 2007a, Faruqui et al., 2012). Many economists acknowledge the need to augment dynamic tariffs with risk management tools like price caps or collars[3], forward contracts, or options (Barbose et al., 2004). Caps and collars are suitable for protecting customers from high price spikes. However, caps are typically too low to incentivize customers to lower consumption sufficiently during scarcity events (Cramton et al., 2025a). Nonlinear derivatives like options are arguably too complex, especially for households and small business consumers. Therefore, multiple

[3]Price collars allow the electricity price for retail consumers to vary, but it cannot be higher than a specified price cap or lower than a price floor (Goldman et al., 2004).

researchers propose forward contracts as an appropriate hedging tool for consumers (Borenstein, 2007a, Wolak & Hardman, 2022, Cramton et al., 2025a, Schlecht et al., 2024).

Since the 1990s, several US utilities have offered commercial customers dynamic tariffs with a forward hedge, known as two-part tariffs (Braithwait & Eakin, 2002, Barbose et al., 2005). While most utilities let customers only buy exactly 100% of their typical consumption forward, some utilities allow them to choose a different hedge share (O'Sheasy, 1998, Stavrogiannis, 2010). In this paper, I compare the effectiveness of a simple 100% hedge tariff used by most utilities to a tariff with an optimal forward hedge.

My paper is closest to Borenstein (2007a), Stavrogiannis (2010), and Winzer et al. (2024). They do not derive optimal hedge shares but simulate how effective various arbitrary hedge shares are in reducing the volatility of electricity bills. For household samples, Stavrogiannis (2010) and Winzer et al. (2024) find that a tariff with a 100% forward hedge effectively reduces bill volatility compared to an unhedged dynamic tariff. For a sample of large industrial consumers in California, Borenstein (2007a) reveals that 77% of consumers can reduce their bill volatility even further by choosing a hedge share above 100%. Such consumers face a strong positive correlation between consumption and prices. These findings suggest that a simple 100% hedge might only be optimal for a few consumers.

My paper differs from Borenstein (2007a) and Stavrogiannis (2010) as they assume consumers to be risk-neutral and price-inelastic. They suppose consumers do not change their consumption when transferred to a dynamic tariff with a forward hedge. In contrast, I study the interaction between price elasticity and hedging. Specifically, I simulate how households adjust their consumption in response to being optimally hedged.

I build on models by Gilbert (1985) and Cowan (2004) to simulate optimal forward hedge shares with quantity shocks. Cowan (2004) derives optimal hedge shares for price-elastic and risk-averse electricity consumers facing quantity shocks. Hedge shares increase in consumers' coefficient of relative price risk aversion and the correlation between price and quantity shocks. Hedge shares decline in consumers' income elasticity of demand. I extend Cowan's (2004) analysis by examining the effect of an exogenous increase in price elasticity on hedge shares. I also study how the ownership of technologies like electric heating and battery storage affects hedging decisions.

Apart from the literature on optimal hedge shares, I also contribute to the literature that studies how price stabilization and hedging influence consumer welfare (Waugh, 1944, Turnovsky et al., 1980, Newbery & Stiglitz, 1981, Gilbert, 1985, Cowan, 2004). These authors emphasize that dynamic prices can increase welfare compared to stabilized prices, even for risk-averse consumers, if they are sufficiently price-elastic. Cowan (2004) reveals that optimal forward hedging always positively impacts consumer welfare under pure price volatility. I extend his analysis by showing that the positive welfare effect of hedging declines the more negatively prices and quantity shocks correlate.

2.3 Model

2.3.1 Optimal hedge shares

Below, I employ a two-stage model based on Gilbert (1985) and Cowan (2004). It allows for deriving optimal hedge shares via backward induction, starting with the second stage.

Second stage: Households choose between electricity consumption x and aggregate consumption of other goods y to maximize the indirect utility function V:

$$V(p, b, \varepsilon, f, h) = \max_{x,y}\{U(x, y, \varepsilon)|xp + y \leq b + (p - f)h^*\} \tag{2.1}$$

Household utility depends on quantity shock ε. The quantity shock captures all factors except prices that impact households' desire to consume electricity, such as weather, national holidays, or a national sports event on TV. Households choose the optimal consumption bundle subject to their income b,[4] dynamic electricity price $p > 0$, forward hedge price f, and the optimal hedge quantity h^* that was chosen in stage 1. Prices for y are normalized to 1. The indirect utility function V is assumed to be concave in budget b with $V_b > 0$ and $V_{bb} < 0$ being its first and second derivatives with respect to b.

There is no uncertainty in the second stage. Households observe electricity prices and quantity shocks.

First stage: Households choose the optimal hedge quantity h^* that maximizes their expected utility given stochastic electricity prices $\tilde{p}$ and quantity shocks $\tilde{\varepsilon}$.

$$h^* = \arg\max_{h} E[V(\tilde{p}, b, \tilde{\varepsilon}, f, h)] \tag{2.2}$$

Following McKinnon (1967) and Lapan and Moschini (1994), prices and quantity shocks can be correlated and are described via a bivariate normal distribution. Lapan and Moschini (1994) find that optimal hedge shares are robust to assumptions on utility

[4] The budget constraint with a forward hedge can be formulated in two equivalent ways: The first formulation is $fh + (x - h)p + y \leq b$. This formulation states that the household pays fh for the forward hedge and the dynamic price p only for the deviation of its consumption x from the hedge quantity h. Rearranging the first formulation leads to the second equivalent formulation $xp + y \leq b + (p - f)h$ as given in equation (2.1) (Braithwait & Eakin, 2002).

functions and distributions of stochastic elements. For risk-averse profit-maximizing firms, they derive similar hedge shares when comparing utility functions with constant absolute risk aversion and constant relative risk aversion, respectively, both under normal and lognormal distributions of prices and stochastic production.

Assuming an interior solution, the first-order condition of equation (2.2) is given by[5]

$$E[V_b(\tilde{p} - f)] = 0 \tag{2.3}$$

Letting $\overline{p} = E[\tilde{p}]$ and $\overline{\varepsilon} = E[\tilde{\varepsilon}]$ represent expectations of prices and quantity shocks, a first-order Taylor approximation of V_b about $(\overline{p}, \overline{\varepsilon})$ yields

$$V_b \approx \overline{V}_b + \overline{V}_{b\tilde{p}}(\tilde{p} - \overline{p}) + \overline{V}_{bb}h(\tilde{p} - \overline{p}) + \overline{V}_{b\tilde{\varepsilon}}(\tilde{\varepsilon} - \overline{\varepsilon}). \tag{2.4}$$

$\overline{V}_b = V_b(\overline{p}, b, \overline{\varepsilon}, f, h)$ is the first derivative of household's baseline utility level, i.e., its utility from choosing the optimal consumption bundle when the random factors $\tilde{p}$ and $\tilde{\varepsilon}$ are at their average. $\overline{V}_{b\tilde{p}}$, $\overline{V}_{bb}$, and $\overline{V}_{b\tilde{\varepsilon}}$ denote derivatives of $\overline{V}_b$ with respect to $\tilde{p}$, b, and $\tilde{\varepsilon}$. In Appendix B.1, I insert equation (2.4) into (2.3) and solve for the optimal hedge quantity h^*.

$$h^* = -\frac{\overline{V}_{b\tilde{p}}}{\overline{V}_{bb}} - \frac{\overline{V}_{b\tilde{\varepsilon}}}{\overline{V}_{bb}}\frac{\sigma_{\tilde{p}\tilde{\varepsilon}}}{\sigma^2_{\tilde{p}}} - \underbrace{\frac{\overline{b}}{\theta}\frac{(f - \overline{p})}{\sigma^2_{\tilde{p}}}}_{\text{Speculation}} \tag{2.5}$$

$\sigma_{\tilde{p}\tilde{\varepsilon}}$ is the covariance between prices $\tilde{p}$ and quantity shocks $\tilde{\varepsilon}$, and $\sigma^2_{\tilde{p}}$ represents the variance of prices. $\theta = -\frac{\overline{V}_{bb}}{\overline{V}_b}b$ is the coefficient of relative risk aversion at baseline.

Equation (2.5) reveals that h^* increases in $\overline{V}_{b\tilde{p}}$ because $\overline{V}_{bb} < 0$. If $\overline{V}_{b\tilde{p}} > 0$, a high price is associated with a high marginal utility of income $\overline{V}_b$. In this case, a large forward hedge quantity h^* increases utility because the forward hedge increases income when prices are high. Increased income causes large utility gains since $\overline{V}_b$ is high when prices are high (Cowan, 2004).

$\overline{V}_{b\tilde{\varepsilon}}$ captures how quantity shock $\tilde{\varepsilon}$ influences the marginal utility of income $\overline{V}_b$. If $\overline{V}_{b\tilde{\varepsilon}} > 0$ and $\sigma_{\tilde{p}\tilde{\varepsilon}} > 0$, a large $\tilde{\varepsilon}$ is associated with both high $\overline{V}_b$ and high prices in expectation. In this case, choosing a large hedge quantity h^* is beneficial because the hedge increases income when both $\overline{V}_b$ and prices are likely high. In contrast, if $\overline{V}_{b\tilde{\varepsilon}} > 0$ and $\sigma_{\tilde{p}\tilde{\varepsilon}} < 0$, the agent hedges less since $\tilde{p}$ is likely low when $\tilde{\varepsilon}$ is large. When a large $\tilde{\varepsilon}$

[5] The second-order condition is $E[V_{bb}(\tilde{p} - f)^2] < 0$. It holds if agents are risk-averse and if there exists some volatility in spot prices (Gilbert, 1985, Cowan, 2004). I assume both conditions hold for domestic electricity consumers on dynamic tariffs.

causes a high $\overline{V}_b$, corresponding low prices raise the household's real income. Therefore, the household reduces h^*.

The last term of equation (2.5) captures speculative motives for hedging. The household wants to speculate more the higher its income b and the smaller its coefficient of relative risk aversion θ. Below, I assume that household customers do not hedge for speculative reasons, i.e., households believe that $f = \overline{p}$. Thus, the last term in equation (2.5) equals 0.

Equation (2.5) describes the optimal hedge quantity h^* for general concave utility functions. To simulate hedge shares for real-world consumers, I assume that households' consumption decisions can be described by a homothetic constant-elasticity-of-substitution (CES) indirect utility function (Kihlstrom, 2009, Lau, 2016).

$$V(p, b, \varepsilon, f, h) = \frac{1}{1-\theta}[b + (p-f)h^*]^{1-\theta}(\varepsilon p^{1-\alpha} + 1)^{\frac{1-\theta}{\alpha-1}} \tag{2.6}$$

where $\alpha \neq 1$ is the elasticity of substitution. $\theta \neq 1$ is the coefficient of relative risk aversion at baseline. CES utility functions are commonly used to model electricity consumption, given their analytical tractability (Herriges et al., 1993, Schwarz et al., 2002, Goldman et al., 2004, Faruqui & Sergici, 2011).[6] Applying Roy's identity to equation (2.6) yields the Marshallian demand for electricity.

$$x^*(p, b, \varepsilon, f, h^*) = \frac{[b + (p-f)h^*]\varepsilon p^{-\alpha}}{(\varepsilon p^{1-\alpha} + 1)} \tag{2.7}$$

The Marshallian demand illustrates how the hedge affects electricity consumption. The hedge adds $(p-f)h^*$ to income b to partially compensate for a change in p. It increases households' disposable budget when prices are high ($p > f$) and decreases budget when they are low ($p < f$) (Braithwait & Eakin, 2002). Since households' income elasticity of demand equals 1, the hedge induces households to increase consumption when prices are above their average $\overline{p} = f$.

[6]This CES specification has a major drawback: For constant prices and quantity shocks, homothecity implies that the share of expenditure households spend on electricity remains constant when income increases. Empirically, electricity's expenditure share typically decreases with income. In the UK, electricity's expenditure share ranges from 4% for the lowest income decile to 1.3% for the highest decile in 2020 (UK ONS, 2021).

In Appendix B.1, I substitute the derivatives of the CES indirect utility function (2.6) into equation (2.5) to obtain the optimal hedge share:

$$\frac{h^*}{\widehat{x}^*} = \left(1 - \frac{1}{\theta}\right)\left(1 + \frac{1}{1-\alpha}\rho\frac{cv(\widetilde{\varepsilon})}{cv(\widetilde{p})}\right) \tag{2.8}$$

The optimal hedge share is the ratio of the absolute optimal hedge quantity h^* and the baseline consumption $\widehat{x}^* = x^*(\overline{p}, b, \overline{\varepsilon})$. $\widehat{x}^*$ defines what I previously called "typical consumption". It is not equal to the household's average consumption but describes how much a household consumes if $\widetilde{p}$ and $\widetilde{\varepsilon}$ are at their average level. Equation (2.7) highlights that $\widehat{x}^*$ does not depend on the hedge quantity h^* since $p = \overline{p} = f$ at baseline. Hence, at baseline, the optimally hedged household consumes as if it was unhedged.

On the right-hand side of equation (2.8), ρ depicts the coefficient of correlation between $\widetilde{p}$ and $\widetilde{\varepsilon}$. $cv(\widetilde{p})$ and $cv(\widetilde{\varepsilon})$ are the coefficients of variation of $\widetilde{p}$ and $\widetilde{\varepsilon}$, respectively. The optimal hedge share increases in the coefficient of relative risk aversion θ. If prices and quantity shocks are uncorrelated, the optimal hedge share approaches 1 when households become infinitely risk-averse. The lower households' risk aversion relative to their income elasticity of demand, the more the optimal hedge share decreases.[7]

A positive correlation ρ between quantity shocks and prices leads to higher optimal hedge shares. For example, if a household uses electric heating, it likely has a high desire to consume electricity when temperatures are low. Electricity prices are typically also high when temperatures are low due to high aggregate demand. Thus, households with electric heating face a positive correlation between electricity consumption and prices. This positive correlation makes households vulnerable to high electricity prices since they consume a lot when prices are high. Therefore, households with electric heating likely have high optimal hedge shares.

An increase in the substitution elasticity α further raises optimal hedge shares when the correlation between price and quantity shocks is positive (at least for reasonably small substitution elasticities, i.e., $\alpha < 1$). For a constant h^*, increasing α makes households' price elasticity of demand γ more negative if the substitution effect of a price change exceeds the income effect.

$$\gamma = \frac{\partial x^*}{\partial p}\frac{p}{x^*} = \frac{h^* p}{[b + (p-f)h^*]} \underbrace{-(1-s)\alpha}_{\substack{\text{Substitution}\\ \text{effect}}} \underbrace{-s}_{\substack{\text{Income}\\ \text{effect}}} \tag{2.9}$$

[7]In general, the first term of equation (2.8) is given as $1 - \frac{\eta}{\theta}$ for every concave utility function. η is the income elasticity of demand at baseline. $\eta = 1$ for the CES utility function in equation (2.6).

$s = \frac{px^*(h^*=0)}{b} < 1$ is electricity's expenditure share for unhedged households. For electricity consumers, the substitution effect typically dominates the income effect since electricity's share in household income is small. Therefore, households' demand becomes more price-elastic when the substitution elasticity α rises. An increase in price elasticity induces households to reduce consumption when prices are high. When $\rho > 0$, households particularly dislike a reduction in consumption since high prices are associated with a high desire to consume electricity. Thus, the household raises the hedge share h^* to reduce its price elasticity. As equation (2.9) illustrates, hedging reduces price elasticity since it mitigates exposure to spot prices. Thereby, the hedge ensures that consumption does not decrease too much when both prices and the desire to consume are high. On the other hand, increasing α reduces hedge shares when $\rho < 0$. A higher α makes the correlation between prices and consumption even more negative since more price-elastic households consume more when prices are low and the desire to consume electricity is high. In this case, the household reduces its hedge share since it wants to be exposed to low spot prices when its consumption is high.

To sum up, with quantity shocks, an increase in price elasticity leads to lower hedge shares only if the correlation between prices and quantity shocks is negative.

2.3.2 Price premia

This section derives the price premia households are willing to pay to receive a dynamic electricity tariff with an optimal forward hedge. I follow the literature on the welfare effect of price stabilization and use a fixed tariff as a benchmark (Hanoch, 1977, Turnovsky et al., 1980, Newbery & Stiglitz, 1981, Gilbert, 1985). Thereby, I can derive price premia for two tariffs: 1) An unhedged dynamic tariff and 2) A dynamic tariff with an optimal forward hedge (optimally hedged tariff). This distinction allows disentangling the two features of optimally hedged tariffs that impact consumer welfare: dynamic pricing and forward hedging. Following Gilbert (1985) and Cowan (2004), I assume that the constant price of the fixed tariff equals the mean dynamic price $\overline{p} = E[\tilde{p}]$.

Price premia for unhedged dynamic tariffs

In the first step, I derive the price premium g_p that households are willing to pay on top of the fixed price to avoid an unhedged dynamic tariff. When unhedged ($h = 0$), household's utility function simplifies to

$$V(p, b, \varepsilon) = \frac{1}{1-\theta} b^{1-\theta} (\varepsilon p^{1-\alpha} + 1)^{\frac{1-\theta}{\alpha-1}} \tag{2.10}$$

Price premium g_p equals the percentage increase in fixed price $\overline{p}$ that makes the household indifferent in expectation between the fixed tariff and the unhedged dynamic tariff with stochastic price $\tilde{p}$ (Gilbert, 1985).

$$E[V((1+g_p)\overline{p}, b, \tilde{\varepsilon})] = E[V(\tilde{p}, b, \tilde{\varepsilon})] \tag{2.11}$$

In Appendix B.2, I apply Taylor approximations to both sides of equation (2.11) and solve for g_p.

$$g_p = \frac{1}{2}\underbrace{(\hat{\gamma} + \beta^u)}_{=\frac{\overline{V}_{pp}}{\overline{V}_p}\overline{p}} cv(\tilde{p})^2 + \rho cv(\tilde{p})cv(\tilde{\varepsilon})[\frac{\beta^u}{1-\alpha} + 1 - \hat{s}] \tag{2.12}$$

If g_p is positive, households want to pay a premium for the fixed tariff. A negative g_p indicates that households would choose the dynamic tariff even if the fixed price was smaller than the average dynamic price. The first term on the RHS of equation (2.12) captures the effect of pure price volatility on consumer welfare. This term increases in $\frac{\overline{V}_{pp}}{\overline{V}_p}\overline{p} = \hat{\gamma} + \beta^u$, which is the absolute value of Turnovsky et al.'s (1980) coefficient of relative price risk aversion.[8] $\hat{\gamma} = -\alpha - (1-\alpha)s$ equals electricity's price elasticity of demand at baseline. $\beta^u = \frac{\partial \overline{V}_b}{\partial \overline{p}} \frac{\overline{p}}{\overline{V}_b} = (\theta - 1)s$ is the price elasticity of the marginal utility of income at baseline for unhedged households. Overall, as long as $\rho = 0$, dynamic prices increase welfare for unhedged households if their price elasticity of demand $\hat{\gamma}$ exceeds the price elasticity of their marginal utility of income β^u in absolute values.

Term 2 of the RHS of equation (2.12) reveals that households do not only care about price volatility but also about the correlation ρ between prices and quantity shocks. Price premium g_p increases if $\rho > 0$. The positive effect of a positive ρ is amplified by an increase of the price-elasticities of demand $\hat{\gamma}$ and of the marginal utility of income β^u as long as $\alpha < 1$ and $\theta > 1$. Hence, households with price-elastic marginal utility of income and high substitution elasticity suffer from dynamic prices if prices and quantity shocks are positively correlated. If ρ is negative, households benefit from dynamic pricing.

Price premia for dynamic tariffs with an optimal forward hedge

In the next step, I derive the price premium g_f that denotes the percentage increase in fixed price $\overline{p}$ that makes the household indifferent in expectation between the fixed tariff and the optimally hedged tariff (Gilbert, 1985).

$$E[V((1+g_f)\overline{p}, b, \varepsilon)] = E[V(\tilde{p}, b + (p - f)h^*, \tilde{\varepsilon})] \tag{2.13}$$

[8]The income elasticity of demand at baseline equals 1 with CES preferences.

Under the optimal hedge tariff, households buy the optimal hedge quantity h^* forward as given in equation (2.8). Appendix B.2 demonstrates that the price premium g_f equals

$$g_f = g_p \underbrace{- \frac{1}{2}\beta^u cv(\tilde{p})^2 - \frac{1}{2}\frac{\beta^u}{1-\alpha}\rho cv(\tilde{p})cv(\tilde{\varepsilon})\left(1 - \frac{h^*}{\hat{x}^*}\right)}_{g_h} \tag{2.14}$$

g_f consists of two terms: g_p captures the welfare effect of dynamic pricing as given in equation (2.12). $g_h = g_f - g_p$ describes the welfare effect of optimal forward hedging. If $g_h < 0$, the optimal hedge increases welfare compared to an unhedged dynamic tariff. g_h will never be positive since households can always choose $h^* = 0$ and stay unhedged.

The optimal forward hedge increases welfare by making the marginal utility of income V_b less elastic to price changes. When optimally hedged, the price elasticity of V_b at baseline is (see Appendix B.2)

$$\beta^{h^*} = \frac{\partial \overline{V}_b}{\partial \tilde{p}} \frac{\overline{p}}{\overline{V}_b} = \beta^u - \theta s \frac{h^*}{\hat{x}^*} \tag{2.15}$$

$\beta^u = (\theta - 1)s$ is the price elasticity of the marginal utility of income for unhedged households. The higher the household's optimal hedge share $\frac{h^*}{\hat{x}^*}$, the more declines the price elasticity of the marginal utility of income β^{h^*}. Inserting the optimal hedge share in equation (2.8) into equation (2.15) yields

$$\beta^{h^*} = -\frac{\beta^u}{1-\alpha}\rho \frac{cv(\tilde{\varepsilon})}{cv(\tilde{p})} \tag{2.16}$$

If $\rho = 0$, the optimal hedge makes the marginal utility of income fully inelastic to price changes ($\beta^{h^*} = 0$). Equation (2.14) reveals that for $\rho = 0$, stabilization of V_b raises consumer welfare whenever $\beta^u > 0$. β^u is positive if households are sufficiently risk-averse with $\theta > 1$. The welfare benefit from hedging rises in the price volatility $cv(\tilde{p})^2$ (Cowan, 2004).

If $\beta^u > 0$, $\rho > 0$ implies $\beta^{h^*} < 0$. In this case, optimal hedging causes the marginal utility of income V_b to fall when prices increase. If $\rho > 0$, the optimal hedge share is very high. Hence, the hedge shifts much income to states of the world when prices are high. This large income compensation reduces V_b since $V_{bb} < 0$. For $\rho > 0$, the income compensation strongly raises welfare because both prices and the desire to consume are high. On the other hand, equation (2.14) reveals that welfare benefit declines in the optimal hedge share $\frac{h^*}{\hat{x}^*}$. The higher $\frac{h^*}{\hat{x}^*}$, the more decreases the benefit from hedging since hedging reduces households' exposure to spot prices and, thereby, its price elasticity of demand. For $\frac{h^*}{\hat{x}^*} > 1$, the negative effect on demand elasticity dominates such that the welfare effect of the second term in g_h becomes negative. If $\rho < 0$, the second term of g_h will always lower welfare since households will never choose a hedge share larger than 1 when $\rho < 0$.

2.4 Simulation

Below, I explain how the model in Section 2.3 allows simulating optimal hedge shares and hedged electricity bills for real-world consumers. First, I define the hedge product that consumers can buy. In the example in the introduction, I assumed that households buy forward a particular share of their daily baseline consumption. However, such a daily hedge ignores predictable intraday patterns of both consumption and prices. Baseline electricity consumption and prices are typically low at night and high during morning and evening peaks. A hedge contract should account for these daily load and price patterns to better protect against bill volatility (Borenstein, 2007a).

To do so, I define a *hedge time segment* k as a set of time intervals t.[9] Households choose a quantity to buy forward for every segment k. In the main specification, I define hedge time segments for every half hour per day for weekdays and weekends. For instance, 8-8:30 am on weekdays is a time segment. Average prices and baseline consumption are similar on weekdays 8-8:30 am, whether on Monday, Wednesday, or Friday (Cramton et al., 2025a). Therefore, households buy the same hedge quantity for all intervals t within 8-8:30 am on weekdays. Prices and baseline consumption in this time segment likely differ from average prices and baseline consumption at midnight or 8-8:30 am on weekends. Hence, households purchase different hedge quantities for those time segments (See Appendix B.3 for a detailed discussion of time segments).

The approach results in 48 half-hourly hedge time segments and 48 corresponding optimal hedge shares for weekdays and weekends, respectively. The advantage of such a fine-grained hedge contract is that it should protect well from bill volatility. On the other hand, a hedge contract with 96 different hedge shares is too complex to be implemented in practice. Therefore, the fine-grained hedge contract should be considered a theoretical benchmark that tests how effectively an optimally hedged tariff can reduce bill volatility. In Appendix B.3, I compare the results of this theoretical benchmark to less fine-grained hedge contracts that contain fewer and longer hedge time segments. Similar to Borenstein

[9] A time interval t is a specific half-hourly period. For instance, "8-8:30 am on Monday, March 1st, 2021" is a time interval. This interval falls within the hedge time segment "8-8:30 am on weekdays".

(2007a), I find that choosing a more granular hedge time segment has only a small effect on optimal hedge shares. However, I find that more fine-grained hedge segments lead to a substantially larger reduction in bill volatility.

The households in my sample are currently unhedged. To simulate their optimal hedge shares, I suppose that households' observable unhedged half-hourly electricity consumption x_t^* in interval t can be described by the Marshallian demand in equation (2.7). When households are unhedged, hedge quantity h equals 0. I invert equation (2.7) to obtain an expression for the unobservable quantity shocks ε_t in interval t as a residual of the unhedged Marshallian demand (Redding & Weinstein, 2019).[10]

$$\varepsilon_t = \frac{x_t^* p_t^{\alpha}}{b_k - x_t p_t} \tag{2.17}$$

I assume that households have a fixed budget b_k for hedge time segment k. I do not observe households' monthly income, let alone their budget for a half-hourly time segment. Surveys reveal that UK households spend on average 2% of their income on electricity (UK ONS, 2021). Therefore, I assume that, in every segment k, households' budget share of electricity $s_t = \frac{x_t^* p_t}{b_k}$ is on average 2%, i.e., $\overline{s}_k = \frac{1}{T_k}\sum_{t\in k} s_t = 2\%$. T_k is the number of time intervals t that fall into segment k. Households' absolute budget for segment k can then be derived as $b_k = \frac{1}{T_k}\sum_{t\in k}\frac{x_t^* p_t}{\overline{s}_k}$.[11]

The assumption that households spend the same average share of their income on electricity in every time segment implies that households have a higher absolute budget for time segments in which they typically consume a lot. For instance, most households have a higher absolute segment budget on weekdays, 8-8:30 am, than at midnight because they typically consume more in the morning. Therefore, the segment budget follows each household's daily consumption pattern.

An important caveat for my analysis is that the assumption regarding the segment budget b_k affects the optimal hedge shares even for homothetic preferences. Changing b_k also changes the derived quantity shocks ε_t. Thereby, the choice of b_k influences the

[10] The Marshallian demand is invertible since x_t and y_t are "connected substitutes", according to Berry et al. (2013). The first condition for connected substitutes is that goods are weak gross substitutes. This condition is satisfied since aggregate consumption of other goods y_t weakly decreases in ε_t for all t and p_t. The second condition for "connected substitutes" requires "connected strict substitution" (Berry et al., 2013) between all goods. This condition requires that any chain of substitution between goods leads to the outside good (i.e., y_t). Since there are only two goods in the given application, electricity consumption can only be substituted with the outside good (Berry et al., 2013, Berry & Haile, 2016).

[11] For a tiny number of time intervals (only 0.006% of all intervals), consumption is extremely high such that the observed expenditure exceeds the segment budget b_k. For these rare intervals, I set the household budget equal to the observed expenditure on electricity to ensure that households' electricity expenditure does not exceed their budget.

correlation between quantity shocks and prices and, ultimately, the optimal hedge shares. As a robustness check, I, therefore, run simulations for various average segment budget shares $\overline{s}_k$ in Appendix B.5. I find that a change in $\overline{s}_k$ only slightly affects optimal hedge shares.

To calculate ε_t in equation (2.17), I also choose the substitution elasticity α as a simulation parameter. Previous studies find that the short-run elasticity of substitution for electricity ranges from 0.07 to 0.21 for domestic electricity consumers (Ericson, 2006, Bartusch, 2011).[12] For my main specification, I conservatively set $\alpha = 0.1$, assuming that households have a relatively low ability to substitute electricity consumption with other goods.

Having defined b_k and α, I calculate the residual quantity shock ε_t in equation (2.17) for every t in the data set using the observed unhedged half-hourly electricity consumption x_t^* and prices p_t. Then, I calculate the following statistics for every segment k: 1) the average quantity shock $\overline{\varepsilon}_k$ and average electricity price $\overline{p}_k$, 2) the coefficients of variation $cv(\tilde{\varepsilon}_k)$ and $cv(\tilde{p}_k)$, 3) the coefficient of correlation ρ_k between price and quantity shocks. These statistics allow calculating the baseline consumption $\hat{x}_k^*$.

Now, I am ready to calculate the optimal hedge share $\frac{h_k^*}{\hat{x}_k^*}$ in equation (2.8) for segment k. For the main specification, I assume that domestic electricity consumers face high levels of relative risk aversion with $\theta = 5$. Most empirical studies estimate θ in the range from 0 to 6 (Cowan, 2004, Lengwiler, 2004). To avoid speculation, I set the forward hedge price equal to the average dynamic price in segment k, i.e., $f_k = E[p_{t\in k}]$.[13] Moreover, I follow Borenstein (2007a) and require the optimal hedge shares to lie between 0% and 200% of baseline consumption to avoid unreasonably extreme hedge shares.

In the next step, I derive monthly electricity bills for optimally hedged households. I consider that the optimal hedge induces income- and price-elastic households to change their consumption compared to the unhedged tariff. Thus, I calculate the hedged households' electricity consumption $x_{t\in k}^*$ with the optimal hedge h_k^* using the Marshallian demand in equation (2.7). The electricity bill B_m for month m results by summing over the expenditure for electricity and the forward hedge for all intervals $t \in m$:

$$B_m = \left(\sum_{t\in m} \sum_{t\in m\cup k} [x_t^* p_t + (f_k - p_t)h_k^*] \right) * \frac{30}{Days_m} \tag{2.18}$$

[12] These estimates for the elasticity of substitution either refer to the substitution between electricity and other energy carriers or the substitution between electricity consumption at different times of day.

[13] For existing two-part tariffs with a forward hedge, some utilities set the forward price even below the average dynamic price to compensate customers for foregone cross-subsidies they would have received on a fixed tariff (Borenstein, 2007a).

where $Days_m$ denotes the number of days the household spends on the tariff in month m. The last term normalizes the electricity bills to a 30-day month to make sure that bill volatility is not caused by variation in the number of days the household spends on the tariff (see Section 2.5 and Borenstein, 2007a).

2.5 Data

The data set contains anonymized smart meter readings of half-hourly electricity price and consumption for 9,718 households in the UK provided by electricity supplier Octopus Energy. 4,066 households are on dynamic tariffs ("dynamic households"), while 5,652 are on fixed ones ("fixed households"). The data contains smart meter readings for each household for up to one year between August 2020 and August 2021 (Cramton et al., 2025a).

Every day between 4-8 pm, Octopus Energy informs dynamic households about the 48 half-hourly dynamic prices for the next day via a smart phone app. Dynamic prices include distribution charges and a peak-time premium. Appendix B.4 explains how dynamic prices are calculated.

Households pay dynamic prices only up to a price cap of 0.35£/kWh. The price cap equals roughly twice the average dynamic price. During the sample period, the cap only binds in 3% of the half-hourly periods. Therefore, the price cap has a negligible effect on my analysis, as I discuss in Section 2.6.5.

Fixed households always receive a constant price for electricity. 8% of fixed tariff customers experience a minor adjustment of their fixed rate once during the sample period. Apart from these one-time adjustments, these households do not face any price volatility. I do not employ households on fixed tariffs for the optimal hedge simulations. Nevertheless, I report the descriptive statistics for these households as a benchmark to analyze how dynamic prices affect consumption and electricity bills.

The households in this sample have been randomly chosen from Octopus Energy's customer base. However, households have self-selected into a tariff type. Hence, households on dynamic tariffs might systematically differ from households on fixed ones. Households can also switch to a different tariff or supplier during the sample period. Figure 2.1 reveals a surprisingly high attrition of households on both tariffs. I drop all households for whom less than eight monthly electricity bills are observable to ensure a meaningful analysis of bill volatility. Moreover, some households spend less than an entire month on a specific tariff because they either leave or join the tariff in the middle of the month. I exclude all observations for months if the household spends less than 15 days on the tariff.

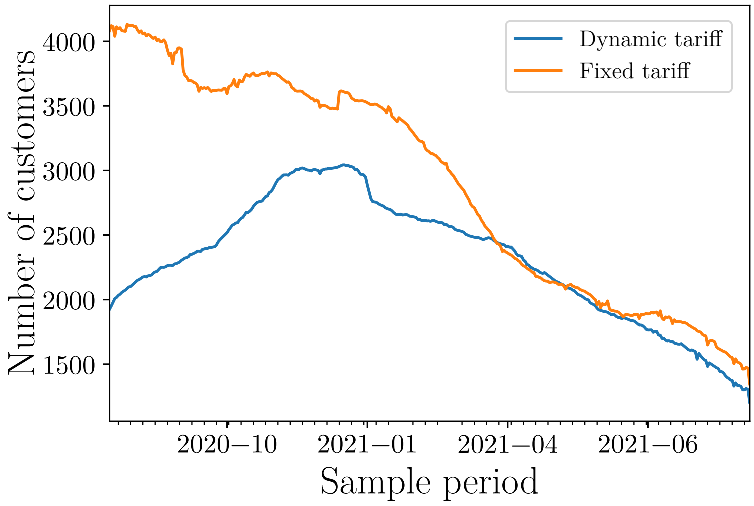

Figure 2.1: Number of customers by tariff type

This ensures that at least 720 half-hourly price-consumption observations are available per household each month. The final dataset contains a relatively large sample of 2,159 dynamic households and 2,540 fixed households.

Tariff type	mean	SD[14]	min	max
	Electricity prices (pence/kWh)			
Dynamic	15.4	7.3	0.1	35
Fixed	15.2	0.1	15.2	15.3
	Monthly consumption (kWh)			
Dynamic	516	406	0	4560
Fixed	334	239	0	5040
	Monthly electricity bill (GBP)			
Dynamic	77	61	0	766
Fixed	50	36	0	798

Table 2.1: Summary statistics

Table 2.1 presents summary statistics for households on dynamic and fixed tariffs. The mean unweighted electricity price is very similar on both tariffs. The standard deviation of dynamic prices is quite large. Differences in fixed prices are driven by different fixed price levels across customers.

[14] Average household-specific standard deviation. Since households can choose between multiple different fixed tariff products with different fixed rate levels, the standard deviation across all fixed prices is slightly larger. In the context of this paper, the household-specific standard deviation is a more relevant indicator of the price volatility that households perceive.

Table 2.1 also reveals that dynamic households have a 54% higher average monthly consumption than fixed households. The higher monthly consumption suggests that households who self-select into dynamic tariffs systematically differ from fixed households. Dynamic customers are first movers who might be exceptionally interested in their electricity usage. Therefore, one should be careful to extrapolate the results of this paper to the general population of domestic energy consumers. This paper is only informative about whether forward hedging is effective for first-movers. Nevertheless, energy suppliers should be particularly interested in how forward hedge tariffs work for first movers as they will likely only sell forward hedges to first-movers in the foreseeable future.

Moreover, Table 2.1 shows statistics on households' monthly electricity bills. I calculate these bills by multiplying consumption and prices every half an hour and adding these half-hourly expenditures every month. Following Borenstein (2007a), I normalize all electricity bills to a 30-day month in order to avoid measuring bill volatility that is purely driven by variation in days per month or the number of days a household spends on the tariff in a given month (see Section 2.4). The higher average consumption of dynamic customers likely drives their higher average electricity bills.

		Percentiles				
tariff type	mean	0.01	0.1	0.5	0.9	0.99
	Coeff. of variation of monthly bills per customer					
Dynamic	0.32	0.10	0.16	0.26	0.57	1.06
Fixed	0.19	0.05	0.08	0.16	0.35	0.65
	Coeff. of correlation between prices and consumption					
Dynamic	-0.007	-0.363	-0.225	-0.005	0.211	0.336

Table 2.2: Variation and correlation statistics

Since this study is concerned with hedging bill volatility for individual customers, Table 2.2 reports the distributions of coefficients of variation for individual customers' monthly bills (Borenstein, 2007a). The coefficient of variation of dynamic customers' electricity bills is, on average, 67% higher than for fixed customers, suggesting that dynamic prices increase bill volatility. While this might be unsurprising, it is striking that fixed customers also face substantial bill volatility. Fluctuations in electricity bills are not only caused by volatile prices but also by volatility in consumption (Borenstein, 2007a).

Moreover, fixed customers in the lowest percentile of the bill volatility distribution have much lower volatility than dynamic customers in the same percentile. When moving up the distribution, the difference in bill volatility between fixed and dynamic tariffs declines. Borenstein (2007a) observes a similar tendency in his study. This suggests that dynamic prices might mainly increase bill volatility for customers who would have low volatility under fixed tariffs (Borenstein, 2007a).

Table 2.2 also reports the distribution of the correlation coefficient between dynamic electricity prices and consumption for dynamic customers. While the correlation between prices and consumption is, on average, close to zero, it varies enormously between customers. Some customers face either a significant negative or large positive correlation.

LCT ownership	Number of households	Mean monthly consumption (kWh)	Mean monthly electricity bill (GBP)
All customers	2159	516	77
Electric heating only	225	352	58
Electring heating + smart thermostat	221	394	65
EV only	411	635	92
Solar only	101	365	51
Battery only	28	769	114
EV + solar	125	497	67
EV + battery	32	1088	156
Solar + battery	84	437	59
EV + solar + battery	82	592	74

Table 2.3: Summary statistics by low-carbon technology (LCT) ownership

Analyzing the same dataset, Cramton et al. (2025a) estimate an average price elasticity of -0.26 for the households in this sample. For a sub-sample, they employ information on low-carbon technology (LCT) ownership to show that technologies influence households' price elasticity. Especially ownership of electric vehicles increases price elasticity.

In this paper, I analyze how LCT ownership impacts optimal hedge shares. Table 2.3 shows the number of customers for which information on technology ownership is available. Octopus Energy conducted a survey to gather ownership information for electric vehicles, solar PV, battery storage, and smart thermostats. For electric heating, ownership is inferred from a lack of a gas contract with Octopus Energy since households in the UK typically purchase electricity and gas from the same supplier (Cramton et al., 2025a).[15]

Electric vehicles (EV) or battery storage owners have above-average monthly electricity consumption and monthly bills. Solar PV owners have lower consumption and bills

[15]This might not be an exact indicator for ownership of electric heating. Some households who are not on a gas tariff with Octopus Energy might have oil heating rather than electric heating or buy gas from another supplier (Cramton et al., 2025a).

since they likely cover a portion of their electricity consumption via self-generated solar electricity. Surprisingly, customers who likely have electric heating consume less electricity than the average household. Given small sample sizes, the analysis for some LCT groups (e.g., battery owners) should be interpreted cautiously.

2.6 Results

2.6.1 Optimal hedge shares

Figure 2.2 presents the distribution of optimal hedge shares that I simulate across all households and time segments. On average, the optimal hedge share is 59% of baseline consumption. There is substantial variation in optimal hedge shares. I set the optimal hedge share to zero for a surprisingly large share of 14% of the time segments. Otherwise, optimal hedge shares would be negative. The desire to hedge a negative quantity is caused by a large negative correlation between prices and quantity shocks in these time segments. Moreover, I set 3% of the hedge shares to 200% to avoid unreasonably large hedge shares.

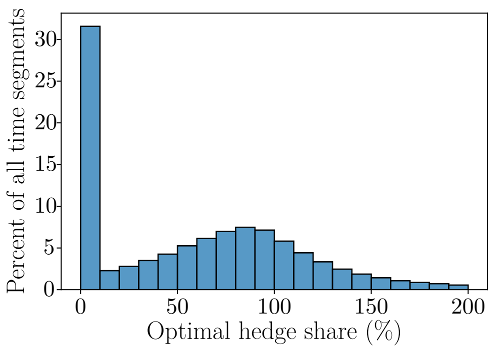

Figure 2.2: Distribution of optimal hedge shares across all time segments

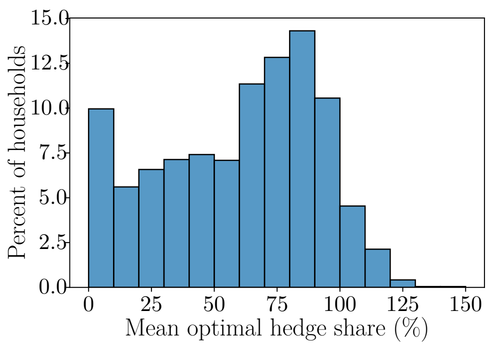

Figure 2.3: Distribution of mean optimal hedge shares by customer

Heterogeneity between households partly drives the considerable variation in optimal hedge shares. Figure 2.3 shows the distribution of the mean optimal hedge share per household. The mean hedge share for a particular household is defined as the average of the household's hedge shares across all time segments. The mean hedge shares differ widely between households. Overall, mean optimal hedge shares are far smaller than the ones Borenstein (2007a) simulates for a sample of price-inelastic industrial customers. Most households find it optimal to hedge on average less than 100% of their baseline consumption. Only 7% of households in this sample overhedge, i.e., hedge more than

100% of their baseline consumption on average. In contrast, Borenstein (2007a) finds that 77% of his inelastic customers should overhedge.

Table 2.4 suggests that ownership of low-carbon technologies (LCTs) partly explains the heterogeneity in mean optimal hedge shares. Households with electric heating or an electric vehicle (EV) have a larger mean optimal hedge share than the average household in the sample. These appliances seem to increase the positive correlation between a household's consumption and prices. Surprisingly, solar PV owners have an average optimal hedge share of only 21%. Solar PV likely results in a strong negative correlation between prices and electricity consumption from the grid. As expected, the small sample of battery storage owners also has a slightly below-average optimal hedge share of 51%.

Given their small sample sizes, the above results should be interpreted with caution for solar PV and battery owners. I also stress that this paper does not establish a causal relationship between technology ownership, optimal hedge shares, and bill volatility. The above results merely provide first suggestive evidence that technology ownership influences optimal hedge shares.

LCT ownership	Households (#)	Optimal hedge share: Average (%)	Optimal hedge share: SD[16] (%)	Average coefficient of bill variation: Unhedged	Average coefficient of bill variation: Optimal hedge	Average coefficient of bill variation: Optimal vs. Unhedged (%)
All customers	2159	59	37	0.34	0.28	-18
Electric heating only	225	76	38	0.25	0.19	-23
Electric heating + smart thermostat	221	76	35	0.24	0.17	-28
EV only	411	68	43	0.26	0.18	-30
Solar only	101	21	29	0.50	0.49	-1
Battery only	28	51	41	0.42	0.42	-2
EV + solar	125	28	38	0.37	0.36	-6
EV + battery	32	68	40	0.33	0.26	-26
Solar + battery	84	17	25	0.65	0.66	1
EV + solar + battery	82	20	31	0.48	0.48	0

Table 2.4: Optimal hedge shares and bill volatility by low-carbon technology (LCT) ownership

Table 2.4 also indicates that heterogeneity between households is not the only reason for the enormous variation in optimal hedge shares. Even for a specific household, optimal hedge shares differ strongly by time of day. On average, the hedge share for a specific household has a standard deviation of 37% across time segments.[16] These differences between time segments are likely caused by significant variations in household-specific daily consumption patterns. In some time segments, a household's consumption pattern

[16] Average of the standard deviation of optimal hedge shares for specific households across time segments.

might align with the aggregate load pattern and is, therefore, positively correlated with prices. The same household's consumption can correlate negatively with aggregate load and prices in other time segments. The fine-grained hedge tariff with small time segments accounts for these household-specific consumption patterns to provide adequate protection from bill volatility.

2.6.2 Bill volatility

In this section, I discuss how effectively the optimal hedge protects customers from volatility of monthly electricity bills. Figure 2.4a compares the coefficients of variations of bills for unhedged and optimally hedged tariffs across mean optimal hedge shares. The optimally hedged tariff proves effective in reducing bill volatility. On average, the optimally hedged tariff reduces the coefficient of variation of a household's monthly bill by 18% compared to an unhedged dynamic tariff. Hedging is even more effective for households who overhedge. It reduces their coefficient of bill variation by 45% on average. On the other hand, hedging hardly reduces bill volatility for households whose mean optimal hedge share is less than 50%. Naturally, hedging does not significantly affect households that only hedge a small share of their baseline consumption. These findings are similar to the ones in Borenstein (2007a).

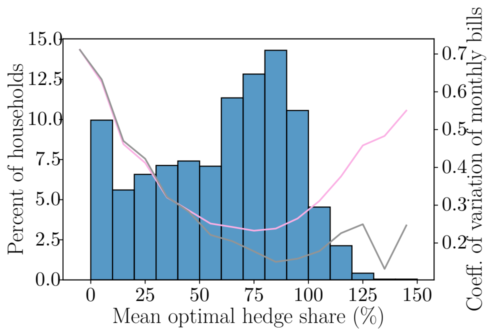

(a) Unhedged tariff vs. optimally hedged tariff

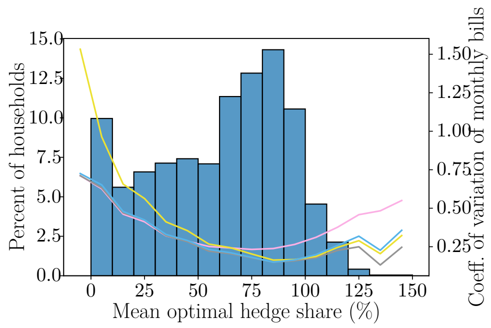

(b) Comparison of various tariff types

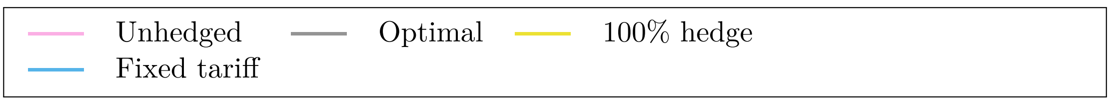

Figure 2.4: Distribution of mean optimal hedge shares by customer and coefficients of variation of monthly bills by tariff type

I also compare the effectiveness of the optimally hedged tariff to a more straightforward tariff in which households always buy 100% of their baseline consumption forward. Figure

2.4b reveals that when always forced to buy a 100% hedge, households with a mean optimal hedge share of less than 50% have a higher bill volatility than under the unhedged tariff. Hence, a simple 100% hedge tariff is unsuitable for households with low mean optimal hedge shares. On average, the 100% hedge tariff leads to a 15% higher coefficient of variation of monthly bills than the optimally hedged tariff.

In addition, I compare the bill volatility of the previous tariffs to the one under a fixed tariff. For customers with mean hedge shares above 100%, the optimally hedged tariff results in lower bill volatility than the fixed tariff, as Figure 2.4b reveals. Borenstein (2007a) reports similar results in his paper. He argues that the hedged tariff protects households more effectively against price and quantity risks. The fixed tariff protects only against price risks. Protection against quantity risk is essential for households who overhedge since they face a positive correlation between prices and consumption.

Table 2.4 reveals that the optimal hedge reduces bill volatility most effectively for households with electric vehicles and heating. For instance, optimal hedging reduces the bill volatility for owners with only an EV by 30%. In contrast, hedging only slightly reduces bill volatility for solar PV and battery owners with low hedge shares.

2.6.3 Price elasticity and risk aversion

Below, I analyze how risk aversion and demand elasticity affect hedge shares and bill volatility. Figure 2.5 confirms that increasing substitution elasticity α increases the variation in mean optimal hedge shares when exogenous quantity shocks remain constant.[17] Raising α leads to lower hedge shares when the correlation between prices and quantity shocks is negative and higher hedge shares when it is positive (see Section 2.3.1).

On average, higher substitution elasticity decreases mean optimal hedge shares. Inelastic households ($\alpha = 0.001$) have an average mean optimal hedge share of 60% with a standard deviation of 36%. Elastic households ($\alpha = 0.6$) show an average mean optimal hedge share of 48% with a standard deviation of 47%. Moreover, a large α increases the share of households that overhedge or would even choose a negative hedge quantity if allowed.

Hedging reduces bill volatility more effectively when households are price-inelastic. For price-inelastic households ($\alpha = 0.001$), the optimal hedge lowers the coefficient of bill volatility by 21% compared to the unhedged dynamic tariff. It does so by mitigating house-

[17]Since quantity shock $\tilde{\varepsilon}_t$ is a residual of the Marshallian demand function, it depends on α. To calculate $\tilde{\varepsilon}_t$ in Section 2.4, I use $\alpha = 0.1$ equal to the α used in the main specification above. Quantity shocks should be considered exogenous shocks to demand. Therefore, when studying the effect of a change in α on hedge shares, I hold quantity shocks constant to their level in the main specification.

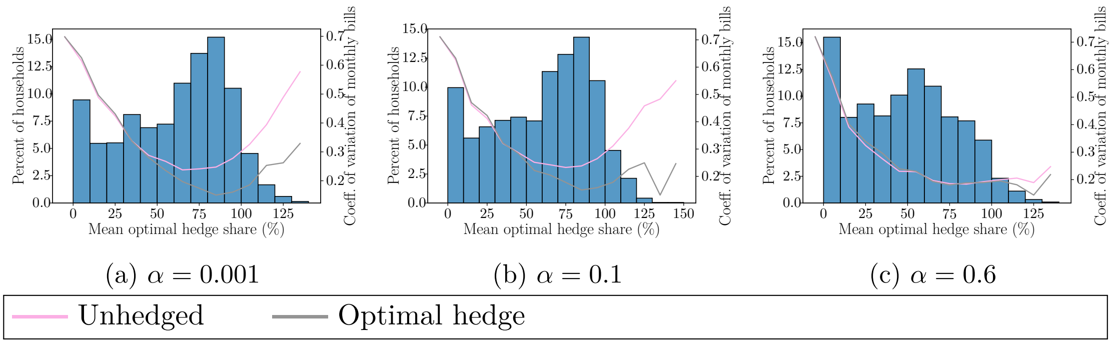

Figure 2.5: Distribution of mean optimal hedge shares by customer and coefficients of variation of monthly bills by tariff type for different substitution elasticity α

holds' exposure to high prices. At the same time, the hedge hardly reduces households' price response compared to the unhedged tariff when households are price-inelastic.

For very price-elastic households ($\alpha = 0.6$), the optimal hedge even increases bill volatility on average by 5% relative to the unhedged tariff. Price-elastic households respond less strongly to prices when optimally hedged relative to being unhedged. Therefore, households consume more when prices are high on optimally hedged tariffs than on unhedged ones. For some households, this leads to even more volatile bills on the optimally hedged tariff than on the unhedged tariff. Households who decrease their hedge when α increases experience on average a higher bill volatility on the optimally hedged tariff relative to the unhedged one. Households who increase the hedge in response to higher price elasticity face lower average bill volatility when optimally hedged.

Moreover, optimally hedged households face an increase in bill volatility by 20% when their price elasticity increases from $\alpha = 0.001$ to $\alpha = 0.6$. This reveals that when price-elasticity rises, the volatility of the optimally hedged bill increases not only relative to the unhedged bill but also in absolute terms. In contrast, the bill volatility of unhedged households declines by 13% when increasing α from 0.001 to 0.6.

Figure 2.6 shows that higher risk aversion results in higher hedge shares. Hardly risk-averse households ($\theta = 1.5$) hedge far less than 100%. Few households still choose substantial hedge shares as protection against a positive correlation between prices and quantity shocks. When risk aversion increases, its impact on hedge shares diminishes. The distribution of mean optimal hedge shares for $\theta = 20$ only moderately differs from the distribution for $\theta = 5$ in the main specification in Figure 2.4b. Similar to an increase in price elasticity, an increase in risk aversion makes hedging less effective in lowering bill volatility relative to the unhedged tariff. Risk-aversion leads to higher hedge shares. These higher hedge shares lower households' exposure to spot prices and, therefore, allow households to consume more when prices are high. For high risk aversion ($\theta = 20$), this leads, on average, even to slightly higher bill volatility on the optimally hedged tariff compared to the unhedged tariff.

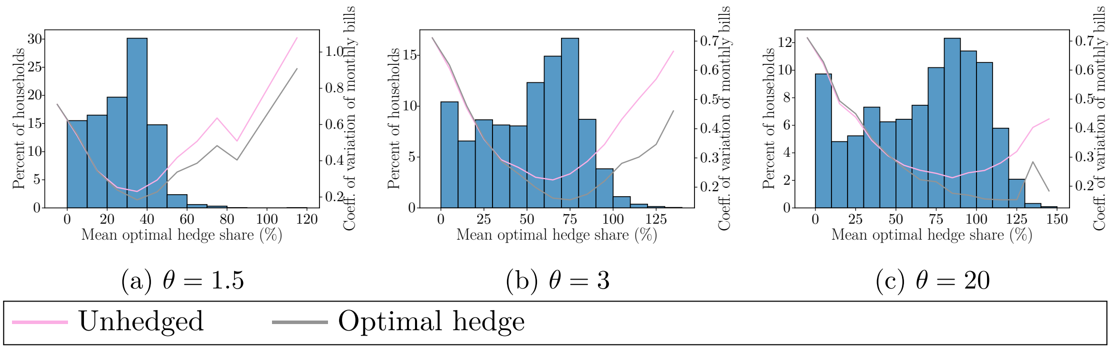

(a) $\theta = 1.5$ (b) $\theta = 3$ (c) $\theta = 20$

Figure 2.6: Distribution of mean optimal hedge shares by customer and coefficients of variation of monthly bills by tariff type for different levels of risk aversion θ

2.6.4 Price premia

Figure 2.7 shows the price premia households are willing to pay to receive the forward hedge by mean optimal hedge share. The negative price premium implies that the hedge increases households' welfare. However, the welfare effect of the forward hedge is small. Compared to an unhedged dynamic tariff, the forward hedge reduces the mean electricity price that makes households indifferent between the unhedged tariff and the optimally hedged tariff by merely 0.3% (for the main specification with $\alpha = 0.1$). The welfare benefit from hedging is most significant for households who hedge close to 100%. The hedge adds the most value for these households by stabilizing the marginal utility of income while still exposing households sufficiently to spot prices to keep demand price-elastic.

One reason for the small welfare effects of hedging might be electricity's small share in household expenditure. In the above simulations, I assumed that households spend on average only 2% of their income on electricity, which is roughly the case in the UK (UK ONS, 2021). Such a small expenditure share implies that a change in electricity prices will have merely a small effect on households' marginal utility of income. When the price elasticity of the marginal utility is already tiny, the hedge can only add little value by stabilizing it.

However, electricity's share in household expenditure will likely increase in the upcoming decades. Households will start using electricity for heating, cooling, and driving cars. To assess how an increasing expenditure share of electricity affects the welfare

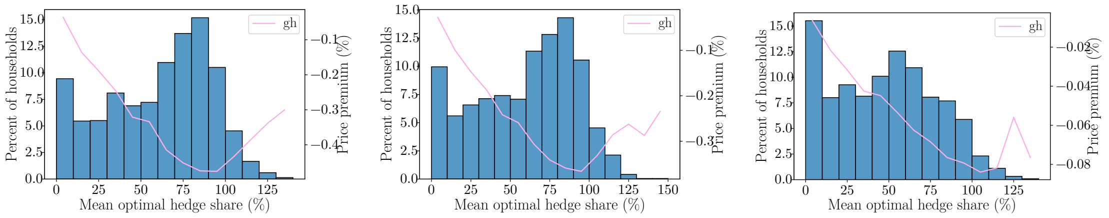

(a) $\alpha = 0.001$ (b) $\alpha = 0.1$ (c) $\alpha = 0.6$

Price premium for the forward hedge (g_h)

Figure 2.7: Distribution of mean optimal hedge shares by customer and forward hedge price premium g_h for different α

benefits of hedging, I run the above simulation assuming that households spend on average 10% of their income on electricity. Figure B.10 in Appendix B.6 reveals that a higher expenditure share of electricity only slightly increases the welfare benefits from hedging. On average, the welfare benefit translates to a 1.6% reduction in the mean electricity price compared to a 0.3% reduction when electricity's expenditure share remains at 2%.

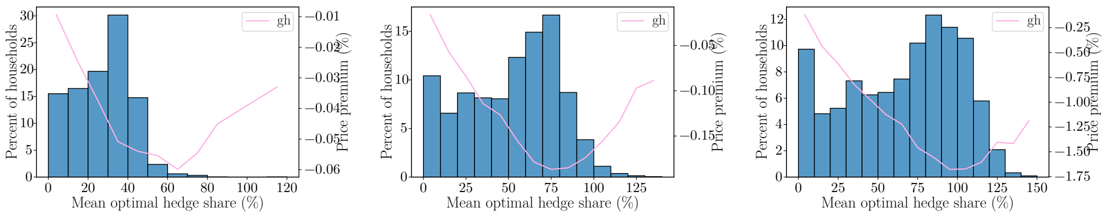

(a) $\theta = 1.5$ (b) $\theta = 3$ (c) $\theta = 20$

Price premium for the forward hedge (g_h)

Figure 2.8: Distribution of mean optimal hedge shares by customer and forward hedge price premium g_h by tariff type for different θ

Figure 2.7 also highlights that increasing substitution elasticity α further diminishes the welfare benefit from hedging. The higher α, the smaller the share of households that hedge close to 100%. Households who hedge roughly 100% benefit most from hedging. When demand is elastic ($\alpha = 0.6$), the hedge only achieves an average welfare benefit of 0.05%. Meanwhile, Figure 2.8 depicts that an increase in risk aversion has a more substantial effect on the welfare gain from hedging. For very small levels of risk aversion ($\theta = 1.5$), the welfare gain from hedging amounts to only 0.04% expressed as a reduction in mean electricity prices. If risk aversion is large ($\theta = 20$) hedging leads to a larger but still surprisingly low 1.2% welfare gain.

Figure 2.9 reveals that dynamic pricing has a much stronger effect on household welfare than forward hedging. The figure shows the price premia for unhedged dynamic tariffs g_p and the optimally hedged dynamic tariffs g_f. g_f hardly differs from g_p since the welfare effect of the forward hedge $g_h = g_f - g_p$ is negligible relative to g_p. Compared to a fixed tariff, the unhedged dynamic tariff leads to an average welfare gain equivalent to a 1.4% reduction in average electricity prices for the main specification ($\alpha = 0.1$). 58% of households achieve higher welfare on the unhedged dynamic tariff than on the fixed tariff. Households that benefit most from dynamic pricing also choose low hedge shares since they have a negative correlation between electricity prices and their desire to consume electricity. Households with a positive correlation are better off on a fixed tariff.

As expected, increasing demand elasticity makes unhedged dynamic tariffs more attractive. When $\alpha = 0.6$, 85% of households prefer dynamic pricing over a fixed tariff. The average welfare gain from dynamic pricing rises to 4%.

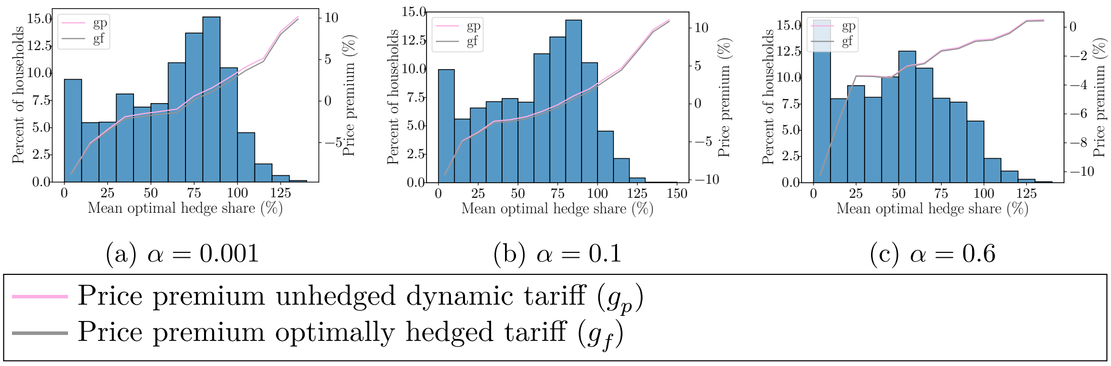

Figure 2.9: Distribution of mean optimal hedge shares by customer and price premia g_p and g_f for different α

2.6.5 Scarcity price event

There are two additional reasons why hedging results in low welfare gains: First, households in my sample are not fully exposed to day-ahead electricity prices but are protected via a price cap. Since the price cap already protects households from high prices, the hedge can only add small additional value. Second, households do not experience an extreme scarcity event like the Texas winter storm during the sample period. The benefit from hedging is larger when households experience worst-case events with extremely high prices.

Figure 2.10b depicts how the performance of the optimal forward hedge changes when removing the price cap and exposing households to fully dynamic prices. Removing the cap has a negligible effect on the distribution of mean optimal hedge shares and bill volatility. Likewise, the welfare benefits from hedging only marginally increase after removing the cap (see Figure 2.11b). The price cap has only tiny effects during my sample period since dynamic prices rarely exceed the cap in only 3% of time intervals. At the same time, price caps are suboptimal since they lower the incentive to respond to high scarcity prices (Appendix B.4 provides statistics on dynamic prices after removing the price cap).

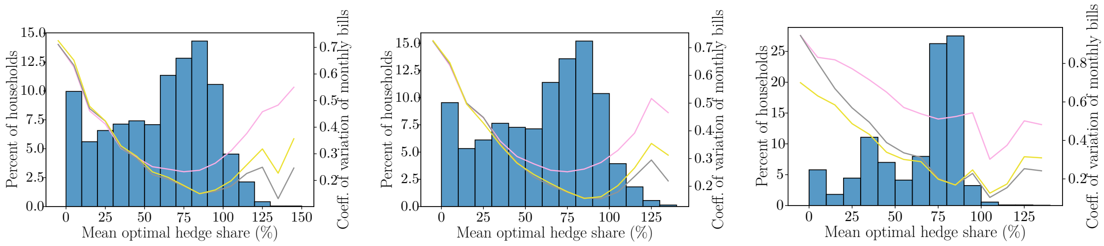

(a) Baseline results with price cap (b) Results without price cap (c) Results without price cap and scarcity event

Unhedged Optimal hedge Fixed tariff

Figure 2.10: Distribution of mean optimal hedge shares by customer and coefficients of variation of monthly bills by tariff type for different levels of price protection

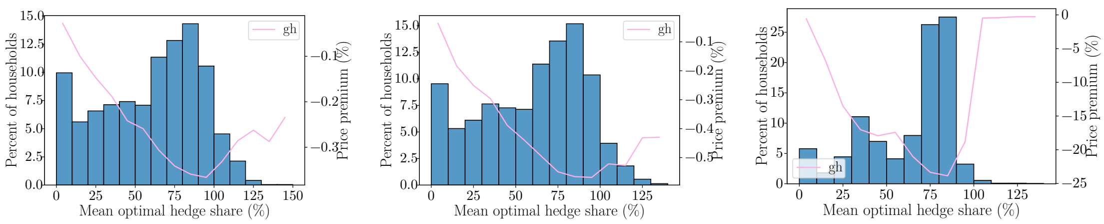

(a) Baseline results with price cap (b) Results without price cap (c) Results without price cap and scarcity event

Price premium for the forward hedge (g_h)

Figure 2.11: Distribution of average optimal hedge shares by customer and forward hedge price premium g_h by tariff type for different levels of price protection

In Figure 2.10c, I study how a scarcity event like the Texas winter storm with extremely high prices over multiple days affects hedging and bill volatility. To do so, I run the above simulations while removing the price cap and simulating a four-day scarcity event: I manually choose a very high dynamic price of 3.5£/kWh, which is ten times the current price cap. I set this very high price for all 192 consecutive time intervals from 12 to 15 January 2021.

Figure 2.10c highlights that mean optimal hedge shares moderately increase to 63% on average when households experience a scarcity event. The scarcity event mainly induces households with initially low hedge shares to hedge more. When households face a scarcity event, the hedge tariff becomes even more effective in reducing bill volatility. The optimal hedge reduces households' coefficients of bill volatility on average by 48% relative to the unhedged tariff with a scarcity event.

Overall, the scarcity event only slightly increases bill volatility on an optimally hedged tariff, as a comparison of Figures 2.10a and 2.10c reveals. The average coefficient of variation of bills is 18% higher for the optimally hedged tariff with a scarcity event compared to the optimally hedged tariff with a price cap. In contrast, bill volatility on an unhedged tariff increases by 108% for unhedged tariffs when moving from an unhedged price-capped tariff to an unhedged, fully dynamic tariff with a scarcity event. Surprisingly, the optimally hedged tariff only results in slightly higher bill volatility than the fixed tariff, even if a scarcity event occurs.

Figure 2.11c points out that the welfare benefits from hedging are substantially larger with a scarcity event. The average welfare benefit with a scarcity event translates to a 19% reduction in mean electricity prices compared to a 0.3% welfare benefit for the price-capped tariff. This illustrates that welfare benefits from hedging might grow in the future since households on dynamic tariffs will likely face more frequent and more extreme scarcity prices.

2.7 Conclusion

Dynamic electricity prices have an enormous potential to make power markets more efficient since they help align fluctuating electricity demand and supply. However, they expose households to very high electricity prices when supply is scarce. In this paper, I study how to augment dynamic electricity prices with a forward hedge. The forward hedge should protect households from high prices while incentivizing them to adjust their consumption to prices elastically.

Using a utility maximization model, I simulate optimal forward hedge shares for a sample of 2,159 UK households exposed to dynamic prices. My main contribution is that I study the relation between optimal hedge shares and households' price elasticity of demand when they face uncertainty about prices and quantity shocks.

The simulations suggest that the households in my sample should hedge on average 59% of their baseline consumption. Optimal hedge shares differ strongly among households and by time of day. These differences are due to significant variations in the correlation between prices and the desire to consume electricity. Ownership of low-carbon technologies like electric vehicles, solar PV, battery storage, and electric heating is correlated with optimal hedge shares and might contribute to their considerable variation.

The central insight of this paper is that an exogenous increase in price elasticity can increase or decrease hedge shares. Higher price elasticity increases hedge shares when prices and the desire to consume electricity are positively correlated. In this case, the household increases the hedge to ensure that it can maintain an acceptable level of electricity consumption when both prices and its desire to consume are high. In contrast, higher price elasticity reduces hedge shares when the correlation between prices and quantity shocks is negative. Households reduce their hedge shares to raise their exposure to spot prices when spot prices are negatively correlated with their desire to consume.

Optimal forward hedging effectively reduces the volatility of monthly electricity bills by an average of 18%. The reduction in bill volatility is more significant when households face a positive correlation between prices and quantity shocks. The optimal forward hedge can even achieve lower bill volatility than a fixed tariff.

However, despite reducing bill volatility, the welfare gains from optimal forward hedging are minimal. In the main specification, these welfare gains amount to just a 0.3% decrease in mean dynamic electricity prices. Welfare gains from hedging remain small for households with low price elasticity, high risk aversion, and even for households with a much larger share of electricity in overall expenditure.

Welfare gains from hedging are much higher and equivalent to a 19% reduction in mean electricity prices when households face a Texas-style scarcity event with extremely high prices. Hedging might become more valuable in the future when extreme weather and price events will occur more often. However, the given model might not fully capture the benefits of protecting households from extreme worst-case events. Real-world households might care more about avoiding extremely high worst-case bills than about maximizing expected consumption utility or minimizing average bill volatility. Further research should therefore analyze how to design optimal forward hedge tariffs that effectively protect against worst-case outcomes.

In addition, the practical implementation of forward hedging requires further study. Most importantly, it needs to be tested whether domestic consumers understand how forward hedging works and which incentives it creates. The experience with the German gas and electricity price breaks suggests that widespread adoption of forward hedging would require large efforts to educate consumers. Future research also needs to examine the optimal length of the hedge time segments. Moreover, it is crucial to gain a better understanding of how to determine households' baseline consumption levels, especially in the absence of historical price-consumption data. The need to determine a baseline consumption level might also reduce competition in retail power markets. Incumbent suppliers have an advantage in estimating baseline consumption based on historical data and customer characteristics. Therefore, regulators should allow households to share their historical price and consumption data with different retailers to ensure all suppliers can offer equally attractive hedged tariffs.

Despite all these challenges, forward contracts will likely play a vital role as a supplement to dynamic electricity prices in the future. Dynamic electricity pricing will become even more relevant for making electricity markets efficient in a world with primarily intermittent electricity supply. However, consumers and politicians have been hesitant to widely adopt dynamic prices up until now. One of the main reasons for this reluctance is that dynamic electricity prices can make households vulnerable to extremely high prices. Forward hedging is a powerful tool to overcome these concerns as it partly shields households from high prices while preserving the incentive to be price-responsive.

Chapter 3

Hedging electricity price spikes with forwards and options

This chapter is co-authored with Peter Cramton, Professor Emeritus of Economics at University of Maryland and Founder of Forward Market Design LLC., Jason Dark, Chief Technology Officer at Forward Market Design LLC., Darell Hoy, Chief Executive Officer at Forward Market Design LLC., and David Malec, a research associate at the University of Maryland and Chief Data Officer at Forward Market Design LLC.

This chapter contains the current version of the working paper Brandkamp et al. (2025).

Abstract

Day-ahead electricity markets expose power generators and load-serving entities (LSEs) to large financial risks due to extreme price spikes and correlated load fluctuations. We study how price spikes shape optimal hedging strategies for a risk-averse generator and LSE. Drawing on a regime-switching and dispatch model for ERCOT's day-ahead market, we simulate optimal combinations of forward contracts and call options for granular delivery periods. We find that large and frequent price spikes in a delivery period can result in extreme worst-case losses. Agents, therefore, take larger call options and smaller forwards positions in such spiky periods. Yet combining forwards and options only slightly reduces profit variance and worst-case losses compared to a forwards-only strategy because profits are almost linear in day-ahead prices. Finally, increasing the option strike price prompts agents to choose larger forwards and smaller options positions. Our findings underscore

that hedging with granular derivative products effectively reduces profit volatility and worst-case losses even when extreme price spikes occur.

Acknowledgments

This research is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1-390838866.

3.1 Introduction

In recent decades, many countries have designed day-ahead electricity markets to procure electricity at the least cost, to integrate fluctuating load and intermittent renewable generation, and to send reliable scarcity signals (Cramton, 2017). A downside of day-ahead markets is that they expose market participants to large financial risks. These risks occur because load and day-ahead prices are extremely volatile. Day-ahead electricity prices are far more volatile than prices for most other commodities and assets (Bessembinder & Lemmon, 2002, Weron et al., 2004).[1]

Extreme price spikes partly drive day-ahead price volatility. Between 2011 and 2022, the average day-ahead price was \$43/MWh in the Texas ERCOT market, yet price spikes between \$200/MWh and \$9,000/MWh occurred in almost 900 hours (ERCOT, 2024a). These price spikes arise from price-inelastic electricity demand, the lack of short-term storage, market power, and the strong convexity of the merit order curve (Lu et al., 2005). In the medium term, price volatility and spikes will likely increase further due to growing intermittent renewables and more extreme weather (Ketterer, 2014, Peura

[1]For instance, the standard deviation of the percentage changes in average daily day-ahead prices was 25% in ERCOT between 2011-22. The standard deviation of the daily percentage changes in the S&P 500 stock prices was only 0.93% during the same period. For the WTI oil price, the figure is 2.67%. (Bessembinder & Lemmon, 2002) provide similar statistics for PJM between 1997 and 2000 (ERCOT, 2024a).

& Bunn, 2021). In the long term, price volatility and spikes might decline if electricity demand becomes more price-elastic or if batteries provide short-term flexibility.

In addition to price risks, market participants face load fluctuations due to changing weather conditions. High load often coincides with day-ahead price spikes (Afanasyev et al., 2015). Because of this positive correlation between load and price spikes, unhedged load-serving entities (LSEs) face large downside risks in day-ahead markets. Unhedged LSEs must procure large quantities at high spike prices, creating a long left tail in their day-ahead profit distribution. For generators, the positive correlation creates an upside opportunity and a right tail in their profit distribution. Hence, the tails in LSEs' and generators' profit distributions have opposite signs. They require different hedging strategies to hedge their respective right and left profit tails.

This paper simulates optimal hedge strategies for a representative generator and an LSE in ERCOT's day-ahead market. The agents can choose between forward contracts and European call options. Our primary research question is how the frequency and size of price spikes in a delivery period affect optimal hedge strategies. In addition, we study how risk aversion shapes these hedge strategies in the presence of price spikes. We also investigate how the magnitude of the call option's strike price impacts hedging decisions and profit distributions.

We begin by simulating the hourly profits of generators and LSEs in ERCOT's day-ahead market between 2011-22. Profits are simulated using the regime-switching model by Coulon et al. (2013), which captures large price spikes. This model simulates joint distributions of hourly day-ahead prices, hourly net load, and daily gas prices. We also incorporate solar and wind generation profiles from ERCOT (2023). For each simulated draw of hourly day-ahead price, gas price, and net load, we run a merit order model to obtain hourly dispatch and profits for a sample of 655 power plants in Texas provided by Mann et al. (2017) and ERCOT (2023). We assume that agents' forwards and option positions do not affect bidding behavior or day-ahead prices.

The regime-switching and dispatch models provide realistic distributions of hourly day-ahead profits. Using these distributions, we simulate optimal quantities of forwards and call options for a representative generator and LSE. We define a pair of forward and option quantities as optimal if it maximizes the agent's constant absolute risk aversion utility.

Following Wolak (2022) and Cramton et al. (2025b), we consider forward and option contracts with granular delivery periods: combinations of year-month-hour-weekend/weekday, such as weekdays 4-5 pm in August 2019. These granular delivery periods allow agents to better manage risks by tailoring hedges to daily and seasonal load or generation cycles (Winzer et al., 2024, Brandkamp, 2025). We assume that agents can buy the granular hedge products at arbitrage-free prices. Ignoring arbitrage opportunities allows us to focus on the risk-management motive for hedging. Our main specification sets a high strike price of $200/MWh for the call option because the option is intended to primarily cover large price spikes.

We simulate optimal forward and option quantities for all 576 unique granular delivery periods in 2019. We show that the frequency and size of day-ahead price spikes differ sharply across these periods. Spikes tend to be small and rare at night and in winter but larger and more frequent during daytime and summer. These varying levels of ”spikiness” allow for the analysis of how price spikes shape hedge strategies.

Large price spikes can lead to extreme worst-case losses. We find that the LSE and the generators trade-off minimizing worst-case losses versus minimizing profit variance when choosing optimal hedge strategies. They do not optimize expected profit because, with arbitrage-free prices, expected profit is constant no matter how many forwards and options the agent buys. In periods with larger price spikes, the generator and the LSE both rely more on options than forwards to mitigate worst-case losses. The LSE uses options more heavily because price spikes create longer negative profit tails for the LSE than for the generator. Agents' risk aversion is crucial for determining how much they focus on worst-case losses or on variance.

Agents who are nearly risk-neutral focus on minimizing profit variance (Min-Var). In low-spike periods, both sides hold forwards and options in comparable proportions. The LSE typically buys around six times as many forwards as options, while the generator sells five times as many forwards as options. This results in variance-minimizing forward volumes of roughly 90-110% of expected generation for the generator and 100-110% of expected load for the LSE. Min-Var option positions are smaller (about 15-20% of expected quantities for the LSE and 20% for the generator)

In periods with heavier price spikes, the nearly risk-neutral generator only slightly lowers its short-forward ratio and keeps its option ratio constant. By contrast, the nearly risk-neutral LSE strongly reduces its long forward ratio to roughly 80% and increases its option ratio to 30-40%, ending up close to a 2:1 forward-to-option mix. Thus, the LSE heavily relies on options to protect against price spikes, even if it is nearly risk-neutral.

The more risk-averse the LSE becomes, the more it shifts focus to ever more extreme parts of its worst-case loss tail. It, therefore, replaces forwards with options, reducing their proportion to almost 2:1 in periods with small spikiness. In spiky periods, the LSE buys even more options than forwards. This strategy reduces tail losses but raises profit variance compared to the Min-Var hedge strategy.

When the generator becomes risk-averse, it reduces its short option position. Shorting options can trigger large losses in spiky delivery periods. For high-risk aversion, the generator, therefore, prefers to buy options rather than sell them. A long option position offsets the substantial losses from shorting forwards when price spikes occur. The downside is that buying call options increases the variance of the generator's profits because it leads to positive payoffs when the generator's day-ahead profit is high. Therefore, the risk-averse generator increases its short forward position to balance the long option payoff in spiky periods. The generator's long option position is smaller than the LSE's because it has a less extreme worst-case loss tail. The reason is that the generator is more flexible in reducing its generation quantities when prices are unfavorable.

Overall, combining forwards and options is effective in reducing profit variance and worst-case risks compared to a fully unhedged strategy. However, we also compare the combined strategy to one that only hedges with forwards. Surprisingly, the combined strategy only marginally reduces profit variance and achieves only small benefits with respect to worst-case losses compared to the forwards-only strategy. The reason is that the generator's and LSE's day-ahead profits are roughly linear in the day-ahead price. Under linear payoffs, linear forwards hedge effectively, and nonlinear options add little value (Lapan et al., 1991). This explains why nearly risk-neutral agents use few options. The small added benefit of reduced worst-case losses justifies higher option holdings for more risk-averse agents.

Finally, we study how the magnitude of the call option strike price affects hedging in a delivery period with large price spikes. We analyze strike prices ranging from 100\$/MWh and $1,000$\$/MWh. The LSE buys more forwards and fewer options when the strike price rises. With a high strike price, the option offers no protection against moderate spikes, making the option less attractive compared to the forwards. More risk aversion reinforces this tendency, with increased forward buying and decreased option buying. The almost risk-neutral generator also increases its short forwards and decreases its short options when the strike rises. If the generator is risk-averse, long option positions increase in the strike price for low strike prices below 400\$/MWh but decline for higher strike prices.

A low strike price slightly increases agents' profit variance. On the other hand, a low strike price enables the LSE to choose hedge strategies that lead to smaller tail losses, at least when its risk aversion is low. The reason is that the low strike price induces the LSE to buy more options, leading to lower tail losses but higher variance. In contrast, the higher option holdings chosen under low strike prices increase tail losses for the generator unless it is very risk-averse.

It, therefore, depends on agents' risk preferences, which strike price the market designer should choose. A smaller strike price has the advantage of offering better spike protection for the LSE. The LSE has less flexibility and is more vulnerable to extreme negative profit tails than the generator. Some researchers also opt for low strike prices because low strikes might improve generators' strategic bidding incentives in day-ahead markets (Willems, 2005, Zhang et al., 2012). Yet, determining the exact strike price level requires further research.

Our paper shows that forwards and options enable participants in day-ahead markets to effectively manage risk created by extreme price spikes. Regulators should encourage hedging by establishing liquid markets for forwards and options with standardized contracts. Obliging LSEs to hedge a particular share of their expected load could also increase liquidity and coordinate trade (Cramton et al., 2025b).

High liquidity will mitigate price premia and arbitrage between day-ahead and derivative markets. Regulators could further reduce price premia by facilitating arbitrageurs and financial institutions to trade in electricity markets. However, some level of price premia will likely persist even in liquid electricity markets. It would be interesting to study how the above optimal hedging strategies change when arbitrage is possible. Finally, future research could study in which market equilibria the generator's and LSE's strategies result.

The following section relates our paper to the literature on optimal hedging strategies in electricity markets. Section 3.3 introduces the generator's and LSE's profit and utility functions. Section 3.4 describes the simulation of the day-ahead market profits, including the regime-switching model and the merit order dispatch model. Section 3.5 defines the forwards and option contracts and calculates arbitrage-free forwards and option prices. Section 3.6 simulates optimal hedging strategies and examines how they depend on price spikes and risk preferences. Section 3.7 concludes.

3.2 Literature review

Hedging is essential in day-ahead markets for electricity generators and load-serving entities (LSEs). Their load and day-ahead prices fluctuate strongly, creating extreme financial risks.

Forward contracts are the most widely used hedging instrument for generators and LSEs (Deng & Oren, 2006). A comprehensive literature identifies the optimal forward quantity a commodity producer or consumer should buy. Danthine (1978), Holthausen (1979), Feder et al. (1980), and McKinnon (1967) show that when price and quantity uncertainties are independent, agents should sell exactly their expected quantity forward. They can sell less than their expected quantity forward if they are flexible to change their quantity in response to fluctuating prices, as most electricity generators are (Moschini & Lapan, 1992, Dionne & Santugini, 2015). McKinnon (1967), Rolfo (1980), Losq (1982), Lapan and Moschini (1994), and Borenstein (2007b) reveal that agents hedge more than their expected quantity when prices and output are positively correlated. LSEs often face such a positive correlation. Jin (2007) and Lien (2010) show that a right-skewed profit distribution also increases the optimal forward position for producers.

Building on the above theoretical insights, many studies identify the optimal quantity of forward contracts for electricity generators (Näsäkkälä & Keppo, 2005, Conejo et al., 2008) and load-serving entities (Woo et al., 2004, Kettunen et al., 2009, Deng et al., 2020). Yet, several researchers argue that linear forward contracts alone are an ineffective hedging instrument if day-ahead market profits are not linear in the day-ahead price. Nonlinearity arises if load and prices are correlated (Sakong et al., 1993), if load or generation is price-elastic (Moschini & Lapan, 1992), or if cost functions are nonlinear.

Various papers demonstrate that nonlinear profits can be hedged effectively by combining linear forwards with nonlinear derivatives such as options (Moschini & Lapan, 1992, Moschini & Lapan, 1995, Oum et al., 2006). Oum and Oren (2010), Zhou et al. (2017), and Hess (2021) describe the optimal combination of forwards and options for LSEs. They emphasize that hedging with options is beneficial for LSEs since they face large non-traded uncertainty in load that is correlated with prices. Oum et al. (2006) and Oum and Oren (2010) derive the optimal payoff function of an exotic option for an LSE and show that this optimal payoff function can be replicated with a portfolio of forwards and standard European options. For generators, Azevedo et al. (2007), Sanchez et al. (2010) and Ocakoglu and Tolga (2018) also stress the benefits of combining options and forwards. Other authors emphasize the advantages of combining electricity forward contracts with exotic swing or Asian options (Hambly et al., 2008, Fanelli et al., 2016), fuel-spread options (Aïd et al., 2013), or weather derivatives (Lee & Oren, 2009, Bhattacharya et al., 2020, Matsumoto & Yamada, 2021).

While options and forwards are useful risk-management tools, agents can also use them to speculate if forward and option prices are not arbitrage-free. Forwards and options can skew the profit distribution to the right if arbitrage opportunities exist (Lapan et al., 1991, Vercammen, 1995). Lien (2010) show that price spikes increase the optimal forward position of producers when the forward price is larger than the expected day-ahead price and reduce the forward position when the forward price is negatively biased. Arbitrage opportunities between day-ahead, and forward and option markets are widespread in electricity markets and are an important driver of hedging behavior (Redl et al., 2009, Botterud et al., 2010). Nevertheless, in this paper, we focus on arbitrage-free prices to study the risk management motive for hedging in isolation.

Large day-ahead price spikes are another important factor that shapes LSEs' and generators' hedge strategies. Price spikes make generators' day-ahead profit distribution right-skewed and LSEs distribution left-skewed. Vercammen (1995) and Barbi and Romagnoli (2018) point out that agents can use forwards and options to reshape their profit distributions. LSEs can go long in forwards and options to offset their negative tail and thereby reduce worst-case losses. For generators, spike prices lead to large upside potential that is reduced by short positions. Vercammen (1995) shows that risk-averse generators might still prefer short forward and option positions to lower profit volatility. However, short positions can cause large worst-case losses if price spikes occur. Bessembinder and Lemmon (2002) suggest that such spikes create a risk premium in forward and option prices since risk-averse generators and speculators require compensation for short positions. LSEs, meanwhile, should be willing to pay this premium to avoid large losses under spike conditions.

Traditionally, researchers model price spikes by assuming lognormally distributed day-ahead prices, enabling closed-form solutions for forward and option positions under mean-variance or CARA utility (Oum et al., 2006). In recent years, however, price spikes have become so extreme that lognormal distributions no longer fit. Numerous authors instead use regime-switching models to describe these extreme price spikes (Hamilton, 1990, de Jong & Huisman, 2002, Weron et al., 2004). Coulon et al. (2013) simulate quantities for forwards and options, respectively, that separately minimize profit variance for an LSE. They also provide suggestive evidence that combining options and forwards will further reduce profit variance. Yet, LSEs and generators may also worry about downside tail risks if prices spike severely, so they might not minimize variance alone. Risk-averse LSEs with large negative profit tails could prioritize worst-case losses over mere variance reduction.

We, therefore, extend Coulon et al.'s (2013) in three ways. First, we derive combined optimal forward and option positions for LSEs and generators under extreme price spikes, assuming CARA utility. We study how the level of risk aversion influences hedging strategies and profit distributions. Supplementing Coulon et al.'s (2013) analysis, we also compare a combined forward–option strategy to a forward-only strategy.

Second, we examine how option and forward quantities change with the "spikiness" of a given delivery period's price distribution, measured via the option-to-forward price ratio, which is strongly correlated with spike size and frequency.

Third, we investigate how the call option's strike price influences hedging effectiveness in high-spike conditions. This analysis builds on Oum et al. (2006), and shows how forward and option quantities depend on the option's strike prices. We study how the strike price affects profit variance and tail profits.

3.3 Hedging with forwards and options

This section describes how electricity generators and load-serving entities (LSE) determine the optimal quantities of forwards and call option contracts. LSE's and generator's (GEN) respective profits from selling electricity in the day-ahead market in hourly period t are

$$\pi_{DA_t}^{LSE} = (R_Y - P_t)q_t \qquad\qquad \pi_{DA_t}^{GEN} = P_t q_t - c_t(q_t)$$

P_t denotes the day-ahead electricity price in hour t. q_t is the load that the LSE procures or the generator produces. R_Y is the fixed retail price in year Y that the LSE charges to its end-consumers. $c_t(q_t)$ represents the generator's cost function. The shape of $c_t(q_t)$ changes over time since generation costs change with fuel prices.

Generators and LSEs can hedge their day-ahead market risk by purchasing forwards or call option contracts. The focus on forwards and call options is not restrictive because the combination of these contracts can replicate positions equal to many other classical derivatives, including put and straddle positions (Cox & Rubinstein, 1985). When buying forwards and call options, agent $i \in \{\text{LSE, GEN}\}$ has the following profit function in each hourly period $t \in M$:

$$\pi_t^i = \pi_{DA_t}^i - (F_M^p - P_t)h_M - (V_M - v_t)z_M$$

h_M is the quantity of forward contract M agent i buys. M defines the delivery period that the forward and option contracts cover. For instance, M could contain all baseload hours in June 2026. If $h_M = 1MW$, this means that the agent would buy $1MW$ of forward energy in all baseload hours in June 2026 (Coulon et al., 2013). F_M^p is the forward price for contract M. V_M is the price for the call option, z_M is the quantity of option contracts, and v_t is the gross value of the call option defined as:

$$v_t = \begin{cases} P_t - K_M & \text{if } P_t \geq K_M \\ 0 & \text{if } P_t < K_M \end{cases}$$

K_M is the option's strike price. Forward and option quantities h_M and z_M can be positive or negative. A negative quantity means that agent i sells electricity forward or sells a call option. It might seem natural that LSEs want to buy electricity forward and generators want to sell it. However, LSEs and generators can buy positive and negative quantities. Importantly, we assume that the chosen forwards and option prices do not affect prices and quantities in the day-ahead markets. Several authors reveal that forward and option positions improve bidding incentives and make day-ahead market equilibria more competitive (Allaz & Vila, 1993, Willems, 2005, Zhang et al., 2012). We abstract from these effects in our day-ahead market simulation.

When agent i chooses the optimal quantities of forwards h_M and options z_M, she faces uncertainty about the day-ahead price P_t and her realized load q_t in period t. The generator's cost function $c_t(q_t)$ is also uncertain when forward and option quantities are chosen because $c_t(q_t)$ depends on fuel prices in period t. Given this uncertainty, agent i chooses the optimal z_M and h_M for contract M that maximize its expected utility:

$$\underset{h_M, z_M}{max}\, E[U(\pi^i_{t \in M})]$$

The first-order conditions are given as:

$$\begin{aligned}
\frac{\partial E[U(\pi^i_t)]}{\partial H} &= E[U'(\pi^i_{t \in M})(F^p_M - P_{t \in M})] = 0 \\
\frac{\partial E[U(\pi^i_t)]}{\partial Z} &= E[U'(\pi^i_{t\ inM})((v_{t \in M} - V_M))] = 0
\end{aligned}$$

If we aimed to solve this problem analytically, we would have to make assumptions about the shape of the utility function and the joint distribution of the random factors prices, load, and costs (Moschini & Lapan, 1995). Several authors solve variants of this problem analytically, assuming a constant absolute risk aversion (CARA) utility function and joint normally distributed load and prices (Lapan et al., 1991, Moschini & Lapan, 1995). We will not follow their approach because load and prices are not jointly normally distributed for electricity generators and LSEs. The main reason is that the electricity price distribution is strongly right-skewed. Assuming a normal distribution of electricity prices is not suitable for us, since we are interested in how to hedge when prices exhibit large spikes. Hence, we will assume a CARA utility function and solve the above problem numerically. To do so, the following section describes a regime-switching model of the ERCOT day-ahead market that allows simulating strongly skewed distributions of prices, load, and costs.

3.4 Day-ahead market simulation

3.4.1 Regime-switching model

This section aims to simulate realistic distributions of generators' and LSEs' day-ahead market profits. We follow Coulon et al. (2013) and build a regime-switching model that allows estimating the joint distribution of the most relevant determinants of day-ahead profits: day-ahead prices, gas prices, net load, and solar and wind generation. To calibrate the model, we use data on hourly day-ahead prices from ERCOT (2024a), and hourly system load data from ERCOT (2024b) for all hours from 2011 to 2022. Moreover, we use daily Henry Hub gas price data for 2011-22 from EIA (2024). We also use hourly generation profiles for 218 wind farms and 189 solar farms that operate in ERCOT. The generation profiles were created by ERCOT (2023). In addition, we use data on hourly solar and wind generation from ERCOT (2024c).

Figure 3.1a depicts the average daily day-ahead electricity price in ERCOT and the daily Henry Hub gas price between 2011 and 2022. Modeling the day-ahead electricity prices is challenging because they experience extreme spikes. Moreover, day-ahead prices follow strong seasonal, weekly, and daily cycles. The literature that aims to model these features can be broadly categorized into two streams. The first stream consists of structural models. These models derive electricity prices from a merit order approach that considers the costs of all power plants and their operational constraints. An advantage of a structural model is that it offers an intuitive explanation of price formation and allows analyzing how capacity expansion and changes in the resource structure affect electricity prices (Carmona et al., 2013). A disadvantage is that structural models struggle to explain extreme price spikes where prices far exceed generators' marginal costs. The second literature stream uses reduced-form models. Such models describe the evolution of electricity prices as a pre-specified stochastic process. They often include price spikes using regime-switching or jump processes. However, reduced models do not capture how price spikes depend on the factors that influence electricity prices, like fuel prices, outages, and load (Carmona et al., 2013). Most importantly, for our purpose of deriving optimal hedge positions, a model of day-ahead profits must capture the fact that price spikes are positively correlated with load and gas prices.

Following Barlow (2002), several authors combined the two above literature streams. They describe electricity day-ahead prices as a function of relevant determinants, especially load and fuel prices. These determinants are assumed to follow correlated stochastic processes. In some models, the determinants influence the probability of a price jump or regime switch to account for price spikes. Most models include load (Barlow, 2002) and

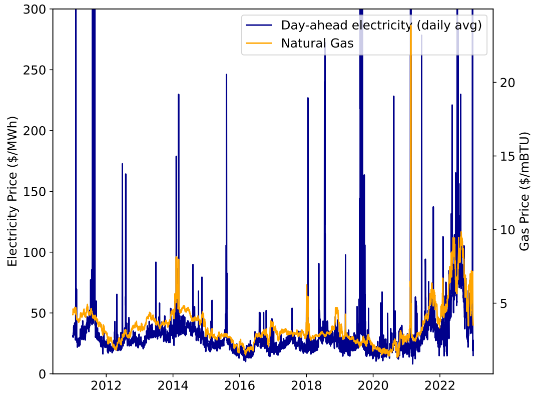

(a) Daily Average Electricity and Gas Prices (Electricity Prices Capped at $300/MWh)

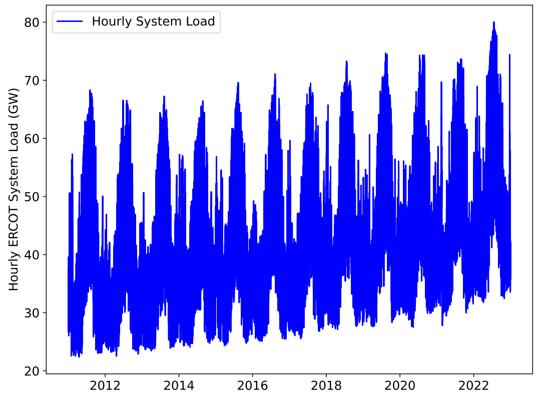

(b) Hourly ERCOT System Load

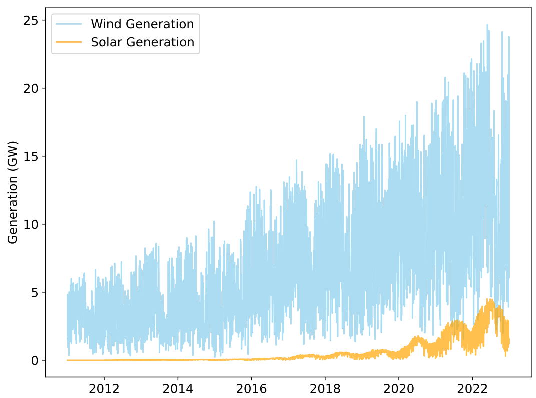

(c) Daily Average Aggregate Wind and Solar Generation in ERCOT

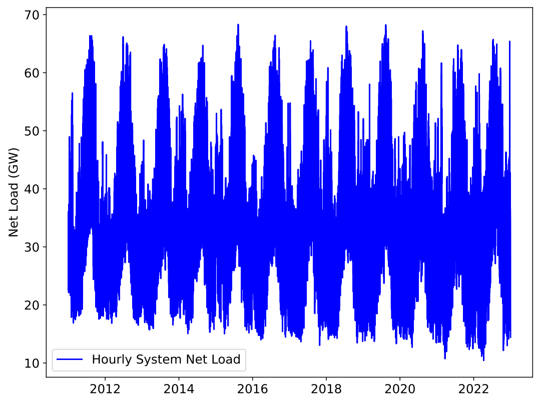

(d) Hourly ERCOT System Net Load

Figure 3.1: Empirical data from ERCOT on prices, renewable generation, and load in 2011-22

fuel prices (Carmona et al., 2013) as crucial determinants of power prices. The literature considers gas as a marginal fuel to explain most variation in electricity prices (Füss et al., 2015). In Carmona et al. (2013), the electricity price depends on multiple fuels like coal and gas. Other authors also include plant outages as a determinant of power prices.

This paper employs such a combined regime-switching model by Coulon et al. (2013). In their model, the following equation describes the hourly day-ahead electricity price P_t:

$$P_t = G_t \exp\left(\alpha_{m_k} + \beta_{m_k} L_t + \gamma_{m_k} X_t\right) \tag{3.1}$$

G_t is the daily gas price, and L_t is the hourly net load, i.e., load minus intermittent solar and wind generation. X_t is a residual process and $m_k \in \{1, 2\}$ indicates whether prices are in the normal price regime 1 or in the spike price regime 2. The above exponential expression for the electricity spot price is commonly used in the literature (Eydeland & Geman, 1999). The advantage of this exponential form is that the resulting day-ahead prices P_t are lognormally distributed in each price regime. This allows for deriving closed-form solutions for forward and option prices (Coulon et al., 2013, Füss et al., 2015).

Following Coulon et al. (2013), we rely on gas as the price-setting marginal fuel. Figure 3.1a reveals that average daily electricity prices are closely linked to daily gas prices, even though daily electricity prices show larger spikes than gas prices. We model logarithmic gas prices G_t as an Ornstein-Uhlenbeck (OU) process (Coulon et al., 2013, Schwartz, 1997).

$$d \log G_t = \kappa_G (m_G - \log G_t) dt + \eta_G dW_t^{(G)} \tag{3.2}$$

In contrast to Coulon et al. (2013), we use net load L_t rather than load as a second determinant of power prices in equation 3.1. Coulon et al. (2013) use load because solar and wind generation was still negligible in Texas during their observation period between 2005 and 2011. Over the last decade, the rising share of wind and solar generation has become an important determinant of day-ahead electricity prices. Solar and wind plants have near-zero marginal costs and are usually dispatched when there is sun or wind and no curtailment. Since intermittent renewable generation shifts the merit order curve, net load rather than load determines prices in day-ahead markets with significant renewable generation (Ketterer, 2014, Peura & Bunn, 2021).

Figure 3.1b reveals that ERCOT's hourly system load shows strong, predictable seasonal cycles and is increasing over time. Figure 3.1c highlights that the aggregate solar and wind generation is also increasing over time. Surprisingly, Figure 3.1d indicates that hourly net load is not increasing even though load and solar and wind generation are all increasing. Even the daily standard deviation of the net load stays relatively stable over time. This suggests a strong positive correlation between load and solar and wind generation in ERCOT.

Following Coulon et al. (2013), we decompose net load L_t into a seasonal component $S(t)$ and a deseasonalized component $\overline{L}_t$.

$$L_t = S(t) + \overline{L}_t \tag{3.3}$$

Seasonal component $S(t)$ is estimated separately for every hour h of the day.

$$S(t) = a_1(h) + a_2(h)\cos(2\pi t + a_3(h)) + a_4(h)\cos(4\pi t + a_5(h)) + a_6(h)t + a_7(h)1_{we} \tag{3.4}$$

a_1 represents the long-term average net load in hour h. a_2 controls the amplitude of seasonal variations, while a_3 sets the starting point of the seasonal cycles. Similarly, a_4 and a_5 capture higher-frequency daily net load cycles. a_6 describes long-term linear trends, and a_7 considers that net load is typically lower on weekends. Deseasonalized net load $\overline{L}_t$ is modeled by the following OU process:

$$d\overline{L}_t = -\kappa_L \overline{L}_t \, dt + \eta_L \, dW_t^{(L)} \tag{3.5}$$

The above process models deseasonalized net load to be normally distributed and mean-reverting. The deseasonalized net load process and the gas price process in equation 3.2 are assumed to be uncorrelated (i.e., $W_t^{(L)}$ and $W_t^{(G)}$ are independent).

While net load and gas prices are crucial, there are additional important determinants of power prices, such as generator or transmission outages. Coulon et al. (2013) do not explicitly model these factors but include them as a residual, unobserved process X_t.

$$X_t = S_X(t) + \overline{X}_t \tag{3.6}$$

$S_X(t)$ is the seasonal component of X_t that accounts for seasonal and daily cycles in the frequency of outages and transmission constraints.

$$S_X(t) = b_1(h) + b_2(h)\cos(2\pi t + b_3(h)) + b_4(h)\cos(4\pi t + b_5(h)) \tag{3.7}$$

As above, the deseasonalized residual $\overline{X}_t$ follows an OU process

$$d\overline{X}_t = -\kappa_X \overline{X}_t \, dt + \eta_X \, dW_t^{(X)} \tag{3.8}$$

The random elements of the deseasonalized net load $W_t^{(L)}$ and the deseasonalized residual load $W_t^{(X)}$ may be correlated. Transmission or generator outages are more likely under extreme weather conditions when net load is typically high. $\overline{L}_t$ and $\overline{X}_t$ are OU processes with mean zero since their means are subtracted by their seasonal components $S(t)$ and $S_X(t)$.

In day-ahead price equation 3.1, indicator m_k determines whether prices are in normal regime 1 or spike regime 2. m_k is chosen by an independent coin flip:

$$m_k = \begin{cases} 1 & \text{with probability } 1 - p_s \Phi\left(\frac{\overline{L}_t - \mu_s}{\sigma_s}\right) \\ 2 & \text{with probability } p_s \Phi\left(\frac{\overline{L}_t - \mu_s}{\sigma_s}\right) \end{cases}$$

The probability of being in spike regime 2 increases in deseasonalized net load $\overline{L}_t$. $\phi(\cdot)$ is the normal cumulative distribution function. p_s gives the probability of being in spike regime 2 when $\overline{L}_t$ goes to infinity. μ_s and σ_s are parameters that govern how the spike probability depends on deseasonalized net load. Following Coulon et al. (2013), we set $\mu_s = 0$ and $\sigma_s = \frac{\eta_L}{\sqrt{2\kappa_L}}$ such that the parameters equal the mean and standard deviation of the stationary distribution of observed deseasonalized net load $\overline{L}_t$. Thereby, the probability of a price spike becomes linear in the quantile of $\overline{L}_t$ (Coulon et al., 2013).

Day-ahead price equation 3.1 has the same exponential form for regimes 1 and 2. Since the parameters α_{m_k}, β_{m_k}, and γ_{m_k} differ between the two regimes, the model allows for a steeper relationship between net load and prices in spike regime 2 (Coulon et al., 2013).

3.4.2 Simulating day-ahead market outcomes

In this section, we simulate hourly day-ahead market states, i.e., we draw values from the joint distribution of day-ahead electricity prices, gas prices, net load, and solar and wind generation. We first estimate the parameters of the day-ahead price equation 3.1 using

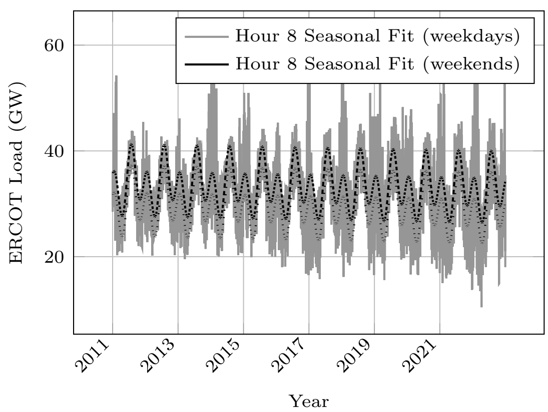

(a) Hour 8 Seasonal Component

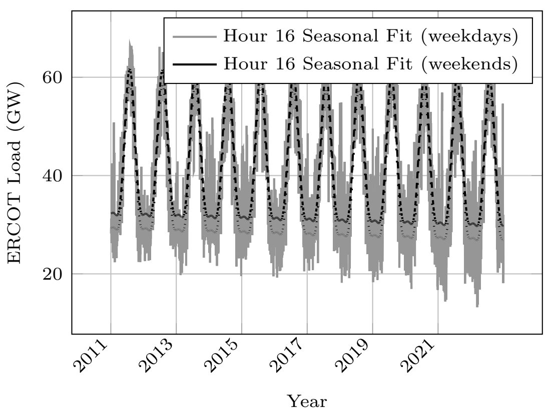

(b) Hour 16 Seasonal Component

Figure 3.2: ERCOT system load seasonal components for hours 8 and 16.

data on hourly day-ahead electricity prices, net load, and daily gas prices between 2011 and 2022. To do so, we fit the coefficients of the seasonal net load component $S(t)$ in equation 3.4. The estimated coefficients can be found in Table C.1 in Appendix C.1 for every hour of the day. Figure 3.2 shows net load and its seasonal components for weekends and weekdays between 8 am and 4 pm.[2] The figure reveals that seasonal variation in net load is much larger for the peak hour 4 pm compared to the less peaky hour 8 am. After calculating the seasonal component of net load, we obtain the deseasonalized net load $\overline{L}_t$ using equation 3.3.

We can now estimate the remaining parameters in price equation 3.1 using the likelihood function:[3]

$$\mathcal{L}^P(\alpha_1, \alpha_2, \beta_1, \beta_2, \gamma_1, \gamma_2, p_2) = \prod_t \Bigg\{ p_s \phi\left(\frac{\overline{L}_t - \mu_s}{\sigma_s}\right) \frac{1}{\gamma_2\sqrt{2\pi}} \exp\left(-\left(\frac{\log\left(\frac{P_t}{G_t}\right) - \alpha_2 - \beta_2 L_t}{\gamma_2}\right)^2\right) \tag{3.9}$$
$$+ \left(1 - p_s \phi\left(\frac{\overline{L}_t - \mu_s}{\sigma_s}\right)\right) \frac{1}{\gamma_1\sqrt{2\pi}} \exp\left(-\left(\frac{\log\left(\frac{P_t}{G_t}\right) - \alpha_1 - \beta_1 L_t}{\gamma_1}\right)^2\right) \Bigg\}$$

The estimation models the deseasonalized residual process X_t as normal random noise (Coulon et al., 2013). The estimated parameters are given in Table 3.1.

α_1	β_1	γ_1	α_2	β_2	γ_2	p_s
1.279	2.39×10^{-5}	0.308	-0.035	7.12×10^{-5}	1.065	0.161

Table 3.1: Estimated parameters for the day-ahead price function in equation 3.1

As expected, $\beta_2 > \beta_1$, which means that the exponential relation between net load and prices is steeper in spike price regime 2 than in price regime 1. $\gamma_2 > \gamma_1$ implies that random shocks like plant outages also have a stronger effect on prices in the spike regime than in the normal regime. $p_s = 16.1\%$ reveals that prices are in the spike regime in $8.05\%(p_s/2)$ of hours because the spike probability $p_s \Phi\left(\frac{\overline{L}_t - \mu_s}{\sigma_s}\right)$ fluctuates symmetrically between 0 and p_s in each hour t (Coulon et al., 2013).

Using the estimated parameters in Table 3.1, we rearrange price equation 3.1 to back out the residual process X_t.

[2]Figure 3.2 was also created by Coulon et al. (2013) and is replicated here for our different observation period.

[3]We follow Coulon et al. (2013) and remove all hourly observations where $P_t/G_t < 0.1$, i.e., the electricity price is very small relative to the gas price. Only 6 out of 105,158 hours are removed. Eliminating these small outliers is supposed to improve the fit of the likelihood estimation (Coulon et al., 2013).

$$X_t = p_s\phi\left(\frac{\overline{L}_t-\mu_s}{\sigma_s}\right)\left(\frac{\log\left(\frac{P_t}{G_t}\right)-\alpha_2-\beta_2L_t}{\gamma_2}\right)+\left(1-p_s\phi\left(\frac{\overline{L}_t-\mu_s}{\sigma_s}\right)\right)\left(\frac{\log\left(\frac{P_t}{G_t}\right)-\alpha_1-\beta_1L_t}{\gamma_1}\right)$$

Afterward, we employ equation 3.7 to fit the seasonal component $S_X(t)$ (see Table C.1 in Appendix C.1 for the estimated coefficients) and obtain the deseasonalized residual process $\overline{X}_t$ using equation 3.6. Having determined $\overline{L}_t$ and $\overline{X}_t$, we now jointly estimate the parameters of the correlated OU processes $d\overline{L}_t$ and $d\overline{X}_t$ in equations 3.5 and 3.8 by maximizing the following likelihood function

$$\mathcal{L}^{OU}(\kappa_L,\eta_L,\kappa_X,\eta_X,\nu)=\prod_{t=2}^{N}\frac{1}{2\pi\sqrt{\det(\Sigma)}}\exp\left(-\frac{1}{2}\mathbf{I_t}^T\Sigma^{-1}\mathbf{I_t}\right)$$

with $\mathbf{I_t}=\begin{pmatrix}d\overline{L}_t\\ d\overline{X}_t\end{pmatrix}$. The variance-covariance matrix of the bivariate normal processes is

$$\Sigma=\begin{pmatrix}\sigma_L^2 & \nu\sigma_L\sigma_X\\ \nu\sigma_L\sigma_X & \sigma_X^2\end{pmatrix} \tag{3.10}$$

ν is the correlation coefficient between $d\overline{L}_t$ and $\overline{X}_t$, $\sigma_L^2=\frac{\eta_L^2(1-e^{-2\kappa_L dt})}{2\kappa_L}$ is the variance of $d\overline{L}_t$, and $\sigma_X^2=\frac{\eta_X^2(1-e^{-2\kappa_X dt})}{2\kappa_X}$ is the variance of $\overline{X}_t$.

Similarly, we maximize the following likelihood function to estimate the parameters of the OU process of the log gas price, assuming that $d\log G_t$ is not correlated with $d\overline{L}_t$ and $d\overline{X}_t$.

$$\mathcal{L}^{G}(\kappa_G,m_G,eta_G)=\frac{1}{\sqrt{2\pi\eta_G^2dt}}\exp\left(-\frac{(d\log G_t-\kappa_G(m_G-\log G_t)dt)^2}{2\eta_G^2dt}\right)$$

The estimated parameters for the three OU processes are shown in Table 3.2.

κ_L	η_L	κ_G	η_G	m_G	κ_X	η_X	ν
125.571	84,329.524	3.524	0.915	1.159	996.966	40.406	0.092

Table 3.2: Estimated parameters relating to the stochastic processes for G_t, L_t, and X_t.

$\kappa_G=3.524$ reveals that gas prices need more than two months to revert to their mean. However, gas prices mean-revert much faster from 2011-22 than from 2005-11, as estimated by Coulon et al. (2013). $\kappa_L=125.571$ indicates that the net load typically requires around two days to revert to its seasonal mean, reflecting net load's dependence on multi-day weather movements. Net load mean reverts slightly slower than load in Coulon et al.'s (2013) earlier time period because net load is likely more weather-dependent than load. Deseasonalized net load is volatile, as the large η_L indicates. $\kappa_X = 996.966$ implies that the residual process reverts to its seasonal mean after roughly 6 hours. Plant or transmission outages likely only affect prices for a couple of hours, as Coulon et al. (2013) also find.

Using the above estimated parameters, we simulate the deseasonalized OU processes for net load, log gas prices, and the residual process. After adding their seasonal components, we insert the simulated values for net load, gas price and the residual process into equation 3.1 to calculate the corresponding day-ahead electricity prices.

To examine how well the simulated prices and quantities match the observed empirical values, we simulate 20 market states for each of the 105,158 observed hourly periods, leaving us with more than 2.1 million simulated market states. Figure 3.3 compares the histograms of the simulated hourly day-ahead prices (left side) with the histograms of the day-ahead prices that were observed in ERCOT between 2011-22 (right side). Figures 3.3a and 3.3b show the distribution of the simulated and observed "normal" prices below \$200/MWh. The simulated and observed distributions below \$200/MWh look similar and resemble a lognormal distribution. The standard deviation of the simulated prices below \$200/MWh is \$23.1/MWh, slightly smaller than the standard deviation of \$25.0/MWh for the observed prices.

Figures 3.3c and 3.3d depict the distribution of price spikes above \$200/MWh. 1.4% of the simulated prices and 1.7% of the observed prices exceed \$200/MWh. In rare cases, the regime-switching model simulates unreasonably large day-ahead prices since the model does not consider the administrative price cap of \$9,000/MWh. In only 0.007% of the simulated hours, prices exceed the \$9,000/MWh cap. To make the simulated prices consistent with the actually observed prices, we also cap prices at \$9,000/MWh. Simulated spike prices are mostly smaller than \$4,000/MWh. In contrast, a larger share of observed spike prices falls between \$4,000/MWh and \$9,000/MWh. This shows that there is room for improving our simulation of the long right-tail of day-ahead prices. However, our model captures price spikes in ERCOT reasonably well to analyze the effects of the price spikes on hedging behavior.

In Appendix C.3, we also compare the simulated and observed distributions for gas prices, net load, and load. To derive load, we also simulate aggregate and plant-level wind and solar generation using plant-level wind and solar generation profiles provided by ERCOT (2023). The model we use to simulate plant-level hourly wind and solar generation is presented in Appendix C.2.

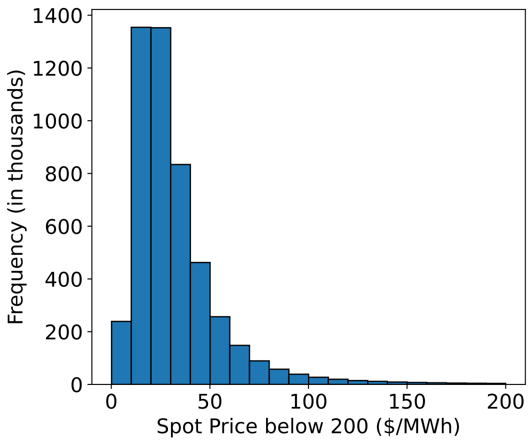

(a) Simulated day-ahead prices below $200/MWh

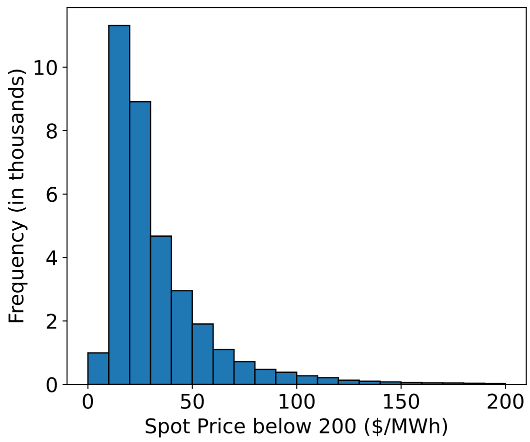

(b) Observed day-ahead prices below $200/MWh

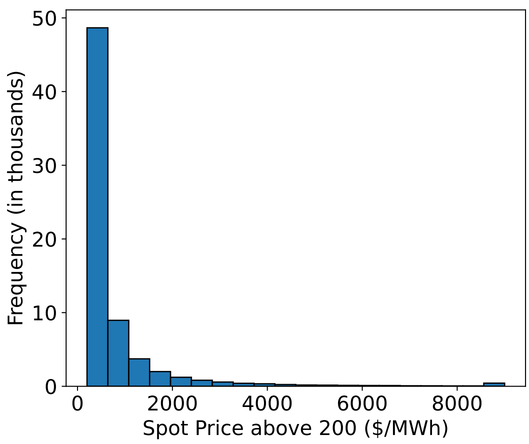

(c) Simulated day-ahead prices above $200/MWh

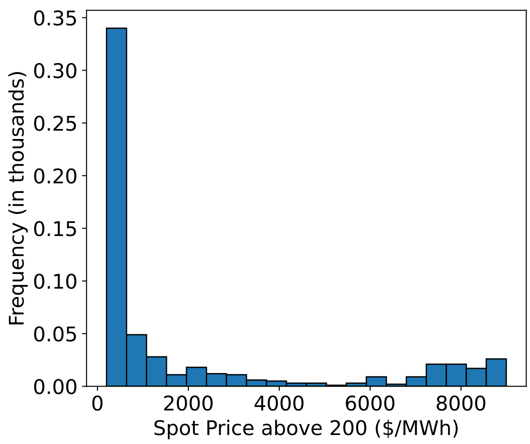

(d) Observed day-ahead prices above $200/MWh

Figure 3.3: Histogram of day-ahead electricity prices below and above $200/MWh for simulated and observed data between 2019-22

3.4.3 Day-ahead market quantities and profits

We calculate the generators' and LSEs' load in every market state based on the above day-ahead market states. For LSEs, load is assumed to be price-inelastic and to be related to aggregate load Q_t as follows:

$$q_t = sQ_t + \varepsilon_t$$

s is a fixed market share of the LSE. ε_t is a random noise process such that the LSE's load fluctuates around its share in aggregate load (Peura & Bunn, 2021).

Technology	**Plant Count**	**Capacity (MW)**	**% in Total Capacity**
Wind	218	39,203	27.29
Solar	189	24,701	17.19
Gas Combined Cycle	68	35,132	24.45
Gas Steam Turbine	46	11,970	8.33
Gas Open Cycle	83	6,438	4.48
Gas Combustion Engine	10	671	0.47
Lignite	12	7,142	4.97
Coal	20	12,637	8.80
Hydro	5	555	0.39
Nuclear	4	4,981	3.47
Biogas	4	93	0.06
Biomass	2	150	0.10

Table 3.3: Overview of plant types, their counts, nameplate capacities, and shares in total capacity

To determine generation quantities and profits for power plants, we employ the merit order dispatch model proposed by Mann et al. (2017). We compiled a dataset of 655 power plants located in Texas. Table 3.3 shows the number of power plants and their nameplate capacity by generation technology. Except for wind and solar plants, plant-level data on capacity, commissioning dates, heat rates, and variable costs are borrowed from Mann et al. (2017). For solar and wind farms, plant-level data on capacity and commissioning dates are obtained from ERCOT (2023) and are supplemented by web searches for some individual plants. Marginal costs for power plant GEN are defined as (Mann et al., 2017).

$$c_t^{GEN} = FP_t^{GEN} * \text{Marginal heat rate}^{GEN} + \text{Variable O\&M Cost}^{GEN}$$

c_t^{GEN} is zero for solar, wind, hydro, and biomass plants (Mann et al., 2017). For all other technologies, marginal costs fluctuate with their fuel price FP_t^{GEN}. Fuel prices for biomass and uranium are taken from Mann et al. (2017) and are held constant over time. Prices for coal and lignite in Texas are assumed to change on a yearly basis and are provided by EIA (2018) and EIA (2023). We allow the gas price to change daily following the Henry Hub wholesale market gas price EIA (2024). Figure 3.1a in section 3.4.1 plots the daily gas price during our simulation period. We ignore fixed costs for all generation technologies. In every period, each plant has a forced outage with a probability of 15%

For the merit order model, we first focus on dispatchable power plants (all plants except wind and solar). Let $N_t = 1, 2, ..., n$ be the list of dispatchable power plants that do not have an outage in period t. Plants in N_t are sorted by their marginal costs such that $c_t^1 \leq c_t^2 ... \leq c_t^n$. Each plant's marginal cost and its position in the merit order list N_t varies over time according to its time-varying fuel price FP_t^{GEN}. In every period t, the goal of the merit order dispatch model is to choose each plant's generation q_t^i such that the aggregate costs of electricity generation are minimized.

$$\min_{\{q_t^i\}} \sum_{i=1}^{n} c_t^i q_t^i$$

subject to the constraints that the aggregate generation of dispatchable plants equals aggregate net load and that each plant respects its capacity limits.

$$\sum_{i=1}^{n} q_t^i = L_t$$
$$0 \leq q_t^i \leq q_{\max}^i, \quad \forall i \in N$$

To meet these constraints and achieve the cost-minimizing dispatch, we start with the first plant in the merit order and dispatch each plant at its maximum capacity until the cumulative dispatched capacity meets or exceeds the net load.

$$q_t^i = \begin{cases} q_{\max}^i, & \text{if } \sum_{j=1}^{i} q_{\max}^j \leq L_t \\ L_t - \sum_{j=1}^{i-1} q_{\max}^j, & \text{if } \sum_{j=1}^{i-1} q_{\max}^j < L_t \leq \sum_{j=1}^{i} q_{\max}^j \\ 0, & \text{if } \sum_{j=1}^{i-1} q_{\max}^j \geq L_t \end{cases}$$

This merit order dispatch model aims to derive reasonable generation quantities for a diverse set of power plants. Notably, the model ignores many operational constraints (e.g., ramp-up times and minimum generation constraints) and non-convexities in the plant's cost functions (e.g., ramp-up costs).

3.5 Forward and option contracts

This section defines the option and forward contracts the generator and LSE can buy. In electricity markets, forward and option contracts have delivery periods, typically a year, quarter, or month. The contract specifies a quantity of electricity, typically $(e.g., 1MWh)$, that has to be sold or bought every hour during the delivery period (Coulon et al., 2013).

Following Cramton et al. (2025b) and Coulon et al. (2013), we derive hedging strategies for options and forwards that cover much more granular delivery periods. The granular periods allow market participants to tailor their forward and option positions to their seasonal and daily load patterns. Thereby, market participants can manage quantity risks more effectively. Moreover, incentives to bid competitively in the day-ahead market are stronger if forward and option quantities are linked to the expected load on an hourly level.

Therefore, we define a delivery period of M as a unique combination of year-month-hour-weekend/weekdays. For instance, market participants can purchase options and forwards covering all hourly periods that fall within $M =$"weekdays 3-4 pm in July 2019". In this paper, we simulate the optimal hedging strategies for all combinations of year-month-hour-weekend/weekdays in 2019. We initialize the simulation on $t =$ December 1, 2018, when the generator and the LSE take all hedging decisions. We assume $\overline{L}_t = \overline{X}_t = 0$, $log(G_t) = m_G$ at decision time t (Coulon et al., 2013). Moreover, we set a risk-free interest rate of 2% and a high strike price for the call option of \$200/MWh. We set such a high strike price since the option should only cover the tail of the day-ahead price distributions (Cramton et al., 2025b).

Since forward contracts do not involve upfront payments, we can derive the arbitrage-free forward price $F(t, M)$ for all hours k that fall in delivery period M using the no-arbitrage relationship under a risk-neutral measure $\mathbb{Q}$ (Coulon et al., 2013).

$$F(t, M) = \mathbb{E}_t^{\mathbb{Q}}[P_{k \in M}]$$

Similarly, the arbitrage-free call option price for delivery period M with $k \in M$ is

$$V(t, M) = \mathbb{E}_t^{\mathbb{Q}}\left[e^{-r(k-t)}(P_k - K)^+\right]$$

Coulon et al. (2013) derive closed-form solutions for the arbitrage-free prices using the parameters estimated in the regime-switching model. We calculate the forward and option prices numerically based on the distribution of the day-ahead prices that we simulated using the regime-switching model. We have to calculate the prices numerically for two reasons: First, we cap the day-ahead prices at $\$9,000$/MWh to make the simulated day-ahead prices consistent with observed prices in ERCOT. Second, we numerically calculate the optimal quantities of forwards and options using a distribution of simulated day-ahead prices. Since the day-ahead price distribution has long tails, we would need to draw extremely large samples to obtain day-ahead price distributions that are consistent and arbitrage-free with the closed-form option and forward prices implied by the regime-switching model. Drawing these large samples is computationally infeasible. Our numerical forward and option prices converge to the closed-form prices implied by the regime-switching models for large samples if we remove the $\$9,000$/MWh price cap.

Numerically, the arbitrage-free forward price in the delivery period M can simply be calculated as the mean of the simulated day-ahead prices in M. Figure 3.4 illustrates the forward prices for each combination of year- month-hour-weekend/weekdays in 2019.[4] Forward prices are below $\$50$/MWh for most hours. In a few afternoon peak hours during summer, forward prices reach more than $\$100$/MWh since day-ahead prices are high on average and face large spikes during these hours.

Forward price ($/MWh)
17 304

	Year / Month 2019											
Hour	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
0	21	21	22	23	24	26	27	27	25	24	22	22
1	19	19	19	19	21	23	24	24	22	20	20	20
2	19	18	17	17	18	21	23	23	21	19	18	19
3	20	19	18	17	18	20	22	22	20	19	19	20
4	22	22	20	19	19	20	22	22	21	20	20	22
5	28	28	26	23	22	22	23	23	23	23	25	27
6	52	52	43	33	27	25	25	26	28	31	37	46
7	50	50	43	34	28	26	26	27	28	31	36	44
8	42	42	38	33	29	28	28	28	29	30	34	39
9	38	39	37	33	31	30	30	31	31	31	33	35
10	34	36	35	34	33	34	35	35	34	32	31	33
11	31	33	35	35	37	40	43	43	40	34	31	30
12	29	31	34	37	42	48	54	54	47	37	31	28
13	27	31	36	41	50	64	79	82	66	44	32	27
14	26	30	37	46	63	93	134	142	99	54	32	26
15	25	31	41	57	89	154	247	268	168	71	35	25
16	26	33	45	67	111	189	286	304	194	83	39	27
17	47	44	45	53	74	105	124	114	90	70	59	53
18	58	62	59	52	51	56	66	70	63	53	49	52
19	40	49	56	55	50	47	50	55	55	48	40	37
20	35	42	50	53	50	47	46	46	45	39	34	32
21	31	34	36	37	38	39	40	39	36	33	31	30
22	27	29	30	31	32	33	34	33	31	29	28	27
23	24	25	25	26	27	29	30	29	28	26	25	24

Figure 3.4: Forward electricity prices for year-month-hour products for weekdays in 2019

[4]The forward prices are calculated from the perspective of $t = 12/01/2018$.

The arbitrage-free price for the call option with strike price K =\$200/MWh for all hours $k \in M$ is given as

$$
\begin{aligned}
V(t, M) &= \mathbb{E}_t^{\mathbb{Q}}\left[e^{-r(k-t)}(P_k - K)^+\right] \\
&= \mathbb{E}_t^{\mathbb{Q}}\left[e^{-r(k-t)} \underbrace{Prob(P_k > K)}_{\text{Spike frequency}} * \underbrace{E[P_k - K | P_k > K]}_{\text{Expected spike size}}\right]
\end{aligned}
\tag{3.11}
$$

Figure 3.5 shows the option prices for the combinations of year-month-hour-weekend products. Option prices differ widely between products because the frequency and expected size of day-ahead price spikes differ by time of day and season. For instance, in a peak delivery period like Weekdays 4-5 pm in August 2019, day-ahead prices exceed the \$200/MWh strike price in 19% of simulated hourly periods and reach a maximum of \$9, 000/MWh. In off-peak hours on Weekends, 1-2 am in July, the day-ahead price exceeds the day-ahead price in only 0.3% of the simulated hourly periods and reaches a far smaller maximum of \$426/MWh.

Option price (\$/MWh)

0,0 205,7

Hour	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
	Year / Month 2019											
0	0.1	0.1	0.1	0.2	0.2	0.3	0.4	0.3	0.3	0.2	0.1	0.1
1	0.0	0.0	0.0	0.0	0.1	0.1	0.2	0.2	0.1	0.1	0.0	0.0
2	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.1	0.1	0.0	0.0	0.0
3	0.1	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.0	0.0	0.0	0.1
4	0.1	0.1	0.1	0.0	0.0	0.1	0.1	0.1	0.1	0.1	0.1	0.1
5	0.5	0.5	0.3	0.2	0.1	0.1	0.1	0.1	0.1	0.2	0.3	0.4
6	7.6	7.5	4.1	1.5	0.5	0.2	0.2	0.3	0.4	0.9	2.2	4.9
7	6.6	6.6	3.8	1.5	0.6	0.3	0.3	0.4	0.5	0.8	1.9	4.3
8	3.5	3.5	2.4	1.2	0.6	0.5	0.4	0.5	0.6	0.8	1.4	2.4
9	2.3	2.6	2.1	1.3	0.8	0.7	0.6	0.7	0.8	0.9	1.1	1.7
10	1.5	1.8	1.7	1.4	1.2	1.2	1.3	1.4	1.2	1.0	1.0	1.2
11	0.9	1.3	1.6	1.8	2.0	2.5	3.1	3.3	2.5	1.5	0.8	0.7
12	0.6	1.0	1.5	2.3	3.4	5.1	7.2	7.3	4.9	2.1	0.9	0.5
13	0.4	0.9	1.8	3.3	6.4	12.5	21.1	22.9	13.7	4.2	1.0	0.4
14	0.3	0.8	2.2	5.1	12.7	31.1	60.7	67.4	35.5	8.3	1.1	0.3
15	0.2	0.9	3.2	9.8	29.7	78.3	155.9	174.0	90.1	18.6	1.7	0.3
16	0.3	1.2	4.8	15.8	45.3	107.5	190.2	205.7	111.8	26.0	2.8	0.4
17	5.2	4.2	4.5	8.0	19.3	39.1	52.5	45.5	29.0	16.8	10.9	7.7
18	10.3	12.7	10.8	7.6	6.6	8.4	12.8	15.0	11.9	7.6	6.0	7.3
19	3.0	6.0	9.2	8.8	6.1	4.7	5.5	7.7	8.2	5.4	2.9	2.2
20	1.5	3.6	6.5	7.8	6.3	4.6	4.0	4.2	3.9	2.6	1.4	1.1
21	0.9	1.3	1.8	2.1	2.2	2.2	2.2	2.1	1.7	1.2	0.9	0.8
22	0.5	0.6	0.8	1.0	1.0	1.1	1.1	1.0	0.9	0.6	0.5	0.5
23	0.2	0.2	0.3	0.4	0.5	0.6	0.6	0.5	0.4	0.3	0.2	0.2

Figure 3.5: Option prices for year-month-hour products for weekdays for a spike price of \$200/MWh in 2019

The primary channel through which day-ahead price spikes affect the optimal quantity of forwards and options is that price spikes change the relative prices of forwards and options. The three graphs in Figure 3.6 show the frequency of day-ahead price spikes ($Prob(P_k > K)$) on the x-axes and the expected spike size ($E[P_k - K|P_k > K]$) on the y-axes for all 576 year-month-hour-weekend/weekday delivery periods in 2019. The graphs reveal a strong positive correlation between the frequency and the average size of the spikes. The color shades highlight the arbitrage-free forward prices (Graph 3.6a) and option prices (Graph 3.6b) increase in spike frequency and size.

However, option prices rise much stronger in spike frequency and size than forward prices. As Equation 3.11 points out, option prices only depend on spike prices. In contrast, forward prices take the expectation across all day-ahead prices such that price spikes have a weaker effect on the forward price than on the option price. Graph 3.6c highlights that the relative price ratio of option-to-forward prices increases in spike frequency and size. Therefore, the relative price ratio can be interpreted as a measure of the "spikiness" of the day-ahead price distribution because it is strongly correlated with spike frequency and size. In the following, we analyze how the spikiness of a delivery period, as measured by the option-to-forward price ratio, affects the demand for options and forwards.

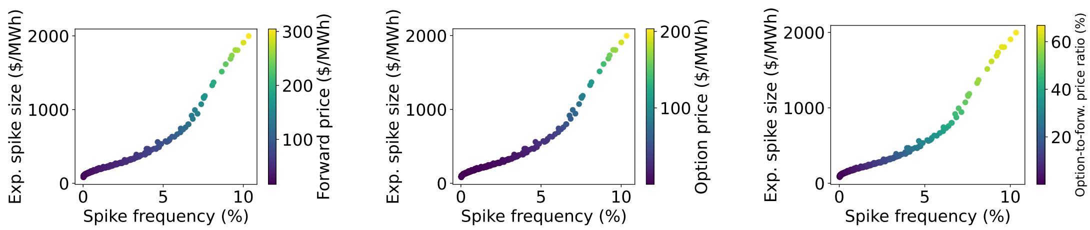

(a) Forwards prices

(b) Option prices

(c) Option-to-forwards price ratio

Figure 3.6: Forwards prices, option prices, and option-to-forwards price ratio for granular delivery periods by spike frequency and expected spike size

3.6 Optimal hedging strategies

In this section, we numerically simulate static optimal hedging strategies for a representative generator and a representative LSE. The generator mimics a large incumbent with 147 power plants and a cumulative operating capacity of 34 GW. The plants were selected from our sample of power plants based on Mann et al. (2017) and ERCOT (2023) in order to form a representative and diverse generation portfolio. Figure 3.7 shows the share of capacity by technology, highlighting that the generator is technologically diversified and owns plants with varying marginal costs. For each of the generator's power plants, we simulate hourly dispatch and profits using the merit order model in section 3.4.3.

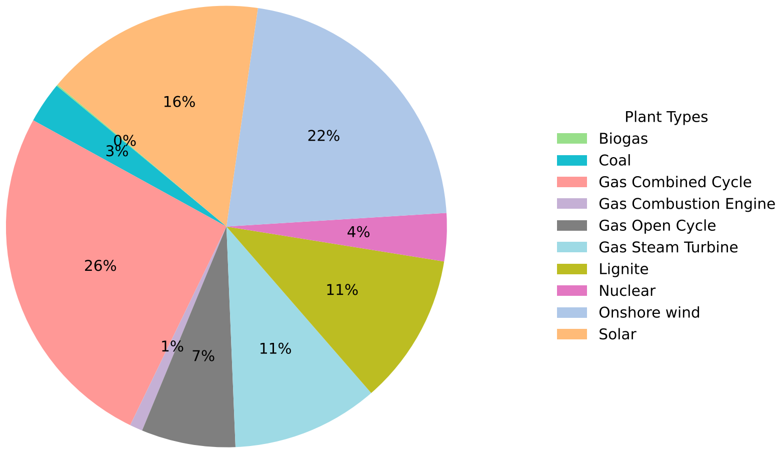

Figure 3.7: Share of technologies in the plant portfolio of a representative generator with 34 GW installed capacity

The LSE is assumed to serve on average 5% of the load in ERCOT. Its hourly load is imperfectly correlated with aggregate load (see section 3.4.3). The retail price equals the average annual day-ahead price plus a 10% retail margin (Peura & Bunn, 2021).

To isolate the distinct roles of forwards and options, we start deriving the optimal quantity of forwards only, assuming that options are not available. In section 3.6.2, we present hedging strategies that combine forwards and options.

3.6.1 Hedging strategies with forwards only

We define a quantity of forward contracts as optimal if it maximizes the agent's expected utility for a given constant absolute risk aversion (CARA) coefficient. This optimal quantity is expressed as a forward hedge ratio, i.e., the optimal quantity of forward contracts as a percentage of the agent's expected load during the contract's delivery period. There are 576 unique granular year-month-hour-weekend/weekday delivery periods in 2019.

Figure 3.8a shows the optimal forward ratio for an LSE for all 576 delivery periods in 2019 as a function of the option-to-forward price ratio in each delivery period. The larger the option-to-forward price ratio, the larger the frequency and size of day-ahead

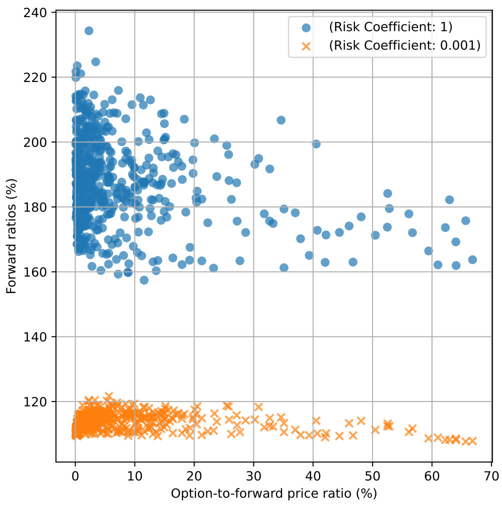

(a) LSE

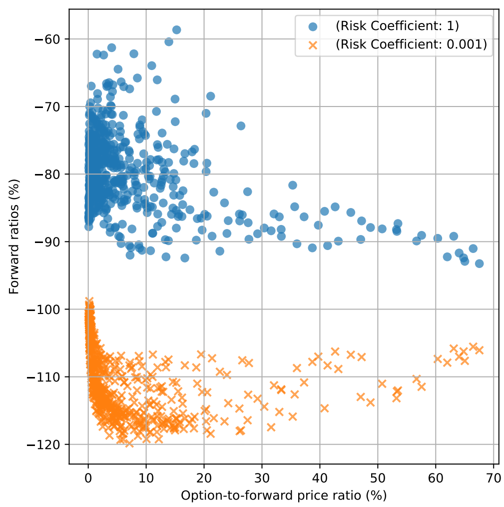

(b) Generator

Figure 3.8: LSE's and generator's optimal forward ratio by option-to-forward price ratio for delivery period weekdays 4-5 pm in August 2019

price spikes in the delivery period. We show the hedge ratios for two levels of relative risk aversion: 0.001 and 1. A relative risk coefficient of 0.001 means the LSE is almost risk neutral, while 1 implies strong risk aversion. Despite using a CARA utility function, we relate hedge ratios to relative risk aversion to make risk preferences comparable between agents with different profit scales.[5]

Figure 3.8a reveals that the frequency and size of price tails have only a tiny impact on the optimal forward ratios for the LSE when only forwards are available. When the LSE is almost risk-neutral with a risk coefficient of 0.01, the forward ratios are not sensitive to price spikes and vary between 108-122%. For a higher risk aversion of 1, the

[5]Raskin and Cochran (1986) demonstrate that absolute risk aversion coefficients are hardly comparable between agents whose profits or wealth differs in scale. Assume agent A's profits fluctuate around an average of 100$ and agent B's profits fluctuate around 1000,000$. The CARA coefficient measures how the agents' marginal utility of profits changes when profits increase by an absolute amount, say 1$. Intuitively, a 1$ increase in profits should change A's marginal utility of profits much more than B's because the 1$ causes a much larger relative profit increase for A. To make risk preferences comparable between A and B, we select a risk preference expressed as a coefficient of relative risk aversion. When calculating hedge ratios, we first assume a relative risk parameter. Then, we divide this relative risk parameter by an agent's average profit during the hedge contract's delivery period. This transforms the relative risk parameter into an absolute risk parameter. The absolute and relative coefficients describe the same preferences, at least locally at the agent's average profit. We assume this absolute risk parameter to be constant across profit levels (CARA). Hence, the CARA parameter represents the same preferences as the relative risk aversion coefficient locally at the agent's average profit. Expressing risk preferences in terms of relative risk in a CARA framework is common in the optimal hedge literature (Newbery, 1989, Lapan & Moschini, 1994).

LSE's forward ratios slightly decline when day-ahead price tails become more severe. However, risk preferences are far more relevant for the LSE's forward ratios than price spikes. When the LSE dislikes risk with a risk coefficient 1, it strongly overhedges with hedge ratios between 160-230%.

The generator's short forward position is also not sensitive to the frequency and size of price spikes for low risk aversion, as Figure 3.8b shows. The nearly risk-neutral generator chooses short forward positions of 98-120% of expected generation. For high risk aversion of 1, the generator chooses smaller short positions of only 58-93%. Its short forward ratio is slightly higher in periods with heavier price spikes.

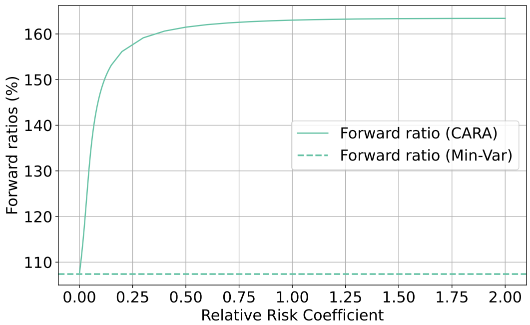

(a) LSE

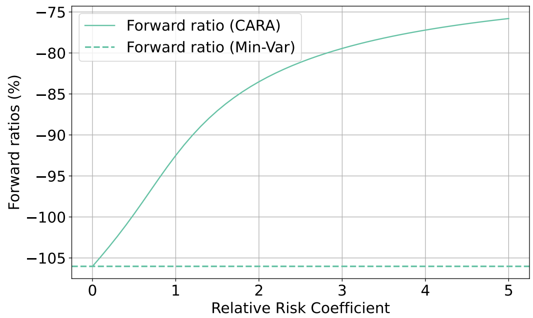

(b) Generator

Figure 3.9: LSE and generator's optimal forward ratio by relative risk aversion for delivery period weekdays 4-5 pm in August 2019

To better understand how price spikes shape the above hedge strategies, we deep dive into peak delivery period Weekdays 4-5 pm in August 2019. This peak delivery period has frequent and large price spikes, which is reflected in a high option-to-forward price ratio of 68%. Figure 3.9 shows the optimal forward ratios for this peak delivery period as a function of the agents' relative risk aversion. The forward ratios for a CARA utility maximizing LSE increase in risk aversion for low levels of relative risk aversion below 1 (Graph 3.9a). If risk aversion increases beyond 1, the hedge ratio remains constant. Irrespective of its risk aversion, the LSE wants to over the hedge, i.e., it wants to buy forward between 7%-64% more than its expected load. In contrast, the generator takes a short forward position that declines in risk aversion (Graph 3.9b). The generator's optimal hedge ratios are also more sensitive to increases in risk aversion, even for large risk coefficients.

As a comparison, the dashed line illustrates the optimal forward ratio if the agents want to minimize their profits' variance (Min-Var). For low-risk aversion, the CARA-optimal forward ratio equals the Min-Var forward ratio. The Min-Var ratio is a 106% short position for the generator and a 107% long position for the LSE. Hence, the generator and the LSE want to hedge roughly the opposite quantities if they are almost risk-neutral. The Min-Var hedge ratio requires agents to slightly overhedge, which indicates that their load is positively correlated with day-ahead prices (McKinnon, 1967, Borenstein, 2007b).

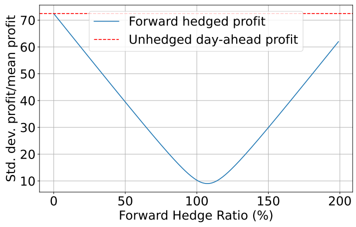

(a) LSE's standard deviation of profits

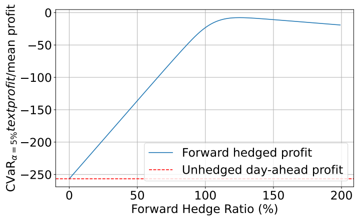

(b) LSE's $CVaR_{\alpha=5\%}$ profit

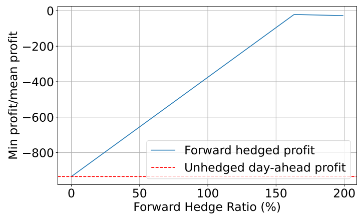

(c) LSE's minimum profit

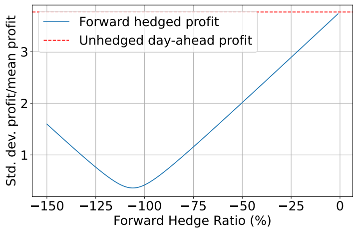

(d) Generator's standard deviation of profits

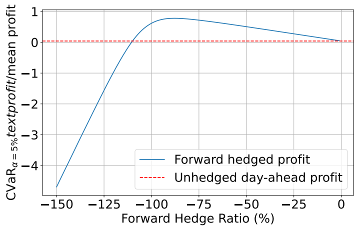

(e) Generator's $CVaR_{\alpha=5\%}$ profit

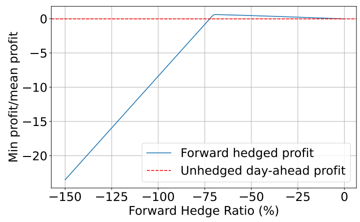

(f) Generator's minimum profit

Figure 3.10: LSE's and generator's profit standard deviation, $CVaR_{\alpha=5\%}$ profit outcome, and minimum profit relative to mean unhedged day-ahead (DA) profit by forward ratio for delivery period Weekdays 4-5 pm in August 2019

To rationalize the above strategies, Figure 3.10 examines how the forward hedge ratio affects the LSE's and generator's profit volatility, their conditional value at risk at the 5% level ($CVaR_{\alpha=5\%}$)[6], and worst-case minimum profit relative to their respective mean unhedged day-ahead profit. The average profit is not affected by the quantity of forwards since the forward price is arbitrage-free (Coulon et al., 2013).

Figure 3.10 highlights that agents' hedge strategies are driven by a trade-off between minimizing profit variance and maximizing worst-case tail outcomes. Both agents choose

[6] $CVaR_{\alpha=5\%}$ is the average of the lowest 5% of the profits in the profit distribution. Suppose profit π takes random values $\{\pi_1, \pi_2, \ldots, \pi_N\}$ with probabilities $\{p_1, p_2, \ldots, p_N\}$. Reorder these values such that $\pi_{(1)} \leq \pi_{(2)} \leq \ldots \leq \pi_{(N)}$ with corresponding probabilities $p_{(1)}, p_{(2)}, \ldots, p_{(N)}$. Let the partial sums be $P_{(k)} = \sum_{i=1}^{k} p_{(i)}$, for $k = 1, \ldots, N$. Define $\alpha = 0.05$ and find the smallest k such that $P_{(k)} \geq \alpha$. Then $\mathrm{CVaR}_{\alpha=5\%} = \frac{1}{\alpha}\left[\sum_{i=1}^{k-1} \pi_{(i)} p_{(i)} + \pi_{(k)} (\alpha - P_{(k-1)})\right]$ (Shapiro et al., 2021).

the Min-Var hedge ratio that minimizes profit variance if they are nearly risk-neutral. However, the generator's Min-Var ratio is a 106% short forward position, which leads to substantial negative tail losses as measured by a negative $CVaR_{\alpha=5\%}$ profit and negative minimum profit (Graphs 3.10e and 3.10f). A short position causes large losses when price spikes occur. The more the generator dislikes risk, the more it reduces the short position to avoid tail losses at the expense of an increase in profit variance (Graph 3.10a).

The risk-averse LSE also aims to reduce its tail losses by choosing a long forward position that exceeds its Min-Var position. The LSE needs to buy more than its expected quantity forward to be well protected against price spikes because spikes often occur when the load is above average. For high-risk aversion greater 1, the LSE chooses a hedge ratio of 164% because this ratio maximizes its worst-case minimum profit (Figure 3.10c). However, the worst-case minimizing ratio of 164% leads to high tail losses $CVaR_{\alpha=5\%}$ and high-profit volatility. Yet, the profit volatility for a hedge ratio of 164% is still much lower than the profit volatility for an unhedged 0% hedge ratio, as the dashed red line in Graph 3.10a reveals.

Noticeably, the LSE's $CVaR_{\alpha=5\%}$ tail loss is already minimized at a forward ratio of 125%. The reason for this large difference between the $CVaR_{\alpha=5\%}$ minimizing ratio and the worst-case minimizing ratio of 164% is that there are few extreme negative outliers in the LSE's day-ahead profit distribution. These outliers are so extreme that the LSE starts focusing primarily on these worst-case outliers, even for moderate risk aversion. In contrast, the generator does not exclusively focus on the minimum profit, even for high risk aversion. Its day-ahead profit distributions do not exhibit such large negative tails because the generator has more flexibility to shut down its plants when prices are unfavorable.

In Appendix C.4 we present forward hedging strategies for an off-peak period with small and rare price spikes: Weekends, 4 am in May 2019. Forward hedging strategies in the off-peak period look very similar to the peak period. When the LSE is very risk-averse, it chooses larger forward ratios in off-peak periods than in peak periods. The generator's hedge ratios are almost the same in peak and off-peak periods. Price spikes only have a small effect on hedge ratios when hedging with forwards only.

Overall, hedging with forwards only is surprisingly effective in reducing volatility and worst-case outcomes for a delivery period with extreme price spikes like Weekdays 4-5 pm in August 2019. The reason is that the generator's and LSE's unhedged day-ahead market profit in the peak period is roughly a linear function of the day-ahead price as

Figure 3.11 highlights.[7] One might have expected a strongly nonlinear relation between day-ahead prices and unhedged profits during a peak period since spike prices are often correlated with high load. However, Figure C.3 in Appendix C.3 reveals only a weak positive correlation between simulated load and prices for the peak delivery period Weekdays 4-5 pm in August 2019. In particular, very high spike prices are only slightly positively correlated with load. Spike prices likely depend much more on net load and random factors like outages than on load in modern electricity markets.

It is widely assumed that adding nonlinear options to a linear forward hedge strategy creates little benefits if profits are linear in prices and if forward and option prices are arbitrage-free (Lapan et al., 1991). In the following, we analyze whether options can still play a role in the generator's and LSE's hedging strategy in the presence of large price spikes.

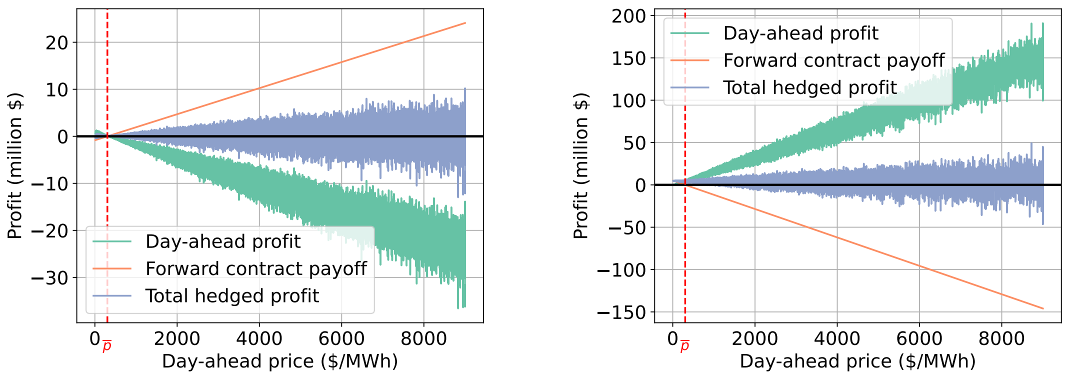

(a) LSE (b) Generator

Figure 3.11: LSE's and generator's day-ahead market profit, forward contract payoffs, and total hedged profit as a function of the day-ahead price for delivery period weekdays 4-5 pm in August 2019

3.6.2 Hedging with forwards and options

When combining forwards and options, hedging strategies become far more sensitive to severe price tails. Figures 3.12a and 3.12b present the optimal forward and option ratios for the LSE's combined hedging strategy for the 576 delivery periods in 2019 as a function of the delivery period's options-to-forward price ratio. A delivery period with a

[7]The forward contract payoff shown in Figure 3.11 is the payoff of a Min-Var hedge strategy chosen by a nearly risk-neutral agent. The Min-Var hedge achieves the lowest profit variance by choosing a forward position that is roughly the opposite of the day-ahead position.

high options-to-forward price ratio has large and frequent price spikes. The LSE's forward ratios decline when a delivery period has more severe price spikes. The more risk-averse the LSE, the more sensitive its forward ratio to price tails.

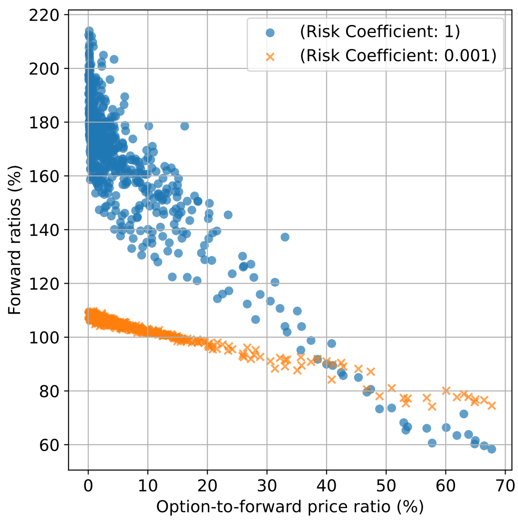

(a) LSE's forward ratios

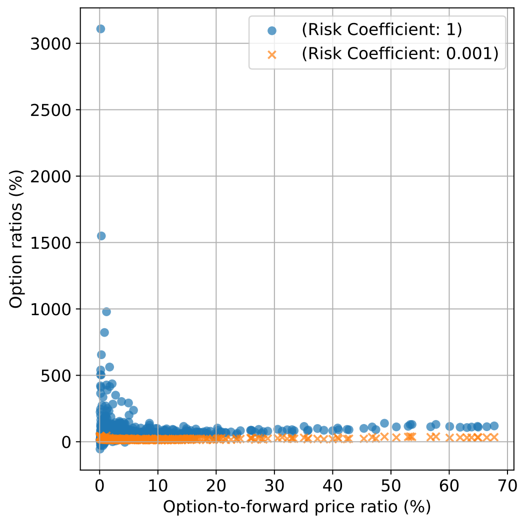

(b) LSE's option ratios

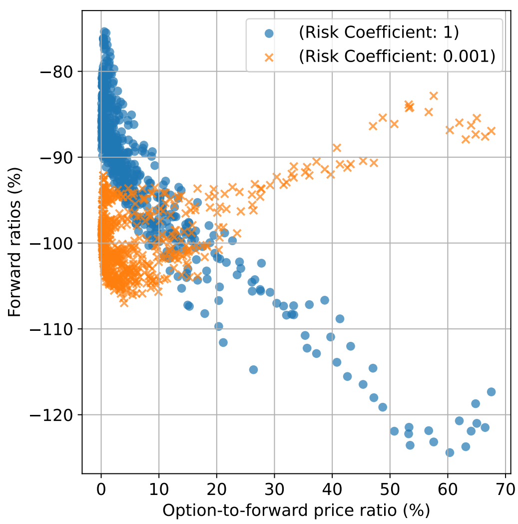

(c) Generator's forward ratios

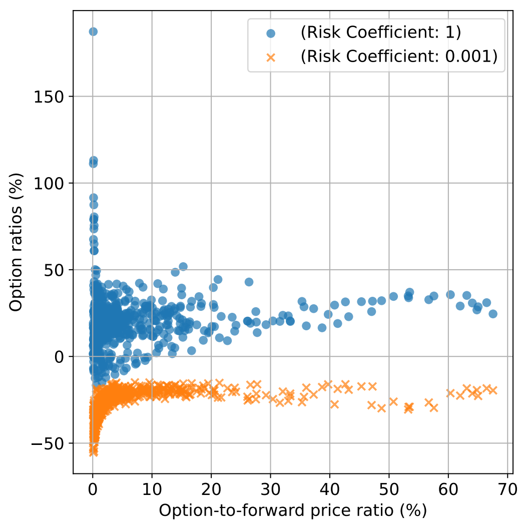

(d) Generator's option ratios

Figure 3.12: LSE's and generator's optimal forward and option ratios by option-to-forward price ratio for delivery period weekdays 4-5 pm in August 2019

Figure 3.12b highlights that the option ratios for the almost risk-neutral LSE only slightly increase from 15-20% to 30-40% when day-ahead price distributions become spikier (relative risk coefficient 0.001). In the high risk aversion case (risk coefficient of 1), the LSE selects much higher option ratios of 74-109% in spiky periods. In periods with rare and small price spikes, the LSE takes smaller option ratios for most periods. However, there are a few outlier periods in which the LSE selects large long positions or small short positions. In these periods, price spikes are extremely rare, and the option price is almost zero. A single extreme price spike can dominate the tail of the price distribution in such a period. Therefore, the option ratio in these low-spike periods is sensitive to outliers in the simulated tail of the profit distribution and is not stable across repeated draws of profit distributions.

Overall, the risk-averse LSE chooses a higher option than forward ratios in periods with large price spikes. In contrast, the almost risk-neutral LSE relies more on forwards than options when price spikes are severe.

The generator's short forward position declines in the option-to-forward price ratio for small risk aversion (0.001), as Figure 3.12c illustrates. For large risk aversion (1), the short forward position increases in price spikiness. Analogously to the LSE, the almost risk-neutral generator takes only small short positions in options, mainly around 20% with a few larger outliers (Figure 3.12d). The short option position remains small if the day-ahead price distribution becomes spikier. The risk-averse generator (1) chooses almost exclusively long option positions, mostly around 20-40%. These long positions slightly increase when price spikes are large and frequent.

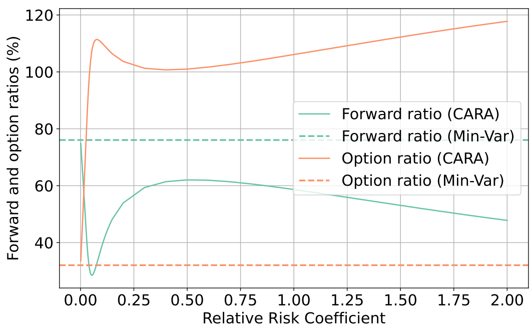

(a) LSE

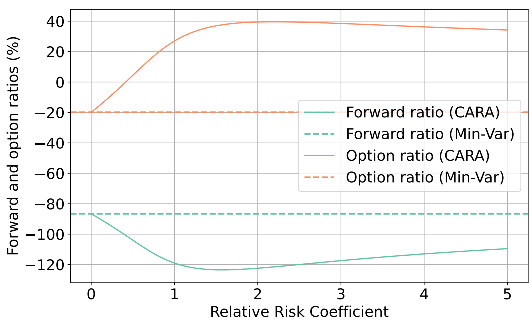

(b) Generator

Figure 3.13: LSE's and generator's optimal forward and option ratio by relative risk aversion for delivery period weekdays 4-5 pm in August 2019

We deep dive again into peak delivery period weekdays 4-5 pm in August 2019. Figure 3.13a shows the LSE's optimal forward and option ratios across risk-aversion levels. For this peak period, the importance of options in the LSE's hedging strategy increases relative to forwards when the LSE becomes more risk-averse. As mentioned above, when the CARA minimizing LSE is almost risk-neutral, it buys the number of forwards and options that minimize the variance of profits (Min-Var). The Min-Var hedge position buys options and forwards roughly in the ratio 2:1. Buying 2 call options and 1 forward contract replicates the payoff of a long straddle. A long straddle position is known to effectively reduce profit variance (Lapan et al., 1991).

The optimal option ratio rapidly increases in the LSE's risk aversion. The forward ratio also sharply declines for low-risk aversion. However, the forward ratios increase for intermediate risk aversion and decline afterward. The LSE's forward ratios are less sensitive to risk preferences than option ratios. The LSE starts buying more options than forwards in this peak delivery period, even for small levels of risk aversion.

For the generator, Figure 3.13b highlights that the almost risk-neutral generator takes a minimum variance (Min-Var) strategy that is similar to the LSE's strategy but with opposite signs since the generator goes short in forwards and options. The generator slightly increases its short position in forwards for intermediate risk levels. At the same time, the short option position decreases more strongly. The generator's option position transitions to a long position for larger risk aversion.

The pink line in Figure 3.14 depicts the agents' standard deviation of profits, $CVaR_{\alpha=5\%}$ tail loss, and worst-case minimum loss as a function of the optimal forward and option ratios that are associated with a relative risk aversion level. For instance, at a risk coefficient of 0.001, the LSE chooses a 76% forwards ratio and 32% options ratio because this Min-Var portfolio leads to the lowest standard deviation of its profits (Graph 3.14a). The nearly risk-averse generator also chooses the Min-Var hedge portfolio (Graph 3.14d).

For small risk aversion, both agents focus more on improving the $CVaR_{\alpha=5\%}$ lowest profit outcomes, which comes at the expense of larger profit variance. The LSE can reduce the $CVaR_{\alpha=5\%}$ tail loss by buying more options and fewer forwards. Going long in options protects more effectively from price spikes than forwards. The generator protects from price spikes by gradually turning its short option position into a long one. Going long in options has the disadvantage of increasing the generator's profit variance because the long option has a high payoff when the generator's day-ahead profits are high anyway. The generator balances the long option position with a larger short forward position to limit the increase in variance.

The more risk-averse the generator and LSE become, the more they focus on ever more extreme parts of the lower end of their profit distribution. The highly risk-averse

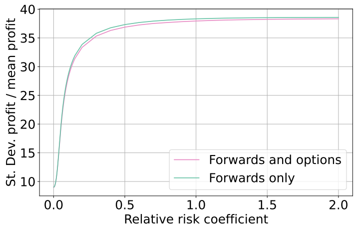

(a) LSE's standard deviation of profits

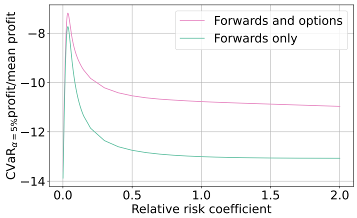

(b) LSE's $CVaR_{\alpha=5\%}$ profit

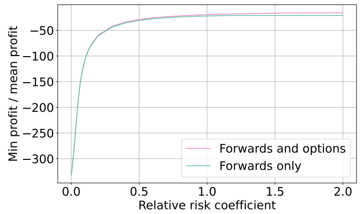

(c) LSE's minimum profit

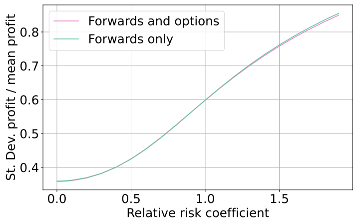

(d) Generator's standard deviation of profits

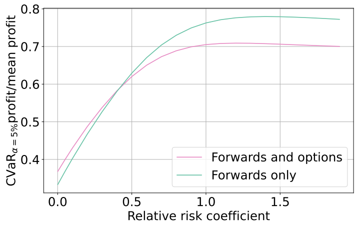

(e) Generator's $CVaR_{\alpha=5\%}$ profit

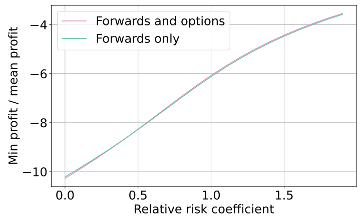

(f) Generator's minimum profit

Figure 3.14: LSE's and generator's profit standard deviation, $CVaR_{\alpha=5\%}$, and minimum profit relative to mean unhedged day-ahead (DA) profit by relative risk aversion level for delivery period Weekdays 4–5 pm in August 2019

LSE buys even more options to reduce the minimum worst-case loss at the expense of a larger $CVaR_{\alpha=5\%}$ tail loss. The generator also reduces its long option and short forward positions to achieve a marginal increase in its lowest worst-case profit.

Figure 3.14 also reveals that the combined forwards and options strategy (pink line) add only limited value compared to the optimal forwards-only strategy (green line). Compared to the forwards-only strategy, the combined strategy achieves only tiny improvements in profit standard deviation and minimum loss. The improvements in terms of $CVaR_{\alpha=5\%}$ tail outcomes are more substantial, but these improvements seem negligible compared to the improvements achieved relative to the unhedged day-ahead $CVaR_{\alpha=5\%}$ outcome, as Figure C.8 in Appendix C.5 reveals. As discussed above, the forwards-only strategy achieves excellent results thanks to the linear relationship between day-ahead prices and profits. Even though options do not add large benefits, the agents still take significant option positions when they are risk-averse because the option achieves small improvements in the tail and worst-case profits. These small improvements can cause large utility increases for risk-averse agents.

Overall, Figure C.8 in Appendix C.5 shows that the combined strategy and the forwards-only strategy achieve enormous reductions in profit variance, $CVaR_{\alpha=5\%}$ tail loss, and worst-case loss compared to the case when the LSE trades fully unhedged in the day-ahead market. Compared to remaining unhedged, hedging with either forwards-only or combining forwards and options proves effective in managing tail risks and profit variance.

In Appendix C.4, we also examine the optimal combined forwards and options strategy for off-peak period Weekends 4-5 am in May 2019, which has small and rare spike prices. Compared to the peak period, the LSE relies far less on long option positions and chooses larger forward quantities in the off-peak period (see Figure C.6 in Appendix C.4). The generator also chooses only very small long option positions for moderate risk levels and small short option ratios for high risk aversion. Choosing a small long option or even a short option position allows the generator to reduce its short forward position for high risk levels in the off-peak periods. This contrasts with the above peak period, where the risk-averse generator selects long option and large short forward positions. In the off-peak period, the combined hedge strategy is more effective in reducing profit variance and the $CVaR_{\alpha=5\%}$ loss compared to the forwards-only strategy, as Figure C.7 in Appendix C.4 emphasizes. The combined strategy only marginally improves worst-case minimum outcomes for high risk levels relative to the forwards-only approach.

3.6.3 Hedging and option strike prices

In the next section, we analyze how the demand for forwards and options change if we choose a different strike price for the call option. So far, we set the strike price to \$200/MWh such that the option covers only the extreme tails of the price distribution. Choosing a lower strike price makes the option more valuable relative to the forward since the option is more frequently in the money if the strike price is low. The higher option value is reflected in a higher arbitrage-free option price.

Figure 3.15 shows the arbitrage-free option price as a function of the strike price for peak delivery period weekdays 4-5 pm in August 2019. The negative relation between option and strike price slightly diminishes in the strike price. The option price is high even for a large strike price of \$1,000/MWh. The reason is that there is a low probability of 6% that the simulated day-ahead prices exceed the high \$1,000/MWh strike price but a high expected spike size of \$2,648/MWh. In contrast, if the strike price is only \$100/MWh, the spike probability rises to 38%, but the expected spike size is only \$590/MWh. This small average spike size explains why the \$100/MWh strike price option has only a moderately higher value than the option with a \$1,000/MWh strike price. For off-peak

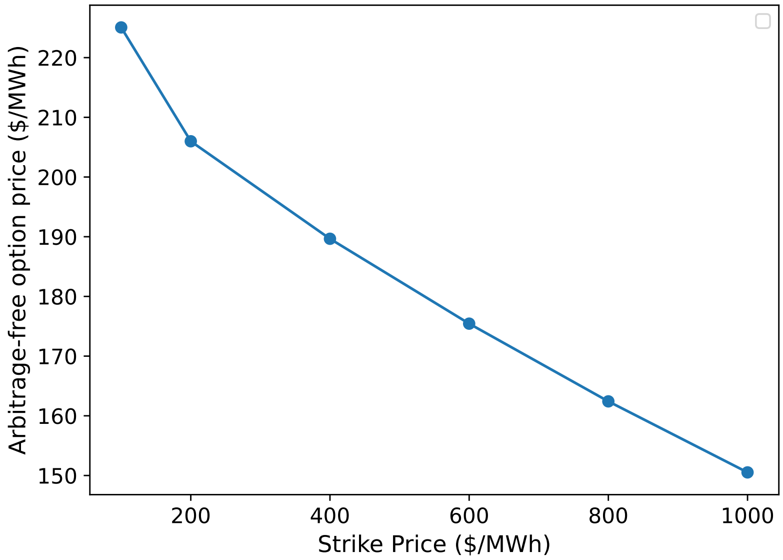

Figure 3.15: Arbitrage-free option price as a function of the call option's strike price for delivery period weekdays 4-5 pm in August 2019

delivery periods, a high $\$1,000$/MWh strike price leads to an option price of zero since the option will never be in the money for such a high strike price.

Figure 3.16 depicts the generator's and LSE's forward and option ratios as a function of the call option's strike price for different risk preferences for peak delivery period weekdays 4-5 pm in August 2019. The LSE buys more forwards and fewer options when the strike price increases. The option becomes less effective in protecting the LSE from price spikes if the strike price is high. The more risk-averse the LSE, the more sensitive is its forward and option demand to higher strike prices.

The generator also sells more forwards and fewer options when the strike price is high for low-risk aversion (0.001). In contrast, for high-risk aversion (1), the generator's short forward position and its long option position decline when the strike price is large.

Figure 3.17 shows the standard deviation of profits, $CVaR_{\alpha=5\%}$ tail outcome, and worst-case-minimum profit as a function of the agent's optimal combination of forwards and options that are linked to a relative risk aversion level for different strike prices. If the strike price is low, agents' standard deviation of profits is larger compared to higher strike prices, especially for high risk levels (Graphs 3.17a and 3.17d). This higher standard deviation is caused by the large option positions the agents hold when strike prices are

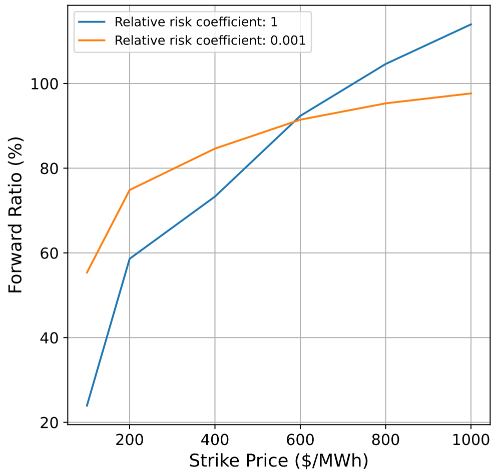

(a) LSE's forward ratios

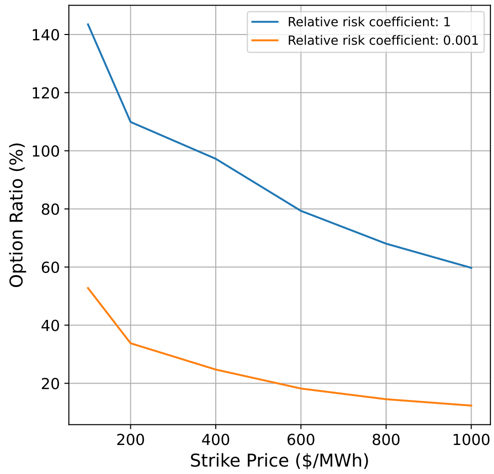

(b) LSE's option ratios

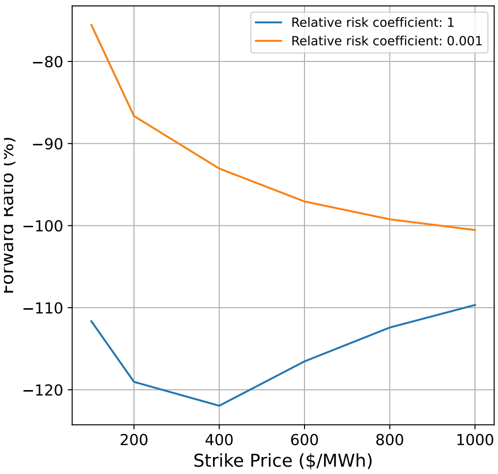

(c) Generator's forward ratios

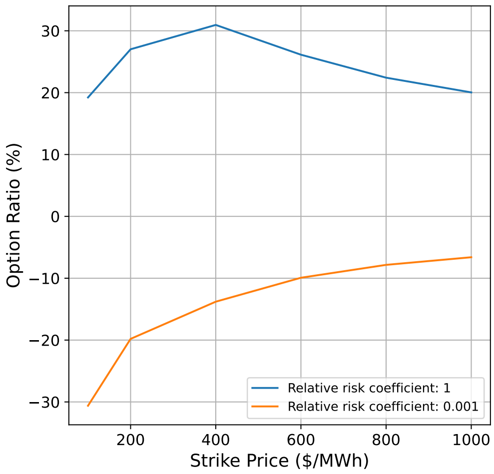

(d) Generator's option ratios

Figure 3.16: LSE's and generator's optimal forward and option ratios by option strike price for delivery period weekdays 4-5 pm in August 2019

small. However, the differences in profit standard deviation and minimum worst-case profits (Graphs 3.17c and 3.17f) among different strike prices are small.

For the LSE, small strike prices lead to small $CVaR_{\alpha=5\%}$ tail losses for small risk aversion levels and high $CVaR_{\alpha=5\%}$ losses for high risk aversion (Graph 3.17b). By contrast, small strike prices lead to smaller positive $CVaR_{\alpha=5\%}$ profits for the nearly

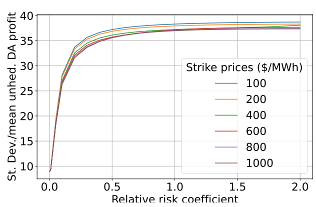

(a) LSE's standard deviation of profits

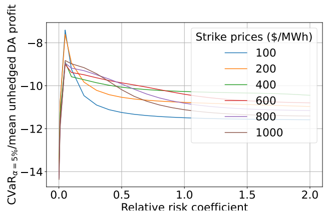

(b) LSE's $CVaR_{\alpha=5\%}$ profit

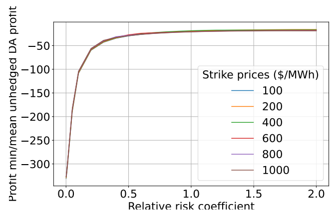

(c) LSE's minimum profit

(d) Generator's standard deviation of profits

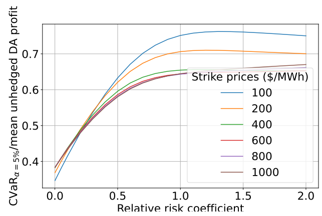

(e) Generator's $CVaR_{\alpha=5\%}$ profit

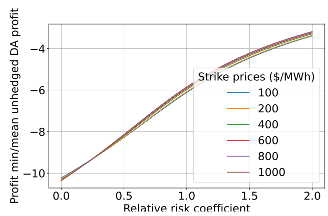

(f) Generator's minimum profit

Figure 3.17: LSE's and generator's profit standard deviation, $CVaR_{\alpha=5\%}$ and minimum profit relative to mean unhedged day-ahead (DA) profit by relative risk aversion level for delivery period Weekdays 4–5 pm in August 2019 for different strike prices

risk-neutral generator and substantially higher $CVaR_{\alpha=5\%}$ profits for the risk-averse generator 3.17f. Hence, it depends on the risk preferences of market participants whether high or low strike prices provide better protection.

3.7 Conclusion

Day-ahead electricity markets exhibit extreme price fluctuations and sudden price spikes. The price spikes are often positively correlated with load. The positive correlation between load and price spikes poses substantial financial risk for load-serving entities (LSEs) because they have to sell large quantities at extremely high prices and sell the electricity at much lower prices to end-consumers. For generators, price spikes create upside opportunities but substantially increase the variance of profits. Generators and LSEs might, therefore, want to hedge parts of their profits.

In this paper, we characterize optimal hedge strategies for a representative generator and an LSE in the ERCOT day-ahead market in Texas. We let the generator and the LSE choose the optimal portfolio of a forward contract and a call option that maximizes their constant absolute risk aversion utility function. Our research addresses the question of how price spikes, risk aversion, and option strike prices jointly shape the selection of forward and call option contracts.

To obtain realistic profit distributions, we first simulate hourly day-ahead prices and aggregate load under a regime-switching model for the Texas ERCOT market based on Coulon et al. (2013). We then apply a merit order dispatch model to simulate hourly generation quantities for a large sample of power plants. Afterward, we define 576 granular delivery periods for the option and forward contracts for each unique combination of year-month-hour-weekend/weekday in 2019. We find that the frequency and size of day-ahead price spikes differ widely between these delivery periods. This allows us to examine how different levels of spikiness in the day-ahead price distribution alter hedge choices for different delivery periods.

We compare the performance of various hedge portfolios, including forwards-only and mixed portfolios of forwards and call options, under differing levels of risk aversion and strike prices. Our results indicate that agents with low-risk aversion primarily rely on forward contracts, adding modest option positions. More risk-averse participants take larger long-option positions to protect their worst-case outcomes.

Surprisingly, the additional risk protection from combining forwards and options remains limited. Using options adds small extra value in terms of reducing variance and improving worst-case outcomes because the generator and the LSE's day-ahead market profits are roughly linear in the day-ahead price. Combined with linear forward contracts, nonlinear options achieve little additional benefits in terms of risk hedging when profits are linear (Lapan et al., 1991). Risk-averse agents still take long option positions because even small improvements in worst-case outcomes can raise utility when risk aversion is high.

Lastly, we show that a higher option strike price induces the generator and the LSE to rely more on forward contracts and less on options. When the strike price is high, the option offers limited protection from relatively high prices that are below the strike price. Therefore, agents shift their portfolio from options to forwards.

Overall, our paper highlights that forwards and option contracts can effectively reduce financial risks for a generator and an LSE in the ERCOT day-ahead market in terms of profit variance and worst-case tail losses. We emphasize that we derived our optimal hedge strategies under the assumption of arbitrage-free prices. It would be interesting to analyze how arbitrage opportunities that often occur in electricity markets affect optimal hedging strategies. In forward markets, positive price premia occur in peak demand delivery periods with high price spikes (Redl et al., 2009, Bunn, 2006). These positive risk premia might hinder LSEs from buying sufficient forwards and option contracts to effectively hedge their large tail risks during peak delivery periods.

Therefore, policymakers should implement measures to reduce price premia in forwards and option markets to facilitate hedging and risk management. Price premia are partly driven by a lack of liquidity in these markets (Bevin-McCrimmon et al., 2018). Policymakers could increase market liquidity by fostering the use of centralized markets for forwards and options with centrally defined delivery periods. They might even oblige LSEs to purchase forward a particular share of their expected demand to increase market liquidity. Price premia could also be reduced by facilitating the participation of financial institutions in forwards and options markets to increase arbitrage and improve price discovery (Cramton et al., 2025b).

The above measures might make hedging in day-ahead markets an even more effective risk management tool for electricity generators and LSEs. When market participants are well hedged, the enormous risk inherent in day-ahead electricity markets is carried on many shoulders. Hedging enables market participants to deal with extreme price spikes that are necessary as scarcity signals in short-term power markets.

Chapter 4

A Forward Energy Market to Improve Reliability and Resiliency

This chapter is co-authored with Peter Cramton, Professor Emeritus of Economics at University of Maryland and Founder of Forward Market Design LLC., Jason Dark, Chief Technology Officer at Forward Market Design LLC., Darell Hoy, Chief Executive Officer at Forward Market Design LLC., David Malec, a research associate at the University of Maryland and Chief Data Officer at Forward Market Design LLC., Axel Ockenfels, Professor of Economics at the University of Cologne and Director at the Max Planck Institute for Research on Collective Goods in Bonn, and Chris Wilkens, Chief Product Officer at Forward Market Design LLC.

This chapter contains the current version of the working paper Cramton et al. (2025b). Parts of the analysis in this paper were also published as related policy white papers under Cramton (2023) and Cramton et al. (2024b).

Abstract

We propose a novel forward electricity market design that enables efficient and granular hedging in the face of growing electrification and intermittent renewables. Market participants can trade thousands of forward contracts and European call options, each tailored to granular delivery windows spanning up to four years ahead. To implement this market, we apply Budish et al.'s (2023) flow trading framework to electricity markets. Flow trading encourages incremental trading and supports liquidity provision via frequent batch auctions. To illustrate our design, we develop a twelve-year simulation of ERCOT's day-ahead market from 2011 to 2022. We derive demand curves for forwards and options

for representative generators and load-serving entities. Our analysis reveals that risk preferences significantly shape these demand curves. Risk preferences determine the price elasticity of demand for forwards and options and influence to what extent market-clearing prices exceed arbitrage-free levels. Moreover, we find that setting a higher option strike price reduces arbitrage positions and makes options less attractive relative to forwards.

Acknowledgments

This research is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1-390838866.

4.1 Introduction

As the world transitions to net zero, nearly all sectors must be electrified, causing the electricity sector to double over the next three decades and spurring rapid innovation in both supply and demand. With an increasing share of intermittent renewables and more frequent extreme weather, balancing supply and demand every second becomes more challenging. A transparent and efficient forward energy market is needed to provide robust price signals for guiding investment and operating decisions consistent with the maximization of social welfare (Cramton, 2017).

Existing forward energy markets are inefficient since they suffer from a lack of liquidity, leading to unreliable forward price signals. This hinders generation companies and load-serving entities (LSEs) from making efficient operational and investment decisions, and complicates risk management in day-ahead and real-time markets (Newbery, 2016, Cramton, 2017). Additionally, existing forward and option contracts feature broad delivery periods that do not align with actual load and generation patterns, limiting their effectiveness as hedging instruments. Insufficient hedging leaves participants vulnerable to extreme events like the 2021 Texas winter storm or the 2001 Californian electricity crisis and can lead to bankruptcies and market failures (Borenstein, 2002, Cramton, 2022).

In this paper, we propose a novel market design for forward electricity. Participants can trade forward contracts and European call options with a high strike price, e.g., $\$1,000$/MWh. At any point in time, agents can trade thousands of these derivative products with granular delivery periods for every combination of year-month-hour-weekend/weekday up to four years ahead. For example, a monthly product could cover weekdays from 8-9 am in August 2018. Daily products for each day-hour combination are also available for the next 30 days ahead. These granular products allow market participants to choose hedge positions that closely follow their load profiles (Cramton et al., 2024a).

We employ the new flow trading technology by Budish et al. (2023) to enable market participants to trade such a large number of granular products. Flow trading makes it easy and computationally feasible to trade numerous products simultaneously. It also incentivizes gradual, small-quantity trades in hourly frequent batch auctions (Budish et al., 2015). Gradual trade guarantees a constant flow of liquidity, mitigates adverse price impact caused by large orders, and reduces opportunities to exercise market power (Budish et al., 2023, Cramton et al., 2024a).

As a proof of concept, we develop a full-scale twelve-year simulation of Texas' ERCOT market between 2011 and 2022. ERCOT provides an interesting case study because it has state-of-the-art day-ahead and real-time markets that send reliable short-term price signals with high scarcity prices up to $\$9,000$/MWh. These scarcity prices can cause enormous downside risks, especially for LSEs.

We simulate ERCOT's day-ahead market using the regime-switching model developed by Coulon et al. (2013) and Brandkamp et al. (2025). The regime-switching model simulates joint hourly distributions of day-ahead prices, net load, and gas prices. Generation profiles for multiple solar and wind farms are simulated following Brandkamp et al. (2025). We then apply Mann et al.'s (2017) merit order dispatch model to calculate hourly dispatch and profits for a large sample of 655 existing power plants that operate in ERCOT. We group these power plants into representative generation portfolios to calculate profits for generation companies. Moreover, we simulate profits for a set of representative load-serving entities (LSEs) that serve a certain share of aggregate load (Brandkamp et al., 2025).

For each representative generator and LSE, we calculate their net demand curves[1] for each granular forward and option product. We obtain agents' net demand curves by numerically deriving the optimal quantities of forwards and options for a given pair of

[1]Net demand is demand minus supply. A net demand curve gives the net demand for each product as a function of the product prices.

forward and option prices. A pair of option and forward quantities is called optimal if it maximizes the agents' expected utility from profits for given prices. We assume agents to be risk-averse with constant absolute risk aversion utility.

Our research aims to use these net demand curves and the flow trading methodology to determine unique market clearing prices and quantities in each hourly frequent batch auction between 2011 and 2022. In this framework, we intend to analyze how equilibrium prices and quantities evolve as they approach their physical delivery periods. We also explore how and why market clearing prices differ from arbitrage-free price levels. In addition, we want to study how agents' risk aversion affects market equilibria and net demand for forwards and options. Finally, we investigate how the forward energy market impacts generators' and LSEs' expected profit, profit volatility, and downside tail risks.

In the current version of our paper, we focus on an analysis of generators' and LSEs' net demand curves. Our main finding is that agents' risk preferences strongly impact the slope of their net demand curves. When generators and LSEs are highly risk-averse, net demand curves are almost vertical since agents are hardly willing to take risks to exploit arbitrage opportunities. The generator is more willing to arbitrage even for high risk levels, because the generator faces less extreme downside risks in the day-ahead market than the LSE (Brandkamp et al., 2025). However, since demand curves are almost vertical for highly risk-averse agents, forward and option prices need to be substantially above their arbitrage-free levels to allow the market to clear.

When agents are less risk-averse, their appetite for arbitrage increases. In particular, the generators' net demand curve becomes less vertical. For reasonably low risk-aversion levels, the market will, therefore, likely clear at prices that deviate less than 10% from their arbitrage-free levels.

When generators and LSEs are nearly risk-neutral, our model results in enormous arbitrage positions. Such large arbitrage positions seem unrealistic since participants in real-world derivatives market face collateral requirements that still need to be incorporated into our model.

We also study the price elasticity of net demand for forwards and options. Notably, net demand for both forwards and options is more sensitive to changes in forward prices than option prices. Net demand is also more elastic to negative biases in the forward price than to positive ones, especially in delivery periods with large day-ahead price spikes. Positive forward price biases encourage agents to take a short position in forwards. Such a short position can lead to extremely large losses when day-ahead price spikes occur. Therefore, agents require a larger price compensation for increasing their short position compared to increasing their long position.

Additionally, we reveal that agents take larger arbitrage positions in peak delivery periods with large price spikes than in off-peak periods with small and rare price spikes. Arbitrage is more lucrative when large price spikes occur. In off-peak periods, net demand curves for forwards and options become almost vertical in the option price. The reason is that the option price is tiny compared to the forward price because an option with a strike price of $1,000/MWh is hardly ever in the money in off-peak periods. Therefore, changes in the tiny option price are almost irrelevant to net demand.

Finally, we examine how net demand curves change if we choose a lower option strike price of $200/MWh. Overall, a lower strike price induces agents to take larger option positions and smaller forward positions. With a lower strike price, the option becomes more valuable as it protects agents against moderate price spikes, as Brandkamp et al. (2025) also demonstrate for arbitrage-free prices. In addition, we show that a lower strike price encourages agents to take larger arbitrage positions compared to a higher $1,000/MWh strike price.

In the next section, we relate our paper to the literature on forward market design. Section 4.3 describes our proposed forward market design and summarizes Budish et al.'s (2023) flow trading approach. Section 4.4 presents our simulation of ERCOT's day-ahead market and calculates day-ahead profits for representative generators and LSEs. In section 4.5, we derive net demand curves for these representative agents and sketch suitable trading strategies in the forward energy market. Section 4.6 presents simulation results and analyzes agents' net demand curves and substitution behavior between forwards and options. Section 4.7 reflects on the political implications of our novel market design. Section 4.8 concludes.

4.2 Related literature

Forward energy markets enhance efficiency by strengthening generators' incentives for competitive bidding in short-term markets (Allaz & Vila, 1993) and by enabling participants to hedge price and quantity risks. However, existing forward markets face significant challenges. The two most prominent challenges are illiquidity and market incompleteness, or the limited number of available products.

Indeed, many forward markets suffer from low liquidity, particularly for contracts beyond one year (Newbery, 2016, Genoese et al., 2016). Illiquidity leads to high price premia and unreliable price signals, as Redl et al. (2009) reveal. Illiquid markets might also make it impossible for load-serving entities (LSEs) to hedge sufficiently to be resilient against extreme weather events. Insufficient hedging on the demand side can lead to bankruptcies and market failures, as seen in the 2001 California electricity crisis and the 2022 European energy crisis (Borenstein, 2002, Joskow, 2008, Cramton, 2022).

Additionally, illiquid forward markets fail to support generators in managing risks from long-term investments. Liquidity for contracts beyond three years rapidly drops to zero (Keppler et al., 2023), discouraging investment in capital-intensive generation assets with long amortization periods. This challenge is particularly relevant for capital-heavy solar and wind assets. Effective hedging can lower the cost of capital for these assets (Genoese et al., 2016).

Another shortcoming of existing forward markets is that they offer only a small number of broad derivative products with monthly, quarterly or yearly delivery periods for base and peak load. These broad products fail to account for the seasonal and daily patterns of generators' and LSEs' load.

To address the above shortcomings, researchers and regulators propose several market design innovations. One approach is offering more granular hedge products for specific clusters of hours. Granular products make risk management more effective as they align hedge positions with seasonal and diurnal load and generation patterns (Borenstein, 2007a, Boroumand et al., 2015, Wolak, 2022, Brandkamp, 2025). However, granular products require complex trading strategies and may further fragment liquidity - at least under the existing trading rules.

An alternative solution is load-following forward contracts, used in some U.S. and Australian markets. The seller of the forward contract agrees to supply a fraction of the buyer's realized demand in every hour of the contract period (Brown & Sappington, 2023). Load-following forward contracts facilitate risk management as they align hedges with load patterns. However, load-following contracts raise average wholesale prices and reduce welfare because they discourage generators from bidding aggressively in day-ahead markets when demand is low, as Brown and Sappington (2023) highlight.

Addressing the lack of liquidity, regulators often mandate generators or LSEs to sell or buy a share of their load forward (de Frutos & Fabra, 2012, Wolak, 2022). In our proposal, we also include mandatory purchase obligations for LSEs to foster liquidity and coordinate trade.

However, mandatory forward obligations alone are not sufficient to guarantee liquidity and reliable price signals. For instance, the Australian electricity market has established mandatory purchase obligations but does not provide LSEs with sufficient means for satisfying these requirements. Many LSEs, especially small retailers, have had trouble finding counter-parties for mandated contracting at reasonable and stable prices (ACCC, 2023).

Another proposal is that regulators purchase long-term electricity contracts via contracts-for-difference as Schlecht et al. (2024) and Fabra (2023) suggest. The regulator's role is to provide liquidity and reduce risk by offering a reliable counter-party. Schlecht et al. (2024) and Fabra (2023) envision that the regulator should purchase long-term contracts via big-event auctions. They acknowledge that these auctions might suffer from a lack of competition and market power. Schlecht et al. (2024) also emphasize that regulator-backed long-term contracts for generators might mute demand response if the regulator does not pass on short-term electricity prices to the demand side.

Our proposal takes a different approach to address the above shortcomings. Our core innovation is to enable market participants to trade thousands of granular products gradually over time. This approach ensures liquidity and reliable price signals while allowing market participants to hedge seasonal and diurnal load variations effectively. As discussed below, some elements of our framework have been developed in earlier studies, especially in research focusing on financial markets. We see our contribution in combining these elements and tailoring them to the unique characteristics of electricity markets.

To make gradual trade of many granular products possible, we apply the flow trading methodology developed by Budish et al. (2023) to electricity markets. Flow trading builds on three principles: 1) Trade gradually in frequent batch auctions, 2) bundle traded products in portfolios, and 3) express preferences over portfolios as piecewise linear demand curves. As opposed to many financial markets, which are continuous in time while clearing discrete quantities, the flow-trading approach relies on discrete timing (the frequent batch auctions) and trading continuous quantities (Budish et al., 2023).

It has long been established that trading gradually over time enhances market efficiency since it allows market participants to minimize adverse price impact (Black, 1971, Kyle, 1985, Vayanos, 1999). Gradual trade is executed in frequent batch auctions that discretize trading to avoid a wasteful arms race for trading speed (Budish et al., 2015). Graf et al. (2024) provide evidence that frequent batch auctions also raise liquidity in electricity markets compared to continuous trade. Kyle and Lee (2017) argue that agents should express their preferences as piece-wise linear demand curves since piece-wise linear demand guarantees the existence of market clearing prices and makes finding these prices computationally fast (Budish et al., 2023). Portfolio trading allows arbitrage among correlated assets and reflects substitution or complementarity between them, as is known from the combinatorial auctions literature (Budish et al., 2023, Cramton et al., 2024a).

Leveraging the power of flow trading, our forward energy market enables participants to trade thousands of time- and location-specific derivatives gradually. Our approach is closest to Wolak's (2022) Standardized Fixed Price Forward Contact (SFPFC) mechanism, which proposes granular hourly forward contracts. The main difference is that Wolak (2022) envisions SFPFCs to be traded in big-event auctions while we propose gradual trade over time. Big event auctions allow market participants to exercise market power and make it more difficult to manage price risk than under gradual trade. In addition, Wolak's (2022) approach proposes far fewer products in the forward energy market since our market includes both forward contracts and call options with high time and locational granularity.

Like Wolak (2022), we advocate mandatory forward purchase obligations for LSEs. However, we propose a linear increase in obligations from zero four years ahead to 100 percent one day ahead. In contrast, Wolak (2022) envisions steep increments for the purchase obligations: 85 percent four years ahead, 87 percent three years ahead, 90 percent two years ahead, and 100 percent one year ahead. A linear schedule allows LSEs to incorporate new information gradually, improving risk management and reducing price impacts. Moreover, gradual trade creates ample liquidity in each frequent batch auction since each market participant trades a rich set of products in every auction.

In a related paper, Cramton et al. (2024a) apply the ideas used in our forward energy market to the market for intersatellite communication capacity.

4.3 The forward energy market

4.3.1 Product definitions

We propose a forward energy market that is centrally operated by the system operator. It is a financial market that trades derivatives of day-ahead energy. Day-ahead energy trades on a nodal level in advanced electricity markets. Our forward energy market aggregates the nodal day-ahead products to load zone levels to limit the number of products while taking the regional characteristics of the electricity system into account (Cramton et al., 2024a).

The forward energy market allows trading three derivative products in each load zone (see Table 4.1): First, financial forward energy contracts that should be the primary risk management tool for market participants. Second, European call options with a very high strike price, e.g., 1000$/MWh. We set such a high strike price because the option's main objective is to manage risks created by large price spikes. The call options are purely financial, unlike capacity markets' reliability options, which bundle a physical component (Cramton et al., 2013). Third, Renewable Energy Certificates (RECs) that let market participants manage jurisdictional renewable requirements.

Types of contract (derivatives of day-ahead prices)	Financial forwards Call options Renewable energy certificates
Time granularity	Monthly products: year–month–hour–we/wd Daily products (only month-ahead): date–hour
Locational granularity	Zonal
Market horizon	4 years ahead until delivery
Auction frequency	Hourly

Table 4.1: Forward market product definition

Trading starts four years before the physical delivery of the energy. The three above products are traded for granular delivery periods defined as combinations of year-month-hour-weekend/weekday (Cramton et al., 2024a). For example, Weekdays 8-9 am in May 2018 is a delivery period. The granular delivery periods enable traders to align their derivative purchases with the seasonal and daily cycles of their physical load. We call these granular products “monthly products”. Starting four years ahead, agents can trade 2,304 monthly products per zone (48 months, 24 hours per day, and 2 weekend/weekday combinations, which gives 48*24*2=2,304 products).

As an illustration, Figure 4.1 shows simulated arbitrage-free forward prices for all year-month-hour-weekend products between 2019 and 2022 in the Texas ERCOT market (Cramton et al., 2024a). The graph reveals that the monthly forward prices follow seasonal and daily load cycles, with low prices in shoulder months and at night, and high prices during summer and afternoon peak hours.

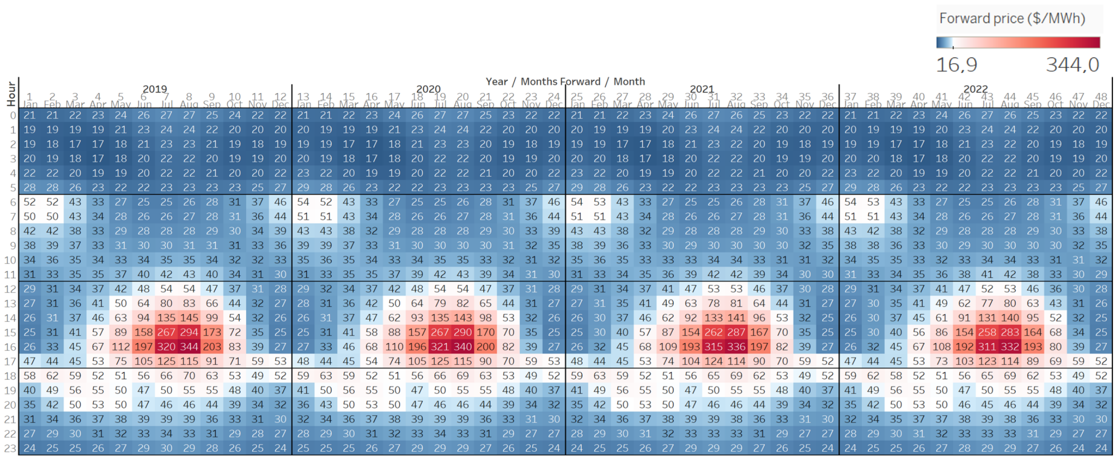

Figure 4.1: Illustrative forward electricity prices for year-month-hour products for weekdays from 2019 to 2022

In addition to the monthly products, there are “daily products” that start trading 30 days before the physical delivery. A daily product is a combination of date and hour of day. For example, 8-9 am on May 15, 2018 would be a daily delivery period. At each point in time, there are 720 (=30*24) daily products for the next 30 days ahead. The daily products should gradually link the coarser monthly products with the underlying day-ahead quantities (Cramton et al., 2024a).

Starting four years ahead of physical delivery, the forward energy market lets agents repeatedly trade every available monthly and daily product in every hour using frequent batch auctions. The bidding window in each batch auction starts one minute after the hour and lasts until the hour’s end. During the bidding window, agents can adjust their orders. The final orders at the end of the bidding window are binding. Only these binding orders are used in the market clearing optimization (Budish et al., 2023, Cramton et al., 2024a).

To increase liquidity and coordinate trade, regulators should implement a modest purchase obligation for LSEs. The obligation starts at 0 percent 48 months ahead and increases linearly to 100 percent of the realized real-time load one day ahead. The LSE can cover its obligation with forward energy or energy options.

The purchase obligation is based on the LSE’s real-time load. Recognizing that real-time load is uncertain, the LSE likely wants to buy most of its anticipated load as forward energy and some as energy options to manage risk. For example, the LSE may buy its expected load as forward energy and then sufficient energy options to cover extreme demand scenarios. The LSE is motivated to buy enough energy and options to cover realized load, even in extreme events. The LSE has two incentives to purchase ahead: 1) it must buy any shortfall at the real-time price, and 2) it pays a penalty introduced by the regulator for under-purchases of its real-time load. The system operator can raise the penalty factor if experience shows that LSEs are purchasing too few options to cover extreme events. Underestimations of the load that exceeds a specified tolerance may also increase collateral requirements.

It may seem extremely complex for market participants to jointly trade thousands of monthly and daily derivative products on an hourly basis. The forward energy market will use the flow trading methodology developed by Budish et al. (2023), which makes trading thousands of products easy for traders and computationally feasible for market operators. The above market design is closely linked to Cramton et al. (2024a) who propose to apply a similar design to the market for intersatellite communication capacity.

4.3.2 Flow trading methodology

Flow trading lets market participants place sustained portfolio orders that incentivize a smooth trade flow among multiple products. Flow trading exploits the power of convex optimization to simultaneously find market equilibria for a large number of products. The ability to trade thousands of products makes the forward energy market more complete. Moreover, granular products combined with hourly frequent batch auctions allow agents to adjust their portfolio positions efficiently as information changes over time (Cramton et al., 2024a).

We summarize the key features of the flow trading bidding language by presenting two theorems from Budish et al. (2023) and an immediate corollary. The following content until the corollary closely follows Budish et al. (2023). Moreover, the below description follows Cramton et al. (2024a), who apply Budish et al.'s (2023) methodology in an analogous way to the market for intersatellite communication capacity.

In each frequent batch auction, flow trading asks traders to submit their orders for a portfolio of products rather than for individual monthly or yearly products. Each order i for the portfolio of products must be expressed as a piece-wise linear net demand curve $D_i(p_i)$:

$$D_i(p_i|\mathbf{w}_i, q_i, p_i^L, p_i^H) := q_i \operatorname{trunc}\left(\frac{p_i^H - p_i}{p_i^H - p_i^L}\right) \qquad \text{where } \operatorname{trunc}(z) = \begin{cases} 1, & \text{for } z \geq 1, \\ z, & \text{for } 0 \leq z < 1, \\ 0, & \text{for } z \leq 0. \end{cases} \tag{4.1}$$

When net demand $D_i(p_i)$ is positive, the trader buys the portfolio while negative net demand indicates that the trader sells it. p_i is the price of the portfolio in order i. $\mathbf{w}_i = (w_{i1}, ..., w_{iN})^T$ is a vector of portfolio weights and w_{in} gives the weight of product n in portfolio order i. A positive weight means that the product is bought, while a negative weight means that it is sold (Budish et al., 2023, Cramton et al., 2024a).

Net demand $D_i(p_i)$ is represented as flows over batch intervals, constrained by a cumulative quantity limit $Q_i^{\max} > 0$. The flow rate $q_i > 0$ gives the maximum quantity of portfolio units the agent wants to buy per batch auction until the cumulative limit $Q_i^{\max}$ is reached. The flow rate allows agents to distribute their trades over time while constantly adjusting the speed at which they trade (Budish et al., 2023, Cramton et al., 2024a).

Moreover, net demand $D_i(p_i)$ is characterized by a lower limit price p_i^L and an upper limit price p_i^H with $p_i^L < p_i^H$. At prices equal to or below p_i^L, the agent wants to buy the portfolio at the maximum flow rate q_i in each batch auction. At prices above p_i^H, the agent has zero demand for the portfolio. Between these bounds $[p_i^L, p_i^H]$, the demand linearly decreases from full quantity at p_i^L to zero at p_i^H. For a sell order, both limit prices are encoded as negative numbers with $p_i^L < p_i^H$ (Budish et al., 2023, Cramton et al., 2024a).

Net demand $D_i(p_i)$ for the portfolio in order i is a function of portfolio price p_i:

$$p_i = \mathbf{w}_i^T \pi = \sum_{n=1}^{N} w_{in} \pi_n. \tag{4.2}$$

$\pi = (\pi_1, \dots, \pi_N)^T$ denotes a column vector containing the prices for all individual products $n = 1, ..., N$. The inner product of the product prices and the portfolio weights w_i give the portfolio price (Budish et al., 2023, Cramton et al., 2024a).

Let $V_i(x_i)$ be the dollar utility of order i for net demand $x_i = D_i(p_i)$ in portfolio units per hour. To determine $V_i(x_i)$, we express agents' marginal utility function $M_i(x_i)$ as an inverse demand curve $p_i = M_i(x_i)$. The inverse demand curve maps order i's net demand $x_i \in [0, q_i]$ into prices $p_i \in [p_i^L, p_i^H]$ (Budish et al., 2023, Cramton et al., 2024a). Rearranging equation 4.1 gives

$$M_i(x_i) = p_i^H - \frac{(p_i^H - p_i^L)}{q_i} x_i, \quad x_i \in [0, q_i] \tag{4.3}$$

$M_i(x_i)$ is the marginal as-bid flow value in dollars per portfolio unit (Budish et al., 2023). The utility function $V_i(x_i)$ is defined as the integral over the marginal utility function in the interval $[0, x_i]$:

$$V_i(x_i) = \int_0^{x_i} M_i(u)\, du. \tag{4.4}$$

$V_i(x_i)$ is quadratic and strictly concave in net demand x_i because marginal utility is linear in x_i (Budish et al., 2023, Cramton et al., 2024a):

$$V_i(x_i) = p_i^H x_i - \frac{(p_i^H - p_i^L)}{2q_i} x_i^2 \tag{4.5}$$

Budish et al. (2023) assume that $V_i(x_i)$ is defined for all $x_i \in \mathbb{R}$ where the constraint $x_i \in [0, q_i]$ is imposed by the order specification in equation 4.1 (Cramton et al., 2024a).

After all market participants have submitted their orders at the end of the frequent batch auction, the market operator's auction platform processes all orders to find the market equilibrium. The equilibrium prices and quantities should maximize as-bid social welfare. ”As-bid” means that the system operator assumes that the net demand curves reflect the participants’ marginal value or marginal cost. Social welfare is maximized when aggregate net demand is zero, and the market clears (Budish et al., 2023, Cramton et al., 2024a).

Budish et al. (2023) formulate the problem of finding market-clearing prices “as two optimization problems: a primal problem of finding quantities that maximize as-bid dollar value and a dual problem of finding prices that minimize the cost of non-clearing prices”. The first-order conditions for these two problems allow for deriving market-clearing prices and quantities (Budish et al., 2023, Cramton et al., 2024a). The system operator acts analogously to a social planner and selects a vector of trade rates $\mathbf{x} = (x_1, ..., x_I)$ for all submitted orders i to maximize aggregate utility:

$$\max_{\mathbf{x}} V(\mathbf{x}) = \sum_{i=1}^{I} V_i(x_i), \quad \mathbf{x} \in \mathbb{R}^I, \tag{4.6}$$

subject to:

$$\sum_{i=0}^{I} x_i \mathbf{w}_i = \mathbf{0} \qquad \text{(market-clearing constraints)} \tag{4.7}$$

$$x_i \in [0, q_i] \quad \text{for all } i \qquad \text{(trade-rate constraints)} \tag{4.8}$$

The objective function $V(\mathbf{x})$ is concave since it is a sum of concave functions. The above problem can be written as a quadratic program since the objective function is quadratic and the constraints are linear. Introducing matrix and vector notation, let $\mathbf{W}$ be the $N \times I$ matrix whose ith column is $\mathbf{w}_i$. Let $\mathbf{p}^H$ represent the column vector whose ith element is p_i^H. Let $\mathbf{D}$ denote the $I \times I$ positive definite diagonal matrix whose ith diagonal element is $(p_i^H - p_i^L)/q_i$ (Budish et al., 2023, Cramton et al., 2024a). Then, the problem in equation 4.6 can be rewritten as:

$$\max_{\mathbf{x}} \left[\mathbf{x}^T \mathbf{p}^H - \frac{1}{2} \mathbf{x}^T \mathbf{D} \mathbf{x} \right] \quad \text{subject to } \mathbf{W}\mathbf{x} = \mathbf{0}, \quad \mathbf{0} \leq \mathbf{x} \leq \mathbf{q} \tag{4.9}$$

Below, we summarize how Budish et al. (2023) demonstrate that there exist quantities that maximize aggregate utility (Theorem 1) and that there exist market-clearing prices (Theorem 2).

Theorem 1 (Budish et al., 2023)**.** There exists a unique vector of trade rates $\mathbf{x}$, which solves the maximization problem in Equation 4.9.

To prove Theorem 1, Budish et al. (2023) use the duality between the problems of finding optimal prices and quantities. They define the following Lagrangian function

$$L(\mathbf{x}, \pi, \lambda, \mu) = \mathbf{x}^T \mathbf{p}^H - \frac{1}{2}\mathbf{x}^T \mathbf{D}\mathbf{x} - \pi^T \mathbf{W}\mathbf{x} + \mu^T \mathbf{x} + \lambda^T (\mathbf{q} - \mathbf{x}) \tag{4.10}$$

The Lagrangian has three constraints: (1) the market clears ($\mathbf{Wx} = 0$); (2) trade rates are greater than or equal to zero ($\mathbf{x} \geq \mathbf{0}$); (3) the trade rates are less than or equal to their maxima ($\mathbf{x} \leq \mathbf{q}$) (Budish et al., 2023, Cramton et al., 2024a). The dual problem related to the primal problem of maximizing aggregate utility in equation 4.9 is given as

$$\widehat{G}(\pi, \lambda, \mu) = \max_x L(x, \pi, \lambda, \mu), \quad \pi \in \mathbb{R}^N, \quad \mu \geq 0, \quad \lambda \geq 0 \tag{4.11}$$

The dual problem is a minimization problem with infimum g given as (Budish et al., 2023, Cramton et al., 2024a):

$$g := \inf_{\pi,\lambda,\mu} \widehat{G}(\pi, \lambda, \mu) \qquad \text{subject to } \pi \in \mathbb{R}^N, \quad \mu \geq 0, \quad \lambda \geq 0. \tag{4.12}$$

The dual problem in equation 4.12 is formulated as an infimum rather than a minimum because we still need to show that there exists a solution (π, λ, μ) that attains the infimum (Budish et al., 2023, Cramton et al., 2024a).

Theorem 2 (Existence of market-clearing, Budish et al., 2023)**.** There exists at least one optimal solution (π, λ, μ) to the dual problem in equation 4.12. The solutions $\mathbf{x}$ and (π, λ, μ) are a primal-dual pair which satisfies the strict duality relationship $g = V(x)$.

Budish et al. (2023) prove Theorem 2 by using the properties that the primal objective function $V(\mathbf{x})$ is strictly concave and bounded since $V(\mathbf{x})$ is the sum of a finite number of concave quadratic and bounded functions. Moreover all constraints are linear and the no-trade constraint ($\mathbf{x} = \mathbf{0}$) is feasible as it clears the market. With a concave primal problem, a finite supremum on the primal problem, feasibility, and linear constraints, there must then be a dual solution that attains the same value, guaranteeing the existence of market-clearing prices (Cramton et al., 2024a).

Budish et al. (2023) stress that Theorem 2 does not imply that market-clearing prices are unique. Market-clearing prices form a convex set and may be unbounded. For instance, if all orders are buy orders and there are no sell orders, any sufficiently high price clears the market at zero trade. When there is a single buy order and a single sell order for the same product with the same maximum rate, and the buyer's lower price limit exceeds the seller's, an interval of prices allows execution (Cramton et al., 2024a).

Uniqueness of prices is important for the transparency of the market-clearing rules (Cramton et al., 2024a). Therefore, Budish et al. (2023) introduce the following tie-breaking rule that guarantees unique prices: “If more than one price vector supports the optimal quantity vector, select the price vector closest to the prior price vector in Euclidean distance” (Budish et al., 2023).

Corollary 1 (Uniqueness of quantities and prices, Cramton et al., 2024a)**.** Prices and quantities are unique with the closest-to-prior-prices rule.

Proof. The set of prices that support the unique optimal quantities is convex. The closest point in a convex set to a given point is unique (Cramton et al., 2024a). □

The closest-to-prior-prices tie-breaking rule ensures a unique mapping of orders into prices and quantities that maximizes as-bid social welfare. It is especially well-suited to frequent batch auctions, where prices shift gradually as persistent orders trade incrementally (Cramton et al., 2024a).

Flow trading makes it computationally feasible to find unique prices and quantities quickly, even for thousands of products. The flow trading problem is an instance of a “global consensus” problem, which can be efficiently solved via the alternating direction method of multipliers (ADMM) (Boyd et al., 2011). Indeed, we can fully parallelize the solution over each participant’s sub-problem, enabling any solution speed with additional GPUs. Furthermore, ADMM methods are trivially warm-started, allowing the reuse of the prior solution to achieve rapid convergence (Cramton et al., 2024a).

Budish et al. (2023) highlight how computation times increase with the number of orders and products. For instance, with a large number of 100,000 orders and 10,000 products, the computation to find unique prices and quantities takes about 100 seconds when running it on a single server. Hence, clearing the market every hour is feasible even with large numbers of orders and products (Budish et al., 2023, Cramton et al., 2024a).

4.3.3 Market settlement and collateral requirements

The market operator publishes the unique prices determined by the above flow trading methodology in the first minute of the hour. It also publicly releases the slope of aggregate net demand for each product as an indication of market liquidity. This makes the forward energy market highly transparent (Cramton et al., 2024a).

The market operator updates each trader’s position according to the quantities implied by prevailing prices. Throughout the bidding window, traders can view and download both current prices and their revised positions. Every hour, the market operator repeats this procedure. Orders stay active unless modified or canceled. If nothing changes, the trader does not need to act. Those same orders will continue to be processed every hour until a trader submits an update (Cramton et al., 2024a).

Every hour, the market operator updates each trader's settlement and alerts the trader whenever its excess collateral falls below a warning threshold. If a trader's collateral becomes negative, any trades that would further increase its collateral requirement are not allowed. For every order, any part that would unbalance the trader's position is removed (Cramton et al., 2024a).

Because the system operator has full visibility into each agent's position, the system operator can implement highly optimized collateral requirements that maximize resiliency against systemic events with minimal collateral. The collateral requirements depend on deviations from balanced positions. The exact market rules for collateral requirements will be based on a to-be developed optimization. This approach aims to maximize market stability while minimizing participants' unnecessary capital commitments (Cramton et al., 2024a).

The essential inputs in determining collateral are 1) the participant's current position, 2) the participant's expected load, and 3) the participant's expected energy production (Cramton et al., 2024a). Each participant reports 2 and 3 to the system operator. Excessive imbalances between estimated and realized load and production increase the participant's collateral requirements, consistent with the higher default risk from larger imbalances. The system operator maintains 1. When the system operator evaluates reported production during net peak load conditions, 3 defines the participant's capacity value or accredited capacity (Cramton et al., 2024a).

Bilateral forward contracts outside the centralized market are allowed, but must be reported to the system operator to take these contracts into account when adjusting collateral requirements.

Based on the above market rules, we will proceed with a proof-of-concept simulation of our forward energy market for the ERCOT market in Texas. As a first step, the next section introduces a model of ERCOT's day-ahead market. The day-ahead market is fundamental for the forward energy market.

4.4 Market simulation

In this section, we describe our model of the ERCOT day-ahead market. We jointly model hourly day-ahead electricity prices, hourly net load, hourly solar and wind generation, and daily gas prices for all hours between 2011 and 2022. After simulating joint draws of day-ahead prices, renewable generation, gas prices, and net load, we employ a merit order dispatch model to obtain hourly generation quantities and profits for a large sample of power plants.

The regime-switching market model employed here was initially developed by Coulon et al. (2013) and simulated for the ERCOT day-ahead market between 2005 and 2011. In Brandkamp et al. (2025), the model was adapted to account for renewable generation and was calibrated to data from 2011 and 2022. In this paper, we use the model and simulation results created by Brandkamp et al. (2025) and only summarize the main features of their simulation. We refer to Coulon et al. (2013) and Brandkamp et al. (2025) for a detailed exposition of the model.

4.4.1 Day-ahead market model

Coulon et al. (2013) and Brandkamp et al. (2025) simulate day-ahead prices P_t in hour t using the following exponential function:

$$P_t \ = \ G_t \ \exp\Big(\alpha_{m_k} \ + \ \beta_{m_k} \, L_t \ + \ \gamma_{m_k} \, X_t\Big) \tag{4.13}$$

G_t is the daily gas price and L_t denotes hourly net load (i.e., load minus wind and solar generation). X_t is a residual process reflecting outages or transmission constraints. Index $m_k \in \{1, 2\}$ indicates that the price is in “normal” regime 1 or in “spike” regime 2. The probability of being in the spike regime increases with net load. Because the exponential function amplifies L_t and X_t differently in each regime, the model captures both typical price variability and rare but severe price spikes. We cap the simulated day-ahead prices at 9,000\$/MWh to account for the price cap in the ERCOT day-ahead market (Coulon et al., 2013, Brandkamp et al., 2025).

We let the logarithm of gas price G_t follow an Ornstein-Uhlenbeck (OU) process:

$$d(\log G_t) \ = \ \kappa_G (m_G - \log G_t) \, dt \ + \ \eta_G \, dW_t^{(G)} \tag{4.14}$$

To capture daily and seasonal cycles, we deseasonalize net load $L_t \ = \ S(t) \ + \ \overline{L}_t$ into a seasonal term $S(t)$ and a mean-reverting residual component $\overline{L}_t$ for every hour of day h:

$$S(t) = a_1(h) + a_2(h)\cos(2\pi t + a_3(h)) + a_4(h)\cos(4\pi t + a_5(h)) + a_6(h)t + a_7(h)1_{we} \tag{4.15}$$

$$d\overline{L}_t = -\,\kappa_L \, \overline{L}_t \, dt \ + \ \eta_L \, dW_t^{(L)} \tag{4.16}$$

where 1_{we} indicates weekends (Coulon et al., 2013). Similarly, residual process $X_t = S_X(t) + \overline{X}_t$, is deseasonalized:

$$S_X(t) = b_1(h) + b_2(h)\cos(2\pi t + b_3(h)) + b_4(h)\cos(4\pi t + b_5(h)) \tag{4.17}$$

$$d\overline{X}_t = -\kappa_X\,\overline{X}_t\,dt \;+\; \eta_X\,dW_t^{(X)} \tag{4.18}$$

X_t captures unobserved supply and transmission constraints and random shocks. The deseasonalized OU processes $\overline{L}_t$ and $\overline{X}_t$ may be correlated, reflecting the fact that extreme demand conditions often coincide with plant and transmission outages due to extreme weather (Coulon et al., 2013, Brandkamp et al., 2025).

Indicator m_k defines if day-ahead prices are in normal regime 1 or spike regime 2 (see Equation 4.13). m_k is realized by an independent coin flip:

$$m_k = \begin{cases} 1 & \text{with probability } 1 - p_s\Phi\left(\frac{\overline{L}_t - \mu_s}{\sigma_s}\right) \\ 2 & \text{with probability } p_s\Phi\left(\frac{\overline{L}_t - \mu_s}{\sigma_s}\right) \end{cases}$$

$\phi(\cdot)$ represents the normal cumulative distribution function (cdf). The probability of switching to spike regime 2 rises with the deseasonalized net load $\overline{L}_t$, where $\mu_s = 0$ and $\sigma_s = \frac{\eta_L}{\sqrt{2\kappa_L}}$ are the mean and standard deviation of the stationary distribution of $\overline{L}_t$. In both regimes, day-ahead prices follow the same exponential form in Equation 4.13, but differ in their parameters $\alpha_{m_k}, \beta_{m_k}, \gamma_{m_k}$. In the spike regime, these parameters produce a sharper sensitivity of prices to net load and residual shocks, capturing rare yet extreme price events (Coulon et al., 2013, Brandkamp et al., 2025).

Brandkamp et al. (2025) calibrate the above regime-switching model to hourly data for ERCOT between 2011 and 2022. Hourly day-ahead electricity prices come from ERCOT (2024a), and hourly system load data from ERCOT (2024b). Daily Henry Hub gas price data is obtained from EIA (2024). In addition, Brandkamp et al. (2025) employ hourly generation profiles of 218 wind farms and 189 solar farms, simulated by ERCOT (2023), and supplement them with aggregate hourly renewable generation data from ERCOT (2024c).

To calibrate the model to the data, Brandkamp et al. (2025) first deseasonalize net load and the residual process X_t. The estimated parameters of the seasonal processes in Equations 4.15 and 4.17 are presented in Appendix C.1 in Brandkamp et al. (2025). Next, maximum likelihood is used to estimate the parameters of the OU processes in Equations 4.14, 4.16, and 4.18, and the exponential day-ahead price Equation 4.13 (Coulon et al., 2013, Brandkamp et al., 2025). The estimated parameters are summarized in Tables 4.2 and 4.3.

κ_L	η_L	κ_G	η_G	m_G	κ_X	η_X	ν
125.571	84,329.524	3.524	0.915	1.159	996.966	40.406	0.092

Table 4.2: Estimated parameters relating to the stochastic processes for G_t, L_t, and X_t.

α_1	β_1	γ_1	α_2	β_2	γ_2	p_s
1.279	2.39×10^{-5}	0.308	-0.035	7.12×10^{-5}	1.065	0.161

Table 4.3: Estimated parameters for the day-ahead price function in equation 4.13

$\beta_2 > \beta_1$ implies that net load and day-ahead prices have a steeper exponential relation in spike regime 2 than in normal regime 1. $\gamma_2 > \gamma_1$ means that random shocks such as plant or transmission outages also have a stronger effect on day-ahead prices in the spike regime. $p_s = 16.1\%$ indicates that day-ahead prices are in the spike regime in 8.05% ($p_s/2$) of hours (Coulon et al., 2013, Brandkamp et al., 2025).

In addition to the above regime-switching model, we also utilize the model of hourly solar and wind generation that Brandkamp et al. (2025) estimate for 218 wind farms and 189 solar farms using their generation profiles provided by ERCOT (2023). The model allows simulating plant-level hourly solar and wind generation. It captures daily and seasonal generation cycles and accounts for the correlation between each plant's hourly generation and aggregate wind and solar generation (Brandkamp et al., 2025).

4.4.2 Merit-order model

Based on the above day-ahead market model, Brandkamp et al. (2025) use a merit order model to simulate dispatch and profits for a sample of 655 power plants that operate in the ERCOT market. Table 4.4 summarizes the number of plants and their capacities by technology. For conventional plants, plant-level data on capacity, commissioning dates, heat rates, and variable costs are provided by Mann et al. (2017). For solar and wind farms, plant-level data on capacity and commissioning dates is obtained from ERCOT (2023).

Employing Brandkamp et al.'s (2025) regime switching model and the renewable generation model, we first simulate a large number of day-ahead market states. We define a market state as a set of hourly day-ahead price P_t, net load L_t, gas price G_t, and hourly

Technology	Plant Count	Capacity (MW)	% in Total Capacity
Wind	218	39,203	27.29
Solar	189	24,701	17.19
Gas Combined Cycle	68	35,132	24.45
Gas Steam Turbine	46	11,970	8.33
Gas Open Cycle	83	6,438	4.48
Gas Combustion Engine	10	671	0.47
Lignite	12	7,142	4.97
Coal	20	12,637	8.80
Hydro	5	555	0.39
Nuclear	4	4,981	3.47
Biogas	4	93	0.06
Biomass	2	150	0.10

Table 4.4: Overview of plant types, their counts, nameplate capacities, and shares in total capacity.

solar and wind generation values. For each set of values, we determine hourly dispatch and profit for each conventional plant using a standard merit order model as presented in Brandkamp et al. (2025). In addition, we use Brandkamp et al.'s (2025) renewable generation model to simulate corresponding generation and profits for the solar and wind plants.

The large number of simulated market states provides us with a rich distribution of hourly day-ahead prices, quantities, and profits. Using these distributions, we continue deriving optimal trading strategies in our forward energy market.

4.5 Trading in the forward energy market

4.5.1 Optimal forwards and option quantities

In this section, we characterize optimal trading strategies for market participants following Brandkamp et al. (2025). We start by modeling day-ahead market profits for a load-serving entity (LSE) and a generator (GEN) in hour t as:

$$\pi^{LSE}_{DA_t} = (R_Y - P_t)q_t \qquad\qquad \pi^{GEN}_{DA_t} = P_t q_t - c_t(q_t)$$

where P_t is the day-ahead price, q_t is the agent's load or generation, R_Y the fixed retail rate in year Y, and $c_t(q_t)$ the generator's cost function. For some generation technologies, costs may vary with fuel prices on a yearly or daily basis (Brandkamp et al., 2025).

Both agents can hedge their day-ahead market risk by buying or selling forwards and European call options. The total hedged profit for agent $i \in \{\text{LSE, GEN}\}$ in hour t is:

$$\pi_t^i \;=\; \pi_{DA_t}^i - (F_M^p - P_t)\, h_M - (V_M - v_t)\, z_M,$$

h_M and z_M are the quantities of forwards and call option contracts for delivery period M. F_M^p is the forward price in delivery period M, V_M is the option price and $v_t = \max\{P_t - K_M,\ 0\}$ is the option payoff with strike price K_M (Brandkamp et al., 2025).

Forwards and option quantities h_M or z_M can be positive or negative. Negative quantities indicate that the agent sells, while positive quantities indicate that it buys. Importantly, we assume that the chosen forwards and option quantities do not affect agents' bidding incentives, and equilibrium prices and quantities in the day-ahead market (Brandkamp et al., 2025).

Agents select the optimal forwards and option quantities h_M and z_M under uncertainty about day-ahead prices P_t, load q_t, and generation costs $c_t(\cdot)$. We define option and forward quantities as optimal if they maximize agent i's expected utility from profits in delivery period M (Brandkamp et al., 2025):

$$\max_{h_M,\, z_M} \; E[\,U(\pi_{t\in M}^i)] \tag{4.19}$$

The optimal quantities $h_M^*(F_M^p, V_M)$ and $z_M^*(F_M^p, V_M)$ that solve the above problem are functions of the prices for forwards and options F_M^p and V_M. Unlike Brandkamp et al. (2025), we let agents express preferences for forwards and options quantities for prices that may deviate from arbitrage-free levels. To integrate this with the flow trading methodology (Section 4.3.2), agents specify their net demand for a combined portfolio of forwards and options as a downward sloping piece-wise linear function of the portfolio price (i.e., a linear combination of forwards and option prices).

To simulate optimal net demand curves, we assume that agents have constant absolute risk aversion (CARA) utility functions. With CARA utility, we require numerical optimization to find optimal forwards and option quantities because day-ahead market profits are not normally distributed due to price-spikes and correlation between prices and load (Lapan et al., 1991, Brandkamp et al., 2025).

4.5.2 Trade-to-target strategies

At any given time, the net demand curve specifies the optimal forward and option quantities based on current information about the distribution of day-ahead prices and load. However, executing the entire net demand in a single order would adversely affect market prices. Instead, large agents typically split their orders into smaller chunks and trade gradually to mitigate price impact (Budish et al., 2023). Trading gradually also has the advantage that agents can constantly adjust their orders in response to new information. Incremental trading, thereby, facilitates risk management (Cramton et al., 2024a).

Flow trading makes gradual trade over time much easier than conventional limit order bids. Flow trading asks agents to express a trade rate, which determines how fast agents want to trade to satisfy their net demand. This makes it easy to split the large total net demand into small fractions that are traded over time (Budish et al., 2023). Agents can adjust trade rate and net demand curves every hour to respond to new information (Cramton et al., 2024a).

A straightforward strategy within this framework is "trade-to-target". Agents set their net demand as a target to be reached by the start of physical delivery. They also specify the trade rate at which they want to move toward the target. For instance, an LSE may choose a net demand target for forwards equal to the expected load it needs to serve during the delivery period. The LSE could select the trade rate such that the fraction of net demand it purchases increases linearly from zero to its target as time moves from 48 months ahead to day-ahead. Its call option target may be enough to cover potential demand surges during extreme weather (Cramton et al., 2024a).

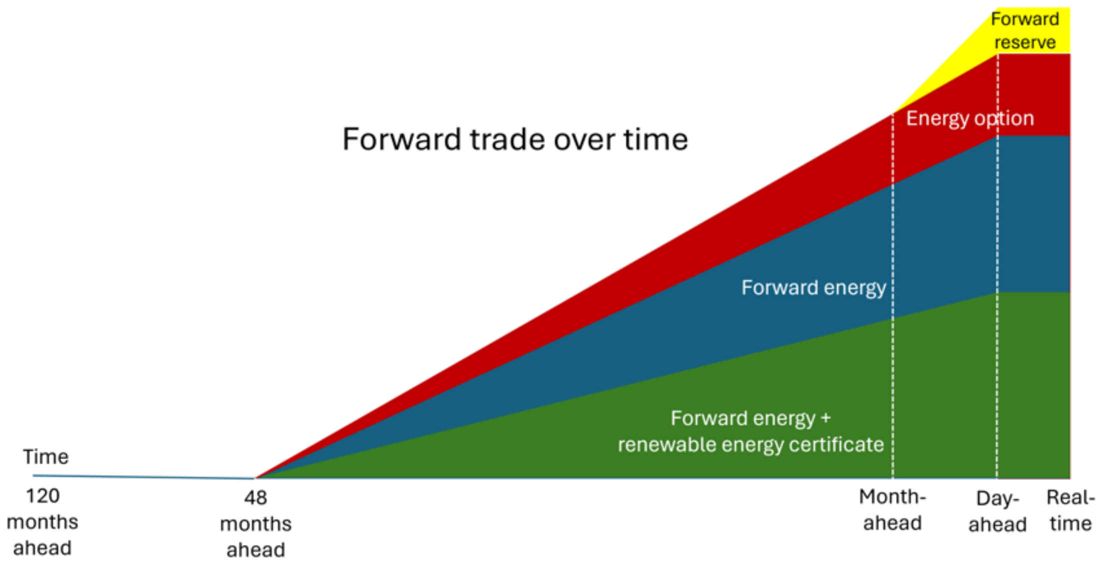

Figure 4.2: Illustration of a trade-to-target strategy for an LSE (not to scale)

Figure 4.2 illustrates a trade-to-target strategy for an LSE that has a 50 percent renewable energy requirement. The LSE purchases 50 percent of its expected demand with renewable energy certificates (RECs, green) and the remaining forward energy without RECs (blue). To cover unanticipated demand, it buys energy options (red). Beginning one month ahead, the LSE may purchase forward reserves (yellow) to get some additional forward energy if required by expected market conditions. Target purchases are completed one day ahead. Agents then respond to new information with intraday adjustments until real time.

Flow trading also makes it easy for the LSE to substitute forwards and options if their relative prices change since net demand is always expressed for a portfolio of forwards and options (Cramton et al., 2024a).

In the next section, we describe how we aim to simulate linear trade-to-target strategies in the forward energy market for representative LSEs and generators with diverse risk preferences, load-serving obligations, and power plant portfolios. We plan to derive forward market clearing prices and quantities using Budish et al.'s (2023) flow trading methodology. We will use these clearing prices and quantities to calculate hourly hedged profits for each market participant.

4.6 Simulation results

In this section, we numerically simulate hedging strategies for a set of representative generators and LSE for all hours in 2019. We initialize the simulation on $t =$ December 1, 2018 and run hourly frequent batch auctions in each hour until December 31, 2019. At the start of the simulation, we initialize all random processes at their long-term average, i.e., $\overline{L}_t = \overline{X}_t = 0$, $log(G_t) = m_G$ (Coulon et al., 2013, Brandkamp et al., 2025). Moreover, we set the risk-free interest rate at 2% and choose a high call option strike price of $\$1,000$/MWh so that the option covers only the tail of the day-ahead price distribution.

We create a set of representative generation companies and load-serving entities (LSEs) to analyze a realistic market structure. For the generation side, we create 10 generation companies by allocating the 655 power plants in our sample to one of these companies such that the resulting power plant portfolios mimic representative generation firms (Mann et al., 2017). As an example, the capacity mix of three generation companies is shown in Figure 4.3. The capacity mixes for all other generation companies are given in Figures D.8 and D.9 in Appendix D.4 (Brandkamp et al., 2025).

Generation owner 1 in Figure 4.3b represents a large incumbent generation firm with 147 power plants and 34 GW installed capacity. Its plant portfolio is technologically

(a) Legend

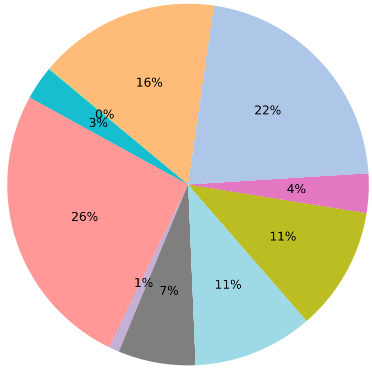

(b) Owner 1 (34.0 GW and 147 plants)

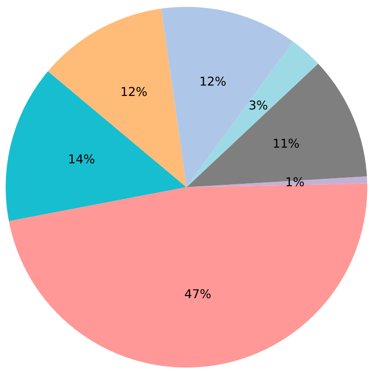

(c) Owner 5 (8.0 GW and 29 plants)

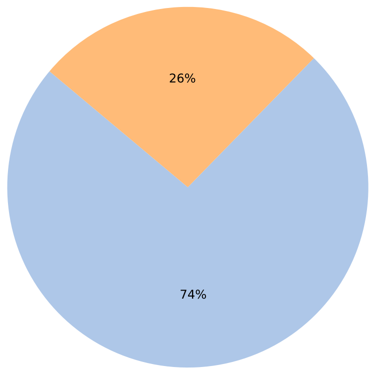

(d) Owner 8 (3.0 GW and 16 plants)

Figure 4.3: Capacity mix per owner in % of total capacity (total capacity in brackets)

well-diversified and representative for the Texan market. Owner 5 in Graph 4.3c is a smaller-mostly conventional generation company with 8 GW installed capacity and 29 plants. More than two-thirds of its portfolio consists of gas plants. Owner 8 in Figure 4.3d mimics a small renewable-only player who has 3 GW installed capacity across 16 solar and wind farms. Due to their different size and technology mix, these generation companies might have very different hedging needs (Brandkamp et al., 2025).

On the demand side, we also create 20 representative LSEs. We assume that each LSE serves, on average, a certain share of aggregate load. Demand market shares range from 2% to 12% of aggregate load. Based on Brandkamp et al. (2025), each LSE's load follows aggregate load while a random noise process causes the LSE's load to occasionally deviate from aggregate load. We assume perfect competition on the demand side, i.e., all LSEs charge their end-customers the same retail price equal to the average annual day-ahead price plus a 10% margin that allows LSEs to cover their fixed costs and to make a profit (Peura & Bunn, 2021, Brandkamp et al., 2025).

Below, we present simulated net demand curves to examine agents' preferences for forwards and options. We focus on simulating hedging strategies for incumbent generation owner 1 (see Figure 4.3b) and for a medium-size LSE with 5% market share.

Figure 4.4 shows matrices of optimal forward and call option ratios for the monthly delivery period on Weekdays 4-5 pm in August 2019. These ratios are expressed as a fraction of the agents' expected load in the delivery period. A 2% option bias indicates that the option price is 2% above its arbitrage-free level. For these matrices, we parameterize agents' risk preferences using a moderate relative risk coefficient of 0.2.[2] The delivery period Weekdays 4-5 pm in August 2019 is a peak period with large and frequent price spikes (Brandkamp et al., 2025)

Graph 4.4b reveals that the moderately risk-averse LSE wants to buy 98% of its expected load in forwards and 57% in options when both forwards and options have zero bias (i.e., prices are at their arbitrage free levels). The generator wants to sell 102% of expected load in forwards and only 2% in options for unbiased prices. With a negative forward bias, the generator reduces its short forward position and the LSE increases its long forward position (Graphs 4.4a and 4.4c). To offset the large forward position, the LSE cuts its long option position while the generator takes large short option positions (Graphs 4.4b and 4.4d). Conversely, a positive forward bias makes buying forwards less attractive such that the LSE reduces its long forward position and the generator goes short in forwards. At the same time, a high forward price prompts both players to buy more options (or at least sell less of them). Similarly, a negative option price bias makes buying options more attractive relative to buying forwards and vice versa.

For both agents, forward and option ratios are more elastic to forward prices than to option prices. Figures 4.5a and 4.5c show the agents' downward-sloping inverse demand curves for forwards when fixing the option price at its arbitrage-free level (zero bias). Plots 4.5b and 4.5d depict the inverse demand curves for options when holding the forward price unbiased. The plots highlight that the own-price elasticity of the option demand is smaller than for the forward demand. Especially, the LSE's option demand is

[2]Following the literature, we transform the relative risk coefficient as an input for CARA utility functions in order to make risk preferences comparable between players who have very different profit scales. To do so, we choose a relative risk coefficient as an exogenous parameter. Then, we divide the relative risk coefficient by the agent's average profit in the delivery period to translate the relative risk coefficient into an absolute risk coefficient. Locally, at the average profit level, the relative risk coefficient and the absolute risk coefficient describe the same risk aversion level. We assume this absolute risk coefficient to be constant across all profit levels to use it as an input for the CARA utility function (Raskin & Cochran, 1986, Newbery, 1989, Lapan & Moschini, 1994, Brandkamp et al., 2025).

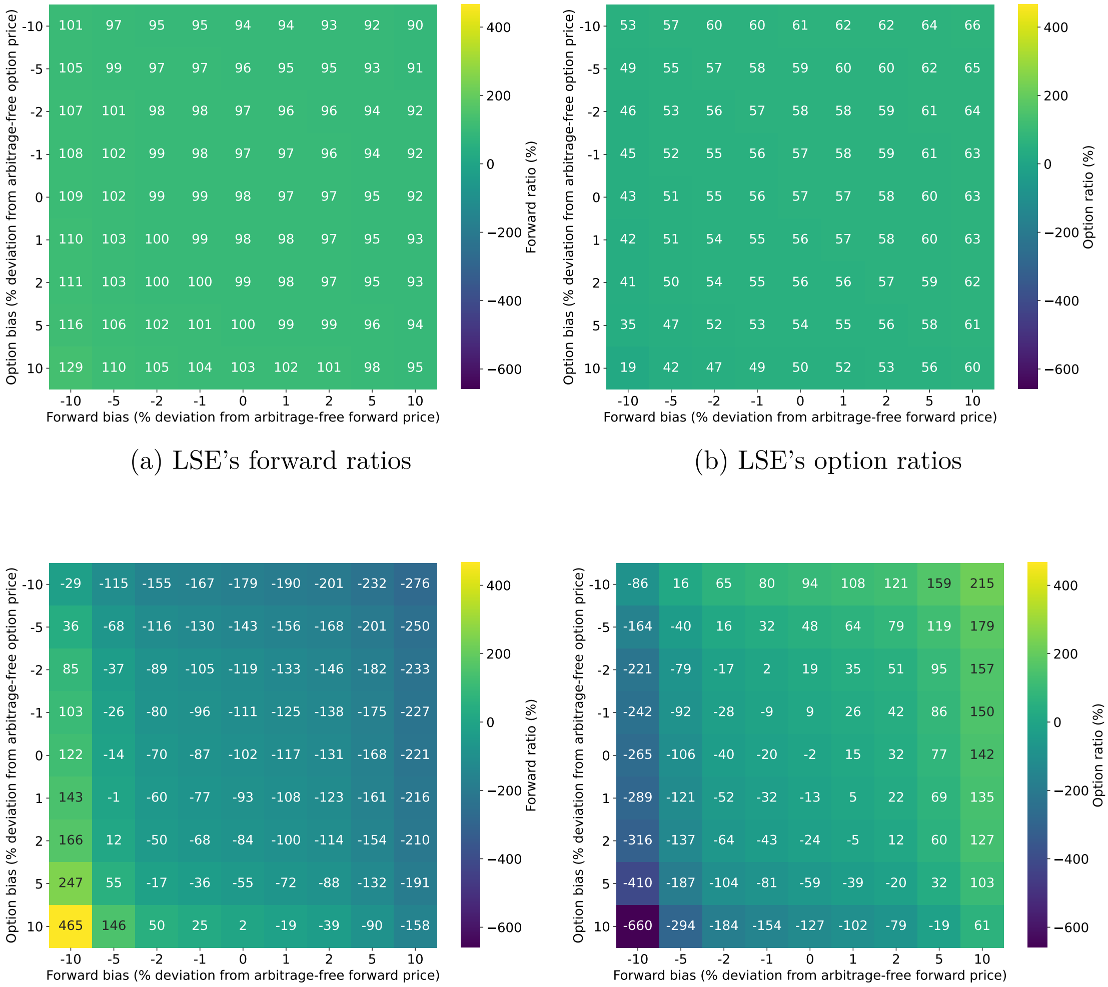

(a) LSE's forward ratios

(b) LSE's option ratios

(c) Generator's forward ratios

(d) Generator's option ratios

Figure 4.4: LSE and generator's optimal forward and option ratios as a function of the bias in forward and option prices for peak delivery period weekdays 4-5 pm in August 2019 for a relative risk coefficient of 0.2

hardly elastic to the option price (Plot 4.5b). In addition, Figure D.1 in Appendix D.1 indicates that the cross-price elasticity of the option demand with respect to the forward price exceeds its own-price elasticity. Option demand rises in the forward price. Demand for forwards is moderately elastic to option prices. However, the forward's own-price elasticity is larger than its cross-price elasticity with respect to the option price, especially for the generator.

The forward-price elasticity of forwards and options is larger for negative biases than for positive ones. This asymmetry arises because a positive forward bias induces a short

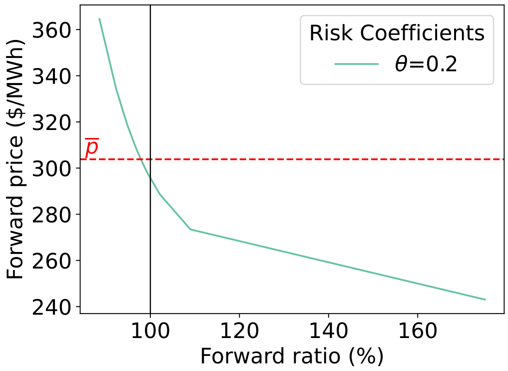

(a) LSE's forward ratios

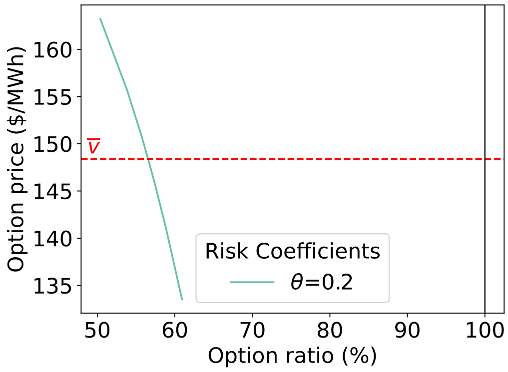

(b) LSE's option ratios

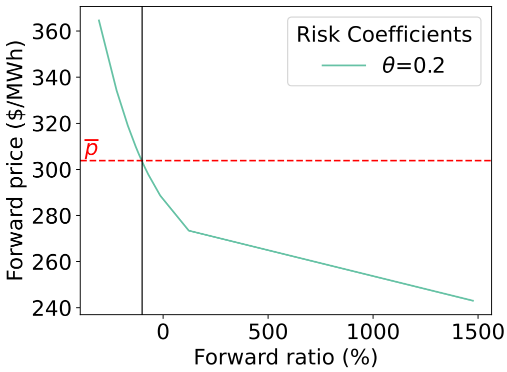

(c) Generator's forward ratios

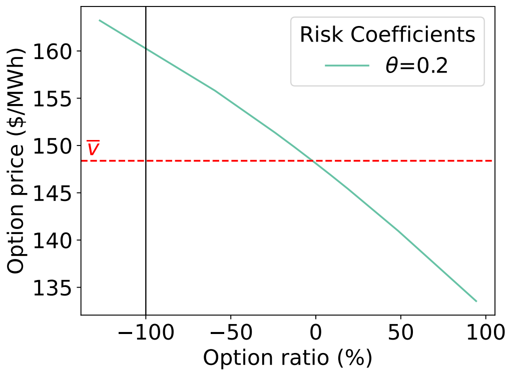

(d) Generator's option ratios

Figure 4.5: LSE and generator's net demand curves for forwards and options for peak period weekdays 4-5 pm in August 2019

forward position. Such short position can incur extreme losses during peak periods with high price spikes (Brandkamp et al., 2025). Consequently, both the generator and the LSE take only small speculative short positions even if there is a large positive bias in the forward price. Negative forward biases induce larger speculative long positions because long positions do not trigger large losses when spike prices occur.

Overall, the generator takes larger arbitrage positions in response to biased prices, likely because the LSE faces more extreme tail risks in the day-ahead market (Brandkamp et al., 2025). Therefore, the LSE is more cautious when speculating. Notably, the generator takes substantial arbitrage positions for relatively small price biases of −10% to 10% for this peak period, where large and frequent price spikes make speculation highly profitable.

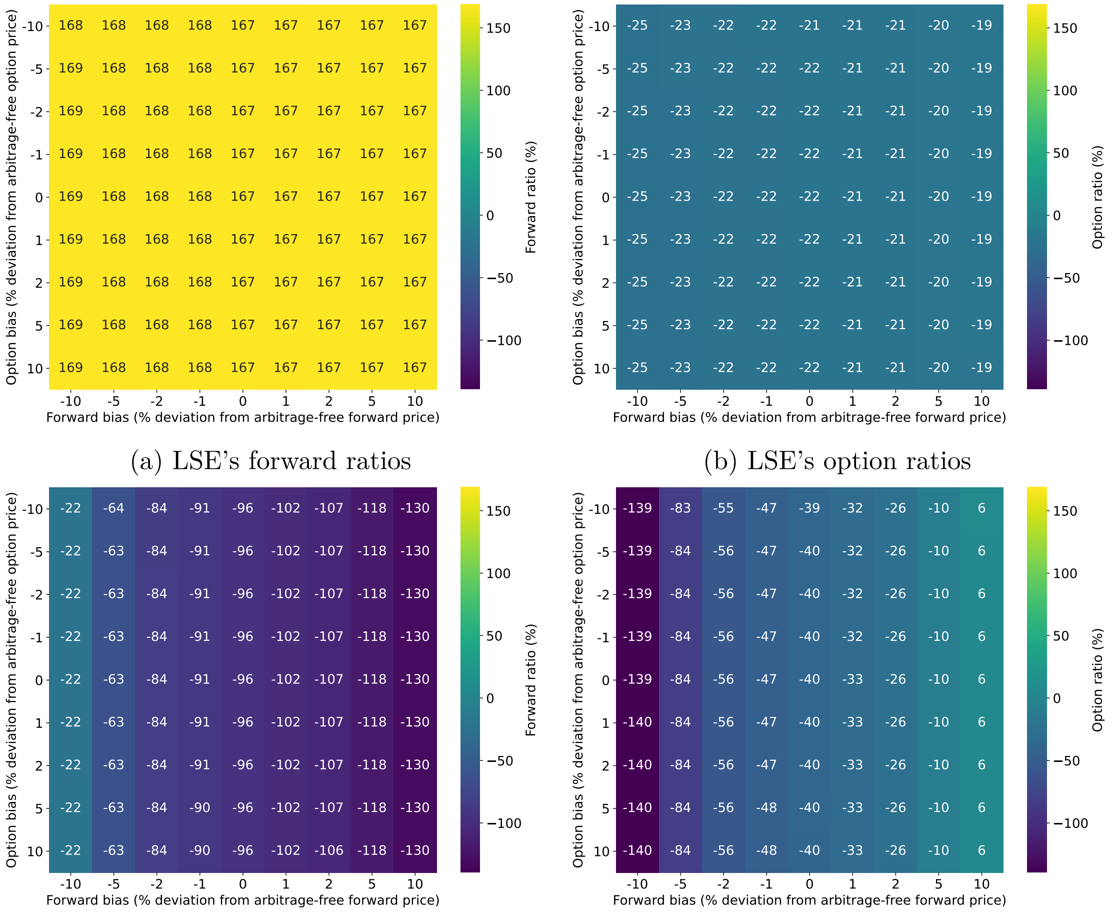

(a) LSE's forward ratios

(b) LSE's option ratios

(c) Generator's forward ratios

(d) Generator's option ratios

Figure 4.6: LSE and generator's optimal forward and option ratios as a function of the bias in forward and option prices for off-peak delivery period weekends 4-5 am in May 2019 for a relative risk coefficient of 0.2

In contrast, Figure 4.6 presents similar matrices for the off-peak period weekends 4-5 am in May 2019, when price spikes are small and rare. Here, both agents take much smaller arbitrage positions. Arbitrage is less attractive in off-peak periods with smaller price spikes. Moreover, in the off-peak period, net demand is almost insensitive to changes in the option price. The option has a very small value relative to the forward because simulated day-ahead prices hardly ever exceed the strike price of $\$1,000$/MWh in this off-peak period. Nonetheless, both agents still opt for substantial short option positions. The LSE finds it optimal to short options to offset its large long forward position.

Remarkably, even with a 10% positive forward bias, the generator only wants to sell at most 130% of expected demand forward while the LSE wants to buy at least 167% forward. If only these two agents participated in the in the market, an even larger forward bias would be required to clear the market.

However, risk preferences play a key role in agents' appetite for arbitrage. Appendix D.2 compares the ratio matrices for nearly risk-neutral agents (relative risk coefficient 0.001) with those for highly risk-averse agents (risk coefficient 1), both for the off-peak and peak period.

Both agents take very large arbitrage positions when they are nearly risk-neutral. Interestingly, the almost risk-neutral LSE wants to speculate more with options in the off-peak period than in the peak period as a comparison of Figures D.2a, D.2b, D.4a, and D.4b in Appendix D.2 highlight. By contrast, the moderately risk-averse LSE (risk coefficient 0.2) wants to arbitrage more in the peak period than in the off-peak period, as shown above.

The arbitrage positions for the nearly risk-neutral agents are likely unrealistically large. Even risk-neutral agents likely take smaller arbitrage positions due to collateral requirements. So far, we have not modeled these requirements. In the future, we will extend the objective function in equation 4.19 in section 4.5.1 such that it contains a penalizing collateral component that increases when agents take speculative imbalanced positions to limit arbitrage.

The figures in Appendix D.2 also reveal that the agents speculate much less when they are highly risk-averse (risk coefficient 1). Especially in the off-peak period, highly risk-averse agents barely arbitrage. It would likely require very large positive biases in options and forward prices to clear the market if all market participants were highly risk-averse. However, it seems unlikely that the majority of generation companies and LSEs are highly risk-averse.

A distinctive feature of our market design is the high strike price of $\$1,000$/MWh for the call option. Appendix D.3 explores the impact of a lower strike price of $200/MWh on net demand. For arbitrage-free prices, Brandkamp et al. (2025) finds that a lower strike price induces agents to choose more options and fewer forwards because the option grants better protection from moderate price spikes if the strike price is lower. In Appendix D.3, we reproduce this finding. In addition, we show that a low strike price of $200/MWh incentivizes larger arbitrage positions compared to the higher $\$1,000$/MWh strike price.

4.7 Discussion

In the above simulations, we have analyzed how generators and LSEs can use forwards and options in our forward energy market to manage risk. The essential innovation of our proposal is that forwards and options are traded gradually. In a next step, our research project aims to demonstrate that gradual trade leads to smooth and slow movements of prices and quantities. Our approach is designed to achieve smooth price movements because it incentivizes market participants to provide a constant flow of liquidity for thousands of granular products. Price signals for forward energy, thereby, become more stable and reliable (Cramton et al., 2024a).

Stable and transparent price signals provide market participants with crucial information for efficient operation and investment. Granular forward prices enable demand and supply resources to create maximum value in their operation throughout the day, season, and year. This value motivates efficient investment, the main driver of competition. Reliable prices also stimulate innovation in new resource types, especially resources that create value by flexibly responding to prices. Batteries and other low-carbon technologies are good examples. Demand-side innovation benefits from transparent prices. LSEs and other service providers can offer valuable services that optimize the use of low-carbon technologies to maximize consumer welfare. Energy efficiency programs also benefit from the multi-year forward price information.

In addition, the forward energy market would improve resiliency and reliability. Robust forward prices would amplify incentives to invest in technologies encouraging price-responsive demand, such as electric vehicles and smart homes. Retail providers would be encouraged to offer dynamic rate plans that allow consumers to create value by being flexible. These dynamic rates could include automatic hedging via forward purchases such that the dynamic rate reduces downside consumer risk relative to a fixed rate (Brandkamp, 2025). Reliable price information four years ahead would give households, service providers, and industry the information necessary to make these investment decisions.

Long-term price signals also support generation companies in managing the substantial risks involved in investing in new generation assets. Thereby, the forward energy market might replace capacity markets and strategic reserves, or at least reduce their relevance. Capacity markets are adopted to strengthen investment incentives. Flaws in the day-ahead and real-time markets may depress spot revenues below desired investment levels. The flaws are 1) a too-low price cap, 2) treatment of nonconvex costs, and 3) unpriced operator decisions, such as intraday commitments for reliability. These flaws create a “missing money” problem (Joskow, 2008, Cramton et al., 2013). Capacity markets address the missing money problem by creating additional revenue streams for generators.

The major problem of capacity markets is market power. Capacity markets use big-event auctions to procure capacity. Typically, a single auction procures 100 percent of the required capacity three years ahead. Since some market participants are dominant, the scope for exercising market power is considerable. Modern capacity markets rely on procedures to mitigate market power such as price caps or minimum offer prices. However, these procedures are controversial and imperfect (Patterson & Reiter, 2016, Macey & Ward, 2021).

By contrast, our forward energy market leaves little room for exercising market power since it encourages granular and gradual trade over many years. Participants trade forward small quantities from near-balanced positions (anticipated load + sales ≈ anticipated production + purchases). Therefore, they have little incentive or opportunity to exercise market power.

At the same time, the forward energy market effectively addresses the three above flaws that capacity markets are supposed to tackle: With a robust forward energy market, the regulator can raise the price cap because market participants are in a nearly balanced position and sufficiently hedged from high real-time prices. Nonconvex costs become less important with robust forward prices that introduce a significant volume of convex arbitrage bids. If the forward energy market was combined with intraday rolling settlement, the intraday reliability decisions would be optimized and efficiently priced. By contrast, existing capacity markets do not directly address these underlying spot market problems.

Ditching capacity markets will not compromise reliability and resiliency. LSEs have a strong incentive to buy forward, especially with options, because the joint occurrence of day-ahead price spikes and high retail demand exposes LSEs to enormous downside risk. In addition, the regulator should introduce a penalty that LSEs pay for any quantity not purchased in advance. The penalty provides an extra incentive to purchase energy options for load that is not expected but may develop in real time. Real-time purchases and imbalance penalties are paid from collateral. The collateral requirement increases with the size of the imbalance, consistent with the higher default risk from larger imbalances. Generation companies’ incentives to sell options stem from LSEs’ willingness to pay a small risk premium to be protected from adverse downside risks. The risk premium provides an additional revenue stream that helps generators to cover their investment costs.

These incentives to buy options should be strong enough for the generators and LSEs to take action to cover their load even in extreme weather events. Resiliency depends on these incentives. The regulator should raise the day-ahead price cap or the penalty factor if the incentives prove too weak. These parameters need to be set to achieve reliability and resiliency goals. The penalty factor is easy to raise if experience demonstrates that LSEs buy insufficient load forward to cover extreme events.

Overall the forward energy market is effective in guaranteeing the reliability and resiliency of the electricity system and in offering attractive risk management tools. Yet, one might wonder if it is necessary that the system operator establishes such a new centrally-organized market. Private exchanges like ICE and CME have already organized established markets for forward energy contracts and options.

However, these private markets have severe problems. Liquidity in private exchanges is poor for products more than one year ahead. There are significant frictions due to the duopoly providers' profit incentives, which include charging high fees for services, such as low-latency data feeds and co-location services, that are only needed because of a flawed trading format (Budish et al., 2015). The trading format provides too little time and location product granularity, and participants need expensive algorithmic tools to slice orders into thousands of pieces to limit adverse price impact. Because of these misaligned incentives of the private exchanges, the forward energy market can and should be managed by the independent system operator, which alone is motivated to adopt an efficient and transparent firm energy market.

Through knowledge of positions, the system operator can establish highly optimized collateral requirements that maximize market resiliency to systemic events with minimal collateral. The collateral requirement depends on deviations from balanced positions. Markets fail when counter-parties become unreliable. Optimized collateral is critical to minimizing this vulnerability at the least cost.

4.8 Conclusion

Electricity markets worldwide need to foster rapid innovation to accommodate the energy transition. Innovation is best accomplished with a market design grounded on first principles. The foundation is the spot market, which optimizes the operation of existing resources to satisfy electricity needs at the least cost. Complementing an efficient spot market, we propose a novel forward market design that allows trading of granular forward contracts and European call options, enabling hedge positions closely aligned with actual load profiles. Our design leverages Budish et al.'s (2023) flow trading technology to ensure liquidity and mitigate price impact through gradual, small-quantity trades in hourly batch auctions.

In this project, we develop a full-scale, proof-of-concept simulation of our forward energy market design within Texas's ERCOT market spanning twelve years. This simulation aims to demonstrate the feasibility of our novel market design. Using a regime-switching model, we simulate the distribution of day-ahead profits for representative generators and LSEs. Based on these day-ahead profit distributions, we derive net demand curves for forwards and options for CARA utility-maximizing agents.

Our simulations indicate that risk preferences critically shape agents' net demand curves. With high risk aversion, both generators and LSEs tend toward nearly vertical net demand curves because agents are unwilling to take on arbitrage risk. Such high risk aversion would force market clearing prices for forwards and options to be substantially above arbitrage-free levels. Generators display a greater willingness to arbitrage as they face lower downside risks in the day-ahead market than LSEs. As risk aversion diminishes, net demand curves become less vertical, especially for generators. Flatter demand curves will likely lead market-clearing prices to converge closer to arbitrage-free benchmarks. However, nearly risk-neutral behavior produces unrealistically large arbitrage positions, underscoring the need to incorporate collateral requirements in our simulations in the future.

We also observe that net demand for both forwards and options is more responsive to changes in forward prices than in option prices. Moreover, net demand demand reacts more acutely to negative forward price biases than to positive ones. A positive bias in the forward price incentivizes large short positions, which can trigger large downside risks.

In peak periods marked by large price spikes, arbitrage positions are considerably larger than in off-peak periods. Net demand for forwards and options becomes almost inelastic to the option price in off-peak periods because the option has a very small value in off-peak periods with rare and small price spikes.

Lastly, lowering the option strike price from $\$1,000$/MWh to $\$200$/MWh shifts agents' behavior by increasing option positions and reducing forward positions, as the lower strike offers better protection against moderate price spikes. Moreover, agents tend to engage more in arbitrage with a lower strike price.

In our future work, the simulated net demand curves above will feed as trading strategies into the flow trading methodology to yield market clearing prices and quantities in each hourly batch auction. We will investigate how clearing prices and quantities evolve as they get closer to their physical delivery periods. We will examine to what extend equilibrium prices deviate from arbitrage-free levels. Moreover, how agents' expected profits, volatility, and downside tail risk are affected by trading in a liquid forward energy market.

To implement our market design in practice, we suggest an incremental approach: Initially, the market operator should introduce an hourly, 30-day ahead forward energy (and reserve) market while retaining existing capacity requirements. The incremental implementation would operate as described above but with a weakening of the mandatory purchase obligation. The obligation starts at 0 percent of the anticipated load 30 days ahead to avoid a discontinuous purchase obligation. LSE obligations could increase linearly from 0 to 100 percent from 30 to 1 day ahead.

The incremental approach allows stakeholders to experience the benefits of liquid forward trading and improved price signals without overhauling current capacity mechanisms. Incremental adoption can occur at a modest cost and with little delay because the changes do not alter the core systems. Subsequent extension to a longer forward window is easy. Forward trading is particularly valuable within the 30 day window before physical delivery due to higher volumes and volatility.

Following this incremental approach, regulators could pilot our forward energy market without interfering with existing market design. When the advantages of our design for price discovery and risk management become apparent, regulators can easily extend the market by including more products and extending the trading period to four years ahead or even longer.

Bibliography

Afanasyev, D. O., Fedorova, E. A., & Popov, V. U. (2015). Fine structure of the price–demand relationship in the electricity market: Multi-scale correlation analysis. *Energy Economics*, *51*, 215–226. https://doi.org/10.1016/j.eneco.2015.07.011

Agency for the Cooperation of Energy Regulators (ACER). (2021). Acer annual report on the results of monitoring the internal electricity and natural gas markets in 2020 - energy retail markets and consumer protection volume. https://acer.europa.eu/Official_documents/Acts_of_the_Agency/Publication/ACER%20Market%20Monitoring%20Report%202020%20%E2%80%93%20Energy%20Retail%20and%20Consumer%20%20Protection%20Volume.pdf

Aïd, R., Campi, L., & Langrené, N. (2013). A structural risk–neutral model for pricing and hedging power derivatives. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, *23*(3), 387–438. https://doi.org/10.1111/j.1467-9965.2011.00507.x

Allaz, B., & Vila, J.-L. (1993). Cournot competition, forward markets and efficiency. *Journal of Economic Theory*, *59*(1), 1–16. https://doi.org/10.1006/jeth.1993.1001

Allcott, H. (2011). Rethinking real-time electricity pricing. *Resource and Energy Economics*, *33*(4), 820–842. https://doi.org/10.1016/j.reseneeco.2011.06.003

Andruszkiewicz, J., Lorenc, J., & Weychan, A. (2019). Demand price elasticity of residential electricity consumers with zonal tariff settlement based on their load profiles. *Energies*, *12*(22), 4317. https://doi.org/10.3390/en12224317

Australian Competition and Consumer Commission (ACCC). (2023). Competition in retail electricity market not delivering for all customers. https://www.accc.gov.au/media-release/competition-in-retail-electricity-market-not-delivering-for-all-customers

Azevedo, F., Vale, Z. A., & de Moura Oliveira, P. B. (2007). A decision-support system based on particle swarm optimization for multiperiod hedging in electricity markets. *IEEE Transactions on Power Systems*, *22*(3), 995–1003. https://doi.org/10.1109/TPWRS.2007.901463

Barbi, M., & Romagnoli, S. (2018). Skewness, basis risk, and optimal futures demand. *International Review of Economics & Finance*, *58*, 14–29. https://doi.org/10.1016/j.iref.2018.02.021

Barbose, G., Goldman, C., Bharvirkar, R., Hopper, N., Ting, M., & Neenan, B. (2005). Real time pricing as a default or optional service for c&i customers: A comparative analysis of eight case studies. *Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States)*. https://doi.org/10.2172/891019

Barbose, G., Goldman, C., & Neenan, B. (2004). A survey of utility experience with real time pricing. *Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States)*. https://doi.org/10.2172/836966

Barlow, M. T. (2002). A diffusion model for electricity prices. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, *12*(4), 287–298.

Bartusch, K. (2011). *Boosting behavioral change in residential electricity consumption* [PhD dissertation]. Mälardalen University. https://mdh.diva-portal.org/smash/record.jsf?pid=diva2%3A452374&dswid=-8119

Berger, S., Ockenfels, A., & Zachmann, G. (2023). Behavioral science can aid household participation in gas savings. *Joule*, *7*(1), 1–4. https://doi.org/10.1016/j.joule.2022.12.009

Berry, S., Gandhi, A., & Haile, P. (2013). Connected substitutes and invertibility of demand. *Econometrica: Journal of the Econometric Society*, *81*(5), 2087–2111. https://doi.org/10.3982/ECTA10135

Berry, S., & Haile, P. (2016). Identification in differentiated products markets // identification in differentiated products markets. *Annual review of Economics*, *8*(1), 27–52. https://doi.org/10.1146/annurev-economics-080315-015050

Bessembinder, H., & Lemmon, M. L. (2002). Equilibrium pricing and optimal hedging in electricity forward markets. *Journal of Finance*, *57*(3), 1347–1382. https://doi.org/10.1111/1540-6261.00463

Bevin-McCrimmon, F., Diaz-Rainey, I., McCarten, M., & Sise, G. (2018). Liquidity and risk premia in electricity futures. *Energy Economics*, *75*, 503–517. https://doi.org/10.1016/j.eneco.2018.09.002

Bhattacharya, S., Gupta, A., Kar, K., & Owusu, A. (2020). Risk management of renewable power producers from co-dependencies in cash flows. *European Journal of Operational Research*, *283*(3), 1081–1093. https://doi.org/10.1016/j.ejor.2019.11.069

Black, F. (1971). Toward a fully automated stock exchange, part i. *Financial Analysts Journal*, *27*(4), 28–35. https://doi.org/10.2469/faj.v27.n4.28

Bobbio, E., Brandkamp, S., Chan, S., Cramton, P., Malec, D., & Yu, L. (2022a). Price responsive demand in britain's electricity market. *ECONtribute Discussion Paper*, *185*. https://www.econtribute.de/RePEc/ajk/ajkdps/ECONtribute_185_2022.pdf

Bobbio, E., Brandkamp, S., Chan, S., Cramton, P., Malec, D., & Yu, L. (2022b). Resilient electricity requires consumer engagement. *ECONtribute Discussion Paper*, *184*. https://www.econtribute.de/RePEc/ajk/ajkdps/ECONtribute_184_2022.pdf

Bobbio, E., Brandkamp, S., Chan, S., Cramton, P., Malec, D., & Yu, L. (2024). Resilient electricity requires consumer engagement. *University of Maryland Working Paper*. https://cramton.umd.edu/papers2020-2024/resiliency-requires-consumer-engagement.pdf

Bollinger, B. K., & Hartmann, W. R. (2020). Information vs. automation and implications for dynamic pricing // information vs. automation and implications for dynamic pricing. *Management Science*, *66*(1), 290–314. https://doi.org/10.1287/mnsc.2018.3225

Borenstein, S. (2002). The trouble with electricity markets: Understanding california's restructuring disaster. *Journal of Economic Perspectives*, *16*(1), 191–211. https://doi.org/10.1257/0895330027175

Borenstein, S. (2005). The long-run efficiency of real-time electricity pricing. *The Energy Journal*, *26*(3). https://doi.org/10.5547/ISSN0195-6574-EJ-Vol26-No3-5

Borenstein, S. (2007a). Customer risk from real-time retail electricity pricing: Bill volatility and hedgability. *The Energy Journal*, *28*(2), 111–130. http://www.jstor.org/stable/41323097

Borenstein, S. (2007b). Wealth transfers among large customers from implementing real-time retail electricity pricing. *The Energy Journal*, *28*(2), 131–149. https://doi.org/10.5547/ISSN0195-6574-EJ-Vol28-No2-6

Borenstein, S., & Holland, S. (2005). On the efficiency of competitive electricity markets with time-invariant retail prices. *RAND Journal of Economics*, *36*(3), 469–493. https://www.jstor.org/stable/4135226

Boroumand, R. H., Goutte, S., Porcher, S., & Porcher, T. (2015). Hedging strategies in energy markets: The case of electricity retailers. *Energy Economics*, *51*, 503–509. https://doi.org/10.1016/j.eneco.2015.06.021

Botterud, A., Kristiansen, T., & Ilic, M. D. (2010). The relationship between spot and futures prices in the nord pool electricity market. *Energy Economics*, *32*(5), 967–978. https://doi.org/10.1016/j.eneco.2009.11.009

Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, *3*(1), 1–122. http://dx.doi.org/10.1561/2200000016

Braithwait, S. (2000). Residential tou price response in the presence of interactive communication equipment. In A. Faruqui & K. Eakin (Eds.), *Pricing in competitive electricity markets* (pp. 359–373). Springer US. https://doi.org/10.1007/978-1-4615-4529-3-22

Braithwait, S., & Eakin, K. (2002). The role of demand response in electric power market design. *Edison Electric Institute*, 1–57. https://scispace.com/pdf/the-role-of-demand-response-in-electric-power-market-design-2g02o3ngl0.pdf

Brandkamp, S. (2025). Hedging households against extreme electricity prices. *University of Cologne, Working Paper Series in Economics*, (106). https://ideas.repec.org/p/kls/series/0106.html

Brandkamp, S., Cramton, P., Dark, J., Hoy, D., & Malec, D. (2025). Hedging electricity price spikes with forwards and options. *University of Maryland Working Paper*. https://cramton.umd.edu/papers2020-2024/brandkamp-et-al-hedging-with-forwards-and-options.pdf

Brown, D., & Sappington, D. (2023). Load-following forward contracts. *The Energy Journal*, *44*(3), 187–222. https://doi.org/10.5547/01956574.44.2.DBRO

Brown, G., & Toft, K. B. (2002). How firms should hedge. *The review of financial studies*, *15*(4), 1283–1324. https://doi.org/10.1093/rfs/15.4.1283

Budish, E., Cramton, P., Kyle, A. S., Lee, J., & Malec, D. (2023). Flow trading. *National Bureau of Economic Research, Working Paper*, (31098). https://doi.org/10.3386/w31098

Budish, E., Cramton, P., & Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, *130*(4), 1547–1621. https://doi.org/10.1093/qje/qjv027

Bunn, D. W. (2006). Risk and electricity price dynamics. *Platts. com News Feature.*

Burger, S. P., Knittel, C. R., Pérez-Arriaga, I. J., Schneider, I., & vom Scheidt, F. (2020). The efficiency and distributional effects of alternative residential electricity rate designs. *The Energy Journal*, *41*(1). https://doi.org/10.5547/01956574.41.1.sbur

Cahana, M., Fabra, N., Reguant, M., & Wang, J. (2022). The distributional impacts of real-time pricing. *CEPR Press Discussion Paper*, (17200). https://cepr.org/publications/dp17200

Carmona, R., Coulon, M., & Schwarz, D. (2013). Electricity price modeling and asset valuation: A multi-fuel structural approach. *Mathematics and Financial Economics*, *7*, 167–202. https://doi.org/10.1007/s11579-012-0091-4

Cohen, J., Agel, L., Barlow, M., Garfinkel, C. I., & White, I. (2021). Linking arctic variability and change with extreme winter weather in the united states. *Science*, *373*(6559), 1116–1121. https://doi.org/10.1126/science.abi9167

Cohn, C. (2021). The cost of saving electricity for the largest u.s. utilities: Ratepayer-funded efficiency programs in 2018. *Washington, DC: ACEEE*. https://www.aceee.org/sites/default/files/pdfs/cost_of_saving_electricity_final_6-22-21.pdf

Conejo, A., García, R., Carrion, M., Caballero, A., & Andres, A. (2008). Optimal involvement in futures markets of a power producer. *Power Systems, IEEE Transactions on*, *23*, 703–711. https://doi.org/10.1109/TPWRS.2008.919245

Coulon, M., Powell, W. B., & Sircar, R. (2013). A model for hedging load and price risk in the texas electricity market. *Energy Economics*, *40*, 976–988. https://doi.org/10.1016/j.eneco.2013.05.020

Cowan, S. (2004). Optimal risk allocation for regulated monopolies and consumers. *Journal of Public Economics*, *88*(1-2), 285–303. https://doi.org/10.1016/S0047-2727(02)00082-8

Cox, J. C., & Rubinstein, M. (1985). *Options markets*. Prentice Hall.

Cramton, P. (2017). Electricity market design. *Oxford Review of Economic Policy*, *33*(4), 589–612. https://doi.org/10.1093/oxrep/grx041

Cramton, P. (2022). Fostering resiliency with good market design: Lessons from texas. *ECONtribute Discussion Paper*, (145). https://www.econstor.eu/handle/10419/252306

Cramton, P. (2023). A forward energy market to improve resiliency: Frequently asked questions. https://cramton.umd.edu/papers2020-2024/cramton-et-al-forward-energy-market-faq.pdf

Cramton, P., Bobbio, E., Brandkamp, S., Chan, S., Malec, D., Ockenfels, A., & Yu, L. (2025a). Resilient electricity requires consumer engagement. *Working Paper*. https://doi.org/10.2139/ssrn.5097987

Cramton, P., Bohlin, E., Brandkamp, S., Dark, J., Hoy, D., Kyle, A., Malec, D., Ockenfels, A., & Wilkens, C. (2024a). An open-access market for global communications.

Telecommunications Policy, *48*(9), 102820. https://doi.org/10.1016/j.telpol.2024.102820

Cramton, P., Brandkamp, S., Chao, H.-p., Dark, J., Hoy, D., Kyle, A., Malec, D., Ockenfels, A., & Wilkens, C. (2024b). A forward energy market to improve resiliency: Paper for decision-makers. *University of Maryland Working Paper*. https://cramton.umd.edu/papers2020-2024/cramton-et-al-forward-energy-market-decision-makers.pdf

Cramton, P., Brandkamp, S., Dark, J., Hoy, D., Malec, D., Ockenfels, A., & Wilkens, C. (2025b). A forward energy market to improve reliability and resiliency. *University of Maryland Working Paper*. https://cramton.umd.edu/papers2020-2024/cramton-et-al-forward-energy-market.pdf

Cramton, P., Ockenfels, A., & Stoft, S. (2013). Capacity market fundamentals. *Economics of Energy & Environmental Policy*, *2*(2), 27–46. https://doi.org/10.5547/2160-5890.2.2.2

Danthine, J.-P. (1978). Information, futures prices, and stabilizing speculation. *Journal of Economic Theory*, *17*(1), 79–98. https://doi.org/10.1016/0022-0531(78)90124-2

de Frutos, M.-Á., & Fabra, N. (2012). How to allocate forward contracts: The case of electricity markets. *European Economic Review*, *56*(3), 451–469. https://doi.org/10.1016/j.euroecorev.2011.11.005

de Jong, C., & Huisman, R. (2002). Option formulas for mean-reverting power prices with spikes. *Energy Global Research Paper*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=324520

Deng, S., & Oren, S. S. (2006). Electricity derivatives and risk management. *Energy*, *31*(6-7), 940–953. https://doi.org/10.1016/j.energy.2005.02.015

Deng, T., Yan, W., Nojavan, S., & Jermsittiparsert, K. (2020). Risk evaluation and retail electricity pricing using downside risk constraints method. *Energy*, *192*, 116672. https://doi.org/10.1016/j.energy.2019.116672

Department for Energy Security and Net Zero (DESNZ). (2023). Transfer statistics in the domestic gas and electricity markets in great britain. https://www.gov.uk/government/statistical-data-sets/quarterly-domestic-energy-switching-statistics

Dertwinkel-Kalt, M., & Wey, C. (2023). Why ”energy price brakes” encourage moral hazard, raise energy prices, and reinforce energy savings. *DICE Discussion Paper*, (407). https://www.econstor.eu/handle/10419/278737

Dionne, G., & Santugini, M. (2015). Production flexibility and hedging. *Risks*, *3*(4), 543–552. https://doi.org/10.3390/risks3040543

Electric Reliability Council of Texas (ERCOT). (2022). 2022 wind and solar profiles update. https://www.ercot.com/files/docs/2022/08/27/4___Updated_Wind_and_Solar_Production_Profiles.pptx

Electric Reliability Council of Texas (ERCOT). (2023). Ercot wind profiles and solar pv profiles, 1980-2022. https://www.ercot.com/gridinfo/resource/2022

Electric Reliability Council of Texas (ERCOT). (2024a). Historical dam load zone and hub prices. https://www.ercot.com/mp/data-products/data-product-details?id=np4-180-er

Electric Reliability Council of Texas (ERCOT). (2024b). Hourly load data archives. https://www.ercot.com/gridinfo/load/load_hist

Electric Reliability Council of Texas (ERCOT). (2024c). Interval generation by fuel reports (2011-22). https://www.ercot.com/gridinfo/generation

Energy Information Administration (EIA). (2018). Electric power annual 2017. https://www.connaissancedesenergies.org/sites/connaissancedesenergies.org/files/pdf-pt-vue/electric_power_annual_2017.pdf

Energy Information Administration (EIA). (2022). Short-term energy outlook supplement: Sources of price volatility in the ercot market. https://www.eia.gov/outlooks/steo/special/supplements/2022/2022_sp_03.pdf

Energy Information Administration (EIA). (2023). Electric power monthly 2022. https://www.eia.gov/electricity/monthly/

Energy Information Administration (EIA). (2024). Henry hub natural gas spot price. https://www.eia.gov/dnav/ng/hist/rngwhhdD.htm

Environmental Protection Agency (EPA). (2021). Global greenhouse gas emissions data. https://19january2021snapshot.epa.gov/ghgemissions/global-greenhouse-gas-emissions-data_.html

EPEX SPOT. (2023). Market data. https://www.epexspot.com/en/market-data

Ericson, T. (2006). Time-differentiated pricing and direct load control of residential electricity consumption. *Statistics Norway, Research Department*. https://EconPapers.repec.org/RePEc:ssb:dispap:461

Espey, J. A., & Espey, M. (2004). Turning on the lights: A meta-analysis of residential electricity demand elasticities. *Journal of Agricultural and Applied Economics*, *36*(1), 65–81. https://doi.org/10.1017/S1074070800021866

European Association of Environmental and Resource Economists (EAERE). (2024). 29th annual conference of the european association of environmental and resource economists: Conference agenda. https://www.conftool.pro/eaere2024/index.php?form_session=83&mode=list&page=browseSessions

European Commission. (2023). Questions and answers on the revision of the eu's internal electricity market design. https://ec.europa.eu/commission/presscorner/detail/en/qanda_23_1593

Eydeland, A., & Geman, H. (1999). Fundamentals of electricity derivatives, 35–43. https://eprints.bbk.ac.uk/id/eprint/32264/

Fabra, N. (2023). Reforming european electricity markets: Lessons from the energy crisis. *Energy Economics*, *126*, 106963. https://doi.org/10.1016/j.eneco.2023.106963

Fabra, N., Rapson, D., Mar, R., & Wang, J. (2021). Estimating the elasticity to real-time pricing: Evidence from the spanish electricity market. *AEA Papers and Proceedings*, (111), 425–429. https://doi.org/10.1257/pandp.20211007

Fanelli, V., Maddalena, L., & Musti, S. (2016). Asian options pricing in the day-ahead electricity market. *Sustainable cities and society*, *27*, 196–202.

Faruqui, A. (2012). Chapter 3 - the ethics of dynamic pricing. In F. P. Sioshansi (Ed.), *Smart grid* (pp. 61–83). Academic Press. https://doi.org/10.1016/B978-0-12-386452-9.00003-6

Faruqui, A., & George, S. (2005). Quantifying customer response to dynamic pricing. *The Electricity Journal*, *18*(4), 53–63. https://doi.org/10.1016/j.tej.2005.04.005

Faruqui, A., Hledik, R., & Palmer, J. (2012). *Time-varying and dynamic rate design*. Regulatory Assistance Project. https://www.raponline.org/wp-content/uploads/2023/09/rap-faruquihledikpalmer-timevaryingdynamicratedesign-2012-jul-23.pdf

Faruqui, A., & Sergici, S. (2010). Household response to dynamic pricing of electricity: A survey of 15 experiments. *Journal of Regulatory Economics*, *38*(2), 193–225. https://doi.org/10.1007/s11149-010-9127-y

Faruqui, A., & Sergici, S. (2011). Dynamic pricing of electricity in the mid-atlantic region: Econometric results from the baltimore gas and electric company experiment. *Journal of Regulatory Economics*, *40*(1), 82–109. https://doi.org/10.1007/s11149-011-9152-5

Faruqui, A., Sergici, S., & Akaba, L. (2014). The impact of dynamic pricing on residential and small commercial and industrial usage: New experimental evidence from connecticut. *The Energy Journal*, *35*(1), 137–160. http://www.jstor.org/stable/24693822

Feder, G., Just, R. E., & Schmitz, A. (1980). Futures markets and the theory of the firm under price uncertainty. *The Quarterly Journal of Economics*, *94*(2), 317–328. https://doi.org/10.2307/1884543

Ferman, M. (2021). Texas legislature approves bill to ban residential wholesale electricity plans — the first major winter storm bill sent to the governor. *The Texas Tribune,* *13 May 2021*. https://www.texastribune.org/2021/05/13/texas-power-grid-failure-legislature/

Füss, R., Mahringer, S., & Prokopczuk, M. (2015). Electricity derivatives pricing with forward-looking information. *Journal of Economic Dynamics and Control*, *58*, 34–57. https://doi.org/10.1016/j.jedc.2015.05.016

Gay, G. D., Nam, J., & Turac, M. (2002). How firms manage risk: The optimal mix of linear and non–linear derivatives. *Journal of Applied Corporate Finance*, *14*(4), 82–93. https://doi.org/10.1111/j.1745-6622.2002.tb00451.x

Gemmill, G. (1985). Forward contracts or international buffer stocks? a study of their relative efficiencies in stabilising commodity export earnings. *The Economic Journal*, *95*(378), 400–417. https://doi.org/10.2307/2233217

Genoese, F., Drabik, E., & Egenhofer, C. (2016). The eu power sector needs long-term price signals. *CEPS Special Report*, (135). https://ssrn.com/abstract=2782355

Gilbert, C. L. (1985). Futures trading and the welfare evaluation of commodity price stabilisation. *The Economic Journal*, *95*(379), 637–661. https://doi.org/10.2307/2233031

Goldman, C., Hopper, N., Sezgen, O., Moezzi, M., Bharvirkar, R., Neenan, B., Boisvert, R., Cappers, P., & Pratt, D. (2004). Customer response to day-ahead wholesale market electricity prices: Case study of rtp program experience in new york. *Ernest Orlando Lawrence Berkeley National Laboratory*. https://doi.org/10.2172/828668

Graf, C., Kuppelwieser, T., & Wozabal, D. (2024). Frequent auctions for intraday electricity markets. *The Energy Journal*, *45*(1), 231–256. https://doi.org/10.2139/ssrn.4080555

Green, R., & Staffell, I. (2017). "prosumage" and the british electricity market. *Economics of Energy & Environmental Policy*, *6*(1), 33–50. https://www.jstor.org/stable/26189570

Hambly, B., Howison, S., & Kluge, T. (2008). Modelling spikes and pricing swing options in electricity markets. *Quantitative Finance*, 1–25. https://doi.org/10.1080/14697680802596856

Hamilton, J. D. (1990). Analysis of time series subject to changes in regime. *Journal of Econometrics*, *45*(1-2), 39–70. https://doi.org/10.1016/0304-4076(90)90093-9

Hanoch, G. (1977). Risk aversion and consumer preferences. *Econometrica: Journal of the Econometric Society*, *45*(2), 413–426. https://doi.org/10.2307/1911218

Harding, M., & Lamarche, C. (2016). Empowering consumers through data and smart technology: Experimental evidence on the consequences of time–of–use electricity pricing policies. *Journal of Policy Analysis and Management*, *35*(4), 906–931. https://doi.org/10.1002/pam.21928

Harding, M., & Sexton, S. (2017). Household response to time-varying electricity prices. *Annual Review of Resource Economics*, *9*, 337–359. https://www.jstor.org/stable/26773561

Herriges, J. A., Baladi, S. M., Caves, D. W., & Neenan, B. F. (1993). The response of industrial customers to electric rates based upon dynamic marginal costs. *The Review of Economics and Statistics*, *75*(3), 446–454. https://doi.org/10.2307/2109458

Hess, M. (2021). Optimal hedging strategies for options in electricity futures markets. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3825720

Holland, S. P., & Mansur, E. T. (2006). The short-run effects of time-varying prices in competitive electricity markets. *The Energy Journal*, *27*(4). https://doi.org/10.5547/ISSN0195-6574-EJ-Vol27-No4-6

Holland, S. P., & Mansur, E. T. (2008). Is real-time pricing green? the environmental impacts of electricity demand variance. *The Review of Economics and Statistics*, *90*(3), 550–561. https://doi.org/10.1162/rest.90.3.550

Holthausen, D. M. (1979). Hedging and the competitive firm under price uncertainty. *American Economic Review*, *69*(5), 989–995. https://www.jstor.org/stable/1813672

Horowitz, S., & Lave, L. (2014). Equity in residential electricity pricing. *The Energy Journal*, *35*(2). https://doi.org/10.5547/01956574.35.2.1

Houthakker, H. S. (1951). Electricity tariffs in theory and practice. *The Economic Journal*, *61*(241), 1–25. https://doi.org/10.2307/2226608

Jessoe, K., & Rapson, D. (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review*, *104*(4), 1417–1438. https://doi.org/10.1257/aer.104.4.1417

Jin, H. J. (2007). Heavy–tailed behavior of commodity price distribution and optimal hedging demand. *Journal of Risk and Insurance*, *74*(4), 863–881. https://doi.org/10.1111/j.1539-6975.2007.00238.x

Joskow, P. L. (2008). Capacity payments in imperfect electricity markets: Need and design. *Utilities Policy*, *16*(3), 159–170. https://doi.org/10.1016/j.jup.2007.10.003

Keppler, J. H., Quemin, S., & Sagan, M. (2023). Why the sustainable provision of low-carbon electricity needs hybrid markets: The conceptual basics. https://hal.science/hal-04231032/

Ketterer, J. C. (2014). The impact of wind power generation on the electricity price in germany. *Energy Economics*, *44*, 270–280. https://doi.org/10.1016/j.eneco.2014.04.003

Kettunen, J., Salo, A., & Bunn, D. W. (2009). Optimization of electricity retailer's contract portfolio subject to risk preferences. *IEEE Transactions on Power Systems*, *25*(1), 117–128. https://doi.org/10.1109/TPWRS.2009.2032233

Kihlstrom, R. (2009). Risk aversion and the elasticity of substitution in general dynamic portfolio theory: Consistent planning by forward looking, expected utility maximizing investors. *Journal of Mathematical Economics*, *45*(9-10), 634–663. https://doi.org/10.1016/j.jmateco.2008.08.008

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, *53*(6), 1315–1335. https://doi.org/10.2307/1913210

Kyle, A. S., & Lee, J. (2017). Toward a fully continuous exchange. *Oxford Review of Economic Policy*, *33*(4), 650–675. https://doi.org/10.1093/oxrep/grx042

Lapan, H., & Moschini, G. (1994). Futures hedging under price, basis, and production risk. *American Journal of Agricultural Economics*, *76*(3), 465–477. https://doi.org/10.2307/1243658

Lapan, H., Moschini, G., & Hanson, S. D. (1991). Production, hedging, and speculative decisions with options and futures markets. *American Journal of Agricultural Economics*, *73*(1), 66–74. https://doi.org/10.2307/1242884

Lau, C. O. (2016). Disentangling intertemporal substitution and risk aversion under the expected utility theorem. *SSRN*, *2989308*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2989308

Lazard. (2020). Lazard's levelized cost of energy analysis, version 14. https://www.lazard.com/media/kwrjairh/lazards-levelized-cost-of-energy-version-140.pdf

Lee, Y., & Oren, S. S. (2009). An equilibrium pricing model for weather derivatives in a multi-commodity setting. *Energy Economics*, *31*(5), 702–713. https://doi.org/10.1016/j.eneco.2009.01.017

Lengwiler, Y. (2004). *Microfoundations of financial economics: An introduction to general equilibrium asset pricing*. Princeton University Press.

Lien, D. (2010). The effects of skewness on optimal production and hedging decisions: An application of the skew–normal distribution. *Journal of Futures Markets: Futures, Options, and Other Derivative Products*, *30*(3), 278–289. https://doi.org/10.1002/fut.20413

Losq, E. (1982). Hedging with price and output uncertainty. *Economics Letters*, *10*(1-2), 65–70. https://doi.org/10.1016/0165-1765(82)90117-3

Lu, X., Dong, Z. Y., & Li, X. (2005). Electricity market price spike forecast with data mining techniques. *Electric Power Systems Research*, *73*(1), 19–29. https://doi.org/10.1016/j.epsr.2004.06.002

Macey, J. C., & Ward, R. (2021). Mopr madness. *Energy Law Journal*, *42*(1), 67–122. https://heinonline.org/HOL/Page?handle=hein.journals/energy42&div=10&g_sent=1&casa_token=&collection=journals

Mann, N., Tsai, C.-H., Gülen, G., Schneider, E., Cuevas, P., Dyer, J., Butler, J., Zhang, T., Baldick, R., & Deetjen, T. (2017). Capacity expansion and dispatch modeling: Model documentation and results for ercot scenarios. *The Full Cost of Electricity. The University of Texas at Austin Energy Institute*. https://energy.utexas.edu/sites/default/files/UTAustin_FCe_ERCOT_2017.pdf

Matsumoto, T., & Yamada, Y. (2021). Simultaneous hedging strategy for price and volume risks in electricity businesses using energy and weather derivatives. *Energy Economics*, *95*, 105101. https://doi.org/10.1016/j.eneco.2021.105101

McDonnel, G., Bogel-Burroughs, N., & Penn, I. (2021). His lights stayed on during texas' storm. now he owes $16,752. *The New York Times*. https://www.nytimes.com/2021/02/20/us/texas-storm-electric-bills.html

McKinnon, R. I. (1967). Futures markets, buffer stocks, and income stability for primary producers. *Journal of political economy*, *75*(6), 844–861. https://doi.org/10.1086/259363

Mnasri, M., Dionne, G., & Gueyie, J.-P. (2017). The use of nonlinear hedging strategies by us oil producers: Motivations and implications. *Energy Economics*, *63*, 348–364. https://doi.org/10.1016/j.eneco.2017.02.003

Moschini, G., & Lapan, H. (1992). Hedging price risk with options and futures for the competitive firm with production flexibility. *International Economic Review*, *33*(3), 607–618. https://doi.org/10.2307/2527128

Moschini, G., & Lapan, H. (1995). The hedging role of options and futures under joint price, basis, and production risk. *International Economic Review*, *36*(4), 1025. https://doi.org/10.2307/2527271

Näsäkkälä, E., & Keppo, J. (2005). Electricity load pattern hedging with static forward strategies. *Managerial Finance*, *31*(6), 116–137. https://doi.org/10.1108/03074350510769721

Newbery, D. (1989). The theory of food price stabilisation. *The Economic Journal*, *99*(398), 1065–1082. https://doi.org/10.2307/2234088

Newbery, D. (2016). Missing money and missing markets: Reliability, capacity auctions and interconnectors. *Energy Policy*, *94*, 401–410. https://doi.org/10.1016/j.enpol.2015.10.028

Newbery, D., & Stiglitz, J. (1981). *The theory of commodity price stabilization: A study in the economics of risk* (Reprinted.). Clarendon Press.

Ocakoglu, K., & Tolga, A. (2018). Effective trading in turkish electricity market: Hedging with options. *Proceedings of the World Congress on Engineering*. https://www.iaeng.org/publication/WCE2018/WCE2018_pp314-318.pdf

Octopus Energy. (2019). Agile pricing explained. https://octopus.energy/blog/agile-pricing-explained/

O'Sheasy, M. T. (1998). How to buy low and sell high. *The Electricity Journal*, *11*(1), 24–29. https://doi.org/10.1016/S1040-6190(98)80020-1

Oum, Y., Oren, S., & Deng, S. (2006). Hedging quantity risks with standard power options in a competitive wholesale electricity market. *Naval Research Logistics (NRL)*, *53*(7), 697–712. https://doi.org/10.1002/nav.20184

Oum, Y., & Oren, S. S. (2010). Optimal static hedging of volumetric risk in a competitive wholesale electricity market. *Decision Analysis*, *7*(1), 107–122. https://doi.org/10.1287/deca.1090.0167

Patterson, D., & Reiter, H. (2016). Chasing the uncatchable: Why trying to fix mandatory capacity markets is like trying to win a game of whack-a-mole. *Public Utilities Fortnightly*, *June*. https://www.proquest.com/docview/1812902078?sourcetype=Trade%20Journals

Peura, H., & Bunn, D. W. (2021). Renewable power and electricity prices: The impact of forward markets. *Management Science*, *67*(8), 4772–4788. https://doi.org/10.1287/mnsc.2020.3710

Poletti, S., & Wright, J. (2020). Real–time pricing and imperfect competition in electricity markets. *The Journal of Industrial Economics*, *68*(1), 93–135. https://doi.org/10.1111/joie.12215

Raskin, R., & Cochran, M. J. (1986). Interpretations and transformations of scale for the pratt-arrow absolute risk aversion coefficient: Implications for generalized stochastic dominance. *Western Journal of Agricultural Economics*, 204–210. https://www.jstor.org/stable/40987816

Redding, S. J., & Weinstein, D. E. (2019). Measuring aggregate price indices with taste shocks: Theory and evidence for ces preferences. *The Quarterly Journal of Economics*, *135*(1), 503–560. https://doi.org/10.1093/qje/qjz031

Redl, C., Haas, R., Huber, C., & Böhm, B. (2009). Price formation in electricity forward markets and the relevance of systematic forecast errors. *Energy Economics*, *31*(3), 356–364. https://doi.org/10.1016/j.eneco.2008.12.001

Reiss, P. C., & White, M. W. (2005). Household electricity demand, revisited. *The Review of Economic Studies*, *72*(3), 853–883. https://doi.org/10.1111/0034-6527.00354

Rolfo, J. (1980). Optimal hedging under price and quantity uncertainty: The case of a cocoa producer. *Journal of political economy*, *88*(1), 100–116. https://doi.org/10.1086/260849

Roques, F. A., Newbery, D., & Nuttall, W. J. (2005). Investment incentives and electricity market design: The british experience. *Review of Network Economics*, *4*(2). https://doi.org/10.2202/1446-9022.1068

Sakong, Y., Hayes, D. J., & Hallam, A. (1993). Hedging production risk with options. *American Journal of Agricultural Economics*, *75*(2), 408–415. https://doi.org/10.2307/1242925

Sanchez, G. A. V., Alzate, J. M., Cadena, A. I., & Benavides, J. M. (2010). Setting up standard power options to hedge price-quantity risk in a competitive electricity market: The colombian case. *IEEE Transactions on Power Systems*, *26*(3), 1493–1500. https://doi.org/10.1109/TPWRS.2010.2089474

Schlecht, I., Maurer, C., & Hirth, L. (2024). Financial contracts for differences: The problems with conventional cfds in electricity markets and how forward contracts can help solve them. *Energy Policy*, *186*, 113981. https://doi.org/10.1016/j.enpol.2024.113981

Schwartz, E. S. (1997). The stochastic behavior of commodity prices: Implications for valuation and hedging. *the Journal of Finance*, *52*(3), 923–973. https://doi.org/10.1111/j.1540-6261.1997.tb02721.x

Schwarz, P. M., Taylor, T. N., Birmingham, M., & Dardan, S. L. (2002). Industrial response to electricity real–time prices: Short run and long run. *Economic Inquiry*, *40*(4), 597–610. https://doi.org/10.1093/ei/40.4.597

Shapiro, A., Dentcheva, D., & Ruszczynski, A. (2021). *Lectures on stochastic programming: Modeling and theory*. SIAM.

Staffell, I., Hamilton, I. G., & Green, R. (2015). *The residential energy sector*. Routledge.

Stavrogiannis, L. C. (2010). *Electricity tariff design and implementation for the smart grid* [Masters Thesis]. University of Southampton. Electronics; Computer Science Department. https://eprints.soton.ac.uk/272365/

Steinberg, D., Bielen, D., Eichman, J., Eurek, K., Logan, J., Mai, T., McMillan, C., Parker, A., Vimmerstedt, L., & Wilson, E. (2017). Electrification and decarbonization: Exploring u.s. energy use and greenhouse gas emissions in scenarios with widespread electrification and power sector decarbonization: Technical report. https://doi.org/10.2172/1372620

Stumpe, V. L. (2022). *Essays in applied microeconomics* [Doctoral Dissertation]. Rheinische Friedrich-Wilhelms-Universität Bonn. https://bonndoc.ulb.uni-bonn.de/xmlui/handle/20.500.11811/9821

The International Association for Energy Economics (IAEE). (2024). 45th iaee international conference 25 -28 june, 2024 istanbul boğaziçi university: Conference proceedings. https://www.iaee2024.org.tr/_files/ugd/b9b2b1_fce8168e91b0427db80-4b66da5646104.pdf

Thomas, N. (2021). Scottishpower warns of energy market 'massacre'. *Financial Times, 21 October 2021.* https://www.ft.com/content/e990af1f-9753-40b2-8947-bb7c65a07733

Train, K., & Mehrez, G. (1994). Optional time-of-use prices for electricity: Econometric analysis of surplus and pareto impacts. *The RAND Journal of Economics*, *25*(2), 263–283. https://doi.org/10.2307/2555830

Turnovsky, S. J., Shalit, H., & Schmitz, A. (1980). Consumer's surplus, price instability, and consumer welfare. *Econometrica: Journal of the Econometric Society*, 135–152. https://doi.org/10.2307/1912022

UK Office for National Statistics (UK ONS). (2021). Family spending workbook 1: Detailed expenditure and trends. https://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/expenditure/datasets/familyspendingworkbook1-detailedexpenditureandtrends

UL Services Group. (2022). Hourly wind and solar generation profiles (1980-2021). https://www.ercot.com/files/docs/2022/12/19/ERCOT_1980-2021_WindSolarGenProfiles_FINAL_public.pdf

Vayanos, D. (1999). Strategic trading and welfare in a dynamic market. *The Review of Economic Studies*, *66*(2), 219–254. https://doi.org/10.1111/1467-937X.00086

Vercammen, J. (1995). Hedging with commodity options when price distributions are skewed. *American Journal of Agricultural Economics*, *77*(4), 935–945. https://doi.org/10.2307/1243816

Wang, F., Xu, H., Xu, T., Li, K., Shafie-Khah, M., & Catalão, J. P. S. (2017). The values of market-based demand response on improving power system reliability under extreme circumstances. *Applied Energy*, *193*, 220–231. https://doi.org/10.1016/j.apenergy.2017.01.103

Waugh, F. V. (1944). Does the consumer benefit from price instability? *The Quarterly Journal of Economics*, *58*(4), 602–614. https://doi.org/10.2307/1884746

Weron, R., Bierbrauer, M., & Trück, S. (2004). Modeling electricity prices: Jump diffusion and regime switching. *Physica A: Statistical Mechanics and its Applications*, *336*(1), 39–48. https://doi.org/10.1016/j.physa.2004.01.008

Willems, B. (2005). Cournot competition, financial option markets, and efficiency. *UC Berkeley: Center for the Study of Energy Markets*. https://escholarship.org/uc/item/65h8p4sb

Winzer, C., Ramírez-Molina, H., Hirth, L., & Schlecht, I. (2024). Profile contracts for electricity retail customers. *Energy Policy*, *195*, 114358. https://doi.org/10.1016/j.enpol.2024.114358

Wolak, F. A. (2010). An experimental comparison of critical peak and hourly pricing: The powercentsdc program. *Department of Economics Stanford University*. https://web.stanford.edu/group/fwolak/cgi-bin/sites/default/files/files/An%20Experimental%20Comparison%20of%20Critical%20Peak%20and%20Hourly%20Pricing_March%202010_Wolak.pdf

Wolak, F. A. (2022). Long-term resource adequacy in wholesale electricity markets with significant intermittent renewables. *Environmental and Energy Policy and the Economy*, *3*(1), 155–220. https://doi.org/10.1086/717221

Wolak, F. A., & Hardman, I. H. (2022). *The future of electricity retailing and how we get there* (1st ed., Vol. 41). Springer International Publishing.

Woo, C.-K., Karimov, R. I., & Horowitz, I. (2004). Managing electricity procurement cost and risk by a local distribution company. *Energy Policy*, *32*(5), 635–645. https://doi.org/10.1016/S0301-4215(02)00317-8

Zhang, S., Fu, X., & Wang, X. (2012). Effects of option contracts on electricity markets: A cournot equilibrium analysis. *Asia-Pacific Power and Energy*, 1–5. https://doi.org/10.1109/appeec.2012.6307417

Zhou, D. P., Sahleh, M. A., & Tomlin, C. J. (2017). Hedging strategies for load-serving entities in wholesale electricity markets. *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. https://doi.org/10.1109/CDC.2017.8263665

Note: All URLs were last accessed February 26, 2025

Appendix A

Appendix to Chapter 1

Table A.1 shows the results for the fixed-effects regression model, including all time leads and lags. The regressions for the low-carbon technology (LCT) groups on the following page contain the consumption of all 5,904 households on fixed tariff as control (See Section 4). Regression results for narrower LCT-group-specific controls are presented in Figure A.1 and Table A.2 below.

Table A.1: Fixed effects regression results with country-wide controls for electricity consumption of customers on fixed rates

Regressors	**All customers**	**winter**	**spring**	**summer**	**fall**	**night (0-6)**	**morning (6-12)**	**afternoon (12-18)**	**evening (18-24)**
const	-0.002***	-0.001	-0.001	-0.000	-0.004***	-0.003**	-0.001	-0.001	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)
-8h	0.042***	0.017	0.026*	0.040**	0.057***	0.037***	0.011	0.004	0.005
	(0.007)	(0.011)	(0.016)	(0.016)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
-7.5h	0.014*	0.024*	0.012	0.027	0.009	-0.001	0.003	0.001	0.014
	(0.008)	(0.014)	(0.020)	(0.019)	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)
-7.0h	-0.000	-0.000	-0.008	-0.009	0.001	0.001	-0.002	-0.001	-0.004
	(0.008)	(0.014)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)
-6.5h	0.015*	0.008	0.024	-0.004	0.015	0.019	0.010	0.007	0.014
	(0.008)	(0.014)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)
-6h	0.003	0.009	0.005	-0.034*	0.000	-0.006	-0.003	0.005	0.037***
	(0.008)	(0.015)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)
-5.5h	0.001	0.004	0.023	0.009	-0.002	0.017	-0.009	0.015	0.019
	(0.008)	(0.015)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)
-5.0h	0.010	0.003	0.032	0.008	-0.000	0.009	0.016	-0.001	0.011
	(0.008)	(0.015)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)
-4.5h	0.005	0.012	-0.013	0.011	0.004	0.013	0.004	0.004	-0.008
	(0.008)	(0.015)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.014)	(0.012)
-4h	0.017**	0.015	0.001	0.014	0.013	0.024*	0.007	0.006	-0.001
	(0.008)	(0.015)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)

Table A.1 (continued)

Regres-sors	**All customers**	**winter**	**spring**	**sum-mer**	**fall**	**night (0-6)**	**morn-ing (6-12)**	**after-noon (12-18)**	**evening (18-24)**
-3.5h	0.002	0.014	-0.005	-0.007	0.007	0.003	0.011	0.022	-0.000
	(0.008)	(0.015)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)
-3.0h	0.003	0.006	-0.006	-0.003	0.004	0.019	0.011	0.004	0.000
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.013)	(0.012)
-2.5h	0.011	0.017	0.029	0.016	0.008	0.015	0.008	0.018	-0.004
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.013)	(0.012)
-2h	-0.000	0.003	0.008	0.002	0.000	0.003	0.009	0.010	-0.012
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
-1.5h	0.005	-0.000	-0.012	-0.001	-0.004	0.004	0.007	0.022	-0.013
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
-1.0h	0.002	0.012	0.029	-0.003	0.002	0.021	0.020	0.020	0.010
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
-0.5h	-0.018**	-0.006	-0.035*	0.003	-0.029**	-0.007	0.025*	-0.011	-0.022*
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
0h	-0.265***	-0.273***	-0.246***	-0.297***	-0.223***	-0.302***	-0.154***	-0.157***	-0.240***
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
0.5h	-0.030***	-0.030**	-0.035*	0.017	-0.029**	-0.060***	0.006	-0.020	0.046***
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
1.0h	-0.002	-0.007	-0.004	0.027	0.002	0.003	0.022*	0.022	0.026**
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
1.5h	0.013	0.010	-0.002	0.016	0.016	-0.011	-0.003	0.001	0.032**
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
2h	0.017**	0.008	0.019	0.026	0.005	0.030**	0.002	0.010	0.042***
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
2.5h	0.018**	0.021	0.001	0.019	0.012	0.019	0.024*	0.009	0.004
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
3.0h	0.012	0.001	0.052**	0.006	0.015	0.009	0.016	0.018	0.003
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.013)	(0.012)
3.5h	0.019**	0.018	0.025	0.026	0.028*	0.012	-0.001	0.007	0.016
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.013)	(0.012)
4h	0.018**	0.005	0.001	0.044**	0.019	0.011	-0.007	-0.004	0.019
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
4.5h	0.010	0.014	0.017	-0.003	0.003	0.005	0.001	-0.001	0.010
	(0.008)	(0.015)	(0.020)	(0.020)	(0.015)	(0.014)	(0.013)	(0.014)	(0.012)
5.0h	0.015*	0.026*	0.003	-0.011	0.010	-0.019	0.014	-0.012	0.017
	(0.008)	(0.014)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.014)	(0.012)
5.5h	0.009	0.004	-0.009	-0.009	0.003	-0.022	0.004	0.001	0.018
	(0.008)	(0.014)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.014)	(0.012)
6h	0.002	0.003	0.011	-0.013	-0.000	0.002	-0.022*	0.019	0.007
	(0.008)	(0.014)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.014)	(0.012)
6.5h	0.001	-0.001	0.013	-0.013	0.003	0.003	0.010	-0.005	-0.005
	(0.008)	(0.014)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)
7.0h	0.008	-0.000	-0.010	0.023	0.012	-0.001	-0.013	0.005	0.007
	(0.008)	(0.014)	(0.020)	(0.020)	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)
7.5h	0.003	-0.004	-0.003	-0.013	0.014	0.012	0.012	0.016	-0.009
	(0.008)	(0.014)	(0.020)	(0.019)	(0.014)	(0.014)	(0.013)	(0.013)	(0.012)
8h	0.008	0.033***	-0.005	-0.010	-0.006	0.015	0.035***	0.003	-0.001
	(0.007)	(0.011)	(0.016)	(0.016)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
delta_ln _watt_fixed	1.008*** (0.009)	1.226*** (0.016)	1.043*** (0.015)	0.838*** (0.021)	0.804*** (0.022)	0.568*** (0.031)	1.112*** (0.014)	1.055*** (0.009)	0.899*** (0.020)
Observa-tions	17,105	4,200	4,247	4,295	4,263	4,233	4,257	4,260	4,257
R2	0.484	0.657	0.570	0.354	0.362	0.434	0.628	0.784	0.415
Adjusted R2	0.483	0.655	0.567	0.348	0.357	0.429	0.625	0.782	0.410
Residual Std. Error	0.070 (df=17070)	0.063 (df=4165)	0.061 (df=4212)	0.054 (df=4260)	0.090 (df=4228)	0.098 (df=4198)	0.054 (df=4222)	0.042 (df=4225)	0.062 (df=4222)
F Statistic (omitting all regressors)	471.336*** (df=34; 17070)	235.063*** (df=34; 4165)	164.230*** (df=34; 4212)	68.553*** (df=34; 4260)	70.542*** (df=34; 4228)	94.636*** (df=34; 4198)	209.497*** (df=34; 4222)	451.391*** (df=34; 4225)	88.134*** (df=34; 4222)

Table A.1 (continued)

Regres-sors	**All customers**	**winter**	**spring**	**sum-mer**	**fall**	**night (0-6)**	**morn-ing (6-12)**	**after-noon (12-18)**	**evening (18-24)**
F Statistic (omitting time leads and lags)	58.241** (df=32; 17070)	21.934** (df=32; 4165)	10.757** (df=32; 4212)	6.900** (df=32; 4260)	13.258** (df=32; 4228)	18.206** (df=32; 4198)	9.938** (df=32; 4222)	11.865** (df=32; 4225)	11.228** (df=32; 4222)

Note:
*p<0.1;**p<0.05;***p<0.01

Table A.1 (continued)

Regressors	No LCT	At least one LCT	Inf. electric heating only	Inf. electric heat. + smart thermost.	EV only	Solar only	Battery only	EV + solar
const	-0.001	-0.002***	-0.001**	-0.001*	-0.002**	-0.003	-0.001	-0.003**
	(0.001)	(0.001)	(0.000)	(0.000)	(0.001)	(0.002)	(0.001)	(0.001)
-8h	-0.003	0.046***	0.011*	-0.003	0.060***	0.028	0.011	0.087***
	(0.013)	(0.007)	(0.006)	(0.006)	(0.010)	(0.022)	(0.018)	(0.018)
-7.5h	0.017	0.011	-0.005	0.014**	0.004	0.024	0.015	0.009
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-7.0h	-0.008	-0.001	-0.003	-0.012*	-0.002	0.004	0.033	0.007
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-6.5h	0.011	0.015*	0.002	0.003	0.012	0.010	-0.007	0.030
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-6h	0.023	0.003	0.004	0.008	0.010	-0.027	-0.026	0.020
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-5.5h	0.005	0.001	-0.002	-0.007	0.008	-0.006	0.017	-0.004
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-5.0h	0.004	0.011	-0.000	0.015**	0.010	0.009	0.003	0.027
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-4.5h	0.010	0.006	-0.002	-0.001	0.011	0.012	-0.001	-0.001
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-4h	-0.026	0.020**	0.010	-0.014*	0.029**	0.015	0.001	0.029
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-3.5h	-0.015	0.003	-0.003	0.002	-0.001	0.001	-0.018	0.011
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-3.0h	0.022	-0.000	0.002	0.014**	-0.008	0.014	0.021	-0.007
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-2.5h	0.016	0.010	0.009	0.006	0.005	0.020	0.026	0.021
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-2h	0.017	0.000	0.000	-0.001	-0.001	0.013	-0.010	0.001
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-1.5h	-0.018	0.006	-0.002	-0.005	0.003	0.011	0.024	0.014
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-1.0h	0.003	0.002	-0.000	0.001	-0.002	0.033	-0.005	-0.009
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
-0.5h	0.006	-0.020**	-0.014*	-0.008	-0.022*	-0.006	-0.039*	-0.044**
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
0h	-0.101***	-0.282***	-0.070***	-0.079***	-0.368***	-0.263***	-0.114***	-0.435***
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
0.5h	-0.024	-0.031***	-0.007	-0.028***	-0.018	-0.060**	-0.041*	-0.047**
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
1.0h	0.014	-0.004	-0.001	-0.002	-0.012	0.026	0.006	-0.001
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
1.5h	0.014	0.014	0.018**	0.000	0.002	0.050*	0.017	0.027
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
2h	0.006	0.017*	0.003	0.013*	0.008	0.043	0.030	0.027
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
2.5h	0.004	0.018**	0.001	0.005	0.016	0.046	0.013	0.033
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
3.0h	0.018	0.012	-0.005	0.000	0.007	0.021	0.056**	-0.005
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
3.5h	0.015	0.019**	0.003	0.003	0.023*	0.021	0.030	0.028
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
4h	0.016	0.019**	0.007	-0.004	0.031**	0.007	0.012	0.023
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
4.5h	-0.011	0.012	0.007	0.011	0.017	-0.007	0.003	0.023
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
5.0h	-0.002	0.017*	0.008	0.001	0.014	0.019	0.005	0.025
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
5.5h	0.014	0.009	-0.003	0.009	0.014	0.016	-0.009	0.020
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
6h	0.008	0.001	-0.011	-0.003	0.005	-0.011	0.024	-0.008

Regressors	No LCT	At least one LCT	Inf. electric heating only	Inf. electric heat. + smart thermost.	EV only	Solar only	Battery only	EV + solar
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
6.5h	0.018	-0.001	-0.005	-0.009	-0.002	-0.015	0.022	0.004
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
7.0h	-0.024	0.009	-0.001	-0.003	0.018	0.010	0.013	0.002
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
7.5h	0.005	0.002	-0.006	-0.002	0.005	-0.013	-0.010	-0.004
	(0.016)	(0.009)	(0.008)	(0.007)	(0.013)	(0.028)	(0.022)	(0.022)
8h	0.010	0.007	0.013**	0.005	0.017*	-0.048**	0.017	-0.012
	(0.013)	(0.007)	(0.006)	(0.006)	(0.010)	(0.022)	(0.018)	(0.018)
delta_ln_ watt_fixed	1.331*** (0.017)	0.965*** (0.009)	0.766*** (0.008)	0.716*** (0.008)	0.709*** (0.014)	2.176*** (0.030)	2.223*** (0.023)	1.183*** (0.024)
Observations	17,049	17,105	17,054	17,105	17,096	17,105	17,049	17,105
R2	0.286	0.455	0.335	0.354	0.266	0.245	0.361	0.193
Adjusted R2	0.285	0.454	0.334	0.353	0.264	0.243	0.360	0.191
Residual Std. Error	0.128 (df=17014)	0.073 (df=17070)	0.065 (df=17019)	0.059 (df=17070)	0.106 (df=17061)	0.233 (df=17070)	0.181 (df=17014)	0.183 (df=17070)
F Statistic (omitting all regressors)	200.905*** (df=34; 17014)	419.567*** (df=34; 17070)	252.189*** (df=34; 17019)	274.979*** (df=34; 17070)	181.800*** (df=34; 17061)	162.757*** (df=34; 17070)	282.433*** (df=34; 17014)	119.738*** (df=34; 17070)
F Statistic (omitting time leads and lags)	6.026** (df=32; 17014)	56.016** (df=32; 17070)	2.351** (df=32; 17019)	3.605** (df=32; 17070)	35.602** (df=32; 17061)	6.944** (df=32; 17070)	10.674** (df=32; 17014)	20.466** (df=32; 17070)
Note: *p<0.1; **p<0.05; ***p<0.01								

Table A.1 (continued)

Regressors	EV + battery	Solar + battery	EV + solar + battery
const	-0.002	-0.003	-0.003*
	(0.001)	(0.002)	(0.002)
-8h	0.029	0.065***	0.073***
	(0.019)	(0.023)	(0.021)
-7.5h	0.007	0.015	0.070***
	(0.023)	(0.028)	(0.026)
-7.0h	0.015	0.025	0.010
	(0.023)	(0.028)	(0.026)
-6.5h	0.031	0.026	0.021
	(0.023)	(0.028)	(0.026)
-6h	-0.006	-0.019	0.013
	(0.023)	(0.029)	(0.027)
-5.5h	0.015	-0.006	-0.005
	(0.023)	(0.029)	(0.027)
-5.0h	-0.004	0.014	0.010
	(0.023)	(0.029)	(0.027)
-4.5h	0.010	-0.008	-0.010
	(0.023)	(0.029)	(0.027)
-4h	0.015	0.005	0.030
	(0.023)	(0.029)	(0.027)

Regressors	EV + battery	Solar + battery	EV + solar + battery
-3.5h	0.011	-0.006	0.003
	(0.023)	(0.029)	(0.027)
-3.0h	0.015	-0.004	0.017
	(0.023)	(0.029)	(0.027)
-2.5h	0.005	0.047	0.018
	(0.023)	(0.029)	(0.027)
-2h	-0.004	0.024	0.016
	(0.023)	(0.029)	(0.027)
-1.5h	0.010	0.022	-0.001
	(0.023)	(0.029)	(0.027)
-1.0h	0.014	0.066**	0.036
	(0.023)	(0.029)	(0.027)
-0.5h	-0.044*	-0.035	-0.019
	(0.023)	(0.029)	(0.027)
0h	-0.266***	-0.357***	-0.561***
	(0.023)	(0.029)	(0.027)
0.5h	-0.066***	-0.046	-0.069**
	(0.023)	(0.029)	(0.027)
1.0h	0.013	0.051*	0.045*
	(0.023)	(0.029)	(0.027)
1.5h	0.016	0.032	0.053**
	(0.023)	(0.029)	(0.027)
2h	0.029	0.045	0.043
	(0.023)	(0.029)	(0.027)
2.5h	0.040*	0.037	0.074***
	(0.023)	(0.029)	(0.027)
3.0h	0.059**	0.086***	0.065**
	(0.023)	(0.029)	(0.027)
3.5h	0.028	0.009	0.049*
	(0.023)	(0.029)	(0.027)
4h	0.020	-0.012	-0.000
	(0.023)	(0.029)	(0.027)
4.5h	0.011	0.026	-0.003
	(0.023)	(0.029)	(0.027)
5.0h	0.016	0.023	0.027
	(0.023)	(0.029)	(0.027)
5.5h	0.015	-0.022	0.014
	(0.023)	(0.029)	(0.027)
6h	0.002	-0.003	-0.012
	(0.023)	(0.029)	(0.027)
6.5h	0.000	0.014	0.000
	(0.023)	(0.028)	(0.026)
7.0h	0.002	0.015	0.008
	(0.023)	(0.028)	(0.026)
7.5h	0.003	0.030	-0.014
	(0.023)	(0.028)	(0.026)
8h	0.004	-0.042*	-0.006
	(0.019)	(0.023)	(0.021)
delta_ln	1.357***	2.516***	1.987***
watt_fixed	(0.025)	(0.031)	(0.028)
Observations	17,101	17,105	17,105
R2	0.174	0.304	0.267
Adjusted R2	0.172	0.302	0.266
Residual Std. Error	0.193 (df=17066)	0.237 (df=17070)	0.220 (df=17070)
F Statistic (omitting all regressors)	105.696*** (df=34; 17066)	218.937*** (df=34; 17070)	183.060*** (df=34; 17070)
F Statistic (omitting time leads and lags)	16.329** (df=32; 17066)	15.527** (df=32; 17070)	28.579** (df=32; 17070)

Note: *p<0.1; **p<0.05; ***p<0.01

Figure A.1 shows the regression results for narrower LCT-group-specific controls while the above Table A.1 shows the regression results for the generic country-wide control groups that are presented in the results in section 1.5.

Figure A.1: Own-price elasticities with LCT-group-specific controls for electricity consumption of customers on fixed rates

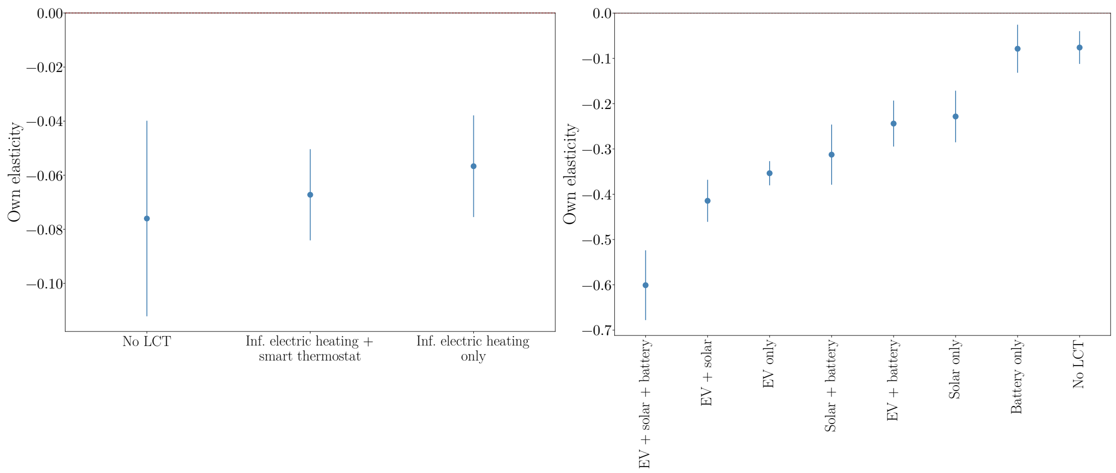

Figure A.2 shows the regression results for narrower LCT-group-specific controls while the above Table A.1 shows the regression results for the generic country-wide control groups that are presented in the results in section 1.5.

Table A.2: Fixed effects regression results with LCT-group-specific controls for electricity consumption of customers on fixed rates

Regressors	No LCT	At least one LCT	Inf. electric heating only	Inf. electric heating + smart thermostat	EV only	Solar only	Battery only	EV + solar
const	-0.002*	-0.002***	-0.001**	-0.001*	-0.002***	-0.004**	-0.002	-0.003**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)
-8h	0.020	0.048***	0.024***	0.007	0.069***	0.083***	0.045**	0.103***
	(0.015)	(0.007)	(0.008)	(0.007)	(0.011)	(0.023)	(0.022)	(0.019)
-7.5h	0.011	0.011	-0.007	0.011	-0.001	0.007	0.004	0.007
	(0.018)	(0.009)	(0.009)	(0.008)	(0.013)	(0.029)	(0.027)	(0.023)
-7.0h	-0.020	-0.001	-0.009	-0.017**	-0.006	-0.006	0.014	-0.003
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-6.5h	0.000	0.016*	-0.004	-0.000	0.007	-0.010	-0.014	0.021
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-6h	0.017	0.002	0.002	0.007	0.009	-0.017	-0.029	0.018
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-5.5h	0.005	0.002	0.000	-0.007	0.011	-0.003	0.012	-0.003
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-5.0h	0.008	0.012	0.002	0.019**	0.013	0.023	0.003	0.030
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-4.5h	0.019	0.006	0.003	0.003	0.015	0.028	0.013	0.006
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-4h	-0.017	0.019**	0.015	-0.011	0.032**	0.003	0.010	0.035
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-3.5h	-0.006	0.003	0.000	0.004	0.001	0.008	-0.003	0.016
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-3.0h	0.025	0.000	0.002	0.015*	-0.005	0.002	0.029	-0.004
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-2.5h	0.019	0.010	0.010	0.008	0.006	0.021	0.019	0.024
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-2h	0.025	0.001	0.005	0.002	0.001	0.026	0.000	0.006
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-1.5h	-0.006	0.007	0.005	0.004	0.011	0.021	0.051*	0.023
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-1.0h	0.025	0.004	0.011	0.012	0.008	0.062**	0.035	0.009
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
-0.5h	0.032*	-0.020**	-0.000	0.006	-0.009	0.028	0.004	-0.023
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
0h	-0.076***	-0.281***	-0.057***	-0.067***	-0.354***	-0.228***	-0.079***	-0.415***
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
0.5h	-0.008	-0.031***	0.003	-0.019**	-0.011	-0.053*	-0.021	-0.034
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
1.0h	0.019	-0.004	0.001	0.002	-0.010	0.034	0.016	0.005
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
1.5h	0.014	0.014	0.019**	0.001	0.005	0.057*	0.024	0.030
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
2h	-0.002	0.018**	-0.004	0.005	0.004	0.017	0.015	0.020
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
2.5h	-0.006	0.018**	-0.003	0.000	0.012	0.008	-0.011	0.022
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
3.0h	0.010	0.013	-0.008	-0.004	0.003	0.010	0.041	-0.015
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
3.5h	0.012	0.018**	-0.002	0.001	0.019	0.001	0.023	0.024
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
4h	0.019	0.019**	0.010	-0.003	0.033**	0.018	0.015	0.027
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)

Table A.2 (continued)

Regres-sors	**No LCT**	**At least one LCT**	**Inf. electric heating only**	**Inf. electric heating + smart thermo-stat**	**EV only**	**Solar only**	**Battery only**	**EV + solar**
4.5h	-0.012	0.013	0.006	0.011	0.014	-0.016	0.004	0.023
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
5.0h	-0.005	0.017*	0.006	-0.001	0.013	0.009	-0.001	0.026
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
5.5h	0.011	0.010	-0.002	0.008	0.014	0.027	-0.014	0.019
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
6h	0.008	0.001	-0.010	-0.003	0.003	0.013	0.024	-0.008
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
6.5h	0.012	-0.001	-0.006	-0.012	-0.004	-0.012	0.017	0.001
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
7.0h	-0.036**	0.009	-0.009	-0.008	0.014	-0.006	-0.001	-0.005
	(0.018)	(0.009)	(0.010)	(0.009)	(0.014)	(0.029)	(0.027)	(0.024)
7.5h	-0.001	0.002	-0.010	-0.004	0.005	-0.017	-0.020	-0.006
	(0.018)	(0.009)	(0.009)	(0.008)	(0.013)	(0.029)	(0.027)	(0.023)
8h	0.026*	0.008	0.022***	0.010	0.021*	-0.009	0.040*	-0.000
	(0.015)	(0.007)	(0.008)	(0.007)	(0.011)	(0.023)	(0.022)	(0.019)
delta_ln	0.011***	0.970***	0.065***	0.051***	0.111***	0.081***	0.104***	0.070***
watt_fixed	(0.001)	(0.010)	(0.003)	(0.002)	(0.005)	(0.003)	(0.004)	(0.004)
Observa-tions	16,951	17,105	17,054	17,105	17,096	15,590	17,049	17,105
R2	0.020	0.437	0.034	0.043	0.171	0.073	0.053	0.090
Adjusted R2	0.018	0.435	0.033	0.041	0.169	0.071	0.051	0.088
Residual Std. Error	0.150 (df=16916)	0.074 (df=17070)	0.078 (df=17019)	0.071 (df=17070)	0.113 (df=17061)	0.237 (df=15555)	0.220 (df=17014)	0.194 (df=17070)
F Statistic	9.924*** (df=34; 16916)	388.987*** (df=34; 17070)	17.852*** (df=34; 17019)	22.485*** (df=34; 17070)	103.435*** (df=34; 17061)	35.780*** (df=34; 15555)	28.165*** (df=34; 17014)	49.479*** (df=34; 17070)

Note:
*p<0.1; **p<0.05; ***p<0.01

Table A.2 (continued)

Regressors	**EV + battery**	**Solar + battery**	**EV + solar + battery**
const	-0.003*	-0.004*	-0.005**
	(0.002)	(0.002)	(0.002)
-8h	0.054***	0.112***	0.095***
	(0.021)	(0.027)	(0.030)
-7.5h	0.005	0.004	0.051
	(0.025)	(0.033)	(0.037)
-7.0h	0.001	0.003	-0.007
	(0.026)	(0.034)	(0.038)
-6.5h	0.024	0.010	0.027
	(0.026)	(0.034)	(0.038)
-6h	-0.010	-0.027	0.016
	(0.026)	(0.034)	(0.037)
-5.5h	0.014	-0.003	-0.022

Regressors	**EV + battery**	**Solar + battery**	**EV + solar + battery**
	(0.026)	(0.034)	(0.037)
-5.0h	0.005	0.023	0.007
	(0.026)	(0.034)	(0.037)
-4.5h	0.019	0.008	0.001
	(0.026)	(0.034)	(0.037)
-4h	0.025	0.021	0.034
	(0.026)	(0.034)	(0.037)
-3.5h	0.018	0.006	0.005
	(0.026)	(0.034)	(0.037)
-3.0h	0.014	-0.002	0.026
	(0.026)	(0.034)	(0.037)
-2.5h	0.006	0.052	0.018
	(0.026)	(0.034)	(0.038)
-2h	0.007	0.037	0.017
	(0.026)	(0.034)	(0.038)
-1.5h	0.029	0.049	0.013
	(0.026)	(0.034)	(0.038)
-1.0h	0.041	0.107***	0.013
	(0.026)	(0.034)	(0.039)
-0.5h	-0.021	0.015	0.031
	(0.026)	(0.034)	(0.039)
0h	-0.244***	-0.313***	-0.601***
	(0.026)	(0.034)	(0.039)
0.5h	-0.047*	-0.015	-0.043
	(0.026)	(0.034)	(0.039)
1.0h	0.017	0.064*	0.059
	(0.026)	(0.034)	(0.039)
1.5h	0.026	0.035	0.060
	(0.026)	(0.034)	(0.039)
2h	0.025	0.023	0.047
	(0.026)	(0.034)	(0.038)
2.5h	0.023	0.015	0.076**
	(0.026)	(0.034)	(0.038)
3.0h	0.049*	0.070**	0.047
	(0.026)	(0.034)	(0.037)
3.5h	0.025	0.003	0.038
	(0.026)	(0.034)	(0.037)
4h	0.029	-0.003	0.001
	(0.026)	(0.034)	(0.037)
4.5h	0.012	0.025	-0.006
	(0.026)	(0.034)	(0.037)
5.0h	0.016	0.018	0.006
	(0.026)	(0.034)	(0.037)
5.5h	0.012	-0.023	0.012
	(0.026)	(0.034)	(0.037)
6h	0.001	-0.002	-0.028
	(0.026)	(0.034)	(0.037)
6.5h	-0.002	0.007	-0.007
	(0.026)	(0.034)	(0.037)
7.0h	-0.009	-0.001	-0.018
	(0.026)	(0.034)	(0.037)
7.5h	-0.001	0.025	-0.011
	(0.026)	(0.033)	(0.037)
8h	0.031	-0.019	-0.043
	(0.021)	(0.027)	(0.030)
delta_ln_	0.005	-0.007**	-0.021***
watt_fixed	(0.004)	(0.003)	(0.002)
Observations	16,173	17,103	11,633
R2	0.033	0.028	0.073
Adjusted R2	0.031	0.026	0.070
Residual Std. Error	0.209 (df=16138)	0.280 (df=17068)	0.244 (df=11598)
F Statistic	16.242*** (df=34; 16138)	14.212*** (df=34; 17068)	26.810*** (df=34; 11598)
Note: *p<0.1; **p<0.05; ***p<0.01			

Appendix B

Appendix to Chapter 2

B.1 Derivation of the optimal hedge share h^*

In this section, I derive the expressions for the optimal hedge share h^* as given in equations (2.5) and (2.8). The derivations follow Gilbert (1985) and Cowan (2004). The first-order condition of the optimization problem in the first stage in equation (2.3) can be rewritten as (Losq, 1982):

$$E[V_b(\tilde{p} - \overline{p})] = E[V_b(f - \overline{p})] \tag{B.1.1}$$

One can approximate V_b by a first-order Taylor approximation of V_b about $(\overline{p}, \overline{\varepsilon})$.

$$V_b \approx \overline{V}_b + \overline{V}_{b\tilde{p}}(\tilde{p} - \overline{p}) + \overline{V}_{bb}h(\tilde{p} - \overline{p}) + \overline{V}_{b\tilde{\varepsilon}}(\tilde{\varepsilon} - \overline{\varepsilon}) \tag{B.1.2}$$

Plugging equation (B.1.2) into (B.1.1) leads to:

$$\overline{V}_{b\tilde{p}}E[(\tilde{p} - \overline{p})^2] + \overline{V}_{bb}hE[(\tilde{p} - \overline{p})^2] + \overline{V}_{b\tilde{\varepsilon}}E[(\tilde{\varepsilon} - \overline{\varepsilon})(\tilde{p} - \overline{p})] = \overline{V}_b(f - \overline{p})$$

where $E[(\tilde{\varepsilon} - \overline{\varepsilon})] = 0$ and $E[(\tilde{p} - \overline{p})] = 0$ because $E[\tilde{p}] = \overline{p}$ and $E[\tilde{\varepsilon}] = \overline{\varepsilon}$.

With $\sigma^2_{\tilde{p}} = E[(\tilde{p} - \overline{p})^2]$ and $\sigma_{\tilde{p}\tilde{\varepsilon}} = E[(\tilde{\varepsilon} - \overline{\varepsilon})(\tilde{p} - \overline{p})]$, one can solve for the absolute optimal hedge quantity h^*:

$$h^* = -\frac{\overline{V}_{b\tilde{p}}}{\overline{V}_{bb}} - \frac{\overline{V}_{b\tilde{\varepsilon}}}{\overline{V}_{bb}}\frac{\sigma_{\tilde{p}\tilde{\varepsilon}}}{\sigma^2_{\tilde{p}}} + \frac{\overline{V}_b}{\overline{V}_{bb}}\frac{(f - \overline{p})}{\sigma^2_{\tilde{p}}} \tag{B.1.3}$$

As discussed in Section 2.3.1, the last term of the above equation equals zero since I assume that households do not hedge for speculative reasons, i.e., households believe that the forward hedge price f equals the mean dynamic price $\overline{p}$.

In the next step, I aim to express the optimal hedge quantity h^* in equation (B.1.3) as a share of the baseline consumption level $\hat{x}^* = x^*(\overline{p}, b, \overline{\varepsilon})$. At baseline, Roy's identity is given as $\hat{x}^* = -\frac{\overline{V}_{\tilde{p}}}{\overline{V}_b}$.[1] Following Turnovsky et al. (1980) and Gilbert (1985), one can differentiate Roy's identity at baseline with respect to b, which leads to

$$\begin{aligned} \overline{V}_{b\tilde{p}} &= -\overline{V}_{bb} - \overline{V}_b \frac{\partial \hat{x}^*}{\partial b} \\ \overline{V}_{b\tilde{p}} &= (\theta - \eta)\hat{x}^* \frac{\overline{V}_b}{b} \end{aligned} \tag{B.1.4}$$

$\eta = \frac{\partial \hat{x}^*}{\partial b}\frac{b}{\hat{x}^*}$ is the income elasticity of electricity demand, and $\theta = -\frac{\overline{V}_{bb}}{\overline{V}_b} b$ is the coefficient of relative risk aversion at baseline consumption. Inserting equation (B.1.4) into (B.1.3) gives

$$h^* = (1 - \frac{\eta}{\theta})\hat{x}^* - \frac{\overline{V}_{b\tilde{\varepsilon}}}{\overline{V}_{bb}} \frac{\sigma_{\tilde{p}\tilde{\varepsilon}}}{\sigma^2_{\tilde{p}}} \tag{B.1.5}$$

For the simulation of optimal hedge shares, I assume that the following CES indirect utility function can describe households' consumption decisions.

$$V(\tilde{p}, b, \tilde{\varepsilon}, f, h) = \frac{1}{1-\theta}[b + (\tilde{p} - f)h^*]^{1-\theta}(\tilde{\varepsilon}\tilde{p}^{1-\alpha} + 1)^{\frac{1-\theta}{\alpha-1}} \tag{B.1.6}$$

Applying Roy's identity, the Marshallian demand functions are given as

$$\begin{aligned} x^* &= \frac{[b + (\tilde{p} - f)h^*]\tilde{\varepsilon}\tilde{p}^{-\alpha}}{(\tilde{\varepsilon}\tilde{p}^{1-\alpha} + 1)} \qquad \text{(B.1.7)} \\ y^* &= \frac{[b + (\tilde{p} - f)h^*]}{(\tilde{\varepsilon}\tilde{p}^{1-\alpha} + 1)} \end{aligned}$$

The first derivative of $V = V(\tilde{p}, b, \tilde{\varepsilon}, f, h)$ with respect to b is

$$V_b = [b + (\tilde{p} - f)h^*]^{-\theta}(\tilde{\varepsilon}\tilde{p}^{1-\alpha} + 1)^{\frac{1-\theta}{\alpha-1}} \tag{B.1.8}$$

[1] In general, Roy's identity with a forward hedge is given as $x^* - h = -\frac{V_{\tilde{p}}}{V_b}$ (Gilbert, 1985).

Taking the derivative of V_b with respect to b and $\tilde{\varepsilon}$, respectively, gives

$$V_{bb} = -\theta[b + (\tilde{p} - f)h^*]^{-\theta-1}(\tilde{\varepsilon}\tilde{p}^{1-\alpha} + 1)^{\frac{1-\theta}{\alpha-1}} \tag{B.1.9}$$

$$V_{b\tilde{\varepsilon}} = [b + (\tilde{p} - f)h^*]^{-\theta}\frac{1-\theta}{\alpha-1}(\tilde{\varepsilon}\tilde{p}^{1-\alpha} + 1)^{\frac{1-\theta}{\alpha-1}-1}\tilde{p}^{1-\alpha} \tag{B.1.10}$$

Evaluating equations (B.1.9) and (B.1.10) at baseline (with $\tilde{p} = \overline{p} = f$ and $\tilde{\varepsilon} = \overline{\varepsilon}$) leads to

$$\frac{\overline{V}_{b\tilde{\varepsilon}}}{\overline{V}_{bb}} = \left(1 - \frac{1}{\theta}\right)\frac{1}{\alpha-1}\hat{x}^*\frac{\overline{p}}{\overline{\varepsilon}} \tag{B.1.11}$$

Moreover, from equation (B.1.7) one can derive the income elasticity of electricity demand at baseline $\eta = \frac{\partial\hat{x}^*}{\partial b}\frac{b}{\hat{x}^*} = 1$. The optimal hedge equation (B.1.5) can then be rewritten as

$$h^* = \left(1 - \frac{1}{\theta}\right)\hat{x}^* + \left(1 - \frac{1}{\theta}\right)\frac{1}{1-\alpha}\hat{x}^*\frac{\overline{p}}{\overline{\varepsilon}}\frac{\sigma_{\tilde{p}\tilde{\varepsilon}}}{\sigma^2_{\tilde{p}}} \tag{B.1.12}$$

$$\frac{h^*}{\hat{x}^*} = \left(1 - \frac{1}{\theta}\right) * \left(1 + \frac{1}{1-\alpha}\rho\frac{cv(\tilde{\varepsilon})}{cv(\tilde{p})}\right)$$

which equals the optimal hedge share in equation (2.8). $\sigma_{\tilde{p}}$ and $\sigma_{\tilde{\varepsilon}}$ denote the standard deviations of $\tilde{p}$ and $\tilde{\varepsilon}$. Moreover, $cv(\tilde{p}) = \frac{\sigma_{\tilde{p}}}{\overline{p}}$ and $cv(\tilde{\varepsilon}) = \frac{\sigma_{\tilde{\varepsilon}}}{\overline{\varepsilon}}$ are the coefficients of variations of $\tilde{p}$ and $\tilde{\varepsilon}$, respectively.

B.2 Derivation of the price premia

In this appendix, I derive the expressions for the price premia g_p and g_f as defined in equations (2.12) and (2.14), respectively.

Derivation of g_p:

g_p equals the percentage increase in fixed price $\overline{p}$ that makes the household indifferent in expectation between the fixed tariff and the unhedged dynamic tariff with stochastic price $\tilde{p}$ (Gilbert, 1985).

$$E[V((1+g_p)\overline{p}, b, \tilde{\varepsilon})] = E[V(\tilde{p}, b, \tilde{\varepsilon})] \tag{B.2.1}$$

When the household is on the unhedged dynamic tariff or the fixed tariff, the household does not hedge, i.e., $h = 0$. A first-order Taylor approximation of the left-hand side of equation (B.2.1) yields

$$\begin{aligned} V((1+g_p)\overline{p}, b, \tilde{\varepsilon}) &\approx \overline{V} + \overline{V}_{\tilde{p}} * E[(1+g_p)\overline{p} - \overline{p}] \\ &\approx \overline{V} + \overline{V}_{\tilde{p}} * g_p \overline{p} \end{aligned}$$

A second-order Taylor approximation of the right-hand side of equation (B.2.1) about $(\overline{p}, \overline{\varepsilon})$ yields

$$\begin{aligned} V(\tilde{p}, b, \tilde{\varepsilon}) \approx{}& \overline{V} + \overline{V}_{\tilde{p}} E[\tilde{p} - \overline{p}] + \overline{V}_{\tilde{\varepsilon}} E[\tilde{\varepsilon} - \overline{\varepsilon}] \\ &+ \frac{1}{2}\overline{V}_{\tilde{p}\tilde{p}} E[(\tilde{p} - \overline{p})^2] + \frac{1}{2}\overline{V}_{\tilde{\varepsilon}\tilde{\varepsilon}} E[(\tilde{\varepsilon} - \overline{\varepsilon})^2] \\ &+ \overline{V}_{\tilde{\varepsilon}\tilde{p}} E[(\tilde{\varepsilon} - \overline{\varepsilon})(\tilde{p} - \overline{p})] \\ \approx{}& \overline{V} + \frac{1}{2}(\overline{V}_{\tilde{p}\tilde{p}} \sigma^2_{\tilde{p}} + \overline{V}_{\tilde{\varepsilon}\tilde{\varepsilon}} \sigma^2_{\tilde{\varepsilon}}) + \overline{V}_{\tilde{\varepsilon}\tilde{p}} \sigma_{\tilde{\varepsilon}\tilde{p}} \end{aligned}$$

The term $\overline{V}_{\tilde{\varepsilon}\tilde{\varepsilon}}$ above can be neglected for comparing the relative welfare between the two price regimes. $\overline{V}_{\tilde{\varepsilon}\tilde{\varepsilon}}$ is not affected when moving from fixed to dynamic prices. Therefore, it does not affect the relative welfare difference between the fixed and dynamic prices measured by g_p (Gilbert, 1985). Inserting the Taylor approximations into equation (B.2.1) and solving for g_p leads to

$$g_p = \frac{1}{2}\frac{\overline{V}_{\tilde{p}\tilde{p}}}{\overline{V}_{\tilde{p}}\overline{p}}\sigma^2_{\tilde{p}} + \frac{\overline{V}_{\tilde{\varepsilon}\tilde{p}}}{\overline{V}_{\tilde{p}}\overline{p}}\sigma_{\tilde{p}\tilde{\varepsilon}} \tag{B.2.2}$$

Using the indirect CES utility function (2.10) for unhedged households ($h^* = 0$), the relevant derivatives can be derived as

$$\begin{aligned} V_{\tilde{p}} &= \frac{1}{\alpha - 1} b^{1-\theta} (\tilde{\varepsilon}\tilde{p}^{1-\alpha} + 1)^{\frac{1-\theta}{\alpha-1} - 1} (1-\alpha)\tilde{\varepsilon}\tilde{p}^{-\alpha} \\ &= -V_b x^* \end{aligned} \tag{B.2.3}$$

$$\begin{aligned} V_{\tilde{\varepsilon}} &= \frac{1}{\alpha - 1} b^{1-\theta} (\tilde{\varepsilon}\tilde{p}^{1-\alpha} + 1)^{\frac{1-\theta}{\alpha-1} - 1} \tilde{p}^{(1-\alpha)} \\ V_{\tilde{\varepsilon}\tilde{p}} &= -b^{1-\theta} (\tilde{\varepsilon}\tilde{p}^{1-\alpha} + 1)^{\frac{1-\theta}{\alpha-1} - 1} \tilde{p}^{-\alpha} [\frac{\beta^u}{1-\alpha} + 1 - \hat{s}] \\ &= \frac{V_{\tilde{p}}}{\tilde{\varepsilon}} [\frac{\beta^u}{1-\alpha} + 1 - \hat{s}] \end{aligned} \tag{B.2.4}$$

Moreover, Turnovsky et al. (1980) and Cowan (2004) show that differentiating Roy's identity $V_{\tilde{p}} = -\hat{x}^* V_b$ w.r.t. $\tilde{p}$ results in

$$V_{\tilde{p}\tilde{p}} = -\frac{\partial \hat{x}^*}{\partial \tilde{p}} V_b - \hat{x}^* V_{b\tilde{p}} \tag{B.2.5}$$

Inserting Roy's identity at baseline $V_b = -\frac{V_{\tilde{p}}}{\tilde{x}^*}$ and equation (B.1.4) into the above equation (B.2.5) gives the absolute value of Turnovsky et al.'s (1980) coefficient of relative price risk aversion

$$\frac{\overline{V}_{\tilde{p}\tilde{p}}}{\overline{V}\tilde{p}} \overline{p} = \hat{\gamma} + \hat{s}(\theta - \eta) \tag{B.2.6}$$

Plugging equations (B.2.3), (B.2.4), and (B.2.6) into equation (B.2.2) and evaluating them at baseline leads to

$$\begin{aligned} g_p &= \frac{1}{2} \frac{\overline{V}_{\tilde{p}\tilde{p}}}{\overline{V}_{\tilde{p}}} \overline{p} \frac{\sigma^2_{\tilde{p}}}{\overline{p}^2} + \frac{\sigma_{\tilde{p}}}{\overline{p}} \frac{\sigma_{\tilde{\varepsilon}}}{\overline{\varepsilon}} \sigma_{\tilde{p}\tilde{\varepsilon}} [\frac{\beta^u}{1-\alpha} + 1 - \hat{s}] \\ &= \frac{1}{2} [\hat{\gamma} + \beta^u] cv(\tilde{p})^2 + \rho cv(\tilde{p}) cv(\tilde{\varepsilon}) [\frac{\beta^u}{1-\alpha} + 1 - \hat{s}] \end{aligned}$$

The above equals the price premium for the unhedged dynamic tariff in equation (2.12). $\beta^u = \frac{\partial \overline{V}_b}{\partial \tilde{p}} \frac{\overline{p}}{\overline{V}_b} = (\theta - 1)s$ is the price elasticity of the marginal utility of income at baseline for unhedged households.

Derivation of g_f:

g_f denotes the percentage increase in fixed price $\overline{p}$ that makes the household indifferent in expectation between the fixed tariff and the optimally hedged dynamic tariff (Gilbert, 1985).

$$E[V((1+g_f)\overline{p}, b, \tilde{\varepsilon})] = E[V(\tilde{p}, b + (p-f)h^*, \tilde{\varepsilon})] \tag{B.2.7}$$

A second-order Taylor approximation of the RHS of equation (B.2.7) about $(\overline{p}, \overline{\varepsilon})$ yields

$$\begin{aligned} V(\tilde{p}, b, \tilde{\varepsilon}) &\approx \overline{V} + \overline{V}_{\tilde{p}}E[\tilde{p}-\overline{p}] + \overline{V}_b E[\tilde{p}-\overline{p}]h^* + \overline{V}_{\tilde{\varepsilon}}E[\tilde{\varepsilon}-\overline{\varepsilon}] \\ &\quad + \frac{1}{2}\overline{V}_{\tilde{p}\tilde{p}}E[(\tilde{p}-\overline{p})^2] + \frac{1}{2}\overline{V}_{b\tilde{p}}E[(\tilde{p}-\overline{p})^2]h^* + \frac{1}{2}\overline{V}_{\tilde{\varepsilon}\tilde{\varepsilon}}E[(\tilde{\varepsilon}-\overline{\varepsilon})^2] \\ &\quad + \overline{V}_{\tilde{\varepsilon}\tilde{p}}E[(\tilde{p}-\overline{p}) * (\tilde{\varepsilon}-\overline{\varepsilon})] \\ &\approx \overline{V} + \frac{1}{2}(\overline{V}_{\tilde{p}\tilde{p}}\sigma^2_{\tilde{p}} + \overline{V}_{\tilde{\varepsilon}\tilde{\varepsilon}}\sigma^2_{\tilde{\varepsilon}}) + \overline{V}_{\tilde{\varepsilon}\tilde{p}}\sigma_{\tilde{p}\tilde{\varepsilon}} + \frac{1}{2}\overline{V}_{b\tilde{p}}h^*\sigma^2_{\tilde{p}} \end{aligned}$$

Following analogous steps as for g_p, g_f is given as

$$g_f = \frac{1}{2}\frac{\overline{V}_{\tilde{p}\tilde{p}}}{\overline{V}_{\tilde{p}}\overline{p}}\sigma^2_{\tilde{p}} + \frac{\overline{V}_{\tilde{\varepsilon}\tilde{p}}}{\overline{V}_{\tilde{p}}\overline{p}}\sigma_{\tilde{p}\tilde{\varepsilon}} + \frac{1}{2}\frac{\overline{V}_{b\tilde{p}}}{\overline{V}_{\tilde{p}}\overline{p}}h^*\sigma^2_{\tilde{p}} \tag{B.2.8}$$

The first two terms of the above expression equal the ones for g_p in equation (B.2.2). However, when optimally hedged, Roy's identity changes to (Gilbert, 1985)

$$V_{\tilde{p}} = -V_b(x^* - h^*) \tag{B.2.9}$$

Differentiating expression (B.2.9) with respect to b yields

$$\begin{aligned} V_{b\tilde{p}} = V_{\tilde{p}b} &= -V_{bb}(x^* - h^*) - V_b\frac{\partial x^*}{\partial b} \\ &= -V_{bb}x^* - \overline{V}_{bb}\frac{\overline{V}_b}{\overline{V}_{bb}b}\frac{\partial x^*}{\partial b}\frac{b}{x^*}x^* + V_{bb}h^* \\ &= -V_{bb}(\frac{\theta-1}{\theta})x^* + V_{bb}h^* \end{aligned}$$

Evaluating the above equation at baseline with $\eta = 1$ and using the optimal hedge quantity $h^* = (1-\frac{1}{\theta})\hat{x}^*\Big(1 + \frac{1}{1-\alpha}\rho\frac{cv(\tilde{\varepsilon})}{cv(\tilde{p})}\Big)$ gives

$$\overline{V}_{b\tilde{p}} = -V_{bb}(\frac{\theta-1}{\theta})\hat{x}^* + V_{bb}(\frac{\theta-1}{\theta})\hat{x}^* + V_{bb}(\frac{\theta-1}{\theta})\hat{x}^*\Big(\frac{1}{1-\alpha}\rho\frac{cv(\tilde{\varepsilon})}{cv(\tilde{p})}\Big) \tag{B.2.10}$$

$$= -\frac{\beta^u}{1-\alpha}\frac{\overline{V}_b}{\overline{p}}\rho\frac{cv(\tilde{\varepsilon})}{cv(\tilde{p})} \tag{B.2.11}$$

where $\beta^u = \overline{V}_{b\tilde{p}} \frac{p}{\overline{V}_b} = (\theta - 1)s$ is the price-elasticity of marginal utility of income for the unhedged household (when $h^* = 0$). As shown above, one can derive an expression for $V_{\tilde{p}\tilde{p}}$ from Roy's identity in equation (B.2.9) as

$$\begin{aligned} V_{\tilde{p}\tilde{p}} &= -x^*_{\tilde{p}} V_b - V_{b\tilde{p}}(x^* - h^*) \\ &= -\gamma \frac{x^*}{p} V_b + V_{b\tilde{p}} \frac{V_{\tilde{p}}}{V_b} \end{aligned}$$

Evaluating the above expression at baseline and using (B.2.11), one gets

$$\overline{V}_{\tilde{p}\tilde{p}} = \gamma \frac{\overline{V}_{\tilde{p}}}{\overline{p}} - \frac{\beta^u}{1-\alpha} \rho \frac{cv(\tilde{\varepsilon})}{cv(\tilde{p})} \frac{\overline{V}_{\tilde{p}}}{\overline{p}} \tag{B.2.12}$$

where I used that at baseline, Roy's identity is $\hat{x}^* = -\frac{\overline{V}_{\tilde{p}}}{\overline{V}_b}$. Hence,

$$\frac{\overline{V}_{\tilde{p}\tilde{p}}}{\overline{V}_{\tilde{p}} \overline{p}} \sigma^2_{\tilde{p}} = \gamma cv(\tilde{p})^2 - \frac{\beta^u}{1-\alpha} \rho cv(\tilde{p}) cv(\tilde{\varepsilon}) \tag{B.2.13}$$

From equation (B.2.11) it also follows that

$$\begin{aligned} \frac{\overline{V}_{b\tilde{p}}}{\overline{V}_{\tilde{p}} \overline{p}} h^* \sigma^2_{\tilde{p}} &= -\frac{\beta^u}{1-\alpha} \frac{\overline{V}_b}{\overline{V}_{\tilde{p}}} \rho cv(\tilde{\varepsilon}) cv(\tilde{p}) h^* \\ &= \frac{\beta^u}{1-\alpha} \rho cv(\tilde{\varepsilon}) cv(\tilde{p}) \frac{h^*}{\hat{x}^*} \end{aligned} \tag{B.2.14}$$

Moreover, one also needs to derive expressions for $\overline{V}_{\tilde{\varepsilon}\tilde{p}}$ when optimally hedged. Following the same steps as for g_p above, one can show that for the optimally hedged tariff, $\overline{V}_{\tilde{\varepsilon}\tilde{p}}$ is equivalent to equations (B.2.4) for the unhedged dynamic tariff when evaluated at baseline. Therefore, one can plug equations (B.2.4), (B.2.13), and (B.2.14) into equation

(B.2.8) to obtain

$$
\begin{aligned}
g_f &= \frac{1}{2}\gamma cv(\tilde{p})^2 + \rho cv(\tilde{p})cv(\tilde{\varepsilon})[\frac{\beta^u}{1-\alpha} + 1 - \hat{s}] \\
&\quad - \frac{1}{2}\frac{\beta^u}{1-\alpha}\rho cv(\tilde{p})cv(\tilde{\varepsilon}) + \frac{1}{2} * \frac{\beta^u}{1-\alpha}\rho cv(\tilde{p})cv(\tilde{\varepsilon})\frac{h^*}{\hat{x}^*} \\
&= \underbrace{\frac{1}{2}[\hat{\gamma} + \beta^u]cv(\tilde{p})^2 + \rho cv(\tilde{p})cv(\tilde{\varepsilon})[\frac{\beta^u}{1-\alpha} + 1 - \hat{s}]}_{g_p} - \frac{1}{2}\beta^u cv(\tilde{p})^2 \\
&\quad - \frac{1}{2}\frac{\beta^u}{1-\alpha}\rho cv(\tilde{p})cv(\tilde{\varepsilon})\left(1 - \frac{h^*}{\hat{x}^*}\right) \\
&= g_p - \frac{1}{2}\beta^u cv(\tilde{p})^2 - \frac{1}{2}\frac{\beta^u}{1-\alpha}\rho cv(\tilde{p})cv(\tilde{\varepsilon})\left(1 - \frac{h^*}{\hat{x}^*}\right)
\end{aligned}
$$

This expression equals the price premium for the optimally hedged tariff in equation (2.14).

B.3 Robustness checks for different hedge time segments

In this appendix, I analyze the effect of choosing different hedge time segments on optimal hedge shares and bill volatility. As explained in Section 2.4, I define a hedge time segment for every half hour per day separately for weekdays and weekends in my main specification. For instance, 8-8:30 am on weekdays is a time segment. 8-8:30 am on weekends is a different time segment. In my main specification, I end up with 96 distinct time segments for all 48 half-hourly periods on weekends and weekdays, respectively. I can define even

Table B.1: Description of different hedge time segments

Time segment definition	**Description**	**Number of segments per customer**
Season and half hour	customer-season-weekend-half hour	384
Main specification	customer-weekend-half hour	96
Season and time of day	customer-season-weekend-time of day	32
Time of day	customer-weekend-time of day	8
Season	customer-season-weekend	8
None	customer	1

more fine-grained time segments by differentiating the main specification by season as shown in Table B.1.[2] For instance, the specification “Season and half an hour” defines 8-8:30 am on weekdays in February as a time segment and 8-8:30 am on weekdays in June as another segment. As Table B.1 highlights, this leads to 384 (4*96) segments per customer.

The specification “Season and time of day” leads to less fine-grained time segments. This specification differentiates by season, but groups periods in broader time of day categories.[3] For instance, 8-8:30 am on weekdays in February would fall in the winter-morning-weekday time segment. This specification results in 32 different time segments.

The specification “Time of day” only differentiates between time of days and weekday/weekends. Thus, 8-8:30 am on weekdays in February would be the same time segment as 8-8:30 am on weekdays in June.

In contrast “Season” differentiates by season and weekday/weekends but not by time of day. In this specification, 8-8:30 am on weekdays in February is the same segment as 20-20:30 am on weekdays in March but a different segment than 8-8:30 am on weekdays

[2]I define “winter” as January to March, “spring” as April to June, “summer” as July to September, and “autumn” as October to December.

[3]I define “night” as midnight to 6:00, “morning” as 6:00 to 12:00, “afternoon” as noon to 18:00, and evening as 18:00 to 24:00.

in June. The specifications “Time of day” and “Season” lead to 8 segments per customer, respectively.

Finally, the specification “None” is the less fine-grained as it counts all time intervals into the same time segment. This assumes that customers buy the same quantity forward for every time interval.

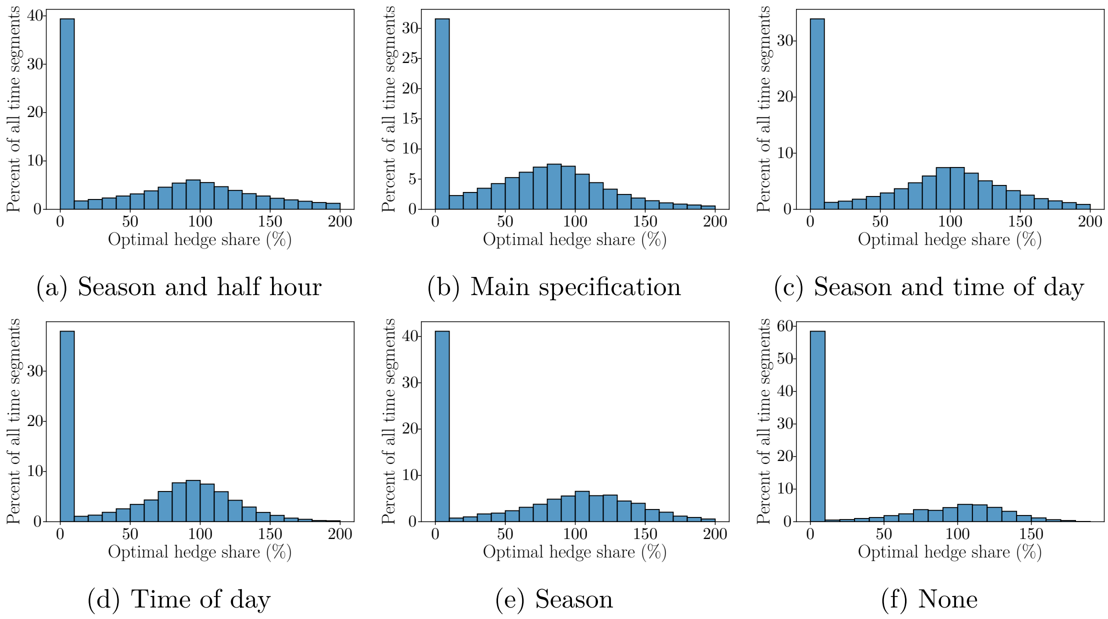

Figure B.1: Distribution of optimal hedge shares across all customers by hedge time segment specification

Figure B.1 highlights that the distribution of optimal hedge shares across all customers is very similar for the different hedge time segment specifications. The average optimal hedge share ranges between 52% (“None”) and 70% (“Season and time of day”). The less fine-grained the hedge specification, the smaller is the share of very large hedge shares (greater 150%) and the larger is the share of hedge shares that I set to zero.

Figure B.2 shows the coefficients of variations of monthly electricity bills for unhedged and optimally hedged tariffs across mean optimal hedge shares per household for the different hedge time segment specifications. This Figure reveals an even clearer trend in mean optimal hedge ratios. The less fine-grained the time segments, the more dispersed are the mean optimal hedge ratios between customers. Without any time segments (“None”) more than 40% of households do not hedge at all while a substantial share of households choos a mean optimal hedge share larger than 100%.

As expected, less fine-grained time segments are less effective in reducing bill volatility. Without any time segments ("None") the optimally hedged tariff tends to reduce the coefficient of variation of monthly bills only by 3% relative to the unhedged tariff. The fine-grained main specification leads to a much larger reduction in bill volatility of 19% when moving from the unhedged to the optimally hedged tariff. Interestingly, the even more fine-grained specification "Season and half hour" achieves only a smaller reduction in bill volatility of 14%. Hence, while an increase in granularity of time segments seems to overall reduce bill volatility, too granular specifications seem to raise volatility.

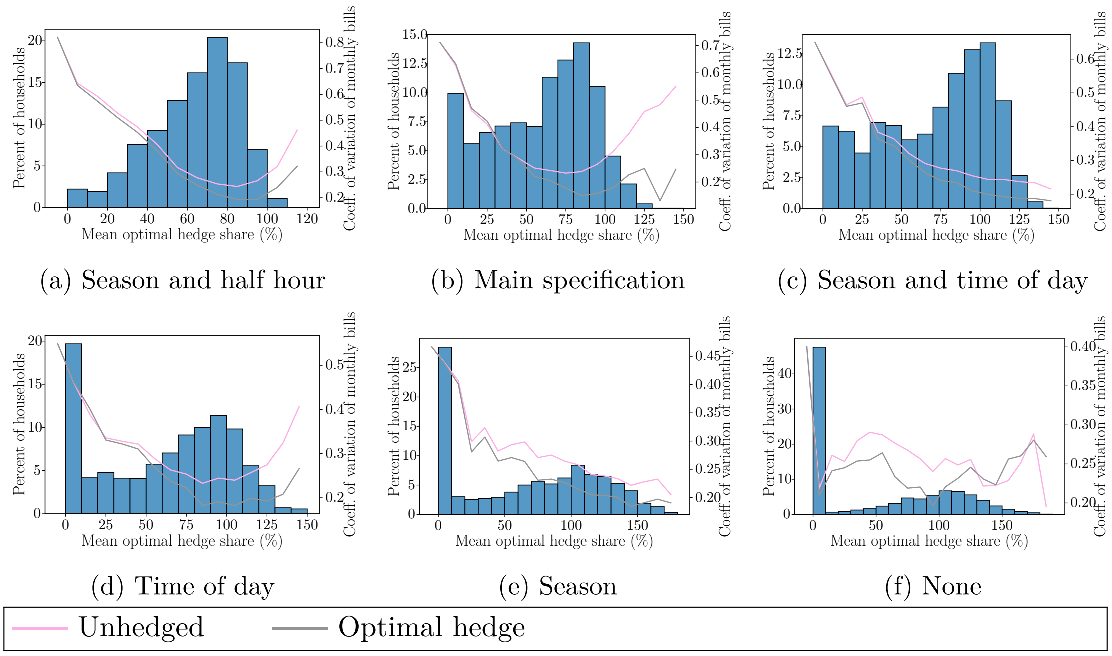

(a) Season and half hour (b) Main specification (c) Season and time of day

(d) Time of day (e) Season (f) None

Figure B.2: Distribution of average optimal hedge shares by customer and coefficients of variation of monthly bills by hedge time segments

B.4 Dynamic prices with removed price cap

This appendix explains how Octopus Energy calculates dynamic prices and provides descriptive statistics for day-ahead and dynamic prices. Dynamic prices are tied to day-ahead electricity prices. Figure B.3 shows the distribution of the half-hourly electricity day-ahead prices for the Great Britain (GB) price zone during the sample period obtained from EPEX SPOT (2023). Day-ahead prices are relatively low and rarely exceed 20 p/kWh. However, while overall volatility is small, there are a few significant price spike hours, in which day-ahead prices are very high. Figure B.4 reveals that day-ahead prices remain on average roughly stable over time during the study period.

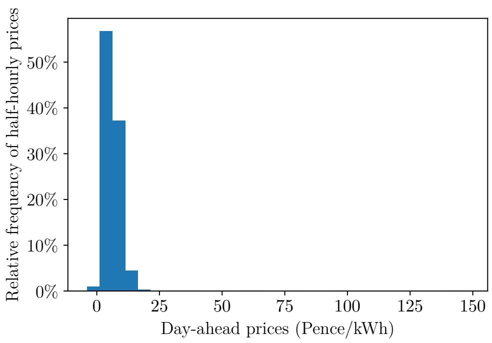

Figure B.3: Relative distribution of day-ahead electricity prices in the Great Britain price zone (EPEX SPOT, 2023)

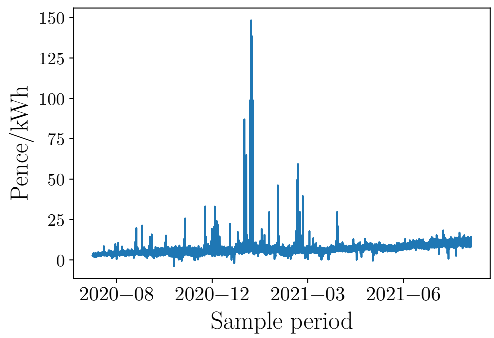

Figure B.4: Half-hourly day-ahead electricity prices in the Great Britain price zone (EPEX SPOT, 2023)

Dynamic prices are based on day-ahead prices. In addition, they contain distribution charges and a peak-time premium. For every half-hourly interval, the day-ahead price is multiplied by a distribution charge multiplier that ranges from 2 to 2.4, depending on the grid supply area the household is located in. Between 4 pm and 7 pm, a peak-time premium is added that ranges from 11p to 14p, depending on the grid supply area. Afterward, VAT is added. The resulting price is the dynamic price for the respective half-hourly interval unless it exceeds the price cap of 35 p/kWh. If the calculated price exceeds the price cap, the dynamic price is set to 35 p/kWh (Octopus Energy, 2019).

For 0.3 percent of the half-hourly intervals, households received weakly negative prices due to excess supply in the day-ahead market. I exclude these rare weakly negative price events since the model in Section 2.3 only applies to positive prices.

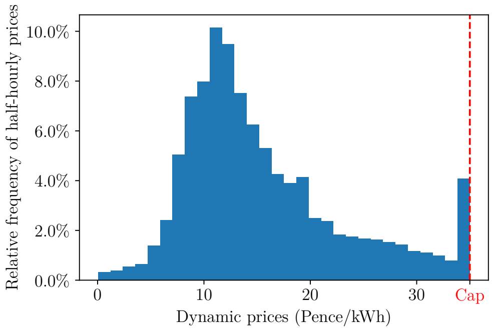

Figure B.5: Relative distribution of dynamic prices with price cap

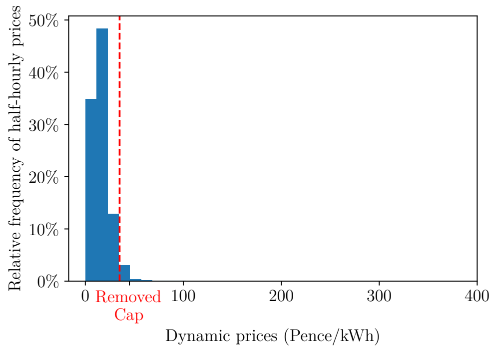

Figure B.6: Relative distribution of dynamic prices without price cap

Figure B.5 shows the distribution of dynamic prices that households on dynamic tariff actually paid with a price cap for the first grid supply area.[4] The graph reveals that dynamic prices are quite volatile at a relatively low level below the price cap. Figure B.6 shows the distribution of dynamic prices when removing the price cap. While dynamic prices only exceed the price cap in 3% of the half-hourly time intervals, the uncapped distribution has a very long tail that consists of a few rare, very high prices. Moreover,

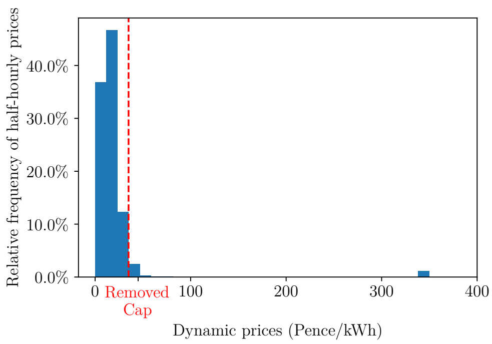

Figure B.7: Relative distribution of dynamic prices without price cap and with scarcity event

Figure B.7 reveals that the distribution of dynamic prices hardly changes when households are exposed to the scarcity event. As explained in Section 2.6.5, I simulate a four-day scarcity event by removing the price cap and manually setting the dynamic price to 350p/kWh for all 192 consecutive time intervals from 12 to 15 January 2021. I choose the

[4]Dynamic prices for other grid supply areas slightly differ due to minor differences in grid charges and peak-time premia.

scarcity price to equal ten times the price cap. These simulated scarcity prices account for only 1% of half-hourly dynamic prices.

Table B.2 compares the distribution of the dynamic prices with and without a price cap and with a scarcity event. The Table highlights that the removal of the price cap and the scarcity event have only small effects on the mean of the dynamic prices but a much larger effect on their standard deviation (SD) due to much higher maximum values.

Tariff type	mean	SD	min	1%	10%	25%	50%	75%	90%	99%	max
Dynamic capped	15.4	7.4	0.0	3.1	8.1	10.4	13.3	18.7	27.3	35.0	35.0
Dynamic removed cap	15.8	11.4	0.0	3.1	8.1	10.3	13.1	18.5	27.0	43.1	386.4
Dynamic scarcity event	20.1	40.0	0.0	3.1	8.1	10.3	13.1	18.7	27.9	350.0	350.0
Fixed	15.1	0.1	10.5								19.9

Table B.2: Distribution of dynamic prices and fixed prices (pence/kWh)

B.5 Additional results for different shares of electricity in household expenditure

This appendix tests the robustness of the optimal hedge shares to vary the assumption with respect to the average share of electricity $\overline{s}_k$ in household expenditure in hedge time segment k. In the main specification, I assume that each household spends on average $\overline{s}_k = 2\%$ of its household budget on electricity in every hedge time segment (see discussion Section 2.4). In this section, I test how the optimal hedge shares and bill volatility change when changing $\overline{s}_k$ to 1%, 5%, and 10%, respectively.

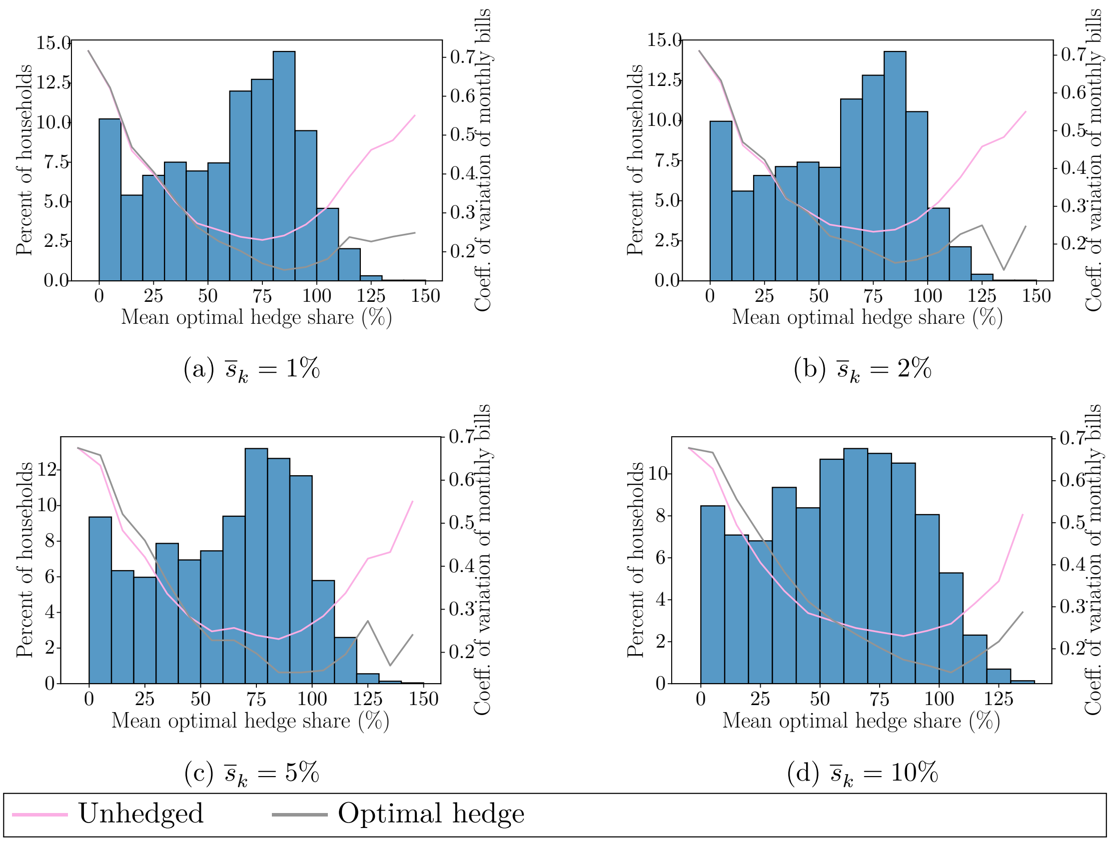

Figure B.8: Distribution of average optimal hedge shares by customer and coefficients of variation of monthly bills by tariff type for different average segment budget shares $\overline{s}_k$

Figure B.8 highlights that the mean optimal hedge shares hardly changes when the average time segment budget share $\overline{s}_k$ rises. For $\overline{s}_k = 1\%$, households' mean optimal hedge share is on average 59% compared to 57% when $\overline{s}_k = 10\%$. However, the variance in mean optimal hedge shares between households slightly rises.

The Figure also shows that the bill volatility for households on the optimally hedged tariff increases relative to the unhedged tariff when $\overline{s}_k$ increases. For small segment budget shares of $\overline{s}_k = 1\%$, the optimal hedged tariff achieves on average a 19% lower coefficient of variation of monthly electricity bills than the unhedged tariff compared to 5% when $\overline{s}_k = 10\%$.

On the other hand, Figure B.9 points out that an increase in budget share raises the welfare benefits from hedging. For small segment budget shares ($\overline{s}_k = 1\%$), the average welfare benefit achieved by the forward hedge only amounts to a 0.13% decrease in mean electricity prices. For large budget shares ($\overline{s}_k = 10\%$), the welfare benefits climb to 1.6%.

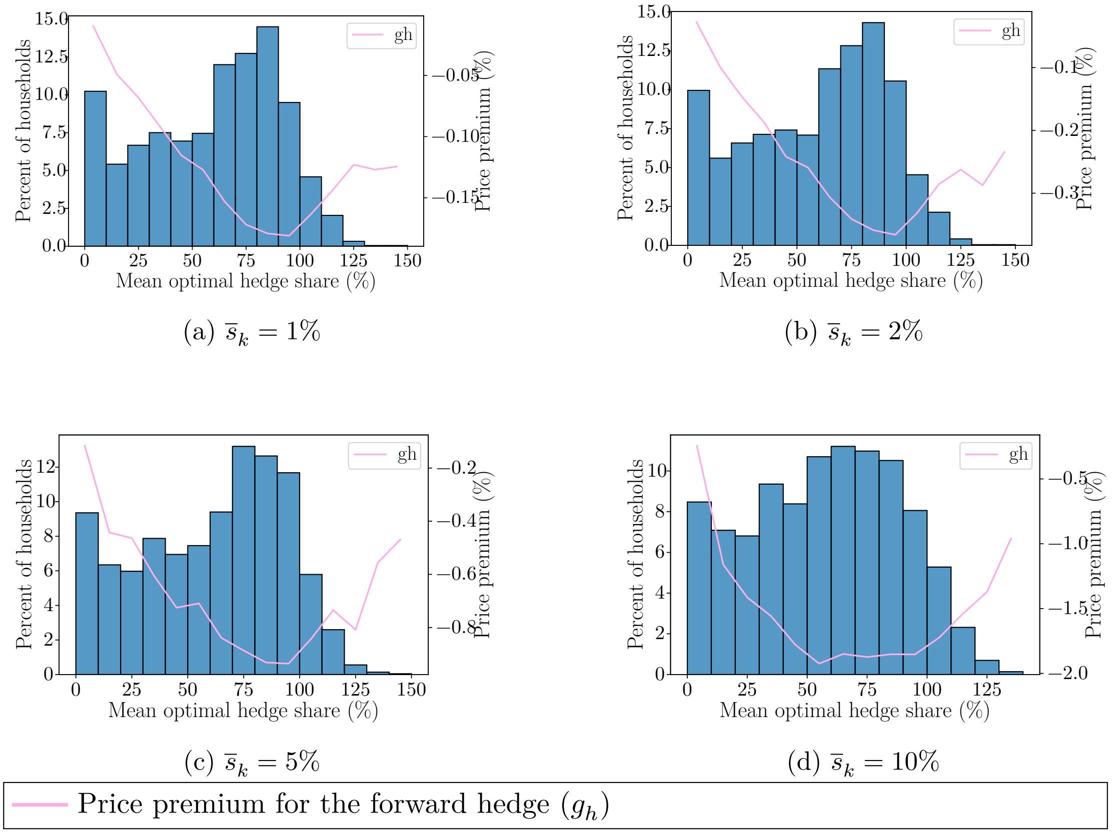

(a) $\overline{s}_k = 1\%$ (b) $\overline{s}_k = 2\%$

(c) $\overline{s}_k = 5\%$ (d) $\overline{s}_k = 10\%$

Figure B.9: Distribution of average optimal hedge shares by customer and hedge price premium g_h for different average segment budget shares $\overline{s}_k$

B.6 Additional results for the price premia when increasing electricity's expenditure share

Figure B.10 shows the price premium households are willing to pay for the forward hedge when electricity's expenditure share increases from $\overline{s}_k = 2\%$ to $\overline{s}_k = 10\%$.

The graph indicates that increasing the portion of expenditure on electricity to $\overline{s}_k = 10\%$ only has a marginal impact on enhancing the welfare advantages of hedging. For the main specification with $\alpha = 0.1$, the larger expenditure share of $\overline{s}_k = 10\%$ leads to a welfare improvement that corresponds to a decrease of 1.6% in the average electricity price, as opposed to a 0.3% decrease observed when the expenditure share on electricity remains at 2%.

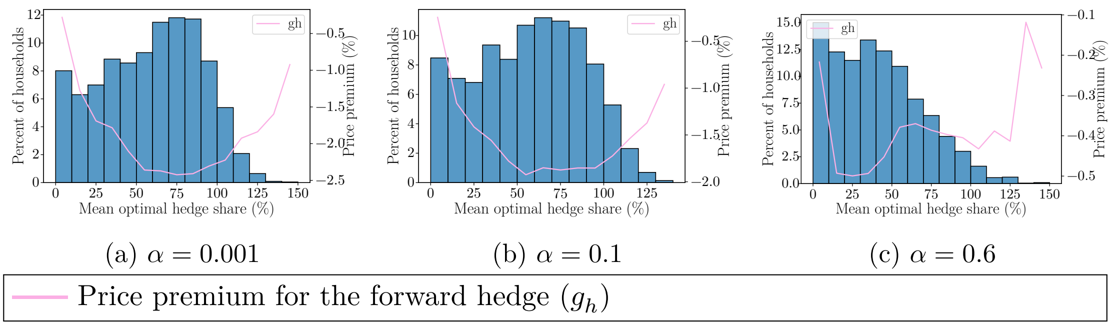

(a) $\alpha = 0.001$ (b) $\alpha = 0.1$ (c) $\alpha = 0.6$

Figure B.10: Distribution of average optimal hedge shares by customer and hedge price premium g_h with an expenditure share of electricity of 10%

Appendix C

Appendix to Chapter 3

C.1 Coefficients of the seasonal components of net load $S(t)$ and the residual process $S_X(t)$

Table C.1: Coefficients for the seasonal components of net load and the residual process (inspired by Coulon et al., 2013, net load is shown in MW)

Hour	**Parameters for $S(t)$**							**Parameters for $S_X(t)$**				
	a_1	a_2	a_3	a_4	a_5	a_6	a_7	b_1	b_2	b_3	b_4	b_5
0	202376.95	-5785.52	-0.61	4662.38	-0.76	-86.43	4.14	-0.31	-0.15	2.12	0.29	2.34
1	162445.84	-4920.87	-0.68	4583.09	-0.76	-67.32	-46.24	-0.53	-0.08	2.02	0.20	2.26
2	116833.06	-4251.69	-0.74	4554.41	-0.76	-45.07	-207.68	-0.65	0.07	-0.52	0.12	2.16
3	68011.24	-3755.37	-0.80	4532.40	-0.76	-20.95	-464.66	-0.66	0.16	-0.35	0.10	2.24
4	16100.24	-3323.39	-0.88	4566.24	-0.75	5.08	-1050.37	-0.54	0.28	-0.48	0.13	2.53
5	24236.35	-2896.22	-1.00	4608.91	-0.76	2.00	-2413.85	-0.26	0.46	-0.57	-0.21	5.68
6	96901.72	-2392.27	5.03	4454.63	-0.81	-32.52	-4495.02	0.26	0.93	-0.49	0.19	2.67
7	77293.43	-2188.93	-1.29	4452.25	-0.78	-22.11	-4885.41	0.17	0.83	-0.53	0.19	2.69
8	262755.59	-3035.20	-0.93	4646.54	-0.76	-113.71	-3932.05	0.07	0.67	-0.58	0.22	2.59
9	528066.45	-4439.00	-0.77	4711.01	-0.77	-244.72	-3261.66	0.03	0.60	5.62	-0.26	5.72
10	624256.69	-6603.83	-0.63	5018.46	-0.78	-291.73	-2969.80	0.00	0.51	5.54	0.28	2.59
11	605929.00	-9059.86	-0.55	5236.95	-0.81	-281.93	-2810.74	0.01	-0.38	2.25	0.32	-3.63
12	545398.21	-11253.77	-0.51	5285.51	-0.86	-251.28	-2748.23	0.06	-0.26	1.90	0.31	-3.62
13	512030.39	-13002.78	-0.49	5164.70	-0.91	-234.16	-2946.51	0.24	-0.22	0.64	0.30	-3.40
14	501096.98	-14226.35	-0.48	4974.34	-0.95	-228.37	-2998.73	0.49	-0.51	-0.02	-0.29	0.06
15	504599.44	-14833.20	-0.48	4759.61	-0.97	-229.89	-2959.78	0.85	0.94	2.97	-0.34	0.24
16	438199.50	-14661.10	-0.48	4627.56	-0.94	-196.78	-2939.22	1.03	1.00	2.99	-0.35	0.08
17	286718.22	-13319.79	-0.53	4862.27	-0.81	-121.55	-2707.58	0.87	0.13	-0.38	0.22	1.82
18	191466.85	-11511.85	-0.56	5261.76	-0.81	-74.47	-2454.90	0.59	0.72	-0.58	-0.23	-0.06
19	94055.47	-10570.78	-0.56	4931.22	-0.91	-26.67	-2277.90	0.34	-0.63	2.29	-0.46	-0.26
20	148836.53	-10014.19	-0.50	4697.71	-0.93	-54.37	-2217.77	0.19	-0.63	2.07	0.44	-3.60
21	285398.54	-9292.12	-0.47	4803.80	-0.82	-123.14	-1882.68	-0.04	0.50	5.47	0.32	2.52
22	311614.80	-8071.18	-0.50	4752.58	-0.77	-137.74	-1349.12	-0.12	0.42	5.49	0.31	2.48
23	328296.86	-6669.25	-0.52	4692.68	-0.80	-147.50	-919.84	-0.21	-0.29	2.38	0.30	2.36

C.2 Solar and wind generation model

In this appendix, we explain our renewable generation model and show some simulated solar and wind generation data.

For our renewables generation model, we use hourly generation profiles for 218 actually existing wind farms and 189 existing solar farms as provided by ERCOT (2023). The generation profiles were created using the Weather Research and Forecasting (WRF) model to simulate historical meteorological conditions. Plant-specific characteristics and observed generation data were incorporated to model operational and planned wind and solar plants. The final profiles were validated against observed generation, capturing seasonal, diurnal, and ramping behaviors. Details on the creation of the hourly generation profiles can be found in ERCOT (2022) and UL Services Group (2022). In addition to plant-level data, we use hourly data on aggregate wind and solar generation from ERCOT (2024c).[1] For all these time series, we have hourly data from 2011-22.

The basic idea of our renewable generation model can be summarized as follows: First, we group the observed hourly data into *year-month-hour* groups to capture seasonal and time-of-day patterns in wind or solar generation. Next, we simulate aggregate solar and wind generation within each group based on its historical mean and variability. Then, we use the plant-specific correlation with the aggregate output to adjust each plant's individual generation around its specific historical mean. Finally, we draw random variation around that conditional mean to randomly simulate hourly generation for each plant.

Let $G_{p,t}$ be the generation of plant p in hour t as given in the ERCOT (2023) generation profile for plant p. Each hour t belongs to a year-month-hour group, which we denote by g. For example, g could represent all observations for January 2015 at 3:00 pm. Grouping the generation captures seasonal and time-of-day cycles in solar and wind generation. Let $\text{GroupMean}_{p,g}$ and $\text{GroupStd}_{p,g}$ be the mean and standard deviation of plant p in group g.

Similarly, let $G_{\text{agg},t}$ denote the observed aggregate wind (or solar) generation at time t. For each group g, we define AggMean_g and AggStd_g to be the mean and standard deviation of the observed aggregate generation. We also calculate the correlation $\rho_{p,g}$ between a plant p's observed generation $G_{p,t}$ and the observed aggregate generation $G_{\text{agg},t}$ in each group g.

For wind and solar generation, respectively, we first simulate aggregate generation values $\widetilde{G}_{\text{agg},t}$ by drawing from a normal distribution centered at zero with standard

[1] We aggregate the 15-minute-level aggregate solar and wind generation to hourly generation

deviation AggStd_g and then shifting by the group mean AggMean_g. Concretely, for each hour t in group g:

$$Z_{g,t} \;\sim\; \mathcal{N}(0,\; \text{AggStd}_g^2),$$
$$\widetilde{G}_{\text{agg},t} \;=\; Z_{g,t} \;+\; \text{AggMean}_g.$$

We set all negative simulated values to zero to ensure that simulated aggregate generation remains nonnegative.

Once the aggregate generation $\widetilde{G}_{\text{agg},t}$ is simulated, we model plant-level generation conditional on $\widetilde{G}_{\text{agg},t}$. Thereby, we capture that a wind plant's generation is typically correlated with aggregate wind generation. For each plant p in group g, we treat the correlation $\rho_{p,g}$ as the slope of a linear relationship between plant level generation $G_{p,t}$ and aggregate generation $G_{\text{agg},t}$.

$$\beta_{p,g} \;=\; \rho_{p,g} \;\times\; \frac{\text{GroupStd}_{p,g}}{\text{AggStd}_g}.$$

The coefficient $\beta_{p,g}$ measures how much plant p's generation changes with a unit change in the aggregate generation. Given the simulated aggregate generation $\widetilde{G}_{\text{agg},t}$, we write the conditional mean of plant p's generation as:

$$\mathbb{E}[G_{p,t} \mid \widetilde{G}_{\text{agg},t}] \;=\; \text{GroupMean}_{p,g} \;+\; \beta_{p,g}\Big(\widetilde{G}_{\text{agg},t} - \text{AggMean}_g\Big).$$

This shifts the plant's mean generation $\text{GroupMean}_{p,g}$ in response to deviations of the simulated aggregate generation from its own aggregate group mean. Here, we take into account that a wind plant's generation is typically lower if aggregate wind generation is low. Assuming a bivariate normal relationship between each plant's generation and the aggregate generation, the conditional variance of $G_{p,t}$ given $\widetilde{G}_{\text{agg},t}$ is

$$\text{Var}[G_{p,t} \;\mid\; \widetilde{G}_{\text{agg},t}] \;=\; \text{GroupStd}_{p,g}^2(1 - \rho_{p,g}^2).$$

Finally, we randomly draw a generation value $G_{p,t}$ for plant p conditional on the simulated aggregate generation $\widetilde{G}_{\text{agg},t}$:

$$G_{p,t}^{(\text{sim})} \;\sim\; \mathcal{N}\Big(\mathbb{E}[G_{p,t} \mid \widetilde{G}_{\text{agg},t}],\; \text{Var}[G_{p,t} \mid \widetilde{G}_{\text{agg},t}]\Big),$$

Again, we set negative generation values to zero. The result $G_{p,t}^{(\text{sim})}$ represents the simulated generation of plant p in hour t, given the simulated aggregate generation $\widetilde{G}_{\text{agg},t}$ and the historical correlation, mean, and variance patterns observed in group g.

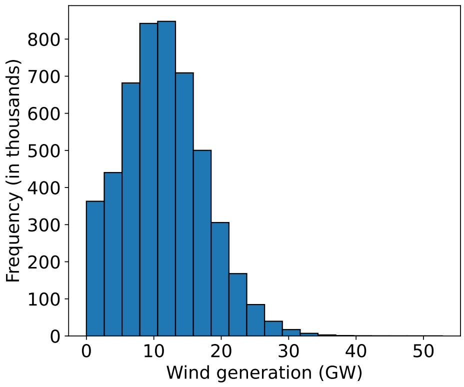

(a) Simulated wind generation (GW)

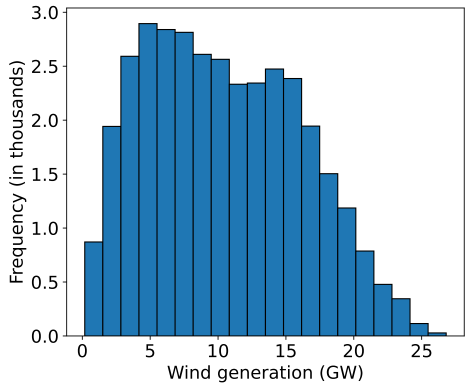

(b) Observed wind generation (GW)

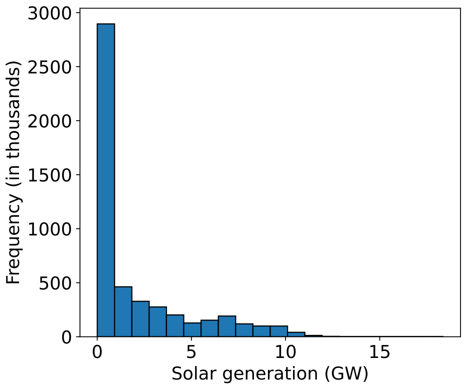

(c) Simulated solar generation (GW)

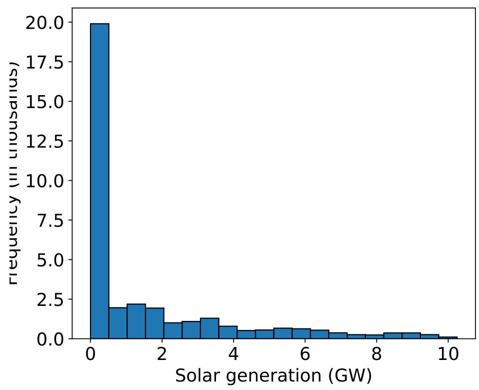

(d) Observed solar generation (GW)

Figure C.1: Histogram of solar and wind generation for simulated and observed data between 2019-22

Figure C.1 shows histograms of the simulated and the observed aggregate solar and wind generation for 2019-2022. The simulated data in Graphs C.1a and C.1c is the histogram of the sum of the simulated hourly generation for all solar and wind plants, respectively. Graphs C.1b and C.1d show the histogram of observed hourly solar and wind generation.

The graphs reveal that our solar and wind generation model creates hourly generation values that have quite similar distributions to the observed aggregate data. The only major difference is that the simulated solar and wind data shows some large outliers that are not observed in our data.

C.3 Histogram of hourly load, net load, and gas prices for simulated and observed data between 2019-22

In this appendix, we show histograms of simulated and observed hourly load, net load, and gas prices. In addition, we plot the relationship between load and day-ahead prices for a generator and an LSE

The left column in Figure C.2 shows the simulated data and the right column shows the respective observed data. Figures C.2a and C.2b show that simulated and observed hourly load have similar distributions. However, observed load is more right-skewed than simulated load. This right-skewness in observed load might be driven by a positive correlation between net load and renewable generation. Our model treats net load and renewable generation as independent, which is a limitation.

Simulated and observed net load have more similar distributions as Figures C.2c and C.2d highlight. Yet, observed net load has a slightly higher standard deviation than load. The distribution of observed and simulated gas prices also look reasonably similar (Graphs C.2e and C.2f).

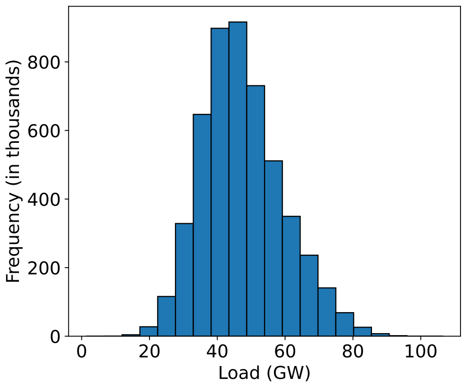

(a) Simulated load (GW)

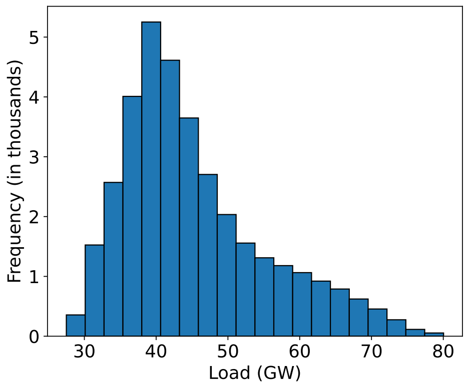

(b) Observed load (GW)

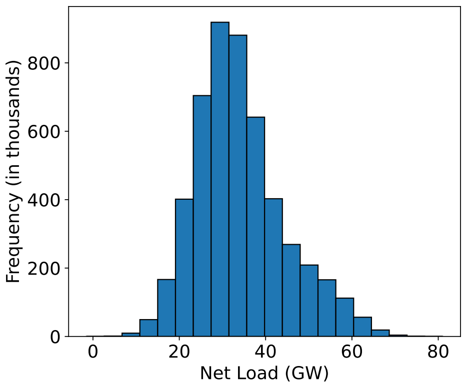

(c) Simulated net load (GW)

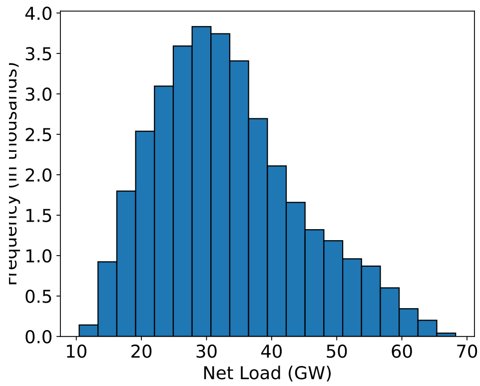

(d) Observed net load (GW)

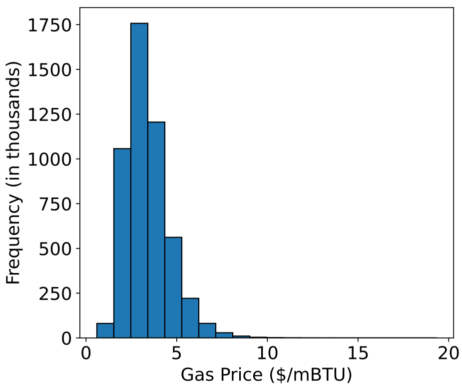

(e) Simulated gas prices ($\$/mBTU$)

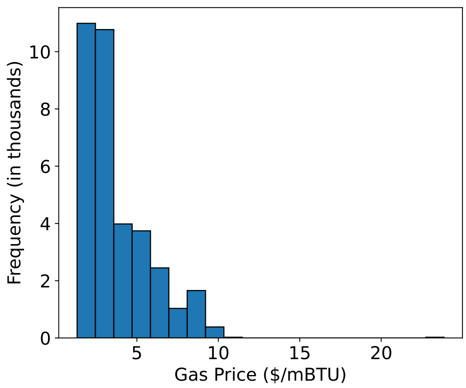

(f) Observed gas prices ($\$/mBTU$)

Figure C.2: Histogram of hourly load, net load, and gas prices for simulated and observed data between 2019-22

In Figure C.3, we also plot the relationship between simulated load and day-ahead prices for an LSE and the generator. In both cases, we only find a moderate positive correlation between load and day-ahead prices.

In particular, it is surprising that very high peak prices even occur when load is below its average. On the other hand, we also sometimes observe low prices when load is really high. These plots reveal that spike prices are not mainly driven by high load, but likely by high net load in combination with random shocks like outages.

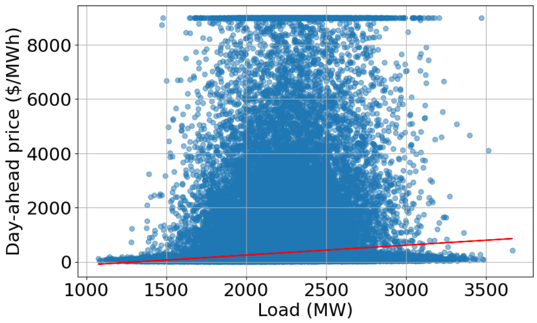

(a) LSE

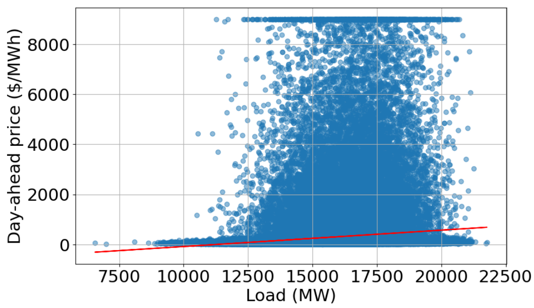

(b) Generator

Figure C.3: Relationship between LSE and generator's load and day-ahead electricity prices in peak delivery period 4-5 pm in August 2019

C.4 Hedge strategies for an off-peak delivery period

In this appendix, we deep dive on off-peak delivery period Weekends 4-5 am in May 2019 which has very rare and small price spikes.

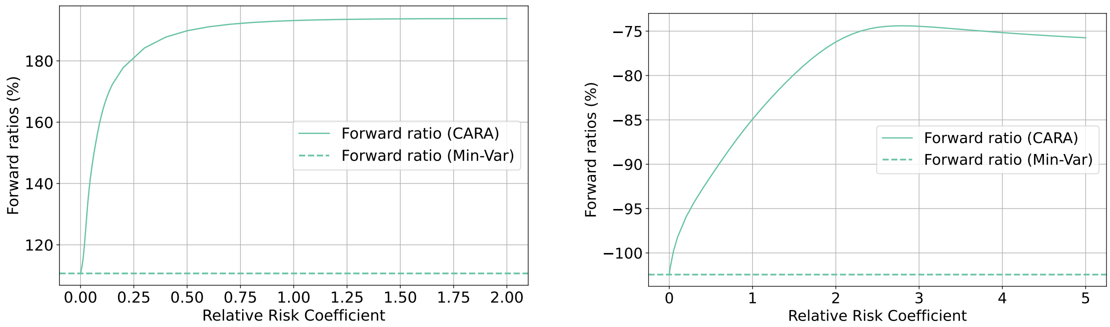

(a) LSE (b) Generator

Figure C.4: LSE and generator's optimal forward ratio by relative risk aversion for off-peak delivery period weekends 4-5 am in May 2019

Figure C.4 shows the optimal forward ratios of the LSE and the generator as a function of their risk-aversion levels for off-peak delivery period weekends 4-5 am in May 2019. The graphs reveal that the hedge strategies are very similar to the strategies in a peak period as shown in Figure 3.8. Interestingly, the LSE chooses higher hedge ratios in the off-peak than in the peak period. This might indicate that the LSE's load experiences larger positive outliers relative to average load in the off-peak than in the peak period. For the generator, the hedge ratios in the peak and off-peak periods are almost identical when hedging with forwards only.

Figure C.5 shows how the forward ratio affects the LSE and generator's profit volatility, their conditional value at risk at the 5% level ($CVaR_{\alpha=5\%}$), and worst-case minimum profit relative to their respective mean unhedged day-ahead profit. The Figure higlights that in the off-peak period, agents face the same trade-off between minimizing variance, and maximizing $CVaR_{\alpha=5\%}$ or worst-case minimum profit as discussed in section 3.6.1.

Figure C.6 below shows the optimal forwards and option ratios for the combined hedge strategy as a function of agents risk aversion for the off-peak period weekends 4-5 am in May 2019. Figure C.6a highlights that the LSE buys more options than forwards for moderate risk aversion in the off-peak period. When the LSE's risk aversion increases further, its option ratio rapidly declines while the forward ratio increases. Thus, the highly risk-averse LSE relies far less on options in the off-peak period compared to the peak period.

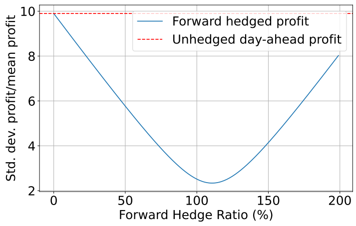

(a) LSE's standard deviation of profits

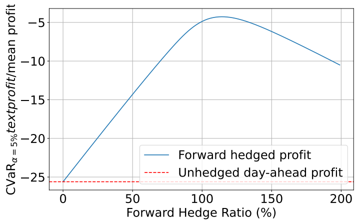

(b) LSE's $CVaR_{\alpha=5\%}$ profit

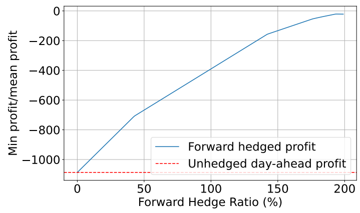

(c) LSE's minimum profit

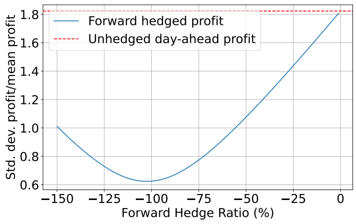

(d) Generator's standard deviation of profits

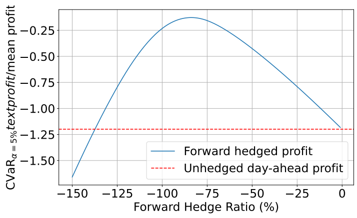

(e) Generator's $CVaR_{\alpha=5\%}$ profit

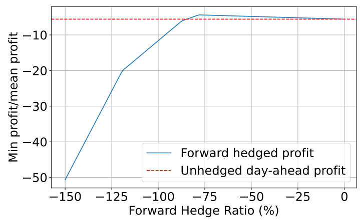

(f) Generator's minimum profit

Figure C.5: LSE and generator's profit standard deviation, $CVaR_{\alpha=5\%}$ profit outcome, and minimum profit relative to mean unhedged day-ahead (DA) profit by forward ratio for off-peak delivery period weekends 4-5 am in May 2019

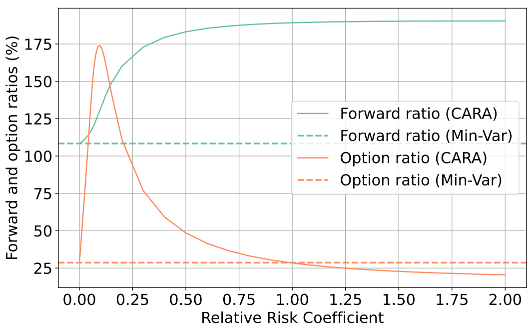

(a) LSE

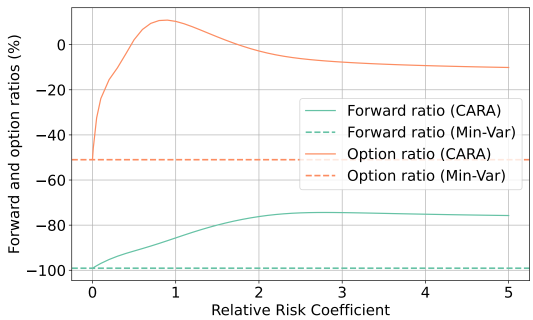

(b) Generator

Figure C.6: LSE and generator's optimal forward and option ratio by relative risk aversion for off-peak period weekends 4-5 am in May 2019

Similar to the LSE, the generator also relies less on long option positions in the off-peak period (Graph C.6b). It only chooses small long option positions for moderate risk aversion levels and short positions for all other levels.

Figure C.7 depicts the agents' standard deviation of profits, $CVaR_{\alpha=5\%}$ tail loss, and worst-case minimum loss as a function of the optimal forward and option ratios that are associated with a relative risk aversion level. For the off-peak period, the Figure points out that the combined option and forward strategy is more successful in reducing profit variance compared to the forwards only strategy. The combined strategy also achieves substantial improvements in $CVaR_{\alpha=5\%}$, especially for the LSE. For the generator, the combined strategy trades off improvements in $CVaR_{\alpha=5\%}$ versus high worst-case losses for low-risk aversion levels. Overall, combining forwards and options hardly improves minimum worst-case losses relative to the forward-only strategy.

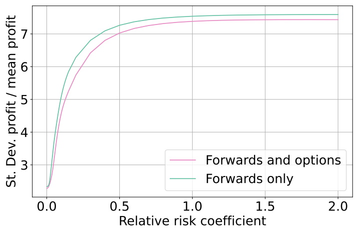

(a) LSE's standard deviation of profits

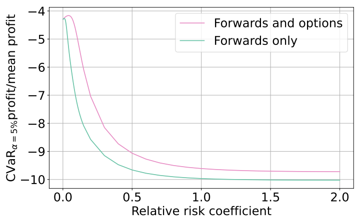

(b) LSE's $CVaR_{\alpha=5\%}$ profit

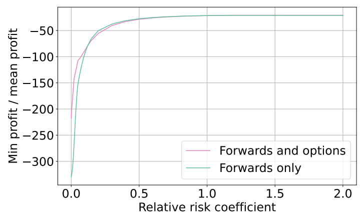

(c) LSE's minimum profit

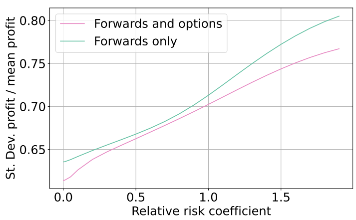

(d) Generator's standard deviation of profits

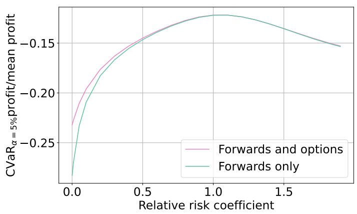

(e) Generator's $CVaR_{\alpha=5\%}$ profit

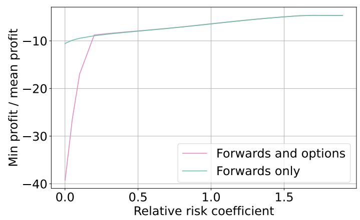

(f) Generator's minimum profit

Figure C.7: LSE and generator's profit standard deviation, $CVaR_{\alpha=5\%}$, and minimum profit relative to mean unhedged day-ahead (DA) profit by relative risk aversion level for off-peak period weekends 4-5 am in May 2019

C.5 Comparison of profit performance metrics of hedging strategies relative to the unhedged day-ahead profit

Figures C.8 shows the standard deviation of profits, $CVaR_{\alpha=5\%}$ tail profit outcome, and worst-case minimum loss for the generator and LSE as a function of the optimal forward and option ratios that are associated with a relative risk aversion level. In contrast to Figure 3.14 in section 3.6.2, this Figure also shows the profit standard deviation, $CVaR_{\alpha=5\%}$, and worst-case minimum loss for the LSE and generator if they trade fully unhedged in the day-ahead market.

All three metrics are significantly improved when hedging with either forwards and options or with forwards only compared to the unhedged day-ahead strategy. In contrast, the added value of combining forwards and options relative to forwards only is almost negligible compared to the added value of hedging with either strategy relative to remaining unhedged.

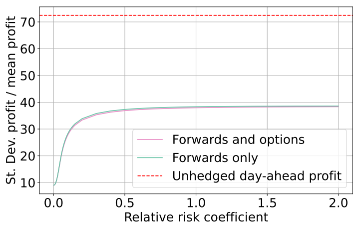

(a) LSE's standard deviation of profits

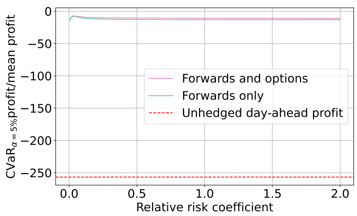

(b) LSE's $CVaR_{\alpha=5\%}$ profit

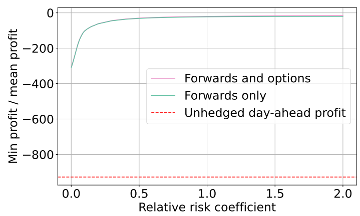

(c) LSE's minimum profit

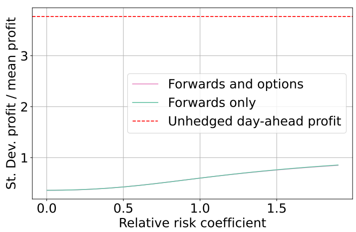

(d) Generator's standard deviation of profits

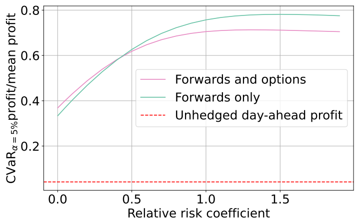

(e) Generator's $CVaR_{\alpha=5\%}$ profit

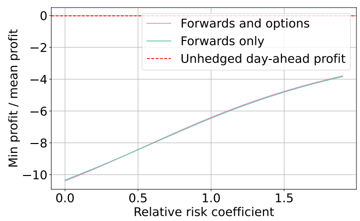

(f) Generator's minimum profit

Figure C.8: LSE and generator's profit standard deviation, $CVaR_{\alpha=5\%}$, and minimum profit relative to mean unhedged day-ahead (DA) profit by relative risk aversion level for delivery period Weekdays 4–5 pm in August 2019

Appendix D

Appendix to Chapter 4

D.1 Net demand curves for forwards and options

In this appendix, we investigate how the net demand for forwards reacts to the option price and how the net demand for options reacts to the forward price. Figures D.1a (LSE) and D.1c (generator) indicate that net demand for forwards increases in the option price when fixing the forward price at arbitrage-free levels. Similarly, demand for options also increases in the forward price when setting the option price to its unbiased values as Figures D.1b (LSE) and D.1d (generator) highlight. For both agents, option demand is highly sensitive to the forward price. By contrast, forwards demand is less sensitive to changes in the option price.

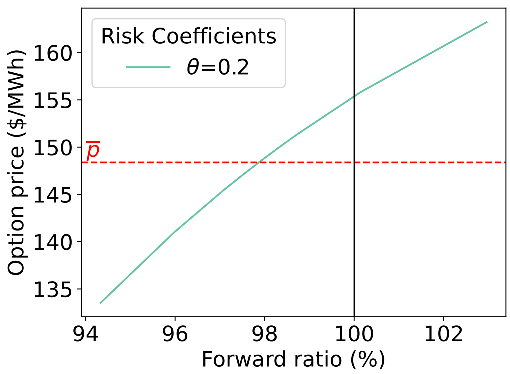

(a) LSE's forward ratios

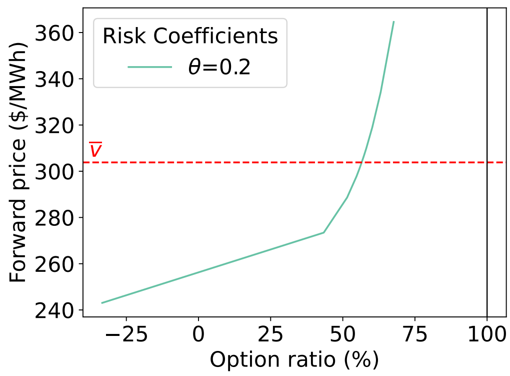

(b) LSE's option ratios

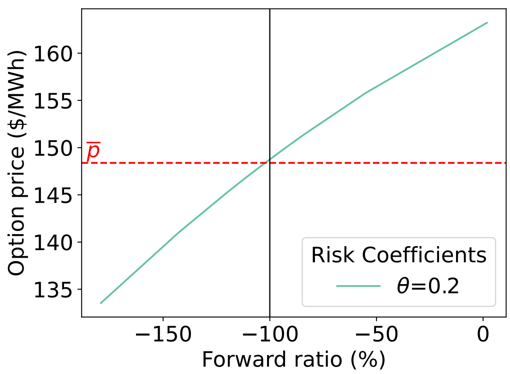

(c) Generator's forward ratios

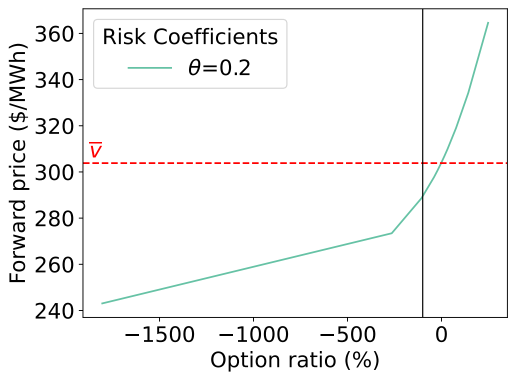

(d) Generator's option ratios

Figure D.1: LSE and generator's inverse net demand curves for forwards as a function the option price and for options as a function of the forward price for peak period weekdays 4-5 pm in August 2019 by risk aversion

D.2 Net demand heat maps by risk preferences

In this appendix, we show additional heat maps for forward and option ratios as a function of forward and option prices to study how agents' risk preferences affect their net demand for forwards and options. The first two figures show the demand matrices for the off-peak period weekends 4-5 am in May 2019 for the LSE (Figure D.2) and the generator (Figure D.3), respectively. Figures D.4 (LSE) and D.5 (generator) present the same matrices for the peak period August 4-5 pm in August 2019.

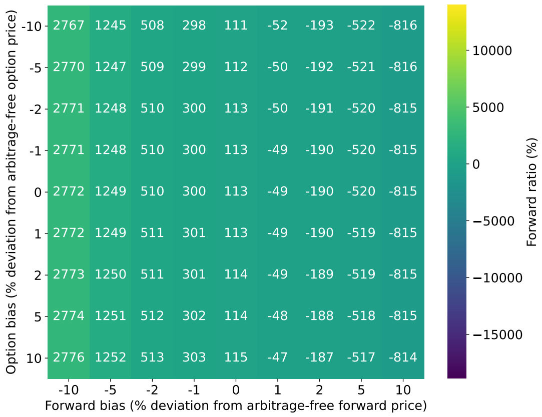

(a) Almost risk-neutral LSE's forward ratios

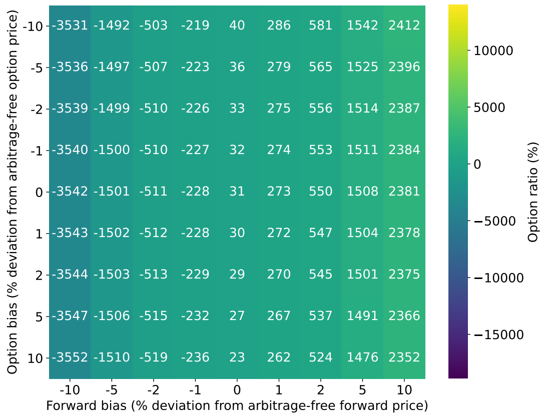

(b) Almost risk-neutral LSE's option ratios

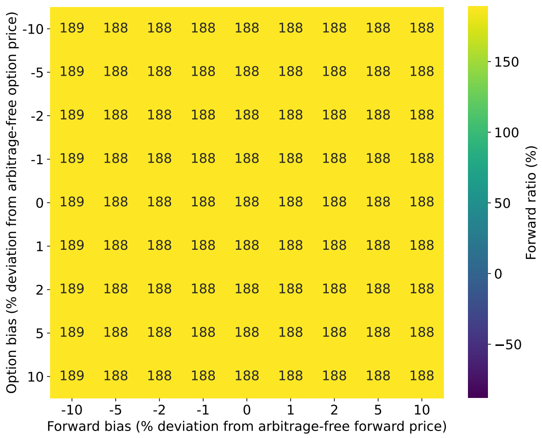

(c) Highly risk-averse LSE's forward ratios

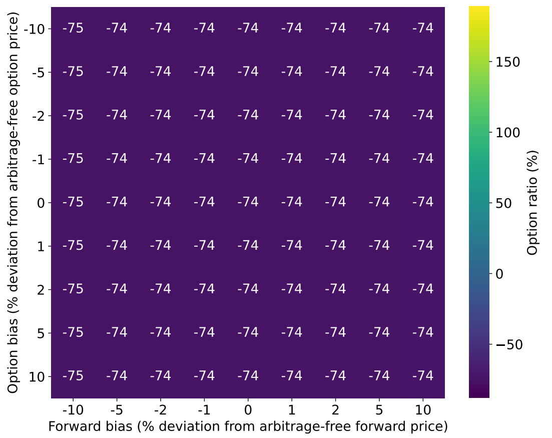

(d) Highly risk-averse LSE's option ratios

Figure D.2: LSE's optimal forward and option ratios as a function of the bias in forward and option prices for off-peak delivery period weekends 4-5 am in May 2019 for near risk neutrality (risk coefficient of 0.001) and high risk aversion (risk coefficient 1)

Figures D.2 (LSE) and D.3 (generator) show the heat maps for almost risk-neutral agents (with relative risk coefficient of 0.001) and highly risk-averse agents (risk coefficient 1). The nearly risk-neutral agents take very large speculative arbitrage positions compared to the moderate risk-aversion level of 0.2 analyzed in Figure 4.4 in 4.5.1. The highly risk-averse generator takes much smaller arbitrage positions than the moderately risk-averse generator while the highly risk-averse LSE hardly arbitrages.

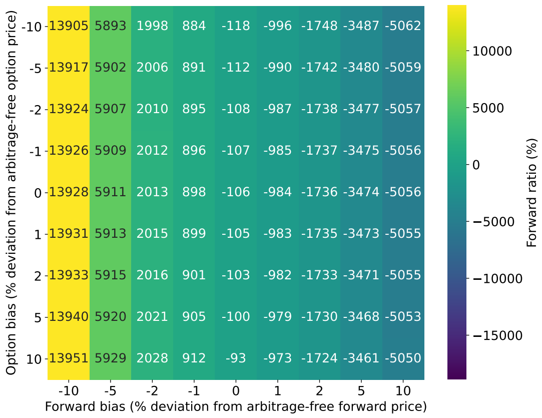

(a) Almost risk-neutral generator's forward ratios

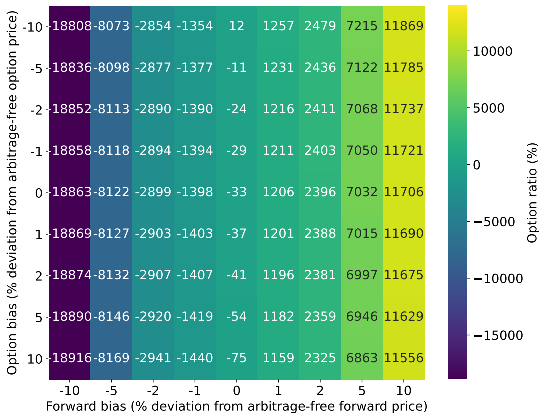

(b) Almost risk-neutral generator's option ratios

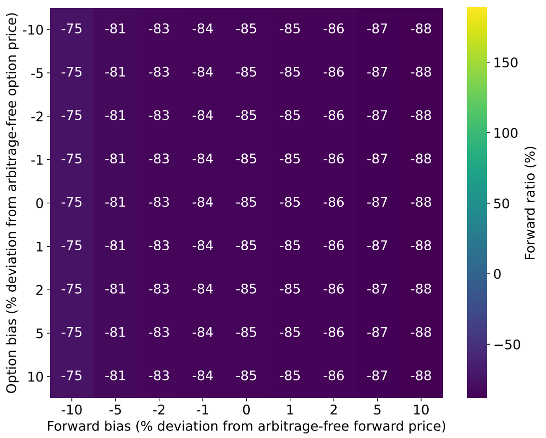

(c) Highly risk-averse generator's forward ratios

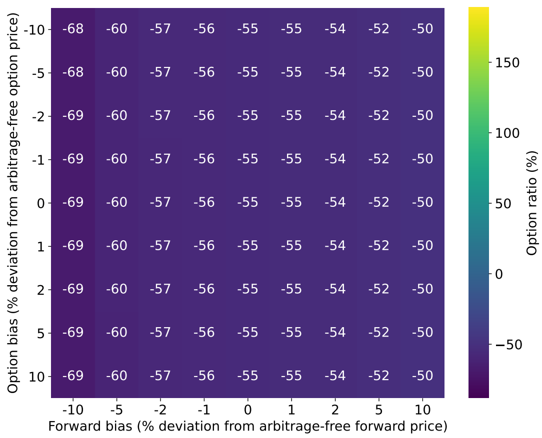

(d) Highly risk-averse generator's option ratios

Figure D.3: Generator's optimal forward and option ratios as a function of the bias in forward and option prices for off-peak delivery period weekends 4-5 am in May 2019 for near risk neutrality (risk coefficient of 0.001) and high risk aversion (risk coefficient 1)

Figures D.4 (LSE) and D.5 (generator) show the analogous heat maps for the peak period. In the peak period, the almost risk-neutral agents also heavily arbitrage, while the risk-averse agents take smaller arbitrage positions.

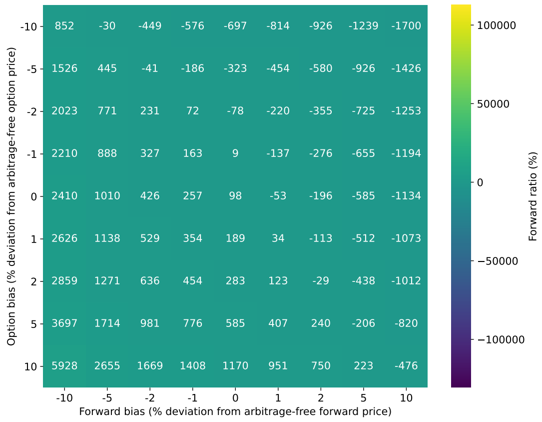

(a) Almost risk-neutral LSE's forward ratios

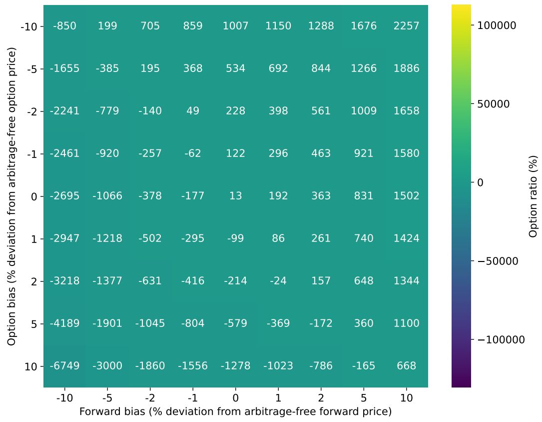

(b) Almost risk-neutral LSE's option ratios

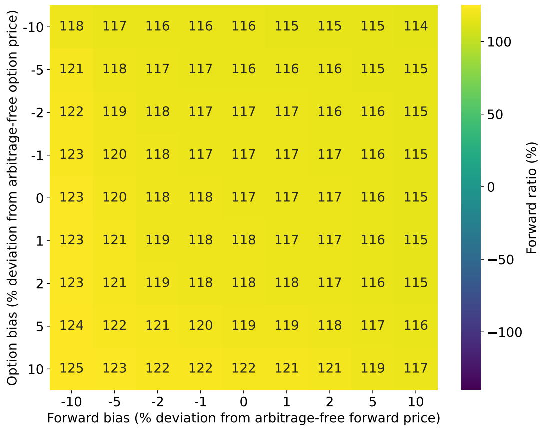

(c) Highly risk-averse LSE's forward ratios

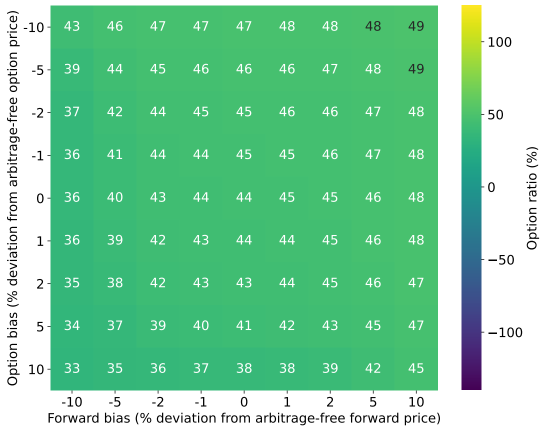

(d) Highly risk-averse LSE's option ratios

Figure D.4: LSE's optimal forward and option ratios as a function of the bias in forward and option prices for peak delivery period weekdays 4-5 pm in August 2019 for near risk-neutrality (risk coefficient of 0.001) and high risk aversion (risk coefficient 1)

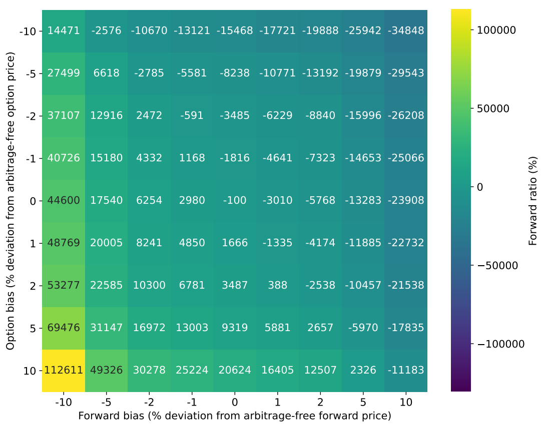

(a) Almost risk-neutral generator's forward ratios

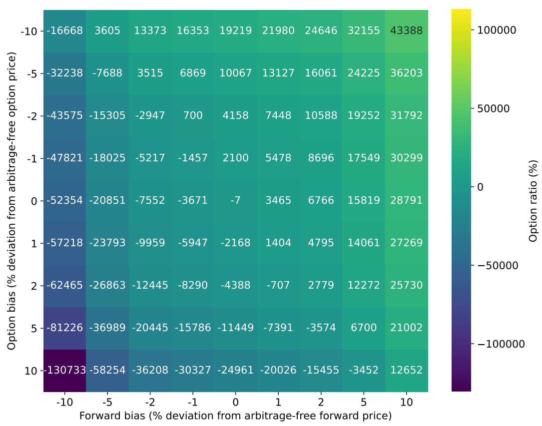

(b) Almost risk-neutral generator's option ratios

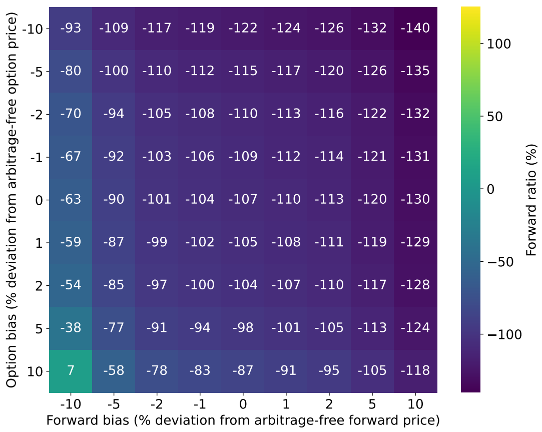

(c) Highly risk-averse generator's forward ratios

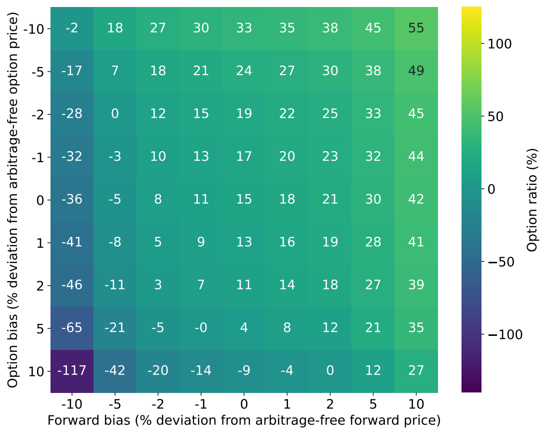

(d) Highly risk-averse generator's option ratios

Figure D.5: Generator's optimal forward and option ratios as a function of the bias in forward and option prices for peak delivery period weekdays 4-5 pm in August 2019 for near risk-neutrality (risk coefficient of 0.001) and high risk aversion (risk coefficient 1)

D.3 Net demand heat maps for a low strike price

In this appendix, the below figures examine the impact of selecting a lower strike price of \$200/MWh on agents' net demand. Figure D.6 depicts the off-peak results and Figure D.7 the peak results. With the low strike price, agents choose larger option positions and smaller forward positions compared to the high strike price of $\$1,000$/MWh, as shown in Graph 4.4 in Section 4.6. Brandkamp et al. (2025) have already established this finding for arbitrage-free prices. Additionally, the below figures demonstrate that a lower strike price of \$200/MWh encourages agents to take larger arbitrage positions compared to a higher strike price of $\$1,000$/MWh.

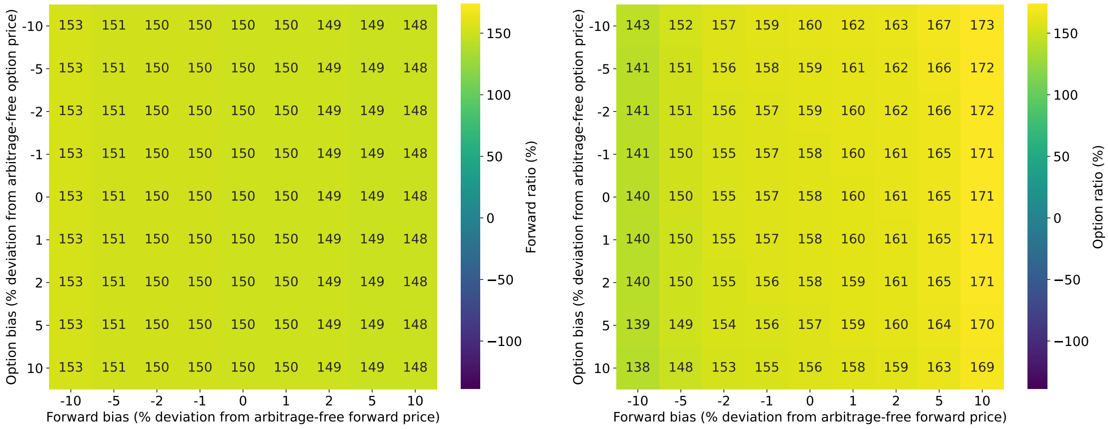

(a) LSE's forward ratios

(b) LSE's option ratios

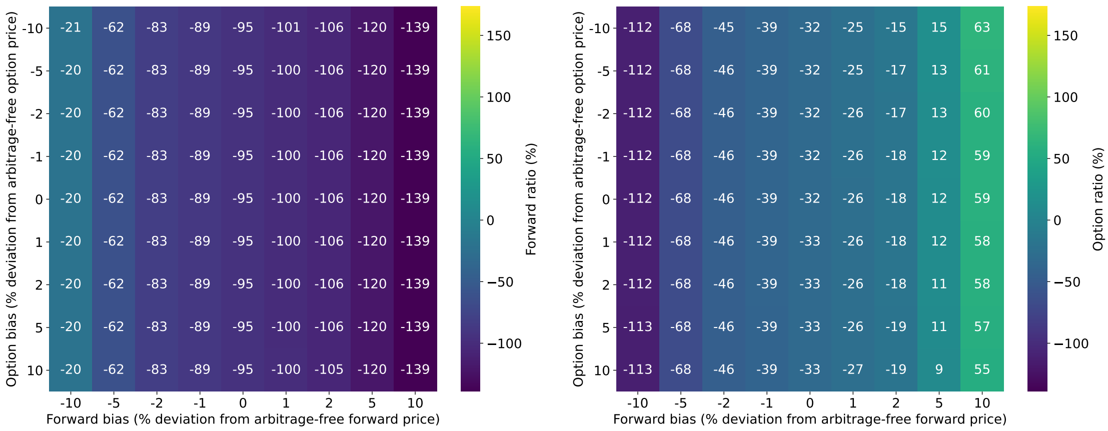

(c) Generator's forward ratios

(d) Generator's option ratios

Figure D.6: LSE and generator's optimal forward and option ratios as a function of the bias in forward and option prices for off-peak period weekends 4-5 am in May 2019 for a relative risk coefficient of 0.2 and a low strike price of \$200/MWh

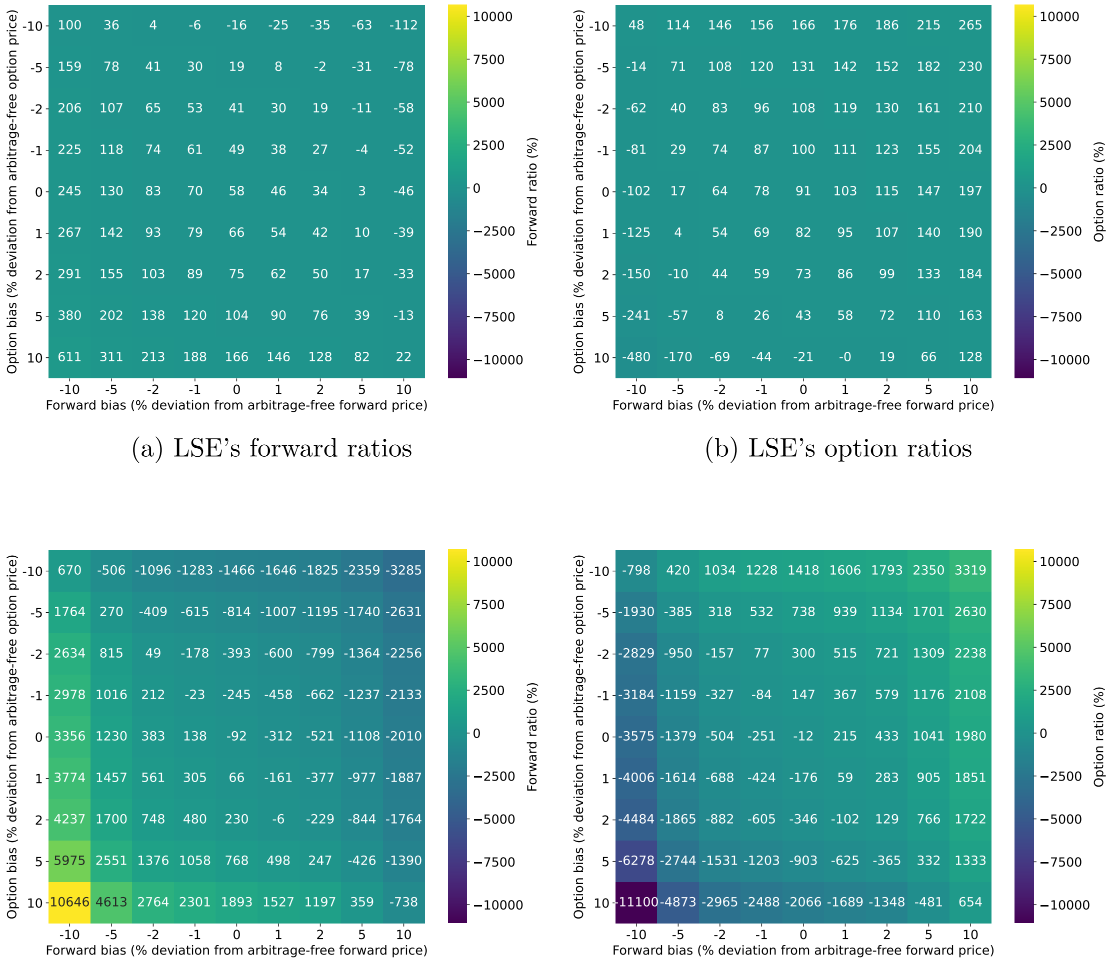

(a) LSE's forward ratios

(b) LSE's option ratios

(c) Generator's forward ratios

(d) Generator's option ratios

Figure D.7: LSE and generator's optimal forward and option ratios as a function of the bias in forward and option prices for peak period weekdays 4-5 pm in August 2019 for a relative risk coefficient of 0.2 and a low strike price of $200/MWh

D.4 Capacity mixes of the representative generators

The below figures show the capacity mixes of the simulated representative generation companies we created based on the power plant datasets by Mann et al. (2017) and ERCOT (2023). For each generation company, its total installed capacity and its number of power plants are given in brackets.

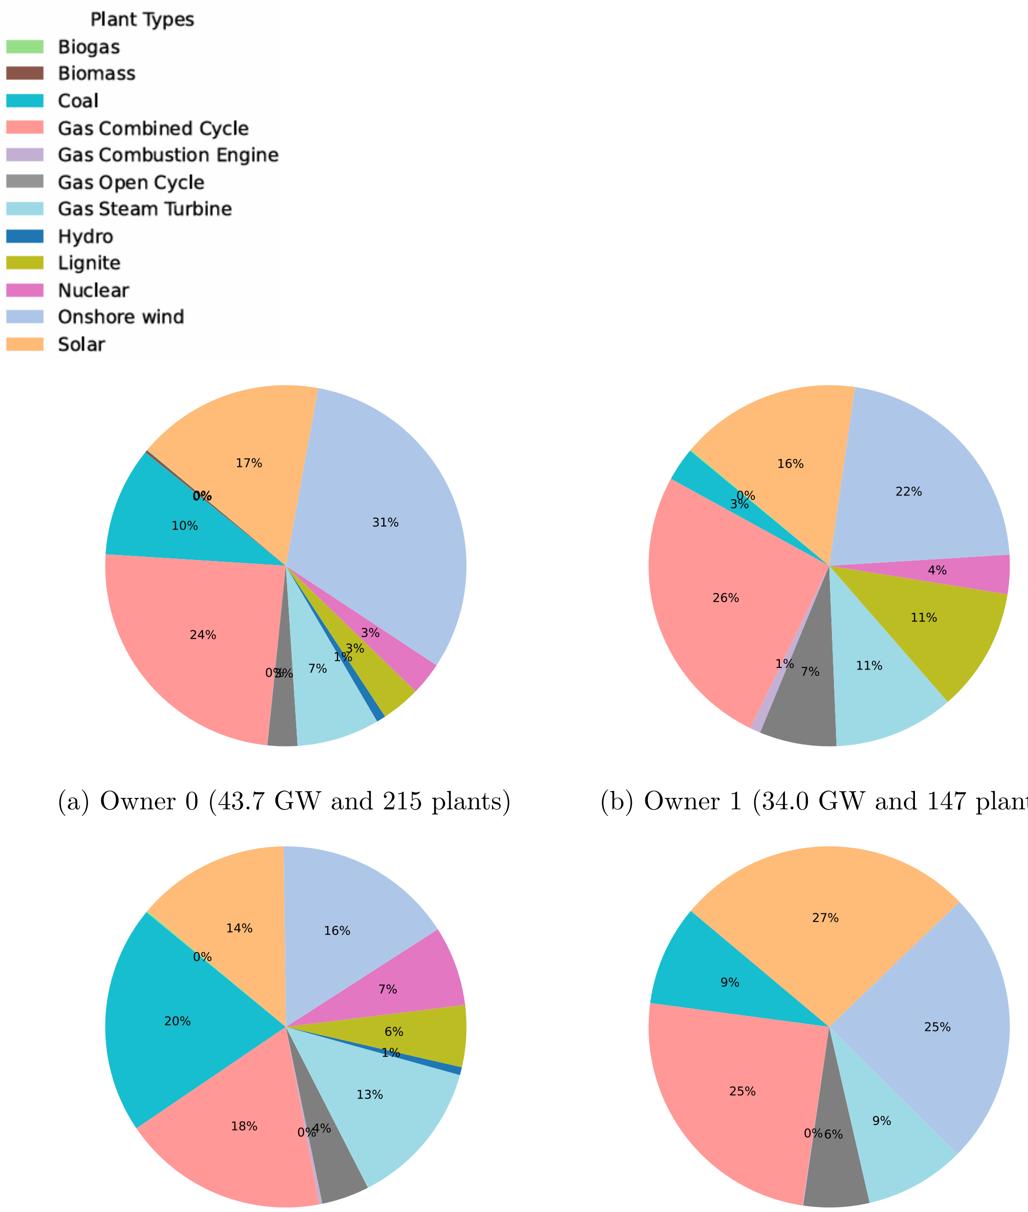

(a) Owner 0 (43.7 GW and 215 plants)

(b) Owner 1 (34.0 GW and 147 plants)

(c) Owner 2 (18.0 GW and 71 plants)

(d) Owner 3 (14.0 GW and 75 plants)

Figure D.8: Capacity mix per owner in % of total capacity (total capacity and number of plants in brackets)

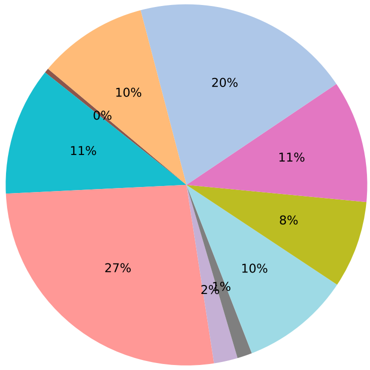

(a) Owner 4 (11.0 GW and 36 plants)

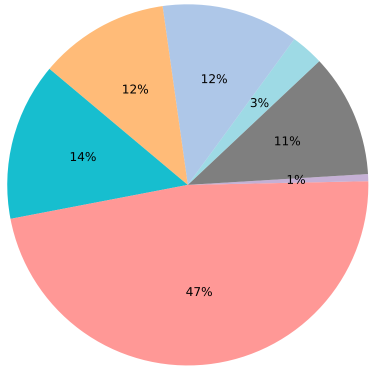

(b) Owner 5 (8.0 GW and 29 plants)

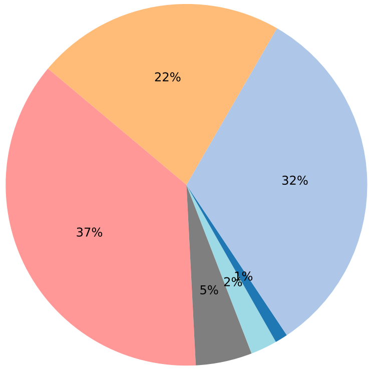

(c) Owner 6 (6.0 GW and 34 plants)

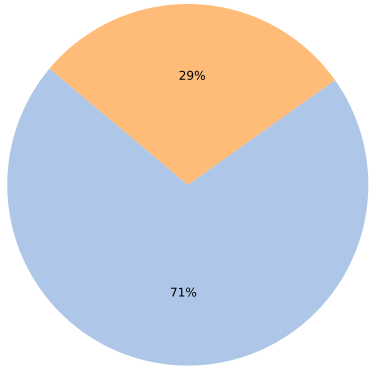

(d) Owner 7 (4.0 GW and 23 plants)

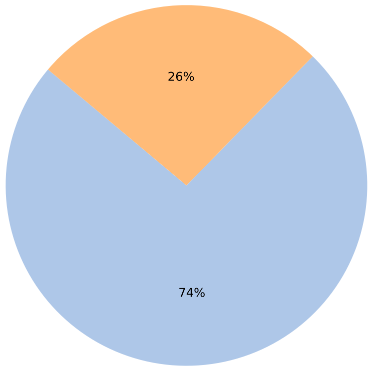

(e) Owner 8 (3.0 GW and 16 plants)

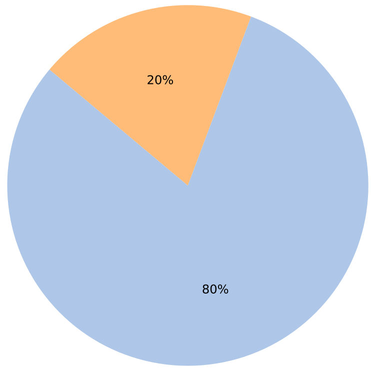

(f) Owner 9 (2.0 GW and 15 plants)

Figure D.9: Capacity mix per owner in % of total capacity (total capacity and number of plants in brackets)(continued)