

Dynamic competition and arbitrage in electricity markets: The role of financial players

Ignacia Mercadal ^{*†}

May 10, 2016 Most recent version [here](#)

Abstract

Over the last decade, many electricity markets have introduced purely financial trading alongside transactions between operators who own physical generation capacity and entities, such as utilities, that serve physical demand. As it has been in other markets, the effect of these financial trades is the subject of an ongoing debate; while they are expected to increase liquidity and informational efficiency, they have also been blamed for higher prices and led to allegations of price manipulation. This paper studies the role of financial trading by examining a natural experiment in the Midwest electricity market. A 2011 regulatory change exogenously attracted more financial players to this market, and a rich dataset on individual behavior allows me to study both physical and financial participants' reaction to it. First, I use a reduced form analysis to show that the regulatory change led to more financial trading, and reduced generators' market power. I then use a structural approach to examine the causal relationship between these two observations, which requires the computation of the residual demand faced by each firm. A major challenge here is that electricity markets are segmented by transmission lines with limited capacity, which creates local markets in which only a subset of the firms competes. I deal with this issue using techniques from machine learning, presenting a new method to study the competitive structure of electricity markets. My findings indicate that financial trading decreases generators market power, but does not fully eliminate it. As a consequence, consumers are better off but productive efficiency might go down.

*I want to thank my advisors Ali Hortaçsu, Michael Greenstone, John Birge, and Brent Hickman for invaluable guidance and support. I am also grateful to Mar Reguant, Frank Wolak, Derek Neal, Ignacio Cuesta, Gunnar Heins, and the participants of seminars at the Berkeley Energy Camp, the Heartland Workshop, the UChicago IO lunch, the UChicago Applications of Economics workshop, the University of Washington, Rice University, the University of Wisconsin-Madison, Carnegie Mellon Heinz, NYU Stern, Columbia SIPA, MIT, MIT Sloan, Georgetown McDonough, Johns Hopkins Carey, UChicago Harris, University of California Santa Barbara, Carlos III in Madrid, BI Norwegian School of Business, EIEF, and the University of Florida for helpful comments and suggestions. Thanks to the Energy Policy Institute at UChicago for financial support. This work was completed in part with resources provided by the University of Chicago Research Computing Center and Social Sciences Computing Center. All errors are my own.

[†]Department of Economics, University of Chicago, ignaciamercadal@uchicago.edu,

1 Introduction

The role of financial traders in commodity markets is controversial. Although they are expected to facilitate risk sharing and increase informational efficiency, distrust of financial traders is so widespread that some politicians have even proposed restrictions and bans on their activity.¹ Among the reasons for this bad reputation is that speculators are frequently blamed for higher and more volatile prices, and accused of market manipulation.² In this paper, I employ a unique dataset to study the role of speculators as competitors of physical producers in the Midwest electricity market.

Typically, it is hard to identify the effect of speculation on a commodity market because only aggregate market outcomes are observed and the physical good is not traded together with its derivatives. Electricity markets provide an excellent setting to study the effects of financial trading since all transactions involving both physical producers and financial players occur in a single market. This paper focuses specifically on the Midwest electricity market (MISO³), which has two additional advantages. First, a regulatory change in 2011 that exogenously attracted more financial traders, which allows me to identify the effect these traders had on the market. Second, I observe individual-level behavior and can separately analyze how buyers, producers, and financial traders reacted to the regulatory change. Exploiting these unique features, this paper shows that financial trading decreases physical producers' market power and increases consumer welfare, but potentially at the cost of lower productive efficiency.

In electricity markets, financial trading takes place in sequential markets, which is how most wholesale electricity markets are organized. There is first a forward market that schedules production a day in advance, and then a spot market that balances demand and supply immediately before operation. Although under certain conditions⁴ the forward price should be equal to the expected spot price, in practice systematic differences between the two have been

¹E.g. see Grossman and Stiglitz (1980); Grossman (1976); Silber (1985) for the benefits. Former Congressman Joseph Kennedy II proposed to ban 'pure' speculators from trading oil futures. He says "Eliminating pure speculation on oil futures is a question of fairness. The choice is between a world of hedge-fund traders who make enormous amounts of money at the expense of people who need to drive their cars and heat their homes, and a world where the fundamentals of life—food, housing, health care, education and energy—remain affordable for all." You can find the Op Ed article [here](#).

²See Juvenal and Petrella (2014); Kilian and Murphy (2014); Knittel and Pindyck (2013) for a discussion on volatile and high prices. On price manipulation, Birge et al. (2014)

³Midwest Independent System Operator.

⁴Perfect competition, risk neutrality and zero transaction costs, for instance.

documented in most electricity markets.⁵ For this reason, most markets have introduced financial traders, expecting them to arbitrage this forward premium down to zero.

By closing the gap between the forward and spot prices, financial traders prevent generators from engaging in intertemporal price discrimination between these two markets. In the spot market, generators face demand from consumers who were not willing to buy at the forward price. Therefore, they have an incentive to sell at a lower price in the spot market, as this will increase their profits without affecting the price received for forward sales.⁶ Speculators arbitrage the resulting premium and thus make the forward market more competitive. As noted by Ito and Reguant (2014), generators will still have market power in the spot market, since financial traders cannot increase electricity production.

In the presence of a forward premium, buyers would be expected to shift their purchases to the spot market in order to pay a lower price. However, in deregulated electricity markets demand comes from regulated utilities that can pass increased costs on to their customers, which decreases their price sensitivity. They can also pass on the cost of hedging, which makes them even less sensitive to prices. Finally, purchases in the spot market are subject to high deviation charges that significantly reduce the amount buyers save by buying there. These factors result in a relatively unresponsive demand⁷, which gives large generators market power and the ability to price discriminate.

In the Midwest electricity market the forward premium persisted despite the presence of financial traders because high transaction costs made arbitrage unprofitable (Birge et al., 2014). A regulatory change lowered these costs significantly in April 2011, after which financial trading increased and the forward premium became smaller. As a consequence, we expect price discrimination in the forward market to decrease, since increased financial trading means more arbitrage. Interestingly, generators not only reacted to the regulatory change by exerting less market power in the forward market, but they did it months before it was implemented.

⁵Bowden et al. (2009) and Birge et al. (2014) find it in the Midwest, Saravia (2003) in New York, Jha and Wolak (2013); Borenstein et al. (2008) in California, Ito and Reguant (2014) in the Iberian market, among others.

⁶Borenstein et al. (2008) show that firms price discriminated in California right before the energy crisis. As opposed to most cases, they find a negative forward premium because market power was on the demand side. Ito and Reguant (2014) show that generators with market power engage in intertemporal discrimination, which results in a forward premium.

⁷Around 85% of the demand bids in the forward market just specify a quantity that they are willing to buy at any price, for instance.

In order to understand the generators' reaction, I use individual bid data to estimate a model of optimal generator behavior, following the approach of Wolak (2000) and Hortaçsu and Puller (2008).⁸ I build a static model of a firm that decides how much to sell in the forward and spot markets. In MISO, these markets are organized as sequential auctions in which firms bid step functions specifying how much they are willing to sell or buy at each price. I extend Hortaçsu and Puller (2008)'s model of optimal bidding in the spot market to the case of a sequential market in which buyers may also have market power in the forward market. As in their model, I include firms forward contract positions as a determinant of profits.⁹ Firms usually hedge by signing contracts for differences that pay sellers (buyers) when the market price is lower (higher) than the price agreed upon in the contract.

In my model, a firm's optimal bid depends on its contract position, as well as on the elasticity of its residual demand, *i.e.* total demand minus the quantity sold by competitors. Since future competition does not affect residual demand today, the model predicts that generators in the Midwest will only lose market power when transaction costs are reduced and financial trading increases. Consequently, the model rationalizes the observed change in behavior as a reaction to changes in current market conditions, *i.e.* residual demand or contract positions. I test the model's optimality condition empirically and find that it does not hold, which means that, rather than reacting to current market conditions, generators changed their behavior in anticipation of increased competition in the future.

I consider two alternative hypotheses, *i.e.* two mechanisms that could explain why firms changed their conduct before market conditions changed. The first is a cooperative equilibrium in a repeated game, which is sustained as long as a player's benefits from continued cooperation outweigh the gains from deviating and stealing the market today. In the context of this paper, increased arbitrage in the future eliminates the benefits from future cooperation, because speculators will arbitrage away any resulting price gap. This mechanism would explain the

⁸There are a number of papers following this approach in electricity markets. Wolak (2007) uses the optimality conditions obtained from static profit maximization to estimate firms marginal costs. He tests the hypothesis of profit maximization and finds no evidence against it. Hortaçsu and Puller (2008), on the other hand, find that while it describes big firms behavior well, small firms are far less sophisticated. Reguant (2014) studies an auction in which firms bids can include complementarities across production hours, which she uses to estimate startup costs. Ryan (2014) estimates marginal costs from the participants bids in the Indian market, taking into account transmission constraints to estimate the consequences of transmission investment.

⁹Wolak (2000) shows that the forward contract position affects a firm's incentives to exert market power.

generators' anticipatory reaction, since the equilibrium unravels as soon as it is known that cooperation cannot be sustained in the future. This is a market in which the same firms interact with each other every day, and have good information about demand and each others' costs. Although there are many firms in the market, a few large ones control most of the production.¹⁰

The second mechanism is entry deterrence. If generators expected financial, or "virtual", traders to enter the market and arbitrage the forward premium, they might have tried to make the market less attractive by lowering the forward premium. Entry deterrence does not seem to be sustainable in equilibrium, as there is no link between periods that could make today's competition affect the entrant's profits in the future.¹¹ Nonetheless, I include this mechanism for the sake of completeness, since the generators' pricing changes might have been a failed attempt to deter the entry of financial traders.

The test I use to evaluate generators' conduct is based on a simple intuition. In a repeated game cooperative equilibrium, firms do not play best response, but behave as if the market were *less* competitive than it is. Under entry deterrence, generators do not play best response either, but they act as if the market were *more* competitive than it is. Therefore, comparing the elasticity of demand actually faced by firms with that implied by their behavior allows me to distinguish between entry deterrence, tacit collusion, and static Nash equilibrium. Although these alternatives do not exhaust the space of alternatives, they can be taken as examples of two different ways in which the null hypothesis of static best response can be rejected, *i.e.* behaving *more* or *less* competitively than what would be optimal under static best response.

This paper also introduces a number of methodological contributions, which I now describe. In structural analysis, optimality conditions are usually imposed on the data and used to obtain an estimate of primitive parameters from the model. Instead, the richness of my data allows me to compute every component of the optimality condition for the forward market, *i.e.* I can construct the residual demand faced by each participant and compute its elasticity. I use the empirical counterpart of the optimality condition obtained from the model to distinguish

¹⁰Evidence of tacit collusion in electricity markets has been found by Fabra and Toro (2005) in the Spanish electricity market.

¹¹Like dynamic demand, for instance, as in Goolsbee and Syverson (2008). In this context, increasing capacity as in Dixit (1980) would not make a firm's threat more credible. Obtaining reputation as a fighter could justify lowering today's profits to deter future entry (Milgrom and Roberts, 1982), but in this case the market became more competitive before entry.

between the three hypotheses that could explain the generators behavior: tacit collusion, entry deterrence, and static Nash (reaction to something different from financial traders that affected residual demand). The test evaluates whether generators' behavior indicates that they perceive residual demand elasticity to be lower than, higher than, or the same as that observed in the data.¹²

The residual demand faced by each generator can in principle be computed by adding up the demand bids and subtracting the supply bids of the generator's competitors. However, this exercise is complicated by the fact that the Midwest electricity market is a nodal market, *i.e.* there may be a different clearing price in each location or node where electricity is generated or demanded. This price represents the marginal cost of supplying energy at that node, which varies significantly because nodes are connected by transmission lines with limited capacity. When the lines reach maximum capacity, demand cannot necessarily be satisfied by the lowest cost generator and is instead satisfied by the lowest cost feasible generator.

To deal with nodal pricing, I assume the MISO market is split into several independent markets. I define these markets empirically by using machine learning techniques that cluster nodes together based on the correlation of their prices. Unlike most applications of these clustering techniques, I build a measure of fit that allows me to choose the market definition that better fits the data. To do this, I simulate the quantities and prices that would clear under each alternative market definition and compare them with those observed in the data. I find that the clusters fit the data fairly well. To the best of my knowledge, this paper is the first to use a structural model to study a nodal market, which is made possible by these market definitions.

My findings indicate that, prior to learning of the impending regulatory change, firms acted as if they were facing a less competitive market than they were, and therefore exerted more market power than would be optimal given the elasticity of the residual demand they actually faced. After learning about the future fall in transaction costs for financial traders, the firms moved closer to a static Nash equilibrium. This reaction is consistent with a repeated game cooperative equilibrium that unravels when future benefits from cooperation disappear. The forward market

¹²Puller (2007) studies the competitiveness of the electricity market in California. Using a Cournot model for the spot market, he simulates the price that would have resulted under perfect competition, Cournot-Nash, and tacit collusion, concluding that the market is well described by the static Nash equilibrium. My approach is similar but I have the advantage of observing all bids. Additionally, I am mainly interested in the effect of financial trading on competition, instead of assessing the overall competitiveness of the market.

becomes more competitive because of increased financial trading, and it does so even before trading goes up. Therefore, this result underscores the importance of considering dynamics when investigating the role of financial traders.

Although my results indicate that generators decreased the amount of market power they exerted in the forward market in response to increased financial trading, increased arbitrage did not eliminate their market power. In the same way that a monopolist who changes from price discrimination to uniform pricing increases prices in the low demand market, I find that firms exerted more market power in the spot market after arbitrage increased. Although this did not reduce total production, since demand is perfectly inelastic in the spot market, it may have increased costs because demand was served by more expensive firms.¹³

Overall, the effect of arbitrage on welfare is ambiguous. Consumers are better off, saving roughly \$1,800,000 per day on average, but producers are worse off because they cannot price discriminate. Nonetheless, even though total quantity does not change, the total effect is not just a transfer from producers to consumers because production costs may change. On the one hand, firms' increased exertion of market power leads to higher costs. On the other, less underscheduling in the forward market results in better planning, which allows cheaper generating units to be scheduled and decreases production costs. Although I do not quantify these effects, the latter is more likely to dominate since most production is scheduled in the forward market.

The rest of this paper is organized as follows: The next section describes the Midwest electricity market and explains its conditions before the regulatory change. Section 3 then describes how the different players in the market reacted to lower transaction costs for financial traders. Section 4 describes the dataset used for empirical analysis. Section 5 presents a static model of generator behavior in a sequential auction, as well as a brief description of the three hypotheses that could explain the generators' behavior. The empirical strategy is described in section 6, and results are presented in section 7. Section 8 concludes.

¹³Ito and Reguant (2014) find that firms indeed exert more market power in the spot market when there is more arbitrage, and this results in increased costs.

2 The MISO energy market

Wholesale electricity markets are different from other markets because electricity cannot be stored, supply needs to meet demand at every moment, and the transmission network that transports electricity from sellers to buyers has limited capacity. As a consequence, proper administration of the transmission grid is essential for reliability and efficiency. For this reason, deregulated electricity markets typically operate under an Independent System Operator (ISO), a non-profit organization that coordinates the use of the transmission grid by the different market participants. In the Midwest, this role is played by the MidContinent ISO¹⁴, which covers 15 U.S. states and the Canadian province of Manitoba. MISO operates an energy market that serves 42 million people and collects US\$20 billion in gross charges per year.

The energy market is organized as an auction in which participants submit bids to buy or sell energy in particular locations; the ISO then clears the market solving a nonlinear programming problem that minimizes cost subject to the capacity constraints imposed by the transmission network. Because the network has limited capacity, electricity supplied at different locations is not a homogeneous good. MISO deals with this heterogeneity by allowing each node or location to clear at a different price, which is known as nodal pricing and described in more detail in Appendix A.

The MISO energy market has over 2000 pricing nodes and often becomes congested (reaches capacity), so in practice there is significant price dispersion among the nodes. Figure 1 presents heat maps of the MISO market in two different moments, which illustrate how prices can substantially differ geographically and over time. There are two reasons why the presence of congestion is relevant for the analysis of this market. First, when lines are at capacity demand cannot always be served by the cheapest generator. In practice, congestion segments the market, creating local markets in which firms have more market power than they would otherwise. Second, congestion poses a challenge for empirical analysis, because the degree of market power enjoyed by a firm depends directly on the level of congestion and its transmission structure. I address this problem by using prices to define independent markets within the MISO market with a machine learning algorithm. This is described in Section 6.1.

¹⁴Midwest ISO until 2013.

2.1 Forward and spot market

Like many deregulated electricity markets, the MISO energy market is structured as a sequential auction. First, there is a day-ahead or forward market that schedules production for the 24 hours of the next day, and then a real-time or spot market that balances demand and supply 30 minutes before each operating hour. Both of these markets are auctions organized by the market operator.

The forward or day-ahead market is a financial market that takes place once a day and clears separately for each hour of the next day. Until 11 a.m. of each day, buyers and sellers submit bids for each of the 24 hours of the next day, starting with the midnight hour. The real-time market is a balancing market and takes place 30 minutes before each operating hour. Only physical supply bids are allowed, and the market is cleared by minimizing the cost of satisfying the forecasted demand subject to transmission constraints. The bulk of demand comes from utilities that sell to most of their final consumers at a fixed price per MWh. This makes demand very inelastic in the short run, and demand bids are therefore not accepted in the real-time market.

A generator can be a seller or a buyer in the spot market, depending on the quantity she schedules in the forward market. Firms are paid for the quantity sold in the forward market regardless of how much they actually produce, but the difference between the forward schedule and the actual production is settled at the spot price. For instance, if a generator schedules 100MWh in the forward market, but then clears 80MWh in the spot market, she receives the forward price for 100MWh but has to pay the spot price for 20MWh, as if she were buying.¹⁵

The rationale behind a sequential market is that generation is cheaper when it is planned, so scheduling forecasted demand in advance decreases production costs. Generators with lower marginal costs generally have high startup costs and cannot adjust the level of production easily. On the other hand, generators that can start and vary production quickly, called peakers, often have high marginal costs. By scheduling production in the forward market, it is easier to satisfy expected demand with cheaper generators and only unanticipated shocks with peakers. Additionally, scheduling the 24 hours of the next day in the forward market increases efficiency by taking into account complementarities across hours, which come from the startup costs faced

¹⁵See Jha and Wolak (2013) for a complete description of how multi-settlement markets work.

by some generating units.¹⁶ The existence of the forward market also allows market participants to face less risk, as price is more volatile in the real-time than in the day-ahead market.¹⁷

Although it would be efficient to schedule enough generation to satisfy all forecasted load in the day-ahead market and only use the real-time market to adjust for unexpected shocks, market participants do not always have incentives to do so. The most costly deviations are those that result in insufficient generation being scheduled in the forward market, because in such cases the market authority needs to quickly cover demand by increasing production, dispatching peakers, and starting inactive plants. This happens either when generation scheduled in the forward market becomes unavailable in the spot market, or when real-time demand is larger than scheduled (for instance, because not enough generation was scheduled as a result of high price offers). Because the clearing price does not cover ramping or startup costs, but only marginal cost, firms that buy in the spot market are subject to deviation charges called Revenue Sufficiency Guarantee (RSG) charges. The revenue collected is then distributed among participants who incurred ramping or startup costs.

2.2 The forward premium and market power

In a perfectly competitive market, the forward market would schedule generation to cover the forecasted demand in full, and the spot market would only be used to manage unexpected shocks. If that were the case, we would expect the forward price and the expected spot price to be the same. Nonetheless, predictable differences between the forward and the spot price have been observed in many wholesale electricity markets, including MISO (Bowden et al., 2009; Birge et al., 2014). Figure 2 shows that the monthly median forward premium between 2009 and 2010 in the MISO energy market was consistently positive during 2009 and 2010. This was also the case in the Iberian market (Ito and Reguant, 2014), New York (Saravia, 2003), and California (Jha and Wolak, 2013).

A positive premium results from market power on the supply side (Ito and Reguant, 2014). Generators with market power have incentives to engage in intertemporal price discrimination by withholding sales in the forward market, thus increasing the forward price and creating a

¹⁶See Reguant (2014) for an analysis of the welfare consequences of allowing complementarities in bids.

¹⁷Additionally, Allaz and Vila (1993) show that sequential markets enhance competition among firms when they compete à la Cournot.

premium. The behavior of generators in the MISO market is consistent with these incentives; on average, most generators increase their production in the spot market.¹⁸

Market power is a concern in many deregulated electricity markets for several reasons. First, both demand and supply are very inelastic. Demand ultimately comes from households that pay a fixed price and are thus insensitive to prices. Supply is inelastic for technological reasons: plants with lower marginal costs are usually unable to make short term production adjustments. Second, electricity cannot be stored, so intertemporal arbitrage is not possible. And lastly, electricity is produced and demanded at particular locations or nodes connected by transmission lines with limited capacity. When capacity is reached, demand cannot always be satisfied by the cheapest generator, effectively reducing the number of competitors for each firm.

Previous empirical research has found that generators have market power in deregulated electricity markets. For example, Borenstein et al. (2002) study the California market between 1998 and 2000, and find that generators had considerable market power even though the market was not very concentrated. Puller (2007) studies this market during the same period and concludes that generators' conduct is consistent with Cournot competition, but not collusion. Ito and Reguant (2014) study the Iberian market and find that firms have market power and exert it by engaging in intertemporal price discrimination, which creates a forward premium. Ryan (2014) finds that increasing transmission capacity can lead to substantial gains in terms of welfare, most of which comes from a reduction in generators' market power. Fabra and Toro (2005) study competition in the Spanish market and find that firms were in a collusive equilibrium that sometimes broke and gave rise to a price war. These studies have found evidence of market power in zonal markets, where prices are allowed to differ. In nodal markets like MISO there are even more opportunities for firms to exert market power.

During 2010 and 2011, the period under study, there were 95 generators actively participating in the MISO energy market.¹⁹ The largest single firm's market share was just 7%, but together the 10 largest firms held 55% of generation capacity, the largest 15 firms held 70%, and the largest 20 held 77%. Because of the limited capacity of the transmission lines that transport electricity,

¹⁸The argument is similar to the one behind the Coase conjecture (Coase, 1972). After selling a given quantity in the forward market, a generator has incentives to increase its sales in the spot market since she will receive a lower price, but it will not affect the price charged in the forward market. Anticipating this, the generator splits its sales between the forward and the spot market.

¹⁹95 had positive sales at least one day during that period.

the MISO market is split into multiple local markets in which concentration is higher and firms have more market power.

Market power on the demand side results in a spot premium, *i.e.* a higher price in the spot than in the forward market. Because buyers are better off when prices are lower, utilities with market power have incentives to withhold purchases in the forward market in order to lower the price. Although less common than a forward premium, such a spot premium was observed by Borenstein et al. (2008) in California in the year 2000. Different price caps for the clearing price in the forward and spot market, along with the absence of penalties for demand not scheduled in the forward market, allowed large utilities to exert market power and lower the forward price. MISO's market rules and market monitoring specifically aim at avoiding under-scheduling of demand.

2.3 Virtual or financial participants

The presence of a forward premium creates opportunities for arbitrage by short selling in the forward market. Even if the market does not allow explicit arbitrage in the form of purely financial transactions, firms have incentives to engage in implicit arbitrage by adjusting their bids when they trade physical energy (Jha and Wolak, 2013).²⁰ Because firms are only allowed to arbitrage at nodes where they have plants, and generators cannot short sell, arbitrage under this circumstances is limited and generators are still able to exert market power in the forward market.

There are efficiency costs associated with intertemporal price discrimination and implicit arbitrage. Firms' production schedules in the forward market depend not only on their costs, but also on whether they are exerting market power or arbitraging, which means that they do not pursue pure cost minimization. In order to avoid these inefficiencies, many deregulated electricity markets have introduced virtual or financial participants, explicit arbitrageurs who profit from differences between the forward and the spot market.

The introduction of financial traders to wholesale electricity markets has been controversial. On the one hand, the forward premium decreased after arbitrageurs were allowed in the California

²⁰For instance, a generator could sell more in the forward market than it intends to sell in the spot market, and then buy the difference in the spot market at a lower price.

(Jha and Wolak, 2013) and New York (Saravia, 2003) markets. On the other hand, Birge et al. (2014) find that arbitrage was limited due to institutional constraints on financial bidding, and that financial bids were used to unlawfully manipulate the price of a related financial instrument used to hedge congestion in the MISO market. In fact, one trader has already agreed to pay a 5 million dollars settlement to avoid a trial on this charge.²¹

Virtual participants have been allowed in the MISO energy market since it first started operating in 2005. These bidders profit from the differences between the forward and the spot price. For instance, selling 1 MWh in the forward market yields profits equal to $P^{forward} - P^{spot}$ because it implicitly requires the purchase of 1MWh in the spot market. In the presence of a forward premium, financial participants have incentives to sell in the forward market. Under perfect arbitrage, these bids would shift forward supply up to the spot market level, neutralizing generators' underbidding and leading to price convergence.

Birge et al. (2014) show that in 2010 the forward price was significantly higher than the spot price. There was limited arbitrage because financial participants were subject to deviation charges that were at least as high as the forward premium. These fees were the RSG charges imposed on spot purchases described in section 2.2. Virtual bidders do not sell any physical energy, so a virtual forward sale entails an equal spot purchase that was subject to RSG charges. The average forward premium was \$0.9, which is the revenue from selling 1MWh in the forward market and buying it in the spot market. Nonetheless, RSG charges per MWh were \$1.8 on average, making the transaction unprofitable.

On April 2011, the Federal Energy Regulatory Commission (FERC) approved MISO's proposal to modify the way RSG were calculated, so that charges were significantly lowered; they went from \$1.8 per MWh to \$0.3 per MWh.²² As a consequence, financial trading significantly increased and the forward premium began to close. In this paper, I use this exogenous change in virtual trading to study the effect of arbitrageurs on the competitiveness of the market.

The change in RSG charges did not come as a surprise to market participants, but instead occurred after a long debate about how to compute RSG charges and who should be subject to them. A committee of market participants discussed the issue and even drafted the proposed rule that was eventually submitted by MISO for FERC approval. The proposal was announced

²¹<http://www.ferc.gov/enforcement/market-manipulation.asp>

²²See Appendix B for computation details.

and submitted to FERC on December 1, 2010, and the market immediately started preparing for implementation of the changes, which they expected to occur in March 2011. MISO began conducting detailed training sessions on the new calculation in January 2011, and the proposal was finally approved in April.²³

RSG charges acted as transaction costs for financial players, since they were subject to them for every MWh they sold in the forward market. They were not an entry barrier, because the entry cost faced by financial trading firms is the cost of becoming a market participant, which is usually low and not affected by RSG charges. Additionally, MISO assigns a credit limit to each firm, which determines how much they can trade every day and depends on the firm's expected capacity to meet the financial obligations derived from virtual trading. This was not affected by the RSG charges either.

The next section describes the effect of this regulatory change on the behavior of the different market participants. As expected, virtual trading increased and price discrimination by generators decreased after entry barriers for financial traders were lowered.

3 Reaction to the regulatory change

3.1 Virtual participants

When RSG charges dropped, profitable arbitrage opportunities appeared. The expected profit from a virtual supply bid, which is equal to the expected premium minus RSG charges, became larger than zero after the drop in RSG charges. In fact, Birge et al. (2014) show that for the first few months after charges were lowered, it was possible to make a profit using simple rules to predict the sign of the forward premium. As expected, these opportunities were quickly closed by increased virtual trading activity.

The top panel of Figure 4 shows the monthly average of the daily volume traded by virtual bidders. The dashed red line indicates the announcement on December 1, 2010 that the proposal to redesign RSG charges had been submitted to FERC. On that date, the market started preparing

²³In MISO, proposals to change market rules are discussed in groups of stakeholders. The change in RSG charges was reviewed by the Revenue Sufficiency Guarantee Task Force, a group specially created for this purpose. The minute from their meeting in December 2010 states that training sessions for all market participants were going to be held in January, while the minute from January 2011 states they expected the proposal to become effective in March, 2011. These are all available in the MISO website.

for the change in RSG charges. The solid red line on April 1, 2011 indicates the date on which the new RSG proposal was actually implemented. The green line is the monthly virtual trade volume, which seems to have increased after RSG charges were reduced.

In order to confirm that there was a change in virtual activity, I look for a structural break in the time series of daily traded virtual volume.²⁴ The standard test for structural break at a known date is the Chow test, which estimates the parameters before and after the break separately, and then tests for equality using an F statistic. As the date of the break is unknown in this case, I compute the F statistic for all dates in the sample. The maximum value is known as Quandt statistic (Hansen, 2001; Quandt, 1960). I use the critical values provided by Andrews (1993) and largely reject the null hypothesis of stable parameter values across the sample.

I follow Bai and Perron (1998) to find the break dates in the time series. They present a sequential method in which it is first determined if there is a structural break and when it occurs; then the sample is split at the break date, and the same method is applied to each of the subsets to determine if there are more breaks. The method to find the date of the break is as follows. For a given break date d , run OLS separately for the samples before and after d and compute the total RSS (whole sample). Do this for all possible break dates except those in the extreme 15% of the sample, and pick the break date as the d at which the RSS reaches a minimum.

The bottom panel of Figure 4 plots the residual sum of squares for each potential break date. The minimum is reached on April 9, 2011, with the confidence interval between April 6, 2011 and April 12, 2011 (Bai and Perron, 1998). This break point confirms the observation that virtual bidders changed their behavior after RSG costs were reduced. The blue line in Figure 4 shows the mean traded volume before and after the breakpoint, indicating that it increased around 40%. This shows that the reduction in RSG charges indeed attracted more financial trading.

Continuing the sequential algorithm to find more breaks indicates that there was another break on January 28, 2010. The confidence interval for this break varies depending on the time period used in the sample, but it is wide (around a month), which indicates that the estimation of this break is not very precise. As it is at the beginning of the sample, it should not affect

²⁴These tests have been used in the applied micro literature before by Greenstone and Hanna (2014); Fabra and Toro (2003). Similar to the spirit of this paper, Fabra and Toro (2003) examine the British electricity market's response to a regulatory change and use breakpoint test to determine whether the response is anticipated, and thus consistent with collusion.

estimation or the conclusions regarding the effect of the regulatory change. Nevertheless, for robustness, whenever results are obtained for a period including this date, they are compared with those starting after this break.

3.2 Generators

On average, generators' spot sales are positive, *i.e.* generators produce a larger quantity of electricity than they schedule in the forward market. This can be observed in the top panel of Figure 5, which shows in green the average daily spot sales for each month. There are a few things worth noting in this figure. The first is that the sales are generally positive, which means that generators, on average, use the spot market to increase their production.

Secondly, Figure 5 shows that spot sales became smaller when RSG charges were reduced. This is in line with expectations if generators are exerting market power. As long as transaction costs for financial participants are high, generators can engage in intertemporal price discrimination between the forward and spot markets, increasing the forward premium. Once these costs are lower, virtual traders will arbitrage this gap between the forward and spot prices, making underbidding by generators less attractive. In this way, financial participants effectively increase the elasticity of the residual demand faced by generators and thus decrease their market power in the forward market.²⁵

Lastly, in Figure 5 it seems that generators reacted before the regulatory change was implemented. As before, the dashed and solid red lines indicate the dates of the announcement and implementation of the regulatory change, respectively. The green line represents generators' spot sales, which appear to have decreased before transaction costs were reduced in April 2011. In fact, using the same tools described in the previous section, I find that there was a structural break on January 10, 2011, with a confidence interval between January 5 and January 15. I do not find other breaks that are robust to changing the sample periods.

The generators' early reaction to the regulatory change is surprising, since their decision about how to split sales between the forward and spot markets today is independent of the same decision in the future. It is possible that the change in the level of spot sales is due to external factors and

²⁵Arbitrageurs do not reduce generators' market power in the spot market. In fact, the following sections will show that producers exert more market power in the spot market after financial trading increases.

not firms' behavior. To the best of my knowledge, there were no important changes in the market clearing algorithm or the market structure around these dates. Wind power became subject to RSG charges in August 2010, and intermittent power sources like wind became dispatchable -*i.e.* able to be turned on or off by the market operator according to demand- in July 2011, but it is not clear how this could affect quantities cleared in the forward and spot market in the observed manner.

I follow a structural approach to determine whether the generators' change in behavior in January 2011 was an anticipatory reaction to a future increase in financial trading, or a static reaction to other factors that made the market more competitive at that moment. In Section 5, I present a model for the generator's decision about how much to bid in the forward and the spot markets. This model allows me to understand what factors affect a generator's behavior, which I later use to empirically determine whether the observed reaction could have been caused by these factors.

Before presenting the model and the empirical strategy that will allow me to better understand the generators' behavior, I describe the reaction of demand to the changes in RSG pricing in the next section.

3.3 Demand

For generators to be able to increase the forward price by withholding sales in the forward market, certain conditions have to be true. First, just as with standard price discrimination, there has to be limited or no arbitrage. As explained above, RSG charges imposed on virtual supply bids were initially high enough to make arbitrage unprofitable.

Second, demand has to be less responsive than supply. If demand reacts by shifting purchases to the spot market, the effect will be the same as that of arbitrage, so generators will not be able to price discriminate. Moreover, if buyers have market power, they will withhold purchases in the spot market to lower the forward price as was observed in California by Borenstein et al. (2008), creating a spot premium. Although the fact that the forward price is larger than the spot price already suggests that the stronger market power is on the supply side, this section describes demand behavior to confirm buyers are not the ones driving the reaction to the regulatory change.

Figure 6 shows spot purchases in the MISO energy market, indicating that, on average, sales in

the spot market are net positive. That is, buyers generally do not schedule enough production in the forward market to meet demand and must cover the difference in the spot market. Although this is consistent with market power on the demand side, it is also what a price-taker buyer facing a forward premium would do to minimize its purchasing cost. As for generators, their behavior in the spot market -*i.e.* whether they choose to buy or sell- provides information about their market power. A price-taker seller facing a forward premium would short sell in the forward market, while a generator exerting market power would reduce forward sales in order to increase the forward price. It is not as simple to infer market power on the demand side because with or without market power, energy buyers are better off by withholding purchases in the forward market. A price-taker buyer wants to buy as little as possible in the forward market because the price is lower in the spot market. A buyer with market power restricts its demand in the forward market in order to lower the price. Therefore, in the presence of a forward premium, purchases in the spot market are expected to be positive. As shown in Figure 6, this is the case in the MISO energy market.

As Figure 6 shows, buyers were initially withholding purchases in the forward market, and spot purchases decreased after RSG charges were reduced. As explained above, not much can be read from this behavior since it can come from both price-takers or firms with market power. Nonetheless, I find a structural break in the net purchases time series on January 26, 2011.²⁶ This indicates that demand reacted before the change in RSG charges was actually implemented, but after generators did, suggesting demand responded to the generators' reaction and not directly to the regulatory change.

Purchasers' late response, as well as the fact that the forward premium was positive both before and after the regulatory change, which is advantageous for sellers, indicates that the premium was being driven by generators rather than purchasers. This may seem surprising because utilities are large companies and are generally expected to have considerable market power. There are a few reasons why demand may not have reacted as much as would be expected. First, many utilities can pass increased costs directly to final consumers, which makes them price insensitive. Second, MISO and the market monitor pay special attention to demand underscheduling. If utilities exerted too much market power by declining to purchase electricity

²⁶ The confidence interval for the break date is between January 20 and February 2.

in the overpriced forward market, they could be sanctioned by the authorities. Third, spot purchases are subject to RSG charges, which makes spot sales expensive for buyers. Lastly, demand may be hedged as there are financial instruments available to hedge the risk of spot price volatility, particularly because hedging costs are generally among the costs that regulated utilities are allowed to recover.²⁷ I am currently working on backing out demand's financial contract positions to determine how relevant this factor is.

3.4 The market's reaction to the regulatory change

This section has so far shown how the different participants in the energy market reacted to the regulatory change that reduced transaction costs for financial participants. From the previous analysis, it appears that financial traders reacted as expected to the forward premium created by generators' market power. Generators, in turn, exerted less market power, which is consistent with increased financial arbitrage making price discrimination more difficult. Purchasers also reduced their net spot purchases as generators exerted less market power. These are the expected reactions from market participants in a static setting.

Although the way in which participants reacted to the regulatory change was expected, the timing of their reactions was not. The fact that generators reacted months before the implementation of the regulatory change does not fit a static model of firm behavior in the energy market. There are two potential causes for this unexpected timing. The first is that unobserved market conditions changed at the same time as the regulatory change, and generators were actually responding to that unknown change. The second alternative is that market conditions remained the same, but the energy market is better understood using a dynamic model in which future changes have effects on present behavior. I will use a structural analysis to distinguish between these two cases.

In a static setting, a generator's decision about how to split sales between the forward and spot markets depends mainly on the residual demand it faces. My structural analysis builds the residual demand faced by each generator and computes its optimal decision given that demand. Then the optimal decision is compared to the firm's observed decision. In the next sections I will

²⁷Regulated utilities are allowed to earn a certain rate of return on capital. To calculate the tariffs that they can charge, estimate costs are subtracted from revenues. Hedging costs are among the costs they can include here.

show that generators initially behaved as if they faced a residual demand less elastic than the one they actually faced. When they learned about the impending entry of virtual traders, their behavior became closer to that predicted by a static model. This is consistent with a cooperative equilibrium in a repeated game in which firms cooperate and reach outcomes better than static Nash. The equilibrium breaks, however, as soon as they learn that the game has a final period.

4 Data

Most of the empirical work in this paper is done using an hourly panel that is publicly available on MISO's website. It contains each participant's bid, as well as the corresponding cleared quantity and price for each hour between 2010 and 2011. The panel has around 100 millions observations, 20 millions from generators' bids, and 80 millions from demand and financial participants' bids.

Demand bids may specify only a quantity (price-taker), or a step function with up to 10 pairs of price and quantity (price-sensitive). Only around 15% of demand bids are price-sensitive, while the remainder simply specify a quantity that the purchaser is willing to buy at any price. Purchasers only participate in the forward market, but MISO publishes aggregate hourly data on total quantities cleared by demand in the forward and spot markets. Table 1 presents summary statistics on demand bids. Most bids are price-takers, with price-sensitive bids being placed by fewer firms and at fewer nodes.

Generators may submit price-taker or price-sensitive bids as well, and they also have the option of submitting an increasing piecewise linear function instead of a step function (see Figure 3 for an illustration). In my sample, 70% of generator bids and 82% of the megawatts hour cleared by generators correspond to piecewise linear bids. I discretize these bids as step functions in intervals of 0.1 MWh in my analysis, which results in residual demands with many steps. Supply bids also include information about the technological restrictions of each plant, such as the minimum/maximum number of hours it needs to operate, ramping times and costs, and startup costs. I do not observe these variables, as MISO only publishes the bid, cleared price and quantity, maximum and minimum production levels under normal and emergency conditions, and the amount a generator sells as a price taker.

The data identify buyers who place bids at multiple nodes, and sellers who own multiple units, but it is not possible to know which participants are vertically integrated utilities, nor whether a generator is also using virtual bids to hedge or arbitrage. Summary statistics on bids are presented in Table 2 for virtual traders, and Tables 3 and 4 for supply bids in the forward and spot markets, respectively. Notice that while around 90% of physical demand bids are cleared in the forward, only around 10% of virtual bids and 50% of physical supply bids are cleared.

Additionally, MISO posts the clearing prices at each pricing node in the market, information I use to match bids, which are not reported by node, to the corresponding nodes. Summary statistics for prices and their components can be found in Table 18. In my data, a node is just a number and a name where one or more participants submit bids. Each node's geographical location is not disclosed, because it is considered a matter of national security.

I use data on prices and volumes of traded Intercontinental Exchange (ICE) futures for the Indiana hub during peak hours. These data are available on the EIA website. Data on oil, coal, and natural gas prices were obtained from the Federal Reserve Bank of St. Louis. They correspond to daily crude oil prices (West Texas Intermediate - Cushing, Oklahoma), the Henry Hub natural gas spot price, and coal prices in two coal regions (Illinois Basin and Powder River Basin).

5 Model

5.1 Static Model

In this section I consider the decision of a generator that sells its production in a sequential auction. For every given hour of electricity production, there is first a forward auction to schedule generation, and then a spot auction to adjust prices and quantities immediately before electricity is needed. I assume bidding decisions for both markets are taken simultaneously by the firm when deciding how much to produce and how to split sales between the forward and spot markets.²⁸

My model modifies that of Hortaçsu and Puller (2008) by introducing sequential markets instead of focusing on the spot market. Additionally, I account for the limited capacity of electrical transmission lines by assuming that both the forward and spot markets are segmented into M

²⁸ I am working on some simple checks to assess the plausibility of this assumption by looking at whether spot bids are affected by the quantities or prices cleared in the forward market.

independent markets. This is a simplifying assumption, since in practice any node can potentially affect any other at a given moment, depending on the level of congestion and the characteristics of the transmission network. I make this assumption for two reasons. First, it makes the model tractable by allowing each market to clear independently. Second, it matches the empirical strategy that I follow to deal with congestion and nodal pricing, which in turn matches the observed data fairly well (see Section 6.1). Empirical papers on wholesale electricity markets have avoided this problem by studying markets in which congestion is adjusted for in a separate market (Ito and Reguant, 2014; Reguant, 2014), looked at hours without congestion (Hortaçsu and Puller, 2008), or studied zonal markets for which congestion data are available (Ryan, 2014). Jha and Wolak (2013) study the effect of financial traders in California, where prices are nodal as well, but use data on prices and do not fit a structural model because CAISO does not publish bid data. To the best of my knowledge, this is the first paper in the economics literature to use a structural model to analyze a nodal market.

Demand

Demand for each market has the same structure in the forward and spot markets. I assume demand in each market m and period t is given by

$$D_{m,t}(p) = d_{m,t}(p) + \epsilon_{m,t} \quad (1)$$

where $d_{m,t}(p)$ is a non-stochastic component and ϵ is a demand shock. I will omit the period subindexes t because I am using a static model, and therefore all equations are the same for every period and there are no connections between periods.

For the spot market, this is a very natural assumption since demand comes from households, who mostly pay a fixed rate per MW and are thus price insensitive. In fact, there are no demand-side bids in the spot market, as enough generation is cleared to cover MISO's short-term load forecast for each hour. For the forward market, this equation is a simplification, since demand is expressed by bids and can be strategic. Nonetheless, under the same assumptions used for generators, optimality conditions for generators remained unchanged when demand is strategic.²⁹ This extended model is presented in Appendix C.

²⁹The residual demand they face will change, but not the condition for the optimality of the generator's bid.

Supply

Generators usually use financial contracts to hedge risk. These contracts specify a certain quantity x and a price h . If the market clearing price \tilde{p} is greater than the contract price, the firm has to pay $(\tilde{p} - h)x$ to the buyer of the contract; if the contract price exceeds the clearing price, the firm is paid $(\tilde{p} - h)x$ by the buyer of the contract (Green, 1999; Wolak, 2000, 2003a). They are settled in terms of the differences between the prices because these contracts are purely financial and do not require physical delivery of energy. It is important to account for forward price hedging contracts in the analysis of generators' decisions, because a firm's financial position determines whether it has incentives to increase or decrease the forward price.³⁰

Often generators hold physical contracts in addition to financial ones. These specify a price and a quantity as well, but in this case energy is delivered to the buyer, who pays the price specified in the contract. These contracts can be treated as sunk costs because they are negotiated in advance and therefore do not affect the generator's decision about how to split sales between the forward and the spot market. Physical contracts affect costs if production costs are not linear, but even in such cases we can simply assume that $C(0)$ in the model is equal to the cost of producing the quantity specified in the physical contract. For this reason, physical contracts are not explicitly included in the model.

Generators decide how much to produce, and how to split sales between the forward and spot markets. Each generator i submits a schedule $Q_i(p^F)$ to the forward market auction, and a schedule $S_i(p^S)$ to the spot market auction. These schedules specify how much a generator is willing to sell at each price. In this section, the quantity cleared in the spot market when the clearing price is \tilde{p}^S , $S_i(\tilde{p}^S)$, is the total quantity produced by generator i , *not* the difference between total production and the quantity scheduled in the forward market. Each generator's strategies $Q_i(p^F, x^F)$ and $S_i(p^S, x^S)$ depend on the firm's contract positions, since these positions affect the firm's preferences for sales in the forward or spot market.

Each generator i has a cost function $C_i(q)$, where q is the quantity cleared in the spot market, *i.e.* the quantity actually produced. I assume generators know each others' cost functions. This is not a strong assumption since the same firms interact with each other over long periods, and the only information required to compute costs are the technical characteristics of the plant, which

³⁰ See Wolak (2000) for the importance of contracts on incentives to exert market power.

do not change over time, and fuel prices, which are easy to observe. Forward hedging positions, on the other hand, are harder to observe because they change over time for each firm, which is why I assume that the hedging positions are private information.

Market clearing

In the forward market, the market clearing price \bar{p} in market m is determined by the forward market clearing condition

$$\sum_{j \in m} Q_j(\bar{p}) = d_m(\bar{p}) + \epsilon_m \quad (2)$$

Market clearing in the spot market is the same; the clearing price is determined by balancing demand and supply: $\sum S(\bar{p}^S) = d(\bar{p}^S) + \epsilon$.

Generator's uncertainty

Each generator i faces uncertainty over the clearing prices \tilde{p}^F and \tilde{p}^S , because she does not know what clearing price will result from submitting different schedules. This uncertainty comes from two sources. First, the demand function has a stochastic component that shifts its level unpredictably. Second, a generator does not know other generators' bids. Although a generator knows her competitors' cost functions, she does not know their financial positions with respect to the forward and spot prices. In other words, the generator is uncertain about the residual demand she faces, because residual demand depends on other firms' bidding behavior.³¹

Bidder i 's uncertainty is represented by $F(x_{-i}, \epsilon | x_i)$, the joint distribution of other firms' contract positions and the demand shock. It is conditional on i 's own position because i 's position may contain information about others' contracts. Correlation between the demand shock and the contract positions of the competitors is allowed, but note that this remains a private value setting since i 's profits do not depend on its competitors' contracts (Hortaçsu and Puller, 2008). To distinguish between the forward and the spot market, I define $F^F(x_{-i}^F, \epsilon^F | x_i^F)$ and $F^S(x_{-i}^S, \epsilon^S | x_i^S)$.

Following Hortaçsu and Puller (2008), I define a probability measure over the realizations of the forward clearing price from the perspective of firm i , conditional on i 's private information

³¹*i.e.* the market demand minus the schedules submitted by all other generators in the market.

about its contract position x_i^F , i 's submission of a schedule $\hat{Q}_i(p, x_i^F)$, and her competitors playing their equilibrium strategies $\{Q_j(p, x_j^F), j \neq i\}$.

$$H(p, \hat{Q}_i(p); x_i^F) \equiv \Pr(\tilde{p}^F \leq p \mid x_i^F, \hat{Q}_i) \quad (3)$$

$H(p, \hat{Q}_i(p); x_i^F)$ represents the uncertainty over the forward market clearing price faced by firm i . It is the probability, given i 's contract position, that generator i will be paid a price p when she sells a quantity $\hat{Q}_i(p)$ and all other generators submit the equilibrium offer functions. The event $\tilde{p}^F \leq p$ is equivalent to the event of excess supply at price p . Using the market clearing condition in Equation 2, H can be written as

$$\begin{aligned} H(p, \hat{Q}_i(p); x_i^F) &= \Pr\left(\sum_{j \neq i} Q_j(p, x_j^F) + \hat{Q}_i(p) \geq D^F(p) \mid x_i^F, \hat{Q}_i\right) \\ &= \int_{x_{-i}^F \times \epsilon^F} 1\left\{\sum_{j \neq i} Q_j(p, x_j^F) + \hat{Q}_i(p) \geq D^F(p)\right\} dF^F(x_{-i}^F, \epsilon^F \mid x_i^F) \end{aligned} \quad (4)$$

Equivalently, generator i 's uncertainty over the clearing price in the spot market can be represented by the probability measure G , defined as

$$\begin{aligned} G(p, \hat{S}_i(p); x_i^S) &\equiv \Pr(\tilde{p}^S \leq p \mid x_i^S, \hat{S}_i) \\ &= \Pr\left(\sum_{j \neq i} S_j(p, x_j^S) + \hat{S}_i(p) \geq D^S(p) \mid x_i^S, \hat{S}_i\right) \\ &= \int_{x_{-i}^S \times \epsilon^S} 1\left\{\sum_{j \neq i} S_j(p, x_j^S) + \hat{S}_i(p) \geq D^S(p)\right\} dF^S(x_{-i}^S, \epsilon^S \mid x_i^S) \end{aligned} \quad (5)$$

The generator's problem

At clearing prices \tilde{p}^F and \tilde{p}^S in the forward and spot market, respectively, the ex-post profits for generator i are given by

$$\Pi_i(\tilde{Q}, \tilde{S}) = \tilde{p}^F \tilde{Q} + \tilde{p}^S [S - Q] - C(\tilde{S}) - [\tilde{p}^F - h^F] x^F - [\tilde{p}^S - h^S] x^S \quad (6)$$

where \tilde{Q} is $Q(\tilde{p}^F, x^F)$ and \tilde{S} is $S(\tilde{p}^S, x^S)$ (the arguments are omitted for clarity). The spot quantity is defined as total sales, *i.e.* the total quantity the generator commits to produce. Note

that this is a different definition than the one used in previous sections, where I used “spot quantity” to refer to the quantity in excess of that sales in the forward market.³² The last two terms of the profits come from the financial position held by the generator in the forward and spot markets. As explained above, these are contracts for differences so a firm gets profits when the market price is lower than the contracted price, and losses if the market price is larger.

A firm chooses schedules $Q_i(p^F, x_i^F)$ for the forward market and $S_i(p^S, x_i^S)$ for the spot market so as to maximize its expected profits. Using the clearing price distributions defined above, the generator’s problem is

$$\max_{Q_i, S_i} \int_{\underline{p}}^{\bar{p}} \int_{\underline{p}}^{\bar{p}} U\left(\Pi_i(Q_i, S_i)\right) dH(p^F, Q(p^F); x_i^F) dG(p^S, S_i(p^S); x_i^S) \quad (7)$$

where $Q_i = Q_i(p^F, x_i^F)$ and $S_i = S_i(p^S, x_i^S)$.

The Euler-Lagrange conditions for an interior solution are as follows (proof in Appendix D). Subindexes i are omitted from now on unless necessary to avoid ambiguities.

$$p^F - p^S = - [Q^*(p^F) - x^F] \frac{H_Q}{H_p} \quad (8)$$

$$p^S - c' = - [S^*(p^S) - Q^*(p^F) - x^S] \frac{G_S}{G_p} \quad (9)$$

where $H_Q = \frac{dH}{dQ}$, $H_p = \frac{dH}{dp}$, $G_S = \frac{dG}{dS}$, and $G_p = \frac{dG}{dp}$. H_p is the density of the clearing price in the forward market when all firms submit optimal schedules. H_Q is the change in the price distribution caused by a change in the bid submitted by i , which can be interpreted as a measure of i ’s market power. G_S and G_p have equivalent interpretations in the spot market.

Because the forward market is purely financial, generators’ sales there have no physical cost. Nonetheless, the spot price is the opportunity cost faced by a generator willing to sell in the forward market, since each unit can be sold in either the spot or the forward market. This becomes clear in Equation 8, which is similar to an oligopolist’s first order condition in which the spot price replaces the marginal cost. The forward premium is then a markup with respect to this opportunity cost. Whether the generator wants to have a positive or negative markup will

³²In practice, the spot market is cleared for total production, as in this section. Generators are paid the spot price only for the difference between the quantity scheduled in the forward market and sold in the spot market. This is the sense in which the forward market is financial.

depend on her hedging contract position, because this determines whether the generator is a net seller or a net buyer in the forward market.

A similar trade-off is present in the spot market. The optimal markup for a generator depends on whether she is a net seller or buyer in the spot market, which depends on both her contract position in the spot market and her forward sales. Additionally, the importance of this position is weighted by the firm's ability to affect prices with bids, G_S .

Hortaçsu and Puller (2008) present a separability condition that allows the optimality conditions to be simplified. Intuitively, the condition is that financial contracts shift the optimal bid, but do not change its slope. Formally, it requires schedules to be additively separable in the two sources of uncertainty, which holds when they can be written as $Q_i(p^F, x_i^F) = \alpha_i(p^F) + \beta_i(x_i^F)$. Figure 9 shows some bids that seem to satisfy this assumption, as they are parallel shifts of each other. Section 7.4 presents some empirical evidence backing up this assumption.

If bids are additively separable, the optimality conditions can be written as directly as a function of the residual demand faced by each firm (see Appendix E for a proof)

$$p^F - p^S = -[Q^*(p^F) - x^F] \frac{1}{R'(p^F)} \quad (10)$$

$$p^S - c' = -[S^*(p^S) - Q^*(p^F) - x^S] \frac{1}{R'(p^S)} \quad (11)$$

Using the separability assumption to write the optimality conditions in terms of the residual demand makes it much easier to obtain its empirical counterpart. The residual demand within a market can be constructed from the bids, while the distribution of prices is harder to compute.

Equations 10 and 11 show the conditions that the optimal schedule submitted by a generator needs to satisfy. The optimal markup will depend both on the forward hedging contract position held by the firm, and on the elasticity of the residual demand it faces. Therefore, observing a smaller difference between the quantities sold in the spot and forward market is not enough to conclude that generators are behaving more competitively. If, after controlling for these factors, the evidence still indicates that generators moved their bidding behavior away from the optimal bids determined by the model before the regulatory change was implemented, then there are dynamic elements in play and the market is not well represented using a static model.

One such dynamic element that may be suggested by the data is a collusive equilibrium. A

tacit collusive equilibrium does not need to be an explicit agreement in which firms sit around a table and agree upon each group member's bid. The equilibrium could take the form of a simple rule of thumb for bids in the forward and spot markets. Firms do, however, have some contact, since the large ones are often MISO stakeholders. These stakeholder firms meet periodically to discuss market design and draft joint proposals for market reform. The likelihood that large firms follow similar strategies is also increased because many of these large firms hire outside firms to do their trading. Furthermore, any collusion between these large firms could have a significant impact on prices, since production is fairly concentrated, with 20% of firms controlling 80% of generation capacity.

5.2 Alternative explanations

As will be shown in the later sections, the generators' behavior is not consistent with the static model described above. In this section I present alternative explanations for the generators' anticipatory reaction to the regulatory change.

Although these alternatives do not exhaust the space of alternatives, they can be taken as examples of two different ways in which the null hypothesis of static best response can be rejected. If not playing static best response, firms either behave as if the market were less competitive than it is, which I will associate to a cooperative equilibrium, or they behave as if the market were less competitive than it is, which I associate to entry deterrence. The empirical test developed in the empirical section is based in this simple intuition.

5.2.1 Cooperative equilibrium in a repeated game

The first hypothesis I consider is a cooperative equilibrium in a repeated game. Cooperation is sustained while the benefits from continued cooperation outweigh the gains from deviating and stealing the market today. In the context of this paper, increased arbitrage in the future eliminates the benefits from future cooperation, because speculators will arbitrage away any resulting price gap. This would explain the anticipated reaction since the equilibrium unravels as soon as it is known that cooperation cannot be sustained in the future.

A repeated game is a plausible representation of firms' interaction in this market, since they bid in the same markets every day, and have good information about demand and each others'

costs. Although coordination seems hard in a market where there are 96 firms, a few large ones concentrate most of the productive capacity: the largest 5 firms concentrate 33% of the market capacity; the largest 9, 52%; and the largest 15, 70%. Moreover, when transmission lines are at capacity the market becomes segmented and competition is restricted to a subset of firms, which results in more market power and facilitates coordination.

A cooperative equilibrium does not need to be an explicit agreement in which firms seat around a table and set the bids for each of the members of the group. In this market, firms have regular contact since the large ones are often MISO stakeholders, and therefore typically part of groups in charge of evaluating different elements of the market design, and submitting proposals to improve it. The equilibrium could take the form of a simple rule of thumb about how to bid in the forward and spot markets. Some of the large firms hire external companies to do their trading, which makes more likely that different firms will follow similar strategies.

Illustration

I will use a simple example to illustrate the central elements of a cooperative equilibrium. Consider two firms repeatedly competing à la Cournot in a market with a forward market. Every period, each generator i chooses a quantity Q_F sold in the forward market, and a quantity Q_S sold in the spot market. Inverse residual demand is given by $P_F(Q_F)$ and $P_S(Q_S)$, respectively, and the stage profits are given by

$$\Pi(Q_F, Q_S) = P_F(Q_F)Q_F + P_S(Q_S)[Q_S - Q_F] - C(Q_S) - [P_F - h_F]x_F - [p_S - h_s]x_s$$

The first order conditions are analogous to those from the generator's problem in section 5.1, where each generator chooses a function or schedule instead of a quantity.

$$\begin{aligned} P_F - P_S &= [x_F - Q_F]P'_F(Q_F) \\ P_S - c' &= [x_S - Q_S - Q_F]P'_S(Q_S) \end{aligned}$$

The stage game has a unique equilibrium (Allaz and Vila, 1993), which means that generators will not change their behavior unless either the contract position or the residual demand changes. From the Folk theorem, we know that in the repeated version of this game any quantity between

the Nash and the monopoly quantities can be sustained in equilibrium. Assume, for simplicity, that firms have two options in every period. They can either cooperate and sell a quantity below the Nash equilibrium in the forward market, or deviate and play their static best response to the other player's strategy. Then, if there are only two firms the game becomes a classic Prisoner's dilemma:

	C	D
C	Π_c, Π_c	$\Pi_{cd} \Pi_{dc}$
D	Π_{dc}, Π_{cd}	Π_d, Π_d

where $\Pi_{dc} > P_{i_c} > \Pi_d > \Pi_{cd}$. For simplicity, this assumes firms are symmetric.

In the one shot game, (D,D) is the unique equilibrium, and both firms sell the Nash quantity in the forward market. In the repeated version, they can sustain the collusive equilibrium as long as the probability of the game ending is low enough. While the discounted benefits from future cooperation exceed the benefits of deviating today, a tacit collusive, or cooperative, equilibrium exists.

Assume generators were in an cooperative equilibrium of this kind before the regulatory change. When the market learns that financial participants will face lower trading costs and thus will arbitrage more, the final period of the game they were playing is defined. Starting on the date in which the regulatory change is implemented, collusion will no longer be possible, since financial participants will close any gap created by the generators. Therefore, the cooperative equilibrium unravels as soon as the likelihood of regulatory change becomes high enough, and producers revert to static Nash.

Notice that under a cooperative equilibrium, firms are not playing their static best response. In fact, to play static best response to the residual demand they face is to deviate. Moreover, they act as if they were facing a demand less elastic than they actually face. These observations will be useful to distinguish between alternative models to rationalize the data.

5.2.2 Entry deterrence

Entry deterrence is another possible explanation for generators' anticipatory reaction to the regulatory change. If generators wanted to convince new financial traders not to enter the MISO market, they would lower the forward premium to make market entry less profitable. This strategy

would make sense if there were some link between different periods that made the generators' threat credible. However, generators' decision to lower the premium before implementation of the regulatory change would have had no effect on the incentives they would face if financial traders decided to enter after transaction costs went down. For entry deterrence to be part of an equilibrium, it is necessary that the incumbent be able to commit to "fighting", *i.e.* going against its own interest to lower the entrant's profits. Even though such entry deterrence is unlikely to be part of a long run equilibrium, the empirical analysis will be flexible enough to include it and test whether it could explain firms' behavior.

Notice that under this hypothesis, firms do not play their static best response either. Instead, they act as if the market were more competitive than it is, in order to discourage financial participants from entering.

5.3 Best response deviation

Define the Best Response Deviation (BRD) as follows:

$$BRD \equiv p^F - p^S - [Q(p^F) - x^F] \frac{1}{|R'(p^F)|} \quad (12)$$

The BRD is the difference between the two sides of the optimality condition described by Equation 10, which implicitly defines the static best response function for a firm. The sign of the BRD can be used to distinguish between the different models that can rationalize the observed generators' behavior.

If the static model is a good representation of the firms' behavior, $BRD = 0$ because the bids satisfy the optimality condition in Equation 10. If firms are in a collusive equilibrium, they will act as if the elasticity of the residual demand were smaller than it is, *i.e.* as if the market were less competitive than it is. In that case, $BRD > 0$ because firms will choose a markup larger than what is best given the elasticity of demand they face. To see this, Equation 10 can be rewritten as follows

$$\frac{p^F - p^S}{p^F} = \frac{Q - x^F}{Q} \frac{1}{\eta} \quad (13)$$

where Q and η are functions of p^F .

Finally, if firms' behavior can be explained by entry deterrence, firms will act as if the market were more competitive than it actually is. This means they will choose a smaller markup than the elasticity of their residual demand implies, and $BRD < 0$.

The different hypotheses have different predictions regarding the evolution of the BRD over time as well. If the market is in a static game equilibrium, the BRD should not change over time. If this is the case, the observed anticipatory reaction of the generators would have been caused by changes in the contract positions or the demand, and the firms would have been playing their static best response to the market conditions they faced at the moment.

If firms were in a collusive equilibrium that broke with the announcement of increased competition in the future, the BRD would be initially positive, and then move toward zero after the announcement of the regulatory change. How close to zero it ends up depends on financial bidders' effectiveness in arbitraging the forward premium. The speed of the adjustment depends on the pace at which the collusive equilibrium breaks. The adjustment could happen all at once when the market learns about the future change, or gradually, beginning at the time of the announcement and finishing when financial trading increases.

Finally, under entry deterrence, the BRD is expected to start at zero before the announcement, become negative when the market learns about future competition, and increase towards zero when generators feel safe from the threat of entry. It is not clear when this last step would happen, as financial participants can increase their trading as soon as generators open the gap enough to make virtual trading profitable.

6 Empirical Strategy

In this section, I describe my method for estimating each of component of the best response deviation (BRD) calculation, which is described in Equation 12. These elements are (1) the elasticity of the residual demand, (2) the expected spot price, and (3) the forward contract position.

6.1 Residual demand and its elasticity

In principle, the residual demand for each generator could be computed directly from the data by just adding up the demand bids and subtracting the supply bids. However, this is not always a close approximation in a market with nodal pricing. When transmission lines are at capacity, the set of generators and physical buyers that enter a given generator's residual demand is a subset of the MISO market. Determining that subset is therefore crucial to correct computation of the residual demand.

Market definition

I define markets using a machine learning technique called hierarchical clustering. In general, clustering techniques group elements of a set into groups or clusters, based on a predefined notion of similarity. The number of clusters is generally determined exogenously.

In hierarchical clustering, each element is initially its own cluster.³³ The first step is to merge the two most similar objects into one cluster, according to the similarity measure. In each of the following steps, the two most similar elements or clusters are joined into one cluster. There are several ways to compute the similarity between two clusters; I use the distance between the centroids of the cluster.³⁴ Figure 7 illustrates the output of the algorithm for the case of 5 elements (nodes in this case).

In my analysis, I use the price correlation between nodes as the similarity measure for the clustering algorithm, since two nodes that belong to the same market should have the same price. Prices among nodes can differ because of congestion and losses, so both need to move together for two nodes to be in the same market. Although in principle it is possible that two nodes that are geographically far from each other have correlated prices, this would only happen if both the congestion and line loss components coincide. Figure 1 shows a heat map of prices in MISO in two different moments. Nonadjacent areas do not seem to have the same color in the two maps, making high price correlation between geographically separate nodes unlikely.³⁵

³³This is the agglomerative algorithm. In the divisive algorithm, all elements start together in one single cluster, and each step splits the most different elements.

³⁴I also tried using complete-linkage clustering, which defines it as the maximum distance between elements of each cluster, but the fit was worse.

³⁵I cannot verify that only prices from geographically adjacent areas are correlated because I do not observe the nodes' geographical location, since it is considered a matter of national security. This is why I present suggestive

The main source of uncertainty in this problem are physical contracts among market participants, which do not affect market clearing prices or quantities, but do use the transmission network. Therefore, they affect flows and network congestion for a given set of observed bids. For this reason, I run the market-definition algorithm over periods in which firms' contractual obligations remain constant. I define markets separately for each month, year, and hour of the day. For instance, I take the prices for all nodes during hour 5 of September 2011 and compute the correlation matrix. I then use these correlation data to define markets for the hour between 5 a.m. and 6 a.m. of September 2011.³⁶

The hierarchical clustering algorithm returns a set of potential market definitions, one for each step of the algorithm. For instance, in Figure 7 there are 5 potential market definitions; there could be only one market $\{1,2,3,4,5\}$, or three markets $\{1\},\{2\},\{3,4,5\}$, etc. Generally, there is no appropriate measure of fit for the clusters, and it is not clear which number of separate markets best represents the data. To remedy this uncertainty, I use bid data to test the ex-post fit of alternative market definitions. To do this, I take a market definition (e.g. 3 markets) and clear each of the market clusters by adding up the demand and supply bids submitted at the nodes belonging to each cluster. For instance, to evaluate the market definition with three clusters, I clear market 1 by crossing aggregate demand and supply bids at node 1. To clear market 3 I add up demand and supply bids from nodes 3, 4, and 5 to obtain aggregate supply and demand, and then clear the market. This process results in a simulated clearing price and quantity for each market under each market definition, which can be compared to the clearing prices and quantities observed in the data.

The difference between the observed and simulated clearing prices and quantities for each market definition is then regressed on a constant to test the null hypothesis that this difference is zero. This is done with both an OLS and a quantile regression for the median.³⁷ All market definitions for which the null is rejected are discarded. The rest are kept, even if there is more than one for each hour, because the different definitions are used to run robustness checks. The

evidence only.

³⁶I also tried accounting for day of the week effects, caused by contracts to deliver electricity during weekdays, for instance. I fed the clustering algorithm the residuals of a regression of prices on day of the week dummies. I do not use those definitions because the resulting fit, as defined below, was bad.

³⁷If a market does not clear in the simulation, because demand's maximum willingness to pay is smaller than supply's minimum price, I assume the cleared prices was 0.

mean difference in prices is below 10% for all hours, and below 5% for the majority of them. Because the inelasticity of demand makes quantities much less variable than prices, I use only price deviations when selecting market definitions.

For some hours and months, the difference between the observed and simulated cleared prices is statistically different from zero for all market definitions. When this is the case, I exclude that hour from the sample. This happens with for at least one month in each of the night hours (hours 0, 1, 2, 3, 4, 5, and 23).

Figure 8 shows an example market obtained using this method. For the hour between 6a.m. and 7a.m. of January 2011, the best fit was obtained defining 17 markets. The plot shows the simulated demand and supply, as well as the clearing prices and quantities, for market 2. As it can be observed, the simulated price and quantity match the observed ones very closely. In this market, there are 37 buyers and 7 sellers. Although it seems not so concentrated, the largest seller controls 50% of the generating capacity, and the next two 20% each. Additionally, this highlights the importance of market definitions, since these firms would not be described as having market power if the market included every firm in the MISO footprint.

This method to define markets is an approximation, because in reality all nodes in the MISO market can affect each other's price. The fact that the simulated clearing price is, on average, not far from the observed one indicates that the ex-post fit of these definitions is good. As long as market conditions remain constant within the month, these definitions can also be used ex-ante to represent generators' rational beliefs about the residual demand they will face.³⁸

Using the generators' physical locations to group nodes into markets may appear to be a simpler way to define markets. However, MISO does not include generators' locations in the dataset because this information is considered a matter of national security. Even if I could obtain location information, it is not possible to infer which firms compete with each other without having more information about the transmission network's capacity. As Figure 1 shows, neighbor nodes may have very different prices and thus, belong to different markets. Finally, even if I had all of the relevant data, I would need to solve a complex optimization problem

³⁸An alternative way to understand the generators' problem is to think they face a distribution of potential markets in which they may be competing each day, where market means group of competitors or residual demand. My exercise allows to compute the empirical distribution of markets by assuming the realization is the market definition with the lowest deviation from observed cleared prices and quantities in the data, and then use this distribution to estimate the generators' best response. This is something I am planning to do in future research.

multiple times for each generator in order to estimate residual demand. This process would be computationally demanding, and most likely far from what firms actually do when they make bidding decisions.

Zheng (2014) also uses clustering tools to define markets in her estimation of an entry game between discount retailers. She splits the market into independent submarkets to lower the computational burden and make the estimation of the model possible. Additionally, the richness of the electricity dataset allows me to define markets and a measure of fit that do not depend on the model of firm behavior.

Residual demand and its elasticity

Because of the richness of the data, market definitions are all is needed to obtain the residual demand faced by each firm. Since I observe every demand and supply bid submitted, I can construct the residual demand faced by each firm simply by adding up demand bids and subtracting the competitors' supply bids. A residual demand is defined for each firm in each market, which is assumed to be the information that each firm uses to make decisions.³⁹

Seventy-five percent of the bids and 82% of the megawatts cleared by generators are to piecewise linear bids, while the rest are step functions. I convert these piecewise linear bids into step functions by splitting them into 0.1 MW increments. As a consequence, residual demand is expressed in step functions with very small steps, which the derivative to be computed by calculating the difference between one step and the next and dividing it by the size of the step. I also fit a cubic spline to the resulting residual demand, and take the derivative to compute the elasticity.

6.2 Expected spot price

The optimality condition for the bid in the forward market in Equation 10 is pointwise optimal, and it is therefore written in terms of P^S instead of the expected value of P^S . For the sake of robustness, I compute the value of the contracts and estimate the best response deviation (BRD) using the expected value of the spot price.

³⁹This assumes that decisions are taken independently by a same company in different markets. Although it seems a strong assumption, given that markets are independent there would not be any gain from making the decisions jointly.

I compute the expected spot price at each node under three different assumptions about agents' expectations. First, it is assumed that agents have rational expectations: generators are forward-looking and use all available information to predict the spot price. To estimate the expected value I run the following regression with data from the prior month

$$p^S = \alpha + \beta_1 f_p^S + \beta_2 f_q^S + \beta_3 p_{lag}^S + \beta_4 p_{lag}^F + \varepsilon \quad (14)$$

where f_p^S is the price and f_q^S the traded volume of the futures for the Indiana hub in peak hours traded in the Intercontinental Exchange (ICE). These are spot price futures traded one day before the underlying production date, so their prices are almost identical to the forward price. p_{lag}^S and p_{lag}^F indicate the lags of the spot and forward price, respectively. The lags used are one, two, and three days before for the same hour and the previous one, plus the price in the previous 12 hours. The same lags are used for the forward and spot prices.

I estimate the coefficients of Equation 14 using data for the month preceding t , the day for which I want to predict the price. Then I predict the spot price for day t using data on day $t - 2$, as bids are submitted on day $t - 1$, while markets for that day are still clearing.

The second assumption is adaptive expectations, under which generators are backward-looking and expect the spot price to continue as it has been. The expected spot price is computed using the average of the price during the last three days. Since the expected spot price affects the amount producers decide to sell in the forward market, I compute it using information available during the hours before the forward market closes. For this reason, the spot price predicted under the adaptive expectations assumption is the average of the last three days preceding the day for which bids are being submitted. The last assumption is perfect foresight, under which generators know exactly what the spot price will be. This is equivalent to using the observed price instead of an expected value.

Table 5 describes the difference between both rational and adaptive expectations and the observed spot price. Although the predictions are not unbiased, on average they are not too far from the spot price. Note, also, that the estimated expected spot prices are much closer to the spot price than the forward price. This observation can be interpreted as a sign of informational efficiency, since in an efficient market the forward price would be the best predictor of the spot price.

6.3 Hedging contracts

I back out the hedging contract position held by each generator from the optimality condition in Equation 8, as in Hortaçsu and Puller (2008). I rewrite the equation as follows for ease of explanation

$$p^F - p^S = -[Q^*(p^F) - x^F] \frac{H_Q}{H_p}$$

The optimal schedule for the forward market is such that when the forward and spot prices are the same, the total quantity offered by each generator in the forward market equals its forward contract quantity, *i.e.* $Q(p^S) = x^F$. From this equation, I obtain the contract position for each generator in each market.

Although this is a condition for pointwise optimality, for robustness I back out the forward contract positions using the expectation of the spot price. I use three estimates of the expected real-time price that correspond to the three different assumptions about how agents form expectations described in the previous sections.

The second estimate assumes agents are not forward-looking and expect the price to be a simple average of the prices observed in the past. The third estimate assumes agents have perfect foresight and base their contract positions on the actual future, *i.e.* $E_{t-1}(p_t^S) = p_t^s$.

The forward hedging contract position can be correctly backed out when the optimality condition holds, *i.e.* under the null hypothesis of static Nash equilibrium. As the hedging positions are correct under the null, the test is valid.

6.4 Computation and analysis of the best response deviation (BRD)

The best response deviation is defined in Equation 12 as the difference between the two sides of the optimality condition for the forward schedule submitted by a generator. I rewrite it here to make exposition easier:

$$BRD = p^F - p^S - [Q(p^F) - x^F] \frac{1}{|R'(p^F)|}$$

The previous sections have shown how to compute each of the components of the BRD: the forward hedging position, the derivative of the residual demand, and the expected spot price.

Each of these elements is obtained separately for each hour, and both the residual demand and hedging position can be obtained for each individual generator. As there are many nodes in each market, each with a potentially different clearing price, I define the market price as the quantity-weighted average. With all these elements, I can build a panel in which I observe the BRD for each generator during each hour in which she was active in the market.

To analyze the evolution of the BRD over time, I define three time periods according to market events related to the change in RSG charges. These periods are the following:

- **Before:** the four months prior to December 1, 2010. On that date, MISO announced that it submitted a proposal to FERC for the redesign of RSG charge and the market began to prepare for the expected implementation of the proposal. The before period therefore provides data about baseline market conditions.⁴⁰
- **Transition:** the four months between December 1, 2010 and April 1, 2011, the date on which the change was implemented. During this period, the market knew the regulatory change was likely to occur, but it had not yet been implemented.
- **After:** the four months between April 1, 2011 and July 31, 2011. This period represents the first four months after the RSG charges were lowered. There were two major events in July 2011: (1) renewable plants became dispatchable, meaning they could be started and stopped by the market operator according to demand like any other plant, and (2) a large producer firm left MISO to join the PJM Interconnection, which serves a market adjacent to MISO's. The latter event changed the market structure because the firm's transmission lines were transferred to PJM as well. Results are robust to the removal of July.

6.5 Market power in the spot market

Although financial participants decrease generators' market power in the forward market, because they cannot intertemporally price discriminate, they do not eliminate it altogether. Rather, firms retain the ability to withhold production in the spot market in order to drive up the spot price. This is analogous to an instance where increased arbitrage forces a monopolist

⁴⁰Training sessions to explain market participants how these costs were going to be computed started in January. A group of market participants was in charge of the redesign proposal. In January, they wrote they expected it to be implemented on March, 2011.

to stop price discriminating between two sets of consumers. Just as the monopolist’s new uniform price will be higher than the original price in the low-demand market, electricity generators will use their market power to raise spot prices after arbitrage decreases the forward premium.

I examine the effect of increased arbitrage on market power looking at firms’ spot-price markups. I back out the spot markup for each firm from the optimality condition in Equation 11. I assume that firms’ hedging position in the spot market is 0, because firms generally hedge with respect to the forward market, since that is where they sell the bulk of their production. For expositional clarity, I rewrite Equation 11 here for the case without hedging

$$p^S - c' = -[S^*(p^S) - Q^*(p^F)] \frac{1}{R'(p^S)}$$

I estimate the markup in the spot market from the right hand side of this equation. I observe the cleared quantities in the forward and spot markets, and estimate the residual demand as I do for the forward market: I define markets using hierarchical clustering of spot prices, and then add up the corresponding bids to compute the residual demand faced by each firm.

Notice that the output from this estimation is the markup with respect to the firm’s opportunity, which is not necessarily equal to actual production cost.⁴¹ Therefore, I cannot quantify the change in production costs using this mechanism. Nevertheless, this equation is sufficient to determine firms’ market power, since the firms use marginal opportunity cost, rather than physical cost, to make their bidding decisions.

7 Results

7.1 Market definitions

7.2 Best response deviation (BRD)

I look at the evolution over time of the best response deviation defined in Section 5.3 and computed according to Section 6.4. To do this, I run the following regression of the BRD on the

⁴¹For instance, hydro plants decide when to sell based on the opportunity cost of using their reserves, the actual cost of production being zero. Additionally, my exercise does not account for complementarities across hours, which also move opportunity costs away from production costs.

time periods defined in Section 6.4

$$BRD_t = \alpha_0 \text{before} + \alpha_1 \text{interim} + \alpha_2 \text{after} + \varepsilon_t \quad (15)$$

where BRD_t is the mean best response deviation for each hour, weighted by the size of the firm. As I have a different BRD for each market definition, and more than one market definition for some hours, I use the market definitions that are most prevalent for each hour as a baseline. For instance, if for hour 7 during February 2010 the market clears well with either 4 or 5 markets, I choose the definition that clears well in more months for hour 7. That is, I count the number of months for which each market definition is a good fit, and for each month select the one with the highest number.

I run the regression in Equation 15 under two specifications of the BRD, one for each method for estimating the derivative of the residual demand. The residual demand is always downward sloping for all firms, but in certain instances the cubic spline method yields a positive-sloping residual demand. I discard these observations. Furthermore, both BRD estimation methods produce a few extreme values that have a disproportionate effect on the results. For this reason, I remove the top and bottom 1% of my observations. I only report results obtained using the cubic spline method for estimating the derivative, but the two methods produce fairly similar results. In order to avoid effects from monthly market fluctuations, I compute the BRD's monthly mean using data on 2010 and 2011. I then remove the month fixed effect from the BRD, and add the mean of the month fixed effects to get a more accurate level.

Results from the regression in Equation 15 are presented in Table 6. The tables show separate columns for the three assumptions about spot price expectations that are used to compute contract position and best response deviation: rational expectations, adaptive expectations, and perfect foresight. Note that it is necessary to add the monthly mean to each number to calculate the baseline BRD.

Table 6 shows that the BRD was positive in the initial “before” period, and then decreased in the next two periods. This indicates that generators reacted to the regulatory change before financial trading actually increased, price discriminating less before arbitrage increased. That is, they decreased their price discrimination before arbitrage increased. Additionally, a positive best response deviation indicates that entry deterrence is not a good explanation for the generators'

anticipatory reaction. I cannot reject a cooperative equilibrium, since evidence is consistent with the presence of a tacit collusive agreement that broke around the time of the announcement.⁴²

Peakers are units that can be started quickly, although typically at a high marginal cost, and they are used to cover last minute increases in demand. Therefore, they are very likely to produce when demand in the spot market exceeds production scheduled in the forward market. Including them in the analysis may add effects coming from technical characteristics instead of from firm behavior. For this reason, I run the same regression excluding peakers from the sample. Results are presented in Table 8, which shows that under rational expectations the Best Response Deviation goes to zero already in the interim period.

Table 7 shows results from running the same regression separately for large and small firms, where large firms are those in the top quintile by production capacity, which together control 80% of production. For small firms, the regression coefficient for the “interim” period is far less significant. This is informative even if the coefficient for the period after implementation is significant because the best response deviation is not as meaningful for firms without market power. Typically, the markup is a good measure of market power because costs are firm specific. In this case, however, the markup is the same for all firms in the market because they all face the same spot price. For this reason, it will tend to overestimate the real value for firms without market power. With the whole sample this is not a big problem because I am weighting by size, but when only the small firms are included, it is a reason to be careful when interpreting the results.

The elasticity of the residual demand faced by the firms increased after financial trading increased, as Table 9 shows, though the effect is not very significant. The sign goes in the expected direction, as increased financial trading reduces generators’ market power.

Table 10 shows the evolution of the backed-out contract positions over time. The change is not significant during the period between announcement and implementation, but the negative sign indicates that firms hedged somewhat less after the announcement. This means that firms were more exposed to the forward price, and therefore had more incentive to engage in price discrimination. Therefore, it is unlikely that the observed change in generators’ behavior came from changes in their hedging contract position.

⁴²Fabra and Toro (2005) find evidence of collusion in the Spanish electricity market, although they observe price wars together with periods of price stability.

7.3 Forward premium

As expected, the forward premium decreased after the announcement as well. Table 12 shows the results of a regression of the forward premium on time-period dummies, using node and month fixed effects. As it is clear from the results, it followed a pattern similar to the best response deviation, decreasing when generators behaved more competitively in the forward market.

7.4 Test of additive separability of the bids

The empirical strategy in this paper relies on the assumption of additive separability of the optimal bid in the hedging contract position and the price. If this assumption holds, changes in the contract position will shift the bid without affecting the slope. I follow Hortaçsu and Puller (2008) and use the data to test the assumption, which is described in section 5.1. The test evaluates whether the slope of the bids changes with variations in the contract position. Under additive separability, contracts should only cause parallel shifts in the bids, with no effect on the slope.

I fit a linear function to the submitted bids to obtain their slope; the fit is around 68%, a decent approximation. I then regress the slope of the bid on the hedging contract position obtained as explained in Section 6.3. The first three columns of Table 11 present the results of this regression, using firm-market fixed effects. The three columns correspond to the three assumptions about the expected spot price that are used to back out the contract position for each firm. The correlation between the slope of a firm's bid and its contract position is not statistically significant, which supports the additive separability assumption.

Because the optimal bid submitted depends on the other players' strategy, I add the slope of the residual demand faced by each firm as a control. I also control for the spot price, since it is the opportunity cost of bidding in the forward market. After controlling for these factors, the forward position is still not significantly correlated with the slope of the bids, as the last three columns of Table 11 show.

7.5 Spot market markups

I examine the effect of increased financial trading on the spot market by estimating the following regression:

$$\widehat{P^S - c'(S)}_{i,t} = \beta_i + \alpha_0\text{before} + \alpha_1\text{interim} + \alpha_2\text{after} + \varepsilon_{i,t}$$

where the left hand side is the estimated markup in the spot market, obtained as explained in section 6.5. The subindex i indicates a firm in a particular market, so the β_i 's are firm-market fixed effects.

Estimation results are presented in Table 15, which indicates that spot markups increased when generators decreased their price discrimination. Like the markup, the coefficients for the time-period dummies are measured in dollars, which implies that markups increased between 4 and 10 cents, which is very small relative to the average spot price of approximately \$30 (see Table 18).

In principle, we would expect this spot markup increase to come mainly from large firms, as they have more market power. For this reason, I split the sample between large and small firms and run the same regression separately for each group. I define large firms as those with a capacity above 1100 MW, which includes the largest 20% of firms, which collectively control 80% of the production. Results from this exercise are presented in Table 16. I find that the change in markup is bigger for the large firms, which is consistent with the hypothesis that large firms had more market power after price discrimination became infeasible. Relative to the baseline, these firms increased their spot markups by 7 to 15 cents in the interim period between announcement and implementation, and between 8 and 16 cents in the period after implementation.

The mean fixed effect in these regression is negative, which could be a concern. Nonetheless, the mean markup is \$1.081 per megawatt hour when the derivative of the residual demand is computed using a cubic spline, and \$2.24 when computed by taking the simple ratio of changes. The median markup is 0 in both cases. The presence of a good number of negative markups is most likely coming from omitting bid complementarities between hours, as shown by Reguant (2014). Table 17 presents results from the same regression as above, but using a sample that only includes firms that are net sellers in the market, because these are the firms whose behavior

is consistent with price discrimination. As expected, markups are larger using this sample, and increase significantly starting after the announcement of the regulatory change.

7.6 Welfare analysis

The effects of increased financial trading on welfare are ambiguous. Consumers are better off because the reduction of the forward premium means that they pay less and total quantity does not change, since final demand is perfectly inelastic. Producers, as a group, are worse off because they are forced to charge a uniform price, instead of acquiring surplus from price discrimination, as they did initially. Nonetheless, the final effect is not just a transfer from producers to consumers, because costs vary depending on whether the same total quantity is produced by the lowest cost firms.

On the one hand, production costs may increase because generators exert more market power in the spot market. As shown in the previous section, generators markups in the spot market increased when arbitrage did, which implies that firms withheld a larger part of their production in this market. As demand is perfectly inelastic, this will not decrease total production but will shift it toward firms with higher costs whose bids would not clear under more competitive conditions. Thus, production cost increases as a consequence of arbitrage. This finding is consistent with that of Ito and Reguant (2014), who reach a similar result by investigating a counterfactual in which financial traders are introduced to the Iberian market.

Although firms' increased exertion of market power in the spot market leads to productive inefficiencies, arbitrage may also decrease production costs by improving scheduling in the forward market. Underscheduling of production in the forward market tends to increase costs for two reasons. First, because more expensive producers will be scheduled in the forward market. Second, because some units will need to be dispatched in the spot market, and the units that can react on short notice often have higher marginal costs. Jha and Wolak (2013) compare production costs and carbon emissions before and after the introduction of financial traders into the California market, and find that both decreased. This mechanism reduces production costs, implying that it is possible for productive efficiency to stay the same or increase when arbitrage goes up.

Overall, the effect on costs is ambiguous. Productive efficiency will increase only if the decrease

in costs from better scheduling in the forward market is larger than the increase in costs from generators exerting more market power in the spot market. Given that 98% of the energy sales happen in the forward market, meaning that most production is scheduled in advance, total costs are more likely to decrease when generators cannot engage in price discrimination. A precise quantification of this effect requires cost data and is therefore left for future research.⁴³

Consumers are unambiguously better off, however, since they pay less for their electricity purchases. To quantify this, I look at changes in total expenditure per MWh over time. After controlling for fuel prices and the forecasted demand level, total expenditure decreased in the period between announcement and implementation of the regulatory change and stayed below the initial level after implementation, as Table 14 shows. The coefficients indicate that total expenditure was 4% higher before the announcement than after implementation. Given that total demand is 1,500,000 MWh a day on average, and the price is around \$30 per MWh, this means that consumers save about \$1,800,000 per day on average. Note, however, that it is important to control for demand and fuel prices. Simply looking at changes in total expenditure, without controls, would indicate that, relative to the period after implementation, total expenditure was 10% lower in the “before” period, and 10% higher in the “interim” period.

Consumer savings come from two sources. The first is the direct effect of financial traders on generators’ ability to engage in price discrimination in the forward market. The second mechanism is the change in the dynamic equilibrium. The evolution of the best response deviation over time indicates that firms’ initial conduct was consistent with more market power than they had, and that after the change, increased arbitrage pushed their conduct closer to the static Nash equilibrium. This effect can be roughly quantified by multiplying the change in the BRD, which is measured in dollars, by the average daily load. Using the lowest estimate for the change in the BRD in the after period, this calculation yields an average savings of about a million dollars a day (\$0.70 times 1,500,000 MWh). This indicates that about half of the reduction in consumer cost is attributable to firms reverting to a static Nash equilibrium.

⁴³I am building a dataset that will match the data published by MISO to plant characteristics. Because the relationship between electricity generation and fuel consumption is stable for each plant, marginal costs can be computed from data on plants’ technical characteristics and fuel prices. This exercise will allow me to determine to which extent demand was covered by more expensive plants, and its consequence on costs. Nonetheless, it will be a lower bound on costs’ decrease from better scheduling in the forward market, since the analysis of marginal costs does not take into account complementarities across hours.

8 Conclusion

This paper studies competition and the role of financial players in electricity markets. I examine a regulatory change that exogenously increased virtual trading and find that financial players made the forward market more competitive. This benefited consumers, but may have reduced productive efficiency because large firms exerted more market power in the spot market. Additionally, my findings indicate that generators were in a tacit collusive equilibrium before the regulatory change, and that cooperation broke as soon as firms learned that traders were going to enter in the future. In fact, generators became more disciplined before the regulatory change was actually implemented, which highlights the importance of dynamic considerations when assessing the impact of financial traders.

Figures

Figure 1: Price dispersion Heat map of prices across the MISO market on September 7, 2011 and April 10, 2012. Prices may differ significantly in a given moment, and over time.

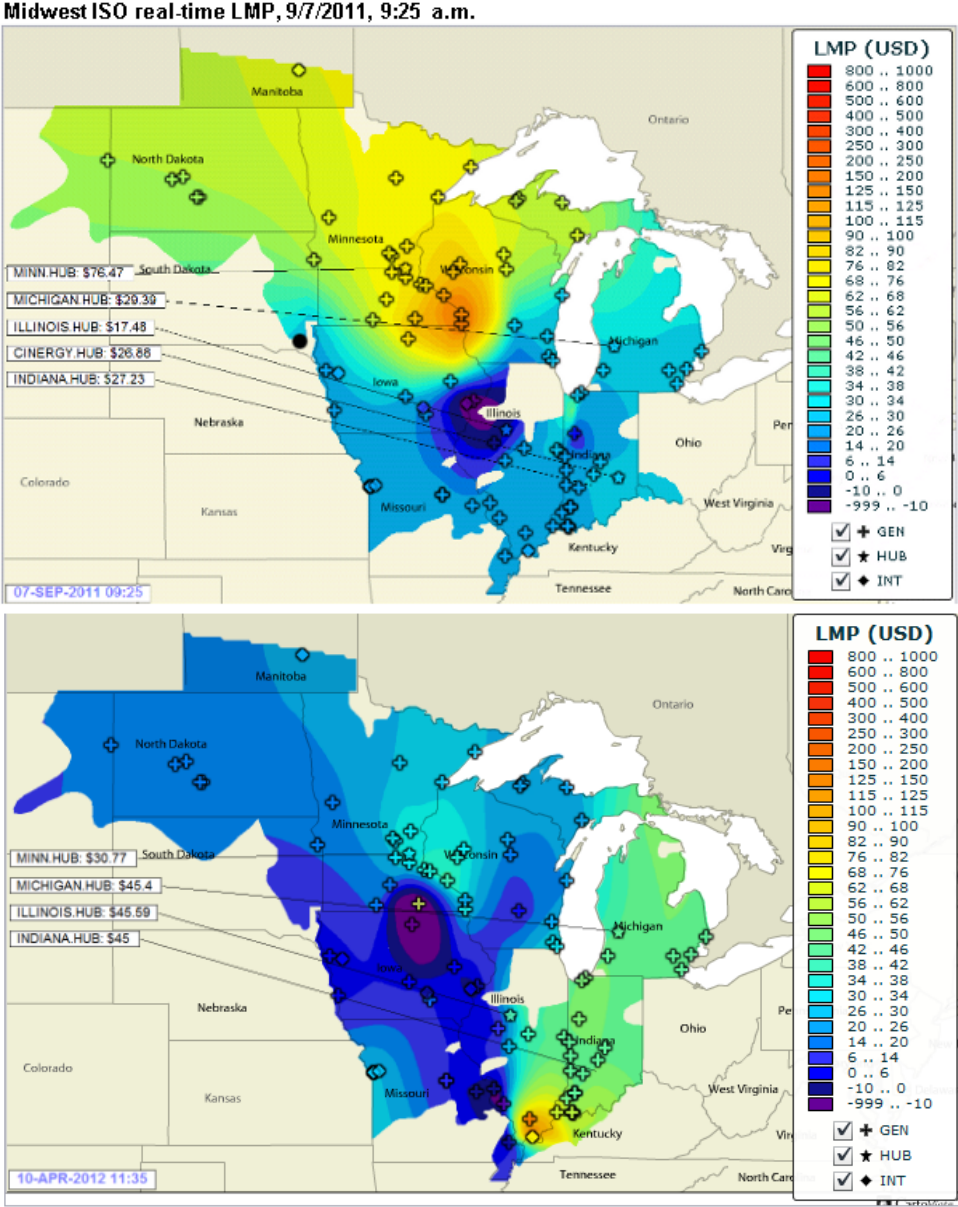


Figure 2: Forward premium over time




Figure 3: Price-sensitive supply bids




Figure 4: Virtual trading over time The green line indicates the monthly average of the daily volume traded by virtual bidders. The first dashed red line indicates the date in which the proposal to redesign RSG charges was submitted to FERC on December 1, 2010. At this point, the market started preparing for the implementation of the new computation proposal, and to explain firms how it was going to work. The solid red line on April 1, 2011 indicates the moment in which the RSG change was actually implemented.




Figure 5: Spot sales over time The green line indicates the monthly average of the daily difference between the quantity cleared in the forward and spot markets. The first dashed red line indicates the date in which the proposal to redesign RSG charges was submitted to FERC on December 1, 2010. At this point, the market started preparing for the implementation of the new computation proposal, and to explain firms how it was going to work. The solid red line on April 1, 2011 indicates the moment in which the RSG change was actually implemented. The structural break occurred on January 10, with a confidence interval between January 5 and January 15.

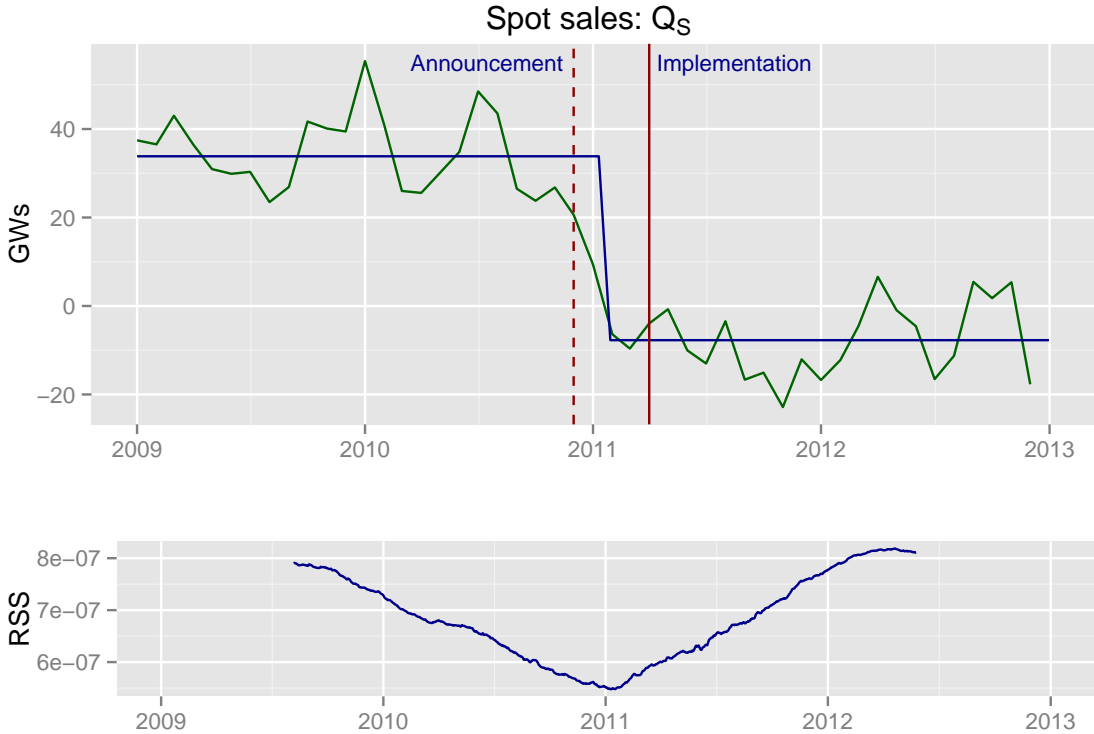


Figure 6: Load cleared in the spot market The green line indicates the monthly average of the daily difference between the quantity cleared in the forward and spot markets. The first dashed red line indicates the date in which the proposal to redesign RSG charges was submitted to FERC on December 1, 2010. At this point, the market started preparing for the implementation of the new computation proposal, and to explain firms how it was going to work. The solid red line on April 1, 2011 indicates the moment in which the RSG change was actually implemented. The structural break occurred on January 26, with a confidence interval between January 20 and February.

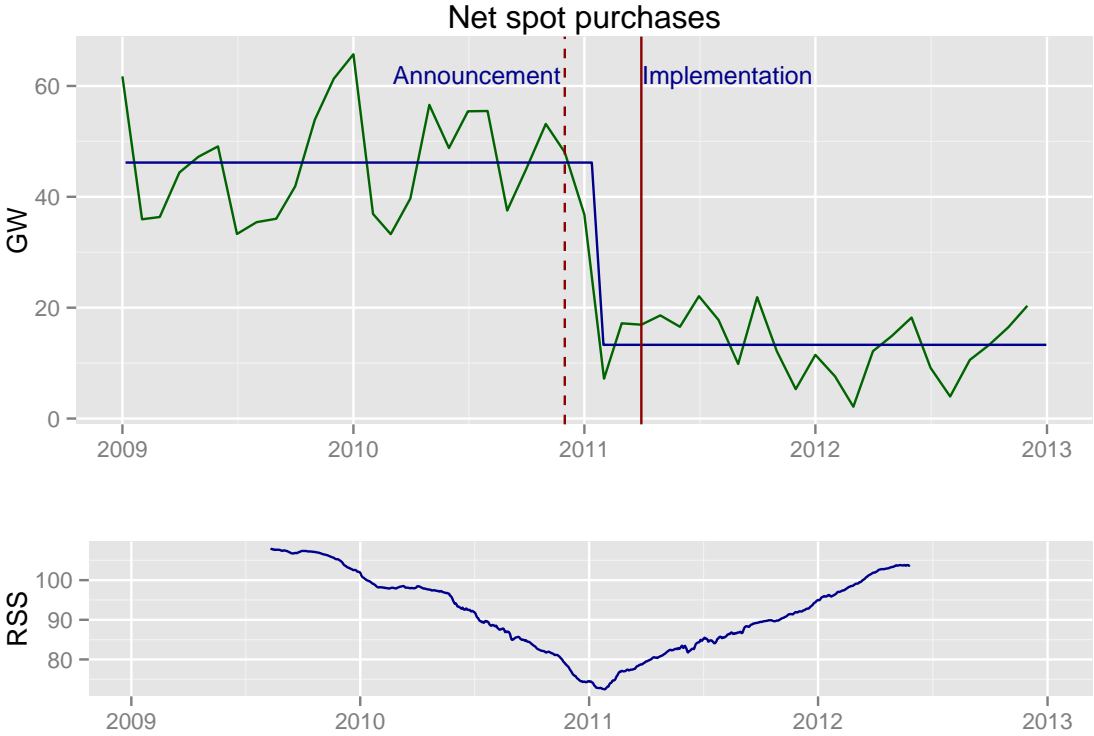


Figure 7: Dendrogram to illustrate hierarchical clustering

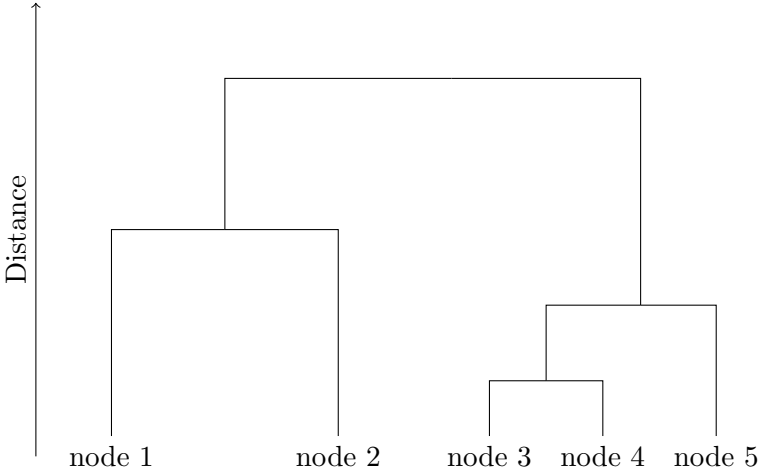


Figure 8: An example of a market The figure shows demand and supply in one of the markets as defined according to hierarchical clustering for January 2011 between 6a.m. and 7a.m. The best fit was found when there are 17 markets. This is market 2 and there are 37 buyers and 7 sellers in it. The largest seller holds 50% of the generating capacity, and the next two hold 20% each.

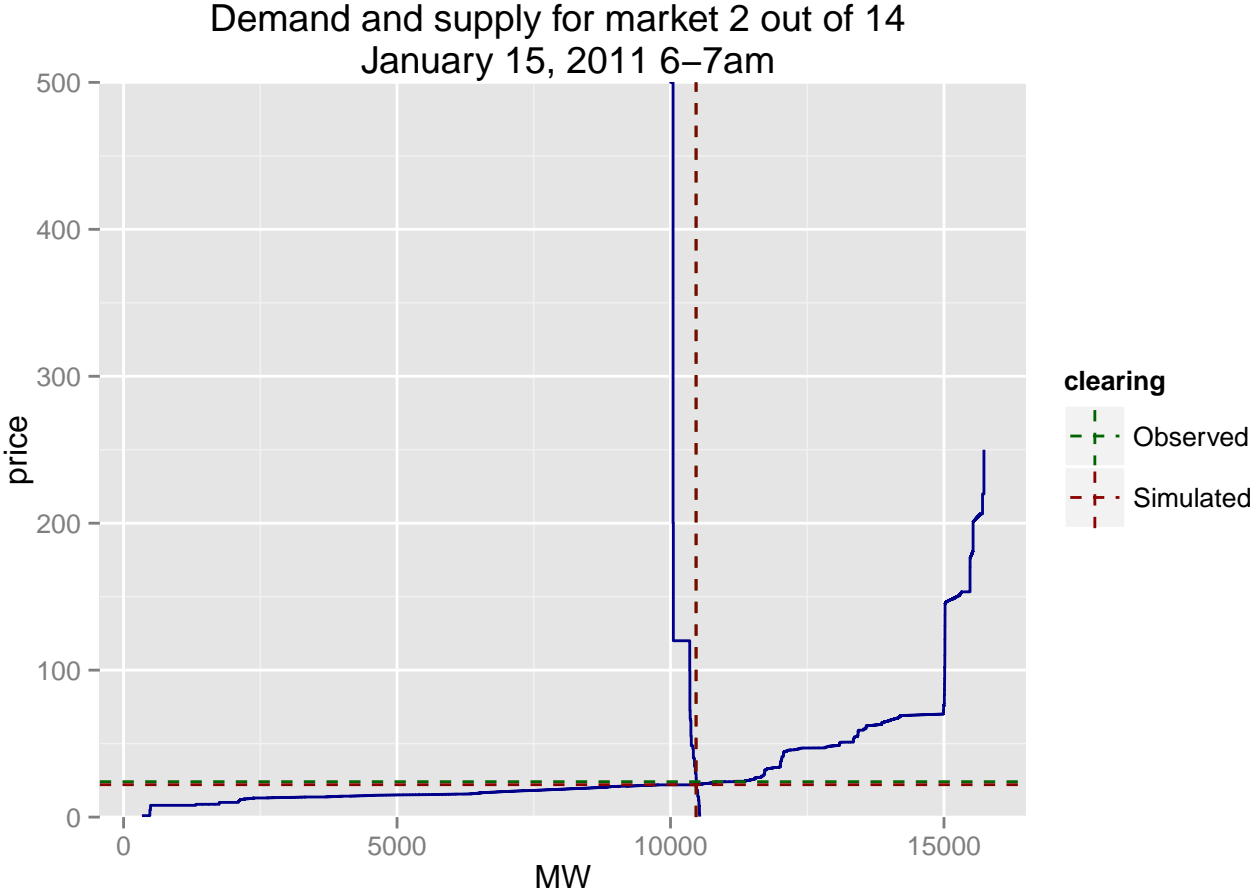




Figure 9: Additive Separability Differences across the bids for a given firm seem to be parallel shifts.



Tables

Table 1: Summary statistics for demand bids Each variable is computed daily. For instance, the number of bids is the total number of bids submitted each day. The sample goes from January 2010 to December 2011.

Statistic	N	Mean	St. Dev.	Min	Max
Price takers					
# bids	730	5,762.379	297.844	5,156	6,299
# nodes	730	228.507	15.725	197	246
# bidders	730	96.155	2.423	90	100
% bids cleared	730	1.000	0.000	1	1
Cleared MW	730	1,478,659	191,083	1,082,308	2,043,150
Price sensitive					
# bids	730	1,015.101	63.5	792	1,152
# nodes	730	42.3	2.7	33	48
# bidders	730	25.2	2.16	18	31
% bids cleared	730	0.901	0.031	0.777	0.985
Cleared MW	730	30,992	5,846	17,030	52,089

Table 2: Virtual bids summary stats Each variable is computed daily. For instance, the number of bids is the total number of bids submitted each day. The sample goes from January 2010 to December 2011.

Statistic	N	Mean	St. Dev.	Min	Max
Virtual Demand					
# bids	730	53,556	18,873	15,240	97,824
# nodes	730	874	274.9	318	1,280
# bidders	730	56.4	6.71	31	77
% bids cleared	730	0.102	0.038	0.028	0.228
Cleared MW	730	86,263	22,058	39,909	161,463
Virtual Supply					
# bids	730	62,313	22,024	16,080	117,384
# nodes	730	993.6	309.4	351	1,378
# bidders	730	50.9	6.34	32	69
% bids cleared	730	0.095	0.032	0.034	0.197
Cleared MW	730	60,983	19,354	23,825	128,022

Table 3: Supply bids in the forward market Each variable is computed daily. For instance, the number of bids is the total number of bids submitted each day. The sample goes from January 2010 to December 2011.

Statistic	N	Mean	St. Dev.	Min	Max
# bids	730	20,717	1,036	18,861	21,886
# nodes	730	927	28.2	883	957
# units	730	1,147	41.172	1,079	1,197
# firms	730	126	4.78	120	132
% bids cleared	730	0.361	0.035	0.288	0.511
Cleared MW	730	1,214,775	162,234	849,110	1,672,726
Price taker MWs	730	163,606	24,390	101,316	212,362
% piecewise linear	730	0.75	0.012	0.72	0.77
MW piecewise linear	730	0.82	0.013	0.78	0.86

Table 4: Supply bids in the spot market Each variable is computed daily. For instance, the number of bids is the total number of bids submitted each day. The sample goes from January 2010 to December 2011.

Statistic	N	Mean	St. Dev.	Min	Max
# bids	730	13,037	1,031	10,607	17,071
# nodes	730	525.2	53.4	432	776
# units	730	603.9	65.1	493	914
# firms	730	100.3	6.24	88	118
% bids cleared	730	0.72	0.027	0.62	0.79
Cleared MW	730	1,447,665	189,301	1,075,636	1,977,326.000
Price taker MWs	730	123,147	27,014	63,248	196,913
% bids piecewise linear	73	0.62	0.03	0.53	0.71
% MW piecewise linear	730	0.81	0.02	0.74	0.86

Table 5: Expected spot prices The first two rows of the table present the mean difference between the expected and the effective spot price, where the expectations are computed assuming rational expectations (RE) or adaptive expectations (AE) as defined in Section 6.2. The third row shows the mean forward premium, and the fourth the mean level of the spot price.

	<i>Dependent variable:</i>	
	2010	2011
$E[P^S]_{RE} - P^S$	0.036*** (0.006)	0.163*** (0.007)
$E[P^S]_{AE} - P^S$	0.001 (0.006)	-0.0004 (0.007)
$P^F - P^S$	1.142*** (0.005)	0.500*** (0.006)
Spot price	31.101*** (0.006)	30.310*** (0.007)
Observations	16,920,576	16,350,480
R ²	0.000	0.000
Adjusted R ²	0.000	0.000
Residual Std. Error	24.145 (df = 16920575)	27.034 (df = 16350479)

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 6: Best Response Deviation Results from a regression of the best response deviation on time-period dummies. The best response deviation is computed as the mean of the BRDs for each hour and market, weighted by the size of the firm. Before computing it, the top and bottom 1% of the sample is removed to avoid extreme values. I control for month effects by defining the dependent variable as the residual from a regression of the BRD on month dummies using 2 years of data, to which I add the mean month fixed effect. In the top table, the baseline is the level of the BRD in the initial period, while the “interim” and “after” coefficients indicate the change with respect to that reference. The sample goes from August 2010 to July 2011, leaving four months before the announcement of the regulatory change, four months between the announcement and the implementation, and four months after implementation. The three specifications correspond to three assumptions regarding expectations over the spot price: rational expectations, adaptive expectations, and perfect foresight. Derivatives are computed using a spline. The bottom table shows the mean BRD during each period.

	BRD over time		
	PF	RE	AE
Baseline (Before)	2.082*** (0.124)	0.848*** (0.115)	1.174*** (0.135)
Interim	-1.086*** (0.171)	-0.107 (0.158)	-0.484*** (0.184)
After	-1.425*** (0.182)	-0.695*** (0.181)	-1.094*** (0.196)
Observations	19,275	19,261	19,275

	Mean BRD		
	PF	RE	AE
Before	2.082*** (0.124)	0.848*** (0.115)	1.174*** (0.135)
Interim	1.020** (0.117)	0.741*** (0.108)	0.723*** (0.123)
After	0.658*** (0.134)	0.122 (0.142)	0.083 (0.141)

Note: *p<0.1; **p<0.05; ***p<0.01

Table 7: BRD large and small Results from a regression of the best response deviation on time-period dummies. The best response deviation is computed as the mean of the BRDs for each hour and market, weighted by the size of the firm. Before computing it, the top and bottom 1% of the sample is removed to avoid extreme values. I control for month effects by defining the dependent variable as the residual from a regression of the BRD on month dummies using 2 years of data, to which I add the mean month fixed effect. In the top table, the baseline is the level of the BRD in the initial period, while the “interim” and “after” coefficients indicate the change with respect to that reference. The standard sample goes from August 2010 to July 2011, leaving four months before the announcement of the regulatory change, four months between the announcement and the implementation, and four months after implementation. The three specifications correspond to three assumptions regarding expectations over the spot price: rational expectations, adaptive expectations, and perfect foresight. Derivatives are computed using a spline. Large firms are the largest 10% of the firms in terms of observed capacity, which have 55% of the capacity. The bottom table shows the mean BRD during each period.

BRD over time						
	PF large	PF small	RE large	RE small	AE large	AE small
Baseline (Before)	2.158*** (0.133)	2.369*** (0.131)	0.810*** (0.126)	0.806*** (0.125)	1.050*** (0.148)	1.229*** (0.147)
Interim	-1.006*** (0.184)	-0.707*** (0.179)	-0.710*** (0.169)	-0.201 (0.175)	-0.541*** (0.197)	-0.199 (0.200)
After	-1.306*** (0.198)	-1.001*** (0.190)	-0.692*** (0.197)	-0.326* (0.194)	-0.680*** (0.216)	-0.824*** (0.210)
Observations	13,257	16,718	13,017	16,619	13,120	16,618

BRD mean						
	PF large	PF small	RE large	RE small	AE large	AE small
Before	2.158*** (0.133)	2.369*** (0.131)	0.810*** (0.126)	0.806*** (0.125)	1.050*** (0.148)	1.229*** (0.147)
Interim	1.178*** (0.126)	1.692*** (0.120)	0.114 (0.112)	0.606*** (0.121)	0.554*** (0.128)	1.056*** (0.135)
After	0.822*** (0.150)	1.360*** (0.138)	0.063 (0.154)	0.453** (0.150)	0.344* (0.158)	0.403** (0.150)

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 8: Best Response Deviation excluding peakers Results from a regression of the best response deviation on time-period dummies. The best response deviation is computed as the mean of the BRDs for each hour and market, weighted by the size of the firm. Before computing it, the top and bottom 1% of the sample is removed to avoid extreme values. I control for month effects by defining the dependent variable as the residual from a regression of the BRD on month dummies using 2 years of data, to which I add the mean month fixed effect. In the top table, the baseline is the level of the BRD in the initial period, while the “interim” and “after” coefficients indicate the change with respect to that reference. The standard sample goes from August 2010 to July 2011, leaving four months before the announcement of the regulatory change, four months between the announcement and the implementation, and four months after implementation. The three specifications correspond to three assumptions regarding expectations over the spot price: rational expectations, adaptive expectations, and perfect foresight. Derivatives are computed using a spline. The sample excludes peakers. The bottom table shows the mean BRD during each period.

	BRD over time		
	PF	RE	AE
Baseline (Before)	1.900*** (0.110)	0.811*** (0.101)	1.200*** (0.120)
Interim	-1.102*** (0.143)	-0.695*** (0.134)	-0.609*** (0.154)
After	-1.415*** (0.155)	-0.791*** (0.153)	-1.066*** (0.169)
Observations	23,287	22,943	23,073

	Mean BRD		
	PF	RE	AE
Before	1.900*** (0.110)	0.811*** (0.101)	1.200*** (0.120)
Interim	0.810*** (0.092)	0.118 (0.087)	0.615*** (0.095)
After	0.487*** (0.110)	0.005 (0.117)	0.130*** (0.120)
Observations	23,287	22,943	23,073

Note: *p<0.1; **p<0.05; ***p<0.01

Table 9: Residual demand elasticity Results from a regression of the hourly average elasticity of the residual demand, weighted by firm size, on time-period dummies. The top and bottom 1% of the sample was removed to avoid extreme values. The baseline level is the level in the initial period, while the “interim” and “after” coefficients indicate the change with respect to that reference. I control for month effects by defining the dependent variable as the residual from a regression of the residual demand elasticity on month dummies using 2 years of data, to which I add the mean month fixed effect. The derivatives are computed both using a spline and the ratio of differences using two points. The standard sample goes from August 2010 to July 2011, leaving four months before the announcement of the regulatory change, four months between the announcement and the implementation, and four months after implementation. This sample excludes peakers.

	Spline	Ratio of differences
Baseline (Before)	-23.628*** (0.415)	-34.772*** (0.576)
Interim	-2.095*** (0.571)	-4.210*** (0.796)
After	-5.900*** (0.634)	-9.153*** (0.901)
Observations	20,566	19,877
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01	

Table 10: Contract positions over time Regression of the hourly mean contract position, weighted by the size of the firm. I control for month effects by defining the dependent variable as the residual from a regression of the contract position on month dummies using 2 years of data, to which I add the mean month fixed effect. The standard sample goes from August 2010 to July 2011, leaving four months before the announcement of the regulatory change, four months between the announcement and the implementation, and four months after implementation. The three specifications correspond to three assumptions regarding expectations over the spot price: rational expectations, adaptive expectations, and perfect foresight. This sample excludes peakers.

	<i>Dependent variable:</i>		
	Forward contract position		
	RE	AE	PF
Baseline (Before)	1,334.3*** (14.246)	1,341.8*** (14.307)	1,304.1*** (13.930)
Interim	-11.907 (19.306)	-33.685* (19.372)	-39.997** (18.760)
After	29.578 (21.399)	-4.893 (21.175)	-13.704 (20.638)
Observations	24,301	23,911	23,970
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01		

Table 11: Test of the additive separability of the bids Results from regressing the slope of the bids submitted by producers on their forward contract position. The latter is computed under three different assumptions about expectations: rational expectations (RE), adaptive expectations (AE), and perfect foresight (PF). Includes owner-market and month fixed effects. The fact that the correlation between the slope and the contract position is not significant supports the additive separability assumption. Controlling for total load or adding a time trend does not affect this result. This sample excludes peakers

	<i>Dependent variable:</i>					
	Slope of the bid					
	RE	AE	PF	RE	AE	PF
Residual demand's slope				0.001** (0.0003)	0.001** (0.0003)	0.001** (0.0003)
Spot price				0.102 (0.265)	0.176 (0.131)	0.090 (0.057)
Contract position	0.001 (0.001)	0.001 (0.001)	0.001 (0.001)	-0.0002 (0.002)	-0.001 (0.002)	-0.001 (0.002)
Observations	195,490	195,490	195,490	194,292	194,292	194,292

Note: *p<0.1; **p<0.05; ***p<0.01

Table 12: Forward premium Results from a regression of the forward premium, using node fixed effects. The premium is measured as a fraction of the forward price.

	<i>Dependent variable:</i>
	premium
interim	-0.094*** (0.001)
after	-0.181*** (0.0004)
Mean FE	0.12
Observations	16,829,313

Note: *p<0.1; **p<0.05; ***p<0.01

Table 13: Quantity weighted forward premium Regression of the hourly mean forward premium, weighted by the quantity traded by each firm. Before computing it, the top and bottom 1% of the sample is removed to avoid extreme values. I control for month effects by defining the dependent variable as the residual from a regression of the forward premium on month dummies using 2 years of data, to which I add the mean month fixed effect. The sample goes from August 2010 to July 2011, leaving four months before the announcement of the regulatory change, four months between the announcement and the implementation, and four months after implementation. Robust standard errors reported.

	<i>Dependent variable:</i>	
	Forward premium	
	Level	Fraction of the forward price
Baseline (Before)	3.380*** (0.121)	0.088*** (0.003)
Interim	-0.641*** (0.172)	-0.014*** (0.005)
After	-0.745*** (0.196)	-0.012** (0.005)
Observations	20,596	20,619
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01	

Table 14: Total expenditure Total expenditure was computed as total purchases in each market, times the average clearing price at demand node. The total is the sum of the purchases in the forward and spot market. Specifications (1) and (2) includes hour and month fixed effects were used. Data is hourly, so a 24 lag is a 1 day lag. The sample goes from August 2010 to July 2011, and each of the periods considered is 4 months long: before the announcement, between announcement and implementation, and after implementation. HAC standard errors reported.

	log(total expenditure)		
	(1)	(2)	(3)
before	0.299*** (0.065)	0.044** (0.021)	-0.102*** (0.015)
interim	0.060 (0.039)	-0.088*** (0.012)	0.106*** (0.013)
log(real-time load)	3.213*** (0.021)	3.210*** (0.021)	
Trend	0.00000 (0.000)		
log(oil price)	0.136 (0.101)	0.138 (0.101)	
log(natural gas price)	-0.131* (0.074)	-0.117 (0.074)	
log(coal price I)	-0.605*** (0.140)	-0.553*** (0.140)	
log(coal price PRB)	-0.417 (0.284)	-0.514* (0.284)	
log(oil price) _{t-24}	0.219** (0.138)	0.199 (0.138)	
log(oil price) _{t-48}	0.533*** (0.096)	0.519*** (0.096)	
log(natural gas price) _{t-24}	-0.490** (0.101)	-0.504*** (0.101)	
log(natural gas price) _{t-48}	0.803*** (0.073)	0.786*** (0.072)	
log(coal price I) _{t-24}	0.279 (0.191)	0.280 (0.193)	
log(coal price I) _{t-48}	0.340 (0.152)	0.275* (0.153)	
log(coal price PRB) _{t-24}	-0.408 (0.340)	-0.403 (0.339)	
log(coal price PRB) _{t-48}	0.172 (0.217)	0.282 (0.217)	
Constant	-36.241*** (3.261)	-23.915*** (0.595)	14.488*** (0.011)
Observations	8,757	8,757	8,757
R ²	0.916	0.916	0.025

Note: 66 *p<0.1; **p<0.05; ***p<0.01

Table 15: Spot market markups Regression of the markup in the spot market on time-period dummies using firm-market fixed effects and clustered standard errors. The sample goes from August 2010 to July 2011, and each of the periods considered is 4 months long: before the announcement, between announcement and implementation, and after implementation. The derivative of the residual demand is computed using a spline and as a ratio of differences using two points (slope). Results are very similar when a time trend is added, but the trend is significant and negative.

<i>Dependent variable: spot markup</i>		
	Spline	Slope
interim	0.014 (0.012)	0.036 (0.031)
after	0.035*** (0.013)	0.101*** (0.030)
Mean FE	-0.024	-0.161
Observations	1,241,481	1,401,455
R ²	0.00001	0.00002
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01	

Table 16: Spot markup for large and small firms Regression of the markup in the spot market on time-period dummies using firm-market fixed effects and clustered standard errors. The sample goes from August 2010 to July 2011, and each of the periods considered is 4 months long: before the announcement, between announcement and implementation, and after implementation. The derivative of the residual demand is computed using a spline and as a ratio of differences using two points (slope). The regression is run separately for large and small firms. Large firms are those with a capacity of 1100 MW or more, which includes 20% of the firms and 80% of the production. Small firms are active less often than large ones, which explains why the number of observations is just roughly twice as large for small firms.

<i>Dependent variable: spot markup</i>				
	Spline		Slope	
	Large firms	Small	Large	Small
interim	0.066** (0.027)	-0.021** (0.010)	0.150** (0.065)	-0.042 (0.026)
after	0.074*** (0.028)	0.009 (0.010)	0.162** (0.066)	0.060** (0.025)
Mean FE	-0.204	0.066	-0.636	0.072
Observations	471,951	769,530	542,269	859,186
R ²	0.00004	0.00002	0.00003	0.00003
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01			

Table 17: Spot markup for large and small firms for spot market including only net sellers
Regression of the markup in the spot market on time-period dummies using firm-market fixed effects and clustered standard errors. The sample goes from August 2010 to July 2011, and each of the periods considered is 4 months long: before the announcement, between announcement and implementation, and after implementation. The derivative of the residual demand is computed using a spline and as a ratio of differences using two points (slope). The regression is run separately for large and small firms. Large firms are those with a capacity of 1100 MW or more, which includes 20% of the firms and 80% of the production. Small firms are active less often than large ones, which explains why the number of observations is just roughly twice as large for small firms. Only firms clearing more in the forward than in the spot market are included.

	<i>Dependent variable: spot markup</i>			
	Spline		Slope	
	Large firms	Small	Large	Small
interim	0.232*** (0.059)	0.093*** (0.027)	0.211 (0.129)	-0.084 (0.062)
after	0.647*** (0.060)	0.240*** (0.027)	1.296*** (0.132)	0.326*** (0.062)
Mean FE	5.77	2.86	11.25	6.11
Median FE	1.91	0.26	4.66	0.89
Observations	212,147	344,930	243,375	387,091
R ²	0.001	0.0002	0.0005	0.0001

Note: *p<0.1; **p<0.05; ***p<0.01

Table 18: Summary statistics for prices Clearing prices are locational marginal prices (LMP) and have three components: marginal cost, which is common to all nodes, the marginal congestion component (MCC), and the marginal losses component (MLC). Although rare, electricity prices can be negative when it is very expensive to stop a plant that has already started. The sample starts in January 2010 and ends in December 2011.

Statistic	N	Mean	St. Dev.	Min	Max
Forward market					
LMP	33,271,056	31.5	14.7	-291.1	500
MCC	33,271,056	-0.79	6.30	-308.8	434.2
MLC	33,271,056	-0.57	1.95	-44.7	37.5
Spot market					
LMP	33,271,056	30.7	26.3	-851	1,888
MCC	33,271,056	-0.80	15.99	-924	1,829
MLC	33,271,056	-0.55	2.14	-148.5	121.8

Table 19: Forward premium Presents results from a regression of the forward premium, computed as proportion of the forward price, on a constant. This was done separately for the period before April 1, 2011, when transaction costs decrease, and the period after this. The total sample covers 2010 and 2011. The second line shows the result from the same regression but using the forward premium net of transaction costs. The computation of the second line after April 2011 indicates the profits of selling 1MW at each node every hour. It is negative, meaning that that trader would lose money with that strategy. Before April 2011, the RSG charge was uniform across nodes, while after the regulatory change the charges vary by node. Therefore a trader could potentially obtain larger profits following a more sophisticated strategy. Robust standard errors reported.

	Before April 2011	After April 2011
Forward premium	0.023*** (0.008)	-0.012 (0.009)
Net forward premium	-0.037*** (0.008)	-0.024*** (0.009)
Observations	21,154,967	12,113,631
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01	

References

- Blaise Allaz. Oligopoly, uncertainty and strategic forward transactions. *International Journal of Industrial Organization*, 10(2):297–308, 1992.
- Blaise Allaz and Jean-Luc Vila. Cournot competition, forward markets and efficiency. *Journal of Economic theory*, 59(1):1–16, 1993.
- Hunt Allcott. Real time pricing and electricity markets. 2013.
- Donald WK Andrews. Tests for parameter instability and structural change with unknown change point. *Econometrica: Journal of the Econometric Society*, pages 821–856, 1993.
- John Asker. A study of the internal organization of a bidding cartel. *American Economic Review*, page 724762, 2010.

- Jushan Bai and Pierre Perron. Estimating and testing linear models with multiple structural changes. *Econometrica*, pages 47–78, 1998.
- Patrick Bajari and Lixin Ye. Deciding between competition and collusion. *Review of Economics and Statistics*, 85(4):971-989, 2003.
- John Birge, Ali Hortaçsu, Ignacia Mercadal, and Michael Pavlin. The role of financial players in electricity markets: An empirical analysis of MISO. 2014.
- John Birge, Ali Hortaçsu, and Michael Pavlin. Inverse optimization for the recovery of market structure from market outcomes: An application to the miso electricity market. *Available at SSRN 2612234*, 2015.
- Severin Borenstein and James Bushnell. An empirical analysis of the potential for market power in california s electricity industry. *The Journal of Industrial Economics*, 47(3):285–323, 1999.
- Severin Borenstein, James Bushnell, and Steven Stoft. The competitive effects of transmission capacity in a deregulated electricity industry. Technical report, National Bureau of Economic Research, 1997.
- Severin Borenstein, James Bushnell, and Frank Wolak. Measuring market inefficiencies in californias restructured wholesale electricity market. *American Economic Review*, 92(5):1376–1405, 2002.
- Severin Borenstein, James Bushnell, Christopher R Knittel, and Catherine Wolfram. Inefficiencies and market power in financial arbitrage: a study of california s electricity markets. *The Journal of Industrial Economics*, 56(2):347–378, 2008.
- N. Bowden, S. Hu, and J. Payne. Day-ahead premiums on the midwest ISO. *The Electricity Journal*, 22(2):64–73, 2009.
- James Bushnell. Oligopoly equilibria in electricity contract markets. *Journal of Regulatory Economics*, 32(3):225–245, 2007.
- Meghan R Busse. Firm financial condition and airline price wars. *Available at SSRN 237374*, 2000.

- Ronald H Coase. Durability and monopoly. *Journal of Law and Economics*, 15:143, 1972.
- Joseph Cullen. Measuring the environmental benefits of wind-generated electricity. *American Economic Journal: Economic Policy*, 5(4):107–133, 2013.
- Joseph Cullen. Dynamic response to environmental regulation in the electricity industry. 2015.
- Avinash Dixit. The role of investment in entry-deterrence. *The economic journal*, page 95106, 1980.
- Ulrich Doraszelski, Gregory Lewis, and Ariel Pakes. Just starting out: Learning and equilibrium in a new market. 2015.
- Natalia Fabra and Juan Toro. The fall in british electricity prices: market rules, market structure, or both? 2003.
- Natalia Fabra and Juan Toro. Price wars and collusion in the spanish electricity market. *International Journal of Industrial Organization*, 23(3-4):155181, 2005.
- Kira R Fabrizio, Nancy L Rose, and Catherine D Wolfram. Do markets reduce costs? assessing the impact of regulatory restructuring on us electric generation efficiency. *The American Economic Review*, 97(4):1250–1277, 2007.
- B Fattouh, L Kilian, and L Mahadeva. The role of speculation in oil markets: What have we learned so far?. *Energy Journal*, 34:733, 2013.
- Austan Goolsbee and Chad Syverson. How do incumbents respond to the threat of entry? evidence from the major airlines. *Quarterly Journal of Economics*, 2008.
- Richard Green. The electricity contract market in england and wales. *The Journal of Industrial Economics*, 47(1):107–124, 1999.
- Michael Greenstone and Rema Hanna. Environmental regulations, air and water pollution, and infant mortality in india. *American Economic Review*, 104(10):3038–72, 2014. doi: 10.1257/aer.104.10.3038. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.104.10.3038>.
- Sanford Grossman. On the efficiency of competitive stock markets where trades have diverse information. *The Journal of finance*, 31(2):573–585, 1976.

- Sanford J Grossman and Joseph E Stiglitz. On the impossibility of informationally efficient markets. *The American economic review*, page 393-408, 1980.
- Bruce E Hansen. The new econometrics of structural change: Dating breaks in us labor productivity. *Journal of Economic perspectives*, pages 117–128, 2001.
- Ali Hortaçsu and David McAdams. Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market. *The Journal of Political Economy*, 118(5):833–865, 2010.
- Ali Hortaçsu and Steven L Puller. Understanding strategic bidding in multiunit auctions: a case study of the Texas electricity spot market. *The RAND Journal of Economics*, 39(1):86–114, 2008.
- Koichiro Ito and Mar Reguant. Sequential markets, market power and arbitrage. Technical report, National Bureau of Economic Research, 2014.
- Akshaya Jha and Frank A Wolak. Testing for market efficiency with transactions costs: An application to convergence bidding in wholesale electricity markets. 2013.
- Paul L Joskow and Jean Tirole. Transmission rights and market power on electric power networks. *The Rand Journal of Economics*, pages 450–487, 2000.
- Luciana Juvenal and Ivan Petrella. Speculation in the oil market. *Journal of Applied Econometrics*, 2014.
- Jakub Kastl. Discrete bids and empirical inference in divisible good auctions. *The Review of Economic Studies*, 78(3):974–1014, 2011.
- Lutz Kilian and Daniel P Murphy. The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454-478, 2014.
- Christopher R Knittel and Robert S Pindyck. The simple economics of commodity price speculation. Technical report, National Bureau of Economic Research, 2013.
- Haifeng Liu, Leigh Tesfatsion, and Ali A Chowdhury. Derivation of locational marginal prices for restructured wholesale power markets. *Journal of Energy Markets*, 2(1):3–27, 2009.

- Henry Louie and Kai Strunz. Locational marginal pricing in north american power systems. *ETG-Fachbericht-Netzregelung und Systemführung*, 2008.
- Erin T Mansur. Measuring welfare in restructured electricity markets. *The Review of Economics and Statistics*, 90(2):369386, 2008.
- Paul Milgrom and John Roberts. Predation, reputation, and entry deterrence. *Journal of economic theory*, 27(2):280312, 1982.
- David M Newbery. Competition, contracts, and entry in the electricity spot market. *The RAND Journal of Economics*, page 726749, 1998.
- John E Parsons, Cathleen Colbert, Jeremy Larrieu, Taylor Martin, and Erin Mastrangelo. Financial arbitrage and efficient dispatch in wholesale electricity markets. *MIT Center for Energy and Environmental Policy Research*, (15-002), 2015.
- Steven L Puller. Pricing and firm conduct in california s deregulated electricity market. *The Review of Economics and Statistics*, 89(1):75–87, 2007.
- Richard E Quandt. Tests of the hypothesis that a linear regression system obeys two separate regimes. *Journal of the American statistical Association*, 55(290):324–330, 1960.
- Mar Reguant. Complementary bidding mechanisms and startup costs in electricity markets. *The Review of Economic Studies*, 81(4):1708–1742, 2014.
- Nicholas Ryan. The competitive effects of transmission infrastructure in the indian electricity market. Technical report, Working Paper, Yale (October), 2014.
- Celeste Saravia. Speculative trading and market performance: the effect of arbitrageurs on efficiency and market power in the new york electricity market. *Center for the Study of Energy Markets*, 2003.
- William L Silber. *The economic role of financial futures*. Salomon Brothers Center for the Study of Financial Institutes, Graduate School of Business Administration, 1985.
- Robert Wilson. Auctions of shares. *The Quarterly Journal of Economics*, pages 675–689, 1979.

- Frank A. Wolak. Identification and estimation of cost functions using observed bid data: an application to electricity markets. In Mathias Dewatripont, Lars Peter Hansen, and Stephen J Turnovsky, editors, *Advances in economics and econometrics: theory and applications, Eighth World Congress*.
- Frank A Wolak. An empirical analysis of the impact of hedge contracts on bidding behavior in a competitive electricity market. *International Economic Journal*, 14(2):139, 2000.
- Frank A Wolak. Identification and estimation of cost functions using observed bid data. In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, volume 2, page 133. Cambridge University Press, 2003a.
- Frank A Wolak. Measuring unilateral market power in wholesale electricity markets: the california market, 1998-2000. *American Economic Review*, pages 425–430, 2003b.
- Frank A Wolak. Quantifying the supplierside benefits from forward contracting in wholesale electricity markets. *Journal of Applied Econometrics*, 22(7):11791209, 2007.
- Frank A Wolak. Measuring the benefits of greater spatial granularity in short-term pricing in wholesale electricity markets. *American Economic Review*, 101(3):247, 2011.
- Catherine D Wolfram. Measuring duopoly power in the british electricity spot market. *American Economic Review*, pages 805–826, 1999.
- Achim Zeileis, Christian Kleiber, Walter Krämer, and Kurt Hornik. Testing and dating of structural changes in practice. *Computational Statistics & Data Analysis*, 44(1):109–123, 2003.
- Fanyin Zheng. Spatial competition and preemptive entry in the discount retail industry. Technical report, Working paper, Harvard University, 2014.

Appendices

A Clearing prices in the MISO market

Typically, the energy market is organized as an auction in which participants submit bids to buy or sell energy at particular locations; the ISO then clears the market solving a linear programming problem that minimizes cost subject to the capacity constraints imposed by the transmission network. Because the latter has limited capacity, electricity supplied at different locations is not a homogeneous good. Therefore, both the cost of a MWh and the willingness to pay for it vary across the market footprint, and it is not obvious how the market should be cleared.

There are two alternative market designs to determine clearing prices in markets in which transmission lines reach capacity. The first is zonal pricing, which divides the market into a few zones and allows for a different price at each zone, but a uniform price within each of them. This makes sense particularly when there is enough capacity within each zone. The second is nodal pricing, in which each location is allowed to be cleared at a different price. Although there were more zonal markets when the deregulation of electricity markets started, today all market in the US have nodal pricing.⁴⁴

MISO uses nodal pricing to clear the energy market. The clearing price at each node or location where energy is produced or demanded represents the marginal cost of bringing 1 MW to that particular node, and it is called locational marginal price (LMP). The LMP has three components: marginal cost, congestion, and losses. The marginal cost component is common across nodes and represents the cost of buying 1 more MW of energy given the supply bids submitted by generators. Moving electricity from one location to another requires some energy, so less than 1 MW arrives to a node when it is produced at a different node. This is captured by the losses component. Lastly, the marginal congestion component of price represents the increase in price required to clear the market when transmission lines are at capacity. For instance, if demand at the marginal cost is larger than what can be transmitted to that node, the price at that node has to increase until there is no excess demand. Summary statistics for prices are presented in Table 18.

⁴⁴See Wolak (2011) for a discussion on the benefits of nodal vs. zonal pricing, and a quantification of the benefits of the former.

To better understand congestion pricing, consider a simple example without losses in which there are only two nodes. At node A there is only demand and it is given by $Q = 120 - P$, and node B only produces energy and has a marginal cost of 10. The transmission line connecting these nodes has a capacity of 100MW. Suppose there are enough generators at node B to have them selling at marginal cost. Demand at that price is 110, but that quantity cannot be brought to A because it exceeds the line's capacity. Therefore, the clearing locational marginal price at node A is 110. The marginal cost component is 100 and the congestion component is 10.

The MISO energy market has over 2000 nodes and often becomes congested, so in practice there is significant price dispersion. Figure 1 presents a heat map of the MISO footprint, and illustrates how prices can substantially differ geographically and over time.

B Revenue Sufficiency Guarantee (RSG) charges

In the MISO market, some eligible generators are guaranteed the full recovery of their production cost when MISO commits them to produce a quantity that differs from their day-ahead schedule. The production cost has three components: the start-up cost, incurred when the generating units start running, the no-load cost, which is the cost of operating and producing zero MWs, and the marginal cost. Only the latter is covered by the market clearing price (LMP), so the eligible generators need to be compensated for their incurred start-up and no-load costs. This is funded by imposing Revenue Sufficiency Guarantee (RSG) charges on deviations from the day-ahead schedule, *i.e.* on differences between the MWs that a market participant cleared in the day-ahead market and what she produces in the real-time market. As virtual participants do not physically buy or sell energy, the total virtual MWs are considered a deviation and are subject to RSG charges.

MISO's treatment of virtual bidders with respect to the RSG has varied over time in a way that affects incentives. When the market was opened to financial participants in April 2005, virtual transactions were not subject to RSG charges. In April 2006, the FERC issued an order according to which virtual offers had to pay RSG charges retroactively until 2005. This was reversed in October of the same year. After a long discussion between MISO, market participants, and the FERC, in November 2008 the latter determined that virtual supply had to pay RSG charges. This applied to future trades as well as retroactively until April 2006. The discussion about what trades

should be subject to the charges and how these should be computed continued until April 2011. During this period, charges were constant across nodes, computed as $RSG_i = MW_i^S \cdot RSG_RATE$, where i is a bid and MWS are MWs of virtual supply. This means that if a virtual bidder was buying 1 MW at a node, her payoff was just the real-time price minus the day-ahead one. For a virtual participant selling 1 MW in the day-ahead market, the payoff was $p^F - p^S - RSG_RATE$. Charges during this period were on average larger than the day-ahead premium (see Tables 1 and 3). On March 2011 the FERC accepted MISO's proposal for a change in the computation of the RSG charges. Since April 1st, 2011, both virtual supply and virtual demand are subject to these charges and their calculation has changed. In addition to a component that is common across nodes, the Day-Ahead Deviation & Headroom Charge or DDC, there is a component that depends on congestion at each specific node called the Constraint Management Charge or CMC. As shown in the formula below, the CMC depends on the sum of deviations weighted by a congestion factor called the Constraint Contribution Factor or CCF which is between -1 and 1. When it is positive, the constraint is relaxed by more demand or less supply, so charges are imposed only on supply; when the factor is negative, only demand has to pay deviation charges. The calculation of the charges for each participant is as follows:

$$\begin{aligned}
 RT_RSG_DIST1_h &= CMC_DIST_h + DDC_DIST_h \\
 CMC_DIST_h &= \sum_n \max \{ (MW_n^S - MW_n^D) \cdot CCF_{h,n}, 0 \} \cdot CMC_RATE_{h,n} \\
 DDC_DIST_h &= \sum_n \max \{ (MW_n^S - MW_n^D), 0 \} \cdot DDC_RATE_{h,n}
 \end{aligned}$$

where h is an hour, MW_n^S and MW_n^D are the virtual supply and demand, respectively, cleared by the participant at node n for hour h .

C Model with strategic demand and supply

This appendix extends the model presented in section 5.1 to include strategic demand. Instead of taking demand given, I model buyers strategically choosing how to distribute their purchases between the spot and the forward markets. Because in wholesale electricity markets most

purchases come from utilities serving downstream consumers, I will refer to buyers as utilities. Additionally, I will assume that firms do not hold hedging contracts for the spot price, i.e. $x^S = 0$. The market subindexes are omitted in this section, but the analysis is always done under the assumption of independent separate markets.

Demand

Unlike generators, utilities' only decision is how to split purchases between the forward and the spot markets. They do not choose how much electricity to buy in the spot market, because final demand is given by households' electricity consumption. Therefore, the spot market is cleared such that there is enough generation to cover the load forecast L , which has a deterministic component l and a random component ϵ . In the forward market, each buyer submits a schedule $D(p^F)$ indicating how much she is willing to buy at each price. The difference between the quantity cleared in the forward market and L has to be purchased in the spot market.

Like generators, buyers may have financial contracts that affect their position in the forward market. I denote the contract terms as above: a firm holds a contract for a quantity x at a price h . Profits from the hedging contract are computed differently from generators though, because utilities are on the other side of the contract. If the clearing price is larger than h , the buyer gets paid the difference; if the clearing price is smaller than h , the buyer pays the difference to the other side (a generator).

Market clearing

The market clearing prices \bar{p}^F and \bar{p}^S are determined by the market clearing conditions below

$$\sum_{j \in \text{Sellers}} Q_j(\bar{p}^F) = \sum_{b \in \text{Buyers}} D_b(\bar{p}^F) \quad (16)$$

$$\sum_{j \in \text{Sellers}} S_j(\bar{p}^S) = l + \epsilon \quad (17)$$

Generators' uncertainty

As before, each generator i faces uncertainty over the clearing prices \tilde{p}^F and \tilde{p}^S , because she does not know what clearing price will result from submitting different schedules. In the spot market, uncertainty comes from the random component of demand, as in the section without strategic demand. In the forward market, it comes from the unknown hedging positions of other firms, which are private information and therefore make the clearing price uncertain. In other words, the generator is uncertain about the residual demand she faces, because residual demand depends on other firms' bidding behavior.

Bidder i 's uncertainty in the forward market is represented by $F_x(x_{-i}|x_i)$, the distribution of other firms' contract positions. It is conditional on i 's own position because i 's position may contain information about others' contracts. Note that this remains a private value setting since i 's profits do not depend on its competitors' hedging positions. In the spot market, uncertainty comes from ϵ , which has distribution $F_\epsilon(\epsilon)$.

As above, I define a probability measure over the realizations of the forward clearing price from the perspective of firm i , conditional on i 's private information about its contract position x_i^F , i 's submission of a schedule $\hat{Q}_i(p, x_i^F)$, and her competitors playing their equilibrium strategies $\{Q_j(p, x_j^F), j \neq i\}$.

$$H(p, \hat{Q}_i(p); x_i^F) \equiv \Pr(\tilde{p}^F \leq p \mid x_i^F, \hat{Q}_i) \quad (18)$$

$H(p, \hat{Q}_i(p); x_i^F)$ represents the uncertainty over the forward market clearing price faced by firm i . It is the probability, given i 's contract position, that generator i will be paid a price p when she sells a quantity $\hat{Q}_i(p)$ and all other generators submit the equilibrium offer functions. The event $\tilde{p}^F \leq p$ is equivalent to the event of excess supply at price p . Using the market clearing condition in Equation 17, H can be written as

$$\begin{aligned} H(p, \hat{Q}_i(p); x_i^F) &= \Pr\left(\sum_{j \neq i} Q_j(p, x_j^F) + \hat{Q}_i(p) \geq \sum_{d \in \text{Buyers}} D_d^F(p, x_d^F) \mid x_i^F, \hat{Q}_i\right) \\ &= \int_{x_{-i}^F} 1\left\{\sum_{j \neq i} Q_j(p, x_j^F) + \hat{Q}_i(p) \geq \sum_{d \in \text{Buyers}} D_d^F(p, x_d^F)\right\} dF^F(x_{-i}^F \mid x_i^F) \end{aligned} \quad (19)$$

The generator's problem

The problem of the firm is to choose forward and spot bids that maximize its expected profits. As in the case without strategic demand, the generator's expected profits are given by:

$$\max_{Q(p^F), S(p^S)} \int_{\underline{p}}^{\bar{p}} \int_{\underline{p}}^{\bar{p}} U\left(\Pi(Q(p^F), x^F), S(p^S)\right) dH(p^F, Q(p^F); x^F) dG(p^S, S(p^S); x^S)$$

The Euler-Lagrange conditions for the bids that maximize the generator's profits are (proof analogous to the one in Appendix D)

$$p^F - p^S = [Q(p^F) - x^F] \frac{H_S}{H_P} \quad (20)$$

$$p^S - c' = [S(p^S) - Q(p^F) - x^S] \frac{G_S}{G_P} \quad (21)$$

Additive separability

If the schedules submitted by both buyers and sellers satisfy additive separability, the optimality conditions can be written in terms of the residual demand or supply. To see this, assume that demand and supply schedules are additively separable and therefore can be written as $D(p) = a(p) + b(x)$ and $Q(p) = \alpha(p) + \beta(x)$. The event of excess supply at price p can then be written

$$\begin{aligned} \sum_{i \in I^S} \alpha_i(p) + \sum_{i \in I^S} \beta_i(x) &\geq \sum_{i \in I^D} a_i(p) + \sum_{i \in I^D} b_i(x) \\ \sum_{i \in I^S} \alpha_i(p) - \sum_{i \in I^D} a_i(p) &\geq \sum_{i \in I^D} b_i(x) - \sum_{i \in I^S} \beta_i(x) \end{aligned}$$

Defining $\theta \equiv \sum_{i \in I^D} b_i(x) - \sum_{i \in I^S} \beta_i(x)$, a random variable with distribution Γ . Then, the expectation of excess supply from the perspective of a generator is

$$\begin{aligned}
H(p, \hat{Q}(p); x_i^F) &= \Pr\left(\sum_{j \neq i} Q_j(p, x_j^F) + \hat{Q}_i \geq D^F(p) | x_i^F, \hat{Q}\right) \\
&= \Pr\left(\sum_{j \in I^D} a_j(p) - Q_i - \sum \alpha_j(p^F) \geq \sum \beta_j(x_j^F) - \sum_{j \in I^D} b_j(x)\right) \\
&= \Gamma\left(\sum_{j \in I^D} a_j(p) - Q_i - \sum \alpha_j(p^F)\right)
\end{aligned}$$

And equivalently for demand. Taking derivatives and simplifying, the optimality conditions can be rewritten as Equations 10 and 11 for sellers and an equivalent one for buyers.

D Derivation of the Euler-Lagrange conditions for the generator's problem

From Section 5.1, the problem of the firm is the following:

$$\max_{Q(p^F), S(p^S)} \int_{\underline{p}}^{\bar{p}} \int_{\underline{p}}^{\bar{p}} U(\Pi(Q, S)) dH(p^F, Q(p^F); x^F) dG(p^S, S(p^S); x^S)$$

We can rewrite $dH(p^F, Q(p^F); x^F)$ and $dG(p, \hat{S}(p); x^S)$ as:

$$\begin{aligned}
dH(p^F, Q(p^F); x^F) &= \frac{dH}{dp^F} dp^F = (H_Q Q' + H_P) dp^F \\
dG(p^S, S(p^S); x^S) &= \frac{dG}{dp^S} dp^S = (G_S S' + G_P) dp^S
\end{aligned} \tag{22}$$

Replacing the above and defining the integrand as $J(Q, Q', p^F, S, S', p^S)$, the integrand now becomes

$$J(Q, Q', p^F, S, S', p^S) \equiv U[H_Q Q' + H_P][G_S S' + G_P]$$

where $U = U(p^F Q(p^F) + p^S [S(p^S) - Q(p^F)] - C(S(p^S)) - [p^F - h^F] x^F - [p^S - h^S] x^S)$. The argument is omitted from now on. The Euler-Lagrange equations are:

$$\begin{aligned}
J_Q &= \frac{\partial}{\partial p^F} J_{Q'} \\
J_S &= \frac{\partial}{\partial p^S} J_{S'}
\end{aligned} \tag{23}$$

Taking derivatives:

$$\begin{aligned}
J_Q &= U'[p^F - p^S][H_Q Q' G_S S' + H_Q Q' G_P + H_P G_S S' + H_P G_P] + \\
&\quad U[H_{QQ} Q' G_S S' + H_{QQ} Q' G_P + H_{PQ} G_S S' + H_{PQ} G_P] \\
J_S &= U'[p^S - c][H_Q Q' G_S S' + H_Q Q' G_P + H_P G_S S' + H_P G_P] + \\
&\quad U[H_{QQ} Q' G_{SS} S' + H_{QQ} Q' G_{PS} + H_{PQ} G_{SS} S' + H_{PQ} G_{PS}] \\
J_{Q'} &= U[H_Q G_S S' + H_Q G_P] \\
J_{S'} &= U[H_Q Q' G_S + H_P G_S]
\end{aligned}$$

$$\begin{aligned}
\frac{\partial}{\partial p^F} J_{Q'} &= U'[Q + p^F Q' - p^S Q' - x^F][H_Q G_S S' + H_Q G_P] + \\
&\quad U[H_{QQ} Q' G_S S' + H_{QP} G_S S' + H_{QQ} G_P Q' + H_{QP} G_P] \\
\frac{\partial}{\partial p^S} J_{S'} &= U'[p^S S' + S - Q - c' S' - x^S][H_Q Q' G_S + H_P G_Q] + \\
&\quad U[H_Q Q' G_{SS} S' + H_Q Q' G_{SP} + H_P G_{SS} S' + H_P G_{SP}]
\end{aligned}$$

After substituting and canceling terms, the Euler-Lagrange conditions are:

$$p^F - p^S = [Q(p^F) - x^F] \frac{H_S}{H_P} \quad (24)$$

$$p^S - c' = [S(p^S) - Q(p^F) - x^S] \frac{G_S}{G_P} \quad (25)$$

E Additive Separability

If schedules are additively separable in the contract position and the price, then the event of excess supply can be written

$$D^F(p^F) - Q_i - \sum \alpha_j(p^F) < \sum \beta_j(x_j^F) - \epsilon^F \quad (26)$$

Define $\theta \equiv \sum \beta_j(x_j^F) - \epsilon^F$, a random variable with distribution $\Gamma(\cdot)$. This variable θ contains the uncertain components determining the clearing price. Using the definition of θ , H can be rewritten as follows

$$\begin{aligned}
H(p, \hat{Q}(p); x_i^F) &= \Pr\left(\sum_{j \neq i} Q_j(p, x_j^F) + \hat{Q}_i \geq D^F(p) | x_i^F, \hat{Q}\right) \\
&\Pr\left(D^F(p^F) - Q_i - \sum \alpha_j(p^F) < \sum \beta_j(x_j^F) - \epsilon^F\right) \\
&1 - \Gamma\left(D^F(p^F) - Q_i - \sum \alpha_j(p^F)\right)
\end{aligned}$$

and an equivalent expression holds for G . Taking derivatives of this expression and simplifying,

$$\frac{H_S}{H_p} = \frac{1}{D'(p) - \sum \alpha'(p)} \quad (27)$$

Notice that the denominator of the right hand side of equation 27 is the derivative of the ex-post residual demand faced by generator i . For a given realization of ϵ and x_{-i} , the residual demand faced by i is

$$R(p) = D(p) + \epsilon - \sum_{j \neq i} \alpha_j(p) - \sum_{j \neq i} \beta(x_j) \quad (28)$$

therefore its derivative is $D'(p) - \sum \alpha'(p)$. Replacing this in the optimality conditions, they become

$$\begin{aligned}
p^F - p^S &= -[Q^*(p^F) - x^F] \frac{1}{R'(p^F)} \\
p^S - c' &= -[S^*(p^S) - Q^*(p^F) - x^S] \frac{1}{R'(p^S)}
\end{aligned}$$

F Market-clearing algorithm

In the MISO market, generators submitted schedules consist of more information than the 10 steps of the bid. They additionally indicate the maximum and minimum quantity that they can produce economically, and under an emergency, as well as whether they act as price-takers. Additionally, they may indicate that the unit is already working, so it must run during that hour but they do not need to pay the start costs. They also provide technical information about the plant like the maximum and minimum temperatures, ramping times and costs, and the number of hours in a row a unit needs to run. The effect of these cost complementarities has been studied by Reguant (2014)

MISO only publishes some of the information provided by the generators at each moment. The main part missing are the complementarities between hours that the market authority must consider when clearing the market. As a simplification, I do not consider this when I clear the markets either, but this does not seem to cause great divergence between my simulated market clearing quantities and prices, and those observed in the data.

I include the step function submitted by each bidder, as well as whether they are price-takers. Additionally, I adjust some bids to reflex other parameters. For instance, a good number of run-of-river and wind units submit offers for 999MW in the second step, even though their capacity, as represented by the economic and emergency maxima, is below this (usually around 10MW).⁴⁵ As keeping this would alter the market clearing results, I modify the bids to reflect the unit's capacity. I generally restrict every step to be below the specified economic maximum. Additionally, when a bid specifies a quantity in the first step, but no prices, I assume they are willing to pay any price for that quantity.

When I compute the measure of fit for the different market definitions, I compute a clearing price for each of the market and compare it to the observed price. I do not observe a single price for any market, as prices differ across nodes. I compute the observed clearing price by taking the quantity weighted average, where quantities are given by the volume cleared by supply. This is better than using the mean of all nodes in the market, since some nodes are hubs used only for financial trading, or not active at all hours. Additionally, the fit is considerably better using quantity weighted average than simple average.

⁴⁵The economic minimum and maximum are part of the bids submitted by generators, and indicate the minimum and maximum quantity that it is profitable to produce. They may be willing to produce more under emergency conditions.