Market Design Opportunities for an Evolving Power System

by

Ian Michael Schneider

Submitted to the Institute for Data, Systems, and Society
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Social and Engineering Systems

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2020

Author . . . .
Institute for Data, Systems, and Society
December 5, 2019

Certified by . . . .
Professor Munther A. Dahleh
Director, Institute for Data, Systems, and Society
Thesis Supervisor

Certified by . . . .
Dr. Mardavij Roozbehani
Principal Research Scientist, Laboratory for Information & Decision Systems
Thesis Supervisor

Certified by . . . .
Professor Paul Joskow
Professor Emeritus, Department of Economics
Doctoral Committee

Accepted by . . . .
Professor Ali Jadbabaie
Associate Director, Institute for Data, Systems, and Society

Market Design Opportunities for an Evolving Power System

by

Ian Michael Schneider

Submitted to the Institute for Data, Systems, and Society
on December 5, 2019, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Social and Engineering Systems

Abstract

The rapid growth of renewable energy is transforming the electric power sector. Wind and solar energy are non-dispatchable: their energy output is uncertain and variable from hour-to-hour. New challenges arise in electricity markets with a large share of uncertain and variable renewable energy. We investigate some of these challenges and identify economic opportunities and policy changes to mitigate them.

We study electricity markets by focusing on the preferences and strategic behavior of three different groups: producers, consumers, and load-serving entities. First, we develop a game-theoretic model to investigate energy producer strategy in electricity markets with high levels of uncertain renewable energy. We show that increased geographic dispersion of renewable generators can reduce market power and increase social welfare. We also demonstrate that high-quality public forecasting of energy production can increase welfare. Second, we model and explain the effects of retail electricity competition on producer market power and forward contracting. We show that increased retail competition could decrease forward contracting and increase electricity prices; this is a downside to the general trend of increased access to retail electricity competition. Finally, we propose new methods for improving demand response programs. A demand response program operator commonly sets customer baseline thresholds to determine compensation for individual customers. The optimal way to do this remains an open question. We create a new model that casts the demand response program as a sequential decision problem; this formulation highlights the importance of learning about individual customers over time. We develop associated algorithms using tools from online learning, and we show that they outperform the current state of practice.

Thesis Supervisor: Professor Munther A. Dahleh
Title: Director, Institute for Data, Systems, and Society

Thesis Supervisor: Dr. Mardavij Roozbehani
Title: Principal Research Scientist, Laboratory for Information & Decision Systems

Acknowledgments

I was very fortunate to have Munther Dahleh and Mardavij Roozbehani as my advisors. They placed a tremendous amount of trust in me to explore research problems that I found compelling. I do not know if I always deserved that trust, but I cannot thank them enough for their willingness to let me explore and stumble my way through research. The breadth of Munther's knowledge is inspiring, and I've learned a lot from his ability to convene and communicate with researchers from substantially different backgrounds. I have valued Mardavij's patience and his careful help. Mardavij has always been willing to dig into the details of a problem with me, to point me in the right direction when I need to learn more, and to reiterate advice that I missed the first or second go-round.

I owe a big thank you to my collaborators. I had wonderful opportunities to work with peers and mentors at MIT and elsewhere, and I had the chance to collaborate with several of them: Scott Burger, Audun Botterud, Ali Kakhbod, Asu Ozdaglar, Cass Sunstein, Chris Knittel, Ignacio Pérez-Arriaga, Frederik vom Scheidt, Jun Jie Joseph Kuan, William Rodriguez, Elaheh Fata, and Igor Kadota.

I have had great colleagues in Munther's group and beyond at MIT. I have enjoyed the camaraderie on the loud side of the office, with Elaheh Fata, Hajir Roozbehani, and Tuhin Sarkar. LIDS has been a friendly home, and I have learned a ton from the brilliant students that surround me. It was fun to watch IDSS blossom into an exciting new entity, and it was a pleasure to learn alongside my colleagues in the first SES cohort. It was wonderful to be part of a vibrant energy community at MIT, including the MIT Energy Initiative and the Center for Energy and Environmental Policy Research (CEEPR). Thank you also to the organizations that funded my work throughout this PhD: the Masdar Institute, Skoltech, the Kuwait Foundation for the Advancement of Sciences (KFAS), the National Science Foundation (NSF), and the Siebel Scholars Program.

My friends—in the Boston area and beyond—provided a tremendous amount of support and fun during graduate school. We have had many adventures and misadventures—skiing on Mt. Washington, skating on Lake Wentworth, dancing at Hong Kong, and even just rock-climbing at the gym.

My parents, Julie and Steve, have provided unshakeable support for as long as I can remember. Carly and Rebecca, my sisters, are inspiring; they are also great skiing and climbing partners. I am lucky to have new parents as well, Jim, Ianthe, and Richard, who welcomed me into their families. Caitlin, my best friend and wife, has been a constant source of love, support, and wisdom. We have had plenty of fun together over the last few years.

Contents

List of Figures

Chapter 1

Introduction

The electric power system is rapidly changing due to the growth of variable renewable generation from wind and solar energy. In the United States, wind and solar energy accounted for 6.5% and 2.2% of total electricity production in 2018, respectively (EIA, 2019b); this represents a five-fold increase from 2009. The growth trend is expected to continue: across a range of policy and macroeconomic scenarios, the U.S. Electricity Information Administration (EIA) projects that renewable generation will grow by 50-115% through 2035 (EIA, 2019a). Electricity generation is also increasingly distributed. Small-scale distributed solar, including rooftop solar, accounted for 25% of net new power capacity in 2018 (EIA, 2019b).

The rapid growth of renewable electricity is driven by public policy, economics, and consumer preference. Tax credits provide federal support for renewable energy, and several states have policies mandating the growth of renewable electricity generation. Wind and solar technology costs continue to fall, and results from the EIA (2019c) suggest that wind and solar energy are cost-competitive with natural gas generation and other fossil fuel technologies. U.S. registered voters broadly support renewable energy policy and growth: 85% of registered voters support requiring electric utilities to use 100% renewable energy by 2050 (Leiserowitz et al., 2018). Policy-makers and voters recognize that increased renewable energy generation can help reduce greenhouse gas emissions and reduce the extent of climate change.

Wind and solar energy sources are uncertain and variable: this poses technical challenges for reliable operation of the electricity grid. The uncertainty of wind and solar energy implies that it is impossible to perfectly predict how much energy supply will be available at a particular time. Furthermore, the energy output from a wind farm or solar array is time- and location-dependent; it changes over the course of a day. These features of variable renewable energy, i.e. wind and solar energy, are critical because the stability of the electricity grid requires a consistent balance of energy injections and withdrawals. Variable renewable energy can increase the need for reserves—generators whose output can be quickly adjusted to account for forecast errors or rapid changes in energy supply or demand (Inman et al., 2013; Bird et al., 2013). Variable renewable energy can also pose challenges to efforts to manage voltage and grid stability (Kroposki et al., 2017). Researchers and electricity market operators are exploring ways to improve renewable energy forecasting, manage energy assets, balance supply and demand, and maintain frequency and voltage stability in electricity grids with large shares of renewable energy (see, e.g., Meyn et al. (2018); Papavasiliou and Oren (2013); Foley et al. (2012)).

The variability and uncertainty of renewable energy also poses economic challenges for the electricity industry. Even though renewable electricity has low cost per unit of electricity produced, it could be expensive to reliably deliver electricity from a system with a large fraction of variable renewable energy (Green and Vasilakos, 2010). The uncertainty of wind and solar energy production is an important cost-driver in markets with a high penetration of renewable resources (Bouffard and Galiana, 2008; Makarov et al., 2009). Given that wind and solar resources are variable, it can be expensive to ensure that sufficient energy is available during periods of low renewable energy production. In order to achieve extremely low emissions in the power system at low-cost, it will be important to reduce the costs of energy storage technologies, to utilize more flexibility from energy consumers, or to increase the flexibility of controllable low-carbon technologies like nuclear generators (De Sisternes et al., 2016; Arbabzadeh et al., 2019). The variability and uncertainty of renewable energy might change the way that producers attempt to exercise market power (i.e. raise prices) in wholesale electricity markets; this could challenge efforts to mitigate market power and to operate efficient markets for coordinating energy resources. Evidence suggests that variable renewable energy can increase price volatility in electric power markets (Rintamäki et al., 2017). This could make it more challenging to finance new generation investments and it could increase the importance of long-term energy contracts.

Solar and wind technologies are creating benefits today, generating energy, decreasing emissions, and satisfying consumer preferences. However, we will require new innovations, regulations, and market designs in order to fully realize the long-term benefits of these technologies. Tools from optimization, game theory, and statistics are invaluable for designing, operating, and researching modern electricity systems. Researchers have used these tools to explain and develop power system theory and practice. New research can help update the ways we manage and regulate the power system to prepare for the growth of renewable energy, and to enable that future at low cost. In this thesis, we tackle three research questions associated with market power and demand participation in a rapidly changing electricity sector. We use analytical tools from statistics, economics, and operations research to study technology and policy opportunities that could help enable a future with reliable, low-cost, and low-carbon electricity. With this research, we try to understand how the independent preferences of consumers and producers can be elicited and coordinated to reduce the cost of a reliable low-carbon power system.

In Chapter 2, we model game-theoretic issues associated with market power and the value of forecast information in systems with stochastic renewable energy. Research on market power traditionally did not focus on generator uncertainty (Joskow et al., 1988; Cardell et al., 1997) because traditional fossil-fuel generators do not have significant resource uncertainty. Improved forecasting capabilities can reduce the uncertainty associated with renewable electricity (Wang et al., 2011; Pinson, 2013), but it becomes important to incentivize accurate and high-quality forecasts in order to minimize power system uncertainty. Market power continues to impact power system regulation, but it takes on new features in markets with high penetration of variable renewable resources. This chapter uses tools from game theory to model and study market power in electricity systems with a large share of stochastic renewable energy, illuminating two key conclusions. First, the level of correlation between energy production in different locations has a major impact on the ability of energy producers to exercise market power and raise prices. Second, information sharing, or high-quality forecasting from market operators, can help reduce market power. These results highlight policy challenges and partial solutions for market power mitigation.

Renewable energy growth will impact, and will be impacted by, forward contracts for energy and financing mechanisms for electricity generating technologies. Forward contracts are agreements to buy and sell assets—in this case, electricity—at a specified future point in time. In Chapter 3, we seek to understand how efforts to introduce retail competition in the electricity sector will impact forward contracting. Forward contracts are an important enabler of new investment in the power sector. The importance of forward contracts will grow as the variance of electricity prices increases, which is an expected outcome of wind and solar growth (Wozabal et al., 2016). Forward contracts can also help reduce market power (Allaz and Vila, 1993). In addition to the renewable energy growth trend, the electricity industry is undergoing deregulation of the electricity retail sector. Supported by changing state laws, individual consumers and municipalities are increasingly able to choose their electricity supplier, or load-serving entity. The term load-serving entity (LSE) encompasses regulated utilities and competitive retail electricity suppliers. Chapter 3 provides an economic analysis that highlights an important interaction between these trends: smaller LSEs, in less concentrated markets, may be less likely to engage in forward contracting. This suggests that additional steps might need to be taken to ensure adequate forward contracting in systems with high renewable energy shares and retail competition.

In Chapter 4, we study issues associated with contracting for incentive-based demand response programs. Growing recognition of the potential benefits of demand response, and the increasing availability of advanced metering and digitally-connected devices, has led to a surge of interest in demand response programs, where participants are paid to reduce or shift demand when doing so reduces power system costs (Chao, 2010a; Faruqui et al., 2017). Time-varying prices for electricity could improve the efficiency of end-use consumption decisions, but for various reasons cost-reflective dynamic prices are not commonly offered for residential and small commercial customers. Demand response programs provide a popular alternative because they incentivize demand reductions but do not penalize overconsumption. However, these features also pose significant challenges to program design. We model the demand response program as a sequential decision problem; an LSE can iteratively improve demand response parameters based on observed customer characteristics. We show that online learning methods could reduce the cost of existing demand response programs and increase customer satisfaction. The proposed design could potentially allow LSEs to offer incentive-based demand response programs by default to all customers. Ultimately, improved demand response programs can help to increase demand flexibility and reduce the costs of efficiently operating a low-carbon power system.

This thesis research exists at the intersection of engineering and economics in power systems. It is part of a growing wave of research that seeks to understand how market interactions influence opportunities for improving power system flexibility and managing variable renewable energy. It treats producers and consumers[1] of electricity as independent, non-cooperative entities. In this research, we try to understand where existing markets might fail, and where new market designs can be developed or adopted, in the service of a more cost-effective, reliable, and low-carbon power system.

The research in this thesis builds on a large body of literature using tools from optimization, game theory, and statistics to develop methods for accommodating variable re-

[1] In the scope of different research questions, a "consumer" could be an end-user of electricity, an aggregator or utility managing electricity consumption on behalf of a large number of end-users, or an algorithm determining consumption on behalf of an aggregator or end-user.

newable energy. We focus on results that are informative and practicable. For example, existing research frequently uses multi-agent optimization models or optimization models with equilibrium constraints. Our focus is to use more heavily stylized models to elucidate fundamental interactions, so the results are informative about key drivers and trade-offs. Existing research in power systems often develops entirely new modes of interaction between power system agents. These results help inspire new options or possibilities, but we prefer to focus on results that could lead to practical improvements in existing systems.

1.1 Summary of Individual Chapters

The following chapters address distinct research challenges, but they share important commonalities in terms of their research focus and theoretical tools. In each chapter, we focus on market design opportunities with the ultimate goal of enabling low-cost, low-carbon power systems. Chapter 2 focuses on energy producers and uses tools from game theory. Chapter 3 describes research that directly implicates both producers and LSEs; it also uses tools from game theory. Chapter 4 focuses on LSEs and consumers; it uses tools from online learning and incorporates elements of strategic behavior. Optimization, uncertainty, and coordinated decision-making are important features of each of the individual chapters.

1.1.1 Strategy and Market Power in Renewable Electricity Systems

Market interactions are a key component of many modern power systems, but the growth of variable renewable energy could pose new challenges for electricity markets. Approximately 70% of U.S. electricity is transacted through wholesale electricity markets. In three large markets,[2] wind and solar energy already account for a quarter of energy production. One primary goal of electricity markets is to coordinate short-term operating decisions with prices, so that the distributed and independent decisions of energy suppliers and consumers will result in an outcome where energy demand is satisfied at least-cost. However, multiple issues compound this challenge: energy storage is expensive, so supply and demand need to approximately match instantaneously; electricity transmission can constrain power flows; the grid operator must procure services to guarantee reliability; and demand is highly inelastic (barely impacted by price) in the short-term. These issues contribute to an additional complication: market power. Market participants, including generators, diversified energy

[2]The three markets mentioned here are ERCOT, which serves most of Texas, CAISO, which serves most of California, and the Southwest Power Pool (SPP), which serves portions of 14 mid-western states.

companies, and energy traders, try to increase electricity prices above competitive levels in order to increase their revenue. Due to the special features of electricity markets, they are often well-positioned to succeed.

Existing research explains how particular features of electricity markets impact strategic behavior, and electricity system operators have practical steps for mitigating market power. Electricity system features like transmission constraints (Cardell et al., 1997), financial transmission rights (Joskow and Tirole, 2000), and market price caps (Joskow and Tirole, 2007) all impact producer market power. System operators have multiple strategies to monitor market power and mitigate its impacts (Pinczynski and Kasperowicz, 2016).

Chapter 2 extends the literature by focusing on producer strategy and market power in systems with high levels of stochastic renewable energy. New energy technologies will impact market power, and they could complicate enforcement efforts to mitigate market power. For example, a regulator might have lower ability to mitigate market power if it has imperfect information regarding production availability of a wind farm (i.e. the amount of energy a wind farm could produce at a particular time) or regarding opportunity costs for a grid-connected battery. Chapter 2 helps explain new drivers of strategic behavior and market power in systems with uncertain production levels due to renewable energy.

In Chapter 2, we model producer decisions in electricity systems with high levels of renewable energy. We use a Cournot model: producers choose a quantity of energy to sell in order to maximize revenue, and their output is constrained by their variable energy availability (e.g., due to changing wind speeds). The model features uncertain availability and asymmetric information; producers do not know how much energy their competitors will be able to produce at a particular time. We use tools from game theory to derive the Bayesian-Nash equilibrium of producer offering strategies when energy availability is uncertain, focusing in particular on the importance of information regarding competitors' production availability.

This chapter provides several contributions that can aid our understanding of producer strategy and market power in systems with high levels of renewable energy. One core contribution of this chapter is to show how the level of correlation between different firms' wind resources impacts strategy and market outcomes. The main insight of the analysis is that increasing heterogeneity in resource availability improves social welfare, as a function of its effects both on improving diversification and on reducing withholding by firms. This insight is robust for common assumptions regarding electricity demand. We present the analysis for a simplified single-period electricity market with wind energy, and then we extend the results to markets with traditional fossil-fuel generation and multiple wind producers. In addition, we analyze the effect of wind resource heterogeneity on opportunities for collusion. Finally, we analyze the impacts of improving public information and weather forecasting; enhanced public forecasting increases welfare, but it is not always in the best interests of strategic producers. This result provides an additional contribution. It extends existing literature on information sharing in Cournot games, which focuses on stochastic production *costs*, by considering information sharing when producers have stochastic production *constraints*.

Practical extensions arise naturally from the results. Greater dispersion of renewable resources is especially valuable when producers have market power, which implies that policy-makers should consider the effects of energy policy on investment location decisions. In addition, high-quality public forecasting can improve efficiency and welfare outcomes in systems with uncertain energy availability and market power. As a result, markets could benefit from high-quality public forecasting or real-time monitoring of wind and solar resources; system operators should consider improving their existing forecast capabilities.

1.1.2 Forward Contracts in Electricity Markets with Concentrated Demand

In Chapter 3, we study how changing features of the retail side of the electricity market impact forward contracting for electricity. Forward contracts are an agreement between two parties to buy and sell an asset at a particular point in time. In the electricity sector, forward contracts can help mitigate price risk for energy suppliers and consumers, and they can help support investment in new generation assets. Forward contracts, including power purchase agreements (PPAs) can provide or enable financing for renewable energy investments (Miller et al., 2018).[3] Forward contracts could be especially important if energy price volatility increases alongside variable renewable energy supply. Forward contracts for electricity can also help reduce market power and increase producer output (Allaz and Vila, 1993), for example in electricity spot markets. Forward contracts are an important tool for investment and risk-management in electricity systems with high levels of renewable energy.

Chapter 3 models an adjacent change in the electricity sector—the increase of retail choice—to understand its impacts on forward contracting. Consumer electricity markets are open to retail electricity competition in many locations, including 17 states. Retail competition has potential benefits for consumers, including lower prices and increased opportunities

[3] On a shorter time-scale, the day-ahead energy market also provides a type of forward contract for the purchase and sale of energy.

to support renewable energy. In states like Massachusetts and California, communities have the option to collectively choose an energy supplier to serve their residents (Faulkner, 2010); as of 2018, over 150 towns and cities in Massachusetts had chosen this approach (Fuller and Berwick, 2018). In 2017, over 13.7 million households participated in retail choice programs, double the number from a decade prior (EIA, 2018). The main contribution of this chapter is to show how retail changes could impact forward contracting and producer market power in the electricity sector.

In Chapter 3, we develop a simple model to show how LSEs will engage in forward contracting in order to minimize their overall costs of serving electricity to their retail customers.[4] In this model, one of the benefits of forward contracting is that it helps reduce market power, thus reducing spot market prices. The main contribution of this chapter is to show how positive externalities for forward contract procurement can arise: the benefits of forward contracting to reduce market power lead to positive externalities because they are shared by all LSEs, not just those who engage in the forward contract. As such, the total forward contracting level and total welfare decrease in the number of LSEs serving the consumer market. We show this dependency using comparative statics of the market equilibrium, based on the number of retail electricity suppliers. This insight suggests new areas for additional research, for instance to study empirical evidence of the described effects or to explore relevant policy interventions. Policies that incentivize or require forward contracting might be a useful tool for reducing market power in electricity systems with retail competition.

1.1.3 Incentive-Based Demand Response

In Chapter 4, we examine new ideas for enabling flexibility in electric power systems through demand response. Demand response is a tool whereby consumers are paid to reduce their energy consumption in times of need, for instance, when energy prices are high or renewable energy supply is low. Demand response will be especially valuable in systems with high levels of renewable energy generation (Roos and Bolkesjø, 2018).

Increased demand-side participation in electricity markets could improve social welfare while reducing the cost of integrating renewable energy. The marginal cost of electricity can fluctuate by two orders of magnitude. Dynamic pricing of electricity, in line with temporal

[4]In practice, there are other reasons that a LSE might engage in forward contracts, including risk reduction. We ignore those issues to focus on the relationship between forward contracting and market power. In the context of this chapter, we can think of retail competition as impacting the residual level of forward contracting, after accounting for other factors that drive forward contract procurement.

changes in the marginal cost of electricity production, can reduce distortions and improve economic efficiency (Borenstein and Holland, 2005; Joskow and Wolfram, 2012). However, dynamic (or 'real-time') prices are not a wide-spread option for residential and small commercial customers. Only 0.2% of U.S. residential customers are offered real-time prices (EIA, 2018). Inertia, consumer preferences, transaction costs (Schneider and Sunstein, 2017), and consumer protection concerns (Burger et al., 2020) might limit the practical usefulness of a real-time electricity price. One alternative mechanism is to pay customers for demand reductions when the marginal cost of energy is high. These mechanisms are called 'demand response' programs.

Incentive-based demand response programs have sizable participation in the United States, but their design implies fundamental challenges. Incentive-based demand response programs, alternatively called 'behavioral demand response,' have 2.8 million participants in the U.S. (Surampudy et al., 2019). Peak-time rebate programs also have sizable participation; this Chapter 4's model encompasses both types of program. These programs are useful because they do not require any special technology on the part of customers, and they do not require direct control by an electric utility. They are popular, in part, because they have no financial downside for customers. However, demand response program operators face substantial challenges: they need to decide how to incentivize customers' reductions, but they cannot perfectly measure reductions because they have imperfect information about what the customer would have consumed in the absence of the demand response incentive. LSEs that operate demand response programs typically determine a baseline threshold for each customer in each program hour; they only pay customers for reductions below the baseline threshold. If this threshold is too high, program operators pay more for the demand response program; if it is too low, they risk customer dissatisfaction or defection. Despite challenges, demand response remains a practical and popular tool for LSEs and regulators.

In Chapter 4, we study demand response procurement in the electricity sector. Consumers have the option to buy electricity at a fixed retail rate, which leads to inefficiently high demand when the marginal cost of producing electricity exceeds the retail rate. An LSE can develop a demand response program to reduce consumption. Typically, the LSE has imperfect information, and the LSE can only incentivize customers; it cannot charge them more than the regulated retail rate. The combination of (weakly) positive incentives and private information makes it challenging to determine the optimal demand response incentive. Typical practice is to set the baseline threshold as an estimate of the customer's counter-factual consumption, and to pay customers based on the difference between their final consumption and that baseline threshold. We consider two customer demand models and investigate natural objective functions for the LSE. From this starting point, it is clear that the baseline should not necessarily be an estimate of the customer's counterfactual consumption. We utilize tools from online learning to explore customer-specific cost functions and iteratively choose better demand response baseline thresholds.

Chapter 4 describes several key contributions that could help enable more effective incentive-based demand response programs. First, we develop natural models of consumer behavior and objectives for demand response programs. We argue that the demand response baseline threshold should not necessarily be an estimate of a customer's counter-factual demand. In order to maximize the LSE's objectives, the baseline threshold will take into account a customer's propensity for demand reductions and the uncertainty (variance) in the LSE's estimate of the customer's counter-factual demand. Second, we show how tools from online learning can be used for the sequential decision problem of choosing customer baseline thresholds. We develop two separate customer models and show how differences in customer responsiveness might impact the design of the sequential decision algorithm for choosing the demand response baseline parameter. We use numerical examples to showcase the potential benefits of these improvements over current practice.

Chapter 2

Selling Wind: Strategic Behavior in Electricity Markets with Substantial Renewable Energy Generation

This work was performed in collaboration with Asu Ozdaglar and Ali Kakhbod.

2.1 Introduction

The market share and total production of renewable electricity is growing rapidly. In 2018, wind energy was responsible for 6.5% of U.S. electricity generation, nearly doubling its market share and total production from five years prior. Renewable electricity is a critical component of global efforts to reduce carbon dioxide emissions, and its cost is rapidly declining.

Prominent sources of renewable electricity - wind and solar energy - have stochastic resource availability: it is not possible to perfectly predict the quantity of wind or solar power available at any given point in time. The associated spatial and temporal variability of renewable energy resources has a significant impact on their value to society (Joskow, 2006; Hirth, 2013; Hirth et al., 2016). Furthermore, since wind production reduces local prices due to the merit order effect, highly correlated local wind energy availability reduces the average value of wind energy produced (Woo et al., 2011; Ketterer, 2014).

Existing literature focuses on strategic behavior in electricity markets without substantial amounts of renewable energy. Research on market power in the electricity sector (Joskow et al., 1988) provided important insight for electricity system deregulation. Electricity system market power research does not traditionally focus on stochasticity because fossil-fuel generators do not have significant resource uncertainty. Instead, it focuses on other key features of the electricity sector that impact market power, like transmission constraints (Cardell et al., 1997), financial transmission rights (Joskow and Tirole, 2000), and market price caps (Joskow and Tirole, 2007). Acemoglu et al. (2017) establish that diverse ownership portfolios of renewable and thermal generation by strategic firms may be welfare reducing, because they can reduce (or even neutralize) the merit order effect. Butner (2018) provides empirical evidence of these effects. Since we model an extreme case of competition with only wind producers, our firms strategically withhold wind energy. In practice, when firms own diverse generation portfolios, they will prefer to withhold output from resources with high marginal costs (Acemoglu et al., 2017). Our model explains how information about production availability influences strategy; it can be combined with the existing literature to help explain producer strategy in systems with diverse ownership portfolios and stochastic, correlated, production constraints for renewable energy. Our model also helps show how public information sharing can improve welfare in systems with strategic behavior and stochastic renewable energy production.

We are interested in how a particular characteristic of renewable energy resources–the stochastic dependence of resource availability across firms–impacts strategic behavior, market power, and welfare. The link between stochastic heterogeneity[1] of resource availability and welfare is an important area for research because various policies impact the investment strategies of wind producers and therefore the stochastic characteristics of the wind energy portfolio in a given region (Kök et al., 2016; Schneider and Roozbehani, 2017b). Common subsidy forms for renewable energy, like the production tax credit (PTC) and state-level renewable portfolio standards (RPS), impact renewable energy investments (Fischer, 2010).

Figure 2-1 shows probability distributions for two different wind farms in the Midwest United States, conditioned on the output of a third wind farm i in the same region. The nature of stochastic dependence is very different for each pair of wind farms. Wind farm i's output is highly correlated with the output of the wind farm displayed on the right, and essentially uncorrelated with the output of the wind farm displayed on the left.[2]

[1]Stochastic heterogeneity refers to the level of stochastic dependence. Throughout the main body of the chapter, we use a term "dispersion" to succinctly refer to the extent of stochastic dependence. In the linear case, high (low) stochastic dependence is equivalent to highly correlated (uncorrelated) stochastic resource availability across different producers.

[2]For the purposes of this graph, we use measured production as a proxy for resource availability. Here and throughout the chapter we stylize resource availability as discrete, i.e. the resource availability at i $w_i = L$ or $w_i = H$. Since real world availability is continuous, for this figure, we say that $w_i = L$ when the wind availability is less than 3% of its maximum availability (bottom 27% of periods) and $w_i = H$ when the wind availability is greater than 67% of its maximum availability (top 20% of periods).

Clearly, policy changes can impact investment strategies for renewable energy and the characteristics of system-wide resource uncertainty. This begs the questions: *Is it important to encourage policies that increase the heterogeneity of stochastic resources? Will investment in wind energy naturally lead to the level of resource heterogeneity that maximizes social welfare?* Just as policy makers seek to limit market concentration in certain industries, they might support policies to increase stochastic heterogeneity of renewable resources in the electric power industry. These efforts have growing import because existing strategies for market power monitoring in electricity markets will be challenged by an influx of renewable generation. Regulators have imperfect information regarding resource availability and risk preferences for firms that own stochastic generation. Regulators also have imperfect information regarding opportunity costs for storage facilities that are proposed to mitigate the variability of renewable resources.[3]

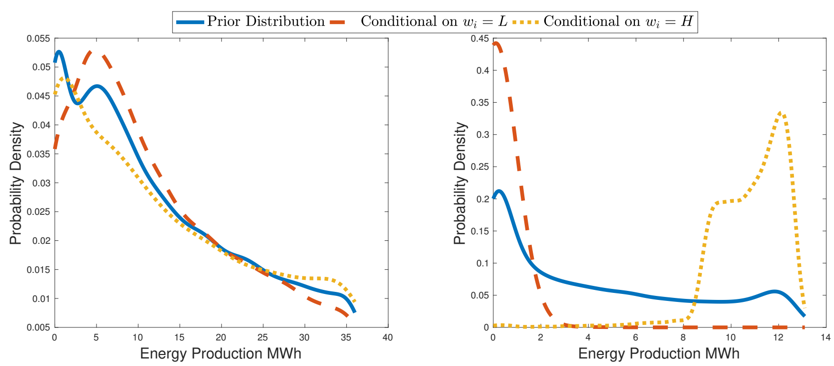

Figure 2-1: Prior and conditional empirical distributions for resource availability from two wind farms in Midwest U.S., based on hourly energy production from 2014 through 2016.

We study strategic firms participating in a Bayesian Game, where firms have private information regarding their realized energy availability, or "state." This energy availability is equivalent to a production constraint, because it limits the extent of production by the firm in any given period. Since the resource availability of wind energy is uncertain, from an individual firm's perspective its competitors' production constraints are stochastic. However, the resource availability of wind energy has a high degree of stochastic dependence; firms can gain important information about their competitors' production constraints from the realization of their own resource availability. As such, the extent of stochastic dependence regarding firms' resource availability becomes an important factor that impacts strategic behavior, market power, and welfare. For clarity, we focus on wind energy, but the insights

[3]Additionally, Munoz et al. (2018) discuss challenges associated with auditing the opportunity costs of traditional generators in markets with physical inflexibilities and non-convex costs. These issues have growing import in systems with high levels of renewable energy.

can be extended directly to solar energy or any other resource with stochastic availability and negligible marginal costs. Since solar and wind resources are not highly correlated, a market with a mix of solar and wind generation probably has greater stochastic heterogeneity than a market with only wind energy and no solar energy. Similarly, a market with a mix of onshore and offshore wind resources probably has greater stochastic heterogeneity than a market with only onshore or offshore wind resources.

We model producer competition as an incomplete information Cournot game with correlated types, where the type refers to the stochastic resource availability (production constraint) that is private information for each individual producer. The base model uses a Cournot duopoly market. We utilize a parameter d to represent the level of heterogeneity amongst wind producers; throughout, we refer to d as the level of dispersion. Intuitively, we can think of dispersion as being similar to geographic distance; research has shown that correlation in wind availability across pairs of wind producers is generally decreasing in geographic distance (Sinden, 2007). This is a useful intuition but not a general rule; distance is only one feature among many (e.g. geography, climate, turbine orientation) that could impact the level of stochastic heterogeneity across wind farms.

The results provide clear insight to explain how stochastic resource heterogeneity can impact welfare in imperfect electricity markets. Increasing heterogeneity in wind resource availability is beneficial for two distinct reasons: it increases the diversification of resources, and it also reduces strategic withholding because it changes the information that a producer's own energy availability provides about the likely energy availability of the other firms in the market. The results of our model imply that imperfect competition in energy markets can affect investment in renewable energy, resulting in a system with sub-optimal levels of resource heterogeneity.[4]

Next, we investigate the effects of public sharing of high-quality weather forecasts, using the limiting case where the true realized energy availability of firms is monitored and shared. Information sharing through improved forecasting is socially beneficial, but it does not always improve producer profits. As such, it will not necessarily be undertaken by producers acting in their own best interest.

This result is conceptually similar to the information sharing literature; we show that it is upheld in our model where production constraints (not costs) are stochastic and corre-

[4]Other features of electricity markets might also reduce heterogeneity of wind resources or distance between wind producers, including the presence of existing transmission, variance in state-level renewable policies, and quantity-based subsidies. However, we focus specifically on market failures due to imperfect competition.

lated across wind farms. Since producers have stochastic availability, high quality weather forecasting can be undertaken publicly in order to maximize social welfare.

Finally, the model is utilized to examine the effects of heterogeneity on collusion and on policies to prevent collusion. If they do not face potential penalties for collusion, firms with stochastic availability will always choose to collude because they benefit from sharing information and from sharing monopolistic profits. Increasing heterogeneity of wind production has a range of impacts on collusion, impacting its value to producers, the costs of collusion on social welfare, and the level of enforcement required to prevent collusion.

These results provide a framework for evaluating policies that impact investment and information-provision in imperfectly competitive markets where firms have stochastic production constraints, like electricity markets with a high penetration of renewable resources. The results can help us understand how policies that impact the dispersion of renewable energy resources, and thus the characteristics of stochastic energy availability, ultimately impact welfare in imperfectly competitive electricity markets.

2.1.1 Literature Review

Literature on wind diversification has focused on the impacts of resource heterogeneity on average electricity prices and the cost of wind integration. Increased heterogeneity of wind resources has at least three impacts on economic surplus and the cost of electricity:

1. Balancing costs for managing wind production. *This impact is well-studied in the literature and not covered in the model in this chapter.*

2. The average benefit of wind production and the average price earned for electricity produced from wind energy. *This is discussed in some existing literature, but we provide a new formal model that provides insights on its impact on welfare.*

3. Strategic curtailment by wind and traditional power producers. *This impact has not been proposed in previous literature. This chapter formalizes the concept and explains its impact on welfare.*

First, increased wind heterogeneity decreases balancing costs because it reduces hour-to-hour fluctuations in total wind energy production and because it reduces net uncertainty of availability in a given hour. This impact is well-studied in existing literature. Fertig et al. (2012) show that increasing diversification reduces the average hourly fluctuation in total power output and increases the equivalent firm power production.[5]

These short-term time-dependent impacts are not covered specifically in our model, which ignores the complexities associated with sequential market clearings.[6] As such, our model ignores issues associated with integrated control management systems and technical issues associated with the management of power systems with high levels of renewable resources. However, the benefits of increased heterogeneity likely have a net positive impact on welfare due to a reduction of balancing costs, so they do not change the general direction of the main welfare results.[7]

Second, increasing heterogeneity increases the average benefit of wind production. Increasing levels of wind generation have been shown to reduce prices in Germany and in West Texas (Ketterer, 2014; Woo et al., 2011). In general, wind has declining marginal benefits because of the convexity of the electricity supply curve from traditional generators, which serve as strategic substitutes for wind energy. We model this effect by assuming that utility obtained from consuming wind energy is concave. Our model mirrors the basic empirical result; when more wind is produced, in a single period or on average, lower prices result. While existing research focuses on the price impacts of additional wind penetration, the stylized model described herein allows us to extend the results by focusing specifically on the impacts of heterogeneity on welfare.

Finally, the third impact of higher wind heterogeneity is its effect on the ability of wind generators to strategically withhold their energy production, when these generators have market power but some uncertainty regarding their competitors' production constraints.[8] To our knowledge, this chapter represents the first time that this effect of resource heterogeneity has been proposed and analyzed.

[5] There are a range of additional costs for managing and controlling grid systems with a high level of variable renewable energy resources (Camacho et al., 2011). Additional research focuses on reducing the costs of wind integration, for instance by improving models for unit commitment in the face of supply and demand uncertainty (Papavasiliou and Oren, 2013; Cerisola et al., 2009; Khazaei et al., 2017).

[6] Prices in sequential markets are also impacted by market power, which helps explain why prices in sequential markets sometimes diverge (Ito and Reguant, 2016).

[7] Other research focuses on how complementary technologies impact wind integration by studying the effects of energy storage on wind energy commitments (Kim and Powell, 2011) and transmission planning (Qi et al., 2015) in markets with wind.

[8] The simplest way that wind producers could withhold their energy production is by changing the blade angle of wind turbines to decrease their energy production below the maximum available energy. This is the context we consider in this paper, and we refer to this as 'curtailment' of wind energy. It implies that available energy is essentially wasted in an effort to increase market prices. The introduction of energy storage devices, coupled with wind farms, could allow operators to withhold energy output in a particular time period by charging a co-located battery; this energy could be sold to the grid in a future period. Thus, the growth of large-scale energy storage on the electricity grid could increase opportunities for short-term withholding of energy production.

This chapter uses a Cournot model to analyze the effects of dispersion on bidding behavior and welfare in a market with stochastic energy availability and private information. The Cournot assumption provides a simple model of imperfect competition, which is an important feature of markets with renewable generation: as firms operate more wind and solar generation, it will become increasingly difficult to prevent the exercise of market power, due to the uncertainty in underlying resource availability. The Cournot model is a useful simplification of the electricity market. In practice, firms submit a supply function that specifies how much energy they are willing to offer at any given price.[9]

We model strategic firms participating in a Cournot-Nash game with incomplete information, which is a specific form of a Bayesian game.[10,11] There is a substantial economics literature on Cournot-Nash games. Szidarovszky and Yakowitz (1977) and Gaudet and Salant (1991) provide useful constructive results on the existence and uniqueness of equilibria in Cournot-Nash games with complete information. Novshek (1985) proves that a Cournot equilibrium exists whenever the marginal revenue for any particular firm declines in the total output of the other firms.

Regarding Cournot-Nash games with *incomplete information*, Einy et al. (2010) survey the literature and explain the conditions for equilibria existence and uniqueness. Nearly all of this literature treat firms' objective functions as stochastic; we focus on the case where production constraints are stochastic. Richter (2013) also study the topic of Cournot games with stochastic production constraints and incomplete information, but they focus on firms that are stochastically independent. Stochastic dependence of the firms' production constraints has major impacts on the results, including the value of information sharing. Our research does not focus on conditions for existence; reasonable assumptions for the electricity sector (including the possibility of negative prices) generally lead to the existence of equilibria. Instead, we focus on developing new results to link the extent of stochastic dependence to strategic behavior and welfare in the equilibria.

Another rich research area discusses information sharing in Bayesian games; this re-

[9]The Cournot setup is considered a good approximation to real-world electricity markets (Hogan, 1997; Oren, 1997; Borenstein et al., 1999; Willems et al., 2009). There are other ways to model producer offers in electricity markets, including supply function offers (Anderson and Philpott, 2002; Holmberg, 2007). Wolfram (1998) and Hortacsu and Puller (2008) offer empirical analyses of strategic bidding in multi-unit electricity auctions. Willems et al. (2009) and Ventosa et al. (2005) discuss the comparative benefits of Cournot models versus the full supply function model.

[10]See Fudenberg and Tirole (1991), chapter 6, for excellent overview of games with incomplete information.

[11]This framework implies that players have full knowledge of the joint distribution governing players' uncertain parameters. This assumption is contested in some models, but it is a fairly natural assumption to make in the electricity sector, where resource availability is largely based on weather, with publicly available weather information and public data on past production.

search is particularly applicable to our Section 2.7. Clarke (1983) studies information sharing in a Cournot game and concludes that firms only have an incentive to share information if they can cooperate on a strategy once information is shared. Vives (1984) discusses information sharing in Cournot and Bertrand games, where players receive imperfect (and correlated) signals about the intercept of the demand function. Sakai (1985), Gal-Or (1986), and Shapiro (1986) study information sharing in the case of private cost functions. Unlike the previous literature, we do not assume that the demand function is linear. In these papers, information sharing between firms does not impact average production and average price; in our model, information sharing does impact average production and average price. Einy et al. (2003) study the value of public information in a Cournot duopoly, with more general forms of demand and cost uncertainty. Unlike the aforementioned research, our model has stochastic production constraints. This could be equivalently modeled as a stochastic cost function with infinite costs for quantities above the constraint, but the existing research focuses on affine cost functions, so existing results can not be directly extended to analyze our model. Richter (2013) considers the case of information sharing with stochastic and independent production constraints. Our results are conceptually similar, but in a model where production constraints are not necessarily independent.

The main idea of our research is to formalize game-theoretic equilibria where producers have stochastic and dependent production constraints, in order to examine the effects of correlation in resource availability on the resulting equilibria. We consider the case of multiple wind producers offering their energy into markets, when their maximum availability is stochastic and correlated amongst producers. Existing research studies energy market equilibria in other ways. For instance, Hobbs and Pang (2007) examine the effects of joint constraints and non-smooth demand functions. Downward et al. (2010) and Yao et al. (2008) study Cournot equilibria in markets with transmission constraints. de Arce et al. (2016) study the effects of Cournot competition on the efficacy of renewable energy incentives. Schneider and Roozbehani (2017a) study strategy and competition in two-stage day-ahead and real-time electricity markets. Gilotte and Finon (2006) and Pineau et al. (2011) study investment in energy markets with Cournot competition. Most similar to our work, Fabra and Llobet (aper) model competition among firms with stochastic production constraints. In their model, firms submit a price, quantity pair; demand is price-inelastic. Our firms compete only with quantities, and we focus on the comparative statics of firm heterogeneity. Compared to the aforementioned literature, we abstract many important notions of real-world electricity systems in order to clearly focus on our question of interest, which is not addressed in the existing literature: *how does heterogeneity in stochastic renewable energy availability impact market power and social welfare?*

Section 2.2 introduces the benchmark duopoly model, and Section 2.3 describes the features of the Cournot equilibrium. Section 2.4 describes the impacts of wind heterogeneity on the diversification of wind energy and on strategic curtailment by wind producers. These impacts drive many of the main results presented herein. Sections 2.5 and 2.6 describe the effects of heterogeneity on social welfare, price, and profits in the duopoly model. Section 2.7 describes how the level of heterogeneity impacts the likelihood that firms will choose to publicly share information, and shows that public information sharing is always socially beneficial. Section 2.8 examines the effect of heterogeneity on collusion and on the cost of efforts to prevent collusion. Sections 2.9 and 2.10 extend the results to the case of an oligopoly market with multiple wind generators and with both wind and traditional generators, respectively.

2.2 Benchmark Model

Consider two wind energy producers engaged in imperfect competition, operating two locally separate wind farms to generate energy. For each producer i, the maximum available wind energy, w_i, is **stochastic** and might be either H (high) or L (low), with $H > L$ and with probability $\Pr\{w_i = H\} = \beta = 1 - \Pr\{w_i = L\} > 0,\ i \in \{1, 2\}$. When $w_i = H$ ($w_i = L$), we say that producer i is in the high (low) state. Let $d \in [0, 1]$ be the **dispersion** between the two wind producers, where the maximum dispersion is normalized to 1. The parameter d captures the extent of **heterogeneity** in terms of wind energy availability for these wind producers. When d is small, wind energy availability is highly correlated amongst wind producers. When one wind producer in the high state, the other wind producer is likely to be in the high state as well (similarly for the low state). However, when d is high these locations become highly heterogeneous in terms of wind availability, so that extent of wind energy available to one producer does not reveal much information about the other wind producer's availability. In the case of high heterogeneity, the extent of wind energy availability is nearly or (when $d = 1$) fully independent across wind producers.

This section models the joint probability distribution of the available wind energy in a simple parameterized form. Precisely, for $i, j \in \{1, 2\}$, the conditional probability of high wind availability is given by (2.1).

$$\begin{aligned}\Pr\{w_i = H | w_j = H\} &= \frac{\beta}{\beta + d(1-\beta)} \\ \Pr\{w_i = H | w_j = L\} &= \frac{d\beta}{\beta + d(1-\beta)}\end{aligned} \tag{2.1}$$

When the wind producers are "far" from each other, $d = 1$, we are in the limiting case of independent production; $\Pr\{w_i = H | w_j = H\} = \Pr\{w_i = H\} = \beta$ and $\Pr\{w_i = H | w_j = L\} = \beta$. On the other hand, when they are locally "close", $d = 0$, we are in the full information case and $\Pr\{w_i = H | w_j = H\} = 1$.[12] Section 2.9 extends the results for arbitrary joint probability distributions for wind availability from multiple producers.

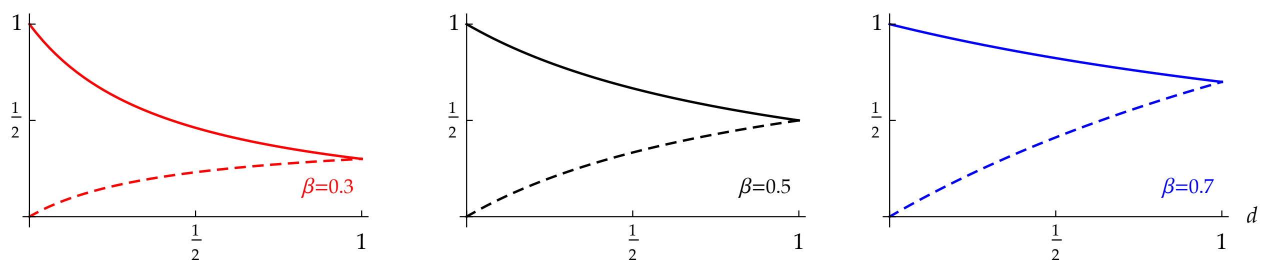

Figure 2-2: The conditional probability distributions from (2.1) for $d \in (0, 1)$, for different values of the prior β. For each graph, the solid line represents $\Pr\{H|H\}$ and the dashed line represents $\Pr\{H|L\}$.

Note that the extent of the difference between the high and low states corresponds to the extent of the variance in wind availability for each individual producer. For instance, if we fix the value of H (e.g. as the capacity value of the wind producer), then variance of wind energy availability $\text{Var}(w_i)$ is monotonically decreasing in L.

Let q_i denote the amount of wind energy generated by producer $i \in \{1, 2\}$. We assume the inverse demand $P : \mathbb{R} \to \mathbb{R}$ as a function of total supply $Q = q_1 + q_2$ is concave and downward sloping, i.e., $P'(Q) < 0, P''(Q) \leq 0$ for all Q. The marginal cost of production via wind energy is negligible. Our model simplifies the electricity market model by focusing on a single real-time market with inverse demand $P(Q)$.[13] We ignore the impacts of short- and long-term forward markets, e.g. day-ahead markets and capacity markets. While these

[12] In this context, we can think of β as a forecast of the energy availability for each firm, using public information or information with negligible cost. When a firm realizes its own (private) energy availability, this new information changes its forecast of the energy available to its competitors, as shown in (2.1).

[13] $P(Q)$ can be thought of directly as inverse demand, or as the net inverse demand arising from price-inelastic demand and a competitive fringe.

markets are undoubtedly important, we focus on the real-time spot market because planned real-time bidding behavior ultimately informs strategy in forward markets.[14]

The producers compete in Cournot fashion. According to its **private information** about its maximum available wind, $w_i \in \{H, L\}$, producer i chooses $q_i(w_i)$ maximizing the expected value of its profit π_i, conditional on its realization of w_i:

$$\begin{aligned} \mathbb{E}_{w_j}[\pi_i | w_i] &= \mathbb{E}\left[q_i \; P(q_i(w_i) + q_j(w_j)) \; | w_i\right], \\ \text{s.t.} \quad & q_i(w_i) \in [0, w_i] \end{aligned} \tag{2.2}$$

2.3 Equilibrium

To focus on the case where wind producers produce at full capacity in the low state (i.e. no curtailment when $w_i = L$), and to avoid equilibria where wind producers produce at full capacity in both states we adopt the following assumption:

Assumption 1. Let $P(\cdot)$ be the inverse demand. Then $P(2L) + LP'(2L) > 0$ and $P(H) + HP'(H) < 0$.

This assumption allows us to focus on equilibria where producers exercise strategic withholding in one state but not in the other. This represents the case of interest where the stochastic nature of the wind resource has important impacts on the equilibrium strategy. Under the above assumption the equilibrium is characterized as follows.

Proposition 1. Let $P' < 0, P'' \leq 0$. Then, there exists a unique symmetric Bayesian Nash Equilibrium (BNE) such that

$$q_i(w_i) = q(w_i) = \min\{w_i, \phi\} \qquad w_i \in \{L, H\},\ i \in 1, 2 \tag{2.3}$$

where $\phi > L$ is the unique root of the following equation

$$\Pr\{L|H\}\big[P(L+\phi) + \phi P'(\phi + L)\big] + \Pr\{H|H\}\big[P(2\phi) + \phi P'(2\phi)\big] = 0. \tag{2.4}$$

The proposition establishes the symmetric Bayesian Nash Equilibrium (BNE) for the benchmark model.[15] In this equilibrium, firms produce L in the low state and $\phi < H$ when

[14] While we refer to w_i as the wind energy availability for firm i, readers can also think of w_i as the realized error, the difference between energy availability in real-time and the day-ahead offer or prior forecast.

[15] Instead of $P'' \leq 0$, weaker assumptions of the form $P'(Q) + QP''(Q) \leq 0, Q \in [2L, 2H]$ would also be sufficient to guarantee the existence of the equilibrium; see Gaudet and Salant (1991). They would also be sufficient for Lemma 1 and the subsequent results in the main body of the chapter. We use the concavity assumption for simplicity in this chapter.

they are in the high state. The proof for this proposition is provided in the chapter appendix, Section 2.12; throughout the chapter, all omitted proofs are in Section 2.12. The intuition is that in the symmetric equilibria producers curtail based on the expected value of the first order condition, given the uncertain state of their competitor and their competitor's equivalent strategy.

Example (Linear inverse demand): To clarify understanding regarding Assumption 1, consider the case of linear inverse demand, i.e. $P(Q) = s - Q$, where Q denotes the sum of the firms' production $Q = q_1 + q_2$. Suppose there is no capacity constraint; then there exists a unique symmetric equilibrium q_C (the Cournot level) in which the optimal production is

$$q_1 = q_2 = q_C \equiv \frac{s}{3} < q_M \equiv \frac{s}{2}, \tag{2.5}$$

where q_M is the corresponding Monopoly level.[16] Thus, with linear inverse demand, Assumption 1 simply says that L is lower than the Cournot level and H is higher than the monopoly level, i.e.

$$L < q_C < q_M < H. \tag{2.6}$$

If the first part of Assumption 1 is violated, wind producers always produce at the Cournot level q_C; the stochastic nature of the wind resource has no impact on the equilibrium strategy. If the second part of the assumption is violated, then wind producers would curtail in any situation, even absent a competitor. Moreover, with linear inverse demand, under Assumption 1, the equilibrium can be explicitly characterized as follows.

Corollary 1. Let the inverse demand be linear, i.e. $P(q_1 + q_2) = s - q_1 - q_2$. Then, there exists a unique symmetric pure-strategy Bayesian Nash equilibrium such that

$$q_i(w_i) = q(w_i) = \min\{w_i, \phi\}, \quad w_i \in \{L, H\},\ i = 1, 2, \tag{2.7}$$

where $\phi \equiv \frac{s\beta + (s-L)(1-\beta)d}{3\beta + 2(1-\beta)d}$.

The subsequent sections introduce key effects that drive the impacts of d on the equilibrium and its resulting impacts on welfare, price, and profits.

[16] Note that q_C is the optimal strategy when $\pi_i = q_i(s - q_1 - q_2)$ and q_M is the optimal monopoly quantity maximizing $\pi = q(s - q)$.

2.4 Strategic Curtailment and Diversification

This section explains useful Lemmas to help illustrate the two major impacts of dispersion d in the strategic setting. Recall that $\phi = q(H)$ is the production when a firm is in the high state. When a firm is in the low state it produces L. Unless otherwise specified, all of the following results hold for concave and downward inverse demand functions (i.e. $P' < 0, P'' \leq 0$) satisfying Assumption 1.

Lemma 1. As d increases, production in the high state increases, i.e. $\frac{\partial \phi}{\partial d} > 0$.

The intuition derives from the fact that the outputs of the wind producers are strategic substitutes because $P' < 0, P'' \leq 0$. Therefore, the best reply for firm i is decreasing in firm j's production, and firm i's best response is a decreasing function of $\mathbb{E}[q_j | w_i = H]$. When d increases, the possibility that the wind producers are in different states increases. Thus, the probability that firm j is in the low state increases, given that firm i is in the high state, and $\mathbb{E}[q_j | w_i = H]$ decreases, increasing ϕ which is firm i's optimal production when it is in the high state.

Lemma 2. As d increases, the expected value of total production increases: $\frac{\partial \mathbb{E}_{w_1,w_2}(Q)}{\partial d} > 0$.

The *a priori* expected value of firm i's production is $\mathbb{E}_{w_1,w_2}(q_1) = \beta\phi + (1-\beta)L$, and the expected value of total production is just the sum of each firm's expected production: $\mathbb{E}_{w_1,w_2}(Q) = \mathbb{E}_{w_1,w_2}(q_1) + \mathbb{E}_{w_1,w_2}(q_2) = 2\beta\phi + 2(1-\beta)L$. Only ϕ on the right-hand side is a function of d; the parameter d has no effect on the prior probability β of being in the high state. Then $\frac{\partial}{\partial \phi}\mathbb{E}_{w_1,w_2}(Q) = 2\beta\frac{\partial \phi}{\partial d}$, with $\beta > 0$, $\frac{\partial \phi}{\partial d} > 0$ (from Lemma 1), which concludes the proof.

Introducing Strategic Curtailment (SC) and Wind Diversification (WD). These two features describe the effects of d on, respectively, high state output ϕ and on the joint probability distribution of the resource availability amongst all producers. The main effects of d, for instance on welfare, are driven by its impacts on strategic curtailment and wind diversification.

- **Strategic Curtailment (SC)**: When d increases it impacts the information available to the wind producers as strategic decision makers. As a result, as d grows, the production of firm i when they are in the high state increases (Lemma 1). Equivalently, this increases the expected value of production (Lemma 2), and decreases the level of **strategic curtailment**, the difference between the expected value of availability and the expected value of production, i.e. $\mathbb{E}[w_i - q_i]$. Thus when d grows the level of **strategic curtailment** decreases because increasing d reduces firm's strategic withholding of available production in the high state.

- **Wind Diversification (WD)**: When d grows the probability of being in different states increases. Consequently, with increasing d firms produce different quantities with a higher probability, improving diversification of the total portfolio of wind producer assets and reducing the variance of the total availability of wind energy $\text{Var}(w_1 + w_2)$. When utility is (weakly) concave, diversification of wind assets (weakly) increases welfare.

In our results, we frequently assess the impact of diversification on various functions (e.g. the welfare function). Here we define a measure of that impact in order to convey the chapter's results more succinctly. Let $f : R^2 \to R$. The impact of diversification on f (denoted by WD_f) is given by the following expression,

$$WD_f \equiv f(x, y) + f(y, x) - f(x, x) - f(y, y) \tag{2.8}$$

where $y > x > 0$. If f is linear, i.e. $\exists\, a, b, c \in R$ such that $f(x, y) = ax + by + c$, then $WD_f = 0$. Furthermore, if f is a concave function of the sum of its arguments, i.e. if $f(x, y) = g(x+y)$ for some $g : R \to R$ where $g'' < 0$ then $WD_f = 2g(x+y) - g(2x) - g(2y) > 0$.

Many of the results presented here are due to the interplay between the effects of d on strategic curtailment and diversification as introduced above. In general, increasing d improves social welfare through its effects on both diversification and strategic curtailment. However, because increasing d decreases strategic curtailment, it can sometimes reduce profits for wind producers. This suggests that the level of heterogeneity sought by profit-maximizing investors can be lower than the welfare-maximizing level.

2.5 Social welfare vs. Dispersion

Since the marginal cost of energy production from wind is negligible, welfare (i.e. firms' surplus plus consumers' surplus) is equivalent to the consumers' net utility of consumption. Let $U(Q)$ be the consumer utility, where $U(0) = 0$ and $\forall Q, U'(Q) > 0, U''(Q) \leq 0$. Note that $U'(Q)$ defines the inverse demand $P(Q)$. The consumer surplus is given by $CS = U(Q) - Qp$, and welfare is $W = \pi_1 + \pi_2 + CS = U(Q)$.

Proposition 2. The expected value of welfare increases in dispersion d.

The expected value of welfare is given by $\mathbf{E}_{w_1,w_2}[W] = \mathbf{E}_{w_1,w_2}[U(q_1 + q_2)]$. By the product rule of differentiation, the total impact of d on welfare is exactly the sum of its impacts on $\mathbb{E}[W]$ through strategic curtailment and wind diversification. Increasing d increases the expected value of welfare because it decreases strategic curtailment and increases diversification, which **both** increase U.

Increasing d reduces strategic curtailment: it increases ϕ, as shown in Lemma 1. This increases q_i whenever $w_i = H$, which also increases $U(Q)$ because $U' > 0$. Increasing d also increases wind diversification: it increases the probability that wind producers are in different states. This increases the probability that Q takes on its middle value, and decreases the probability that it takes on an extreme value. Since U is concave, the impact of diversification on U is weakly positive, i.e. $WD_U \geq 0$, as shown above. Figure 1 illustrates these effects. This intuition becomes clear with the following proof.

Proof. Since $W = U(q_1 + q_2)$, the expected value of social welfare is given by:

$$\mathbb{E}_{w_1,w_2}[W] = \Pr\{L, H\}U(L + \phi) + \Pr\{L, L\}U(2L) + \Pr\{H, H\}U(2\phi) + \Pr\{L, H\}U(L + \phi). \tag{2.9}$$

In addition, define $\frac{\partial \Pr\{L,H\}}{\partial d} \equiv \zeta$. By taking the derivatives of the probability values with respect to d, we have that $-\frac{\partial \Pr\{L,L\}}{\partial d} = -\frac{\partial \Pr\{H,H\}}{\partial d} = \frac{\partial \Pr\{L,H\}}{\partial d} \equiv \zeta > 0$ (the calculation is provided in Section 2.12, equation (2.38)). That is, when d increases the probably of being in different states increases. So,

$$\begin{aligned} \frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[W] = & \underbrace{\underbrace{\zeta}_{>0} \underbrace{(2U(L+\phi) - U(2L) - U(2\phi))}_{=WD_U \geq 0,\text{ by concavity of } U}}_{\textbf{wind diversification}} \\ & + 2 \underbrace{\frac{\partial \phi}{\partial d}}_{\substack{>0,\ \textbf{reduction of} \\ \textbf{strategic curtailment}}} \left(\underbrace{\Pr\{L, H\}P(L + \phi) + \Pr\{H, H\}P(2\phi)}_{>0} \right). \end{aligned} \tag{2.10}$$

Concavity of U implies that $WD_U = 2U(L+\phi)-U(2L)-U(2\phi) \geq 0$. Thus, wind diversification has a (weakly) positive impact on welfare. In addition, by increasing d, the production in the high state increases; $\frac{\partial \phi}{\partial d} > 0$ by Lemma 1. Therefore, the reduction of strategic curtailment (due to increasing d) has a positive impact on welfare. Overall, increasing dispersion increases the expected value of social welfare. Figure 2-3 shows these effects.

□

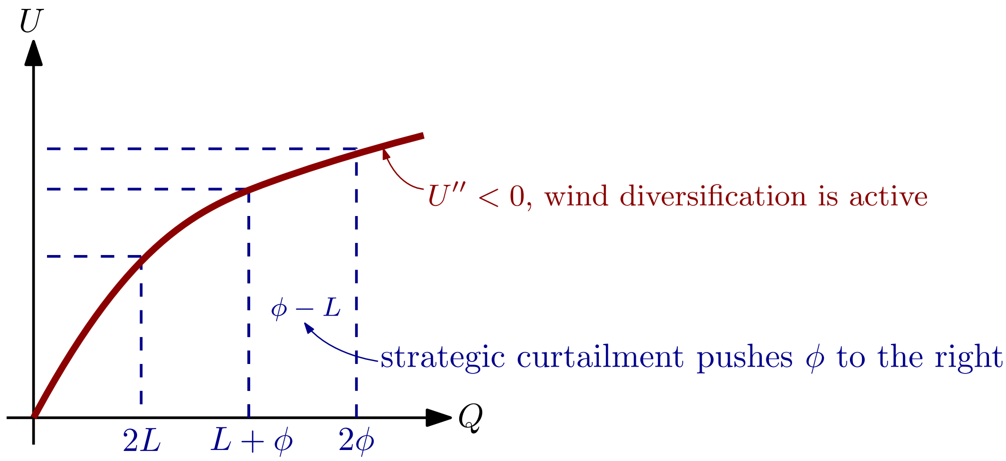

Figure 2-3: Wind diversification increases U and information effects decrease strategic curtailment which also increases U.

2.6 Price and Profit vs. Dispersion

How does extent of heterogeneity/dispersion affect average price and profit? We show the effect in general is *ambiguous*, i.e. profit is not necessarily a weakly increasing function of d nor a weakly decreasing function of d, regardless of the system conditions and exogenous parameters. This occurs because the impacts of diversification and of changing levels of strategic curtailment on average price and profit are *not* aligned. To understand this, we analyze how average price responds to changes in dispersion. Figure 2.6 shows these effects.

Proposition 3. The general impact of dispersion d on the expected value of price is ambiguous. In the case of linear inverse demand, increasing d decreases the expected value of the price.

Proof. Let $P'' < 0$. Since

$$\mathbb{E}_{w_1,w_2}[P(q_1(w_1)+q_2(w_2))] = 2\Pr\{L,H\}P(L+\phi)+\Pr\{L,L\}P(2L)+\Pr\{H,H\}P(2\phi), \tag{2.11}$$

thus

$$\begin{aligned} \frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[P] = &\underbrace{\underbrace{\zeta}_{>0}\underbrace{(2P(L+\phi)-P(2L)-P(2\phi))}_{WD_P>0,\text{ by strict concavity of }P}}_{\textbf{wind diversification}} \\ &+2\underbrace{\frac{\partial\phi}{\partial d}}_{>0,\ \textbf{reduction of strategic curtailment}}\left(\underbrace{\Pr\{L,H\}P'(L+\phi)+\Pr\{H,H\}P'(2\phi)}_{<0,\ P\text{ is downward}}\right). \end{aligned} \tag{2.12}$$

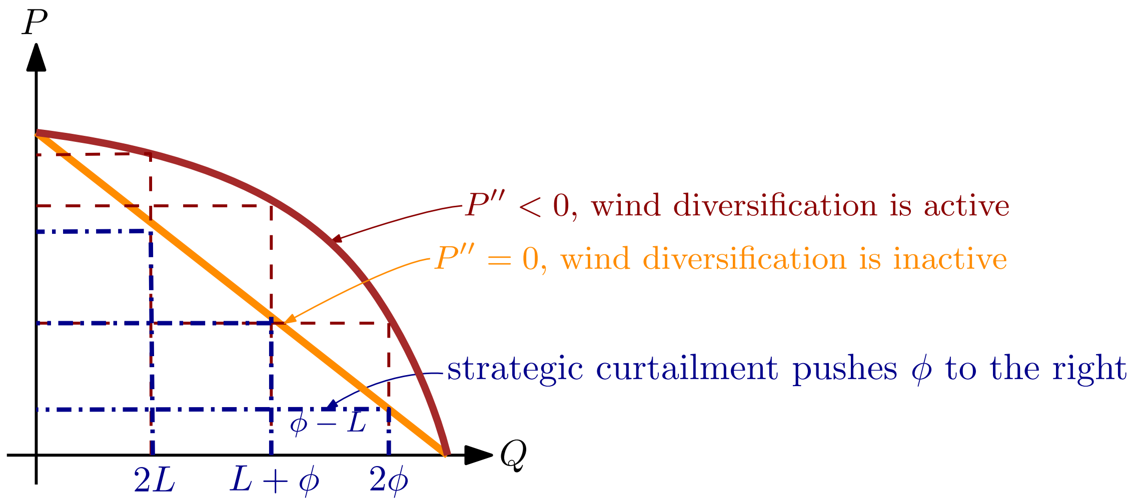

Figure 2-4: Interplay between the effects of wind diversification and strategic curtailment on average price. Wind diversification increases the average price when $P'' < 0$ and is inactive when $P'' = 0$. The impacts of increasing d on strategic curtailment always decrease the average price.

Higher dispersion reduces strategic curtailment, which decreases the average price because inverse demand is downward, i.e. $P' < 0$. However, diversification increases the average price because of concavity in inverse demand, i.e. $WD_P = 2P(L+\phi) - P(2L) - P(2\phi) > 0$. The net effect is ambiguous.

Note that when inverse demand is linear, i.e. $P'' = 0$, then $WD_P = 2P(L+\phi) - P(2L) - P(2\phi) = 0$. Thus, the effect of diversification is completely inactive. As a result, because of the impacts of d on strategic curtailment, the expected value of the market price decreases in d in the case of a linear inverse demand.

□

Like the average price, the impact of increasing dispersion on profit is in general ambiguous. When d increases, it increases ϕ. This decreases profit under the outcome where $w_1 = w_2 = H$ because 2ϕ is greater than the monopoly output. However, increasing d increases the probability that the two producers have different resource availability, $\Pr\{w_1 \neq w_2\}$, which increases the expected value of profit because diversification has a positive effect on profit. We can again characterize the effect of d on profit completely through its effects on strategic curtailment and diversification. The overall impact of dispersion on profit is ambiguous, as shown in the following Example.

Example 1. Let $P' < 0$, $P'' < 0$. As d increases, the expected value of producer profit increases due to diversification and decreases due to reduced strategic curtailment. Thus, in general, the impact of heterogeneity on profit is ambiguous.

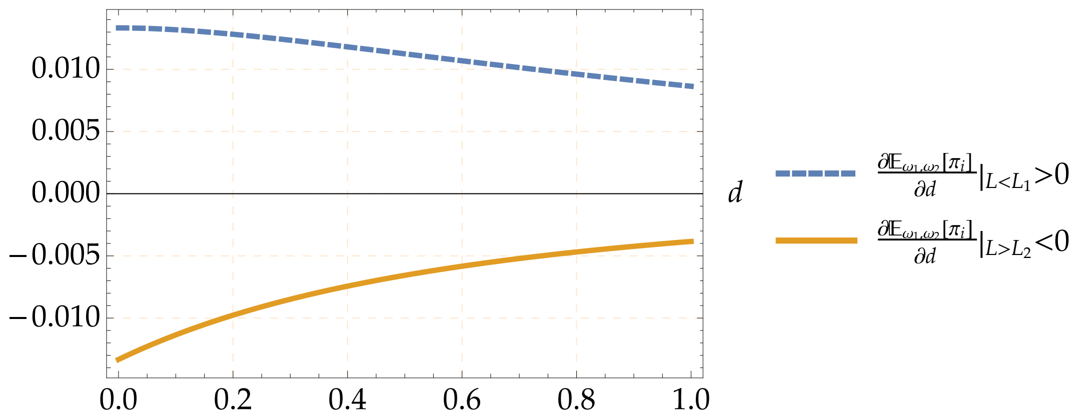

Figure 2-5: Wind diversification (heterogeneity) increases profit if L is sufficiently small, and it decreases profit if L is sufficiently large. The y-axis is the derivative of profit with respect to d, evaluated over the range of $d \in (0, 1]$ (x-axis). Plot parameters: $s = 3, \beta = \frac{1}{2}$, for the dashed line $L = 0.6$ and for the solid line $L = 0.8$.

In general, increasing dispersion d can increase or decrease the expected value of profit. However, in the case of linear inverse demand, we can obtain sharp insights based on the absolute value of the low state energy availability L. This is because the extent of L affects the *strength* of diversification and changing strategic curtailment levels on profit in *opposite* directions. As such, for sufficiently high L, increasing heterogeneity d reduces profits. See Figure 2-5. The following Proposition summarizes:

Proposition 4. Let $P(q_1 + q_2) = s - q_1 - q_2$, then there exists two thresholds L_1 and L_2, with $L_1 = \frac{2s}{9} < L_2 = \frac{2s}{8}$, such that

(i) When $L < L_1$ the impact of diversification dominates the strategic curtailment effects, thus $\frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[\pi_i] > 0$. Consequently, it is beneficial for firms to place their wind farms *far* from each other, i.e.

$$\arg\max_{d\in[0,1]} \mathbb{E}_{w_1,w_2}[\pi_i] = 1. \tag{2.13}$$

(ii) When $L > L_2$, strategic curtailment dominates diversification, thus $\frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[\pi_i] < 0$. Consequently, it is beneficial for firms to place their plants *close* to each other, i.e.

$$\arg\max_{d\in[0,1]} \mathbb{E}_{w_1,w_2}[\pi_i] = 0. \tag{2.14}$$

When L is sufficiently high, Proposition 4 shows that the expected value of profit for each producer is decreasing in d. In this case, investors prefer lower d even though higher d improves overall welfare, as shown in Proposition 2. This suggests that profit and welfare motives may sometimes be misaligned, since dispersion uniformly improves social welfare but may not improve profit. For example, a regulator may propose policies to increase d by encouraging investment far from existing sites, but firms might find more value in investing close to existing sites.

The next two Sections study how dispersion impacts information sharing and collusion in the duopoly model of wind competition. Section 2.9 extends the original model to the case of multiple wind generators. Section 2.10 extends the original model to an economy where two wind generators compete with traditional generation.

2.7 Public Forecasting: Who Benefits?

This section focuses on the benefits of public sharing of information under the assumption that wind producers do not collude. It investigates the benefits of publicly providing high-quality short-term weather forecasts or real-time wind speeds for all wind-generating locations. It suggests that public forecasting always improves welfare, but it does not always benefit producers. This suggests that producers will not provide public forecasting, but that policy makers should consider funding forecasting efforts to improve the quality of public information.

The results in this section show that information sharing always improves social welfare. However, we also show that when L is sufficiently large (as a function of dispersion d), wind producers do not choose to share information. The limit on L is increasing as a function of dispersion d. The results suggest that policies to implement public weather forecasting can improve welfare, because profit-maximizing producers will not always share weather information even though doing so always improves social welfare.

In order to understand the effects of information sharing on social welfare and producer profit, we compare the baseline model (see Section 2.2), where wind availability is private information, to the case where wind energy producers **ex-ante** commit[17] to share their private information about their available energy, given that the extent of wind producer heterogeneity is $d \in (0, 1]$. We assume inverse demand is linear, i.e. $P(q_1, q_2) = 1 - q_1 - q_2$.

[17] We assume wind producers are committed and there is no room for adverse selection. For instance, there could be automatic equipment for weather monitoring that shares information publicly.

Under this assumption, the net welfare obtained by consuming $Q = q_1 + q_2$ units of energy is $U(Q) = \int_0^Q P(q)dq = \int_0^Q (1-q)dq = Q - \frac{1}{2}Q^2$.

Is sharing information between wind producers *socially* beneficial? Information sharing has both positive and negative effects on welfare. It helps prevent producers in the high state from inefficiently withholding their output when the other producer is in the low state, but it also introduces additional costs to welfare due to the reduction in welfare when producers producer at the Cournot output when they are both in the high state. In general, however, these impacts are in favor of the benefits of sharing information, as the following proposition summarizes.

Proposition 5. Sharing information between wind producers is *socially* ex-ante beneficial.

Throughout this section, we let K denote the equilibrium outcomes when wind producer share private information (or that information is made public), and we let K^c denote equilibrium outcomes when producers compete without sharing information, as in Section 2.3.

To understand this result, consider the following. Let $W = \pi_1 + \pi_2 + CS = U(Q)$ denote welfare. Then, consider the expected value of the welfare benefits of information sharing, as follows

$$\begin{aligned}\mathbb{E}_{w_1,w_2}[W(K,K^c)] = \Pr\{L,H\}W_{L,H}(K,K^c) + \Pr\{H,L\}W_{H,L}(K,K^c) \\ + \Pr\{L,L\}W_{L,L}(K,K^c) + \Pr\{H,H\}W_{H,H}(K,K^c)\end{aligned} \tag{2.15}$$

where the benefit to social welfare of sharing information between wind producers at state $\{w_1, w_2\} \in \{H, L\}^2$ is

$$W_{w_1,w_2}(K,K^c) \equiv W^K_{w_1,w_2} - W^{K^c}_{w_1,w_2} = Q^K_{w_1,w_2} - \frac{1}{2}\left(Q^K_{w_1,w_2}\right)^2 - \left(Q^{K^c}_{w_1,w_2} - \frac{1}{2}\left(Q^{K^c}_{w_1,w_2}\right)^2\right), \tag{2.16}$$

where $Q^K_{w_1,w_2} - \frac{1}{2}\left(Q^K_{w_1,w_2}\right)^2$ is the social welfare at state $\{w_1, w_2\}$ when wind producers share their private information. Similarly, $Q^{K^c}_{w_1,w_2} - \frac{1}{2}\left(Q^{K^c}_{w_1,w_2}\right)^2$ denotes the social welfare when wind producers compete without sharing information.

Equation (2.16) highlights the fact that information sharing has mixed effects on social welfare in different states. In particular, it increases total output quantity (and welfare) when only one producer is in the high state, but it decreases output quantity and social welfare when both producers are in the high state. However, since total production is relatively lower (and therefore $U'(Q)$) is relatively higher, when the producers are in opposite states, the net expected value of information sharing is in favor of the benefits gained when producers are in opposite states.

As d increases, the benefits in the $\{H, L\}$ and $\{L, H\}$ states weakens and the costs incurred in the $\{H, H\}$ state increase, but the probability of being in the same state also decreases, so the proportional impact of the costs in state $\{H, H\}$ declines. Overall, information sharing improves social welfare for any β, d when Assumption 1 is satisfied.

Next, we consider the benefits of information sharing for producers' profits and show that in general they are not always aligned with the benefits for social welfare.

Is sharing information beneficial for *wind producers*? While sharing information always improves social welfare, it is not always beneficial for wind producers. We show the answer depends on the extent of wind energy at the low state, which directly affects the variance in the aggregate output. When wind in the low state is sufficiently small, sharing information is extremely beneficial for a generator that is in its high state. As a result, ex-ante wind producers prefer to share information when L is sufficiently small.

Proposition 6. There exists a threshold $L^*(d, \beta)$ that is increasing in d and decreasing in the prior β so that sharing information is ex-ante beneficial for wind energy producers if only if $L < L^*(d, \beta)$.

Let $D(K, K^c)$ represent the change in profits due to information sharing. The result aims to characterize the sign of (2.17), which represents the expected value of the benefits of information sharing for producer profits.

$$\begin{aligned}\mathbb{E}_{w_1,w_2}[D(K, K^c)] = &\Pr\{L, H\}D_{L,H}(K, K^c) + \Pr\{H, L\}D_{H,L}(K, K^c) \\ &+ \Pr\{L, L\}D_{L,L}(K, K^c) + \Pr\{H, H\}D_{H,H}(K, K^c)\end{aligned} \tag{2.17}$$

The benefit of sharing information at state $\{w_1, w_2\} \in \{H, L\}^2$ is

$$D_{w_1,w_2}(K, K^c) = \pi^K_{1_{w_1,w_2}} + \pi^K_{2_{w_1,w_2}} - \pi^{K^c}_{1_{w_1,w_2}} - \pi^{K^c}_{2_{w_1,w_2}}, \tag{2.18}$$

where $\pi^K_{i_{w_1,w_2}}$ denotes i's profit, $i \in \{1, 2\}$, at state $\{w_1, w_2\}$ when firms share their private information and $\pi^{K^c}_{i_{w_1,w_2}}$ denotes i's profit when firms compete with no information sharing.

To understand the effects, first consider the benefits of information sharing in the $\{H, H\}$ and $\{L, L\}$ states. In the $\{H, H\}$ state sharing information is always beneficial because of improved coordination. In the $\{L, L\}$ state the benefit of sharing information is

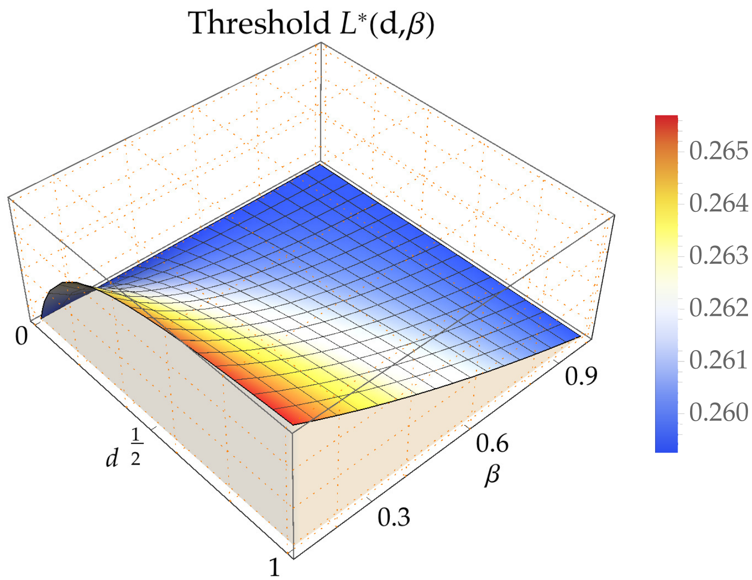

Figure 2-6: The information-sharing threshold $L^*(d, \beta)$ is increasing in the dispersion d and decreasing in the prior β.

always zero because both firms produce at the low level regardless of information sharing.

Now, suppose wind producer (WP) 1 is in the low state and WP 2 is in the high state. With information sharing, WP 2 achieves a best response to $w_1 = L$ and produces more energy compared to the case in which they do not share information. This benefits WP 2, because they achieve a best response based on improved information, but it hurts WP 1 because the price is reduced since WP 2 increases its output quantity. These effects favor information sharing when L is relatively lower. Low L improves the value of information sharing to WP 2 (because its overall adjustment is larger). Furthermore, low L decreases the cost of information sharing to WP 1, because the price effect impacts a lower quantity of production since L is small.

Overall, considering all the cases together implies that the expected benefit of sharing information is controlled by a threshold on the amount of wind energy in the low state. Therefore, sharing information is ex-ante beneficial for producers when wind energy at the low state is sufficiently small; there is a threshold $L^*(d, \beta)$, where sharing information is ex-ante beneficial for wind energy producers if only if $L < L^*(d, \beta)$. This suggests that when the variance of wind availability for individual generators is high, wind producers tend to benefit individually from information sharing; when the variance of their energy availability is low, information sharing reduces profits even though it improves social welfare. Furthermore, by increasing heterogeneity (i.e. the dispersion between the wind producers) this threshold increases, which incentivizes more wind energy producers to share their information. Figure 2-6 shows the impact of d and β on the threshold $L^*(d, \beta)$.

2.8 Collusion, Prevention, and Dispersion

This section investigates potential collusion between wind producers and studies the effect of increased heterogeneity on the presence of collusion. It focuses on linear inverse demand for simplicity, and shows that collusion is always possible (incentive compatible) among wind producers when there are no penalties for collusion. This is a straightforward result, given our modeling assumptions.[18] The section also examines the case where firms may face random penalties for engaging in collusion, so the threat of sanctions poses a random cost on the decision to collude. The level of dispersion d impacts the size of the penalty required to prevent collusion, but in a non-monotonic way.

Consider two wind producers that are willing to collude in order to increase profits. They set up a contract to produce at the monopoly level when possible and share profits depending on their stated availability. The true availability of wind is private information, so a wind producer in the high state can lie about their state and produce the amount of wind appropriate for a producer in the low state.[19]

Let π_M be the combined monopoly profits and π_L be the profits when both producers are in the low state. Since $H > q_M$, the producers can jointly achieve monopoly profits whenever at least one producer is in the high state. In the case of linear inverse demand, where $P(q_1, q_2) = s - q_1 - q_2$, the optimal output for a monopoly producer is $q_M = \frac{s}{2}$.

$$\pi_M = P(q_M)q_M = \frac{s^2}{4} \qquad \pi_L = P(L, L)L = (s - 2L)L \tag{2.19}$$

There is an exogenous cost to collusion $\gamma \geq 0$, to explain a situation where the government tries to identify and penalize collusion. We can think of γ as being the government's penalty for a firm engaged in collusion, times the probability of detection. The government might undertake various efforts to identify collusion, for instance by reviewing price trends, measuring the difference between wind forecasts and outputs, or monitoring information

[18] The quantity constraints imposed by the stochastic wind availability prevent the wind producers from deviating from the collusive output by increasing their production.

[19] We assume the the contract is enforceable with regards to production quantities, which are publicly verifiable. Therefore, if the producer announces that they are in the H (or L) state, then in any equilibrium they will produce the agreed upon amount for a producer in that state, regardless of their true state. However, it is not possible for a firm to verify the true state of their competitor (which is private information); out of equilibrium, a producer could choose to lie about its production constraint.

exchange between competing firms.

Colluding firms jointly produce at the monopoly level when at least one of them is in the high state. If the producers are both in the high state, they each receive $\frac{\pi_M}{2}$. If the producers are both in the low state, they each produce L and receive π_L. Additionally, the firms set up a transfer scheme where firms that are in the high state pay $t\pi_M$ to firms that are in the low state.[20]

Collusion is possible whenever there exists a monetary transfer $t\pi_M$ satisfying the incentive compatibility (IC) constraint, which implies that high state producers will not lie and pretend they are in the low state, and which satisfies the individual rationality (IR) constraints, which implies that firms will know *ex-ante* that they would like to participate regardless of their unknown state. The incentive compatibility constraint is

$$\Pr\{H|H\}\frac{\pi_M}{2} + \Pr\{L|H\}(\pi_M - t\pi_M) \geq \Pr\{H|H\}t\pi_M + \Pr\{L|H\}\pi_L. \tag{2.20}$$

The IC constraint says that the expected value of the profit for a colluding producer i in the high state is greater than the expected value of the profit they would receive if they lied and declared that they were in the low state. The individual rationality constraints for high and low state producers are, respectively,

$$\Pr\{H|H\}\frac{\pi_M}{2}+\Pr\{L|H\}(\pi_M-t\pi_M)-\gamma \geq \Pr\{H|H\}\phi(s-2\phi)+\Pr\{L|H\}\phi(s-L-\phi) \tag{2.21}$$

$$\Pr\{H|L\}t\pi_M + \Pr\{L|L\}\pi_L - \gamma \geq \Pr\{H|L\}L(s-\phi-L) + \Pr\{L|L\}\pi_L. \tag{2.22}$$

As before, the conditional probability $\Pr\{H|L\}$ refers to $\Pr\{w_j = H|w_i = L\}$ (this is the same for other combinations of the states H and L). Equation (2.21) explains that a producer in the high state would prefer to collude than to participate in the strategic equilibrium from Section 2.3. Equation (2.22) explains that a producer in the low state would prefer to collude than to participate in the strategic equilibrium from Section 2.3. Both of these constraints must hold; otherwise, a firm would not participate ex-ante because they would recognize that they would terminate the collusion agreement if they were in one state, revealing their availability to their competitor and reducing their profit.

Proposition 7. If there is no cost to collusion, i.e. $\gamma = 0$, then there is always an available transfer satisfying the IC and IR constraints. That is, when $\gamma = 0$, $\exists t \in \mathbb{R}$ that satisfies

[20]The transfer fraction t represents the fraction of monopoly profits given to the low state firm; since arbitrary $t \in \mathbb{R}$, and $\pi_M > 0$, any real number is a feasible transfer; the total transfer is written as a product of t and π_M (as opposed to a single parameter) because it simplifies the exposition.

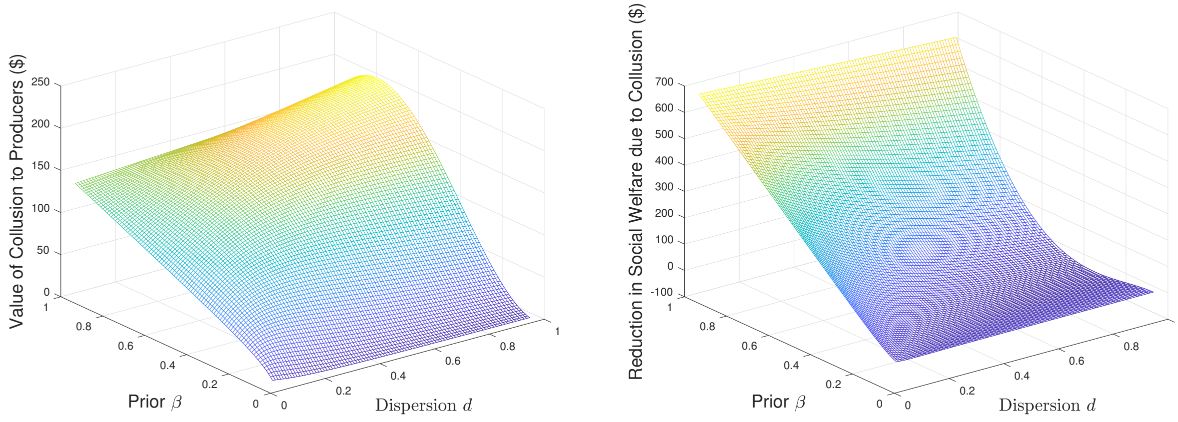

(a) (b)

Figure 2-7: The impact of dispersion on various features of collusion. (a) shows the impact of dispersion on the value of collusion to producers. (b) shows the impact of dispersion on the costs of collusion in terms of a reduction of social welfare.

(2.20), (2.21), and (2.22). Therefore, when $\gamma = 0$, producers can always increase profits by colluding.

The intuition is that a transfer is always possible when $\gamma = 0$ because the sum of profits from the generators strictly improves when they collude and when one producer is in the high state, so the benefit to producers in the high state is larger than the cost to producers in the low state, and thus there is a feasible transfer that allows collusion to be beneficial for producers *ex-ante*. Next, we examine the effect of d on various features of collusion.

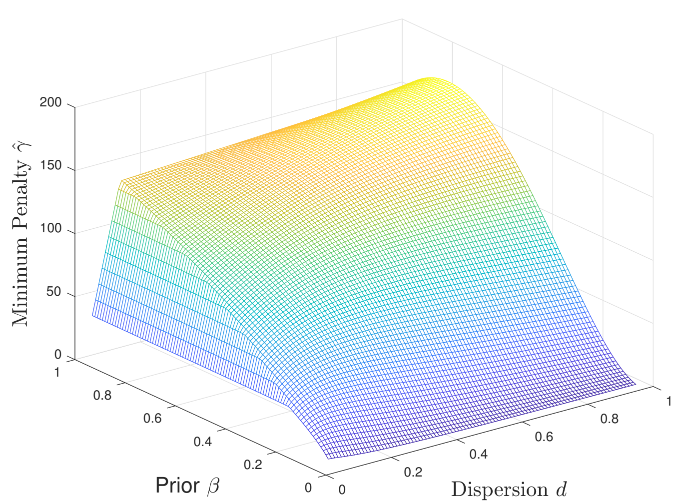

Figure 2-8: The impact of dispersion on the level of enforcement required to prevent collusion.

How does dispersion d impact collusion?

In general, we find that dispersion d does not have generalizable impacts on collusion in our model. Dispersion does not have monotonic impacts on the value of collusion to producers. It also does not monotonically impact the change in welfare due to collusion by producers. Figure 2-7 summarizes these two effects.

We can also estimate the impact of d on policies intended to prevent collusion. Let $\hat{\gamma}$ represent the minimum γ such that (2.20), (2.21), and (2.22) imply a contradiction. Variable $\hat{\gamma}$ represents the minimum expected value of a collusion penalty such that enforcement is sufficient to prevent collusion; if $\hat{\gamma}$ is very high, this implies that collusion must have a high probability of being punished and/or that the punishment must be very severe in order to prevent collusion. We find that dispersion does not have monotonic impacts on $\hat{\gamma}$. Figure 2-8 displays this effect.

2.9 Multiple Wind Generators with a Generic Joint Distribution of Wind Availability

This section shows that the main results of the chapter extend to markets with multiple wind generators. We demonstrate a parsimonious way to extend the notion of dispersion d to markets with an arbitrary number of wind producers, and we show that high state output and welfare are still increasing in d due to its effects on strategic curtailment and diversification.

Consider a market with $N+1$ wind generators, each with prior probability $\mathbb{P}(w_i = H) = \beta$, separated by dispersion d. Here, d gives a proxy for the level of correlation among the states of different producers, where as before high d implies that the stochastic resource availabilities of different producers are more independent. We define the state of producer i as $s_i = \mathbf{1}_{\{w_i=H\}}$. Let $S_{-i} = \sum_{j\neq i} s_j$, the number of producers besides producer i that are in the high state. Let $S = \sum_i s_i$, the total number of producers in the high state.

Consider the random vector $\mathbf{s}^d$ for $d \in (0,1]$ and assume $\beta > 0$. The probability distribution of $\mathbf{s}^d$ is the joint probability distribution $\Pr\{s_1, s_2, ..., s_{N+1}; d\}$. Each of s_i are random variables, as are the sums, and therefore both

$$\Pr\{S = j; d\} \quad j \in \{0, ..., N+1\} \tag{2.23}$$

$$\Pr\{S_{-i} = k | w_i = H; d\} \quad k \in \{0, ..., N\} \quad i \in \{1, ..., N+1\} \tag{2.24}$$

are well defined. Moving forward, we use S^d and S^d_{-i} as the random variables of the sum of states generated by distributions parameterized by dispersion d. We assume that distributions are symmetric; $\forall i, j$, the probability law for S^d_{-i} is equal to the probability law for S^d_{-j}.

As before, we assume that L is sufficiently small such that producers never curtail in the low state, i.e. $P((N+1)L) + LP'((N+1)L) > 0$. This is equivalent to the first part of Assumption 1 in the duopoly case. The first order optimality condition for ϕ is given by

$$\mathbb{E}_{S_{-i}}[P(\phi + S_{-i}\phi + (N - S_{-i})L) + \phi P'(\phi + S_{-i}\phi + (N - S_{-i})L)|w_i = H] = 0, \tag{2.25}$$

where the expectation is evaluated using the conditional probability distribution in (2.24). We assume there exists some $v < H$ that solves (2.25) when $\phi = v$. This corresponds to the second part of Assumption 1 for the oligopoly case, but it is a weaker requirement. It simply ensures that the equilibrium is of interest; otherwise, $q_i(w_i) = w_i$ and players always produce their full energy availability. Under these assumptions, there is a unique root ϕ that solves (2.25), with $L < \phi < H$, and the unique symmetric BNE is given by $q_i(w_i) = \min\{w_i, \phi\}$. We adopt these assumptions for the remainder of this section, and let ϕ refer to the unique root of (2.25).

Next we characterize two sufficient conditions on the effect of the parameter d on the joint and conditional distributions.[21] These conditions allow for the extension of our results on strategic curtailment and welfare to any arbitrary inverse demand curve with $P' < 0$, $P'' \leq 0$ in a market with $N + 1$ producers.

Assumption 2. For all $d' > d$, for each i, conditional on $w_i = H$, $S^d_{-i} \succeq_{FOSD} S^{d'}_{-i}$ That is, $\forall i$, $\forall j \in \{0, ..., N\}$

$$\Pr\{S_{-i} > j | w_i = H; d\} \geq \Pr\{S_{-i} > j | w_i = H; d'\}. \tag{2.26}$$

Assumption 3. For all $d' > d$, $S^{d'} \succeq_{SOSD} S^d$ That is, $\forall m$,

$$\sum_{j=0}^{m} \big(\Pr\{S > j; d'\} - \Pr\{S > j; d\} \big) \geq 0. \tag{2.27}$$

From the perspective of a single producer i in the high state, Assumption 2 requires that more competitors are likely to be in the high states when dispersion d is lower. The idea is that when dispersion d is small, producer i being in the high state provides a stronger

[21] The conditions are based on first- and second-order stochastic dominance, see Shaked and Shanthikumar (2007).

signal that competitors are also more likely to be in the high state.

Assumption 3 says that when d is higher, the sum of wind availability has at least as high a mean and less weight in the tails of its distribution. When d is high, the resource availabilities of different producers are nearly independent. When d is low, there is high correlation between producers. Both Assumptions 2 and 3 are satisfied by the duopoly model in Section 2.2.[22]

Proposition 8. For general $N \geq 1$, given Assumption 2, the output of producers in the high state ϕ is (weakly) increasing in d. Therefore, as in the duopoly case, increasing d (weakly) decreases strategic curtailment.

The left hand side of the first order condition (2.25) is in general decreasing in the output of other producers. The intuition is that the expected value of the output of other producers, with ϕ fixed, is decreasing in d, from the perspective of a producer whose output is high. Therefore, higher d increases the left hand side of (2.25). Lower ϕ also increases the left hand side. Thus, as d increases, a lower ϕ cannot possibly solve the first order condition because both higher d and lower ϕ increase the left hand side of (2.25).

Proposition 9. For general $N \geq 1$, given Assumptions 2, 3, and $P(\phi^1(N+1)) \geq 0$,[23] the expected value of welfare $\mathbb{E}_{\mathbf{s}^d}[W]$ is increasing in d.

Consider $d' > d$. We aim to show that $\mathbb{E}_{\mathbf{s}^{d'}}[W] > \mathbb{E}_{\mathbf{s}^d}[W]$. Let ϕ^d refer to the equilibrium curtailment level as described by (2.25), for the random availability vector $\mathbf{s}^d$ indexed by d. Consider a given realization of the resource availability for each producer, and let $S = \tilde{s}$, for some $\tilde{s} \in \mathbb{Z}$ with $0 \leq \tilde{s} \leq N+1$. We can describe welfare as a function $W(\tilde{s}, \phi)$. The full proof in Section 2.12 explains that under the first-order conditions of the equilibrium described by (2.25), W is concave and increasing in $\tilde{s}$. Welfare W is also increasing in ϕ.

The distributions of S satisfy Assumption 3, so the distribution with higher d has total wind availability S that second-order stochastically dominates the original distribution. The definition of second-order stochastic dominance implies that the dominating random variable leads to higher expected value for increasing concave functions. Therefore, holding

[22] Consider the duopoly model in Section 2.2, but with general probability distributions $\Pr\{w_i = H | w_j = H\} \equiv f(d, \beta)$ and $\Pr\{w_i = H | w_j = L\} \equiv g(d, \beta)$. In Section 2.2, specific functional forms are provided in (2.1) for $f(d, \beta)$ and $g(d, \beta)$ in order to motivate the exposition. For generic conditional probabilities in the duopoly model, Assumption 2 establishes that $f(d, \beta)$ is weakly decreasing in d. Assumptions 2 and 3 together establish that $g(d, \beta)$ is weakly increasing in d.

[23] The variable ϕ^1 represents the high state production when $d = 1$. This assumption implies that equilibrium prices will not become negative. In practice, we see negative prices arise in markets with high penetrations of wind resources, but this is due to the presence of subsidies, and non-convexities associated with traditional generation, not a result of producer strategy in the face of uncertainty.

ϕ constant, higher d increases the expected value of welfare: $\mathbb{E}_{\mathbf{s}^{d'}}[W(\cdot, \phi^{d'})] > \mathbb{E}_{\mathbf{s}^d}[W(\cdot, \phi^{d'})]$. Furthermore, using Assumption 2, Proposition 8 shows that ϕ is increasing in d. Since W is increasing in ϕ, $\mathbb{E}_{\mathbf{s}^d}[W(\cdot, \phi^{d'})] > \mathbb{E}_{\mathbf{s}^d}[W(\cdot, \phi^d)]$; together, the two inequalities establish that W is increasing in d.

2.10 Competition with Traditional Generation

This section considers Cournot competition between two wind producers and a traditional generator. The traditional generator models fossil fuel generators, and also some renewable energy facilities, like biomass generators. These traditional generators are controllable; unlike wind and solar facilities, their production is not constrained by stochastic resource availability.

In this model, the wind producers with dispersion d and availability β have identical characteristics to those described in Section 2.2. The traditional producer can output any quantity $x \in \mathbb{R}^+$ with constant marginal cost $c \geq 0$; it has no information regarding the availability of the wind generators.

This section extends the existing results on the impact of d on welfare. As before, welfare is increasing in d. The models used in this section and Section 2.9 could be used to analyze markets with multiple wind producers and multiple traditional generators, but the analysis in this section is sufficient to highlight the main insights. The behavior of the thermal generator is straightforward because the traditional producer's output decreases when the wind farm's output increases.

Let $\underline{x} \geq 0$ be the solution to $\mathbb{E}[P(w_1, w_2, \underline{x}) + \underline{x}P'(w_1, w_2, \underline{x})] = c$. This value represents the lower bound for the total energy production by the traditional generator in an equilibrium. The existence of the equilibrium for the market with two wind producers and a traditional producer requires the following assumption:

Assumption 4. Let $P(\cdot)$ be inverse demand and c be the marginal cost of traditional generation. Assume $c < P(2H)$, which guarantees that the traditional generator produces a positive quantity. Assume $P(3L) + LP'(3L) > 0$ and $P(H + L + \underline{x}) + HP'(H + L + \underline{x}) < 0$.

Assumption 4 extends Assumption 1 to the case of three players with a traditional generator. It guarantees that we have a solution of interest, so we avoid explaining the cases whereby L is sufficiently high that wind producers might always curtail, where H is too low so that wind producers might never curtail, or where c is sufficiently high that the traditional producer will never produce.[24]

Proposition 10. The Cournot equilibrium for generic inverse demand $P(\cdot)$, with $P' < 0$, $P'' < 0$ satisfies the following first order conditions, where (2.28) is the first order condition for wind producers and (2.29) is the first order condition for the traditional producer.

$$\Pr\{L|H\}(P(L+\phi+x)+\phi P'(L+\phi+x))+\Pr\{H|H\}(P(2\phi+x)+\phi P'(2\phi+x))=0 \quad (2.28)$$

$$\Pr\{L,L\}(P(2L+x)+xP'(2L+x))+2\Pr\{L,H\}(P(L+\phi+x)+xP'(L+\phi+x)) \\ +\Pr\{H,H\}(P(2\phi+x)+xP'(2\phi+x))-c'(x)=0 \quad (2.29)$$

The result follows exactly from Proposition 1 with the addition of the traditional generator whose output satisfies the first order condition described in (2.29). Equation (2.28) describes the equilibrium high state output ϕ for wind producers to maximize their profit, contingent on the equilibrium behavior of the other wind producer and the traditional generator. Equation (2.29) describes the equilibrium output x of the traditional generator, with $c'(x) = c$ in our example.

Example 2. Consider a market with linear inverse demand, $P(q_1, q_2, q_3) = s - q_1 - q_2 - q_3$. Then the unique high state output of the wind generators ϕ and the production output of the traditional generator x are given by:

$$\phi = \frac{\frac{1}{2}(s+c)(\beta+d(1-\beta))+L\beta(1-\beta)(1-d)}{3\beta+2d(1-\beta)-\beta^2-\beta d(1-\beta)} \quad (2.30)$$

$$x = \frac{1}{2}(s-c)-\phi\beta-L(1-\beta). \quad (2.31)$$

The Example is explained in Section 2.12. It is obtained by solving the first-order conditions (2.28) and (2.29) in terms of ϕ and x in the case of linear inverse demand.

Next we consider the impact of heterogeneity on strategic curtailment by wind producers $\frac{\partial\phi}{\partial d}$ and quantity withholding by the traditional producer $\frac{\partial x}{\partial d}$ in the case of linear inverse demand. We can take the derivative of ϕ with respect to d, using the form of the equation in (2.30).

$$\frac{\partial\phi}{\partial d} = \frac{(s+c-4L)\beta(1-\beta)}{2(3\beta+2d(1-\beta)-\beta^2-\beta d(1-\beta))^2} \quad (2.32)$$

[24]The assumption establishes an upper limit on c. When c is lower, the output of the traditional generator increases because they have lower marginal costs of production.

Under our assumptions, this is always positive. Equation (2.32) is always positive when $s+c-4L>0$, which is always satisfied by Assumption 4. Therefore, the output of the wind generators is increasing in d, $\frac{\partial\phi}{\partial d}>0$, as in the original market.

Then, taking the derivative of x using the first order condition in (2.31), $\frac{\partial x}{\partial d}=-\beta\frac{\partial\phi}{\partial d}<0$. Therefore, the output of the traditional generator is decreasing in d, so the traditional generator withholds more when the wind generators have less information about the other wind producer's state.

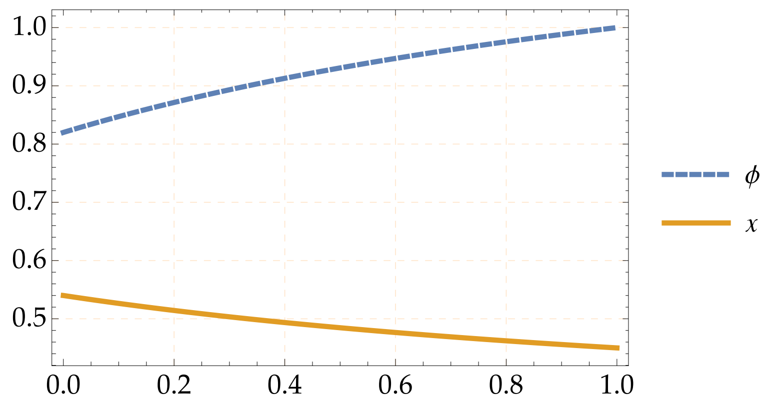

Figure 2-9: This chart shows that the traditional firm's output x decreases in the diversification d, but the average output (and the high state production ϕ) of the wind generators is increasing in d. Plot parameters: $s=3, \beta=\frac{1}{2}, L=0.1$ and $c=1$.

Next, we consider the effects of heterogeneity on welfare. Increasing dispersion d improves welfare in the market that includes a traditional generator.

Proposition 11. In the three player market with two wind producers and a traditional producer, and a linear inverse demand $P(q_1,q_2,q_3)=s-q_1-q_2-q_3$, the expected value of welfare is increasing in dispersion d.

In this model, as before, increasing d still reduces the strategic curtailment of wind producers $\frac{\partial\phi}{\partial d}>0$, and improves wind diversification. However, when a fossil fuel generator has market power, the fossil fuel generator responds by withholding more of their own output due to strategic substitutability with $\mathbb{E}[q_1+q_2]$, which increases; thus $\frac{\partial x}{\partial d}<0$. With a linear inverse demand, the FOCs (2.28) and (2.29) imply that the sum of the welfare impacts, due to the changes in the equilibrium values of ϕ and x, is 0. Thus, increasing d only impacts the expected value of welfare through the change in wind diversification, which positively impacts welfare.

Finally, we show that in a market with traditional generation and linear inverse demand, d decreases the expected value of price. This result extends earlier results to the case of a market with traditional generators and highlights the potential benefits of increased heterogeneity for reducing market prices.

Proposition 12. The expected value of the market price, given by $\mathbb{E}_{w_1,w_2}[P(q_1(w_1)+q_2(w_2)+x)]$, satisfies $\frac{\partial \mathbb{E}_{w_1,w_2}[P]}{\partial d} = -\beta\frac{\partial \phi}{\partial d} < 0$.

The expected value of total energy production is increasing in d. Its derivative with respect to d is given by $2\Pr\{H\}\frac{\partial \phi}{\partial d} + \frac{\partial x}{\partial d} = (2\beta - \beta)\frac{\partial \phi}{\partial d} > 0$, where the equality is because $\frac{\partial x}{\partial d} = \beta\frac{\partial \phi}{\partial d}$, as explained in the Example 2 above. Under linear inverse demand, the expected value of the market price is decreasing in d.

Since the production by the traditional generator is uniformly decreasing in d, increased dispersion reduces profits for the traditional generator. On the other hand, the effects of d on wind producer profits are ambiguous, as was the case in the original model.

2.11 Conclusion

This research links the heterogeneity in wind producer availability to social welfare in electricity markets with strategic behavior. It introduces the idea that the level of correlation in wind farm energy availability impacts strategic behavior. It shows that increasing heterogeneity decreases the strategic incentive of individual wind producers and increases welfare. This impact could become increasingly important as renewable energy penetration grows, especially because of the difficulties associated with monitoring market power when resource availability is not deterministic.

The results show that increasing heterogeneity is generally beneficial because of its positive impacts on increasing diversification and on decreasing strategic curtailment. Some existing policies and subsidy models for wind energy, like state-level renewable portfolio standards, have been shown to impact the optimal investment locations for new projects; these effects should be further reviewed in light of these results. The research also highlights the benefits of publicly sharing high-quality real-time weather information, even when it is not in the best interest of producers. As such, policy makers should consider the potential benefits of improved public forecasting and publicly sharing real-time energy output data, understanding that welfare-improving policies may be opposed by electricity generators.

2.12 Proofs Omitted from the Main Text

Proof of Proposition 1. Since $P' < 0, P'' \leq 0$, firm i's profit $\pi_i(q_i, q_j) = q_i P(q_i, q_j)$ is concave in q_i regardless of the production q_j by its competitor. Let firm i be in the high state, i.e. $w_i = H$. By Assumption 1, $P(H) + HP'(H) < 0$. Furthermore, $P(x) + xP'(x)$ is decreasing in x. Therefore, the resource availability in the high state does not bind, i.e. $q_i(H) = \phi \leq H$. The optimal output $q_i(H) = \phi$ belongs to $\arg\max_{q_i \in \mathbb{R}} \mathbb{E}_{w_j}[\pi_i | w_i = H]$. Due to concavity of $\pi_i(q_i, q_j)$ in q_i, the first order optimality condition (the necessary and sufficient condition for optimality) implies that ϕ should satisfy the following

$$\Pr\{L|H\}\big[P(L+\phi) + \phi P'(\phi + L)\big] + \Pr\{H|H\}\big[P(2\phi) + \phi P'(2\phi)\big] = 0, \tag{2.33}$$

given firm j strategy is $q_j(w_j) = \min\{w_j, \phi\}$. Next, with the following Claims we show ϕ indeed satisfies (2.33) and verify that $q(L) = L$. Subsequently, we prove the symmetric equilibrium is unique.

Claim 1 There exists a unique ϕ satisfying (2.33). Moreover, $L < \phi < H$.

Proof Let us define $f(x) \equiv \Pr\{L|H\}\big[P(L+x)+xP'(L+x)\big]+\Pr\{H|H\}\big[P(2x)+xP'(2x)\big]$. The derivative $f'(x) < 0$ because $P' < 0, P'' \leq 0$, $x \geq 0$. Moreover, $f(L) > 0$, from Assumption 1. Furthermore,

$$\begin{aligned} f(H) &= \Pr\{L|H\}\big[P(L+H) + HP'(H+L)\big] + \Pr\{H|H\}\big[P(2H) + HP'(2H)\big] \\ &< (P(H) + HP'(H))[\Pr\{H|H\} + \Pr\{L|H\}] < 0 \end{aligned}$$

where the first inequality follows since $P(x+y) + xP'(x+y)$ is decreasing in y, and the second inequality follows by Assumption 1. Since $f(L) > 0$, $f(H) < 0$, and $f'(x) < 0$ thus there exists a unique ϕ for which $f(\phi) = 0$, with $L < \phi < H$.

Claim 2 When $w_i = L$ then $q_i(L) = L$, given that firm j's strategy is $q_j(w_j) = \min\{w_j, \phi\}$.

Proof Let $g(x) \equiv \Pr\{H|L\}xP(\phi + x) + \Pr\{L|L\}xP(L+x)$. We aim to show that $x = L$ maximizes $g(x)$ when $x \in [0, L]$; this follows in a straightforward way from the necessary Karush–Kuhn–Tucker (KKT) condition for x:

$$\Pr\{H|L\}\big[P(x+\phi) + xP'(x+\phi)\big] + \Pr\{L|L\}\big[P(x+L) + xP'(x+L)\big] = \mu_L - \mu_0 \tag{2.34}$$

with $\mu_L \geq 0$, $\mu_0 \geq 0$, $\mu_L(x - L) = 0$, $\mu_0 x = 0$, and $x \in [0, L]$. The constants μ_L and μ_0 are the KKT multipliers associated with $x \leq L$ and $x \geq 0$. Since $P(x+y) + xP'(x+y)$ is decreasing in y, $f(\phi) = 0$ implies that $P(L+\phi) + \phi P'(L+\phi) > 0$, and since also $x \leq L < \phi$

and $P' < 0$, $P(x+\phi)+xP'(x+\phi) > 0$. Then again, since $L < \phi$, $P(x+L)+xP'(x+L) > 0$. Therefore, the left-hand side of (2.34) is strictly positive, which implies that $\mu_L > 0$ and therefore that $x = L$.

Claim 3 The equilibrium described in Proposition 1 is the unique symmetric equilibria.

Proof The poof follows by contradiction. Suppose, by contrary, firm j produces $q_j(L) = \tilde{L}$, where $\tilde{L} < L$. We show firm i has incentive to deviate by producing more than $\tilde{L}$ in the low state. Suppose $q_j(H) = q_i(H) = \tilde{\phi}$; thus, $\tilde{\phi}$ (according to first order optimality condition) solves the following:

$$\Pr\{L|H\}\big[P(\tilde{L}+\tilde{\phi}) + \tilde{\phi}P'(\tilde{\phi}+\tilde{L})\big] + \Pr\{H|H\}\big[P(2\tilde{\phi}) + \tilde{\phi}P'(2\tilde{\phi})\big] = 0. \tag{2.35}$$

By following the arguments from Claim 1, there is a unique $\tilde{\phi}$, where $\tilde{L} < \tilde{\phi} < H$, satisfying (2.35). Now, let $w_i = L$. Then, evaluating firm i's marginal expected profit when $w_i = L$ and $q_i(L) = \tilde{L}$, given firm j's strategy, implies

$$\begin{aligned} \frac{\partial}{\partial q_i}\mathbb{E}_{w_j}[\pi_i(q_i,q_j)|w_i = L]\big|_{q_i=\tilde{L}} = &\Pr\{H|L\}\big[P(\tilde{\phi}+\tilde{L}) + \tilde{L}P'(\tilde{\phi}+\tilde{L})\big] \\ &+ \Pr\{L|L\}\big[P(2\tilde{L}) + \tilde{L}P'(2\tilde{L})\big] \\ > &\, 0, \end{aligned} \tag{2.36}$$

where the last inequality is due to the following. By Assumption 1, $P(2L) + LP'(2L) > 0$. Also $\tilde{L} < L$ (by the above assumption) and $P(2x) + xP'(2x)$ is decreasing in $x \geq 0$. Thus $P(2\tilde{L}) + \tilde{L}P'(2\tilde{L}) > P(2L) + LP'(2L) > 0$. In addition, since $P' < 0$, and $\tilde{\phi} > \tilde{L}$, thus $P(\tilde{\phi}+\tilde{L}) + \tilde{L}P'(\tilde{\phi}+\tilde{L}) > P(\tilde{\phi}+\tilde{L}) + \tilde{\phi}P'(\tilde{\phi}+\tilde{L}) > 0$. The inequality (2.36) establishes a contradiction, because firm i has incentive to deviate, and produce more than $\tilde{L}$ when $w_i = L$. This completes the proof. □

Proof of Corollary 1. Let $q_i(w_i) = \min\{w_i, \phi\}$. Consider $i = 1$. The objective is to find ϕ. Thus, writing the first order optimality condition implies that ϕ satisfies the following equality

$$\begin{aligned} \phi &= \frac{s - \mathbb{E}_{w_2}[q_2|w_1 = H]}{2} \\ &= \frac{s - [L\Pr\{L|H\} + \phi\Pr\{H|H\}]}{2} \\ &= \frac{s - [L(1 - \frac{\beta}{\beta+d(1-\beta)}) + \phi(\frac{\beta}{\beta+d(1-\beta)})]}{2} \end{aligned}$$

where $\Pr\{L|H\} = \Pr\{w_2 = L|w_1 = H\} = \frac{(1-\beta)d}{\beta+d(1-\beta)}$ and $\Pr\{H|H\} = \Pr\{w_2 = H|w_1 = H\} = \frac{\beta}{\beta+d(1-\beta)}$. The above equality gives $\phi \equiv \frac{s\beta+(s-L)(1-\beta)d}{3\beta+2(1-\beta)d}$, completing the proof. □

Proof of Lemma 1. As shown in Proposition 1, production in the high state, i.e. ϕ, solves

$$\Pr\{L|H\}\big[P(L+\phi)+\phi P'(\phi+L)\big] + \Pr\{H|H\}\big[P(2\phi)+\phi P'(2\phi)\big] = 0. \tag{2.37}$$

Furthermore, according to (2.1), since $\Pr\{L|H\} = \frac{d(1-\beta)}{\beta+d(1-\beta)}$ and $\Pr\{H|H\} = \frac{\beta}{\beta+d(1-\beta)}$ thus

$$\frac{\partial}{\partial d}\Pr\{L|H\} = \frac{\beta(1-\beta)}{(\beta+d(1-\beta))^2} > 0$$
$$\frac{\partial}{\partial d}\Pr\{H|H\} = \frac{-\beta(1-\beta)}{(\beta+d(1-\beta))^2} < 0.$$

Now, taking a derivative from (2.37) with respect to d and taking into account that $\frac{\partial}{\partial d}\Pr\{H|H\} = -\frac{\partial}{\partial d}\Pr\{L|H\} < 0$ gives

$$\begin{aligned}
0 =& \frac{\partial \Pr\{L|H\}}{\partial d}[P(L+\phi)+\phi P'(\phi+L)] + \Pr\{L|H\}\left[2\frac{\partial\phi}{\partial d}P'(L+\phi)+\phi\frac{\partial\phi}{\partial d}P''(\phi+L)\right]\\
&+\Pr\{H|H\}\left[3\frac{\partial\phi}{\partial d}P'(2\phi)+2\phi\frac{\partial\phi}{\partial d}P''(2\phi)\right] + \frac{\partial \Pr\{H|H\}}{\partial d}[P(2\phi)+\phi P'(2\phi)]\\
=&\frac{\partial\phi}{\partial d}\left\{\Pr\{L|H\}\big[2P'(L+\phi)+\phi P''(\phi+L)\big]+\Pr\{H|H\}\big[3P'(2\phi)+2\phi P''(2\phi)\big]\right\}\\
&+\frac{\partial \Pr\{L|H\}}{\partial d}[P(L+\phi)+\phi P'(L+\phi)-P(2\phi)-\phi P'(2\phi)].
\end{aligned}$$

Therefore,

$$\begin{aligned}
\frac{\partial\phi}{\partial d} =& -\frac{\frac{\partial \Pr\{L|H\}}{\partial d}[P(L+\phi)+\phi P'(L+\phi)-P(2\phi)-\phi P'(2\phi)]}{\Pr\{L|H\}\big[2P'(L+\phi)+\phi P''(\phi+L)\big]+\Pr\{H|H\}\big[3P'(2\phi)+2\phi P''(2\phi)\big]}\\
&> 0,
\end{aligned}$$

where the inequality follows because: (i) $\frac{\partial \Pr\{L|H\}}{\partial d} > 0$, (ii) $P' < 0, P'' \le 0$, implying the denominator is negative, (iii) $P' < 0, P'' \le 0$ and $L < \phi$, implying that $P(L+\phi) > P(2\phi), P'(L+\phi) \ge P'(2\phi)$. □

Proof of Example 1. From (2.1), with prior probability $\Pr\{H\} = \beta$, we have

$$-\frac{\partial \Pr\{L,L\}}{\partial d} = -\frac{\partial \Pr\{H,H\}}{\partial d} = \frac{\partial \Pr\{L,H\}}{\partial d} \equiv \zeta = \frac{\beta^2(1-\beta)}{(\beta+d(1-\beta))^2} > 0. \tag{2.38}$$

The derivatives of the respective outcome probabilities are labeled as ζ and $-\zeta$ according to (2.38). By definition $\pi_i = q_i P(q_1 + q_2)$. Therefore $\mathbb{E}_{w_1,w_2}[\pi_i] = \Pr\{L, H\}[LP(L+\phi) + \phi P(L+\phi)] + \Pr\{H, H\}\phi P(2\phi) + \Pr\{L, L\}LP(2L)$. Taking the derivative of average profit with respect to d implies

$$\begin{aligned}
\frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[\pi_i] = & \underbrace{\zeta}_{>0} \underbrace{\big[LP(L+\phi) + \phi P(L+\phi) - LP(2L) - \phi P(2\phi)\big]}_{\equiv WD_\pi, \text{ wind diversification}} \\
& + \underbrace{\frac{\partial \phi}{\partial d}}_{>0} \left\{ \underbrace{\Pr\{L,H\}[LP'(\phi+L) + P'(\phi+L)\phi] + \Pr\{H,H\}[2P'(2\phi)\phi]}_{\equiv T_2, \text{ effects of } d \text{ on price through its impact on strategic curtailment}} \right\} \\
& + \underbrace{\frac{\partial \phi}{\partial d}}_{>0} \left\{ \underbrace{\Pr\{L,H\}[P(\phi+L)] + \Pr\{H,H\}[P(2\phi)]}_{\equiv T_3 \text{ the value of additional production due to reduced strategic curtailment}} \right\}.
\end{aligned}$$

WD_π represents the effects of wind diversification, which is positive because

$$\begin{aligned}
WD_\pi &= L[P(L+\phi) - P(2L)] + \phi[P(L+\phi) - P(2\phi)] \\
&> L\big[2P(L+\phi) - P(2L) - P(2\phi)\big] \\
&\geq 0
\end{aligned}$$

where the first inequality follows because $\phi > L$ and $P(L+\phi) - P(2\phi) > 0$ and the second inequality follows because of concavity in P, i.e. $P'' \leq 0$. Thus, profit increases due to increased diversification. Note that unlike the effect of diversification on average price, which is inactive when $P'' = 0$, diversification improves profit even when the inverse demand curve is linear.

Furthermore, the impacts of d on strategic curtailment has two effects on profit, because reducing strategic curtailment lowers the average price but also increases the aggregate quantity; these impacts are labeled as T_2 and T_3, respectively. Since inverse demand is downward, i.e. $P' < 0$, the impact of increasing d on markup through its effects on strategic curtailment is negative, i.e. $T_2 < 0$. The impact of reducing strategic curtailment on quantity is, expectedly, positive, i.e. $\frac{\partial \phi}{\partial d} > 0$ and $T_3 > 0$, because higher d results in lower information and less extensive strategic curtailment. However, the overall impact of d, through its impacts on strategic curtailment, is to reduce price. This is because its effect on average price is greater than its effect on average quantity; i.e. $T_2 + T_3 < 0$ because

$$
\begin{aligned}
T_2 + T_3 &= \Pr\{H,L\}\big[P(\phi+L) + (L+\phi)P'(\phi+L)\big] + \Pr\{H,H\}\big[P(2\phi) + 2\phi P'(2\phi)\big] \\
&= \Pr\{H,L\}\big[P(\phi+L) + \phi P'(\phi+L)\big] + \Pr\{H,H\}\big[P(2\phi) + \phi P'(2\phi)\big] \\
&\quad + \Pr\{H,L\}LP'(\phi+L) + \Pr\{H,H\}\phi P'(2\phi) \\
&= \Pr\{H,L\}LP'(\phi+L) + \Pr\{H,H\}\phi P'(2\phi) \qquad (2.39) \\
&< 0 \qquad (2.40)
\end{aligned}
$$

where (2.39) follows from the first order condition (2.33), and (2.40) follows because inverse demand is downward, i.e. $P' < 0$. Therefore, the effects of d on diversification increase profits, and the effects of d on strategic curtailment decrease profits. The overall impact of heterogeneity is ambiguous. Figure 2-5 provides examples showing that profit can be increasing or decreasing in d. □

Proof of Proposition 4. By definition $\pi_i(w_1, w_2) = q_i(w_i)(s - q_1(w_1) - q_2(w_2))$ where $q_i(w_i)$ is explicitly given by Corollary 1, for $w_i \in \{L, H\}$ and $i \in \{1, 2\}$. The expected value of profit for producer i is given by (2.41).

$$
E_{w_1,w_2}[\pi_i] = \Pr\{L,H\}[\pi_i(L,H) + \pi_i(H,L)] + \Pr\{H,H\}\pi_i(H,H) + \Pr\{L,L\}\pi_i(L,L) \qquad (2.41)
$$

As before, from (2.1), $\Pr\{L,L\} = (1 - \frac{d\beta}{\beta+d(1-\beta)})(1-\beta)$, $\Pr\{L,H\} = (1-\beta)\frac{d\beta}{\beta+d(1-\beta)}$, and $\Pr\{H,H\} = \beta\frac{\beta}{\beta+d(1-\beta)}$. In addition,

$$
\begin{aligned}
\pi_i(L,H) &= L(s - L - \phi) \qquad (2.42) \\
\pi_i(H,L) &= \phi(s - L - \phi) \qquad (2.43) \\
\pi_i(L,L) &= L(s - 2L) \qquad (2.44) \\
\pi_i(H,H) &= \phi(s - 2\phi) \qquad (2.45)
\end{aligned}
$$

where, as shown in Corollary 1, $\phi = \frac{s\beta+(s-L)(1-\beta)d}{3\beta+2(1-\beta)d}$. By plugging (2.42)-(2.45) into (2.41), the total (ex-ante) wind producers' surplus becomes

$$
\mathbb{E}_{w_1,w_2}[\pi_i] = \frac{\beta}{4} + L(1-2\beta) + L^2(\frac{15}{4}\beta - 2) - \frac{\beta^2(s-3L)(s-4L)}{2(3\beta+2d(1-\beta))} + \frac{\beta^3(s-3L)^2}{4(3\beta+2d(1-\beta))^2}. \qquad (2.46)
$$

Next, we characterize how d affects profits. The derivative of (2.46) with respect to d is

$$\begin{aligned}\frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[\pi_i] &= \frac{-\beta^3(s-3L)^2(1-\beta)}{(3\beta+2d(1-\beta))^3} + \frac{\beta^2(1-\beta)(s-3L)(s-4L)}{(3\beta+2d(1-\beta))^2} \\ &= \frac{\beta^2(1-\beta)(s-3L)}{(3\beta+2d(1-\beta))^2}\left[\frac{-\beta(s-3L)}{3\beta+2d(1-\beta)} + s - 4L\right] \\ &= \underbrace{\frac{\beta^2(1-\beta)(s-3L)}{(3\beta+2d(1-\beta))^3}}_{>0}\left[\beta(2s-9L)+d(1-\beta)(2s-8L)\right]. \qquad (2.47)\end{aligned}$$

From the last equality we obtain: If $L < \frac{2s}{9} \equiv L_1$, then $2s-9L>0$ and $2s-8L>0$; thus, $\frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[\pi_i] > 0$ and, consequently, $\arg\max_{d\in[0,1]}\ \mathbb{E}_{w_1,w_2}[\pi_i] = 1$. If $L > \frac{2s}{8} \equiv L_2$ then $2s-8L<0$ and $2s-9L<0$; thus $\frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[\pi_i] < 0$ and, consequently, $\arg\max_{d\in[0,1]}\ \mathbb{E}_{w_1,w_2}[\pi_i] = 0$.

In sum, (2.47) implies that if $L<L_1$ then $\frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[\pi_i] > 0$ and thus $\max_{d\in[0,1]}\mathbb{E}_{w_1,w_2}[\pi_i]$ happens at $d=1$. Similarly, if $L>L_2$ then $\frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[\pi_i]<0$ and thus $\max_{d\in[0,1]}\mathbb{E}_{w_1,w_2}[\pi_i]$ occurs at $d=0$. For the sake of completeness, we further note that $\arg\max_{d\in[0,1]}\ \mathbb{E}_{w_1,w_2}[\pi_i] \in \{0,1\}$ for *any* $L<\frac{s}{3}$.[25]

□

Proof of Proposition 5. Let $BR_i(\zeta) = \frac{1-\zeta}{2}$ denote i's best reply when $q_j=\zeta$. We aim to characterize the expected value of information sharing, which is given by

$$\begin{aligned}\mathbb{E}_{w_1,w_2}[W(K,K^c)] &= \Pr\{L,H\}W_{L,H}(K,K^c) + \Pr\{H,L\}W_{H,L}(K,K^c) \\ &\quad + \Pr\{L,L\}W_{L,L}(K,K^c) + \Pr\{H,H\}W_{H,H}(K,K^c), \qquad (2.48)\end{aligned}$$

where (according to (2.1)), $\Pr\{L,L\} = (1-\beta)(1-\frac{d\beta}{\beta+d(1-\beta)})$, $\Pr\{H,L\} = \Pr\{L,H\} = (1-\beta)\frac{d\beta}{\beta+d(1-\beta)}$, and $\Pr\{H,H\} = \beta\frac{\beta}{\beta+d(1-\beta)}$.

The benefit of cooperation/sharing information at state $\{w_1,w_2\}\in\{H,L\}^2$, denoted by $W_{w_1,w_2}(K,K^c)$, is given by (2.16), and $Q^K_{w_1,w_2}$ denotes total output at state (w_1,w_2) when firms cooperate and share their private information. Similarly, $Q^{K^c}_{w_1,w_2}$ denotes total output when firms compete with no shared information. We consider four separate cases as follows:

Case 1: $\{L,H\}$. In this case WP 1 is in the low state and WP 2 is in the high state. Information sharing increases total output because WP 2 can produce more energy, knowing for certain that WP 1 can only produce L units rather than $\mathbb{E}_{w_1}[q_1|w_2=H] = L\Pr\{L|H\}+$

[25]This is because any interior $\tilde{d}\in(0,1)$ such that $\frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[\pi_i]=0$ implies that $\frac{\partial}{\partial d}\mathbb{E}_{w_1,w_2}[\pi_i]\big|_{d>\tilde{d}}>0$. Thus any $d\in\{0,1\}$ for any d that maximizes profits.

$\phi \Pr\{H|H\} > L$. Therefore,

$$Q^K_{L,H} = L + BR_2(L) = L + \frac{1-L}{2} = \frac{1+L}{2}$$

whereas $Q^{K^c}_{L,H} = L + \phi$.

Case 2: $\{H, L\}$ This case by symmetry is identical to Case 1.

Case 3: $\{H, H\}$ In this case, information sharing reduces total output because the producers learn that the opposing producers have the ability to produce at the Cournot level, since information sharing eliminates the possibility that the other producer is in the low state (which causes them to overproduce, given that the other producer is in the high state). Under information sharing, both producer produce at the Cournot level. Therefore, $Q^K_{L,H} = 2q_C = \frac{2}{3}$. In the absence of information sharing $Q^{K^c}_{L,H} = 2\phi$.

Case 4: $\{L, L\}$ In this case WP 1 and WP 2 are both in the low state. Thus there is no difference between cooperation and competition since both produce at the L level, meaning that $W_{L,L}(K, K^c) = 0$.

Plugging these results into (2.48) and (2.16), we have that

$$\begin{aligned}\mathbb{E}_{w_1,w_2}[W(K, K^c)] = &\frac{d\beta(1-\beta)}{\beta + d(1-\beta)}\left(2\frac{1+L}{2} - \left(\frac{1+L}{2}\right)^2 - 2(L+\phi) + (L+\phi)^2\right)\\ &+ \frac{\beta^2}{\beta + d(1-\beta)}\left(\frac{2}{3} - \frac{1}{2}\left(\frac{2}{3}\right)^2 - 2\phi + \frac{1}{2}(2\phi)^2\right).\end{aligned}$$

By rearranging the above equation, we have that

$$\mathbb{E}_{w_1,w_2}[W(K, K^c)] = \Gamma(\beta, d, L)\bigg(39\beta + 28d(1-\beta) - 60Ld(1-\beta) - 81\beta L\bigg). \tag{2.49}$$

The common factor

$$\Gamma(\beta, d, L) = \frac{\beta^2 d(1-3L)(1-\beta)}{36\left(\beta + d(1-\beta)\right)\left(3\beta + 2d(1-\beta)\right)^2}$$

is positive because $L < 1/3$ (equivalently, $L < s/3$ for general s) by Assumption 1, $(1-\beta) \in (0, 1)$. Similarly, since $L < 1/3$, the additive terms of (2.49)

$$\begin{aligned}39\beta + 28d(1-\beta) - 60Ld(1-\beta) - 81\beta L &> 39\beta + 28d(1-\beta) - 20d(1-\beta) - 27\beta\\ &= 12\beta + 8d(1-\beta) > 0.\end{aligned}$$

The social welfare benefit of information sharing $\mathbb{E}_{w_1,w_2}[W(K,K^c)]$ is the product of two positive terms, and therefore $E_{w_1,w_2}[W(K,K^c)] > 0$. $\square$

Proof of Proposition 6. Let $BR_i(\zeta) = \frac{1-\zeta}{2}$ denote i's best reply when $q_j = \zeta$. We aim to characterize the following

$$\begin{aligned}\mathbb{E}_{w_1,w_2}[D(K,K^c)] = &\Pr\{L,H\}D_{L,H}(K,K^c) + \Pr\{H,L\}D_{H,L}(K,K^c) \\ &+ \Pr\{L,L\}D_{L,L}(K,K^c) + \Pr\{H,H\}D_{H,H}(K,K^c)\end{aligned}$$

where the benefit of cooperation/sharing information at state $\{w_1, w_2\} \in \{H, L\}^2$ is

$$D_{w_1,w_2}(K,K^c) = \pi^K_{1_{w_1,w_2}} + \pi^K_{2_{w_1,w_2}} - \pi^{K^c}_{1_{w_1,w_2}} - \pi^{K^c}_{2_{w_1,w_2}}$$

and $\pi^K_{i_{w_1,w_2}}$ denotes i's profit at state (w_1, w_2) when firms cooperate and share their private information. Similarly, $\pi^{K^c}_{i_{w_1,w_2}}$ denotes i's profit when firms compete with no information sharing (no cooperation). We consider four possible cases separately as follows.

Case 1: $\{L,H\}$. In this case WP 1 is in the low state and WP 2 is in the high state. Thus, information sharing is highly beneficial for WP 2 (and detrimental for WP 1). This is due to the fact that producer 2 will strategically overproduce and thus price goes down, hurting producer 1.

This overproduction is beneficial for producer 2, even though it results in a decrease in the equilibrium price. In this case, cooperation is beneficial for WP 2 and detrimental for WP 1, compared to competition with no information sharing. It is intuitive and important to note that the extra benefit to WP 2 from information sharing is **particularly high** when L is small. More precisely,

$$\begin{aligned}\pi^K_{1_{L,H}} &= L[1 - (L + BR_2(L))] \\ \pi^K_{2_{L,H}} &= BR_2(L)[1 - (L + BR_2(L))],\end{aligned}$$

where $BR_2(L) = \frac{1-L}{2}$. With no cooperation, each WP supplies according to the original equilibrium. Thus,

$$\begin{aligned}\pi^{K^c}_{1_{L,H}} &= L[1 - (L + \phi)] \\ \pi^{K^c}_{2_{L,H}} &= \phi[1 - (L + \phi)]\end{aligned}$$

with $\phi = \frac{\beta+(1-L)(1-\beta)d}{3\beta+2(1-\beta)d}$ by Corollary 1. With algebra we can show

$$\begin{aligned} D_{H,L}(K, K^c) &= \pi^K_{1_{H,L}} + \pi^K_{2_{H,L}} - \pi^{K^c}_{1_{H,L}} - \pi^{K^c}_{2_{H,L}} \\ &= \left(\frac{\beta(1-3L)}{2(2d(1-\beta)-3\beta)}\right)^2 - L\left(\frac{\beta(1-3L)}{2(2d(1-\beta)-3\beta)}\right). \end{aligned}$$

Case 2: $\{H, L\}$ This case by symmetry is similar to Case 1. Thus,

$$\begin{aligned} D_{H,L}(K, K^c) &= \pi^K_{1_{H,L}} + \pi^K_{2_{H,L}} - \pi^{K^c}_{1_{H,L}} - \pi^{K^c}_{2_{H,L}} \\ &= \left(\frac{\beta(1-3L)}{2(2d(1-\beta)-3\beta)}\right)^2 - L\left(\frac{\beta(1-3L)}{2(2d(1-\beta)-3\beta)}\right). \end{aligned}$$

Case 3: $\{H, H\}$ In this case WP 1 and WP 2 are both in the high state. Thus, cooperation is highly beneficial for both of them because by reducing uncertainty they both optimally coordinate and produce at the corresponding Cournot level, i.e. $q_C = \frac{1}{3}$. Thus the profit with information sharing is characterized as follows:

$$\pi^K_{1_{H,H}} = \pi^K_{2_{H,H}} = q_C[1 - 2q_C]$$

With no cooperation each WP best replies to her belief; thus,

$$\pi^{K^c}_{1_{H,H}} = \pi^{K^c}_{1_{H,H}} = \phi[1 - 2\phi]$$

where (as specified above) $\phi = \frac{\beta+(1-L)(1-\beta)d}{3\beta+2(1-\beta)d}$. Using algebra, we can show

$$\begin{aligned} D_{H,H}(K, K^c) &= \pi^K_{1_{H,H}} + \pi^K_{2_{H,H}} - \pi^{K^c}_{1_{H,H}} - \pi^{K^c}_{2_{H,L}} \\ &= 2q_C[1 - 2q_C] - 2\phi[1 - 2\phi] \\ &\geq 0. \end{aligned}$$

The function $f(x) = x(1 - 2x)$ is concave in x and is maximized at $x = \frac{1}{4}$. The last inequality follows because $\phi > q_C = \frac{1}{3}$ (because of the overproduction of each WP producing ϕ in this state due to uncertainty and mis-coordination), and thus $f(q_C) > f(\phi)$, because $\phi > q_C = \frac{1}{3} > \frac{1}{4}$.

Case 4: $\{L, L\}$ In this case WP 1 and WP 2 are both in the Low state. Thus there is no difference between cooperation and competition since both produce at the L level, meaning that $D_{L,L}(K, K^c) = 0$.

Plugging the results of the above cases into (2.17) implies that

$$\mathbb{E}_{w_1,w_2}[D(K,K^c)] = \frac{\beta^2 d(1-3L)(1-\beta)}{(3\beta+2d(1-\beta))^2(\beta+d(1-\beta))}\big[21\beta+16d(1-\beta)-L(81\beta+60d(1-\beta))\big].$$

As a result there exists a *unique* $L^*(d,\beta)$ such that $\mathbb{E}_{w_1,w_2}[D(K,K^c)] > 0$ if and only if

$$L < L^*(d,\beta) = \frac{21\beta+16d(1-\beta)}{81\beta+60d(1-\beta)}.$$

In the above expression, note that $L^*(d,\beta) < \frac{1}{3}$. □

Proof of Proposition 7. We seek to prove that there always exists a suitable t that satisfies (2.20), (2.21), and (2.22) when $\gamma = 0$. Let $\gamma = 0$ by assumption. Rearranging the IC condition gives

$$t \le \frac{1}{2}\Pr\{H|H\} + \Pr\{L|H\}(1 - \frac{\pi_L}{\pi_M}). \tag{2.50}$$

Rearranging the IR-H condition provides another upper bound on t:

$$t \le 1 + \frac{\beta}{2d(1-\beta)} - \frac{\beta}{d(1-\beta)}\frac{\phi(s-2\phi)}{\pi_M} - \frac{\phi(s-\phi-L)}{\pi_M} - \frac{\gamma}{\Pr\{L|H\}\pi_M}. \tag{2.51}$$

Rearranging the IR-L condition provides a lower bound for t:

$$t \ge \frac{L(s-\phi-L)}{\pi_M} + \frac{\gamma}{\Pr\{H|L\}\pi_M}. \tag{2.52}$$

The proof follows by showing that the lower bound for t, the right-hand side (RHS) of (2.52) is always less than or equal to the upper bounds for t from the RHS of (2.50) and (2.51) when $\gamma = 0$. Thus, there is always a nonempty feasible subset of $\mathbb{R}$ from which a transfer t can be selected that satisfies the criteria for collusion.

First, with $\gamma = 0$, the RHS of (2.52) is less than the RHS of (2.50). The RHS of (2.52),

$$\frac{L(s-\phi-L)}{\pi_M} < \frac{L(s-2L)}{\pi_M} \le 1/2,$$

where the first inequality is due to $\phi > L$ and the second is because $L(s-2L)$ is maximized at $\frac{s^2}{8}$ when $L = s/4$, and because $\pi_M = \frac{s^2}{4}$. Furthermore, from the RHS of (2.50),

$$\frac{1}{2}\Pr\{H|H\} + \Pr\{L|H\}(1-\frac{\pi_L}{\pi_M}) \ge \frac{1}{2}\Pr\{H|H\} + \Pr\{L|H\}(1-\frac{4}{8}) = 1/2$$

where the inequality is because $\pi_L \leq \frac{s^2}{8}$ and the equality is because the expression is a weighted probabilistic sum of two values equal to 1/2. Therefore, the lower bound on t, (2.52) is at most $\frac{1}{2}$, and one upper bound on t, (2.50), is at least $\frac{1}{2}$.

Now, it remains to be shown that the RHS of (2.51) (the other upper bound on t) is at least as great as the RHS of (2.52). Equivalently, their difference T is greater than or equal to zero:

$$\begin{aligned} T &= 1 + \frac{\beta}{2d(1-\beta)} - \frac{\beta}{d(1-\beta)}\frac{\phi(s-2\phi)}{\pi_M} - \frac{\phi(s-\phi-L)}{\pi_M} - \frac{L(s-\phi-L)}{\pi_M} \\ &= \frac{\beta}{d(1-\beta)}(\frac{1}{2} - \frac{\phi(s-2\phi)}{\pi_M}) + \frac{\pi_M - (\phi+L)(s-\phi-L)}{\pi_M} \geq 0. \end{aligned}$$

The second line is a rearrangement of the first. The first term in the second line is positive because $\pi_M = \max_\phi 2\phi(s-2\phi)$ by the definition of monopoly profits; therefore, $\frac{\phi(s-2\phi)}{\pi_M} \leq \frac{1}{2}$. The second term is positive since $\pi_M = \max_{\phi,L}(\phi+L)(s-\phi-L)$, again by the definition of Monopoly profits. Therefore, none of the aforementioned constraints contradict, and there always exists some t that allows the producers to collude. □

Proof of Proposition 8. Let ϕ^d be the equilibrium high state production that for each $i \in \{1, ..., N+1\}$ satisfies

$$\mathbb{E}_{S^d_{-i}}[P(\phi^d + S^d_{-i}\phi^d + (N - S^d_{-i})L) + \phi^d P'(\phi^d + S^d_{-i}\phi^d + (N - S^d_{-i})L)|w_i = H] = 0. \quad (2.53)$$

The left hand side of (2.53) is decreasing in ϕ^d. Its derivative with respect to ϕ^d is

$$\mathbb{E}_{S^d_{-i}}[(2+S^d_{-i})P'(\phi^d+S^d_{-i}\phi^d+(N-S^d_{-i})L)+\phi^d P''(\phi^d+S^d_{-i}\phi^d+(N-S^d_{-i})L)|w_i = H] < 0 \quad (2.54)$$

which follows because $2 + S^d_{-i} > 0$, $\phi^d > 0$, $P' < 0$, and $P'' \leq 0$.

Now, consider $d' > d$ and assume towards a contradiction that $\phi^{d'} < \phi^d$, where ϕ^d and $\phi^{d'}$ satisfy the first order condition in (2.53), with the expectations taken according to their respective random variables S^d_{-i} and $S^{d'}_{-i}$:

$$\begin{aligned} 0 &= \mathbb{E}_{S^d_{-i}}[P(\phi^d + S^d_{-i}\phi^d + (N - S^d_{-i})L) + \phi^d P'(\phi^d + S^d_{-i}\phi^d + (N - S^d_{-i})L)|w_i = H] \\ &< \mathbb{E}_{S^d_{-i}}[P(\phi^{d'} + S^d_{-i}\phi^{d'} + (N - S^d_{-i})L) + \phi^{d'} P'(\phi^{d'} + S^d_{-i}\phi^{d'} + (N - S^d_{-i})L)|w_i = H] \\ &\leq \mathbb{E}_{S^{d'}_{-i}}[P(\phi^{d'} + S^{d'}_{-i}\phi^{d'} + (N - S^{d'}_{-i})L) + \phi^{d'} P'(\phi^{d'} + S^{d'}_{-i}\phi^{d'} + (N - S^{d'}_{-i})L)|w_i = H]. \end{aligned}$$

The first inequality is due to (2.54), with $\phi^{d'} < \phi^d$. The second inequality is due to Assumption 2 with $F(x) = P(\phi^{d'} + x\phi^{d'} + (N-x)L) + \phi^{d'} P'(\phi^{d'} + x\phi^{d'} + (N-x)L)$ decreasing in x. But the result implies that $\phi^{d'}$ does not satisfy the first order condition of the equilibrium, so we have a contradiction. Therefore, ϕ^d is (weakly) increasing in d. □

Proof of Proposition 9. By definition, $W = U(Q)$ where $Q = \sum_{i=1}^{N+1} q_i$. Note that $U'(Q) = P(Q) \geq 0$ for any equilibrium Q, and $P' < 0$, $\phi > 0$. Furthermore, note that we can write Q as a function of d given any realization of availability, $Q(d) = \sum_i s_i(\phi^d - L) + (N+1)L = (\phi^d - L)S + (N+1)L$, which is increasing and linear in S. This implies that the expectation $\mathbb{E}[W]$, which is taken over the random states of all of the producers, is fully defined by the probability distribution of S. Let $\mathbf{s}$ be the random vector of states of each of the producers, i.e. $\mathbf{s} = [s_1, s_2, ...s_{N+1}]$. Then for all i, s_i, with $S = \sum_i s_i$, we have that $\mathbb{E}_{\mathbf{s}}[W(Q(d))] = \mathbb{E}_S[W(Q(d))]$. Then for $d' > d$:

$$\begin{aligned}
\mathbb{E}_{\mathbf{s}^{d'}}[W] - \mathbb{E}_{\mathbf{s}^d}[W] &= \mathbb{E}_{S^{d'}}[W(Q(d'))] - \mathbb{E}_{S^d}[W(Q(d))] \\
&= \sum_{k=1}^{N+1} U(Q(d')) \Pr\{S^{d'} = k\} - \sum_{k=1}^{N+1} U(Q(d)) \Pr\{S^d = k\} \\
&\geq \sum_{k=1}^{N+1} U(Q(d))(\Pr\{S^{d'} = k\} - \Pr\{S^d = k\}) \\
&= \mathbb{E}_{\mathbf{s}^{d'}}[U(Q(d))] - \mathbb{E}_{\mathbf{s}^d}[U(Q(d))] \geq 0.
\end{aligned}$$

The first line is because S provides equivalent information for the expectation as explained above. The second line is an expansion of the expectations for the discrete random variables, and the third line is because of Proposition 8 with $Q(d)$ increasing in d and $U(Q)$ increasing in Q. The fourth line rewrites the third as a difference of expectations, which is non-negative by the definition of second-order stochastic dominance with U increasing and concave in S. U is increasing and concave; $Q(S) = (\phi - L)S + (N+1)L$ is linear and increasing in S. □

Proof of Example 2. For the linear inverse demand, conditions (2.28) and (2.29) are represented by (2.55) and (2.56).

$$\Pr\{L|H\}(s - L - 2\phi - x) + \Pr\{H|H\}(s - 3\phi - x) = 0 \tag{2.55}$$

$$\Pr\{L, L\}(s - 2L - 2x) + 2\Pr\{L, H\}(s - L - \phi - 2x) + \Pr\{H, H\}(s - 2\phi - 2x) - c = 0 \tag{2.56}$$

Under linear inverse demand, (2.55) is the first order condition for the wind generators, and (2.56) is the first order condition for the traditional generator with constant marginal cost *c*. Equation (2.55) allows us to write the expression for ϕ in equilibrium, analogous to the result from Corollary 1 but including the effect of x.

$$\phi = \frac{(s-x)\beta + (s-x-L)d(1-\beta)}{3\beta + 2d(1-\beta)} = \frac{s\beta + (s-L)d(1-\beta) - x(\beta + d(1-\beta))}{3\beta + 2d(1-\beta)} \tag{2.57}$$

We can rearrange the generator's first order condition (2.56) to obtain equation (2.31). By the assumption, with $\phi < H$, the traditional generator chooses to participate in the market, i.e. (2.31) is solved by some $x \geq 0$. Combining equations (2.57) and (2.31), we have the result in (2.30). The uniqueness of the symmetric[26] equilibrium is clear from the fact that the result for ϕ in (2.30) does not depend on x. □

Proof of Proposition 11. Average output is given by

$$\begin{aligned}\mathbb{E}[q_1(w_1) + q_2(w_2) + x] &= 2\Pr\{L,H\}(L+\phi+x) + \Pr\{L,L\}(2L+x) + \Pr\{H,H\}(2\phi+x) \\ &= x + 2\beta\phi + 2(1-\beta)L.\end{aligned}$$

Taking the derivative with respect to d, we have that

$$\frac{\partial \mathbb{E}[q_1(w_1) + q_2(w_2) + x]}{\partial d} = \beta \frac{\partial \phi}{\partial d} > 0,$$

due to the linearity of expectation and because $\frac{\partial \mathbf{E}[q_i(w_i)]}{\partial d} = \beta \frac{\partial \phi}{\partial d}$, for $i \in \{1,2\}$, and $\frac{\partial x}{\partial d} = -\beta \frac{\partial \phi}{\partial d}$. Now, average welfare is given by

$$\mathbb{E}_{w_1,w_2}[W] = 2\Pr\{L,H\}U(L+\phi+x) + \Pr\{L,L\}U(2L+x) + \Pr\{H,H\}U(2\phi+x) - c(x).$$

Taking the derivative with respect to d,

$$\begin{aligned}\frac{\partial \mathbb{E}_{w_1,w_2}[W]}{\partial d} &= \zeta(2U(L+\phi+x) - U(2L+x) - U(2\phi+x)) + \frac{\partial x}{\partial d}\Pr\{L,L\}P(2L+x) \\ &\quad + 2(\frac{\partial \phi}{\partial d} + \frac{\partial x}{\partial d})\Pr\{L,H\}P(L+\phi+x) + (2\frac{\partial \phi}{\partial d} + \frac{\partial x}{\partial d})\Pr\{H,H\}P(2\phi+x) - c\frac{\partial x}{\partial d} \\ &= \Gamma + 2\frac{\partial \phi}{\partial d}\Pr\{L,H\}P(L+\phi+x) + 2\frac{\partial \phi}{\partial d}\Pr\{H,H\}P(2\phi+x) + \frac{\partial x}{\partial d}x \\ &= \Gamma + \beta\frac{\partial \phi}{\partial d}(2\phi - x) \geq \Gamma \geq 0.\end{aligned}$$

The second and third lines replace the first term with $\Gamma = \zeta(2U(L+\phi+x) - U(2L+$

[26] The equilibrium is symmetric in the sense that the wind producers have identical strategies; the traditional producer has a different objective and a different strategy from the wind producers.

$x) - U(2\phi + x))$ to concatenate the expression; Γ is the impact of wind diversification on welfare, and it is weakly positive due to the concavity of U, as explained in Proposition 2. The second equality uses the first order condition from (2.56). The third equality uses the conditional probabilities $\Pr\{L, H\} = \Pr\{L|H\}\beta$ and $\Pr\{H, H\} = \Pr\{H|H\}\beta$, along with the first order condition from (2.55). The expression $2\phi - x$ is minimized when $c = 0$ and when $d = 0$. Therefore, by using (2.30) and (2.31), with $c, d = 0$, we confirm that $2\phi - x \geq 0$. This fact and $\Gamma \geq 0$ establish the inequalities in the final line and complete the proof. □

Proof of Proposition 12. Average price can be expressed as

$$\mathbb{E}_{w_1,w_2}[P] = 2\Pr\{L, H\}P(L + \phi + x) + \Pr\{L, L\}P(2L + x) + \Pr\{H, H\}P(2\phi + x).$$

Taking the derivative with respect to d gives:

$$\begin{aligned}
\frac{\partial \mathbb{E}_{w_1,w_2}[P]}{\partial d} &= \zeta(2P(L + \phi + x) - P(2L + x) - P(2\phi + x)) - 2\frac{\partial \phi}{\partial d}(\Pr\{L, H\} + \Pr\{H, H\}) - \frac{\partial x}{\partial d} \\
&= \zeta 0 - 2\beta\frac{\partial \phi}{\partial d} + \beta\frac{\partial \phi}{\partial d} = -\beta\frac{\partial \phi}{\partial d}.
\end{aligned}$$

This is due to the fact that $P(x)$ represents linear inverse demand, so the first term sums to 0 and so $\forall x\ P'(x) = -1$, and also due to the fact that $\Pr\{L, H\} + \Pr\{H, H\} = \beta$. This completes the proof. As in the two-producer case, for a linear inverse demand curve, average price is monotonically decreasing in d.

□

Chapter 3

Forward Contracting in Electricity Markets with Retail Deregulation

This work was performed in collaboration with Audun Botterud and Mardavij Roozbehani.

3.1 Introduction

Capacity markets for electricity provide a regulated market setting through which generating units are compensated for their contribution to power system reliability, the ability of the power system to meet peak demand. In many regulated markets, for instance in the U.S., the independent system operator (ISO) sets a demand curve for capacity for the region and charges load-serving entities (LSEs) based on their contribution to the peak system load. As such, the capacity market essentially serves as a market for a specific type of long-term contract that consumers are required to purchase. The type of forward contract varies, but can be modeled similarly as a type of capacity certificate or reliability option, as demonstrated by an analytical comparison of forward contract types (Léautier, 2016).

While capacity markets have diverse forms and requirements, compelled participation of demand is a key feature shared in many markets: customers, or LSEs acting on their behalf, are required to engage in a specified level of contracting by paying for the forward capacity quantity that has been determined in an auction process based on the ISO's demand curve.

The main rationale for capacity markets is to help generators achieve revenue sufficiency in a market with price caps, which are used to mitigate generator market power. However, capacity markets also provide many of the benefits of financial forward contracts, including risk reduction. Changing energy markets and increased penetration of variable renewable resources have strained capacity markets or led to apparent capacity shortages or excesses in some markets. This has driven increased focus on the benefits and costs of capacity markets (Bushnell et al., 2017), as well as additional efforts to define the market failures that capacity markets should and can seek to address (Cramton et al., 2013).

Some researchers have advocated for scaling back capacity markets, arguing that social welfare would be better served by direct efforts to reduce market power (Léautier, 2016). Others suggest that a combination of low price caps and higher capacity payments can actually increase market concentration (Elberg and Kranz, 2014). Hogan (2005) has advocated for market design without capacity markets, and ERCOT has implemented an energy-only market design. In existing capacity market designs, administrators determine the capacity value of renewable energy; this leads to inefficiencies because regulators, not the market, determine the value of new renewable energy projects. As the penetration of renewable resources increases, the design of capacity markets will become increasingly important for economic efficiency in the power sector.

Many of the benefits of forward contracting in a capacity market can clearly be achieved by optional financial forward contracts in an energy-only market design. For instance, financial forward contracts can help reduce risks associated with price uncertainty and counterparty risk, and they can ease financing of lumpy generation investments. Forward contracting among market participants can still be expected to be an important component of an energy-only market design (Hogan, 2005). Most of the benefits of forward contracting are not coupled with market failure problems that would compel mandatory participation in forward contracts or capacity markets; instead, we should expect LSEs to rationally select the appropriate level of forward contracts, taking into account the benefits they provide.

However, while research suggests that one significant benefit of forward contracting is its ability to help reduce generator spot market power, this work argues that spillover effects, or positive externalities, related to forward contracting and producer market power might limit the extent of discretionary forward contracting and reduce social welfare. This effect is not necessarily limited to electricity markets, but this research focuses in particular on the characteristics of forward and spot markets for electricity.

A rich literature suggests that forward contracting can reduce market power in electricity spot markets. Allaz and Vila (1993) provided analytical evidence that forward contracting can impel producers to offer higher quantities in real-time markets. Wolak (2000) provides empirical evidence in support of this conclusion, using data from the Australian market to show that forward hedging can reduce generator market power.

In a two-stage model with reliability options, a specific form of forward contract, Léautier (2016) also shows that producers that have sold forward contracts have less ability to exercise market power in spot markets. Chao and Wilson (2004), Cramton and Stoft (2008), and Ausubel and Cramton (2010) argue that one of the major benefits of capacity markets using reliability options is their ability to help reduce generator market power. Some research has questioned the impact of forward contracting on market power, for instance, when firms in a duopoly compete in prices instead of quantities (Mahenc and Salanié, 2004) or when firms are capacity constrained (Dappe, 2008). Harvey and Hogan (2000) question whether the model by Allaz and Vila (1993) is a useful model for the California Electricity market.

Despite uncertainty regarding the extent to which forward contracting mitigates market power, Allaz and Vila (1993) present a useful model for understanding the ways in which this impact occurs. The research in this chapter takes those effects as given in order to study how demand side market concentration (i.e. the number of LSEs) impacts the extent of forward contracting and market power.

The main contribution of this research is to show how competition amongst load-serving entities (as buyers in the forward contract market) impacts social welfare in electricity markets. It is the first research, to our knowledge, that studies spillover effects of forward contract purchase: forward contract purchases that reduce market power benefit all consumers, not just those represented by the individual LSE making the forward contract purchasing decision.

This research extends the literature by focusing specifically on how demand-side characteristics impact forward contracting. It is relevant because policy choices dictate the characteristics of the demand-side of the market. Following electricity restructuring, many states allow consumers to choose their energy supplier in a competitive marketplace.[1] In some U.S. states, municipalities also have the option to choose an energy supplier for their residents. In 2017, over 13.7 million U.S. customers participated in retail choice programs (EIA, 2018); we say that these customers are using competitive or retail energy suppliers, and the remaining customers are using a default supplier. In this chapter, competitive energy suppliers and default suppliers (usually regulated utilities) are collectively referred to as LSEs.

[1]Some potential benefits including lower energy prices for consumers, increased product offerings (e.g. energy supply with a high percentage of renewable energy), and innovative business models that compensate consumers for demand flexibility. However, the extent of consumer benefits is uncertain, and critics have also noted several downsides associated with competitive energy suppliers.

The model is based on two-stage procurement of energy, first through forward contracts and then in real-time. As in the model by Allaz and Vila (1993), producers have market power in the second stage. This research presents new contributions regarding the effects of firm concentration on the demand-side. In our research, LSEs try to maximize welfare on behalf of their consumers; in a departure from previous research, the LSEs have demand-side market power. This is an important modeling choice because it reflects real-world electricity markets, where oligopolistic energy suppliers compete for profits and customers.

The focus on demand-side competition leads to new modeling and analytical challenges. We model demand-side market power in forward contracting by allowing LSEs to set their level of forward contracting. We find the equilibrium market conditions, including total energy consumption, total forward contracting, energy price, and social welfare. We use comparative statics to examine how demand-side competition impacts these equilibrium conditions.

This research shows that spillover benefits associated with forward contracting serve to decrease the total forward contracting level and social welfare. In markets with several sufficiently large LSEs, welfare is decreasing in the number of LSEs. The research posits that positive externalities associated with forward contract procurement, under certain assumptions, provide a rational for requiring consumer participation in forward contracting or capacity markets, or for internalizing the public benefits of forward contracting.

3.2 Model

We present an electricity market model that is significantly simplified for ease of tractability and clarity, but which retains the core features of demand uncertainty, competition among producers, and the presence of multiple retail electricity suppliers. The producer side of the model is similar to (Allaz and Vila, 1993), but we allow producers to have generic valuations of forward contracts and then determine the level of forward contracting preferred by consumers in equilibrium.

Consider a model where producers $i \in \{1, 2, ..., N\}$ sequentially offer q_i units of energy in a forward contract and then sell x_i total units of energy in a given hour. Throughout the chapter, we consider q_i as the quantity of energy offered in each hour, so that all variables are equivalently in units of energy (kWh or MWh).

For simplicity, we ignore the distinction between short-term markets, like day-ahead and real-time markets, and we make the simplifying assumption that the real-time wholesale price is given by the linear inverse demand function $P(X) = t - bX$, where X is the total supply or demand, $X = \sum_{i \in N} x_i$, and with $b \in \mathbb{R}^+$. The consumers' utility from consuming X units at demand level t is $V(X,t) = \int_0^X (t - bz)dz$.

The parameter t represents the realization of a random variable T; it corresponds to total net energy demand in a given period, based on the difference between the maximum energy demand and the realized renewable energy availability. Specifically, $\frac{t}{b}$ is the maximum net energy demand, the energy demand at a price of 0 minus the renewable energy supply. This model allows for the analysis of regions with large renewable energy supply, if renewable supply exceeds max demand, then $\frac{t}{b} < 0$ or some renewable energy is curtailed and $\frac{t}{b} = 0$ in the particular period under consideration.

Renewable energy generators are assumed to bid competitively at zero marginal cost, providing all available energy in every period with a positive price. The parameter $b \in \mathbb{R}$ is fixed and T is varied as a random variable, and realized in real-time. In practice, the inverse demand curve need not be linear, and its slope could vary with demand in addition to its intercept. In the forward contracting stage, uncertainty can be modeled by the cumulative distribution function $F(t)$ and the associated density function $f(t)$. Assume that T has finite support on $[\underline{t}, \bar{t}]$ with $\underline{t}, \bar{t} \in \mathbb{R}$. The expectations over T are all taken with respect to future demand over a generic time-period, so they only affect the forecast for the distribution $F(t)$. Since $F(t)$ is exogenous to the results, and considered fixed at the start of the contracting period, the length of the forward contract period is entirely generic; it does not impact the model and associated results described herein.

In our model, generators have a linear cost c_i for production of electricity in each hour; in practice, their production can include nonlinearities, for instance due to startup costs. Let $s_i = x_i - q_i$ represent producer i's net output sold in real-time. Let $p_W = P(x)$ be the wholesale spot market price for electricity in a given hour with production x. Then the profit for producer i in a given hour, who has previously sold q_i units of a forward contract at price p_F, is given by

$$\pi_i^W = p_W s_i + p_F q_i - c_i(s_i + q_i). \tag{3.1}$$

Based on the impact of forward contracts on their profits, and also on additional features like their own risk preferences, generators have a generic price that they demand per unit of forward contract sold. This price, for instance, could represent a small or a substantial discount from the average real-time price, or it could be higher than the real-time price in the case of much more risk-aversion for demand than for generators.

Furthermore, the model features LSEs $j \in \{1, 2, ..., M\}$ that procure forward contracts

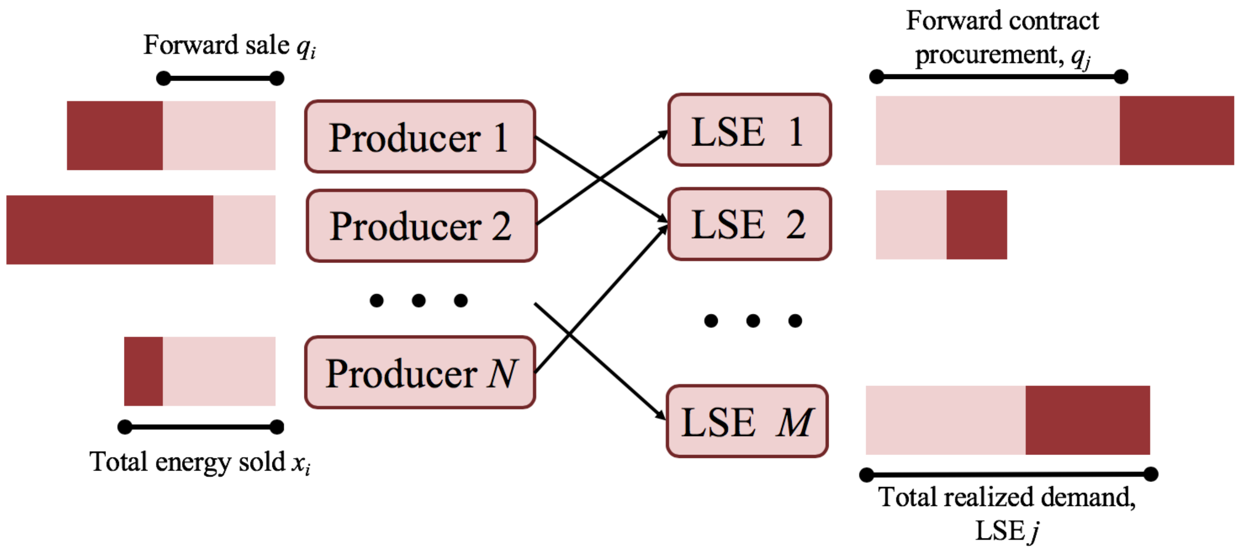

Figure 3-1: Relationships and key variables in the forward contracting energy market model

and spot-market electricity in order to maximize the welfare of consumers. Specifically, each LSE chooses a forward contract quantity q_j in order to maximize the expected welfare for its fraction of consumers. Each LSE services $\alpha_j \in (0, 1]$ fraction of consumers; we assume that the demand fractions are constant in time, so if the total consumption is $Z(t)$ in a specific period with $T = t$, then the consumption by LSE j is $\alpha_j Z(t)$. This is a reasonable assumption when the demand of consumers across retailers does not vary significantly.

Figure 3-1 depicts the model described in this section. N producers sell energy to M load serving entities, through both forward contracts and the spot market. Each producer $i \in \{1, 2, ..., N\}$ sells q_i units in the forward market and x_i total units of energy. In the first stage, with imperfect information about future demand, each LSE $j \in \{1, 2, ..., M\}$ procures q_j units of forward contracts. The sum of forward contracts is balanced, i.e. $\sum_{i=1}^{N} q_i = \sum_{j=1}^{M} q_j$. In the second stage, 'real-time,' the demand parameter is realized and producers have the opportunity to sell more energy in a spot market with Cournot competition. Each LSE's total quantity demanded is based on the demand parameter and the quantities offered in the spot market $\{x_1 - q_1, x_2 - q_2, ..., x_N - q_N\}$ by strategic producers; the equilibrium levels for the spot market quantities are described in the subsequent Section 3.3.

3.3 Equilibrium Spot Market Output by Producers

In this section, we find the symmetric Nash equilibrium strategy for producers in the spot market who have already sold a fixed level of forward contracts, in the case of linear inverse demand.

Proposition 1. The symmetric Nash equilibrium spot market output for all producers $i \in \{1, 2, ..., N\}$ market with prior forward commitments $(q_1, q_2, ..., q_N)$ is given by:

$$s_i = \frac{t - bQ + C}{b(N+1)} - \frac{c_i}{b} \tag{3.2}$$

where $Q = \sum_{i=1}^{N} q_i$ and $C = \sum_{i=1}^{N} c_i$.

Proof. Each producer seeks to maximize (3.1). Let $S = \sum_{i=1}^{N} s_i$. The spot price $p_W = P(S+Q)$. In the case of linear inverse demand, the first order condition of (3.1) is given by:

$$bs_i = t - bS - bQ - c_i. \tag{3.3}$$

Summing over all i, this condition requires that

$$S = \frac{Nt - bNQ - C}{b(N+1)}. \tag{3.4}$$

By substituting this result into the original condition, and combing terms, we have (3.2). Note that the second derivative of profit with respect to net output s_i

$$\frac{\partial^2 \pi_i^W}{\partial s_i^2} = -2b < 0. \tag{3.5}$$

Therefore, the solution to each individual producer's optimization problem is unique. The equilibrium in (3.2) is also unique; note that s_i is not impacted by the decision s_j for any $j \neq i$. □

In the case where generators each have equivalent marginal costs, i.e. $(c_1, c_2, ...c_N) = (c, c, ..., c)$, then (3.2) simplifies to

$$s_i = \frac{t - bQ - c}{b(N+1)}. \tag{3.6}$$

Going forward, we focus on generators with equal marginal costs.

3.4 Equilibrium Forward Contracting by Multiple Consumers

Now we consider the equilibrium forward contracting level by multiple entities representing consumers of electricity. Each LSE $j \in \{1, 2, ..., M\}$ chooses a forward contract level to maximize their consumer surplus,

$$\pi_j^F = \mathbb{E}_T[\alpha_j V(X,T) - (\alpha_j X - q_j)p_W(T) - q_j p_F] \tag{3.7}$$

$$= \int_{\bar{t}}^{\underline{t}} \left(\int_0^X \alpha_j(t - bz)dz - (\alpha_j X - q_j)(t - bX) - q_j p_F \right) f(t)d(t) \tag{3.8}$$

which is $\alpha_j V(X,t)$ the consumers' utility of energy consumption at demand level t for the LSEs fraction of customers α_j, minus the LSEs real-time demand times the real-time price, minus the LSEs chosen level of forward contracts times the forward contract price. Note that the output / consumption level $X = S + Q$, where S is set by the Cournot competition amongst producers in real-time according to (3.2), and is itself is a function of t. Furthermore, we assume that $q_j \geq 0$ and focus on cases where LSEs each procure a positive amount of forward contracts in equilibrium.

Proposition 2. Assume that each firm $j \in M$ procures an unconstrained quantity of forward contracts, i.e. $q_j > 0$. Then, the equilibrium forward contract level for the sum of forward contracts from M LSEs is

$$Q = \frac{\mathbb{E}[T]}{b} + \frac{(MN(N+1) - N)c - M(N+1)^2 p_F}{b(MN + M + N)}. \tag{3.9}$$

Proof. Each LSE simultaneously chooses the forward contract amount q_j to maximize their consumer surplus in (3.8). The individual first order conditions satisfy

$$\frac{\partial \pi_j^F}{\partial q_j} = \frac{\partial}{\partial q_j} \mathbb{E}_T[\frac{\alpha_j}{2} bX^2 + q_j(p_W - p_F)] \tag{3.10}$$

$$= \mathbb{E}_T[\frac{\partial}{\partial q_j}(\frac{\alpha_j}{2} bX^2 + q_j(p_W - p_F))] \tag{3.11}$$

$$= \mathbb{E}_T[\alpha_j bX \frac{\partial X}{\partial q_j} + p_W - p_F + q_j \frac{\partial p_W}{\partial q_j}] = 0. \tag{3.12}$$

The first line is due to (3.8), the second line is because we can move the integral inside the expectation because t and its functions are bounded, and the third line is a computation of the derivative.

Next, we solve explicitly for S, p_W, and their respective derivatives for any t in the support of T. Real-time production / consumption X is given by the sum of (3.6) over all N plus the total forward contract quantity, i.e.,

$$X = S + Q = \frac{N(t - bQ - c)}{b(N+1)} + Q = \frac{Nt + bQ - Nc}{b(N+1)} \tag{3.13}$$

where S is given by (3.6). Therefore, $\frac{\partial X}{\partial q_j} = \frac{1}{N+1}$. Given X, we can find p_W, which is given by

$$p_W = t - bX = \frac{t - bQ + Nc}{N+1}. \tag{3.14}$$

Therefore, $\frac{\partial p_W}{\partial q_j} = \frac{-b}{N+1}$. Subbing this into (3.12), and noting that the expectation is now linear in t, we have that

$$0 = \alpha_j \frac{N\mathbb{E}[T] + bQ - Nc}{(N+1)^2} + \frac{\mathbb{E}[T] - bQ + Nc}{N+1} - p_F + q_j \frac{-b}{N+1}. \tag{3.15}$$

Summing the above over j, and multiplying each term by $(N+1)^2$ gives

$$0 = N\mathbb{E}[T] + bQ - Nc + (N+1)(M\mathbb{E}[T] - MbQ + MNc - M(N+1)p_F - bQ). \tag{3.16}$$

Rearranging terms,

$$b(M(N+1) + N)Q = (M(N+1) + N)\mathbb{E}[T] + (MN(N+1) - N)c - M(N+1)^2 p_F. \tag{3.17}$$

The final result (3.9) is due to a simple rearranging of the above. □

Remark 1. The first order condition for individual producers, (3.12), allows for additional insight into the minimum population fraction requirement for participators in the M producer equilibrium. The optimal equilibrium quantity $q_j > 0$ iff

$$\mathbb{E}\left[\alpha_j X + \frac{N+1}{b}(p_W - p_F)\right] = \mathbb{E}\left[\alpha_j X + \frac{N+1}{b}(T - bX - p_F)\right] > 0. \tag{3.18}$$

This is equivalent to the requirement that

$$\alpha_j > (N+1)\mathbb{E}\left[1 + \frac{p_F - T}{bX}\right]. \tag{3.19}$$

The market price equation $p_W = t - bX$ and Assumption 5 imply that $\mathbb{E}\left[\frac{1}{bX}\right] > \frac{1}{\mathbb{E}[T]-c}$. This, and the fact that expected maximum willingness to pay $\mathbb{E}[T]$ is greater that p_F provide a sufficient condition for α_j,

$$\alpha_j > (N+1)\frac{p_F - c}{\mathbb{E}[T] - c}. \tag{3.20}$$

If N is very high this will be hard to satisfy because the benefits of forward contracts for reducing market power are proportionally smaller. Similarly, if $p_F >> c$ this will be hard to satisfy because forward contracts will be expensive for consumers.

In general, however, markets that fulfill or approximately fulfill this condition represent typical (not extreme) examples. For instance, if there are 5 producers, the forward contract price is 20% higher than the cost of electricity, and the maximum peak willingness to pay is 10 times the marginal cost of electricity, then each LSE with at least $\frac{1}{9}$ of the population would procure a positive quantity of forward contracts. Our model focuses on the example where each LSE is sufficiently large; e.g., with up to 9 LSEs each representing at least $\frac{1}{9}$ of the total customers.

3.5 Demand Competition and Forward Contracting

This section presents two simple result that represents the key ideas of the chapter. The equilibrium level of forward contracting is decreasing in M, suggesting that a competitive group of LSEs would, in equilibrium, purchase a lower total level of forward contracts than a single buyer. Furthermore, due to this decreasing level of forward contracting, welfare is also decreasing in M.

The proofs makes use of a single additional assumption requiring that prices support market participation for suppliers with cost c.

Assumption 5. Profitable Market: The market supports positive profits for suppliers with marginal cost c, i.e. either $\mathbb{E}p_W > c$ or $p_F > c$.

Either the average real-time price or the price of forward contracts must exceed the producers' marginal costs of production. If this assumption was not true we would either be witnessing a perfectly competitive market or we would expect to see suppliers exit the markets. However, we know that suppliers are engaged in oligopolistic competition. Therefore, we expect to see some (potentially very small) supplier profits and this assumption follows naturally from the basic problem assumptions. Note that using the results (3.6) and (3.9) we can make the assumption $\mathbb{E}[p_W] > c$ explicit in terms of the problem parameters.

$$\mathbb{E}p_W - c = \mathbb{E}\left[\frac{t - bQ - c}{N+1}\right] = \frac{M(N+1)p_F - M(N+1)c}{MN + M + N} > 0, \tag{3.21}$$

which is clearly true iff $p_F > c$. Therefore, the two statements are equivalent and the Assumption 5 is equivalent to the requirement that $p_F > c$.

Proposition 3. Given a market that satisfies Assumption 5, the sum of forward contracts procured by demand entities Q is decreasing in M, $\frac{\partial Q}{\partial M} < 0$.

Proof. We simply compute the derivative of the expression for the equilibrium level of forward contracting

$$\frac{\partial Q}{\partial M} = \frac{(N(N+1)c - (N+1)^2 p_F)b(MN+M+N)}{b^2(MN+M+N)^2} - \frac{(N+1)b((MN(N+1)-N)c - M(N+1)^2 p_F)}{b^2(MN+M+N)^2}. \tag{3.22}$$

By simplifying the expression, it is clear that

$$\frac{\partial Q}{\partial M} = \frac{N(N+1)^2(c - p_F)}{b(MN+M+N)^2} < 0 \tag{3.23}$$

where the inequality is due to Assumption 5, $p_F > c$. □

3.5.1 Forward Contracting and Social Welfare

Now consider the effects of the forward contracting level and demand competition on welfare. Expected welfare $\mathbb{E}_T \Pi$ is given by

$$\mathbb{E}[\Pi] = \mathbb{E}[V(X,T) - (X-Q)p_W(t) - Qp_F + (X-Q)p_W(T) + Qp_F - Xc] \tag{3.24}$$

$$= \mathbb{E}[\int_0^X (T - bz)dz - Xc] = \mathbb{E}[TX - \frac{b}{2}X^2 - Xc] \tag{3.25}$$

which is the sum of the expected value of consumption for consumers with expectation taken over demand T minus the expected cost of energy at production level X, which is itself a function of T, again taken over demand T. The middle terms in the first line represent the transfers from consumers to producers for real-time energy and forward contracts, respectively; in terms of expected social welfare, these net to zero. Next we show that the expected welfare is decreasing in the number of LSEs.

Proposition 4. Under the competitive market Assumption 5, $\frac{\partial \mathbb{E}\Pi}{\partial M} < 0$.

Proof. The derivative of equilibrium social welfare with respect to Q, is given by:

$$\begin{aligned}\frac{\partial \mathbb{E}\Pi}{\partial Q} &= \mathbb{E}[\frac{\partial}{\partial X}(tX - \frac{b}{2}X^2 - Xc)\frac{\partial X}{\partial Q}] \\ &= \frac{1}{N+1}\mathbb{E}[t - bX - c] \\ &= \frac{1}{N+1}(\mathbb{E}[p_W] - c) > 0\end{aligned} \tag{3.26}$$

where the first line is due to the boundedness of the terms in the expected value, the second line is a computation of the derivative, and the final line is due to the definition of p_W and Assumption 5, since the equivalence of the two statements in the assumptions requires that each is true. Therefore,

$$\frac{\partial \mathbb{E}\Pi}{\partial M} = \frac{\partial \mathbb{E}\Pi}{\partial Q}\frac{\partial Q}{\partial M} < 0. \tag{3.27}$$

The first term of the product is positive due to (3.26) and the second is negative, due to Proposition 3. This completes the proof. □

Note that this proof is not intended to show generically that welfare is decreasing in the level of competition amongst LSEs. Competition amongst LSEs could provide many additional benefits to consumers through, for instance, better sorting into preferred product types and more competitive prices. However, the proof implies that due specifically to its effect on reducing the equilibrium forward contract level Q, which subsequently reduces total electricity production, increasing the number of LSEs serves to reduce average welfare.

3.6 Consumer Market Power: A Mitigating Effect

This section generalizes the results to the case when the price for forward contracts increases in the number of forward contracts procured. In this case, consumers have market power in the forward market because their consumption impacts the price of forward contracts. The results show that the presence of consumer market power reduces the total number of contracts purchased. This is especially true when the number of LSEs is small and producer quantity choices have a big effect on price.

When consumer purchases in the forward market impact price, then increasing the number of LSEs serves to increase the quantity of forward contracts by reducing the ability of individual producers to exercise buyer market power. This mitigates the effects of positive externalities described above, acting on the total quantity of forward contracts in the opposite direction.

When the impact of forward contract procurement on price is not too high, or if the forward contract price is very high, the effect of positive externalities dominates and welfare decreases in the number of producers. If forward contract quantities have a big effect on the forward contract price, then the market power effects dominate and welfare increases in the number of producers.

We define the derivative of the forward contract price with respect to the quantity procured $p'_F(Q) = 0$. Next, we generalize the results from Section 3.4 to account for the case where $\exists Q\mathbb{R}^+$ s.t. $p'_F(Q) > 0$. We assume that the price curve is weakly convex, i.e. $p''_F(Q) \geq 0$.

Proposition 5. The equilibrium level of forward contracting Q in the case of a generic demand curve is given by:

$$Q = \frac{(MN + M + N)\mathbb{E}[T] + (MN(N+1) - N)c - M(N+1)^2 p_F}{b(MN + M + N) + (N+1)^2 p'_F(Q)}. \tag{3.28}$$

The proof follows from (3.4) but with the additional term since $p'_F(Q)$ is possibly nonzero. Note that $p'_F(Q) > 0$ at optimal Q necessarily implies that the forward contract price sensitivity reduces total forward contracting levels, since it increases the (positive) value of the denominator. Specifically, let $\tilde{Q}$ refer to the equilibrium forward contracting level for M producers as described in Section 3.4. Then $Q < \tilde{Q}$. When consumers' forward contract consumption levels have a proportionally bigger impact on the price of forward contracts, the consumers reduce their total procurement of forward contracts in equilibrium.

Next, consider the effects on the number of LSEs M on the forward contract quantity and on welfare.

Proposition 6. The effect of the number of LSEs on the total equilibrium forward contract quantity $\frac{\partial \Pi}{\partial M}$ is ambiguous. For the forward contract quantity to increase in M, i.e. $\frac{\partial \Pi}{\partial M} > 0$, it is necessary that $p'_F(Q)$ is sufficiently large for some Q.

Proof. By differentiating (3.28), observe that

$$\begin{aligned} Q(b(N+1) + (N+1)^2 p''_F(Q)\frac{\partial Q}{\partial M}) + (b(MN + M + N) + (N+1)^2 p'_F(Q))\frac{\partial Q}{\partial M} \\ = (N+1)\mathbb{E}[T] + N(N+1)c - (N+1)^2 p_F(Q) - M(N+1)^2 p'_F(Q)\frac{\partial Q}{\partial M}. \end{aligned} \tag{3.29}$$

Rearranging the above shows that

$$\frac{\partial Q}{\partial M} = N + 1 \frac{\mathbb{E}[T] + Nc - (N+1)p_F(Q) - Qb}{b(MN + M + N) + (N+1)^2 p_F'(Q) + (N+1)^2 p_F''(Q) + M(N+1)^2 p_F'(Q)}. \tag{3.30}$$

The results in Section 3.5.1 still hold exactly in the case described in this Section, because they do not depend on the forward price or the equilibrium result for Q. Therefore, $\frac{\partial \Pi}{\partial M} > 0 \iff \frac{\partial Q}{\partial M} > 0$.

The denominator in (3.30) is positive because each term is positive. Therefore, $\frac{\partial \Pi}{\partial M} > 0$ iff the numerator is greater than zero. Taking into account the result in (3.28) for Q, the numerator is equivalent to

$$\frac{(N+1)^2 p_F'(Q)(\mathbb{E}[T] + Nc - (N+1)p_F(Q)) + bN(N+1)(c - p_F(Q))}{b(MN + M + N) + (N+1)p_F'(Q)}. \tag{3.31}$$

Again, the denominator is positive, so the sign is dependent on the numerator. By rearranging the numerator, observe that the requirement that the numerator is greater than zero is equivalent to

$$(N+1)p_F'(Q)(\mathbb{E}[T] + Nc - (N+1)p_F(Q)) > bN(p_F(Q) - c). \tag{3.32}$$

The sign of the left hand side is ambiguous. This implies two necessary conditions. First, $p_F(Q) < \frac{\mathbb{E}[T]+Nc}{N+1}$, i.e. the forward price is not too high. Second,

$$p_F'(Q) > \frac{bN(p_F(Q) - c)}{(N+1)(\mathbb{E}[T] + Nc - (N+1)p_F(Q))}. \tag{3.33}$$

Therefore, if the forward price is not too high, but its derivative with respect to Q is fairly high, then the output Q and welfare Π can be locally increasing in M. In the case where the price influence of individual purchasing decisions is high, the influence of increasing M has a larger impact on reducing buyer market power (and therefore increasing the quantity of forward contracts) than it does on increasing the effects of positive externalities (which decreases the quantity of forward contracts).

□

These results explain how the price-effects of forward contracting decisions can impact the equilibrium quantity of forward contracts purchased. Buyer market power introduces a mitigating effect versus previous results. Specifically, if buyer market power is sufficiently

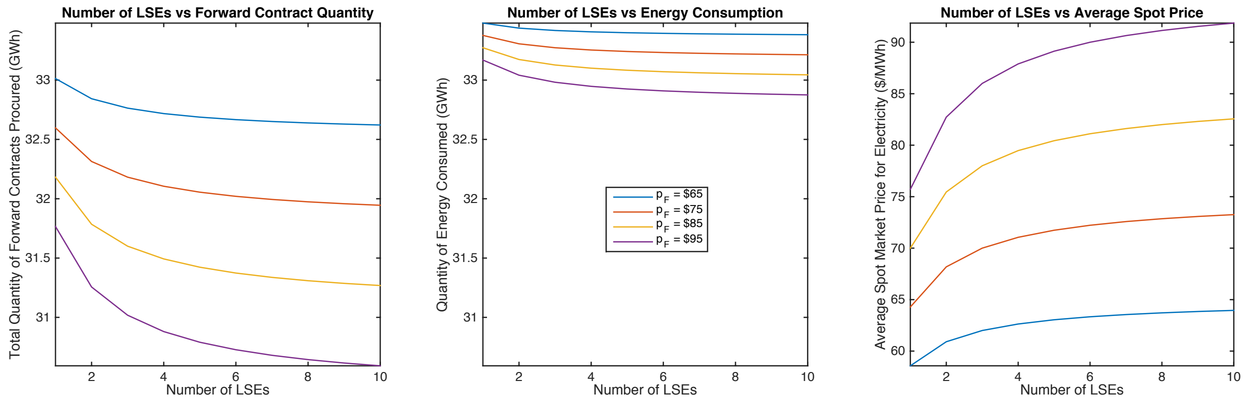

Figure 3-2: The effect on the number of LSEs, M, on forward contracting, energy consumption, and spot market prices, over a range of forward contract prices.

high, then increasing the number of LSEs may have a net positive effect on welfare because the equilibrium quantity effects due to a reduction of market power by purchasers of forward contracts outweighs the quantity effects due to the positive externalities of such contracts.

3.7 Case Study

This section presents a simple case study of the results. It helps illustrate the results of the analysis and the sensitivity of forward contracting and spot prices to the number of generators, the number of producers, and to the forward contract price and price slope. For the computations in this study, we use $b = 0.055$, based on the average of the low and high values for real-time price sensitivity found by Lijesen (2007). The parameter T was fixed at $T = 1900$; it was chosen so that net consumption at the reference case was approximately equal to average net consumption in an hour in the ERCOT system, 34 GWh. This average hourly consumption was calculated as 349 TWh, the total ERCOT demand in 2016, minus 54 TWh, the total ERCOT renewable energy production in 2016, divided by 8760 hours per year. For the reference case the marginal cost of energy production is 50 \$/MWh, the forward contracting price is 75 \$/MWh, there are $N = 3$ producers, and the number of LSEs varies from one to ten, but these are varied in the sensitivity analyses.

In Figures 3-2 and 3-3, the forward contract has a fixed price; the results are based on the analysis from Sections 3.2-3.5. In each of the examples, the quantity of forward contracts and the total amount of energy purchased are decreasing in the number of LSEs, as predicted by the analysis. As was shown previously, welfare also decreases in M because

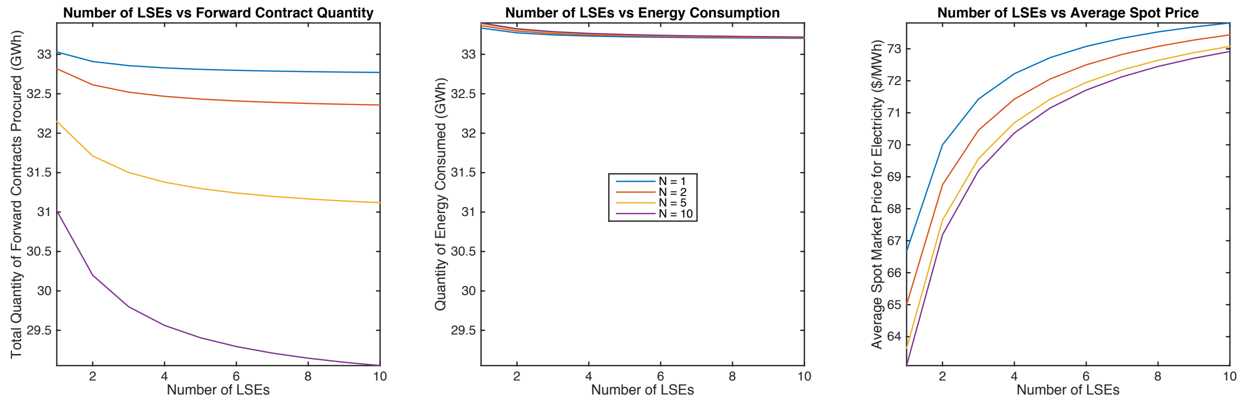

Figure 3-3: The effect of the number of LSEs on forward contracting, energy consumption, and spot market prices, as the number of producers, N, is varied.

its derivative with respect to M has the same sign as $\frac{\partial Q}{\partial M}$. Furthermore, the average spot price for electricity is increasing in the number of LSEs.

In Figure 3-2, the forward contract price varies from \$65 to \$95 / MWh. As the forward contracting price increases, the total amount of forward contracting decreases and the average spot price increases. The effect of the number of LSEs on total contract procurement is especially pronounced when the forward contract price is much higher than the marginal cost of electricity.

In Figure 3-3, the number of producers varies from $N = 1$ to $N = 10$. The number of producers has a major effect on the extent of forward contracting, because LSEs contract at much higher levels when the number of producers is low in order to mitigate the higher levels of producer market power. The effect of N on total production and average spot price is less pronounced. As expected, increasing the level of producer competition decreases the average spot price.

Finally, Figure 3-4 models the case where the price of forward contracts increases in the quantity of forward contracts procured, and varies the extent of this effect. This figure mirrors the analysis of Section 3.6. Increasing the supply curve slope also increases the average forward contract cost, which adds a second order effect to the changes depicted in the figure. To compensate for this, we adjusted the intercept of the forward price inverse supply curve such that the average forward contract price for each slope, taken over the range of the number of LSEs, is 124 \$ / MWh. As shown in Figure 3-4, when the slope of the forward contract inverse supply curve is sufficiently high, contract quantity, total production, and welfare no longer decrease in the number of LSEs. In these cases, the effects

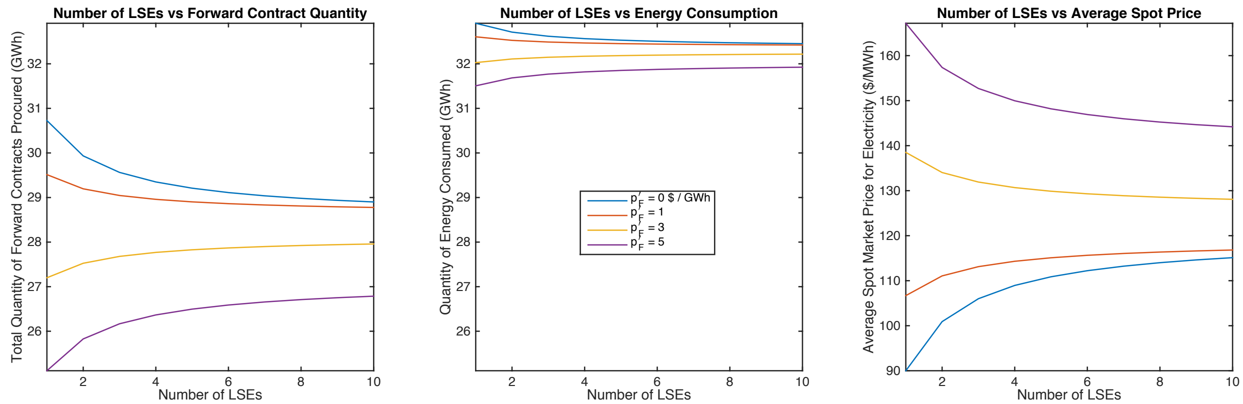

Figure 3-4: The effect on the number of LSEs on forward contracting, energy consumption, and spot market prices, over a range of slopes for the forward contract price supply curve.

of reducing buyer market power in the forward contracts market dominate, and increasing the number of LSEs reduces their ability to withhold forward contract purchases to keep the price low. Furthermore, when the inverse supply curve slope is sufficiently high, the average spot market price is also decreasing in the number of LSEs, as they procure a higher level of forward contracts.

3.8 Conclusion

Forward contracting can help reduce market power and supply withholding by producers. However, due to the positive externalities of LSE engagement in forward contracts, whereby the benefits of forward contracting and increased real-time supply are shared by all consumers, the level of forward contracting decreases in the number of load-serving entities. Therefore, in a competitive marketplace with many load-serving entities, each of whom pursue their own forward contracts, the level of forward contracts is below the level that maximizes social welfare. This implies that regulation may be required, for instance through mandated forward contracting or participation in capacity markets, in order to achieve the optimal level of forward contracting. This argument provides more compelling support for mandated forward contracting for electricity than the oft-repeated statement that forward contracting can help reduce market power, which itself does not imply a coordination problem nor compel regulation that is intended to increase the total level of forward contracting.

This work is based on a number of simplifying assumptions. Future work can extend this model for more generic inverse demand curves for real-time electricity and especially for forward contracts. It could also consider the risk preferences of the producers and LSEs in order to more accurately model the difference between forward and real-time prices and to examine the influence of risk-preferences on the features described here. Furthermore, future work could attempt to compare the effects discussed here to the risk-reduction benefits of forward contracting, or to estimate the level of the impact described in a real-world market.

Chapter 4

Learning Better Baselines for Electricity Demand Response

4.1 Introduction

Demand response programs create value for electricity systems by reducing or shifting electricity demand during specific time periods. The cost of delivering electricity varies substantially from hour to hour; in some systems, over 20% of energy costs are driven by just 2% of peak hours. As such, targeted demand reductions during particular time periods can provide substantial value. Demand response (DR) can improve power system reliability and reduce costs for network and capacity investments (Siano, 2014). Demand response can provide rapid response for ancillary services, it can help shift demand to lower priced hours for economic benefit, and it can help reduce peak demand to decrease the need for new planned investments (O'Connell et al., 2014; Lee et al., 2013).

Experts anticipate that the value of demand response will continue to grow as renewable energy penetration increases. Major renewable energy sources like wind and solar energy are intermittent and uncertain; these attributes can increase temporal volatility of the marginal cost of electricity and they can increase the value of flexible demand. As penetration of renewable resources grows, demand response could provide increasing benefits to mitigate the variability of renewable resources (O'Connell et al., 2014; Nolan and O'Malley, 2015).

Existing demand response markets are substantial. Shariatzadeh et al. (2015) describes the state of demand response opportunities in U.S. wholesale markets. In the U.S., about 23 GW of demand response were enrolled in U.S. wholesale markets as of 2017 (Surampudy et al., 2019). Participants in the Surampudy et al. (2019) Utility Survey reported over 20

GW of utility demand response; over 2.7 million customers are engaged in the types of DR programs that are the core focus of this work.

There are two overarching paradigms that define the market for demand response programs. In the first, demand response provides direct value to load-serving entities (LSEs), including regulated utilities or competitive retail electricity suppliers, that have a regulatory or contractual requirement to buy energy on behalf of end-use customers. In this context, demand response is used by load-serving entities like utilities and retail electric companies to reduce their costs associated with serving reliable electricity. The LSE can engage demand response by asking energy consumers to reduce their consumption at specific times. A DR program can help the LSE avoid high-priced energy purchases or transmission tariffs that stem from wholesale market participation, or it can mitigate or defer the need for investment in new distribution system components that would otherwise be required to serve growing demand. Throughout this chapter, we use the term LSE to refer to the organization that sells energy to end-use customers. This can be a regulated utility or a competitive retail electricity supplier. The described approach might align most naturally with the incentives of a competitive retail electricity supplier, who must honor an existing contract with an end-use customer; a retail electricity supplier in this position could directly reduce their costs by reducing energy consumption during high price periods. In our formulation, demand response is a type of principal-agent problem: the LSE is the principal and the end-use customer is the agent.

In an alternative market for demand response, aggregators offer demand response services directly into wholesale electricity markets, with aggregators serving as the middleman between wholesale market operators (ISOs / RTOs) and consumers with demand flexibility. The end-use customer in these programs is typically a large commercial or industrial consumer. In this work, we focus on the paradigm of an LSE that is creating a demand response program for residential and small commercial customers. The analysis could also help wholesale market-facing DR aggregators improve their contracting with end-use consumers.

Demand response programs are also differentiated by the attributes of customer engagement in the programs. The work in this chapter focuses on a subset of those programs that include baseline measurements to determine compensation for demand reductions. Demand response programs are often categorized as incentive-based programs or price-based programs (Deng et al., 2015). Incentive-based programs include direct load control, interruptible/curtailable load programs, demand bidding and buyback, and emergency demand reductions. Price-based programs include critical peak-pricing and real-time pricing. (Deng et al., 2015). Vardakas et al. (2015) provide an overview for the different types of demand response programs and optimization algorithms for implementing demand response.

This chapter focuses on a subset of demand response programs that share the following features: a customer-specific baseline is calculated, customers have voluntary load reductions, and utilities have no direct load control (e.g. of thermostats or electric vehicles).[1] These programs share a fundamental challenge: they seek to offer compensation for beneficial participation, but they do not penalize customers for high consumption. Furthermore, the LSE cannot perfectly calculate a customer's reduction because they do not precisely know the customer's counterfactual consumption. These programs are frequently called 'incentive-based demand response programs' or 'behavioral demand response.' The Utility Survey by Surampudy et al. (2019) shows that over 2.8 million customers in the United States are engaged in these demand response programs. Our framework also directly encompasses one type of price-based program, commonly called a peak-time rebate, where customers have a certain tariff for electricity consumption, but also can earn a rebate during peak hours when they reduce their demand below an established baseline. For simplicity in this research, we refer to the combination of incentive-based programs and peak-time rebate programs as *incentive-based demand response*. Hogan (2010) refers to the same class of programs as imputed demand response programs, because the baseline consumption must be imputed. The aforementioned programs are the focus of this research.

From the perspective of economic efficiency, incentive-based demand response is not first-best. In the power system, the efficient energy price at a given time is the short run marginal cost of delivering power to a particular location at that time, adjusted for losses, congestion, and the potential for scarcity (Rivier and Pérez-Arriaga, 1993) (Hogan, 2013). One solution is to directly charge customers this time-varying short-run marginal cost; this is often called the 'real-time price.' Joskow and Wolfram (2012) explain progress and challenges for dynamically pricing electricity.

Given consumer preferences, risk-aversion and transaction costs, certain customers will not prefer to pay the real-time price of electricity. Even if a real-time price is first-best, there are multiple challenges that limit its use, especially for residential and small commercial customers. From an economic perspective, transaction costs and costs of acquiring information

[1]Direct load control programs are also very popular in the academic literature and in practice; cumulatively, direct load control programs had over 6 million enrolled customers in 2018 (Surampudy et al., 2019). The programs that we focus on provide a valuable complement to direct load control programs for customers who do not have direct control capabilities or who do not wish to provide those capabilities to their LSE. The results in this chapter are aimed at incentive-based programs but could be useful for decision-making and contract-design in some direct load control programs.

might limit the benefits of a real-time price (Schneider and Sunstein, 2017). Regulators have not been keen to impose real-time prices on consumers by default. Less than 0.2% of U.S. residential customers currently pay a real-time price for electricity. Time-varying electricity prices can also increase electricity bills for certain customers or increase month-to-month bill volatility; these impacts can be especially harmful for low- and fixed-income customers (Burger et al., 2019, 2020).

When real-time dynamic prices are not an option, or when they are only a partial solution, incentive-based demand response is a potential second-best option for improving economic efficiency. Incentive-based demand response programs can provide the optimal marginal incentive for demand reductions without placing any price risk on consumers. When dynamic rates are employed, they tend to imperfectly track the marginal cost of delivered energy. Common time-varying tariffs, like time-of-use (TOU) prices or critical-peak prices (CPP), can be supplemented by a well-designed demand response program to improve economic efficiency. Due to economic and regulatory challenges associated with real-time prices, demand response programs are very common in practice. Based on a survey covering about 80% of the U.S. (Surampudy et al., 2019), approximately one in ten residential consumers participates in some form of demand response program. In 2018, utilities spent over $300 million in customer incentives for residential demand response programs, and approximately $235 million administering residential demand response programs (EIA, 2018).

Time-varying prices are most effective when they are provided as the default rate, but for various reasons they are typically offered on an opt-in basis. The default rates defines the electricity price structure that residential and small consumers will face if they make no active decisions about their electricity supply; if you have not changed your electricity supplier or opted-in to a time-varying price plan, you are most likely paying the default rate. Demand response programs are also typically offered on an opt-in basis, where consumers have the option to participate in the program but are not automatically enrolled. Incentive-based DR could potentially be offered on a default basis, because it does not require a digital connection for direct load control. Cappers et al. (2016) show, using a randomized control trial in Sacramento, that a default time-of-use rate decreases consumption during peak hours. This experience suggests that DR-by-default could also help reduce overall consumption during peak hours. However, as we will show, the high cost of current programs might limit the practicality of this approach. Our methodology helps to solve this problem by reducing incentive payments for low-performing DR customers.

The first contribution of this chapter is to describe customer response models and objective functions for utilities developing incentive-based demand response programs. The second contribution is to present an argument that underlies the overall approach of this chapter: reasonable objective functions for the LSE imply that the baseline threshold for demand response should not necessarily be an estimate of the counter-factual consumption. Then, we present two online learning algorithms that can help manage the sequential decision problem of choosing customer baselines in incentive-based DR programs. We provide numerical data to showcase their value, explaining how they achieve better performance than current practice.

Section 4.2 provides a review of the literature on demand response. Section 4.3 explains the basic model of customer demand and the demand response incentive. Section 4.4 investigates a customer model where customers do not alter their behavior based on the chosen baseline threshold; it explains the LSE's objective under this customer model. Section 4.4.2 develops an online learning approach, based on the Upper Confidence Bound (UCB) algorithm, that can be used to choose demand response baselines and customer participants. Section 4.5 provides an alternative model, where customers observe and respond directly to the baseline threshold; Section 4.5.2 provides an online learning approach to sequentially choose customer baselines under this model of customer behavior. Section 4.6 provides numerical examples, and Section 4.7 concludes.

4.2 Literature Review

Researchers and industry practitioners use the term 'demand response' to refer to a wide variety of programs that derive value by managing or shifting electricity demand during specific time periods. Vardakas et al. (2015) and Deng et al. (2015) provide an overview of demand response program categories. Section 4.1 explained that we focus on a subset of demand response programs, often called incentive-based demand response programs, that share two common features: (1) they offer users an incentive for reducing energy demand during specific time periods, and (2) they do not charge or financially penalize consumers for non-performance.

Incentive-based demand response programs for small consumers face a fundamental problem: they seek to provide a weakly-positive (≥ 0) incentive to participating customers, but they have imperfect information regarding customer demand and preferences. Retail electricity tariffs allow customers to buy any quantity of energy at the retail rate, so programs would face significant moral hazard if customers could buy any quantity of energy and

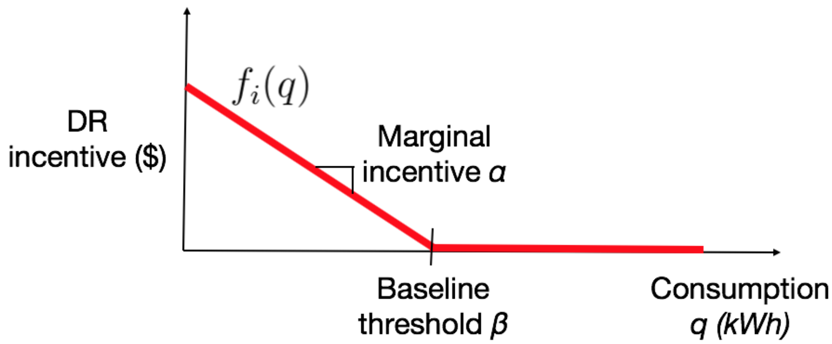

Figure 4-1: Example of demand response incentive function

offer it as demand response. The practical solution is generally to compute a customer-specific baseline that estimates what the customer would have consumed in the absence of the demand response offering. This baseline serves as a threshold in a piecewise demand response program; customers can be paid for demand response to the extent that their consumption is below the imputed baseline (see Figure 4-1). According to (Chao, 2010a), "the customer baseline is the estimated level of "normal" or counterfactual consumption during the time period against which demand reductions are measured and payments are determined... The customer baseline is conjectural (i.e., it is not directly observable and is generally estimated from data that represent past customer behavior using statistical estimation methods)." The estimation of customer baselines for determining compensation provides a workable model that is widely used in practice.

Baseline uncertainty is a core and ongoing challenge associated with incentive-based demand response programs; one strain of existing literature attempts to quantify the costs of existing uncertainty. Issues associated with imputed baselines have major impacts on the efficiency and value of demand response programs. Chao (2010a) highlights several problems with demand response using administratively-determined customer baselines: baseline manipulation, inefficient price formation, and generation relocation and load shifting behind the meter. Baseline manipulation can occur because of moral hazard or adverse selection. Customers might seek to impact their baseline estimates, or they could selectively participate when their baseline estimate is high, reducing the overall efficiency of the program. Additionally, uncertainty in baseline measurement can impact the overall evaluation of the benefits of a demand response program (O'Connell et al., 2014; Nolan and O'Malley, 2015). Addy et al. (2013) explain how the effect of baseline modeling choices can have a significant impact on the performance analysis of a demand response program. They detail the sensitivity of baseline estimates to various modeling choices. Chao (2010b) provides a simple economic example that highlights the inefficiency of demand response programs that pay the full locational marginal price (LMP) for reductions in demand.

There are two prominent streams of literature that seek to address the baseline problem. The first focuses on methods to improve baseline estimation. The second focuses on alternatives to existing demand response programs that avoid baseline estimation problems, which can be especially pernicious in the case of adversarial actors. Separately, an additional group of literature highlights problems associated with economic incentives for demand response, even in the absence of baseline uncertainty.[2] A fourth group of literature focuses on demand response with direct load-control, for instance of customer thermostats. Besides the academic interest, these programs are effective in practice and provide a valuable complement to the incentive-based DR programs described here; incentive-based DR is most useful in the absence of direct control capabilities. This literature review focuses primarily on the the first two groups of literature described in this paragraph, which offer different approaches to improving incentive-based DR.

Since baseline estimation is a core component of demand response programs, extensive efforts seek to estimate and reduce associated prediction errors. Standard industry practice is highly variable but often simple. The North American Energy Standards Board (NAESB) has defined five types of baseline methodologies, appropriate for different types of demand response (Rossetto, 2018). These include baseline type-1, which is essentially the estimated counterfactual consumption. Mohajeryami et al. (2016) explains several different methods for calculating customer baselines (CBLs) and argues that CBL calculations are more challenging for residential consumers than for industrial consumers. Park et al. (2015) describe a method to improve estimates of CBLs and display the benefits of their methodology in terms of a reduction of mean squared error. Nexant Inc. (2017) convened a Baseline Accuracy Working Group to investigate topics related to the use of alternative baseline methods and the use of control groups for estimating DR. Todd et al. (2019) estimate the extent to which spillover effects bias baseline measurements. Wijaya et al. (2014) estimate the effects of baseline estimation error and prediction bias on the profits of different stakeholders in DR programs.

Given the issues associated with imputed demand response, research has provided alternative methods for contracting for demand response. Chao (2011) describes and contrasts contractual and imputed approaches to determining CBLs. He suggests three contract types

[2] Early wholesale demand response programs often paid customers the full LMP for reductions in demand (Ruff, 2002), but this is inefficient in the case of imputed demand response programs because customers are essentially double-paid to reduce demand: they receive the LMP for demand reductions, and they also save at the retail rate by not consuming energy (Ruff, 2002) (Hogan, 2010).

that engage demand while avoiding problems with existing DR. Contracted demand response programs are a very promising option for engaging demand response. They avoid many of the inefficiencies associated with imputed demand response. However, the types of programs described in Chao (2011) relax a key constraint that characterizes most incentive-based DR programs: the programs in the paper by Chao (2011) can penalize or increase costs for some customers. We focus on demand response programs without additional contracts, where participating consumers are compensated for demand reductions but can never be charged a positive amount by the demand response program.

Additional literature provides alternative methods for managing demand response, but compared to the bulk of this literature I prefer to focus on more incremental approaches for improving existing demand response methods. Meir et al. (2017) describe a VCG mechanism for guaranteeing a minimum level of demand reduction from individual users when baselines are known. Zhong et al. (2013) research a method whereby customers are offered rebates to reduce their demand, and the offer is updated if an insufficient amount of demand responds. Khezeli and Bitar (2017) investigate a risk-sensitive demand response mechanism when the underlying demand curve for electricity is assumed to be affine and subject to unobservable random shocks. Li et al. (2015) consider demand response markets where consumers submit supply functions for reducing demand to meet available supply and compare competitive to oligopolistic equilibria. The aforementioned research provides helpful context and new directions for demand response, but it does not explicitly study the principal's baseline decision problem in incentive-based DR programs.

Mohsenian-Rad et al. (2010) study a group of users who schedule energy consumption to minimize system cost, and they employ a VCG-like mechanism to coordinate energy consumption. While this method could incentivize more honest baseline reporting, it could lead to new inefficiencies because consumers share the total costs associated with system consumption, rather than paying their own marginal costs of consumption. Muthirayan et al. (2019) provide a method for incentive-compatible baseline reporting. However, they allow the demand response program to levy a penalty on users; we focus on programs that are constrained to only offer (weakly) positive incentives to participating users. The method described by Muthirayan et al. (2019) also increases consumption in customers who are not called for DR service, diminishing the total demand reduction during peak periods. Dobakhshari and Gupta (2018) consider the joint problems of adverse selection and moral hazard in demand response incentive based programs. They consider a specific form for the inventive contract, which features a payment for the reported demand reduction and for profit sharing, and they find the optimal parameters for that specific contract structure.

We apply well-known results from online learning to the baseline decision problem. Other work utilizes online learning to study demand response, but only our work focuses on how online learning can help determine the baseline parameter for demand response contracting. Kalathil and Rajagopal (2015) consider a method of online learning for demand response, where the aggregator can learn about customer responsiveness over the course of their participation in the program, accounting for the fact that customer responses diminish if they are frequently asked to participate. Wang et al. (2014) utilize online learning to pick the "best" loads to deploy in each time-step; they do not focus on contracting or paying these loads. In Section 4.4.2, we are also interested in choosing a subset of loads to deploy, but our decision is motivated by economic incentives. Other research utilizes tools from online learning to understand electricity demand management, but it is not particularly applicable to research on incentive-based demand response programs (e.g. (Bahrami et al., 2017)).

Our research focuses on the challenge of contracting for demand response and of choosing optimal customer baselines. We use online methods to simultaneously learn the cost function for the baseline decision problem and to exploit existing information by choosing high-value baselines and requesting participation from high-value customers. To our knowledge, this is the first research effort that uses tools from online learning to explicitly help with baseline setting and contracting for incentive-based demand response.

4.3 Basic Model

In this section we present the basic notation regarding customer demand and the demand response incentive. Then, we motivate the work by explaining the current practice for the decision problem of choosing a customer baseline for the demand response incentive. Next, we present two distinct formulations for explaining how customers react to a demand response incentive. These formulations lead to the crux of the argument that underlies the work in this chapter: in each formulation, the cost function associated with the baseline decision variable β is asymmetric, and heavily influenced by each customer's propensity for demand reductions.

In the first formulation, a customer's random utility function determines their response to a demand response incentive; the demand response baseline or threshold does not immediately impact their consumption. However, the DR program operator (the LSE) ultimately bears some cost if they set the baseline threshold too low, for example due to customer dissatisfaction. In the second formulation, the chosen demand response baseline interacts directly with the customer's random utility function, which determines the customer's response in each demand response period.

In the subsequent sections, we investigate the demand response decision problem for each of these formulations. We formulate the demand response program as an online optimization problem where the LSE learns about its customers over time and iteratively chooses new demand response baselines.

4.3.1 Customer Model and Demand Response Incentive

Consider a customer i with a fixed utility function associated with the consumption of electricity at a particular time t, i.e. for i, t, $U_{it} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$. Customer i receives $U_{it}(q)$ units of utility from consuming q units of electricity in time period t. There is a fixed retail cost $R > 0$ for purchasing electricity. Therefore, the objective function for customer i in period t is $U_{it}(q) - qR$. Assume in the base-case that the customer is welfare maximizing, with full knowledge of the retail rate R. Furthermore, assume that U is twice-differentiable, with $U'(q) > 0$, $U''(q) \leq 0$ for any $q \in \mathbb{R}^+$. Let $Q_{it}(p) = (\frac{\partial U_{it}}{\partial q})^{-1}(p)$ be the demand function, where $p \in \mathbb{R}^+$ is the price per unit of electricity. In the base-case, customer i consumes $d_{it} = Q_{it}(R)$ units of electricity in time period t.

Throughout this chapter we consider a simple approach to demand response (DR), where demand response is offered as a non-negative incentive alongside an existing electricity rate. In our model, as is typical in practice, the DR incentive function $f_{it} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ determines the payment to customers who are participating in the demand response program, as a function of q. The demand response incentive $f_{it}(q)$ is non-negative because the demand response program can only provide an incentive, not a penalty. This is a typical feature of demand response programs. In particular, we consider a piecewise linear demand response incentive

$$f_{it}(q) = \alpha_t(\beta_{it} - q)^+ \tag{4.1}$$

Throughout this chapter, $(x)^+ = \max\{x, 0\}$. In the demand response program, α_t is the payment per unit of reductions in period t, and β_{it} is the customer-specific, "baseline", or the threshold below which customer i is paid for demand response reductions. Generally, α_t is chosen at the outset of the program, or chosen in each period from a limited set of pricing options. We focus mainly on the baseline β_{it}, generally treating α_t as constant or fixed in advance.

From the perspective of the LSE or demand-response program operator, customer i's utility function is unknown. The convention in practice is to choose the baseline β_{it} as an estimate of the counterfactual consumption $\mathbb{E}[d_{it}]$, where the principal has imperfect information regarding customer i's utility function and (therefore) counterfactual consumption in period t. Under this paradigm, the problem of choosing β_{it} is an estimation problem; $\mathbb{E}[d_{it}]$ is the estimand. In practice, demand response programs will use simple heuristics to create estimators for $\mathbb{E}[d_{it}]$, e.g. consumption during four of the last five business days, at the same hour of day as the current demand response period.

In reality, the cost of baseline errors—the cost when the baseline threshold is not equal to the counter-factual consumption—may not be symmetric, which implies that current practice may be suboptimal. A reasonable cost function of errors, e.g. of the form $g(\beta_{it} - d_{it})$, may not be symmetric, and it may be heavily influenced by other features of the customers' utility functions, which can be learned over time. Goldberg and Agnew (2013) and Rossetto (2018) describe the potential costs of biased baseline estimates in terms of customer satisfaction. Private interviews with utilities suggest that utilities recognize customer satisfaction benefits from upwardly biasing baselines; this could be achieved, for instance, by creating a biased estimator s.t. $\mathbb{E}[\beta_{it} - d_{it}] > 0$. There are no current efforts that use first principles or customer data to determine the appropriate size of that bias, which may differ across customers. We attempt a more thorough treatment that explicitly recognizes the trade-offs implied by asymmetric costs and how they might vary across different customers. We argue that the choice of β_{it} is ultimately a decision problem that seeks to maximize welfare or to achieve some combination of objectives for the principal and the agents.

In the following sub-sections, we present two different formulations for considering how customers respond to demand response incentives $f_{it}(q)$. These different formulations might be appropriate for different types of customers or customer-device pairs; an individual receiving a text message about demand response will respond differently than an electric vehicle charger that might receive a digital signal about a demand response event with full information regarding the demand response incentive. Despite their differences, the formulations provide complementary arguments that in general utilities should find β_{it} that optimizes some objective function, not simply choose β_{it} as an estimate of $\mathbb{E}[d_{it}]$, and they provide new intuition for optimizing β_{it} to improve demand response programs.

Section 4.4 covers the case where customer demand response is not explicitly impacted by the demand response baseline β_{it} that is chosen by the principal. This mirrors the typical assumptions in practice, where β_{it} is selected as an estimate of counter-factual consumption.

Section 4.5 covers the case where the specific baseline threshold may β_{it} influence the level of customer demand response; this follows naturally from the utility-maximization framework, but it assumes that the customer has access to information about β_{it} in real-time.

4.4 Online Learning for Demand Response

This section shows how an LSE can learn about the optimal customer baseline in a setting where the baseline does not directly impact customer participation. Section 4.4.1 presents the customer model. Section 4.4.2 presents a natural objective function for an LSE operating under this model of customer behavior. We present an online learning model for iterative decision making by the LSE; the baseline is chosen in each period based on all available information, but the LSE faces exploration trade-offs in deciding what customers to include in the demand response program.

4.4.1 Baseline Agnostic Response Model

Customers in residential-style demand response programs are unlikely to track the particular details of their baseline incentive. This formulation assumes that customers respond directly to α_t, the maximum marginal benefit of reductions under the baseline incentive, without directly accounting for β_{it}. Therefore, in the 'baseline agnostic' response framework, customer demand is given by:

$$q_{it} = \begin{cases} Q_{it}(R + \alpha_t) \text{ if } p_{it} = 1 \\ Q_{it}(R) \text{ otherwise} \end{cases} \tag{4.2}$$

The variable p_{it} is an indicator that equals one when the customer is asked to participate in a demand response program in period t. In the base case, the customer responds to the marginal price R. In demand response periods, the customer responds to the marginal price $R + \alpha_t$. In a specific period, customer i's response $r_{it} = Q_{it}(R) - Q_{it}(R + \alpha_t)$. With that notation, we have an alternative form for (4.2)

$$q_{it} = \begin{cases} d_{it} - r_{it} \text{ if } p_{it} = 1 \\ d_{it} \text{ otherwise} \end{cases} \tag{4.3}$$

Generally, $\alpha_t > 0$, which, along with the monotonicity of $Q_{it}(p)$, implies that $r_{it} > 0$. Under this framework, when the principal sets β_{it} too low, there is no indirect cost in terms of reduced reductions. Therefore, in Section 4.4.2, we incorporate these costs directly into the principal's objective function. Our idea is that customers respond to demand response programs in the short-run even if they are not paid for their participation (because they don't have knowledge of the fact that β is too low), but that low β will ultimately drive costs for the LSE in the form of dissatisfaction or customer churn. In Section 4.4.2, we examine a candidate objective function for an LSE operating under this customer response model and present an algorithm to sequentially set β_{it} over time.

4.4.2 LSE Objective and Online Learning Procedure

Consider a particular customer engaged in a demand response program that is offered by the LSE. For convenience in this section, we begin by focusing on a single customer i, whose response is defined according to Section 4.4.1, and drop the subscript i from the notation. In a particular hour $t \in \{1, 2, ...T\}$ the customer has a baseline demand d_t, where $d_1, d_2, ..., d_T$ are iid and distributed according to a known probability distribution represented by the cumulative distribution function (CDF) $F(x) = \mathbb{P}(d \leq x)$[3]. Furthermore, the customer has a potential demand response r_t, where $r_1, r_2, ..., r_T$ are iid with unknown probability distribution. The observed demand is $q_t = d_t - p_t r_t$, where $p_t \in \{0, 1\}$ is an indicator variable indicating whether the customer has been asked to participate in the demand response event in a particular hour. For simplicity, fix $\alpha_t = \alpha$ for every period t.

To enhance intuition, imagine that these demand parameters d_t and r_t are generated by an underlying utility function that is randomly selected from a class of suitable utility functions; the specific utility function in a particular period is unknown to the principal. The underlying uncertainty in the customer utility function in each period t generates the randomness in the demand parameters d_t and r_t; d_t and r_t are not necessarily independent.

We consider the problem of choosing customer participation p_t and customer baseline β_t in each period t in order to maximize the expected value of a stochastic objective function. For simplicity, assume that the value of demand response is equal in every period under consideration; in practice, it would vary from period to period; this is a natural extension of these results.

We employ a simple, but generally non-convex objective function $g(\beta, d, r)$ that describes the costs associated with choosing some baseline β when faced with customer demand

[3]In practice, customer demand would vary over hours of the day and months of the year. For example, demand will be higher for many customers, on average, during the afternoon hours than the middle of the night. We assume that these effects are negligible, or that they have already been accounted for. For example, d_t could be residuals after performing a multi-variate time series forecast

parameters d and r.[4] The function $g(\beta, d, r) : \mathbb{R}^3_+ \to \mathbb{R}$ combines the following objectives:

1. The benefit from the demand response reduction r.
2. The cost associated with under-compensating customers for demand response reductions, i.e. based on the difference $r - (\beta - q)^+$.
3. The cost associated with overpaying customers, i.e. based on $(\beta - d)^+$.

The following is an example of one such objective function with $\mu, \lambda > 0$:

$$g(\beta, d, r) = \lambda r - \mu(r - (\beta - (d - r))^+) - (\mu + \alpha)(\beta - d)^+. \tag{4.4}$$

The first term λr is the benefit associated with the demand response action. The second term $\mu(r - (\beta - (d - r))^+)$ is the dissatisfaction cost associated with underpaying the customer. Once the baseline threshold is fixed, the LSE pays customers based on the calculated reduction $(\beta - (d - r))^+$; if the customer reduces r units, then the customer incurs a cost based on the difference between their actual reduction r and the calculation reduction; this cost, in our formulation, is passed on to the LSE, for instance due to customer dissatisfaction. The final term is the cost of overpaying a customer, i.e. from offering more than their actual baseline consumption d. The leading parameter is $\mu + \alpha$ so that $\frac{\partial g}{\partial \beta}|_{\beta > d} = \alpha$; if the demand response incentive is $\alpha(\beta - q)^+$, then if β is too high the marginal cost to the LSE is α. The parameters λ and μ measure the benefit of demand reductions and the cost of underpaying customers, respectively. We utilize the function (4.4) throughout this section.

We seek to minimize the expected value of regret, where average regret after T periods is defined as:

$$R_T = \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{d,r}[p^* g(\beta^*, d, r)] - p_t g(\beta_t, d_t, r_t). \tag{4.5}$$

At each period t, we choose β_t and p_t with knowledge of past consumption d_τ and r_τ for $\tau \in \{1, ...t - 1\}$. The optimal values of β^* and p^* represent the optimal responses if the probability law for random variables d and r is known:

$$\beta^*, p^* = \arg\max_{\beta, p} \mathbb{E}_{d,r}[p g(\beta, d, r)]. \tag{4.6}$$

[4] Non-convexity implies that it could be challenging to minimize $g(\beta, d, r)$, or $\mathbb{E}[g(\beta, d, r)]$. However, the minimization problem is one-dimensional, and in practice β belongs to a discrete set for some level of decimal precision, so here we ignore the additional challenges posed by non-convexity. In practice, $\mathbb{E}[g(\beta, d, r)]$ will often be quasi-convex, further simplifying the optimization problem.

From the perspective of a principal trying to minimize regret, a priori we only know the probability law for the random baseline d, and not the random response r.[5] We need to deploy $p_t = 1$ a sufficient number of times to learn about the distribution of r and to decide how best to set the baseline parameter β. Over time, we gain confidence regarding the value of that customer; if the maximum value looks to be less than 0, we can stop asking the customer to participate by more frequently setting $p_t = 0$.[6]

For any β, let $f(\beta) = \mathbb{E}_{d,r}[g(\beta, d, r)]$. At time t, consider a sequence of observations of $q = d - r$ in periods where $p_t = 1$, i.e. $\{d_1, ...d_{n_t}\}$ and $\{q_1, ...q_{n_t}\}$, where $n_t = \sum_{j=1}^{t} p_t$. Furthermore, let t_n be the minimum t s.t. $n_t \geq n$; i.e t_n is the time period corresponding to the n$^{\text{th}}$ observation. Note that $t_{n_t} = t$. For any β, let

$$\overline{f_t}(\beta) = \frac{1}{t} \sum_{\tau=1}^{\tau=t} \lambda(\mathbb{E}d - q_\tau) - \mu(\mathbb{E}d - q_\tau - (\beta - q_\tau)^+) - (\mu + \alpha)\mathbb{E}(\beta - d)^+. \tag{4.7}$$

Note that $\mathbb{E}d$ and $\mathbb{E}(\beta - d)^+$ are known (for any β), since the distribution of d is assumed to be known. Furthermore, note that for any β, t, $\overline{f_t}(\beta)$ depends only on the available data $\{q_i\}_{i=1}^{i=t}$ (i.e. it doesn't require knowledge of d_t or r_t). Finally, note that $\mathbb{E}_{d,r}[\overline{f_t}(\beta)] = f(\beta)$, by the linearity of expectation and the definition of q_t. Therefore, by using the data $\{q_\tau\}_{\tau=1}^{\tau=t}$, we can learn about $\mathbb{E}_{d,r}[\overline{f_t}(\beta)]$ towards optimizing $f(\beta)$.[7] The proposed algorithm follows: choose

$$\beta_t = \beta_t^* \equiv \arg\max_\beta \overline{f_t}(\beta) \qquad p_t = \begin{cases} 1 \text{ if } V_t + w\sqrt{\frac{\ln t}{n_t}} \geq 0 \\ 0 \text{ otherwise} \end{cases} \tag{4.8}$$

where $V_t = \max_\beta \overline{f_t}(\beta)$ so $\overline{f_t}(\beta_t^*) = V_t$. Similarly, let β^* be some maximizer of $f(\beta)$. Let $w \in \mathbb{R}$ be the width of the bounded interval $[a, b]$ that strictly bounds $f(\beta)$, for any β. The bounds here are guaranteed because a customer has a maximum level of power demand, determined by their connection to the grid; this maximum power demand implies a bound on the set of feasible estimates for β and the support of d and r. Collectively, this implies that $g(\beta, d, r)$ is bounded. Since $\overline{f_t}(\beta)$ is the average of n_t samples drawn from a distribution

[5]Demand response periods are infrequent, so we have an order of magnitude more samples of d than r. This suggests that we can have a fairly accurate probability representation of d in comparison to r.

[6]What does it mean for a customer to have negative value? Typically, it means that the LSE spends more money incentivizing reductions than they earn from the customer's demand response participation; this could happen, for instance, if $\mathbb{E}[r]$ is very low and/or if the variance of d is very high.

[7]Note that it would be comparatively difficult to directly learn the distribution or density of r; this is a distribution or density deconvolution problem (see Gaffey (1959); Carroll and Hall (1988); Fan (1991); Meister (2009); Dattner et al. (2011)).

with mean $f(\beta)$, then for all β, by the Chernoff-Hoeffding bound, $\forall \delta \geq 0$

$$\mathbb{P}\left(|\overline{f}_t(\beta) - f(\beta)| \geq \delta\right) \leq 2\exp(-\frac{2n_t\delta^2}{w^2}). \tag{4.9}$$

Now, let $\delta = w\sqrt{\frac{\ln t}{n_t}}$. Then, the Chernoff-Hoeffding bound (4.9) tells us that for each value of β and each time t:

$$\mathbb{P}\left(|\overline{f}_t(\beta) - f(\beta)| < w\sqrt{\frac{\ln t}{n_t}}\right) \geq 1 - \frac{2}{t^2}. \tag{4.10}$$

By applying the Hoeffding bound (4.10) at β^*,

$$\mathbb{P}\left(\overline{f}_t(\beta^*) + w\sqrt{\frac{\ln t}{n_t}} > f(\beta^*)\right) \geq 1 - \frac{2}{t^2}. \tag{4.11}$$

This following theorem shows that the algorithm (4.8) guarantees that the expected value of the average regret (4.5) converges to 0. This result follows from the literature on Upper-Confidence Bound algorithms in multi-armed bandit problems (Auer et al., 2002; Agrawal, 2018).

Theorem 1. If β_t and p_t are selected according to (4.8), the expected value of the average regret converges to 0, i.e. $\lim_{T\to\infty} \mathbb{E}[R_T] = 0$.

Proof. We separately consider the case of a customer with positive value $f(\beta^*) > 0$ and with (weakly) negative value $f(\beta^*) \leq 0$. First, consider a customer with positive value, $f(\beta^*) > 0$. For this customer, the expected value of regret is

$$\mathbb{E}\left[TR_T\right] = \mathbb{E}\left[\sum_{t=1}^{T}(1-p_t)f(\beta^*) + \sum_{t=1}^{T}p_t(f(\beta^*) - f(\beta_t^*))\right]. \tag{4.12}$$

First we calculate the regret from not selecting a customer that has positive value $f(\beta^*) > 0$. By the definition of β_t^*, $\overline{f}_t(\beta_t^*) \geq \overline{f}_t(\beta^*)$. Therefore, considering (4.11) and the definition of p_t in (4.8), this implies that for customers with $f(\beta^*) > 0$, $\mathbb{P}(p_t = 1) \geq 1 - \frac{2}{t^2}$. Expanding the first term in (4.12):

$$\mathbb{E}\left[\sum_{t=1}^{T}(1-p_t)f(\beta^*)\right] = \sum_{t=1}^{T}\mathbb{P}(p_t = 0)f(\beta^*) \leq \sum_{t=1}^{T}\frac{2}{t^2}f(\beta^*) \leq \frac{\pi^2}{3}f(\beta^*). \tag{4.13}$$

Next, we investigate the regret from choosing the wrong customer baseline, the second term in (4.12). By applying (4.9) with $\delta = w\sqrt{\frac{\ln n_t}{n_t}}$,

$$\mathbb{P}\left(|\overline{f}_t(\beta) - f(\beta)| < w\sqrt{\frac{\ln n_t}{n_t}}\right) \geq 1 - \frac{2}{n_t^2}. \tag{4.14}$$

This holds true at any β. Therefore, define events A and B as

$$A = \{\overline{f}_t(\beta^*) + w\sqrt{\frac{\ln n_t}{n_t}} > f(\beta^*)\} \qquad B = \{f(\beta_t^*) + w\sqrt{\frac{\ln n_t}{n_t}} > \overline{f}_t(\beta_t^*)\}. \tag{4.15}$$

From (4.14), we know that $\mathbb{P}(A) \geq 1 - \frac{2}{n_t^2}$ and $\mathbb{P}(B) \geq 1 - \frac{2}{n_t^2}$. Now $\mathbb{P}(A \cap B) = 1 - \mathbb{P}(\overline{A} \cup \overline{B}) \geq 1 - \mathbb{P}(\overline{A}) - \mathbb{P}(\overline{B}) \geq 1 - \frac{4}{n_t^2}$ by the nature of complements and the union bound. Furthermore, by the definition of β_t^*, $\overline{f}_t(\beta_t^*) \geq \overline{f}_t(\beta^*)$. Therefore, with probability at least $1 - \frac{4}{n_t^2}$,

$$f(\beta^*) - f(\beta_t^*) < 2w\sqrt{\frac{\ln n_t}{n_t}}. \tag{4.16}$$

Now, investigating the upper-bound on the right-hand term in (4.12), and incorporating (4.16),

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{T} p_t(f(\beta^*) - f(\beta_t^*))\right] &= \mathbb{E}\left[\sum_{n=1}^{n_T}(f(\beta^*) - f(\beta_{t_n}^*))\right] \\
&\leq \mathbb{E}\left[\sum_{n=1}^{n_T} \frac{4}{n^2}w + 2w\sqrt{\frac{\ln n}{n}}\right] \\
&\leq \frac{2w\pi^2}{3} + \mathbb{E}\left[\sum_{n=1}^{n_T} 2w\sqrt{\frac{\ln n_T}{n}}\right] \\
&\leq \frac{2w\pi^2}{3} + \mathbb{E}\left[\int_{n=0}^{n_T} 2w\sqrt{\frac{\ln n_T}{n}}dn\right] \\
&= \frac{2w\pi^2}{3} + 2w\mathbb{E}\left[2\sqrt{n_T \ln n_T}\right] \leq \frac{2w\pi^2}{3} + 4w\sqrt{T\ln T}.
\end{aligned} \tag{4.17}$$

The first equality rewrites the expression in terms of periods where $p_t = 1$. The first inequality expands the initial expression by noting that it is upper-bounded by the case where $f(\beta^*) - f(\beta_t^*) = 2w\sqrt{\frac{\ln n_t}{n_t}}$ with probability $1 - \frac{4}{n_t^2}$ and the case where $f(\beta^*) - f(\beta_t^*) \leq \sup f(\beta^*) - f(\beta_t^*) = w$ with probability $\frac{4}{n_t^2}$. The second inequality upper-bounds the first summation and uses the inequality $\ln(n) \leq \ln(n_T)$ for $n \leq n_T$ to simplify the right-hand term. The fourth line upper-bounds the summation by a Riemann integral, since $2w\sqrt{\frac{\ln n_T}{n}}$

is decreasing in n. The fifth line evaluates that integral. Note that the expectation is carried through because n_t is a random variable. The final inequality upper-bounds the expected value because with probability 1, $n_T \leq T$.

Therefore, adding the final terms in (4.13) and (4.17), $\mathbb{E}[TR_T] = O(\sqrt{T \ln T})$, for customers with $f(\beta^*) > 0$.

Next, consider a customer without positive value, $f(\beta^*) \leq 0$.[8] For this customer, the expected value of regret is

$$\mathbb{E}\left[TR_T\right] = -\mathbb{E}\left[\sum_{t=1}^{T} p_t f(\beta_t^*)\right]. \tag{4.18}$$

From (4.10), with probability at least $1 - \frac{2}{t^2}$,

$$f_t(\beta_t^*) < f(\beta_t^*) + w\sqrt{\frac{\ln t}{n_t}}. \tag{4.19}$$

When $n_t \geq \frac{4w^2 \ln t}{f(\beta^*)^2}$, this implies that with probability at least $1 - \frac{2}{t^2}$

$$f_t(\beta_t^*) < f(\beta_t^*) + \frac{|f(\beta^*)|}{2}. \tag{4.20}$$

The inequality follows from the lower bound on n_t and the fact that $w\sqrt{\frac{\ln t}{n_t}}$ is decreasing in n_t. Therefore, when $n_t \geq \frac{4w^2 \ln t}{f(\beta^*)^2}$, with probability at least $1 - \frac{2}{t^2}$,

$$\begin{aligned} f_t(\beta_t^*) + w\sqrt{\frac{\ln t}{n_t}} &< f(\beta_t^*) + \frac{|f(\beta^*)|}{2} + w\sqrt{\frac{\ln t}{n_t}} \\ &\leq f(\beta_t^*) + |f(\beta^*)| \\ &\leq f(\beta^*) + |f(\beta^*)| = 0. \end{aligned} \tag{4.21}$$

The first inequality follows directly from (4.20), and the second inequality uses the same mechanism as (4.20) to bound $w\sqrt{\frac{\ln t}{n_t}}$, given the lower bound on n_t. The third inequality comes from the definition of β^* as the maximizer of $f(x)$. The final equality is because $f(\beta^*) \leq 0$. Altogether, (4.21) implies that

$$\mathbb{P}\left(p_{t+1} = 1 | n_t \geq \frac{4w^2 \ln t}{f(\beta^*)^2}\right) \leq \frac{2}{t^2}. \tag{4.22}$$

[8]Imagine a customer with very low average response $\mathbb{E}[r]$ or with significant variability in baseline consumption d.

Therefore,

$$
\begin{aligned}
\mathbb{E}[n_T] &= 1 + \mathbb{E}\left[\sum_{t=1}^{T} \mathbb{1}(p_{t+1} = 1)\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \mathbb{1}\left(p_{t+1} = 1, n_t < \frac{4w^2 \ln t}{f(\beta^*)^2}\right) + \sum_{t=1}^{T} \mathbb{1}\left(p_{t+1} = 1, n_t \geq \frac{4w^2 \ln t}{f(\beta^*)^2}\right)\right] \\
&\leq \frac{4w^2 \ln T}{f(\beta^*)^2} + \sum_{t=1}^{T} \mathbb{P}\left(p_{t+1} = 1, n_t \geq \frac{4w^2 \ln t}{f(\beta^*)^2}\right) \qquad (4.23) \\
&= \frac{4w^2 \ln T}{f(\beta^*)^2} + \sum_{t=1}^{T} \mathbb{P}\left(p_{t+1} = 1 | n_t \geq \frac{4w^2 \ln t}{f(\beta^*)^2}\right) \mathbb{P}\left(n_t \geq \frac{4w^2 \ln t}{f(\beta^*)^2}\right) \\
&\leq \frac{4w^2 \ln T}{f(\beta^*)^2} + \sum_{t=1}^{T} \frac{2}{t^2} = \frac{4w^2 \ln T}{f(\beta^*)^2} + \frac{\pi^2}{3}.
\end{aligned}
$$

The first equality uses the law of total probability. The first inequality utilizes the upper-bound on n_t in the first term of the sum. The second equality rewrites the joint probability using the conditional probability chain rule. The second inequality utilizes (4.22) and $\mathbb{P}(A) \leq 1$ for any event A. The final equality evaluates the summation. Now the expected value of regret

$$
\mathbb{E}\left[TR_T\right] = \mathbb{E}\left[-\sum_{t=1}^{T} p_t f(\beta_t^*)\right] \leq \mathbb{E}\left[\sum_{t=1}^{T} p_t w\right] \leq w\mathbb{E}n_T \leq \frac{4w^3 \ln T}{f(\beta^*)^2} + \frac{w\pi^2}{3} = O(\ln T). \quad (4.24)
$$

Therefore, for any customer, with $f(\beta^*) > 0$ or $f(\beta^*) \leq 0$, $\mathbb{E}[TR_T] = O(\sqrt{T \ln T})$. The expected value of average regret converges to 0: $\lim_{T\to\infty} \mathbb{E}[R_T] = 0$.

□

Note that one outcome of this learning procedure is often that $\beta_t < \mathbb{E}d$, especially for small t. We could tweak the algorithm to bias β_t upwards when t is small, to ensure that we don't underpay customers when he haven't yet seen significant evidence of $\mathbb{E}r > 0$.

One potential concern of this algorithmic approach is that the regret bounds focus on the growth of regret; regret could be very high (due to a large leading constant) for small T. In Section 4.6 we provide numerical examples of the performance of the algorithm. The algorithm performs better than the existing approach even for very small T (i.e. $T = 10$). This can be explained in part by the fact that the algorithm uses p_t to explore customers, but it doesn't need to vary β_t to explore rewards because the available information provides a noisy version of the entire cost function (i.e. for any β) in each round.

One shortcoming of this approach is that it ignores the possibility that r_t is impacted in real-time by our choices of $\beta_1, \beta_2, ..., \beta_t$. Section 4.5 develops a customer model with those characteristics and investigates online learning for demand response baselines in that setting.

4.5 Online Learning for DR with Responsive Customers

The previous model assumed that the demand response baseline threshold for customer i in period t, β_{it}, did not directly impact a customer's level of demand response r_t. This section considers a utility-maximizing customer who is faced with a demand response incentive of the form described in (4.1). Under this model, with full information, it is clear that β_{it} can impact q_{it}. Section 4.5.1 explains the model, and Section 4.5.2 describes an online learning algorithm that can be leveraged to set baselines under the aforementioned model of customer behavior.

4.5.1 Full Information Response Model

In this framework, customers have full knowledge of the demand response incentive $f_{it}(q)$, and they only reduce their demand when β_{it} is sufficiently high. Given the demand response incentive, and perfect information on the part of the customer, then customer i's objective function in period t is

$$U_{it}(q) - qR + f_{it}(q). \tag{4.25}$$

The customer chooses $q_{it} = \arg\max_q U_{it}(q) - qR + f_{it}(q)$. As before, the "baseline" d_{it} is the customer's counterfactual consumption, i.e. what they would have consumed in the absence of the demand response program. The customer's maximization problem is non-convex, because it is the sum of $U_{it}(q)$ (concave) and $f_{it}(q)$ (convex). As before, let $f_{it}(q) = \alpha_t(\beta_{it} - q)^+$. The solution to the customer's non-convex maximization problem is given by:

$$q_{it} = \begin{cases} d_{it} - r_{it} & \text{if } U_{it}(d_{it} - r_{it}) + f_{it}(d_{it} - r_{it}) \geq U_{it}(d_{it}) + f_{it}(d_{it}) \\ d_{it} & \text{otherwise} \end{cases} \tag{4.26}$$

where, as before, $r_{it} = Q_{it}(R) - Q_{it}(R + \alpha_t)$. The value $d_{it} - r_{it}$ is the optimal response to the part of the incentive function where $f_{it}(q) > 0$, while the value d_{it} is the optimal response (i.e. no response) along the part of the incentive function where $f_{it}(q) = 0$. The discontinuity is due to the non-convexity of the customer's decision problem, which arises from the fact that $U_{it}(q)$ is concave but $f(q)$ is convex.

The customer response model (4.26) implies the existence of a specific threshold such that the customer only reduces their demand q_{it} if β_{it} is above the threshold. Let h_{it} be that threshold; from (4.26), $d_{it} - r_{it} < h_{it} < d_{it}$. With this notation,

$$q_{it}(\beta) = d_{it} - r_{it}\mathbb{1}_{\{h_{it} \leq \beta\}}. \tag{4.27}$$

Note that $\forall \beta' \geq \beta$, with probability 1, $q_{it}(\beta') \leq q_{it}(\beta)$.

4.5.2 Learning Baselines with Responsive Customers

In this section we consider the case where customers adjust their demand using full information of the demand response incentive and baseline. As explained in Section 4.5.1, under this paradigm customers reduce their demand if and only if the baseline threshold is sufficiently high. This section considers that type of customer behavior and develops online tools for setting customer baselines under this model of behavior. The online learning algorithms and bounds in this section come from (Kleinberg, 2005) and Auer et al. (1995).

This approach has several benefits: it is robust to non-stationarity of underlying demand parameters, it does not require a separate estimate of the counter-factual consumption, and it incurs sub-linear regret even when the customer is 'adversarial.' We consider the case where behavior in each round is determined by an 'adversary,' a customer who can choose their utility function in each period, not necessarily drawn from a stochastic model. Under this approach, the expected value of average regret converges to 0, where average regret compares actual performance to the optimal baseline that would be chosen in hindsight. Due to this 'adversarial' framework, this approach is suitable for non-stationary models of customer behavior.

The downside of this approach is that it learns slowly about customer behavior. It might not be practical for demand response programs with less than 50 events per year. However, typically one might expect that 'responsive' customers have digitally connected devices. Compared to an individual, these devices are more likely to account for the baseline parameter β in determining their response. The devices are also better suited to demand response programs with frequent events. This design might be suitable for programs with frequent demand response periods and digitally connected devices.

Under this paradigm, it is natural to assign some benefit for demand response reductions, which will be proportional to the actual size of the reduction. On the other hand, there is a cost associated with providing demand response payments to customers. The following function describes a natural objective function for the LSE:

$$g(\beta, d, r, h) = \lambda r \mathbb{1}_{h \leq \beta} - \alpha(\beta - q)^+ \tag{4.28}$$

where q is determined by d, r, h, and β, as described in (4.27). The LSE earns value proportional to λ for demand response reductions, and it bears costs due to the demand response transfer $\alpha(\beta - q)^+$.

Note that the LSE cannot directly observe $g(\beta, d_t, r_t, h_t)$ in any period t because they only observe q_t, not directly d_t, r_t, or h_t. Let β_M be the maximum baseline threshold for the customer under consideration. For each $t \in \{1, 2, ..., T\}$ we have a cost function $f_t : [0, \beta_M] \to S$, as follows:

$$f_t(\beta) = \lambda q_t + \alpha \left(\beta - q_t\right)^+ . \tag{4.29}$$

Note that $f_t(\beta) = \lambda d_t - g(\beta, d_t, r_t, h_t)$, and d_t is unaffected by β. Therefore, $\arg\min_\beta f_t(\beta) = \arg\max_\beta g(\beta, d_t, r_t, h_t)$ and $\arg\min_\beta \sum_{t=1}^{T} f_t(\beta) = \arg\max_\beta \sum_{t=1}^{T} g(\beta, d_t, r_t, h_t)$. Furthermore, note that the compactness of the domain $[0, \beta_M]$ and of the support of q_t implies that S is compact. Assume S is nonempty. Then $S_0 = \min S$ and $S_1 = \max S$ exist. Let $w = S_1 - S_0$.

Next, we present an algorithm that chooses $\beta_1, \beta_2, ..., \beta_T$ in order to bound the expected value of average regret, where average regret is

$$R_T = \frac{1}{T} \sum_{t=1}^{T} \left(g\left(\beta^*, d_t, r_t, h_t\right) - g\left(\beta_t, d_t, r_t, h_t\right)\right) = \frac{1}{T} \sum_{t=1}^{T} \left(f_t(\beta_t) - f_t(\beta^*)\right). \tag{4.30}$$

The optimal-in-hindsight parameter $\beta^* = \arg\max_\beta \sum_{t=1}^{T} g\left(\beta, d_t, r_t, h_t\right)$. Note that we are interested in the expected value of the average regret because the algorithm could make random choices of β_t. Our proposed algorithm follows from Kleinberg (2005):

The phrase MAB refers to the appropriate multi-armed bandit algorithm with a finite number of arms. An appropriate algorithm for MAB in this case is EXP3 from the paper by Auer et al. (1995). The function $\Gamma : S \to [0, 1]$ is a normalizing function for f_t, i.e. $\Gamma(x) = \frac{1}{w}(S_1 - x)$.

Theorem 2. Under Algorithm 1, $\mathbb{E}\left[T R_T\right] = O(T^{2/3} \log^{1/3} T)$. The expected value of average regret converges to 0, i.e. $\lim_{T\to\infty} \mathbb{E}[R_T] = 0$.

Algorithm 1 DR1

$n \leftarrow 1$
while $n \leq T$ **do**
  $K \leftarrow \left(\frac{n}{\log n}\right)^{1/3}$
  Initialize MAB with strategy set $\{\beta_M/K, 2\beta_M/K, ..., \beta_M\}$
  for $t = n, n+1, ..., \min(2n-1, T)$ **do**
    Get strategy β_t from MAB
    Play β_t and discover $f_t(\beta_t)$
    Feed $\Gamma(f_t(\beta_t))$ back to MAB
  end for
  $n \leftarrow 2n$
end while

Proof. The proof follows from Kleinberg (2005). In our case, the functions $f_1, ..., f_T$ are not uniformly locally Lipschitz due to the discontinuity of $f_t(\beta)$ near $\beta = h_t$ (due to the discontinuity of $q_t(\beta)$). However, as we will explain, this is addressed by the special form of f_t. The functions $f_1, ..., f_T$ also do not necessarily take values in $[0, 1]$, but they are bounded as described above. It suffices to show that the regret in the inner (for) loop of Algorithm 1 is $O(n^{2/3} \log^{1/3} n)$. This is because

$$\sum_{r=0}^{\lfloor \log_2 T \rfloor} (2^r)^{2/3} \log^{1/3}(2^r) \leq \sum_{r=0}^{\lfloor \log_2 T \rfloor} (2^r)^{2/3} \log^{1/3}(T) \leq \frac{T^{2/3} - 1}{2^{2/3} - 1} \log^{1/3}(T) \leq 2T^{2/3} \log^{1/3}(T). \tag{4.31}$$

The first expression sums the regret for each iteration of the outer loop, where $n = 2^r$ in the r^{th} iteration. The first inequality is because $2^r \leq T$. The second inequality computes the geometric sum with constant ratio, and the final inequality simplifies the expression. Therefore, if the inner loop is $O(n^{2/3} \log^{1/3} n)$ then the outer loop is $O(T^{2/3} \log^{1/3} T)$.

Next, consider a single iteration of the inner loop with n fixed and K defined as in Algorithm 1, i.e. $K = \left(\frac{n}{\log n}\right)^{1/3}$. Then, K points partition the domain $[0, \beta_M]$. Let β' be the smallest element in $\{\beta_M/K, 2\beta_M/K, ..., \beta_M\}$ s.t. $\beta' \geq \beta^*$. Then

$$\mathbb{E}\left[\sum_{t=1}^{n} f_t(\beta') - f_t(\beta^*)\right] \leq \alpha \frac{n\beta_M}{K} = O(n^{2/3} \log^{1/3}(n)). \tag{4.32}$$

This is due to the fact that for any t, $\beta' \geq \beta$, $q_t(\beta') \leq q_t(\beta)$ from the form of q_t as explained in (4.27). Therefore, $f_t(\beta') - f_t(\beta^*) \leq \alpha|\beta' - \beta^*| \leq \alpha \frac{\beta_M}{K}$. The equality is due to the form of K, described above.

Then, the next step is to show that $\mathbb{E}\left[\sum_{t=1}^{n} f_t(\beta_t) - f_t(\beta')\right] = O(n^{2/3}\log^{1/3}(n))$. Let β_n^* be the optimal arm in the inner loop that begins with $t = n$. Then

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^{n} f_t(\beta_t) - f_t(\beta')\right] &\leq \mathbb{E}\left[\sum_{t=1}^{n} f_t(\beta_t) - f_t(\beta_n^*)\right] \\ &= O(\sqrt{nK\log K}).\end{aligned} \tag{4.33}$$

The inequality says that the regret of β_t versus β' is no higher than the regret of β_t versus β_n^*. This follows from the definition of β_n^*. The equality follows from the results by Auer et al. (1995), where the MAB function is selected as the EXP1 algorithm from their paper. Compared to their result, cumulative regret is multiplied by the constant w here, but this constant does not affect the limiting behavior.

Due to our definition of K, $\mathbb{E}\left[\sum_{t=1}^{n} f_t(\beta_t) - f_t(\beta')\right] = O(n^{2/3}\log^{1/3}(n))$ and therefore $\mathbb{E}\left[\sum_{t=1}^{n} f_t(\beta_t) - f_t(\beta^*)\right] = O(n^{2/3}\log^{1/3}(n))$.

Including the previous argument about the sum of the geometric series of inner loops, the expected value of cumulative regret $\mathbb{E}\left[TR_T\right] = O(T^{2/3}\log^{1/3}(T))$. Therefore, the expected value of average regret converges to 0: $\lim_{T\to\infty} \mathbb{E}[R_T] = 0$. □

4.6 Case Studies and Examples

This section shows how the learning algorithms presented in Sections 4.4.2 and 4.5.2 can improve demand response program trade-offs on simulated customers.

First, consider a group of $N = 2000$ customers participating in multiple demand response periods. For all customers and time periods, $\forall i, t$, $d_{it} \sim N(10, 1.5^2)$. Customer response $r_{it} \sim B_i Z$, where $Z \sim N(2, 0.5^2)$, and $B_i \sim \text{Ber}(p_i)$ where for each i, p_i is one draw of the uniform random variable on $[0, 1]$.

Figure 4-2 highlights the fact that a decision-based framework can improve the terms of a trade-off between multiple outcomes of interest in a demand response program. The x-axis describes the average price paid per unit of demand-response reduction ($ / kWh), where α_t is normalized to 1 in all periods. The y-axis describes the average shortfall versus the actual DR for consumers; this explains the extent to which consumers are underpaid relative to their actual reductions in certain periods. There is a fundamental trade-off between minimizing each of these individual attributes. Even with a fairly limited number of demand response periods ($T = 10$ periods) we learn enough about the customer to improve versus the base case, where the baseline threshold is an estimate of the mean or a uniformly biased estimate

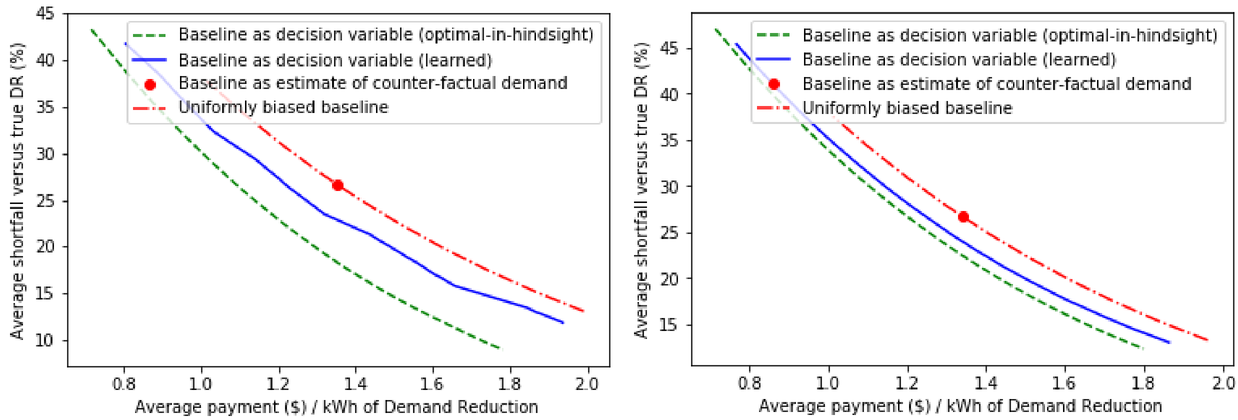

(a) 10 DR periods

(b) 50 DR periods

Figure 4-2: Trade-offs for demand response procurement.

of the mean. With perfect information, the optimal point that could be achieved is $(1, 0)$.

The dot near $(1.35, 26)$ in Figures 4-2a and 4-2b displays the outcome that would be achieved by setting $\beta_{it} = \mathbb{E}d_{it}$ for all customers in all time periods; this represents the base-case. The dot-dashed line shows how various trade-offs could be achieved by uniformly biasing the baseline estimate for all customers. The dashed line shows the optimal outcome that could be achieved in hindsight; the range of the line refers to multiple outcomes with λ and μ in different proportions ($\lambda = 1$ and μ ranges from 0.9 to 5). The solid line shows the outcome that can be achieved using the online learning framework in Section 4.4.2, with $T = 10$ periods in Figure 4-2a. As T grows, the solid line converges to the dashed line; see Figure 4-2b. Even with just 10 demand response periods to learn about customers, our approach drives a 10% improvement in the outcome variables of interest.

These benefits are even more pronounced in a demand response program where most customers do not respond to demand response signals. This example helps explain the case where demand response is offered by default, or where demand elasticities vary significantly from customer to customer. In this example, assume that all variables are generated the same way as in Figure 4-2 except for p_i. In this example, 10% of customers have a $p_i = 0.8$; these customers respond to a demand response event 80% of the time. Assume the other 90% of customers never respond to a DR event; for these customers $p_i = 0$. Figure 4-3 shows the trade-offs achieved by our algorithm in this type of setting and compares it to the base-case where the baseline threshold is calculated as an estimate of counterfactual consumption.

Next, consider a numerical example corresponding to Section 4.5. The online learning procedure in Section 4.5 directly tries to maximize demand response reductions and to min-

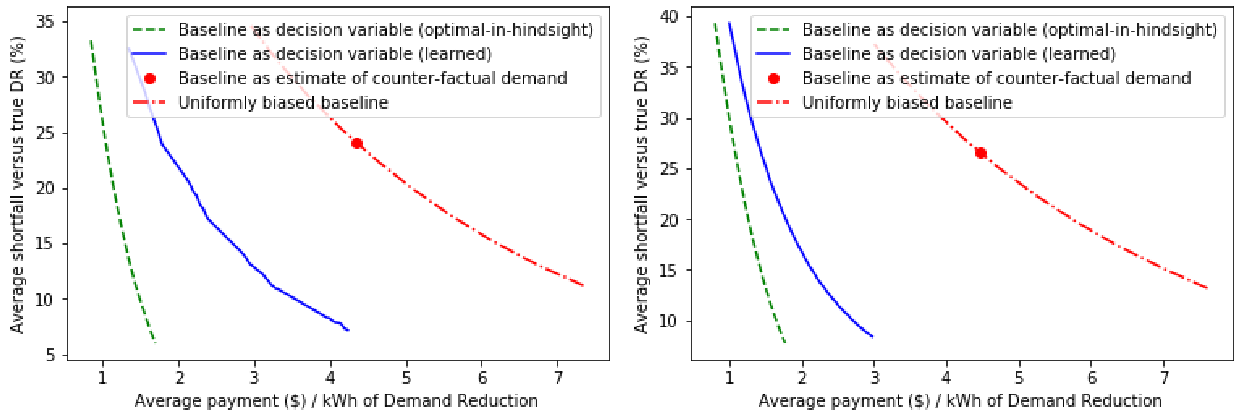

(a) 10 DR periods

(b) 50 DR periods

Figure 4-3: Trade-offs for demand response procurement, with low customer response rates.

imize costs; these objectives are not aligned, so there is a natural trade-off between each of the objectives.

Consider a group of $N = 1000$ customers participating in a demand response program. As before, for all customers, in all time periods, $\forall i, t$, $d_{it} \sim N(10, 1.5^2)$. Customer response $r_{it} \sim B_i Z$, where $Z \sim N(2, 0.5^2)$, and $B_i \sim \text{Ber}(p_i)$.

Figure 4-4 shows the average cumulative cost per unit of reduction, when α_t is normalized to 1 in all periods, as learning occurs in sequential DR periods. The x-axis is the total number of DR periods. The y-axis represents the average cumulative cost per unit of DR reduction. In a particular period, this value is the total price paid in the demand response program, divided by the total extent of demand response reductions (in kWh). Figure 4-4a represents the uniform case; for each i, p_i is one draw of the uniform random variable on $[0, 1]$. Figure 4-4b represents the low-participation case; $p_i = 0.8$ for 10% of customers, and $p_i = 0$ for the remaining customers. In both cases, the online-learning program quickly learns to reduce costs per unit of DR reduction, versus the base case where $\beta_{it} = \mathbb{E}d_{it}$ (represented by the dot-dashed line).

This amount of learning required to surpass the base case is higher than in the examples for the model and algorithm from Section 4.4; the number of periods required for effective learning still seems practical given that this model is most appropriate for highly-responsive customers, like customers with automated devices. These types of customers are better equipped to respond to 40+ demand response periods in a year.

Note that the long-run improvement of the online learning framework for the demand response program is guaranteed by the upper-bounds on regret in Section 4.5.2. However,

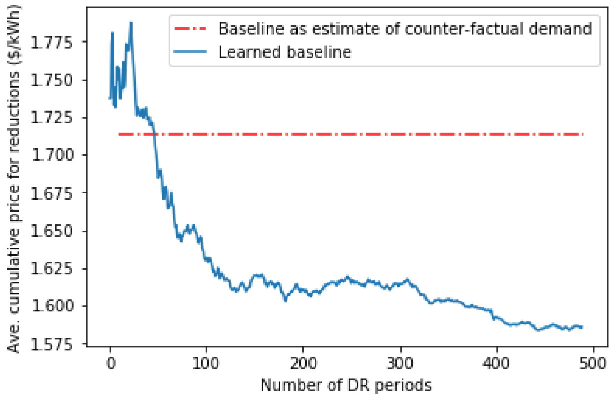

(a) Random participation levels

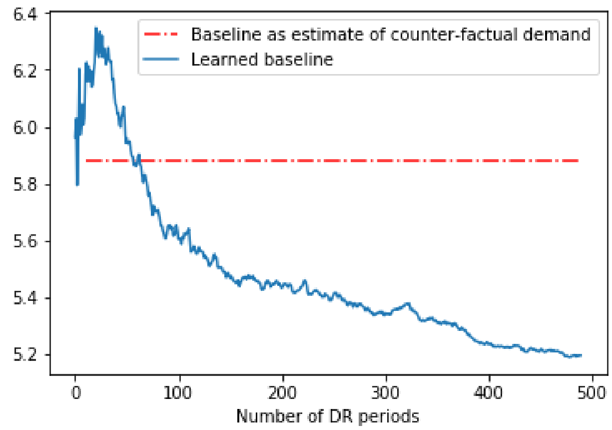

(b) Low participation

Figure 4-4: Demand response learning with responsive customers.

the practical behavior of the program for a small, finite number of DR periods is strongly dependent on the initialization. For these examples, we assumed that $\beta_{it} \in [7.5, 10.5]$ and tested the case where $\lambda = 1.5$. A wider range for permissible β_{it} has a high cost in terms of regret in the short-term, because the online learning algorithm needs to spend a significant number of demand response periods exploring low-value parameters. Therefore, it is practically important to carefully initialize the possible range for β_{it}, taking into account historical programs and expected performance. This issue would be an interesting area for future work.

Note that here and in Section 4.6 we fixed $\alpha_t = 1$ for all t for simplicity. This is higher than a typical DR incentive; a reasonable range for a demand response incentive is $0.3 - 0.6$ $/kWh above the retail rate, so the actual cost per unit of reduction would be about half of the normalized values reported in the Figures in these sections, assuming the same levels of reduction.

4.7 Conclusion

This chapter investigates the sequential decision problem of choosing customer baselines in incentive-based demand response programs. We considered two different models of customer demand and of LSE's objective functions for demand response programs. In a reasonable objective function, the LSE will incorporate (1) the cost of incentivizing customers, (2) customer satisfaction, and (3) the total level of demand reductions. These objectives are often in conflict.

Under these customer and LSE models, it is clear that the optimal demand response baseline is not necessarily an estimate of the customer's counterfactual demand, which is the typical practice. We show how an LSE can use tools from online learning to sequentially set customer baseline thresholds, offering customers a linear payment for demand response reductions below the chosen threshold. This allows them to optimize the demand response program to suit their particular objectives, overcoming uncertainty about customer counter-factual consumption and demand reduction and even lack of knowledge regarding the underlying distributions of those customer-specific variables. In simple examples, we validate our approach and show that it can outperform current practice.

Chapter 5

Conclusion and Future Work

The growth of variable renewable energy will continue to challenge and inspire change in the electricity sector. In this thesis, we investigate challenges that could test energy market design as renewable energy generation continues to grow. Using tools from optimization, decision sciences, probability, and statistics, we study problems related to market power, forward contracting, and demand-side participation in future electricity markets.

In the first chapter, we investigate producer strategy and market power in energy markets with high levels of stochastic renewable energy. We model the market equilibrium as a Bayesian Nash Equilibrium; the features of the equilibrium are heavily influenced by the probabilistic relationship between the energy outputs of renewable resources. We show how correlation between random energy availability can impact market power and welfare, and we consider the value of public information sharing of high-quality forecasts.

In the second chapter, we model the impact of forward contracting on market power in electricity spot markets. We investigate incentives for forward contracting and their downstream impacts on producer market power. We show that changes in the retail side of the electricity sector—specifically a reduction in concentration, or an increase in the number of LSEs serving end-use consumers—can reduce forward contracting. In many states, including Massachusetts and California, retail electricity choice and municipal aggregation are increasing competition in the retail electricity sector. The results in this chapter suggest that these impacts could reduce incentives for forward contracting. Forward contracting is an important tool for financing investment in renewable energy, and it helps reduce market power in short-term electricity markets. Therefore, this chapter highlights a potential downside of increased retail competition.

In the third chapter, we study the iterative decision process for choosing customer baseline thresholds in incentive-based demand response programs. Incentive-based demand response programs can help reduce electricity demand from residential and small commercial customers during time periods when the cost of delivered electricity is high. These programs are popular because they provide no downside risk to customers: they pay customers for reductions, but do not increase prices for consumption. Despite its practical popularity, this model creates challenges for designing an effective compensation mechanism. We show how tools from online learning can be used to iteratively choose baseline thresholds for customers, offering them a demand response incentive with a linear reward for consumption below the baseline threshold. This approach learns and utilizes customer-specific information in order to most effectively allocate program funds to reduce electricity demand. In simple examples, this methodology outperforms the current practice, where the baseline threshold is an estimate of the customer's counter-factual consumption.

There are exciting opportunities for research in optimization, economics, statistics, and control theory to improve electricity markets and grid operations in order to enable a low-cost, low-carbon grid. Given the scope of potential research activities, we consider here just a few interesting possibilities for new research extending or validating the specific ideas in this thesis.

In order to aid our understanding of market power in future electricity systems, research efforts could test new modeling frameworks and consider practical questions of market power monitoring in high-renewable energy systems. Theoretical research on producer strategy and market power could directly model joint ownership of fossil-fuel and renewable technologies, focusing on the impacts of uncertainty and resource heterogeneity under that scenario. Additional theoretical work could consider the impact of forward markets and forward contracts, including day-ahead markets and long-term contracts. Future work should focus on practical questions for market power monitoring in electricity systems where the regulator has imperfect information regarding energy availability from renewable generators.

In addition, further research could focus on strategic behavior of energy storage assets and the ability of energy storage operators to manipulate market prices. Existing research focuses on the energy arbitrage capabilities of storage or its potential value for ancillary services. When energy storage is co-located with renewable energy assets, it could enhance the ability of producers to exercise market power; if a strategic wind producer withholds energy in a particular time period, this energy could potentially be used to charge an energy storage device and sold at a later period, instead of wasted at zero value. Research could investigate the impacts of low-cost storage on market power. The dynamic optimization problem of operating storage for energy arbitrage is very interesting; the potential for strategic behavior adds new complexity and research interest. Finally, research could also focus on practical questions for market power monitoring when regulators have imperfect information regarding opportunity costs for storage.

Research efforts could continue to investigate the link between retail electricity competition and forward contracting in electricity markets. Researchers could develop a more detailed model that incorporates multiple incentives for forward contracting, including price-risk reduction. It would also be valuable to gain better understanding of current forward contracting behavior by energy suppliers, corporate customers, and financial intermediaries. An empirical analysis could estimate the magnitude of the potential impacts described in Chapter 3. Besides the issues raised in this thesis, a related concern is that increased retail competition could drive free-loading for reliability and resource adequacy products: small retailers might not engage in forward contracts because grid operators can not preferentially shut-off their customers during shortage periods or reliability events. Future research could investigate the impacts of increased retail competition on forward contracting by focusing on these market failures.

Well-designed incentives can improve the responsiveness of demand to changing grid conditions and reduce the cost of supplying reliable, low-carbon electricity. Researchers could extend our results by improving the methods and by considering enhancements that improve the rate of learning. Tools from machine learning, especially reinforcement learning, could be utilized to choose baseline thresholds during sequential demand response events. Our approach has zero average regret in the limit; other approaches could have better practical results even if they do not have provable bounds. Our method learns about customers individually, but clustering methods could be useful to the extent that many customers have similar usage patterns and behaviors. Research could use clustering methods to identify similar customers and learn simultaneously about average parameters for a group of similar customers. In each demand response period, we could learn more quickly about customer groups proportional to the size of each cluster. New research can tackle additional challenges related to incentive-based demand response. Researchers could consider the joint problem of setting the baseline threshold (our focus) and also setting the demand-response incentive, which could vary from hour to hour in some DR programs. Our approach helps mitigate adverse selection, by reducing payments to low-performing or low-value customers. New research could seek to extend the literature on the problem of moral hazard in DR programs, which could be especially pertinent in programs with larger numbers of automated devices.

Finally, research should continue to focus on alternative methods for demand response, including direct load-control, and on understanding and mitigating barriers to cost-reflective pricing.

Bibliography

Acemoglu, D., A. Kakhbod, and A. Ozdaglar (2017). Competition in Electricity Markets with Renewable Energy Sources. *The Energy Journal 38*(SI1), 137–155.

Addy, N., J. L. Mathieu, S. Kiliccote, and D. S. Callaway (2013). Understanding the effect of baseline modeling implementation choices on analysis of demand response performance. In *ASME 2013 International Mechanical Engineering Congress and Exposition*, pp. V011T06A024—-V011T06A024. American Society of Mechanical Engineers.

Agrawal, S. (2018). Lecture notes on multi-armed bandits and reinforcement learning.

Allaz, B. and J.-L. Vila (1993). Cournot competition, forward markets and efficiency. *Journal of Economic Theory 59*(1), 1–16.

Anderson, E. J. and A. B. Philpott (2002). Using supply functions for offering generation into an electricity market. *Operations Research 50*(3), 477–489.

Arbabzadeh, M., R. Sioshansi, J. X. Johnson, and G. A. Keoleian (2019). The role of energy storage in deep decarbonization of electricity production. *Nature communications 10*(1), 1–11.

Auer, P., N. Cesa-Bianchi, and P. Fischer (2002). Finite-time analysis of the multiarmed bandit problem. *Machine learning 47*(2-3), 235–256.

Auer, P., N. Cesa-Bianchi, Y. Freund, and R. E. Schapire (1995). Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pp. 322–331. IEEE.

Ausubel, L. M. and P. Cramton (2010). Using forward markets to improve electricity market design. *Utilities Policy 18*, 195–200.

Bahrami, S., V. W. Wong, and J. Huang (2017). An online learning algorithm for demand response in smart grid. *IEEE Transactions on Smart Grid 9*(5), 4712–4725.

Bird, L., M. Milligan, and D. Lew (2013). Integrating variable renewable energy: Challenges and solutions. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States).

Borenstein, S., J. Bushnell, and C. R. Knittel (1999). Market power in electricity markets: Beyond concentration measures. *The Energy Journal 20*(4), 65–88.

Borenstein, S. and S. Holland (2005). On the efficiency of competitive electricity markets with time-invariant retail prices. *RAND Journal of Economics 36*(3), 469–493.

Bouffard, F. and F. Galiana (2008, may). Stochastic Security for Operations Planning With Significant Wind Power Generation. *IEEE Transactions on Power Systems 23*(2), 306–316.

Burger, S., I. Schneider, A. Botterud, and I. Pérez-Arriaga (2019). Fair, equitable, and efficient tariffs in the presence of distributed energy resources. In F. Sioshansi (Ed.), *Consumer, Prosumer, Prosumager: How Service Innovations will Disrupt the Utility Business Model*, pp. 155. Academic Press.

Burger, S. P., C. R. Knittel, I. J. Pérez-Arriaga, I. Schneider, and F. Vom Scheidt (2020). The efficiency and distributional effects of alternative residential electricity rate designs. *The Energy Journal 41*(1), 199–239.

Bushnell, J., M. Flagg, and E. Mansur (2017). Capacity Markets at a Crossroads. Technical report, Energy Institute at Haas, UC Berkeley.

Butner, M. (2018). Gone with the wind: Consumer surplus from renewable generation. In *Evolving Energy Realities: Adapting to What's Next, 36th USAEE/IAEE North American Conference, Sept 23-26, 2018*. International Association for Energy Economics.

Camacho, E. F., T. Samad, M. Garcia-Sanz, and I. Hiskens (2011). Control for renewable energy and smart grids. *The Impact of Control Technology, Control Systems Society*, 69–88.

Cappers, P., C. A. Spurlock, A. Todd, P. Baylis, M. Fowlie, and C. Wolfram (2016). Time-of-use as a default rate for residential customers: issues and insights. In *Lawrence Berkeley National Laboratory*. LBNL-1005704.

Cardell, J. B., C. C. Hitt, and W. W. Hogan (1997). Market power and strategic interaction in electricity networks. *Resource and energy economics 19*(1-2), 109–137.

Carroll, R. J. and P. Hall (1988). Optimal rates of convergence for deconvolving a density. *Journal of the American Statistical Association 83*(404), 1184–1186.

Cerisola, S., Á. Baíllo, J. M. Fernández-López, A. Ramos, and R. Gollmer (2009, feb). Stochastic Power Generation Unit Commitment in Electricity Markets: A Novel Formulation and a Comparison of Solution Methods. *Operations Research 57*(1), 32–46.

Chao, H.-p. (2010a). Demand management in restructured wholesale electricity markets. Technical report, ISO New England.

Chao, H.-p. (2010b). Price-responsive demand management for a smart grid world. *The Electricity Journal 23*(1), 7–20.

Chao, H.-p. (2011, feb). Demand response in wholesale electricity markets: the choice of customer baseline. *Journal of Regulatory Economics 39*(1), 68–88.

Chao, H.-P. and R. Wilson (2004). Resource Adequacy and Market Power Mitigation via Option Contracts. Technical report, EPRI.

Clarke, R. N. (1983). Collusion and the incentives for information sharing. *The Bell Journal of Economics*, 383–394.

Cramton, P., A. Ockenfels, and S. Stoft (2013, apr). Capacity Market Fundamentals. *Economics of Energy & Environmental Policy 2*(2).

Cramton, P. and S. Stoft (2008, sep). Forward reliability markets: Less risk, less market power, more efficiency. *Utilities Policy 16*(3), 194–201.

Dappe, M. H. (2008). Market power, forward trading and supply function competition. In *Briding Energy Supply and Demand: Logistics, Competition and Environment,,*. International Association for Energy Economics.

Dattner, I., A. Goldenshluger, A. Juditsky, et al. (2011). On deconvolution of distribution functions. *The Annals of Statistics 39*(5), 2477–2501.

de Arce, M. P., E. Sauma, et al. (2016). Comparison of incentive policies for renewable energy in an oligopolistic market with price-responsive demand. *The Energy Journal 37*(3), 159–198.

De Sisternes, F. J., J. D. Jenkins, and A. Botterud (2016). The value of energy storage in decarbonizing the electricity sector. *Applied Energy 175*, 368–379.

Deng, R., Z. Yang, M.-Y. Chow, and J. Chen (2015). A survey on demand response in smart grids: Mathematical models and approaches. *IEEE Transactions on Industrial Informatics 11*(3), 570–582.

Dobakhshari, D. G. and V. Gupta (2018). A contract design approach for phantom demand response. *IEEE Transactions on Automatic Control 64*(5), 1974–1988.

Downward, A., G. Zakeri, and A. B. Philpottt (2010). On Cournot Equilibria in Electricity Transmission Networks. *Operations Research 58*(42), 1194–1209.

EIA (2018). Form eia-861. https://www.eia.gov/electricity/data/eia861/. Accessed: 2019-11-20.

EIA (2019a). Annual Energy Outlook 2019. Technical report, EIA, Washington, D.C.

EIA (2019b). Electric Power Monthly - September 2019. Technical report, U.S. Energy Information Administration, Washington, D.C.

EIA (2019c). Levelized cost and levelized avoided cost of new generation resources in the annual energy outlook 2019. Technical report, U.S. Energy Information Administration, Washington, D.C.

Einy, E., O. Haimanko, D. Moreno, and B. Shitovitz (2010). On the existence of bayesian cournot equilibrium. *Games and Economic Behavior 68*(1), 77–94.

Einy, E., D. Moreno, and B. Shitovitzc (2003). The value of public information in a cournot duopoly. *Games and Economic Behavior 44*, 272–285.

Elberg, C. and S. Kranz (2014). Capacity Mechanisms and Effects on Market Structure. Technical Report 04.

Fabra, N. and G. Llobet (working paper). Auctions with unknown capacities: Understanding competition among renewables.

Fan, J. (1991). On the optimal rates of convergence for nonparametric deconvolution problems. *The Annals of Statistics*, 1257–1272.

Faruqui, A., S. Sergici, and C. Warner (2017). Arcturus 2.0: A meta-analysis of time-varying rates for electricity. *The Electricity Journal 30*(10), 64–72.

Faulkner, K. (2010). Community choice aggregation in california. Technical report, Berkeley.

Fertig, E., J. Apt, P. Jaramillo, and W. Katzenstein (2012). The effect of long-distance interconnection on wind power variability The effect of long-distance interconnection on wind power variability. *Environ. Res. Lett 7*, 34017–6.

Fischer, C. (2010). Renewable portfolio standards: when do they lower energy prices? *The Energy Journal*, 101–119.

Foley, A. M., P. G. Leahy, A. Marvuglia, and E. J. McKeogh (2012). Current methods and advances in forecasting of wind power generation. *Renewable Energy 37*(1), 1–8.

Fudenberg, D. and J. Tirole (1991). Game theory, 1991. *Cambridge, Massachusetts 393*(12), 80.

Fuller, R. and A. Berwick (2018, December). Municipal electricity aggregation really works.

Gaffey, W. R. (1959). A consistent estimator of a component of a convolution. *The Annals of Mathematical Statistics 30*(1), 198–205.

Gal-Or, E. (1986). Information transmission—cournot and bertrand equilibria. *The Review of Economic Studies 53*(1), 85–92.

Gaudet, G. and S. W. Salant (1991). Uniqueness of cournot equilibrium: new results from old methods. *The Review of Economic Studies 58*(2), 399–404.

Gilotte, L. and D. Finon (2006). Investments in generation capacities in an oligopolistic electricity market. *The Energy Journal*.

Goldberg, M. L. and G. K. Agnew (2013). Measurement and verification for demand response. Technical report, US Department of Energy: Washington, DC, USA.

Green, R. and N. Vasilakos (2010). Market behaviour with large amounts of intermittent generation. *Energy Policy 38*(7), 3211–3220.

Harvey, S. M. and W. W. Hogan (2000). California Electricity Prices and Forward Market Hedging. Technical report.

Hirth, L. (2013). The market value of variable renewables: The effect of solar wind power variability on their relative price. *Energy economics 38*, 218–236.

Hirth, L., F. Ueckerdt, O. Edenhofer, et al. (2016). Why wind is not coal: on the economics of electricity generation. *The Energy Journal 37*(3), 1–27.

Hobbs, B. F. and J. S. Pang (2007, feb). Nash-Cournot Equilibria in Electric Power Markets with Piecewise Linear Demand Functions and Joint Constraints. *Operations Research 55*(1), 113–127.

Hogan, W. (2010). Demand response pricing in organized wholesale markets. *ISO/RTO Council Comments on Demand Response Compensation in Organized Wholesale Energy Markets 13*.

Hogan, W. W. (1997). A market power model with strategic interaction in electricity networks. *The Energy Journal 18*(4), 107–141.

Hogan, W. W. (2005). On an "Energy Only" Electricity Market Design for Resource Adequacy. Technical report.

Hogan, W. W. (2013). Electricity scarcity pricing through operating reserves. *Economics of Energy & Environmental Policy 2*(2), 65–86.

Holmberg, P. (2007). Supply function equilibrium with asymmetric capacities and constant marginal costs. *The Energy Journal*, 55–82.

Hortacsu, A. and S. L. Puller (2008). Understanding strategic bidding in multi-unit auctions: a case study of the texas electricity spot market. *The RAND Journal of Economics 39*(1), 86–114.

Inman, R. H., H. T. Pedro, and C. F. Coimbra (2013). Solar forecasting methods for renewable energy integration. *Progress in energy and combustion science 39*(6), 535–576.

Ito, K. and M. Reguant (2016). Sequential markets, market power, and arbitrage. *American Economic Review 106*(7), 1921–57.

Joskow, P. and J. Tirole (2007, mar). Reliability and competitive electricity markets. *The RAND Journal of Economics 38*(1), 60–84.

Joskow, P. L. (2006). Competitive Electricity Markets and Investment in New Generating Capacity. Technical report.

Joskow, P. L., R. Schmalensee, and Others (1988). Markets for power: an analysis of electrical utility deregulation. *MIT Press Books 1*.

Joskow, P. L. and J. Tirole (2000). Transmission rights and market power on electric power networks. *The Rand Journal of Economics*, 450–487.

Joskow, P. L. and C. D. Wolfram (2012). Dynamic pricing of electricity. *American Economic Review 102*(3), 381–85.

Kalathil, D. and R. Rajagopal (2015). Online learning for demand response. In *Communication, Control, and Computing (Allerton), 2015 53rd Annual Allerton Conference on*, pp. 218–222. IEEE.

Ketterer, J. C. (2014, jul). The impact of wind power generation on the electricity price in Germany. *Energy Economics 44*, 270–280.

Khazaei, J., G. Zakeri, and S. S. Oren (2017). Single and Multisettlement Approaches to Market Clearing Under Demand Uncertainty. *Operations Research 65*(5), 1147–1164.

Khezeli, K. and E. Bitar (2017). Risk-Sensitive Learning and Pricing for Demand Response. *IEEE Transactions on Smart Grid*, 1–1.

Kim, J. H. and W. B. Powell (2011). Optimal energy commitments with storage and intermittent supply. *Operation Research 59*(6), 1347–1360.

Kleinberg, R. D. (2005). Nearly tight bounds for the continuum-armed bandit problem. In *Advances in Neural Information Processing Systems*, pp. 697–704.

Kök, A. G., K. Shang, and Ş. Yücel (2016). Impact of electricity pricing policies on renewable energy investments and carbon emissions. *Management Science 64*(1), 131–148.

Kroposki, B., B. Johnson, Y. Zhang, V. Gevorgian, P. Denholm, B.-M. Hodge, and B. Hannegan (2017). Achieving a 100% renewable grid: Operating electric power systems with extremely high levels of variable renewable energy. *IEEE Power and Energy Magazine 15*(2), 61–73.

Léautier, T.-O. (2016). The visible hand: ensuring optimal investment in electric power generation. *The Energy Journal 37*(2).

Lee, M., O. Aslam, B. Foster, D. Kathan, J. Kwok, L. Medearis, R. Palmer, P. Sporborg, and M. Tita (2013). Assessment of demand response and advanced metering. *Federal Energy Regulatory Commission, Tech. Rep*.

Leiserowitz, A., E. Maibach, S. Rosenthal, J. Kotcher, A. Gustafson, P. Bergquist, M. Ballew, and M. Goldberg (2018). Energy in the american mind-december 2018. Technical report, New Haven, CT: Yale Program on Climate Change Communication: Yale University and George Mason University.

Li, N., L. Chen, and M. A. Dahleh (2015, jul). Demand Response Using Linear Supply Function Bidding. *IEEE Transactions on Smart Grid 6*(4), 1827–1838.

Lijesen, M. G. (2007). The real-time price elasticity of electricity. *Energy Economics 29*, 249–258.

Mahenc, P. and F. Salanié (2004, jun). Softening competition through forward trading. *Journal of Economic Theory 116*(2), 282–293.

Makarov, Y. V., C. Loutan, J. Ma, and P. De Mello (2009). Operational impacts of wind generation on california power systems. *IEEE transactions on power systems 24*(2), 1039–1050.

Meir, R., H. Ma, and V. Robu (2017, may). Contract Design for Energy Demand Response.

Meister, A. (2009). Deconvolution problems in nonparametric statistics: lecture notes in statistics. *Lecture Notes in Statistics 193*.

Meyn, S., T. Samad, I. Hiskens, and J. Stoustrup (2018). *Energy Markets and Responsive Grids: Modeling, Control, and Optimization*, Volume 162. Springer.

Miller, L., R. Carriveau, and S. Harper (2018). Innovative financing for renewable energy project development–recent case studies in north america. *International Journal of Environmental Studies 75*(1), 121–134.

Mohajeryami, S., M. Doostan, and A. Asadinejad (2016, feb). An investigation of the relationship between accuracy of customer baseline calculation and efficiency of Peak Time Rebate program. In *2016 IEEE Power and Energy Conference at Illinois (PECI)*, pp. 1–8. IEEE.

Mohsenian-Rad, A.-H., V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia (2010, dec). Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid. *IEEE Transactions on Smart Grid 1*(3), 320–331.

Munoz, F. D., S. Wogrin, S. S. Oren, and B. F. Hobbs (2018). Economic inefficiencies of cost-based electricity market designs. *The Energy Journal 39*(3).

Muthirayan, D., E. Baeyens, P. Chakraborty, K. Poolla, and P. P. Khargonekar (2019). A minimal incentive-based demand response program with self reported baseline mechanism. *arXiv preprint arXiv:1901.02923*.

Nexant Inc. (2017). California ISO Baseline Accuracy Work Group Proposal. Technical report.

Nolan, S. and M. O'Malley (2015). Challenges and barriers to demand response deployment and evaluation. *Applied Energy 152*, 1–10.

Novshek, W. (1985). On the existence of cournot equilibrium. *The Review of Economic Studies 52*(1), 85–98.

O'Connell, N., P. Pinson, H. Madsen, and M. O'Malley (2014). Benefits and challenges of electrical demand response: A critical review. *Renewable and Sustainable Energy Reviews 39*, 686–699.

Oren, S. (1997). Economic inefficiency of passive transmission rights in congested electrical system with competitive generation. *The Energy Journal 18*(1), 63–83.

Papavasiliou, A. and S. S. Oren (2013, jun). Multiarea Stochastic Unit Commitment for High Wind Penetration in a Transmission Constrained Network. *Operations Research 61*(3), 578–592.

Park, S., S. Ryu, Y. Choi, J. Kim, and H. Kim (2015, sep). Data-Driven Baseline Estimation of Residential Buildings for Demand Response. *Energies 8*(9), 10239–10259.

Pinczynski, M. and R. Kasperowicz (2016). Overview of electricity market monitoring. *Economics & Sociology 9*(4), 153.

Pineau, P.-O., H. Rasata, and G. Zaccour (2011). A dynamic oligopolistic electricity market with interdependent market segments. *The Energy Journal*, 183–217.

Pinson, P. (2013). Wind energy: Forecasting challenges for its operational management. *Statistical Science 28*(4), 564–585.

Qi, W., Y. Liang, and Z.-J. M. Shen (2015, dec). Joint Planning of Energy Storage and Transmission for Wind Energy Generation. *Operations Research 63*(6), 1280–1293.

Richter, J. (2013). Incomplete information in cournot oligopoly: The case of unknown production capacities. Technical report, EWI Working Paper.

Rintamäki, T., A. S. Siddiqui, and A. Salo (2017). Does renewable energy generation decrease the volatility of electricity prices? an analysis of denmark and germany. *Energy Economics 62*, 270–282.

Rivier, M. and I. J. Pérez-Arriaga (1993). Computation and decomposition of spot prices for transmission pricing. In *Proc. 11th PSC Conf.*

Roos, A. and T. F. Bolkesjø (2018). Value of demand flexibility on spot and reserve electricity markets in future power system with increased shares of variable renewable energy. *Energy 144*, 207–217.

Rossetto, N. (2018). Measuring the intangible: An overview of the methodologies for calculating customer baseline load in pjm. Technical report, European University Institute.

Ruff, L. E. (2002). Economic Principles of Demand Response in Electricity. *Edison Electric Institute*.

Sakai, Y. (1985). The value of information in a simple duopoly model. *Journal of Economic Theory 36*(1), 36–54.

Schneider, I. and M. Roozbehani (2017a). Energy market design for renewable resources: Imbalance settlements and efficiency–robustness tradeoffs. *IEEE Transactions on Power systems 33*(4), 3757–3767.

Schneider, I. and M. Roozbehani (2017b). Wind capacity investments: Inefficient drivers and long-term impacts. Technical report, MIT CEEPR.

Schneider, I. and C. R. Sunstein (2017). Behavioral considerations for effective time-varying electricity prices. *Behavioural Public Policy 1*(2), 219–251.

Shaked, M. and J. G. Shanthikumar (2007). *Stochastic orders*. Springer Science & Business Media.

Shapiro, C. (1986). Exchange of cost information in oligopoly. *The review of economic studies 53*(3), 433–446.

Shariatzadeh, F., P. Mandal, and A. K. Srivastava (2015). Demand response for sustainable energy systems: A review, application and implementation strategy. *Renewable and Sustainable Energy Reviews 45*, 343–350.

Siano, P. (2014). Demand response and smart grids: A survey. *Renewable and sustainable energy reviews 30*, 461–478.

Sinden, G. (2007). Characteristics of the uk wind resource: Long-term patterns and relationship to electricity demand. *Energy Policy 35*(1), 112–127.

Surampudy, M., B. Chew, M. Keller, and T. Gibson (2019). 2019 utility demand response market snapshot. Technical report.

Szidarovszky, F. and S. Yakowitz (1977). A new proof of the existence and uniqueness of the cournot equilibrium. *International Economic Review*, 787–789.

Todd, A., P. Cappers, C. A. Spurlock, and L. Jin (2019). Spillover as a cause of bias in baseline evaluation methods for demand response programs. *Applied Energy 250*, 344–357.

Vardakas, J. S., N. Zorba, and C. V. Verikoukis (2015). A Survey on Demand Response Programs in Smart Grids: Pricing Methods and Optimization Algorithms. *IEEE Communications Surveys & Tutorials 17*(1), 152–178.

Ventosa, M., A. R. A. Baill, and M. Rivier (2005). Electricity market modeling trends. *Energy Policy 33*(7), 897–913.

Vives, X. (1984). Duopoly information equilibrium: Cournot and bertrand. *Journal of economic theory 34*(1), 71–94.

Wang, Q., M. Liu, and J. L. Mathieu (2014). Adaptive demand response: Online learning of restless and controlled bandits. In *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 752–757. IEEE.

Wang, X., P. Guo, and X. Huang (2011). A Review of Wind Power Forecasting Models. *Energy procedia 12*, 770–778.

Wijaya, T. K., M. Vasirani, and K. Aberer (2014). When bias matters: An economic assessment of demand response baselines for residential customers. *IEEE Transactions on Smart Grid 5*(4), 1755–1763.

Willems, B., I. Rumiantseva, and H. Weigt (2009). Cournot versus supply functions: What does the data tell us? *Energy Economics 31* (1), 38–47.

Wolak, F. (2000). An Empricial Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market. *International Economic Journal 14* (2).

Wolfram, C. D. (1998). Strategic bidding in a multiunit auction: An empirical analysis of bids to supply electricity in england and wales. *The Rand Journal of Economics*, 703–725.

Woo, C.-K., I. Horowitz, J. Moore, and A. Pacheco (2011). The impact of wind generation on the electricity spot-market price level and variance: The texas experience. *Energy Policy 39* (7), 3939–3944.

Wozabal, D., C. Graf, and D. Hirschmann (2016). The effect of intermittent renewables on the electricity price variance. *OR spectrum 38* (3), 687–709.

Yao, J., I. Adler, and S. S. Oren (2008). Modeling and computing two-settlement oligopolistic equilibrium in a congested electricity network. *Operation Research 56* (1), 34–46.

Zhong, H., L. Xie, and Q. Xia (2013, may). Coupon Incentive-Based Demand Response: Theory and Case Study. *IEEE Transactions on Power Systems 28* (2), 1266–1276.