

Limits to Arbitrage in Electricity Markets: A Case Study of MISO

JOHN BIRGE, ALI HORTAÇSU,
IGNACIA MERCADAL AND MICHAEL PAVLIN

JANUARY 2017

CEEPR WP 2017-003

Limits to Arbitrage in Electricity Markets: A case study of MISO

John Birge

University of Chicago

Ali Hortaçsu

University of Chicago

Ignacia Mercadal

MIT

Michael Pavlin

Wilfred Laurier University

January 21, 2017*

Abstract

As in most commodities markets, deregulated electricity markets allow the participation of purely financial traders to enhance informational and productive efficiency. The presence of financial players is expected, among other things, to help eliminate predictable pricing gaps between forward and spot prices, which may arise in the presence of market power and are linked to productive inefficiency. However, we find that the impact of financial players on reducing pricing gaps has been limited, even using credibly exogenous variation in financial activity to address potential endogeneity. A forward premium persists. We show that financial traders effect on the premium was limited by two barriers. First, arbitrageurs do not have unlimited access to capital. Trading was reduced during the financial crisis, when capital availability was restricted. The second is regulation, as high transaction costs imposed by the regulator restricted arbitrage. Moreover, during this period we observe that some financial players appear to be betting in exactly the opposite direction of the pricing gap, sustaining large losses while doing so. We find evidence consistent with participants using forward market bids to affect congestion and thus increase the value of their Financial Transmission Rights (FTR), i.e. these financial players incur losses with one financial instrument to make larger profits with another, introducing artificial congestion to the system.

*We would like to thank Dallas Burtraw, Mar Reguant, and seminar participants at the IIOC 2014, Energy Policy Center at UChicago, and ASSA Meetings 2017.

1 Introduction

In most commodity markets, prices are determined by the interaction of physical traders that actually consume or produce the good, and financial traders that do not typically own physical assets but use derivatives to hedge or arbitrage. Though these financial players are expected to be beneficial, introducing liquidity and enhancing informational efficiency, they have also been blamed for higher commodity prices and accused of price manipulation in several markets. These concerns have led to proposals to either ban financial speculators from some markets, or to restrict their activity, e.g. by imposing a tax on financial trades.

This paper is a case study of the role of financial players in the wholesale electricity market of the American Midwest (MISO¹). An empirical analysis of this market has a number of advantages. First, the regulator exogenously imposed high transaction costs on financial traders for over two years, and then significantly reduced them. These changes resulted in exogenous variation that allows us to identify the effect of financial traders on market outcomes. Additionally, we are able to investigate the potential consequences of restricting speculators activity by comparing their effect in the periods with high and low transaction costs.

Second, we have a detailed picture of the strategy of each financial bidder: we observe the strategy followed by each trader over a large number of hourly interactions as well as the revenues from each trade. Third, we have a very complete picture of the market: physical and financial firms interact in just one market and, because electricity is not storable, we know how financial participants can profit from arbitrage and how prices are expected to change as a result of financial trading. This allows us to examine whether financial firms are behaving as expected, and if not, to find explanations that can credibly rationalize their behavior. We also have detailed data on a related derivative market, the financial transmission rights (FTR) market, which is ostensibly used to hedge congestion risks but may be profitable to financial players.

Financial traders have been introduced to several deregulated wholesale electricity markets in order to arbitrage away predictable price differences that result in inefficiencies [Jha and Wolak, 2013]. While in other markets the introduction of financial traders has decreased these gaps [Jha and Wolak, 2013, Saravia, 2003], we find that in MISO this has not been successful. Even after using exogenous variation in trading from capital restrictions and high transaction costs, we find that financial players have at most a weak

¹Midcontinent Independent System Operator

effect. We show that instead of arbitraging the premium, some players have bid in the opposite direction, effectively buying expensive and selling cheap, and consistently made losses from it. This behavior is consistent with using financial trading in the energy market in order to increase local prices, which increases the value of a related financial instrument that allows to bet on local price differences. We also find that evidence consistent with market manipulation is stronger when transaction costs for financial traders were high, which suggests that competition can discipline financial players better than constraints to their activity.

Wholesale electricity markets are able to include financial players because most of them are organized as sequential markets. There is first a *day-ahead* forward market that schedules production a day in advance, and then, right before operation, a *real-time* spot market to balance demand and supply. The day-ahead price is on average larger than the real-time price², which results from generators withholding supply in order to increase the price [Ito and Reguant, 2016], and in turn creates incentives for physical participants to arbitrage by underscheduling demand. As last minute changes in the generation level are generally expensive, productive inefficiencies arise when utilities or generators choose different quantities in the spot and forward markets, either to affect price or to benefit from an existing forward premium [Jha and Wolak, 2013]. For this reason, financial trading has been introduced into deregulated wholesale energy markets to prevent the inefficiencies that arise from a positive forward premium, as well as to provide a hedging instrument for physical participants.

Financial traders, referred to as virtual traders in MISO, buy or sell energy along with physical participants at all locations in the day-ahead market. These virtual transactions are then reversed in the real-time market. For instance, if a virtual trader buys 1 MWh in the forward market, it must be sold at the spot price and their profit is the difference between the two prices. The introduction of financial or virtual bidders is expected to make day-ahead and spot prices converge, and thus to eliminate the related inefficiencies. Additionally, financial bids provide an opportunity for physical participants to arbitrage expected price differences or hedge against the spot market risk. See Jha and Wolak [2013] for an excellent discussion of the motivations for incorporating virtual bidders into the market design.

²A positive day-ahead premium has been documented in several markets. Bowden et al. [2009] in the MISO, Saravia [2003] in New York, Jha and Wolak [2013], Borenstein et al. [2008] in California, Ito and Reguant [2016] in the Iberian market, among others.

In MISO, a positive day-ahead premium persisted despite the presence of virtual traders.³ Motivated by this fact, we investigate the effect of financial traders on the gap between the day-ahead and real-time prices.

We show that virtual participants' effect on the premium has been limited. We identify two mechanisms that restrict arbitrage in the MISO electricity market. First, arbitrage requires capital: virtual volume decreased during the financial crisis because bidders' ability to trade and volume of allowed transactions depend on their financial position. Additionally, if capital is restricted, active arbitrageurs will have market power and thus will not bid enough to make the profits go to zero [Borenstein et al., 2008]. Second, regulation imposed barriers to entry and uncertainty over market rules. Between November 2008 and April 2011, virtual supply bids were subject to transaction charges as high as the premium and, additionally, there was significant uncertainty over how these charges would be calculated.

The first two panels of Figure 1 illustrate how financial trades reacted to changes in regulation and capital availability. The volume of virtual transactions fell with the rise of the TED spread, a measure of perceived credit risk that spiked during the crisis, and after the change in regulation in November 2008. On this date, transaction charges on virtual sales in the day-ahead market were significantly increased when deviation charges imposed on unscheduled demand were imposed on virtual transactions. When this happened, virtual volume significantly decreased, and it increased again when they were lowered in April 2011 (Figure 1).

These exogenous changes in the volume traded by financial players allowed us to test the causal effect of financial participants on the forward premium, which is usually hard to determine because a larger premium attracts more financial participants. Previous studies have found that financial traders reduce the premium by comparing the market before and after speculators can participate [Jha and Wolak, 2013, Saravia, 2003]. We instead focus on the effect of increased financial trading. Our findings suggest that virtual bidders reduce the day-ahead premium, but the effect is statistically weak and comes only from virtual supply bids, i.e. those that profit from a positive premium.

The day-ahead premium may persist not only because high transaction costs reduced profits from arbitrage, but also because this resulted in restricted competition and thus affected incentives for virtual players. We find evidence that during the period in which virtual activity was restricted by high transaction costs, some large bidders were not

³The day-ahead premium is about 3% of the day-ahead price between 2008 and 2010.

maximizing profits in the virtual energy market, but rather using their bids to increase the profits of a related financial instrument. Despite the day-ahead premium, virtual traders *bought* more energy in the day-ahead market than they sold. Some virtual traders also accumulated negative profits steadily for longer than a year without leaving the market. These virtual participants were also more likely to submit bids far from the clearing price, consequently not contributing to closing the gap. This is puzzling, because virtual bidders are ostensibly speculators and should take advantage of arbitrage opportunities if they exist and leave the market if they steadily lose money.

Our rationalization of this consistently money-losing behavior by virtual bidders is based on strategic incentives created by the closely connected financial transmission rights (FTR) market. FTRs are financial instruments that allow market participants to bet on local price differences in the day-ahead market, which arise because of the limited capacity of the transmission lines. As it is not always possible to transport electricity from the cheapest producer to cover demand, the resulting local markets clear at different prices. Congestion occurs when a transmission line is at capacity causing prices to differ between locations. The FTR market exists to distribute the rents generated by these price differences. An FTR is defined between a source node and a sink node or location, and results in profits when the day-ahead price at the sink is larger than at the source, i.e. more congested, and losses otherwise.

Although virtual bids are purely financial in the sense that they are not related to any physical energy, they are treated identically to physical bids and can affect day-ahead scarcity in local markets. Virtual bids can be used to artificially increase day-ahead demand at the sink node, which will lead to an increase in price in order to clear the market, particularly when it is not possible to bring electricity to that node due to the limited capacity of the transmission line. In this manner, a virtual bidder can create artificial congestion to increase the profits from an FTR. The profitability of this strategy depends on demand and supply elasticity, but if the transmission line is close to its limits, a few MWs can be enough to change the price dramatically. If the FTR is large in terms of MWs, the strategy would be profitable even if the bidder closes at a loss in the energy market. This strategy is described theoretically by Ledgerwood and Pfeifenberger [2012].

We present evidence consistent with the use of virtual bids to create congestion and increase the value of FTRs. Although the data do not allow us to identify the virtual bidders who hold FTRs, evidence suggests that the FTR and virtual markets are correlated. Nodes with FTRs have a higher ratio of virtual demand to supply, and lower virtual demand

profits.

This sort of price manipulation would not be possible in a market in which there is strong competition among financial traders. Whenever a trader uses virtual demand to increase the day-ahead price at a particular location, that would create incentives for a “pure” financial trader, one that does not hold FTR contracts, to profit from bidding in the opposite direction. In line with this expectation, we find evidence consistent with price manipulation particularly for the period in which there were high transaction costs for financial players.

Manipulation of FTRs using virtual bids has been the subject of a recent FERC case. In January 2014, FERC published a notice of alleged market manipulation by Louis Dreyfus Energy Services. The firm was accused of using virtual demand bids to increase congestion at a node where they held FTRs. The case was settled and the firm agreed to pay seven million dollars in remedies and penalties. We have verified that the facts that we present as suggestive of manipulation of FTR prices were present in the case of alleged manipulation by Louis Dreyfus Energy Services.

The next section describes the market in more detail and the following section presents the data. Section 4 shows that capital restrictions and regulation limit arbitrage and use this exogenous variation to present evidence for the effect of virtual bidders on the premium. Section 5 explains how regulation affected financial bidders, describes their behavior before and after policy changes, and presents evidence consistent with the use of virtual bids to increase the value of FTRs. Finally, we conclude in Section 6.

2 The MISO Electricity Market

The Midcontinent Independent System Operator (MISO) is one of the nine deregulated electricity markets in North America which together are responsible for the delivery of electricity to over 60% of the population of the United States and Canada. While each market has idiosyncrasies, the features of MISO which are studied in this paper are representative of most markets. The need for centralized markets comes from particular features of electricity; namely it is difficult to store and it must be delivered reliably across a transmission grid with limited capacity. To continually meet demand, the central authority coordinates transactions and flows to achieve an efficient and reliable use of the transmission network. MISO coordinates the electric power market in 15 states of the U.S. and the

province of Manitoba in Canada, providing electricity to more than 40 million people.⁴

Over 400 firms participate in the energy auction organized by MISO, where they may act in one or more of three distinct capacities: as a physical seller or generator, as a physical buyer, or as a virtual or financial bidder. Virtual sales or purchases are different from physical ones in that they do not require an energy injection or withdrawal, they are just financial instruments that firms can use to hedge or arbitrage price differences. Among virtual traders, there are firms with physical assets as well as purely financial ones ranging from small companies with one or two employees to large ones like Morgan Stanley. In order to understand how virtual bids take place, we will first briefly describe how the energy market operates.

Energy Market Operation

Wholesale electricity markets are designed to accommodate fluctuating demand and generation technologies which vary in their marginal costs and ability to change production levels. For these reasons, most deregulated wholesale electricity markets are organized as sequential markets: a day-ahead market allows careful planning across hours of the efficient mix of generating units and a real-time market which is able to respond to unpredictable demand. This is the case for MISO. First, a day-ahead market determines prices and schedules hourly production and delivery following operating day. Then, a real-time market receives generation bids 30 minutes before operation, and uses these bids to clear last-minute demand adjustments every five minutes. Between 98% and 99% of the total production is scheduled in the day-ahead market and the remaining 1% to 2% clears in the real-time market, which balances production to adjust for unexpected changes in demand or supply (i.e. a very hot day or an emergency plant shutdown).

Energy prices

Despite being a homogeneous good, a megawatt hour (Mwh) can have different prices throughout the market footprint due to transportation capacity limits. The pricing system in MISO, called locational marginal pricing, takes location into account to capture these differences in value. Under locational marginal price, the clearing price, or LMP, at each node is equal to the marginal cost of supplying one additional Mwh at that particular node

⁴Prior to 2013, when most of Arkansas, Mississippi, and Louisiana joined the RTO, MISO stood for Midwest Independent Transmission System Operator.

or location. Over the study horizon, for each day-ahead hour, MISO reports prices for between 1800 and 2000 nodes.

The LMP may be decomposed into three components: a marginal cost of energy, a marginal cost of congestion (MCC), and a marginal cost of losses (MLC). The MLC exists because some energy is naturally lost when it moves over transmission lines. Congestion occurs when demand at a node cannot be supplied from the generator with the lowest offered price without exceeding the capacity of a transmission line.


Table 1 provides summary statistics on day-ahead and real-time LMPs. On average, the day-ahead price is larger than the real-time price, which is also much more volatile. As expected, the day-ahead price is also more predictable: a regression of price on its lag from the previous day has an R-squared larger than 0.8 for the day-ahead price, while for the real-time price it is below 0.3.⁵

Congestion exists because transmission lines have limited capacity: when the demanded Mwhs exceed the capacity of the transmission line, after considering supply at the node, market clearing requires the price at the node to be higher and this increment is the congestion component of the LMP. While a more detailed treatment of LMPs in North American energy markets can be found in the tutorial Louie and Strunz [2008], let us illustrate congestion with a simple example. Consider a setting where there are only two nodes, A and B , connected by a transmission line with a capacity of 100 MW. At node A , all suppliers have a marginal cost of 20 and the market is perfectly competitive; so the energy is sold at marginal cost. At node B , there is only demand given by $Q(P) = 120 - P$. If there are no losses, the market clears with a price equal to the marginal cost and there is no congestion.

Now suppose that the marginal cost of producers at node A is \$10/MW instead of \$20/MW. In this case, consumers at B would demand 110 MW if the price was equal to the marginal cost; however, the transport of this electricity would require exceeding the capacity of the transmission line. Therefore, the price at node B has to increase to \$20/MW in order to clear the market and the extra \$10/MW over the marginal production cost corresponds to the congestion component of the LMP. This is illustrated in Figure ???. For an example with three nodes, see Louie and Strunz [2008].

⁵The results from this regression can be found in Table B1 in the online appendix

Example of a 2 node market for electricity.



Virtual bidding

Virtual bids are financial instruments that allow market participants to bid on the difference between the day-ahead and the real-time prices at different locations throughout the market. Players may virtually buy or sell energy in the day-ahead market, and revenues are computed as if the transaction were reversed in the real-time market. For instance, if a firm buys $10MWh$ in the day-ahead market, she earns profits for $10(P^{RT} - P^{DA})$, where P^{DA} and P^{RT} are, respectively, the day-ahead and real-time locational marginal prices. In terms of volume, virtual bids represent between 8% and 12% of the quantity purchased in the MISO market.

There are several reasons why virtual bids have been introduced into deregulated electricity markets. Importantly, virtual bidders, often without physical assets, are expected to help arbitrage systematic differences between the day-ahead and real-time prices, which eliminates incentives for physical players to game between the two markets and thus leads to better generation scheduling and lower production costs and emissions [Jha and Wolak, 2013]. Lastly, as we explain in more detail below, virtual bidding restricts the generators' ability to exert market power [Ito and Reguant, 2016, Mercadal, 2016]. Virtual bids also allow physical participants to hedge the real-time risk and to move the payoff of financial contracts which are often based on the day-ahead price to the real-time market.

Congestion and the FTR market

Congestion produces differences in prices, which in turn create revenues for MISO because the price paid by demand is higher than the one received by producers. MISO is a nonprofit organization; so these revenues have to be distributed among the market participants. This is done using two financial instruments: auction revenue rights (ARR) and

financial transmission rights (FTR). FTRs are financial instruments whose value depends on the difference in the day-ahead MCC between two nodes. An FTR consists of a volume (V), a source node (A) and a sink node (B), and will pay the difference between the MCC at the sink and the source: $V \times (MCC_B - MCC_A)$. If there are no losses in the network, the payoff can be calculated as the difference in the day-ahead LMP between both nodes.

The FTR instruments are auctioned in yearly and monthly auctions, and the revenues of the yearly auction are distributed among the holders of ARRs, which are in turn assigned to market participants depending on their historical use of the transmission network. Any market participant with enough credit can participate in FTR auctions; so they can either be used to hedge against congestion or just to speculate. Table 1 shows some statistics for FTRs sold in the monthly and yearly auctions. Each observation is an FTR held by a market participant. As can be observed from the tables, the range of FTR profits is wide. Capital requirements to participate in this market are large, because a participant needs to be able to pay large amounts in the case that profits are negative.

3 Data

The data used in this paper is publicly available at the MISO website.⁶ The sample contains data on the energy market between January 2008 and December 2012. The data includes bids, cleared quantities, clearing prices, and other bid characteristics for each participant at a node and hour level for both the day-ahead and real-time markets. Bid data is available for both physical and virtual bids and offers starting in August 2009, but the generator identifiers cannot be matched to physical demand and virtual participants; so it is not possible to identify generators who submit virtual bids or companies that own both generation and load assets. The dataset also includes the results of the FTR auction between 2009 and 2013, but we cannot identify holders of FTRs who are active in the energy market. Table 1 presents summary statistics of the data.

The TED spread used in Section 4.1 was constructed using treasury bills and LIBOR time series obtained from the website of the Federal Reserve Bank of St. Louis.⁷

⁶<https://www.misoenergy.org>

⁷Board of Governors of Federal Reserve System (US), 3-Month Treasury Bill: Secondary Market Rate [DTB3], retrieved from the website of the Federal Reserve of St. Louis on June 24, 2016.

4 Virtual Bidding in the MISO Electricity Market

Financial participants, trading on a similar footing to physical participants and labelled *virtual bidders*, have been introduced into MISO and other major electricity markets⁸ to improve price convergence between the day-ahead and the real-time prices. Price convergence is expected to increase market efficiency since production can be scheduled more accurately when the day-ahead price is a good predictor of real-time conditions.

Previous research has shown that the introduction of virtual trading contributes to a lower premium by comparing markets before and after financial participants are allowed to enter. Saravia [2003] studies how the introduction of virtual bidders affects the efficiency of the market. Using data on the New York ISO electricity market and a model that matches its structure, she finds that after virtual trading was introduced the premium decreased and profits earned by generators with market power fell.

Jha and Wolak [2013] investigate the effect of virtual bidding on the premium in the electricity market of California. They argue that though the premium did go down after virtual trading was introduced, virtual bidders will not necessarily bring the gap down to zero because there are transaction costs. In fact, besides the capital required to backup financial transactions, finding arbitrage opportunities requires some technical skills to analyze a complex market and large amounts of data. To estimate these costs, they propose three tests in which the implied marginal transaction cost is the minimum cost that would make profits from arbitrage equal to zero. They find that the implicit costs have decreased since the introduction of virtual bidding, which allows market participants to explicitly arbitrage instead of altering their physical bids for this purpose.

A lower premium has been associated to higher efficiency mainly because it allows for better planning. Jha and Wolak [2013] find that both production costs and emissions went down in California after financial traders were introduced. Nonetheless, not all virtual transactions improve market performance. Parsons et al. [2015] show that virtual bidders may also profit from discrepancies between the day-ahead and real-time models in ways that do not increase market efficiency, even if they contribute to a smaller gap between the day-ahead and the real-time prices.

The day-ahead premium has persisted in the MISO market for years despite the fact that virtual activity has been allowed since the market opened in 2005. Table 2 shows

⁸Markets using virtual bidding include PJM, and the markets in California (CAISO), Texas (ERCOT), and New York (NYISO)

the results obtained from regressions of the day-ahead premium on yearly dummies, which confirm the existence of the day-ahead premium previously documented in this and other markets⁹. This observation points to the question of whether virtual bidders are effectively arbitraging the day-ahead premium in this market, and if not, why. In the following section we investigate the first part of this question using an instrumental variable regression. In contrast to the findings of Saravia [2003] and Jha and Wolak [2013], we find that financial trading has a weak effect on the premium, if any. In Section 4.2 we discuss regulation and capital restrictions as potential limits on virtual bidder arbitrage.

4.1 Virtual Bidders and the Day-Ahead Premium

Answering these questions is not straightforward: the effect of financial participants on the premium cannot be assessed by simply looking at the correlation between the premium and the virtual volume because a larger premium will attract more virtual bidders; so we might observe a positive correlation between the volume of financial trading and the forward premium even if the causal correlation is negative. In fact, we see in Table 2 that the average premium is larger at nodes with virtual sales than in the market as a whole, which in principle can be either because virtual activity is concentrated at nodes with a higher premium or because financial trading somehow results in a higher premium.

We address this endogeneity problem using two sources of exogenous variation in the volume of virtual transactions as instrumental variables. The first is the limited access to credit during the financial crisis of 2007-2008, which is likely to have restricted the volume of virtual transactions by decreasing capital availability.¹⁰ In addition to the minimum capital required to become a market participant, the total MWs each virtual bidder can trade depend on credit limits imposed by MISO according to the participant's capital.¹¹ We capture this effect using the TED spread, which is the difference between the 3 month

⁹See Bowden et al. [2009] for the Midwest market, Saravia [2003] for New York, Jha and Wolak [2013], Borenstein et al. [2008] for California, and Ito and Reguant [2016] for the Iberian market, among others.

¹⁰See the State of the Market Report 2008, 2009, by the independent market monitor Potomac Economics.

¹¹To buy or sell energy in the MISO market, it is necessary to first register as a market participant. While details of the formal registration process vary depending on the participant's intended market activities, it generally requires the submission of supporting documents and a credit check. A company who wants to participate in the energy market needs to demonstrate total assets of \$5,000,000 or a tangible net worth of \$500,000. If not able to provide audited financial statements at the required level, the participant must post a financial security of \$50,000, of which \$25,000 are restricted. Requirements increase to a minimum financial security of \$500,000 for firms who want to participate in the FTR market. Besides minimum requirements for participation, MISO runs a credit check to determine the number of MWs a participant is allowed to trade.

LIBOR - London Interbank Offered Rate- and the 3 month U.S. Treasury Bill rate. The TED spread has been used as a measure of the perceived risk of credit, and it spiked during the financial crisis, as it can be observed in Figure 1.

The second exogenous shock is a regulatory change that took place in in November 2008 and lead to an increase in transaction costs for financial participants. This regulatory change affected Revenue Sufficiency Guarantee (RSG) charges, which are imposed on firms that deviate from their day-ahead schedule in the direction of increasing net real-time demand, i.e. producing less or consuming more than scheduled. These deviations are costly because plants need to be quickly ramped up or started in order to cover demand, which is typically expensive given the available technologies. This is the reason why deviations are subject to RSG charges, which are then used to compensate the plants that incurred in higher start-up or ramping costs due to unexpected real-time demand.

When the MISO market first opened in 2005, virtual transactions were not subject to RSG charges. After a long discussion that started in 2006, in November 2008 FERC determined that virtual supply transactions contributed to higher production costs and should therefore be subject to RSG charges. As virtual transactions never inject energy, the whole volume sold in the day-ahead market counted as a deviation, and thus these charges acted like a per unit transaction cost. This did not only apply to future transactions, but also to those going back to April 2006. In January 2009 MISO started the resettlements from the application of these charges to past transactions, which worsened the financial conditions of some firms to the point where they started defaulting in their payments to MISO. FERC then decided against the resettlements, but the discussion over the charges continued until April 2011, when the calculation of the charges was modified in a way that significantly lowered them for virtual transactions and the uncertainty was resolved. As a consequence, virtual volume rose again. These changes are marked with vertical lines in Figure 1.

Results from the instrumental variables regression of the day-ahead premium on the virtual volume are presented in Table 3, which shows separate regression for virtual demand, virtual supply, and the total volume traded by virtual participants. The instruments are the TED spread and its 5, 10, 15, 30, 60, 90, and 120 days lags, and a dummy for the period before November 2008, when the order that made virtual offers pay RSG charges came out. We have also included physical demand as a control, because market conditions are largely influenced by the level of demand. The IV regressions indicate that virtual volume has a negative effect on the day-ahead premium, but it is not statistically significant.

The bottom part of Table 3 presents some diagnostic tests for the IV regressions. The F-tests from the first stage regression indicate that the instruments are not weak (Results from the first stage are presented in Table 5). The Sargan test of over-identifying restrictions indicates the instruments are valid. However, the Wu-Hausman test fails to reject the null that the IV and OLS estimators are equally consistent for total virtual volume and virtual supply, suggesting that endogeneity is not an important problem and therefore OLS is preferable, since it is more efficient.

Results from the OLS regressions are presented in Table 4. Though standard errors are smaller, virtual volume is still not statistically significant for the total volume and the virtual demand. The coefficient on virtual supply is weakly significant, with a p-value below 10% but above 5%. These results seem to indicate that there is an arbitraging effect of virtual bidding, and it comes from the supply bids, as we would expect in the presence of a positive day-ahead premium. On the other hand, the effect is weak and not very robust to changes in the specification. In the next section, we will examine virtual bidders' behavior in detail in order to understand why they are not closing the day-ahead premium.

4.2 Limits to Virtual Bidder Arbitrage

Although virtual bidders seem to contribute to convergence between the day-ahead and the real-time prices, the day-ahead premium still persists in the MISO market. Limited capital availability may restrict entry to the market or the volume that a participant can trade. Moreover, if entry is limited, existing participants will not arbitrage until the premium is zero, just like a monopolist does not choose a quantity that drives the price to the marginal cost [Borenstein et al., 2008].

Besides capital requirements, we identify two additional sources of limited arbitrage in the MISO market between 2010 and 2012. The first comes from the market design: deviation or RSG charges acted as transaction charges for virtual trades and therefore reduced profits from arbitraging the premium. The second reason is that some virtual bidders appear to be using their bids to affect congestion prices and increase the price of FTRs instead of arbitraging in the energy market.

Figure 2 shows that the level of virtual activity has had two major changes in levels, one by the end of 2009 and the other in the spring of 2011. These changes coincide with the two major changes in the rules regarding RSG charges, though as discussed above the financial crisis had an effect on virtual trading as well. In November 2008, RSG charges were

imposed for the first time on virtual supply bids, and in April 2011 the computation of the charges changed in a way that effectively reduced them significantly. As the figure shows, the traded volume responded to this changes by falling when they were initially imposed and going up when they were decreased. This section describes the mechanism whereby RSG charges restricted virtual activity. In order to do this, we will distinguish between two periods. The first period goes from January 2010 to March 2011, and the second period starts on April 2011, when RSG charges were modified, and ends in December 2012.¹²

When the day-ahead premium is on average positive, arbitrageurs profit from selling in the day-ahead market. Nonetheless, supply bids were not necessarily profitable during the first period because they were subject to high RSG charges. Table 9 shows the day-ahead premium net of RSG charges. As it can be observed, the net premium was negative during the first period, which means negative profits for supply bids because they sell in the day-ahead and buy in the real-time market. When the sample is restricted to nodes in which there were virtual supply bids, the premium net of RSG charges is statistically not different from zero, which is consistent with virtual bidders exploiting arbitrage opportunities as far as possible. In the second period the premium decreases, but RSG charges are low enough to make some virtual offers profitable, as we will show in the remaining of this section.

Figure 3, which shows the cumulative day-ahead premium net of RSG charges, which is equivalent to the profits an arbitrageur would have made by selling 1MW at each node in every hour. Before April 2011, indicated by the vertical line, this was clearly not a profitable strategy. Nonetheless, cumulative profits seem to increase from April 2011 until the second half of 2012, when arbitrageurs had had enough time to enter and close the largest arbitrage opportunities. From April 2011 to December 2012, profits would have been \$7,604,298.

Just the fact that buying at every node would have made large profits is not enough to establish that there were profitable opportunities, since this could not necessarily have been predicted. Consequently, we test two arbitrage strategies based on simple predictions of the premium, and find some profitable opportunities after the change in RSG charges that quickly disappear after a year or so. Figures 4 and 5 show the cumulative profits from these strategies in the first and second period, respectively. The top panels show what an

¹²We exclude 2009 because there was uncertainty over how the market was going to work. First, because the ancillary services market opened that year. And second, because at the beginning of the 2009 the discussion over whether past financial transactions were going to be subject to charges was still active, and the uncertainty over the charges might have played a role in addition to the charges as transaction charges.

arbitrageur would have made by following a very simple rule: buying if the premium was negative 48 hours ago, and selling if it was positive. This strategy does not work well in the first period, but it is profitable after RSG charges are reduced. A participant entering the market in April 2011 would have made \$4,150,737 by the end of 2012 if she had also left the market for two days whenever she lost \$10,000 or more on a given day. The Sharpe ratio of this strategy is 1.04.¹³

The second arbitrage strategy buys when the predicted premium is below -1 and sells when it is above 3. The prediction is made for the next 24 hours using a regression of the premium on its own lags with data on the past 30 days (details in the appendix). On average, the arbitrage strategy sells in the day-ahead, as expected. Cumulative profits are shown in the bottom panel of Figures 4 and 5. The blue line shows profits from a strategy that uses the prediction and buying strategy above, and leaves the market for 2 days if total losses on a given day are higher than \$10,000. As before, this precaution seems to increase profits by avoiding large losses when unexpected shocks raise real-time prices. Compared to the simpler strategy described above, using a linear regression to predict the premium reduces the losses in the first period, but also earns lower profits in the second period. Until mid-2012, the Sharpe ratio was 0.61 and the strategy had made profits of over a million dollars. However, in the second half of 2012 there were substantial losses erasing its earlier profits which supports the idea that arbitrage opportunities disappeared about a year after the policy change. Although the initial cost of both strategies is low—for the second it is \$129 to buy the portfolio for the first day and \$2,332 for the first week—they require enough capital to have losses of up to \$10,000 in just one day.

Jha and Wolak [2013] propose a method to compute the implicit transaction costs faced by bidders who arbitrage between a forward and a spot market. For an asset equal to 1 MW each hour of a day at a certain node, the implicit transaction cost is the minimum cost such that the expected profits from arbitrage are zero. There are arbitrage opportunities when the mean premium during the 24 hours is statistically different from zero, which can be determined using three different statistical tests. We run the same tests for the MISO market and find lower values for the implied transaction costs as would be expected in a market that has had virtual traders for some years. Additionally, we find that the implied transaction costs decreased in the second period (Table 6).

Tables 7 and 8 present some descriptive statistics for the financial bidders in the first

¹³The Sharpe ratio is calculated as $\frac{\bar{X}\sqrt{365}}{\sigma_X}$, where \bar{X} is the mean of daily total profits and σ_X the standard deviation of daily total profits.

and second period, respectively. If we take the implicit transaction costs from Jha and Wolak [2013] as daily costs per node and MWh, and use the lowest of the average costs from Table 6, an average participant would have faced a total implicit transaction cost of \$1,100,000 during the first period and \$1,700,000 during the second. Implicit transaction costs decrease from the first to the second period, but the average participant gets larger. In the first period, this back of the envelope calculation assumes the average participant bids 32MW (average bid size), in 168 nodes, and is active 147 days. These numbers are taken from Table 7 for the supply bids, which, because of the positive day-ahead premium are on average likely to be arbitrage.

5 Virtual Bidders and the FTR market

RSG charges acted as a barrier for financial participants, making it unprofitable to bid against the premium and thus to close it. When they were lowered, virtual volume increased and participants took advantage of the arbitrage opportunities that appeared. As a consequence, the premium decreased. During times when arbitrage is limited, market power is greater, and some participants can use their bids to manipulate prices without having an arbitrageur bidding in the opposite direction. Particularly in the first period, we observe some virtual participants not acting as speculators would be expected to. We first rationalize this behavior as using financial bids to create artificial congestion and increase the value of FTR instruments, and then present evidence supporting this hypothesis. Note, that when there is a limited pool of virtual bids due to capital constraints, if participants substitute virtual bidding moves due to the attractiveness of increasing the value of FTR positions, then, this might further reduce price convergence as these virtual bids are not used for arbitrage.

5.1 FTRs and Market Power

Virtual bidders may not close the gap between the day-ahead and the real-time price partially because some of them are using virtual bids to change their FTR position instead of trying to arbitrage. We are aware of two strategies that relate virtual trades and FTRs. The first one is to transform a standard FTR, which pays on the difference in day-ahead congestion between two nodes, into a real-time FTR, i.e. to make the position depend on the difference in real-time congestion between the two nodes. The second, is to manipulate

the FTR payoff by using virtual bids to strategically to increase congestion.

Moving an FTR to the real-time market is a strategy frequently utilized by physical players to hedge real-time congestion risk. To understand this, return to the two node example illustrated in Figure ???. Consider an individual who holds an FTR for 1 MW from node A to node B and wants to move her position to the day-ahead. She can do this by virtually buying 1 at node B and selling 1 at node A with two virtual bids. Then her profits would be:

FTR	$LMP_B^{DA} - LMP_A^{DA}$
Sell 1 at A	$LMP_A^{DA} - LMP_A^{RT}$
Buy 1 at B	$LMP_B^{RT} - LMP_B^{DA}$
Total	$LMP_B^{RT} - LMP_A^{RT}$

This strategy does not have manipulation as an objective, yet it creates artificial congestion because bidders inject MWs at the source node and demand at the sink node. Moreover, as these bids are far from the clearing price, they do not contribute to price convergence between the day-ahead and the real-time market.

We have identified all pairs of bids that could potentially be FTR moves and they amount to about 0.2% of cleared bids, so they do not seem to be a common strategy in this market.

In the second strategy, virtual bids are used to increase the volume in sufficient quantities to create congestion and thus increase the FTR payoff. Consider again the example from Section 2 where the marginal cost of generating at node A is 10 and a financial participant bought an FTR for 100 MW between nodes A and B at a price of \$10 per MW. With demand determined by $MWh = 120 - P$, the day-ahead LMP at node A is $LMP_A^{DA} = 10$ and node B is $LMP_B^{DA} = 20$ resulting in an FTR yield of zero payoff. Now if, using a virtual bid, the FTR holder buys 10 MWh at node B as a price-taker (or bidding a very high price), the demand would be shifted to $MWh = 130 - P$ so the LMP at node B has to increase to \$30. Suppose that the real-time LMP at node B is $LMP_B^{RT} = 8$ is unchanged by this action. Then, the final profits for this participant are the following:

Virtual bids	$(LMP^{RT} - LMP^{DA}) MWh = (8 - 30) 10 =$	-220
FTR	$(LMP_B^{DA} - LMP_A^{DA}) MW - P^{FTR} MW = (30 - 10) 100 - 10 \cdot 100 =$	1000
Total	$1000 - 220 =$	780

Although this participant is losing money on its virtual bids, the increase in the FTR profits compensates for the virtual loss. In reality, more things have to be considered: FTRs are costly and have a duration of a month or a quarter; so it is necessary to manipulate congestion during a longer period and to have enough capital to sustain large virtual losses as well as to buy the FTR. For this strategy to be profitable, it is necessary to have sufficiently inelastic demand to produce congestion with only a small volume virtual bid.

5.2 Evidence

This section presents empirical evidence consistent with the hypothesis that speculators are using virtual bids to increase the value of financial transmission rights. In principle, a correlation between virtual bidding and FTRs would not be expected. Virtual bids payoffs depend on the difference between the day-ahead and the real-time price for a certain node, while FTR profits are computed from the difference in the day-ahead price between two different nodes. Moreover, the day-ahead premium is not a congestion premium. If we subtract the congestion and losses components from the day-ahead premium, we find almost the same values as for the average day-ahead premium. And if we repeat the instrumental variable exercise of Section 4.1 to test the effect of virtual participants on the congestion premium, we find the traded virtual volume has no effect. In spite of this, many aspects of these two markets seem to be correlated, as we show below. We will first describe the aspects of the virtual bidders' behavior that are hard to rationalize within the energy market and then present the connection to the FTR market.

To illustrate the evidence of manipulation, the different pieces of evidence will be connected to the case of alleged market manipulation by Louis Dreyfus Energy Services L.P. FERC published a notice of alleged manipulation on January 2014, and a month later they reached a settlement in which the company agreed to pay a penalty of four million dollars and a disgorgement of three million dollars. Additionally, the trader, who wrote his graduate thesis on FTR manipulation with virtual bids, paid a civil penalty of 300,000. The company was accused of placing demand bids at the sink node of some of its FTRs, incurring losses in virtual bids to increase the day-ahead price. Trading activity decreased

shortly after the FTRs held by Dreyfus expired.

Virtual bidders bought more than what they sold

As Figure 2 shows, virtual purchases accounted for a larger volume than virtual sales during most of the period under study. This is not what we would expect, since with a positive day-ahead premium virtual demands lose money on average. Though the profitability of virtual supply bids is restricted by RSG charges, we find that *ex-post* virtual sellers have positive and higher profits while virtual buyers lose money in both periods (Tables 7 and 8). Therefore, demand bids were not generally aimed at particular nodes with a negative premium. Demand bids were also on average larger in volume. This is especially prominent in the first period, which may indicate that demand and supply bids are of different types.

A big part of the difference between virtual demand and virtual supply comes from virtual bids placed by physical buyers (Table 10). Moreover, about 60% of the virtual demand volume cleared by physical buyers corresponds to one large firm focusing its bids on three nodes. This firm's strategy was to buy physical energy at a load zone or demand node as a price-taker, and to place price sensitive virtual demand bids at two or three generation nodes in the same area, clearing on average about 5% of the quantity bought with the physical bid. Most likely, this is a pretty expensive hedging strategy. As utilities do not know exactly how much demand they will need to serve, they may complement physical purchases with price sensitive virtual demand bids to hedge real-time risk. Unfortunately, MISO does not publish real-time purchases at the firm level, so we cannot test whether virtual demand bids help to reduce cost volatility for utilities. For this reason, we exclude virtual bids placed physical buyers.

After excluding physical buyers, the difference between virtual demand and supply is not as large, though still significantly larger than zero, and it seems to have remained around the same level after RSG charges were reduced (Table 10). Nonetheless, when we look at the individual firm level we find a change in the second period. Figure 6 shows the cumulative distribution of the ratio of virtual demand to total virtual volume by firm, which moved to the left after RSG charges were reduced. Firms moved from clearing slightly more demand than supply to slightly more supply than demand, as well as towards more balanced positions.

Louis Dreyfus was alleged to have used demand bids to create artificial congestion at a node called Velva, which was the sink node of some FTRs the company held. The FTRs

were valid from December 2009 until the end of February 2010. Starting by the end of November, Dreyfus placed virtual demand bids that on average lost \$132 per bid, totaling losses for \$273,000.¹⁴ During that period the average day-ahead premium was \$1.28, yet it was \$6.73 at the node where Dreyfus was active and during the hours in which Dreyfus cleared bids. This is consistent with their bids pushing the price upwards. The company cleared a similar volume in demand and supply bids during the period under FERC scrutiny, but its total losses were over \$1,600,000 from the demand bids and around \$170,000 from supply bids. Figure 7 shows the company's profits since it became active until the end of 2012. It did not start making consistently positive profits over time until after the RSG charges were reduced.

Some participants steadily lost money

Louis Dreyfus Energy Services was not the only company with large losses. Some participants steadily lost money over long periods without leaving the market. Figure 8 plots the number of days in which a participant was active against the total profits made by that participant during the whole sample period. As it can be observed, some bidders stayed for more than a year in the market even though they were having large losses. The cumulative profits of some of the ten largest virtual losers among virtual bidders consistently decreased (see Figure 9), indicating that some players stayed in the market despite having steady losses. This is not consistent with profit maximization by pure speculators, especially considering that they are experienced bidders in a complex market. Even a naive individual would leave the market after accumulating losses for 100 days if there are no other sources of profit. Nonetheless, losses would be justified if they were the cost of creating artificial congestion to increase the value of FTRs. As long as the cost is lower than the increased FTR profits, the observation would be consistent with profit maximizing behavior.

These observations suggest that some virtual bidders were doing something different from maximizing profits in the electricity market. Although in principle this could be

¹⁴MISO does not publish information on how to match the IDs in the FTR market to the ID's in the energy market. Although we can identify the company in the FTR market from the name, in the energy market we just inferred the ID by looking at participants who placed virtual demand bids at the Velva node from November 2009 to February 2010. The participant chosen was active the most days during the period, was not active after that (as FERC stated in its report), only submitted demand bids, and was the one clearing the most MWs. Additionally, the losses add up to \$390,000, as in the FERC report. This participant cleared 60,000 MW during the period between November 2009 and February 2010, with the next one clearing 7,000 MW, making it likely to attract the market monitor's attention. No other participant matched the facts described in the FERC's report as well as this one.

due to some physical participants' hedging, the price of the insurance would be high. The largest loser, for instance, accumulated 15 million dollars in losses, and was a physical buyer most likely using this bids to hedge. Though there are some physical buyers among those consistently losing from virtual bids, there are also other firms that can be either generators or purely financial traders. The hedging explanation cannot be fully discarded for these firms, because the data does not allow to know us whether big losers have generating capacity. The cumulative profits look almost the same when they are drawn using profits before discounting the RSG charges, implying that the losses are not caused by the RSG charges. The large losses sustained by some participants are not a general feature of the market either. As Figure 10 shows, participants with the largest profits earned them by steadily accumulating positive profits.

Large participants

Using virtual bids to increase the value of an FTR has large capital requirements. In order to trade FTRs, a firm needs to show that it is able to cover for potential losses in case congestion does not go in a favorable direction, and the stakes are often large (see 12). Moreover, FTRs are traded in auctions held monthly and the secondary market is not very liquid. Additionally, such a strategy involves submitting virtual bids during several hours of the day for the duration of the FTRs, and as we show above capital constraints can limit virtual activity. For these reasons, we would expect only relatively large firms to be able to use virtual bids to increase FTR's value.

In line with this hypothesis, we find that larger firms, in terms of average cleared quantity every day, clear relatively more demand than supply. Louis Dreyfus Energy Services is, in fact, a large bidder. They were active in 296 nodes, cleared 11,650 MWs a day on average, and held FTRs for 95,000 MW during the period of the accusation.

Price insensitive bids

Virtual bidders can submit price sensitive bids that specify up to ten combinations of price and quantity, or just a quantity that they are willing to but at any clearing price. To create congestion, a virtual demand bidder needs to increase the number of MWs cleared in that node; so he would be more likely to bid a high price or to be a price-taker. The price will be higher than without congestion, but it is very likely that the price will be

lower than what he or she is offering unless supply is very inelastic.¹⁵ Therefore, we would expect to observe some participants submitting bids much higher than the clearing price, even if they do not determine the price. This can be used to guess which participants are trying to increase congestion, although a high bid does not necessarily imply congestion manipulation. In line with this reasoning, we observe that a group of participants placed bids far from the clearing price and that the largest losers of the market were in this group.

Figure 11 shows a plot of the relation between virtual profits and the bid premium for price-sensitive bids. The bid premium is the difference between the highest price offered in the bid step function and the clearing price. Figure 11 appears to confirm the reasoning above by showing that there are two types of bidders: those who are, on average, offering a price close to the clearing price—the true virtual participants—and those who bid high prices. Moreover, large losses seem to be associated with a large bid premium. This is confirmed by the regressions of virtual profits on the premium, which also show that the effect is larger for demand bids (Table 11).

In the MISO market, the average bid premium is \$43 for demand bids and \$30 for supply bids in the first period; in the second period it is \$49 for demand and \$26 for supply. The average day-ahead price is \$29 (Table 1).

In the case of Dreyfus, the average bid premium during the period of the accusation is \$65. This is not so high, but it is because Dreyfus was setting the price 6% of the times. Looking at all the nodes at which they participated during the same period, they only set the price 2% of the time. Taking the same period, less than 1% of the bids clearing a positive quantity set the price.

Although the bid premium seems to be distinguishing two types of bidders, those with a high premium are not necessarily using virtual bids to manipulate congestion. In the next section, we will describe the different ways in which virtual transactions can be related to the FTR market.

FTR profits are enough to cover virtual losses

The FTR payoff depends on the difference in day-ahead congestion between the source node and the sink node; so to increase FTR profits it is necessary to increase either demand at the sink node or supply at the source node. Table 12 shows some summary statistics for

¹⁵A virtual bidder will offer cheap MWs at the source node. Equivalently, the clearing price is unlikely to be as low as his offer, but it will be lower and this will increase demanded MW at the sink node.

FTR profits and virtual profits at the sink or the source node. Because we cannot match the participant identifiers in the FTR market to those in the energy market, we compute the virtual profits made by each participant at the sink and source node, take the mean and the minimum across participants, and compare that to the profits made by each FTR. The idea is to know whether FTR profits would be enough to cover the losses sustained by the mean participant at that node, and also by the participant with the largest losses. The table shows that the addition of virtual demand profits lowers total profits, especially at the sink node, but FTR profits still appear sufficient to potentially compensate for virtual losses.

FTRs and virtual bids

The value of an FTR is higher when congestion is higher at the sink node, or when it is lower at the source node. Therefore, virtual bids intended to create congestion are either demand bids at the sink node or supply bids at the source node. We do not have enough information to check whether participants submit virtual bids at the nodes where they hold FTRs, so we examine whether there is correlation between FTR and virtual trading. In this exercise, as before, we exclude virtual bids submitted by physical buyers since they submit a large number of virtual demand bids but seem to be hedging.

First, we find that nodes with more FTR MWs are associated to a larger ratio of cleared virtual demand to virtual supply at the sink node in the first period (Table 13), but find no significant correlation at the source node. This is what we would expect if virtual demand bids are used to increase the day-ahead price at the energy price in order to increase the FTR's payoff. Moreover, one of the reasons why Louis Dreyfus was investigated was that they increased their virtual purchases at a node where they had FTR positions, and then stopped trading there once these positions decreased in size.

This correlation between the ratio of virtual demand to supply and FTR MWs is only significant during the first period, i.e. when transaction charges prevented financial traders from fully arbitraging the day-ahead premium. In a perfectly competitive market for virtual traders it is not possible to increase the value of FTRs using virtual bids, since whenever a trader manages to artificially increase the price at one node, there would be a new profitable arbitrage opportunity for another trader to exploit. Therefore, the fact that the correlation disappears is consistent with increased virtual activity restricting players' ability to game the FTR market using virtual bids.

If some firms are using virtual demand bids to increase congestion and FTRs value instead of arbitraging profitable opportunities, we would expect a negative correlation between virtual demand profits at the sink and FTR profits. Table 14 shows that this is what we find during the first period: nodes with higher FTR profits are associated to lower virtual demand profits at the sink node. This correlation disappears in the second period as well, which is again consistent with increased competition making manipulation harder. We do not find any significant correlation between virtual profits and FTR profits at the source node.

The case of Louis Dreyfus Energy Services is presented in Table 15. During the period of alleged manipulation, we look at FTRs that had the Velva node as either a sink or a source. We find that FTR profits are associated with lower virtual demand profits at the sink, as we would expect from a firm placing demand bids to increase the price even at the cost of virtual losses.

Congestion and virtual profits

We find that nodes with more congestion in the day-ahead market are associated with lower profits from virtual demand, higher profits from virtual supply, and a higher ratio of virtual demand to virtual supply (See Table 16). This is consistent with the use of financial bids to create artificial congestion in sink nodes and to decrease artificial congestion in source nodes in a market with a positive day-ahead premium. Demand bids are submitted at sink nodes, where they increase congestion but can earn negative profits. Offers are submitted to decrease congestion at source nodes, and on average earn positive profits because the premium is positive. In the second period, these correlations are still significant, but much smaller in magnitude, which is again consistent with increased competition making it harder to use virtual bids to manipulate congestion prices. Finally, these results do not come from congestion being the driver of the day-ahead premium, since in fact the ratio of virtual demand to supply is not significantly correlated to the day-ahead congestion premium.

6 Conclusion

This paper contributes to the discussion about the controversial role of financial participants in commodity markets. Using detailed data on trader behavior, we show that financial participants contribute to lowering the forward premium, but that their effect is

limited by capital availability and regulation. Further, when financial participants ability to arbitrage the premium is limited, large actors may find it more profitable to use their bids to increase the value of related instruments. Evidence points to this behavior in the MISO markets where financial traders appear to be acting to increase the value of FTR instruments using virtual bids in the electricity market. This behavior reduces the pool of arbitraging bidders and can act to reinforce the forward premium. Proper regulation and market design appear key to achieving the efficiency gains that financial participants are expected to bring. The evidence of manipulation of related financial instruments in MISO markets disappears when regulation better supports arbitrage of the forward premium. Our findings point to the challenges of achieving market efficiency through purely financial players when the underlying physical market is segmented and subject to the exercise of market power in each segment.

References

- Hendrik Bessembinder and Michael L Lemmon. Equilibrium pricing and optimal hedging in electricity forward markets. *the Journal of Finance*, 57(3):1347–1382, 2002.
- 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. Diagnosing market power in california’s restructured wholesale electricity market. Technical report, National Bureau of Economic Research, 2000.
- 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.
- Potomac Economics. 2008 state of the market report for the miso electricity markets, 2009.
- Potomac Economics. 2009 state of the market report for the miso electricity markets, 2010.
- Potomac Economics. 2010 state of the market report for the miso electricity markets, 2011.
- Potomac Economics. 2011 state of the market report for the miso electricity markets, 2012.
- Potomac Economics. 2012 state of the market report for the miso electricity markets, 2013.
- David Hirshleifer. Residual risk, trading costs, and commodity futures risk premia. *Review of Financial Studies*, 1(2):173–193, 1988.
- Koichiro Ito and Mar Reguant. Sequential markets, market power, and arbitrage. *American Economic Review*, 2016.
- 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.
- S.D. Ledgerwood and J.P. Pfeifenberger. Using virtual bids to manipulate the value of financial transmission rights. 2012.
- Haifeng Liu, Leigh Tesfatsion, and AA Chowdhury. Locational marginal pricing basics for restructured wholesale power markets. In *Power & Energy Society General Meeting, 2009. PES'09. IEEE*, pages 1–8. IEEE, 2009a.
- 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, 2009b.
- Francis A Longstaff and Ashley W Wang. Electricity forward prices: a high-frequency empirical analysis. *The Journal of Finance*, 59(4):1877–1900, 2004.
- Henry Louie and Kai Strunz. Locational marginal pricing in north american power systems. *ETG-Fachbericht-Netzregelung und Systemführung*, 2008.

- Ignacia Mercadal. Dynamic competition and arbitrage in electricity markets: The role of financial players. 2016.
- Andrew L Ott. Experience with pjm market operation, system design, and implementation. *Power Systems, IEEE Transactions on*, 18(2):528–534, 2003.
- 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.
- Celeste Saravia. Speculative trading and market performance: The effect of arbitrageurs on efficiency and market power in the new york electricity market, 2003. URL <http://www.escholarship.org/uc/item/0mx44472>.
- Andrei Shleifer and Robert W Vishny. The limits of arbitrage. *The Journal of Finance*, 52(1):35–55, 1997.
- Frank A Wolak. Measuring unilateral market power in wholesale electricity markets: the california market, 1998-2000. *American Economic Review*, pages 425–430, 2003.
- 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.

Tables

Table 1: **Summary Statistics.** For the FTR, the clearing price is the price paid per MW of the financial right.

Variable		N	Mean	St. Dev.	Min	Max
LMP	DA price	48,547,536	29.9	14.4	-977.1	756.9
	RT price	48,547,536	29.2	27.0	-1,318	3,225
Virtual Supply	MW	7,400,827	9.9	23.7	0.1	2,550
Virtual Demand	MW	6,816,498	14.3	51.3	0.1	1,888
FTR	MW	597,104	6.440	14.955	0.100	642.9
	Price	597,104	432.0	1,349	-22,938	21,020
	Profits	597,104	698.3	33,400	-2,389,008	5,382,571
RSG	RSG1	10,920	1.8	2.3	0.0	21.0
	CCF	10,470,096	-0.01	0.1	-1.0	1.0
	CMC	10,470,096	2.2	12.5	0.0	961.3
	DDC	10,470,096	0.9	1.5	0.0	27.4

Table 2: **Day-ahead premium over the years:** Results from a regression of the hourly day-ahead premium on year dummies. In the first three columns, the hourly premium is computed as the simple average across nodes, the median across nodes, and a weighted average using the total quantity traded by physical participants as weights. In the last three columns, the sample is restricted to nodes with virtual activity and the day-ahead premium is computed as a weighted average, and the three columns present three different weights: virtual demand in MWh (VD), virtual supply in MWh (VS), and total virtual volume in MWh (VV). HAC standard errors in parenthesis.

	Mean	Median	Weighted	VD	VS	VV
2008	1.38** (0.55)	1.54*** (0.56)	1.42*** (0.54)	0.77 (0.56)	2.29*** (0.56)	1.42** (0.56)
2009	0.86*** (0.26)	1.14*** (0.25)	0.84*** (0.25)	-0.01 (0.27)	1.89*** (0.29)	0.72*** (0.26)
2010	1.14*** (0.30)	1.42*** (0.29)	1.26*** (0.29)	0.53 (0.32)	2.37*** (0.33)	1.24*** (0.31)
2011	0.52 (0.36)	1.13*** (0.34)	0.69* (0.36)	-0.03 (0.38)	1.48*** (0.38)	0.64* (0.38)
2012	0.39 (0.34)	0.83** (0.33)	0.49 (0.34)	-0.06 (0.39)	1.28*** (0.37)	0.53 (0.37)
N	43,848	43,848	43,848	43,824	43,824	43,824
R^2	0.003	0.005	0.003	0.0005	0.01	0.003

Notes: ***Significant at the 1 percent level., **Significant at the 5 percent level,
*Significant at the 10 percent level.

Table 3: **IV regressions: Day-ahead premium and virtual bidding** Results from regressions of the median daily day-ahead premium on the virtual volume and total daily demand (both in logs of Mwhts). The instruments are the TED spread and its 5,10,15,30,60, 90, and 120 days lags, and a dummy for the period before November 2008, when the order that made virtual offers pay RSG charges came out. The first column uses total virtual volume, while the second and third include only supply and demand, respectively. HAC standard errors reported in parenthesis.

	Median day-ahead premium		
Total virtual volume	-1.97 (1.61)		
Virtual supply		-1.71 (1.32)	
Virtual demand			-1.89 (1.68)
Physical demand	14.00*** (5.19)	13.29*** (4.73)	13.97** (5.48)
Constant	-173.39*** (65.33)	-167.72*** (62.72)	-175.00** (68.27)
Weak instruments	131.42***	86.6***	80.71***
Wu-Hausman	1.35	0.2	7.62**
Sargan	5.07	4.87	5.4
N	245	267	245
R^2	0.03	0.04	0.02

Notes:

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Table 4: **OLS regressions: Day-ahead premium and virtual bidding** Results from an OLS regression of the hourly median day-ahead premium on the virtual volume and total daily demand (both in logs of Mwhts). The first column includes total virtual volume, while the second and third include only supply and demand, respectively. HAC standard errors in parenthesis.

Median day-ahead premium			
Total virtual volume	-1.21 (1.45)		
Virtual supply		-1.98* (1.14)	
Virtual demand			0.47 (1.55)
Physical demand	12.39** (4.86)	13.86*** (4.76)	8.74* (4.81)
Constant	-159.45** (63.02)	-172.70*** (63.74)	-127.59** (62.04)
N	267	267	267
R^2	0.03	0.04	0.03

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 5: **First stage: Virtual bidding on the TED spread and RSG changes**
 Results from an OLS regression of the log of the hourly volume traded by virtual bidders in Mwhts on the TED spread and its 5,10,15,30,60, 90, and 120 days lags, and total daily demand in logs of Mwhts. The first column uses total virtual volume, while the second and third include only supply and demand, respectively. HAC standard errors reported.

	Virtual volume	Virtual supply	Virtual demand
TED spread	-0.01 (0.03)	-0.02 (0.04)	-0.005 (0.03)
TED spread _{t-5}	0.04 (0.05)	0.07 (0.06)	0.02 (0.04)
TED spread _{t-10}	0.01 (0.02)	-0.02 (0.04)	0.03 (0.02)
TED spread _{t-15}	-0.03 (0.03)	-0.04 (0.04)	-0.04 (0.03)
TED spread _{t-30}	0.11*** (0.03)	0.12*** (0.04)	0.11*** (0.03)
TED spread _{t-60}	-0.03 (0.02)	-0.08** (0.03)	-0.001 (0.03)
TED spread _{t-90}	-0.11*** (0.02)	-0.03 (0.03)	-0.18*** (0.02)
TED spread _{t-120}	0.02 (0.02)	0.02 (0.03)	0.01 (0.02)
RSG Nov 2008	0.54*** (0.06)	0.72*** (0.08)	0.41*** (0.06)
Physical demand	0.76*** (0.19)	0.44 (0.33)	0.97*** (0.16)
Constant	1.11 (2.71)	4.65 (4.73)	-2.35 (2.30)
<i>N</i>	267	267	267
R ²	0.89	0.82	0.85

Notes:

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Table 6: **Implicit transaction costs** These costs represent the minimum marginal transaction cost that would make arbitrage unprofitable, and were calculated using the method proposed by Jha and Wolak [2013]. The three tests correspond to different tests for the null of no arbitrage opportunities and are computed separately using data from the first and second periods. The first panel shows results using every node in the MISO market, and the second one including only hubs

Test	Mean	Max	Min	SD
Test 1 - First period	3.50	23.66	1.81	1.35
Test 1 - Second period	3.76	23.07	1.55	2.50
Test 2 - First period	1.40	4.84	0.73	0.45
Test 2 - Second period	1.31	8.52	0.54	0.74
Test 3 - First period	1.59	3.45	0.91	0.39
Test 3 - Second period	1.36	4.68	0.59	0.53

Test	Illinois	Michigan	Minnesota	Indiana
Test 1 - First period	3.84	3.94	2.30	3.19
Test 1 - Second period	3.97	3.01	2.98	2.76
Test 2 - First period	1.63	1.43	0.98	1.29
Test 2 - Second period	1.61	1.12	0.77	1.02
Test 3 - First period	1.68	1.84	1.35	1.35
Test 3 - Second period	1.43	1.35	0.77	1.14

Table 7: **Summary statistics for virtual participants in the first period** Each row was computed by participant, using all the cleared bids. Then the summary statistics among participants were computed. For instance, mean MW is the average size of a bid in MWs, and it is equal 54.2 for demand bids cleared by the average participant.

Statistic	N	Mean	St. Dev.	Min	Max
Total MW	142	257,322	689,601	2.0	4,827,232.
Mean MW	142	54.2	114.9	0.1	856.2
SD MW	141	38.1	82.8	0.0	492.8
Total profits	142	-75,850	742,437	-5,813,273	2,380,380
Mean profits	142	-36.3	499.8	-1,635	5,029
SD profits	141	1,378	3,160	5.0	30,997
Active hours	142	2,108	2,884	1	10,735
Active days	142	155.7	154.6	1	455
# nodes	142	168.0	311.0	1	1,229
Supply	N	Mean	St. Dev.	Min	Max
Total MW	138	167,436.9	352,043.2	1.6	2,194,033.0
Mean MW	138	32.6	42.8	0.2	271.8
SD MW	137	21.1	29.4	0.0	199.7
Total profits	138	91,834.5	378,104.8	-937,480.6	1,921,904.0
Mean profits	138	5.4	257.0	-1,036.2	2,280.5
SD profits	137	823.2	974.1	2.4	7,170.9
Active hours	138	2,062.8	2,721.2	1	10,870
Active days	138	147.4	150.5	1	454
# nodes	138	168.0	322.4	1	1,349

Table 8: **Summary statistics for virtual participants in the second period** Each row was computed by participant, using all the cleared bids. Then the summary statistics among participants were computed. For instance, mean MW is the average size of a bid in MWs, and it is equal 28.1 for demand bids cleared by the average participant.

Statistic	N	Mean	St. Dev.	Min	Max
Total MW	154	408,542	1,083,090	1.0	8,441,286
Mean MW	154	43.9	80.8	0.1	486.9
SD MW	152	28.2	48.6	0.0	321.6
Total profits	154	-18,342	1,534,969	-13,638,828	3,938,031
Mean profits	154	-46.7	653.3	-6,336	3,434
SD profits	152	1,151	1,838	1.4	11,905
Active hours	154	3,684	4,408	1	15,295
Active days	154	230.0	217.8	1	641
# nodes	154	192.3	302.2	1	1,303
Supply	N	Mean	St. Dev.	Min	Max
Total MW	154	326,712	728,613	1.0	4,813,957
Mean MW	154	29.1	35.2	0.1	161.5
SD MW	153	18.9	25.8	0.0	142.9
Total profits	154	458,214	1,036,782	-595,834	6,045,100
Mean profits	154	46.8	225.9	-885.9	1,077
SD profits	153	963.5	1,604.8	4.8	10,933x
Active hours	154	3,574	4,057	1	15,293
Active days	154	231.4	211.7	1	641
# nodes	154	199.0	333.9	1	1,379

Table 9: **Net day-ahead premium over the years:** Results from a regression of the hourly day-ahead premium, net of RSG charges, on year dummies. In the first three columns, the hourly premium is computed as the simple average across nodes, the median across nodes, and a weighted average using the total quantity traded by physical participants as weights. In the last three columns, the sample is restricted to nodes with virtual activity and the day-ahead premium is computed as a weighted average, and the three columns present three different weights: virtual demand in Mwhts, virtual supply in Mwhts, and total virtual volume in Mwhts. HAC standard errors reported.

	Average	Median	W-average	V demand	V supply	All virtual
	(1)	(2)	(3)	(4)	(5)	(6)
Period 1	-0.91*** (0.27)	-0.60** (0.26)	-0.75*** (0.26)	-1.44*** (0.29)	0.27 (0.30)	-0.76*** (0.28)
Period 2	0.53** (0.25)	1.07*** (0.24)	0.64** (0.26)	-0.01 (0.28)	1.45*** (0.28)	0.63** (0.27)
<i>N</i>	26,304	26,304	26,304	26,304	26,304	26,304
<i>R</i> ²	0.002	0.003	0.002	0.002	0.004	0.001

Notes: ***Significant at the 1 percent level, ** at 5 percent, * at 10 percent.

Table 10: **Differences between virtual demand and supply** Results from a regression between the difference between the quantities cleared by virtual demand and virtual supply (in logs of MWhs) on the interaction of two dummy variables: One for the period after April 2011, and another for bids cleared by participants that also submit physical bids to purchase electricity. Quantities were aggregated at the daily level, the ratio of demand to supply is the difference of their logs. HAC standard errors reported.

	Ratio DS
Period 1, no physical buyers	0.06** (0.02)
Period 2, no physical buyers	-0.01 (0.03)
Period 1, only physical buyers	0.98*** (0.12)
Period 2, only physical buyers	-0.32** (0.14)
N	2,192
R^2	0.40

Table 11: **Bid premium and virtual profits** The bid premium is the absolute value of the difference between the maximum willingness to pay or minimum willingness to receive and the cleared price. The table presents the results of regressing the profits earned by a virtual bid and its bid premium, excluding bids submitted by physical buyers from the sample. Standard errors are clustered at the participant and node level. Bids that do not specify a price are excluded as well.

<i>Dependent variable:</i>				
Profits				
	Virtual demand		Virtual supply	
Bid premium	-0.15** (0.07)	-0.17* (0.09)	-0.02* (0.01)	0.06*** (0.01)
Constant	-0.68 (1.92)	6.65*** (1.55)	8.99*** (1.15)	12.91*** (0.77)
Observations	1,907,910	4,053,446	2,215,876	4,151,093
R ²	0.0000	0.0001	0.0000	0.0000

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

Table 12: **Total FTR and virtual profits during the first period.** The tables show summary statistics for the sum of virtual and FTR profits at sink and source nodes. Each row is an FTR traded in the FTR quarterly or monthly auction that MISO runs. This means that if two different participants hold FTRs between the same two nodes for the same period, possibly for different MWs, they are different observations. Virtual profits are taken as either the minimum or the mean profits made by a participant at that node during the period in which the FTR was valid. This is the most conservative approach given that we cannot match virtual participants and FTR holders. The last row describes FTR profits without adding the virtual profits. The sample excludes FTR moves and virtual bids submitted by physical buyers.

First period	N	Mean	St. Dev.	Min	Max
V demand at sink - mean	244,221	241	31,949	-1,889,587	3,653,247
V demand at the source - mean	244,221	502	31,422	-1,875,329	3,674,827
V supply at sink - mean	244,221	1,142	27,029	-1,872,853	3,652,959
V supply at the source - mean	244,221	1,175	27,273	-1,876,974	3,652,929
V demand at sink - min	244,221	-3,909	57,528	-1,960,655	3,653,247
V demand at the source - min	244,221	-2,393	49,722	-2,524,247	3,653,236
V supply at sink - min	244,221	-363	27,327	-1,880,317	3,652,959
V supply at the source - min	244,221	-1,498	28,131	-1,878,508	3,652,697
FTR profits	244,221	1,059	26,795	-1,875,420	3,653,247
Second period	N	Mean	St. Dev.	Min	Max
V demand at sink - mean	449,349	222	35,267	-2,204,800	2,910,730
V demand at the source - mean	449,349	819	30,996	-2,113,448	2,911,725
V supply at sink - mean	449,349	1,457	27,974	-1,257,254	2,912,936
V supply at the source - mean	449,349	1,758	28,269	-1,232,456	2,909,822
V demand at sink - min	449,349	-2,411	45,133	-2,236,815	2,909,764
V demand at the source - min	449,349	-1,445	39,575	-2,468,148	2,907,973
V supply at sink - min	449,349	-320	28,345	-1,283,461	2,909,605
V supply at the source - min	449,349	-1,024	28,743	-1,275,252	2,902,643
FTR profits	449,349	1,139	27,513	-1,255,307	2,910,177

Table 13: **FTR and virtual volumes** Regression of the ratio of cleared virtual demand to virtual supply at the sink or source node on the log of FTR MWhs having that node as a source. Does not include virtual bids that could potentially be FTR moves nor bids placed by physical buyers.

	Ratio of cleared virtual demand to virtual supply MWhs			
	Sink node		Source node	
	First period	Second period	First period	Second period
FTR sink MWhs	7.81*** (2.74)	2.10 (1.67)		
FTR source MWhs			-2.86 (2.33)	-0.82 (1.18)
Constant	9.63** (4.61)	29.64*** (8.29)	45.87*** (12.90)	32.85*** (4.82)
Observations	33,382	53,394	32,171	52,589
R ²	0.0005	0.0000	0.0000	0.0000

Note:

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

Table 14: **FTR and virtual demand profits** Regression of total FTR profits on profits accumulated by cleared virtual demand bids at the sink and source nodes, respectively. A dummy for peak hours is included. SE clustered by FTR sink node and FTR source node. Excludes virtual bids that could potentially be FTR moves and bids submitted by physical buyers.

<i>Dependent variable:</i>				
	FTR profits			
	First period	Second period	First period	Second period
Sink profits	−0.02*** (0.01)	−0.002 (0.03)		
Source profits			0.03 (0.03)	0.04 (0.04)
Peak	4,909*** (636)	5,900*** (818)	5,091*** (645)	5,892*** (815)
Constant	−734*** (248)	−1,067*** (288)	−709*** (245)	−1,072*** (288)
Observations	142,073	277,098	142,073	277,098
R ²	0.003	0.002	0.003	0.003

Note:

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

Table 15: **Correlation between virtual profits and FTRs for Louis Dreyfus.** The table presents the correlation between FTR profits and virtual bids or offers at the sink and source node. A dummy indicating whether the FTR was valid during peak hours was added as a control as in regressions for the whole market, though results are basically the same when it is excluded. The sample includes only FTRs valid between November 2009 and February 2010, for which either the sink or the source node is the Velva node, where the alleged manipulation occurred.

	FTR profits			
Demand at sink	-4.97***			
	(1.64)			
Demand at source		0.40		
		(0.49)		
Supply at sink			0.83	
			(1.92)	
Supply at source				-5.75
				(5.96)
Peak	15,871.80	146,261*	138,037*	173,084*
	(29,663)	(86,779)	(76,726)	(98,402)
Constant	-23,563	17,850*	21,819	3,138
	(16,244)	(10,316)	(14,098)	(17,146)
Observations	67	67	67	67
R ²	0.60	0.04	0.04	0.07

Note:

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

Table 16: **Congestion and virtual activity** The first two columns present results from a regression of virtual profits on the congestion component of the day-ahead price, ran separately for virtual demand and virtual supply. The third and fourth columns presents results from the ratio of virtual demand to virtual supply on day-ahead congestion, and the day-ahead congestion premium, respectively. Bids submitted by physical buyers were excluded from the sample. Standard errors are clustered at the node level.

	First Period			
	VD prof	VS prof	MWvd/MWvs	
DA cong	-2.31*** (0.47)	1.67*** (0.23)	0.13*** (0.02)	
C premium			-0.003 (0.002)	
Constant	-9.56** (4.09)	1.99** (0.92)	4.75*** (0.56)	4.84*** (0.57)
Observations	1,208,556	1,414,051	1,576,762	1,576,762
R ²	0.0005	0.001	0.0004	0.0000
	Second Period			
	VD prof	VS prof	MWvd/MWvs	
DA cong	-1.46*** (0.26)	0.46* (0.24)	0.05*** (0.01)	
C premium			-0.01*** (0.002)	
Constant	-6.35 (4.01)	6.36*** (1.18)	3.18*** (0.26)	3.21*** (0.27)
Observations	2,777,257	2,856,464	2,856,464	2,856,464
R ²	0.0002	0.0000	0.0002	0.0000

Note:

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

Figures

Figure 1: **Virtual volume, the TED spread and the premium.** The top panel shows the 120 days lag of the daily TED spread over time. The TED spread is the difference between the 3 month LIBOR - London Interbank Offered Rate- and the 3 month U.S. Treasury Bill rate, and it is a measure of perceived risk. When it is higher, capital is more restricted and so is arbitrage. The middle panel shows the trajectory of the daily cleared virtual volume in the MISO market. The bottom panel shows the 45 days moving average of the day-ahead premium over time. The red lines indicate regulation changes. The first one, on November 2008, marks the FERC order according to which virtual offers have to pay deviation (RSG) charges. Virtual volume decreased as a consequence. The second red line, on April 2011, indicates a change in the computation of deviation charges that made them lower and thus encourage an increase in financial participation.

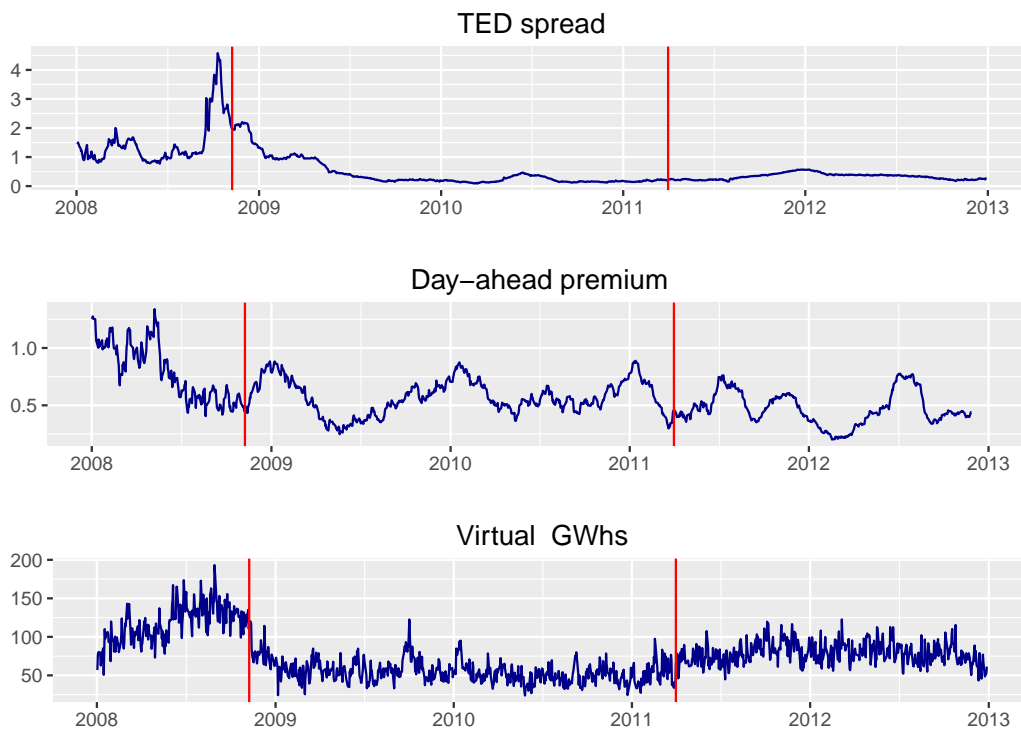


Figure 2: **Volume cleared by virtual buyers and sellers over time** The green dashed lines indicate the dates in which the rules regarding RSG charges and virtual bidders were changed. These charges were first imposed on virtual sales on November 2009, and were then significantly decreased on April 2011.




Figure 3: **Cumulative premium after discounting RSG charges.** The plot shows the cumulative day-ahead premium minus the RSG charges applied on virtual offers. It is equivalent to the cumulative profits that a participant would make by selling 1 MW at every node in every hour. The vertical line indicates April 1st, 2010, the first day in which the new RSG charges were applied.






Figure 4: **Cumulative profits from arbitrage in the first period.** The first panel shows profits from a simple strategy that buys if the premium at that node in that hour was negative, and sells if it was positive. The bottom panel shows profits from a strategy that buys when the predicted premium is larger than -1, and sells when it is larger than 3. The premium is predicted using a regression of the premium on lags with data for the past month. The three lines in each graph represent three different modifications of the initial strategy strategies. The red line is the one described above. The blue and green lines represent profits from the same strategy, but leaving the market for 2 and 5 days, respectively, if losses one day exceed \$10,000.



(a) Simple arbitrage: prediction based on the sign of the past premium.

(b) Arbitrage based on a prediction made with a regression of the premium on lags.

Figure 5: **Cumulative profits from arbitrage in the second period.** The first panel shows profits from a simple strategy that buys if the premium at that node in that hour was negative, and sells if it was positive. The bottom panel shows profits from a strategy that buys when the predicted premium is larger than -1, and sells when it is larger than 3. The premium is predicted using a regression of the premium on lags with data for the past month. The three lines in each graph represent three different modifications of the initial strategy strategies. The red line is the one described above. The blue and green lines represent profits from the same strategy, but leaving the market for 2 and 5 days, respectively, if losses one day exceed \$10,000.



(a) **Simple arbitrage: prediction based on the sign of the past premium.** The blue line ends up with \$4,150,737 in profits and has a Sharpe ratio of 1.04. The green line accumulates profits for \$3,069,150 and has a Sharpe ratio of 0.77; the red line gets \$598,181 and a Sharpe ratio of 0.10.

(b) Arbitrage based on a prediction made with a regression of the premium on lags.

Figure 6: **Fraction of virtual demand to virtual volume** The plots show the change of the fraction of cleared virtual demand to total virtual volume after RSG charges were lowered on April 2011. The sample goes from 2010 to 2012. The first plot shows the density of the ratio of virtual demand to virtual supply for each participant each day. The second plot shows the cumulative distribution of the mean ratio of virtual demand to total virtual volume by participant, i.e. each observation is a participant. The sample excludes bids placed by physical buyers.

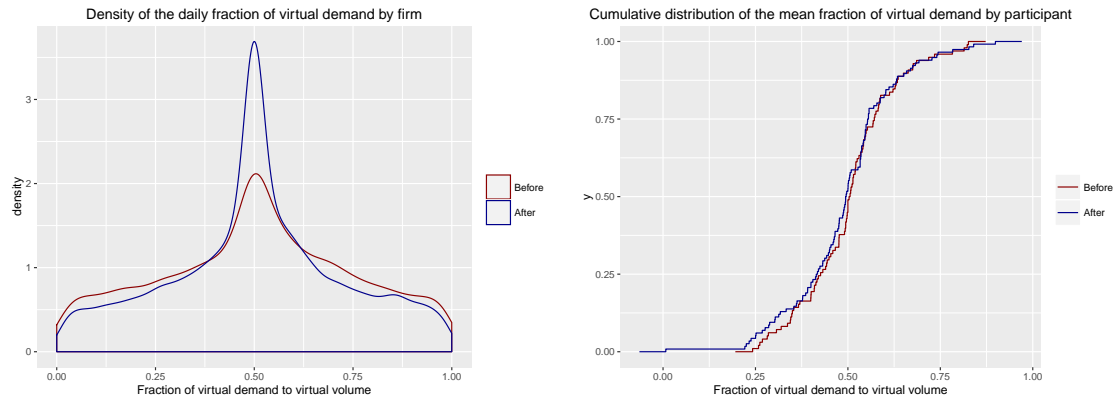


Figure 7: **Cumulative profits Louis Dreyfus Energy Services**




Figure 8: **Virtual profits and active days by period.** Each dot represents a participant, and it indicates the number of days in which she was active, and the total profits she made. Profits consider RSG charges. All participants are included, not only those without physical load capacity.




Figure 9: Cumulative profits of the largest losers. RSG charges were already subtracted.

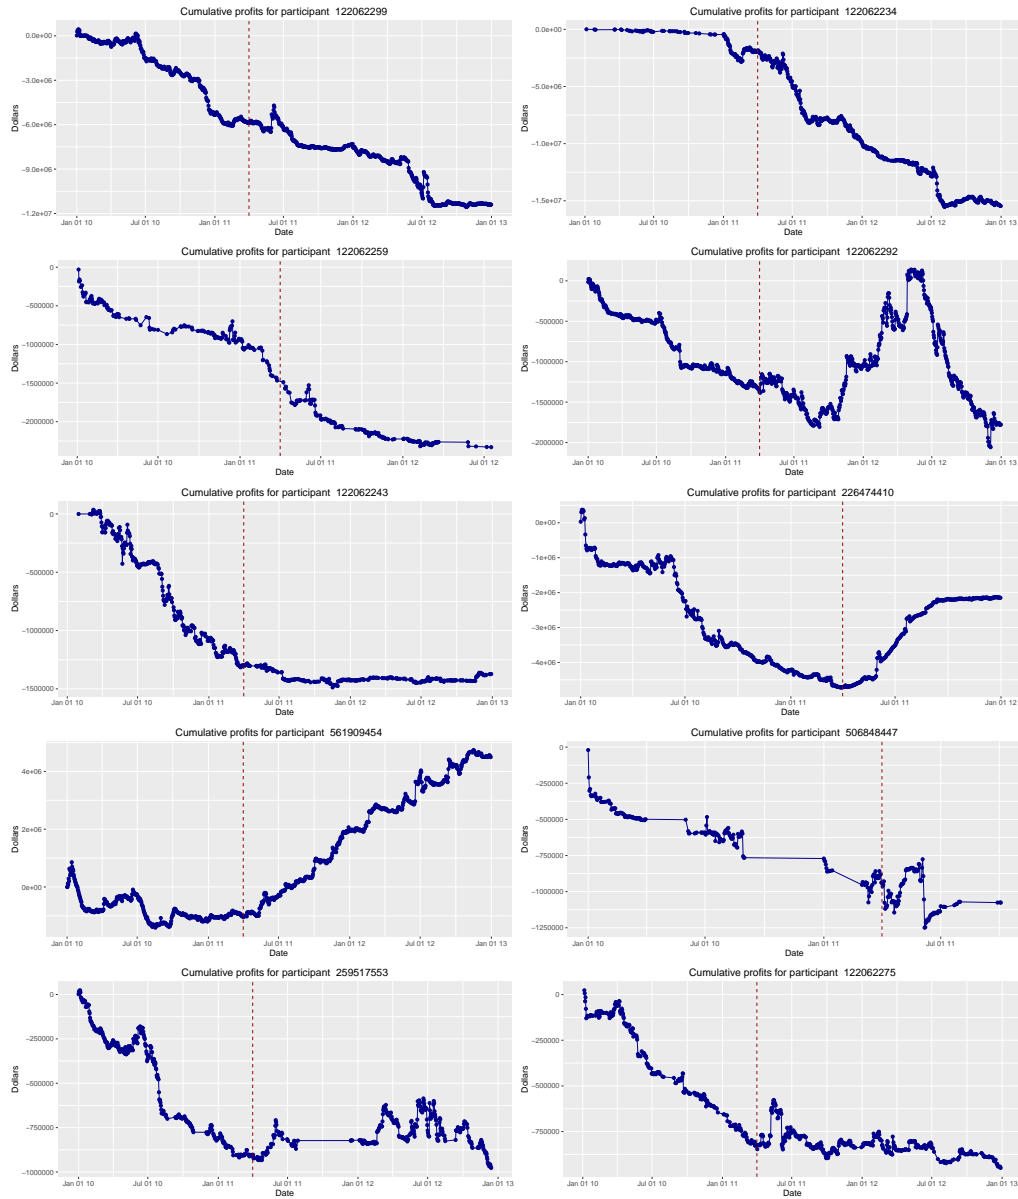


Figure 10: **Largest winners' profits.** The graph shows the cumulative profits of the virtual participants who earned the most during the period 2010-2012. As

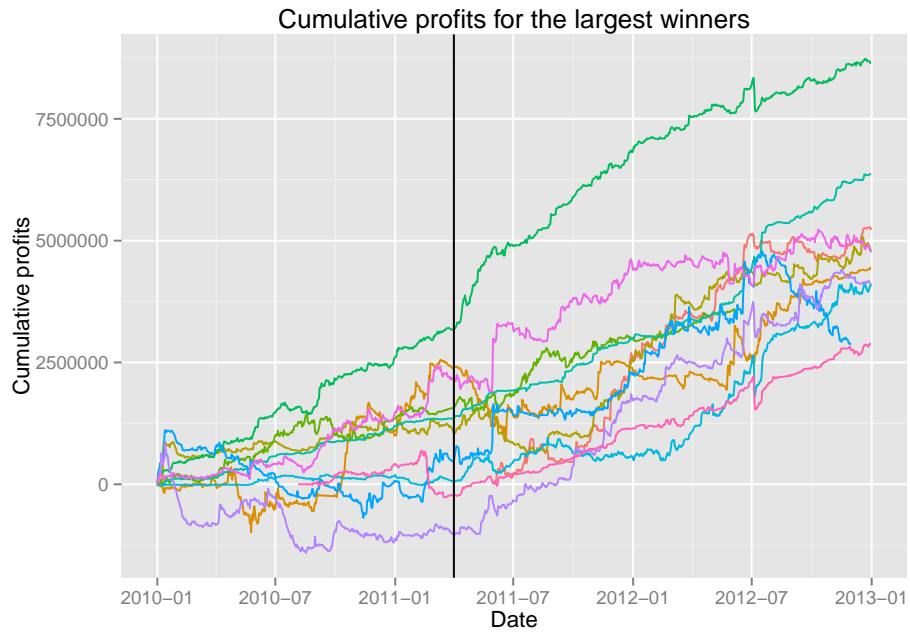


Figure 11: **Bid Premium and virtual profits without potential FTR moves** For each bid, the graph shows the profits it made against the bid premium, which is the difference between the maximum submitted willingness to pay and the actual cleared price (for supply bids, it is the difference between the clearing price and the minimum price the bidder is willing to accept). The graph indicates that in the first period there is a group of bidders submitting bids very far from the cleared price. Price-taking bids in which only a quantity was submitted were excluded from the sample. Virtual bids suspected of being used to move an FTR from the day-ahead to the real-time were excluded, i.e. all pairs of a bid and an offer cleared for the same MWs by the same participant.





Figure 12: **Cumulative day-ahead premium after RSG charges.** The plot shows the cumulative day-ahead premium after RSG charges, starting on 2005. These are the profits that a virtual bidder selling 1MW at each node would have made. As it can be observed, there were profitable opportunities until November 2008, when the RSG calculation changed and they were imposed on virtual offers as well.



Appendices

A The Day-Ahead Premium and Market Power

A positive day-ahead premium can be explained by strategic behavior of participants with market power. Even though the market may not be very concentrated—in MISO the largest generation owner produces 13% of the energy output on average—congestion creates local market power because it isolates a node or a group of nodes from the rest of the market [Borenstein et al., 1997]. When a transmission line is at capacity, it may not be feasible to deliver energy from the lowest priced generator to a buyer, allowing better located generators to affect the price by strategically choosing the quantities they offer in the day-ahead and real-time markets [Saravia, 2003]. In this section we present the argument that intertemporal price discrimination by generators with market power can drive a day-ahead premium but do not consider the incentives to create or avoid congestion [Borenstein et al., 1997, Joskow and Tirole, 2000].

Consider a generator facing a residual demand who decides how much to sell in the day-ahead and how much in the real-time market. We model the generator's decision as if she could choose both the real-time and the day-ahead sales. Most firms are not as flexible, but their problem is a particular case of the one presented here.

The generator solves the following problem

$$\max_{\{Q_{DA}, Q_{RT}\}} \Pi = P(Q_{DA})Q_{DA} + P(Q_{DA} + Q_{RT})Q_{RT} - C(Q_{DA} + Q_{RT}) \quad (1)$$

where Q_{DA} and Q_{RT} are the quantities sold in the day-ahead and real-time markets, respectively, $C()$ is the total cost function, and $P()$ is the inverse residual demand function. This set-up assumes that demand firms do not have market power, so generators expect the same demand in the day-ahead and the real-time market. Assuming that demand is not strategic is reasonable in the MISO market because most bids are price insensitive and day-ahead withholding is small. In this case demand can be satisfied in either market, so the day-ahead quantity causes a parallel shift in the effective real-time demand. After rearranging the two first order conditions, we get:

$$P(Q_{DA}) - P(Q_{DA} + Q_{RT}) = -\frac{\partial P(Q_{DA})}{\partial Q_{DA}} Q_{DA} \quad (2)$$

This equation is very similar to the standard pricing equation of a firm with market power; here the real-time price replaces the marginal cost and the margin is the day-ahead premium. Equation 2 along with downward sloping demand, implies that in equilibrium firms will withhold production in the day-ahead market. If a generator has market power, it will choose its day-ahead sales such that the day-ahead price is larger than the real-time price, i.e. $Q_{RT} > 0$. If a generator does not have market power she faces a constant price in both markets and will be indifferent between the two markets. However, the existence of charges imposed on deviations from the day-ahead

schedule, a common industry practice, will make it more attractive to sell the whole production in the day-ahead market and choose $Q_{RT} = 0$. Note that for the fixed production level version of this problem, where the sum of Q_{DA} and Q_{RT} are restricted to the total production of the plant, similar results continue to hold.

Consistent with this explanation for the day-ahead premium, Mercadal [2016] examines the MISO market and finds that generators consistently sell less in the day-ahead than they end up producing and that the day-ahead premium is lower when firms withhold less. In the Iberian market, Ito and Reguant [2016] present evidence consistent with the theoretical predictions of a model in which the forward premium results from firms' market power.

We would not expect a day-ahead premium to persist in a market where virtual transactions are allowed. Financial participant would exploit existing profitable arbitrage opportunities and therefore close the day-ahead premium. In the MISO market, virtual trading has been allowed since the market started in 2005. In spite of this, Table 2 shows that the day-ahead price has been systematically higher than the real-time price at least until 2010. The reasons behind this persistence will be explored in the following sections.

B RSG charges and virtual bidders

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, a fixed 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 actually produces in the real-time market. Currently, as virtual participants do not physically generate or consume 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, 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 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 MW^S 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 $P^{DA} - P^{RT} - RSG_RATE$. Charges during this period were on average larger than the day-ahead premium (see Tables 1 and ??).

On March 2011 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 CCF is positive, the transmission 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:

$$RT_RSG_DIST1_h = CMC_DIST_h + DDC_DIST_h$$

$$CMC_DIST_h = \sum_n \max \{ (MW_n^S - MW_n^D) \cdot CCF_{h,n}, 0 \} \cdot CMC_RATE_{h,n}$$

$$DDC_DIST_h = \sum_n \max \{ (MW_n^S - MW_n^D), 0 \} \cdot DDC_RATE_{h,n}$$

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 .

The change in the computation of RSG charges divides our sample in two periods. The first period ranges from January 2010 to March 2011, while the second goes from April 2011 to December 2012. This distinction will be maintained throughout the paper.

Some summary statistics of RSG charges are presented in Table 1, where RSG 1 is the node invariant charge used before April 2011. The charges decreased significantly after their modification. Moreover, they stopped offsetting the day-ahead premium as they did before that (see Table ??). Even though during the first period the day-ahead price was on average larger than the real-time price, virtual bidders could not profit from this gap by virtually selling in the day-ahead because RSG charges were larger than the profits. As we show in section 4.1, virtual traders indeed help to reduce the difference between the day-ahead and real-time prices. Nonetheless, the premium persisted in the MISO market because these deviation charges prevented virtual entry, as we will

argue in more detail in the following sections.

C Arbitrage strategies

Virtual participants submit their bids in the day-ahead market, which closes at 11am and clears bids and offers for the 24 hours of the next day, starting at midnight. Therefore, although participants can learn the clearing day-ahead and real-time prices immediately after the real-time closes, they cannot use the 24 hours lag to predict the premium on the next day. This restriction in the available information is the reason why the simplest strategy we considered depends on the sign of the premium 48 instead of 24 hours ago. It buys when it is negative, and sells when it is positive. The second arbitrage strategy is more complex; so it is described in more detail below.

Consider a potential bidder who runs rolling regressions using data on the past 30 days to predict the day-ahead premium at every hour of the next day. To predict the premium on day 31 (tomorrow, starting today at midnight), the bidder uses data on the 24 hours of days 1 to 29 and the first 8 hours of day 30 (today). Using the coefficients of the regression and the last 24 hours of data (the first 8 hours of 30 and the last 16 hours of 29), the bidder predicts the premium for the 24 hours of 31, starting on day 30 at midnight. This means that the first 8 hours of day 31 are predicted using the first 8 hours of day 30, while hours 9 to 24 are predicted using the same hours of day 29. The prediction is done using a regression of the premium on lags:

$$\begin{aligned}
 P_t = & \alpha_0 + \alpha_1 P_{t-24} + \alpha_2 P_{t-25} + \alpha_3 P_{t-26} + \alpha_4 P_{t-27} + \alpha_5 P_{t-28} + \alpha_6 P_{t-29} \\
 & + \alpha_7 P_{t-30} + \alpha_8 P_{t-31} + \alpha_9 P_{t-32} + \alpha_{10} P_{t-33} + \alpha_{11} P_{t-34} + \alpha_{12} P_{t-35} \\
 & + \alpha_{13} P_{t-36} + \alpha_{14} P + \alpha_{15} P_{t-48} + \alpha_{16} P_{t-71} + \alpha_{17} P_{t-72} + \varepsilon_t
 \end{aligned}$$

where P_t is the premium at hour t and P_{t-24} is the premium at the same hour the day before. These regressions are run separately for each of the nodes in the MISO footprint. When the prediction is above 3 for a certain node, she sells in the virtual market at that node; when it is below -1 , she buys.

The day-ahead market closes at 11 AM; so we assume that the bidder has three hours to find the optimal portfolio. This should be enough time because she only has to run one regression for each node. MISO publishes prices data in real-time; so all the required data are available at the moment.

The cumulative profits from this strategy are shown in the bottom panel of Figures 4 and 5 for the first and second period, respectively. The three lines represent three different strategies. The red line is the one described above. The blue and green lines represent profits from the same strategy, but leaving the market for 2 and 5 days, respectively, if losses one day exceed \$10,000. These more conservative versions do better, specially in the first period.

A Online Appendix

A The day-ahead premium as a risk premium

The day-ahead premium has also been explained as a risk premium (Bessembinder and Lemmon [2002], Longstaff and Wang [2004], among others). In fact, electricity prices are very volatile, and the real-time price is much more volatile than the day-ahead one (see Table 1). Moreover, because electricity is not storable and last minute adjustments to production are expensive, the spot price presents large spikes when there is an unexpected shock to demand or generation. Table ?? shows the variation across hours in the first three moments of the day-ahead and real-time prices in the MISO market. Prices vary across the day, and the distribution of the spot price is generally more volatile and positively skewed. Additionally, the day-ahead price is more predictable than the real-time price. Table B1 shows the correlation between the price at certain hour and the price at the same hour the day before. For the day-ahead price, correlation is higher and the lags explain more than 80% of the variation, while for the real-time price the regression explains less than 30% of the variation.

Bessembinder and Lemmon [2002] present a general equilibrium model in which the premium results from demand uncertainty and convex production costs with risk-averse firms. The premium can be positive or negative depending on whether the risk is higher for generators or retailers. Longstaff and Wang [2004] test for time variation in the forward premium, finding that it is related to risk factors like the volatility of unexpected changes in demand, the spot price, and total revenues. They interpret these results as meaning that the premium is the result of the behavior of rational risk averse agents, still in the perfectly competitive setting.

Bessembinder and Lemmon [2002] present four testable hypothesis that derive from their model. The first two are that the premium decreases with the variance of the spot price and increases with its skewness. They did not have enough data to test them because the market had just started, but Longstaff and Wang [2004] do it for the PJM market and their findings are consistent with the predictions of the model. They regress the mean premium on the variance and skewness of the spot price, where these variables are computed by pooling 30 months of hourly observations and separating each of the 24 hours of the day. We repeat the exercise, but computing the moments from all observations of a certain hour and node during a month. As shown in Table ??, we find negative correlation between the premium and both the variance and skewness of the spot price. The remaining testable hypothesis are that the premium is convex in the variability of power demand, initially decreasing and then increasing, and that the premium increases in expected demand. We find the latter to be true, but the premium initially increases and then decreases in the variability of demand (Table ??). There are no virtual bidders in the model presented by Bessembinder and Lemmon [2002]. As the spot price risk is unlikely to be correlated with the overall risk in the economy, the rejection of the predictions from the model can be due to the fact that with financial participation, even if limited, the risk premium is not a first order effect.

B Tables and Figures

Table B1: **Serial correlation of electricity prices:** Results from a regression of the day-ahead or real-time price at a node and a certain hour on three lags, including node fixed effects. The lags are 24 hours apart, so the 24 hours lag is the price at the same node during the same hour on the previous day. The first period ranges between January 2010 and March 2011, the second one between April 2011 and December 2012.

	Day-ahead price		Real-time price	
	First period	Second period	First period	Second period
24 hours lag	1.282*** (0.0002)	1.235*** (0.0002)	0.191*** (0.0001)	0.423*** (0.0002)
48 hours lag	-0.377*** (0.0003)	-0.253*** (0.0003)	0.090*** (0.0001)	0.098*** (0.0002)
72 hours lag	-0.020*** (0.0002)	-0.075*** (0.0002)	0.056*** (0.0001)	0.117*** (0.0002)
N	21,159,585	27,334,708	21,159,585	27,334,708
R^2	0.870	0.891	0.261	0.294

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

Table B2: **FTR and virtual supply profits** Regression of total FTR profits on profits accumulated by virtual supply bids at the sink and source nodes, respectively. A dummy for peak hours is included. SE clustered by FTR sink node and FTR source node. Excludes virtual bids that could potentially be FTR moves and bids placed by physical buyers.

<i>Dependent variable:</i>				
FTR profits				
	First period	Second period	First period	Second period
Sink profits	-0.01 (0.02)	0.07 (0.06)		
Source profits			-0.01 (0.02)	0.05 (0.04)
Peak	4,990.57*** (662.34)	5,879.04*** (819.90)	4,987.42*** (662.10)	5,880.41*** (812.71)
Constant	-707.22*** (241.73)	-1,168.34*** (297.28)	-706.84*** (247.91)	-1,189.41*** (295.89)
Observations	142,073	277,098	142,073	277,098
R ²	0.002	0.003	0.002	0.003

Note:

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



MIT Center for Energy and Environmental Policy Research

Since 1977, the Center for Energy and Environmental Policy Research (CEEPR) has been a focal point for research on energy and environmental policy at MIT. CEEPR promotes rigorous, objective research for improved decision making in government and the private sector, and secures the relevance of its work through close cooperation with industry partners from around the globe. Drawing on the unparalleled resources available at MIT, affiliated faculty and research staff as well as international research associates contribute to the empirical study of a wide range of policy issues related to energy supply, energy demand, and the environment.

An important dissemination channel for these research efforts is the MIT CEEPR Working Paper series. CEEPR releases Working Papers written by researchers from MIT and other academic institutions in order to enable timely consideration and reaction to energy and environmental policy research, but does not conduct a selection process or peer review prior to posting. CEEPR's posting of a Working Paper, therefore, does not constitute an endorsement of the accuracy or merit of the Working Paper. If you have questions about a particular Working Paper, please contact the authors or their home institutions.

**MIT Center for Energy and
Environmental Policy Research**
77 Massachusetts Avenue, E19-411
Cambridge, MA 02139
USA

Website: ceepr.mit.edu

MIT CEEPR Working Paper Series is published by
the MIT Center for Energy and Environmental
Policy Research from submissions by affiliated
researchers.

Copyright © 2017
Massachusetts Institute of Technology

For inquiries and/or for permission to reproduce
material in this working paper, please contact:

Email ceepr@mit.edu
Phone (617) 253-3551
Fax (617) 253-9845