The Role of Financial Players in Electricity Markets: An Empirical Analysis of MISO

John Birge, Ali Hortaçsu, Ignacia Mercadal and Michael Pavlin

October 11, 2013

Abstract

As in most commodities markets, deregulated electricity markets allow the participation of purely financial (virtual) 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 (day-ahead) and spot prices, which may arise in the presence of market power by physical suppliers or buyers. However, we find that the impact of financial players on reducing pricing gaps has been limited. A forward premium persists. Surprisingly, some large 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. 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.

1 Introduction

The Midwest Independent Transmission System Operator (MISO) is a regional transmission organization (RTO) that coordinates the wholesale electricity market in 12 states in the U.S. and the Canadian province of Manitoba. RTOs allow electricity markets to be decentralized by providing an independent platform where generators and load entities interact. Their stated goal is to organize the market to achieve both low prices and high reliability. The system has many different components, but in this paper, the focus will be on two parts: the energy market and the financial transmission rights (FTR) market.

In the energy market, megawatts (MWs) of electricity are traded at particular locations throughout the market footprint (nodes) for each hour. Participants submit demand or supply bids for an hour of energy at specific nodes in the network and the market is cleared at the level of the node. Prices at nodes are referred to as locational marginal prices (LMPs) and have a common marginal cost component but vary between nodes depending on congestion and transmission losses. Congestion occurs when demand at a certain node exceeds the capacity of the transmission lines connecting the node to lower cost generators, making it necessary to increase the price until the market clears. Transmission losses occur because part of the electricity is lost as heat as it travels over transmission lines from one node to another.

The energy market has two stages: the day-ahead or forward market, which schedules production or demand for each participant over the next day, and the real-time market, where final adjustments are made 30 minutes before the operating hour. In an informationally efficient, competitive market, one may expect the day-ahead and real-time prices to be the same on average since differences can only result from unpredictable shocks to supply or demand. However, MISO is far from a frictionless market due to the congested nature of the transmission network: congestion effectively creates local oligopolists, sometimes monopolists at the nodes that are affected [Borenstein et al., 2008]. Figure 1 illustrates what local market power can lead to: here we analyze the simple case of a monopolist with fixed capacity, facing non-strategic buyers. The generator has Q_T MWs of capacity (fixed capacity is reasonable for big producers with coal or oil plants, which are the most likely to have market power). If she sells all of this in the DA market, she gets P_T for each MW. If she only sells Q_{DA}, those are sold in the day-ahead at P_{DA} and the remaining $(Q_T - Q_{DA})$ are sold in the real-time market at P_T. The generator has no incentive to sell all the units in the day-ahead because then she would get P_{DA} as well. How many to sell in each market will depend on the elasticity of demand, as for a monopolist. In the above example, the local monopolist is effectively doing intertemporal price discrimination between the day-ahead and real-time markets, leading to a positive day-ahead price premium.[1]

Indeed, because the day-ahead price was observed to be consistently higher, MISO opted to include financial (called "virtual" by MISO) bidders in the energy market. (See Jha and Wolak [2013] for an excellent discussion of the motivations for incorporating virtual bidders into the market design.) Virtual bidders are participants

[1]This example presents the case of a generator who has market power in a net importer node, i.e. one in which local demand exceeds local supply at the clearing price. Congestion creates market power because the generator becomes the only alternative to buy in that node. The case for an exporter will be made in section 4.

without physical load or generation capacity and are permitted to bid at any MISO node. Virtual bidders are expected to act as arbitrageurs because from their point of view, virtual MWs are simply a financial instrument that pays the difference between the day-ahead and the real-time price to a buyer and the negative of that to a seller. Virtual bidders should in principle close any predictable differences between the day-ahead and real-time prices by bidding against it.

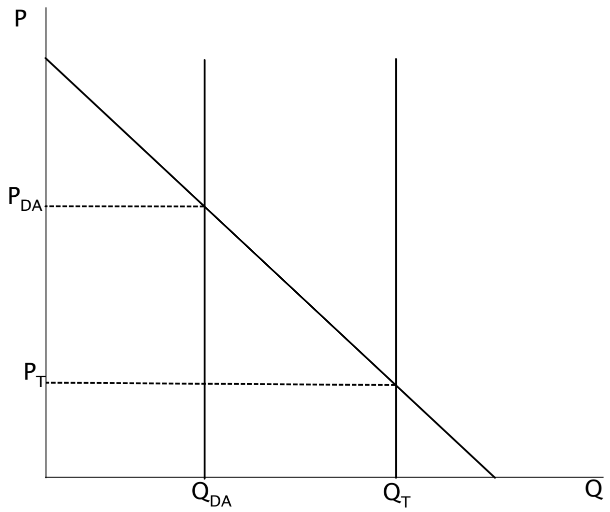

Figure 1: Intertemporal price discrimination

Despite the participation of virtual bidders, between 2010 and 2012, the day-ahead price on MISO was higher than the real-time price. Moreover, virtual bidders *bought* more energy in the day-ahead market than they sold. Moreover, some participants accumulated negative profits steadily for longer than a year without leaving the market. 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. These virtual participants were also more likely to submit bids far from the clearing price, consequently not contributing to closing the gap.

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. The FTR market exists to distribute the congestion rents generated by the pricing system in the day-ahead market. Suppose, for instance, that the transmission line from node A to node B is congested, so the price is 10 at node A and 20 at node B. Generators at node A receive 10, but buyers at node B pay 20. If the capacity of the transmission line is 100 MW, there is a rent equal to $\$100(20-10) = \$1000$ that has to go somewhere. MISO is a non-profit organization, thus it distributes rights over these revenues among participants who pay for congestion.

These rights are financial instruments for which the payoff depends on the price difference between two nodes. For example, an FTR for 10 MW between nodes A and B with respective price LMP_A and LMP_B will pay the holder $10\,(\text{LMP}_A - \text{LMP}_B)$. FTRs are auctioned and anyone with sufficient funds can participate in the auction. The auction revenues are then distributed among the market participants who pay congestion charges.

Although virtual bids are purely financial in the sense that they are not related to any physical energy, they are treated by the market mechanism identically to physical bidders in the day-ahead market and affect congestion in the day-ahead market. For instance, consider the example above and suppose that there is only demand at node B, only generation at node A and the transmission line has a capacity of 100MW. At a price of 10, demand in B is only 95MW so the line is not congested. In this case, the market clearing price at B is 10. If a virtual participant bids 10MW at a high enough price and the demand at B is not too elastic, the cleared quantity will be 100MW and the price will increase until the market clears. In this way, a virtual bidder could create artificial congestion to increase the profits from an FTR between node A and node B. The profitability of this strategy depends on demand elasticity, but if the transmission line is close to its limits, a few MWs can be enough to change the price. If the FTR is large in terms of MWs, even a large marginal loss in the virtual market could be profitable. This strategy is described theoretically by Ledgerwood and Pfeifenberger [2012].

In what follows, we show that data is consistent with the use of virtual bids to create congestion and increase the value of FTRs. Although the data does not allow us to identify the virtual bidders who hold FTRs, evidence suggests that the FTR and virtual markets are correlated. Virtual profits are lower at nodes with larger FTR profits, more FTR MWs, as well as at nodes with higher congestion. Notice that this need not be the case because virtual profits depend on the difference between the day-ahead and the real-time price in a given node, while FTR profits depend on the difference in the day-ahead price between two nodes.

Our findings show that the failure to converge can be partially due to the fact that financial participants are entering the market just to change their position in the FTR market. Nonetheless, we find that there are still some arbitrage opportunities that could have attracted more virtual participants. Jha and Wolak [2013], on the other hand, conclude that there were no profitable strategies in the CAISO market after financial participants were introduced. They run three different mean tests on the premium and obtain the implicit transaction cost that would make arbitrage unprofitable, arguing that as the premium does not depend on its past realizations for half of the hours, more complex strategies based on lags will not do better. We repeat the same tests and find similar implicit transaction costs. In our case, however, the white noise tests are rejected for most of the nodes, suggesting a predictable component to forward premia. In fact, we show that a strategy taking linear predictions of the premium using lags can yield substantial positive profits. We speculate that the lack of entry can be due to credit restrictions, and show that virtual volume effectively decreased during the financial crisis.

The next section describes the market in more detail and the following presents a description of the data. Section 4 explains why market power can give rise to a day-ahead premium, section 5 details the puzzling observations from the virtual market, and section 6 shows the correlation between virtual bids and the FTR market. In section 7 we discuss why we do not observe entry even though there were profitable opportunities for virtual bidders during the period of study. Finally, we conclude the paper in Section 8.

2 Market description and institutional details

The Midwest Independent Transmission System Operator (MISO) is a regional transmission organization (RTO) that coordinates the electric power market in 12 states of the U.S. and the province of Manitoba in Canada. MISO is a non-profit organization regulated by the Federal Energy and Regulatory Commission (FERC) and its aim is to have low prices while maintaining a high level of reliability. It operates the transmission network by coordinating, controlling, and monitoring the use of the electric transmission system by utilities, generators and other participants. MISO does not own transmission assets and it is independent from all market participants. The MISO market has 2,168 pricing nodes and 368 market participants, and it serves more than 40 millions people.

Market participants

To 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.

Market participants may act in three distinct capacities in the energy markets: as a generator (supply), as a load serving entities (demand) and/or virtual bidders. Generators and load serving entities are linked to physical resources at particular pricing nodes whereas a virtual bidder requires no physical capacity. A generator will only sell energy and load serving entities will only buy energy at their respective nodes whereas virtual bidders may sell or buy energy at any node in the MISO footprint. These transactions are accomplished through participation in the MISO moderated auctions. The participant either submits a price sensitive bid consisting of a step function with up to ten combinations of price and quantity, or they may operate as price-takers by submitting only a quantity that they are willing to buy or sell at the clearing price. In the MISO market, bids submitted by buyers are referred to as *bids* and those submitted by sellers submit are referred to as *offers*. We will maintain this distinction from this point forward.

Any market participant can submit virtual bids or offers provided that she has sufficient approved credit. In the sample, 39 of 203 virtual participants are load entities, but from the way identifiers are assigned to participants it is not possible to know how many generators are there among virtual bidders. Table 1 has summary statistics on virtual bids and offers and Table 2 indicates the volumes traded in the energy market during the sample period, as well as the relative size of virtual participants.

The day-ahead and real-time markets

Energy is traded in the energy and operating reserves market, which has two parts: the day-ahead market, which is only financial though transmission constraints are taken into account when the market is cleared, and the real-time market, which is the balancing spot market. Participants with physical generation capacity submit bids into both of these markets to buy or sell energy, while virtual bidders and load entities are only allowed in the day-ahead market. A virtual bidders cleared MWs are reversed in the real-time market. I.e., if a virtual bidder buys 1 MW in the day-ahead market, it is understood that she will sell this energy in the real-time market, which leaves her payoff as the difference between real-time and day-ahead prices.

The day-ahead market accepts bids and offers between 11:00 and 15:00 hours of the preceding operating day. The market clearing results in an hourly schedule for each of the 24 hours of the operating day which begins at hour 00:00 (midnight). The real-time market accepts price sensitive supply offers up to 30 minutes prior to the hour and then uses these offers to balance demand requirements (equivalent to price-taking bids) every five minutes. The results are communicated five minutes before the operating time. Although the real-time market is cleared every five minutes, the price actually payed by or charged to market participants is the hourly average. Both day-ahead and real-time prices are locational marginal prices (LMP), i.e. the clearing price at each node is equal to the marginal cost of supplying one additional

MW at that node. LMPs may be decomposed into three components: the marginal cost of energy, the cost of congestion, and the cost of losses which exists because some energy is naturally lost when it moves over transmission lines. Congestion occurs when demand at a node cannot be supplied from a generator with a lower offered price without exceeding the capacity of a transmission line. As losses are typically small[2], the price difference between nodes is mainly due to congestion. Summary statistics for day-ahead and real-time LMPs in the MISO market can be found in Table 1.

Congestion exists because transmission lines have limited capacity: when the demanded MWs exceed the transmission line capacity, 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 there is only supply and the market is perfectly competitive, so the energy is sold at marginal cost. At node B, there is only demand and it is given by $Q(P) = 150 - 5P$. If there are no losses, the market clears with a price equal to the marginal cost and there is no congestion. This is illustrated in Figure 2 under the assumption that the marginal cost per MW is \$10 .

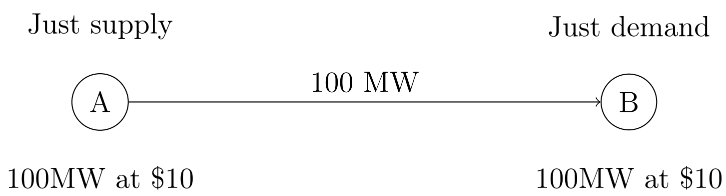

Figure 2: Two nodes without congestion

Now suppose that the marginal cost of producers at node A is \$5/MW instead of \$10/MW. In this case, consumers at B would demand 125 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 \$10/MW in order to clear the market and the extra \$5/MW over the marginal production cost corresponds to the congestion component of the LMP. This is illustrated in figure 3. For an example with three nodes, see Louie and Strunz [2008].

[2] 2% on average, without much variation.

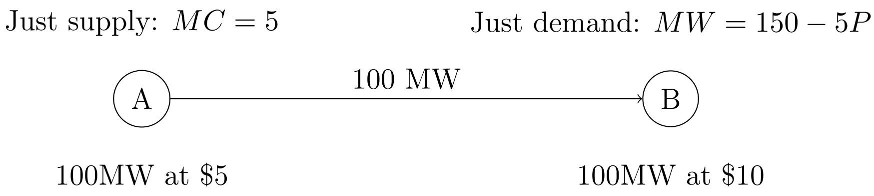

Figure 3: Two nodes with congestion

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

MISO's treatment of virtual bidders with respect to the RSG has varied over time in a way that affects incentives. Prior to April 1st 2011, RSG charges applied only to virtual supply bids and were constant across nodes. They were computed as $\mathrm{RSG}_i = \mathrm{MW}_i^S \cdot \mathrm{RSG\_RATE}$, 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 LMP minus the day-ahead one. For a virtual participant selling 1 MW in the day-ahead market, the payoff was $\mathrm{LMP}^{\mathrm{DA}} - \mathrm{LMP}^{\mathrm{RT}} - \mathrm{RSG\_RATE}$.

Since April 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. Notice that now RSG charges depend on the portfolio of nodes held by a participant, as cleared bids and offers can compensate each other for certain combinations of congested and uncongested nodes. The calculation of the charges for each participant is as follows:

$$\text{RT\_RSG\_DIST1}_h = \text{CMC\_DIST}_h + \text{DDC\_DIST}_h$$

$$\text{CMC\_DIST}_h = MAX\left\{\sum_n \left(MW_n^S - MW_n^D\right) \cdot \text{CCF}_{h,n}, 0\right\} \cdot \text{CMC\_RATE}_{h,n}$$

$$\text{DDC\_DIST}_h = MAX\left\{\sum_n \left(MW_n^S - MW_n^D\right), 0\right\} \cdot \text{DDC\_RATE}_{h,n}$$

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.

Figure 5 shows daily mean RSG charges per MW payed by virtual bidders over time. As it can be observed, charges significantly decreased after the policy change. Some summary statistics of RSG charges are presented in table 1. where RSG 1 is the charge until April 2011.

Congestion and the FTR market

The fact that prices differ among nodes depending on congestion creates revenues that have to be distributed among market participants because MISO is a nonprofit organization. This is done using two financial instruments: auction revenue rights (ARR) and financial transmission rights (FTRs). FTRs are financial instruments whose value depends on the difference in day-ahead congestion price between two nodes. For instance, an FTR for 10 MW between nodes A and B will pay the difference in congestion between these two nodes, yielding a positive payoff when B has a higher congestion component of the LMP. If there are no losses in the network, this payoff can be calculated as the difference in the day-ahead LMP between both nodes. These 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 used to hedge against congestion or just to speculate. Table 1 shows some statistics for FTRs sold in the monthly and yearly auctions. The clearing price is the price paid per MW of the financial right.

3 Data description

The data used in this paper is publicly available at the MISO website. The sample contains data on the energy market and the FTR auctions from January 2010 to September 2012.

There is data for bids, cleared quantities, clearing prices, and other bid characteristics for each participant at a node and hour level. These are available for both physical and virtual bids and offers, although the generator identifiers cannot be matched to physical demand and virtual participants, so it is not possible to know if generators submit virtual bids. Additionally, there is data available for the RSG charges that participants have to pay when there is a difference between their day-ahead commitment and their real-time actual quantity. Table 1 presents some description of the data.

4 Congestion, market power, and the day-ahead premium

Under nodal pricing, congestion creates local market power because it isolates a node or a group of nodes from the rest of the market. When one of the transmission lines connecting a node is at capacity, local buyers cannot buy energy at market level marginal cost because it is not possible to bring energy from the cheapest producer. Local generators can thus affect the price by playing with the quantities they offer in the day-ahead and real-time markets.

Consider first the case of a generator at an importer node, i.e. one at which demand at the lowest price available exceeds the capacity of the transmission lines connecting the node. As it was explained in the introduction, a monopolist with fixed capacity will maximize profits by intertemporally discriminating between the day-ahead and the real-time markets. Half of the energy generated in the MISO footprint comes from coal and quantities produced in these plants cannot be adjusted easily, which makes it reasonable to assume that at the day level they have fixed capacity. The demand is assumed to be unresponsive because only 2% of the demand bids are price sensitive, the rest comes from price insensitive bids which buy a fixed quantity at the market clearing price.

Take now the case of a generator at an exporter node B, which is the cheapest alternative for buyers at node A (given transmission constraints). Nonetheless, the line connecting B to A has limited capacity and when it becomes congested, the generator has to sell her energy in the competitive market at a lower price. In this case, the generator has incentives to withhold production from the day-ahead market to avoid day-ahead congestion in the line connecting A and B whenever this is possible. This allows her to get a higher price in the day-ahead -the price paid by the demand at A-, and to get the lower price only for the MWs sold in the real-time market. A more detailed example of this case, specific to the New York ISO, can be found in Saravia [2003]

If generators consistently underschedule their production in the day-ahead market and demand does not respond, the real-time price is lower than the day-ahead price and there is a day-ahead premium. We run regressions of the premium on the HHI, and the types of generation present in the market. Table 3 presents the results of regressions ran defining market as a node. As this is likely to be a too narrow definition of a market, we also define market as all nodes with the same real-time and day-ahead price. Results are in Table 4. The HHI is not always significant, but it has the expected sign when markets are defined by a unique price. Regarding generation types, the omitted one is steam turbine, the dominant in MISO, which uses coal and is permanently active because it is expensive to stop it or start it suddenly (fixed capacity in the short run). Peakers like pumped storage or diesel units are associated with higher premiums, likely because they are active in periods of higher demand and congestion. Intermittent resources like wind and run of river are correlated with a lower premium because they do not have much control of production and are generally small.

To prevent the abuse of market power and to increase the convergence between day-ahead and real-time prices, which increases the efficiency of the market, in 2005 MISO opened the energy market to financial traders without physical assets. As their profits depend on the day-ahead premium, they are expected to bid at nodes with expected price differences, thus making them smaller. Previous research comparing nodal markets before and after the introduction of virtual bidders has found that it decreases the premium in CAISO and NYISO [Jha and Wolak, 2013, Saravia, 2003] . Nonetheless, they do not show that they close the gap completely. Although in the MISO market the premium has decreased since 2005, we still find predictable differences between the day-ahead and the real-time market. Virtual bidders' behavior will be described in the next section.

5 Virtual bidders in the MISO market

This section describes the virtual bidders' behavior in the MISO market. It shows that some virtual participants seem to act in a way that is not consistent with profit maximizing behavior if profits are assumed to come from the virtual market only. In particular, three facts invite to a closer look: Virtual demand is much higher than virtual supply in spite of the positive day-ahead premium, some virtual bidders consistently loose money over time, and some virtual bidders bid to clear MWs at any price although this ends up in large losses.

Virtual bidders buy more than what they sell, despite the premium

Figure 6 shows the daily difference between virtual demand and supply in MWs over the sample period. Virtual bidders buy more than what they sell, and the difference is relevant in magnitude, as it is observed in Table 5.

There is no obvious reason to expect more demand than supply from virtual bidders, specially because the day-ahead premium, measured as the difference between the day-ahead and the real-time market for each node, is on average positive. This means that a virtual demand bid loses money on average because the buying price -the day-ahead LMP- is higher than the selling price -the real-time LMP. Table 6 shows the results obtained from regressions of the day-ahead premium on an intercept, which confirm the existence of a day-ahead premium, as has been previously documented by Bowden et al. [2009]. When only nodes with virtual demand are considered, the day-ahead premium is smaller, but still significantly positive. Figures 7 and 8 show that the premium is positive and it keeps accumulating, but sometimes it has big negative jumps.

In spite of the positive premium, supply bids are not necessarily profitable because they have to pay RSG charges. Table 6 shows the average premium before and after discounting RSG charges. A negative value means negative profits for supply bids because they sell in the day-ahead and buy in the real-time market. During the first period, buying and selling are equally attractive. In the second period, demand bids look more attractive because the premium is not significantly different from zero when RSG charges are included. Nonetheless, this computation assumes that the different bids are independent, thus ignoring the potential RSG discounts that come from bidding at different nodes. As we will show in the last section, there were profitable arbitrage opportunities during this period.

The existence of a positive day-ahead premium and an effective negative or zero premium for supply bids (after RSG charges are considered) means that, in principle, there is no reason to prefer demand or supply virtual bids. Nonetheless, as Table 7 shows, demand bids were as profitable as supply bids, or more. We observe as well that demand bids are considerably larger in volume than supply ones.

As observed, the day-ahead premium is significantly positive, but it has decreased over the years. Virtual volume has increased, as it is observed in Figure 9. Nonetheless, this is not enough to conclude that virtual bidders participation decreases the premium because the change in the RSG charges calculation affected the incentives for all the participants in the market and could have changed the equilibrium.

Some virtual bidders are consistently losing money

We observe that some participants steadily lose money over long periods, without leaving the market. This cannot be consistent with profit maximization by pure speculators, specially considering that they are experienced bidders in a complex market. Even a naive individual would leave the market after having large losses if there are no other sources of profit. The hypothesis of this paper is that the low profits are justified because they are the cost of increasing the value of FTRs by creating artificial congestion. Therefore, as long as losses are lower than FTR profits they can still be maximizing profits.

Figure 10 plots of 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. In fact, the cumulative profits of the ten largest virtual losers among virtual bidders consistently decreased (see Figure 11), indicating that some players were staying in the market despite having steady losses.

For some participants, the cumulative profits from Figure 11 look flatter after the change in the RSG charges computation, which is indicated by the vertical line. In principle, this could just be the consequence of lower RSG charges instead of showing a change in behavior. However, this is not the case. In the graphs, the blue line represents the actual cumulative profits earned by the respective participant, while the red one shows the part of the profits obtained from the difference between the forward and the spot price, without considering the RSG charges. The two lines move closely together and the red one does not continue to decrease after April, 2011. Moreover, some curves get flatter at the beginning of 2012, 8 months after the RSG change, which makes other explanations more likely. This is not necessarily evidence against the hypothesis of price manipulation because it may just be the consequence of the decrease in the premium (Table 6).

These observations suggest that some virtual bidders are doing something different from maximizing profits in the virtual market. Although in principle this could be due to some physical participants hedging, the price of the insurance would be high. The largest loser, for instance, lost 15 million dollars. The hedging explanation cannot be fully discarded because the data does not allow to know whether big losers have generating capacity. However, three of the ten largest losers are load entities (physical demand) and thus cannot be hedging.

There are different types of bidders

Virtual bidders can submit price sensitive bids that specify up to ten combinations of price and quantity, or just a quantity (price-taking bids). In the sample, a big part of the virtual bids are price-takers (see Table 9). 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 submit a bid with 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 the demand is very inelastic.[3] 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, it is observed that a group of participants placed bids far from the clearing price and that the largest losers of the market where in this group.

Figure 13 shows a plot of the relation between virtual profits and the bid premium for price-sensitive bids.[4] The bid premium is the difference between the highest price offered in the bid step function and the clearing price. As it is observed, Figure 13 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 the regressions of virtual profits on the premium, which also show that the effect is larger for supply bids (Table 8).[5] Figure 14 extends the sample by adding price-takers as well by assigning them a price premium of 1000. As observed, profits seem to be more variable among price-takers and bids with large bid premiums.

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.

[3] A virtual supply 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.

[4] Profits do not include RSG charges. In this way, they clearly reflect the difference between spot and forward case.

[5] Profits here do not consider RSG charges. These can only decrease profits.

6 Virtual bidders and FTRs

Virtual bidders do 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. There are at least 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. i to make their position depend on the difference in real-time congestion between the two nodes. This is a strategy frequently utilized by physical players to hedge real-time congestion risk. To understand this, consider an individual who holds an FTR for 1 MW from A to B and wants to move her position to the day-ahead. She can do this by virtually buying 1 in B and selling 1 in A. Then her profits would be:

FTR	$\text{LMP}_B^{DA} - \text{LMP}_A^{DA}$
Sell 1 at A	$\text{LMP}_A^{DA} - \text{LMP}_A^{RT}$
Buy 1 at B	$\text{LMP}_B^{RT} - \text{LMP}_B^{DA}$
Total	$\text{LMP}_B^{RT} - \text{LMP}_A^{RT}$

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.

The second way in which the FTR and virtual markets are correlated is when virtual bids are used to increase the volume enough to create congestion and thus increase the FTR payoff. Consider again the example presented above with nodes A and B for the case in which the marginal cost is 10, so there is no congestion. Imagine that there is a participant who bought an FTR for 100 MW between A and B at a price of \$1 per MW. As there is no congestion, initially this instrument yields a zero payoff. Now if this individual buys 10 MW at node B as a price-taker (or bidding a very high price), the demand would be shifted to $MW = 160 - 5P$. The demand at the marginal cost exceeds the 100 MW capacity of the transmission line between A and B, so the LMP at node B has to increase to 12. This is shown in figure 4.

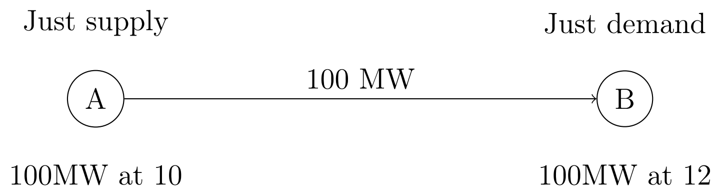

Figure 4: Using virtual bids to create congestion

Suppose that the real-time LMP at node B is 8. Then, the overall profits for this participant are the following:

	Calculation		Profits
Virtual Market	$\left(P^{RT} - P^{DA}\right) MW$	$(8 - 10)\,10$	-20
FTR	$\left(P_B^{DA} - P_A^{DA}\right) MW - P^{FTR} MW$	$(12 - 10)\,100 - 1 \cdot 100$	100
Total		$100 - 20$	80

Although this participant is losing money in the virtual market, the increase in the FTR profits greatly compensates 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 a demand inelastic enough to produce congestion with only a few MWs.

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 the virtual market 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. Nonetheless, many aspects of these two markets seem to be correlated in the MISO market, as we show below.

First, we establish that 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 10 shows some summary statistics for daily FTR profits, which are computed as total profits divided by the duration of the instrument in days, and for the daily average of FTR profits plus the virtual profits at the sink and source node. As it can be observed, the addition of the virtual profits increases the variability of the profits, but the gains from the FTRs seem to be enough to cover losses in the virtual market. This is computed by adding all the virtual profits or losses at a certain node to the FTR profits, because it is not possible to identify individual participants.

Besides the fact that FTR profits are enough to cover virtual losses, we observe that virtual profits are correlated to congestion, FTR profits, and FTR MWs at the node level. The latter is shown in Table 13 which reports the results from a regression of virtual profits on FTR MWs, in which each observation is a financial transmission right. Nodes with more FTRs, measured in MWs, have lower virtual profits and this is true both at the source and at the sink node. This makes sense if virtual bids are being used to create artificial congestion because larger FTRs make it possible to incur in larger virtual losses to increase congestion. Although the existence of an FTR at a certain node does not imply that a virtual bid would be unprofitable, virtual participants are expected to avoid unprofitable nodes when they are just speculating in the energy market, which would explain the correlation observed in Table 13. The results shown in this table come from a selected sample because it only includes nodes with FTRs in some period.

Additionally, nodes with larger FTR profits have lower virtual profits, as shown by the results of regressing FTR profits on virtual profits at the sink and source node presented in Table 11. The coefficients are small, but significant, and they show that nodes with higher FTR profits have larger virtual profits at the source and lower at the sink. Each observation is an FTR, so again the sample is restricted to nodes in which there is an FTR at some point in the sample period. Profits were computed as just the difference between day-ahead and real-time prices, without considering RSG charges. This is enough to detect irregular virtual behavior as RSG charges can only decrease profits and will not provide an incentive to bid. When virtual transactions that appear to be FTR moves are removed, correlations have the same sign and significance. 12 presents the results of the same regressions, but excluding virtual bids that appear to be used to move an FTR from the day-ahead market to the real-time one. We remove all pairs of demand and supply bids at different nodes, for the same quantity, at the same hour and by the same participant.

7 Potential arbitrage opportunities and entry

We have shown that there is a day-ahead premium in the Midwest ISO market and that virtual bidders have not closed it. Only part of them are actually arbitrageurs, the rest are using virtual trades to change their FTR positions and therefore submit price insensitive bids that do not improve convergence. The natural question then is why there is not more entry of pure speculators who could close the gap, since simple strategies could have made profits. As it can be observed in figure 8, just buying one MW at every node between July, 2011 and April, 2012 could get a trader $1,300,000 in profits. Moreover, we find that simple linear regressions have enough predictive power to build profitable arbitrage opportunities.

Our results differ from the findings of Jha and Wolak [2013] in the California ISO market, who conclude that there are no arbitrage opportunities once transaction costs are accounted for. They define an asset as a vector with the premium at each of the 24 hours, and use different tests to find the implicit transaction cost that would reject the null of profitable arbitrage opportunities. They argue that as the premium does not depend on past realizations, the mean test is enough because conditioning on lags is not likely to be profitable. We run the same tests for the MISO market and find similar values for the implied transaction costs. Using the most conservative test, the intersection-union, they find a transaction costs between 9 and 11 in the three aggregation nodes, and between 10 and 100 when all the nodes are considered. For the five hubs in MISO, where prices are aggregated to provide participants with price indexes, we find implicit trading costs between 3.7 and 4.2, depending on the hour and the hub; and values between 2.8 and 100 when all the nodes are considered. Jha and Wolak [2013] perform the other two tests, the square-based and the ellipsoidal. The square-based test actually rejects the null of arbitrage opportunities for all nodes with a trading cost of 10, and for all but 1% of the nodes with a transaction cost of 3 (the maximum value found by Jha and Wolak [2013] after the introduction of virtual bidding). The third one, the ellipsoidal test, rejects the null for 80% of the nodes with an implicit transaction cost of 3, and for almost all with a cost of 10. Results for the five trading hubs are not so different from Jha and Wolak [2013] : the square-based test implies trading costs between 0.75 and 1.5, while the elliptical test costs are between 1.48 and 2.2. Even though these tests reject the null of profitable arbitrage opportunities, in the MISO market the null hypothesis of white noise is rejected for a big part of the nodes in most hours (see Table 15). This indicates that lags of the premium can help to build a profitable strategy, which we find by using rolling regressions of the premium on lags.

Consider a potential entrant who runs rolling regressions using the past 30 days to predict the premium at every hour of the next day. The arbitrageur uses data for the last 30 days to estimate the coefficients of the following regression:

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

Where P_t is the premium at hour t and P_{t-24} is the premium at the same hour the day before. She uses this estimates to predict the premium on day 31 and buys (sells) whenever the predicted premium for an hour is lower than -5 (5). These regressions are run separately for each of the nodes in the MISO footprint. This strategy would have yielded profits, specially after April, 2011, when RSG charges considerably decreased. If the arbitrageur had bought 1 MW per hour at nodes with a predicted premium above the threshold of 5, she would have made up to $\$88,800$ in the most profitable node between April, 2011 and December, 2012. Considering all the nodes, she would have had profits for $\$22,218,333$ with an initial cost of $\$340,456$ (spent buying energy the first day). The mean of daily profits is $\$36,000$ and the standard deviation is $93,000$, which gives an annualized Sharpe ratio of 7.4. Figure 16 shows the daily cumulative profits that an arbitrageur could have made using this strategy.

Although this simple strategy would have yielded some profits, virtual trading has significant credit requirements. In fact, as figure 17 shows, virtual trading showed negative correlation with the TED spread during the financial crisis, indicating that credit restrictions may decrease virtual participation.

8 Conclusion

We present empirical evidence showing that the interaction between the FTR and the virtual energy market can explain the lack of convergence between the forward and the spot price, as well as some seemingly puzzling behavior of the financial participants. It seems that the virtual energy market is attracting many investors who want to either change their position in the FTR market, or use virtual energy bids to increase the value of their FTR instrument. Moreover, we show that there may be profitable arbitrage opportunities, which indicates that the efficiency of the market can be increased if what prevents entry is removed. We suggest that the low entry can be related to credit requirements, but further exploration is necessary.

Tables

Table 1: **Summary Statistics.** Data between 2010-01-01 and 2012-10-01. Prices are in \$/MW and volume is measured in MWs.

Variable	Observations	SD	Min	Mean	Max
DA LMP	44,633,058	14.68	-977	30.11	756.85
RT LMP	44,633,058	27.34	-1,318	29.36	3,225
Virtual Bid MW	12,442,024	40.98	0.1	12.47	2550
V. Bid DA Price	12,442,024	21.86	-977	30.99	756.85
V. Demand DA Price	6,501,716	17.37	-977	28.4	500
V. Supply DA Price	7,277,465	24.33	-98.33	33.02	757
RSG 1	10,531	2.27	0.01	1.88	20.99
CMC factor	8,543,258	0.11	-0.9993	-0.0143	1.00
CMC price	48,531	28.24	0.0	9.7	961
DDC price	13,177	1.51	0.0	0,78	27.39
FTR MW	563,746	18.56	0.1	8.44	1,351
FTR Clearing Price	563,746	1,057	-22,055	275.42	16,937

Table 2: **Total cleared quantities** from January 2010 to September 2012

	MW	%
Cleared Physical Demand	1,472,347,300	94.15%
Cleared Physical Supply	1,435,564,250	95.24%
Cleared Virtual Demand	88,892,660	5.85%
Cleared Virtual Supply	66,769,240	4.76%

Table 3: **Market Power and the premium**. A market is defined as a node. The reference energy type is steam turbine.

	2010	2011	2012
Intercept	1.68***	0.63***	0.77***
HI	−0.39**	0.05	−0.18
Combine Cycle ST	4.67***	7.48***	−51.32***
Combustion Turbine	2.11***	2.88***	2.03***
All Diesel Units	−0.48	1.03*	7.91
Run of River	−0.11*	−0.45***	−0.29***
Pumped Storage	3.54***	2.85***	1.85**
Combine Cycle CT	−3.01***	−2.26**	50.90***
Combine Cycle Aggregate	1.59***	1.08***	0.19·
Wind	−0.22***	−0.22***	−0.27***
Other Fossil	0.38***	−0.41***	0.10
Other Peaker	4.41***	6.03**	3.38**
Demand Response Type I		9.62***	
Demand Response Type II	−0.38*	−0.63**	0.44

Table 4: **Market Power and the premium**. Market defined as all nodes with the same day-ahead and real-time prices. The reference energy type is steam turbine.

	2010	2011	2012
Intercept	1.85^{***}	0.54^{***}	1.00^{***}
HI	-0.61^{***}	0.09	-0.47^{***}
Combine Cycle ST	4.68^{***}	7.48^{***}	-0.69
Combustion Turbine	1.77^{***}	2.28^{***}	1.29^{***}
All Diesel Units	-0.41	1.08^{*}	13.7^{*}
Run of River	-0.07	-0.41^{***}	-0.27^{***}
Pumped Storage	3.36^{***}	2.70^{***}	1.70^{*}
Combine Cycle CT	-2.95^{***}	-2.21^{*}	15.8^{***}
Combine Cycle Aggregate	1.58^{***}	1.10^{***}	0.16
Wind	-0.23^{***}	-0.17^{***}	-0.28^{***}
Other Fossil	0.44^{***}	-0.36^{**}	-0.02
Other Peaker	4.48^{***}	6.06^{**}	3.45^{**}
Demand Response Type I		9.67^{***}	
Demand Response Type II	$-0.32^{\cdot}$	-0.58^{*}	$0.49^{\cdot}$

Table 5: **Average daily virtual demand MW to virtual supply MW ratio and standard deviation of the ratio.**

Year	Virtual Demand to Virtual Supply	SD
2010	1.60	0.49
2011	1.40	0.37
2012	1.19	0.24

Table 6: **Day-Ahead Premium.** The mean is computed by running regressions of the premium on a constant. Each observation is the difference between the day-ahead and the real-time price at a particular node and hour. Two samples were used: all the 2,262 pricing nodes in the whole sample, and the only 1,485 nodes that have virtual bids.

	All		Only virtuals	
	First period	Second period	First Period	Second Period
Day-ahead premium	0.9 (0.005)	0.53 (0.005)	0.84 (0.004)	0.51 (0.004)
DA premium - RSG	-0.91 (0.006)	0.17 (0.16)		
Observations	$21,105,786$	$23,527,272$		

In the second period, RSG charges depend on the portfolio of nodes where the participant places bids. The day-ahead premium minus the RSG charges is computed assuming independent bidders at each node, i.e. it is a lower bound on the effective premium.

Table 7: **Virtual demand and supply bids**: mean profits, revenues and volume. Payoffs are revenue from virtual transactions without considering RSG charges, profits are revenues minus RSG charges. Both are computed per MW, i.e. the profits made by each bid divided by the bid's cleared MWs. In the second period, RSG charges are computed as if each bid were from a different participant (actual RSG charges are lower because there are portfolio effects).

	First period		Second period	
	Supply bids	Demand bids	Supply bids	Demand bids
Payoff per MW	3.05	1.11	2.13	0.88
Profits per MW	1.22	1.11	-0.65	0
Volume in MW	9.47	16.73	12.75	16.79

Table 8: **Bid Premium and Virtual Profit**. 4,163 observations for the supply profits, and 4,437,093 for the demand profits. This regressions do not include price takers.

Virtual Demand Profits:

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)	13.5322	0.7301	18.53	0.0000
Bid Premium	-0.3018	0.0064	-46.86	0.0000

Virtual Supply Profits:

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)	-12.8115	24.4117	-0.52	0.5997
Bid Premium	-9.7798	1.8003	-5.43	0.0000

Table 9: **Price-taking cleared virtual MWs.** A bid is defined as price-taking if the difference between the bid price and the cleared price is larger than 20. This definition is taken from the market reports on the MISO markets made by Potomac economics.

	Cleared Bids	
	Supply	Demand
Total MWs		97,096,349
MWs of price-taking bids	45,309,552	60,765,836

Table 10: **Daily FTR profits and virtual profits** at the sink and source node. We divide total profits over the number of days during which the FTR is valid to compute daily profits. the price payed for the FTR in the auction is included.

	Profits	Mean	SD	Min	Max
Whole Sample	FTR	59.81	596	-49,960	48,150
	FTR + Sink V	0.01	4,180	-183,700	120,300
	FTR + Source V	52.07	3,471	-267,400	114,500
First Period	FTR	74.92	691	-49,960	48,150
	FTR + Sink V	-41.9	6,176	-183,700	120,300
	FTR + Source V	-24,82	5,323	-183,600	114,500
Second Period	FTR	48.97	536	-40,820	46,310
	FTR + Sink V	-35.72	3,885	-250,800	63,510
	FTR + Source V	86.55	1,988	-267,400	46,310

Table 11: OLS regression of FTR profits on virtual profits at the sink and source node during the period in which the FTR was valid. Observations are FTRs.

	All	First Period	Second Period
Intercept	4,181	4,832	3,567
	(55.35)	(97.00)	(69.39)
Virtual profits at sink	-0.006	-0.007	-0.003
	(0.0008)	(0.0009)	(0.002)
Virtual profits at source	-0.003	-0.002	-0.007
	(0.001)	(0.001)	(0.004)

Table 12: **OLS regression of FTR profits on virtual profits** at the sink and source node during the period in which the FTR was valid. Observations are FTRs. Virtual bids suspicious of being FTR moves are excluded. These are a demand and a supply bid at the same hour, for the same MWs, and submitted by the same participant.

	All	First Period	Second Period
Intercept	4, 199	4, 851	3, 589
	(55.6)	(97.3)	(70.5)
V Profits at Sink	-0.007	-0.008	-0.003
	(0.001)	(0.001)	(0.002)
V Profits at Source	-0.002	-0.001	-0.005
	(0.001)	(0.001)	(0.004)

Table 13: OLS regressions of virtual profits at the sink and source node of FTRs on the FTR MWs at those nodes. Each observation is an FTR, so this sample only includes nodes with active FTRs.

	At the sink source	At the source node
Intercept	$160,068$	$212,329^{*}$
	$(80,621)$	$(83,266)$
FTR MW	$-2,300^{*}$	$-2,480^{*}$
	$(1,172)$	$(1,067)$
Observations	$55,265$	$55,146$

Table 14: Congestion Indicator Regression. Each observation is a node at a particular hour. Results are similar for the two periods.

	Estimate	Std. Error
(Intercept)	-19.135^{***}	1.565
congDA	-0.939^{***}	0.042
congRT	1.025^{***}	0.028

Table 15: Box-Pierce test results. Tests are perform for the premium at a certain hour and at a certain node. The table presents the percentage of nodes for which the null of white noise is rejected at 5% significance. As in Jha and Wolak [2013], we use 20 lags.

	% of nodes where the premium is not white noise
Hour 1	0.51
Hour 2	0.87
Hour 3	0.70
Hour 4	0.72
Hour 5	0.81
Hour 6	0.59
Hour 7	0.35
Hour 8	0.63
Hour 9	0.83
Hour 10	0.69
Hour 11	0.80
Hour 12	0.67
Hour 13	0.90
Hour 14	0.18
Hour 15	0.25
Hour 16	0.22
Hour 17	0.38
Hour 18	0.29
Hour 19	0.29
Hour 20	0.58
Hour 21	0.88
Hour 22	0.17
Hour 23	0.47
Hour 24	0.57

Figures

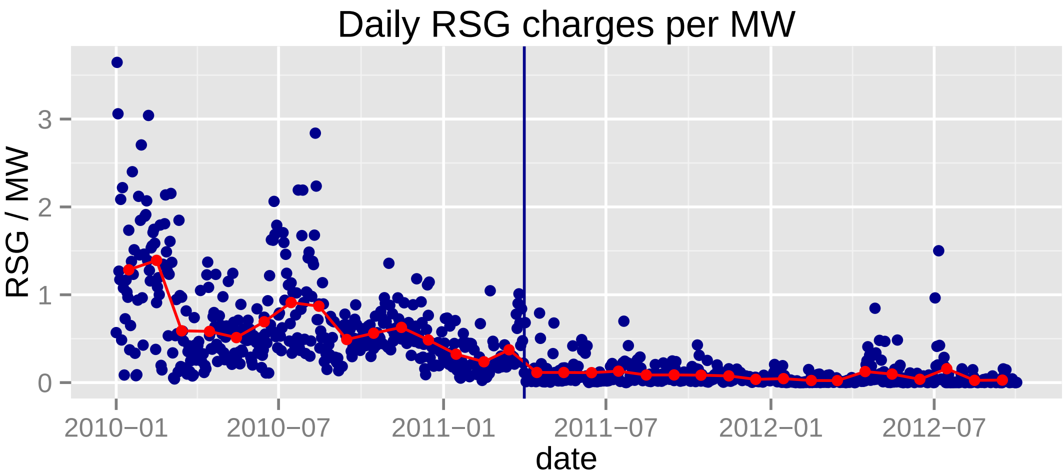

Figure 5: **RSG charges over time.** Each dot is total RSG charges payed by virtual bidders on a certain day, divided by the total cleared MWs during the day.The read line connects the monthly averages.

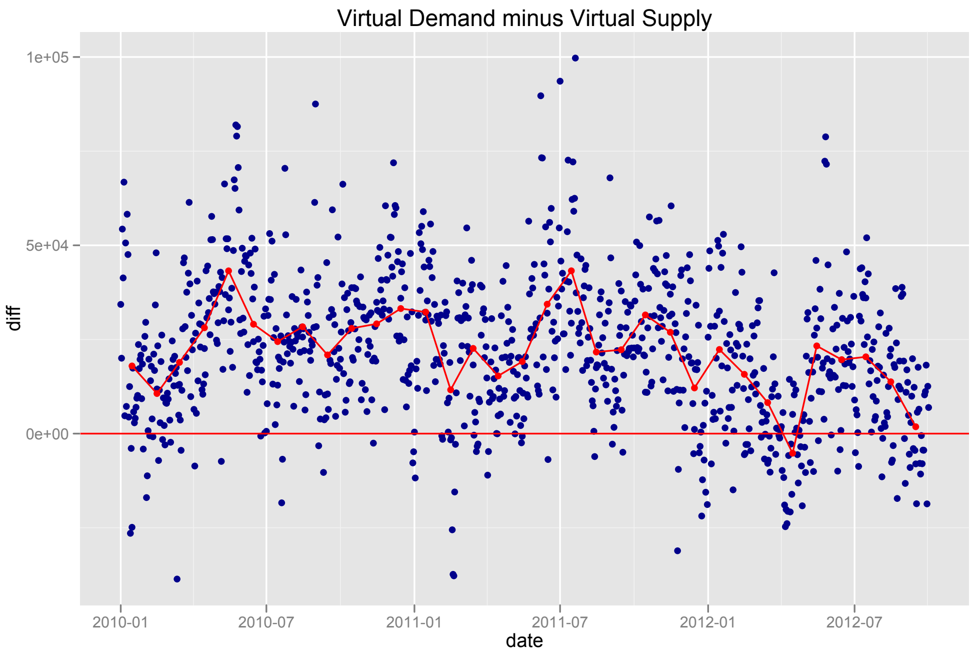

Figure 6: **Virtual bidders buy more than what they sell.** Each dot represents the difference between daily virtual demand and daily virtual supply in MWs. The line shows the monthly average.

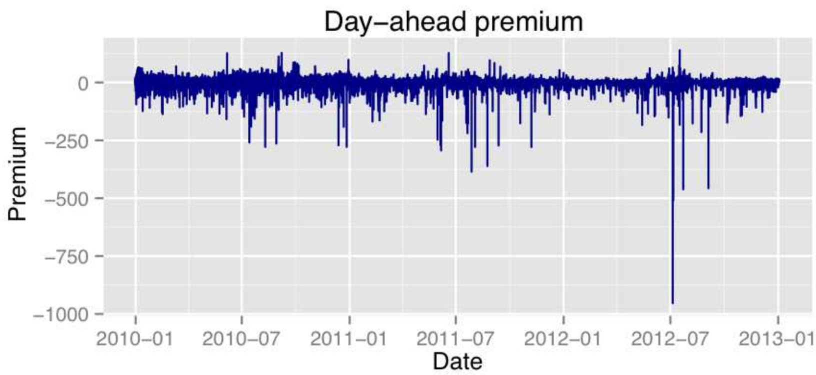

Figure 7: **Day-ahead premium over time.** Each observation is average of the day-ahead premium at an hour across all nodes.

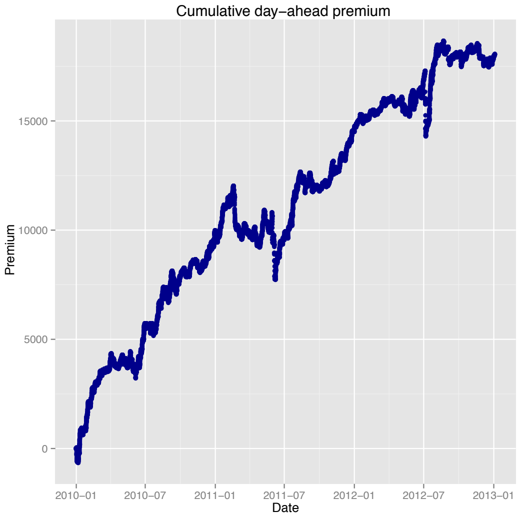

Figure 8: **Cumulative average hourly premium.** The average is taken across nodes.

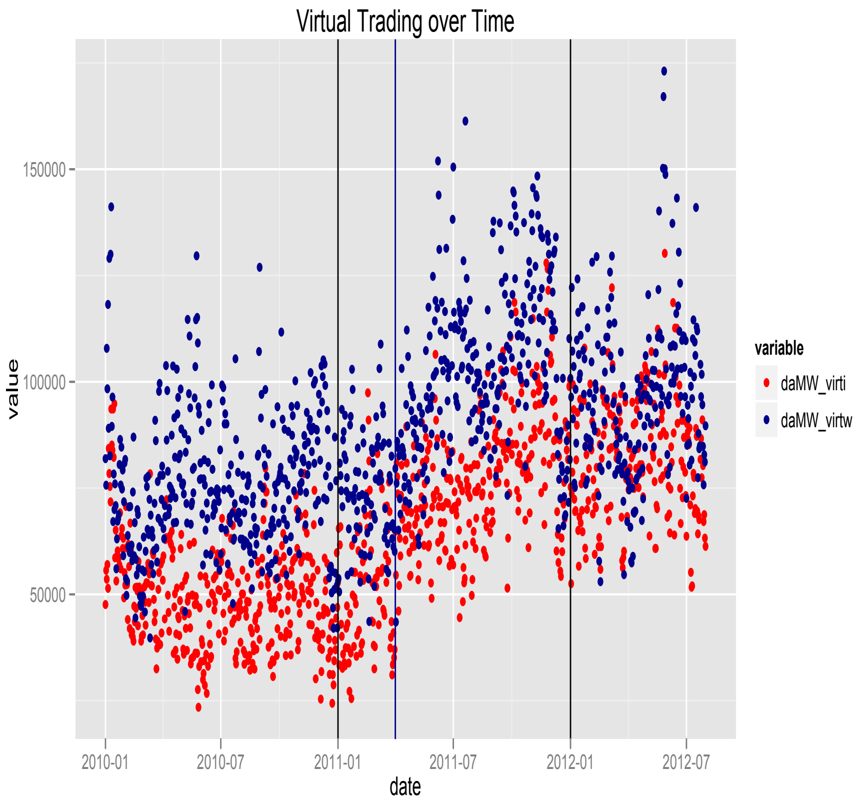

Figure 9: **Virtual Trading over time.** Each dot represents the total MWs cleared that day by virtual demand (blue) and virtual supply (red). The vertical lines indicates the change of year and the change in RSG charges in April 2011.

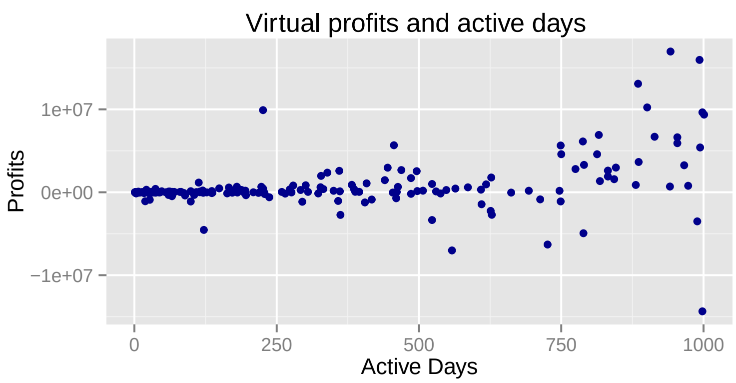

Figure 10: Virtual Profits and Active Days. The largest loser is excluded because it distorts the scale, she stayed for 700 days.

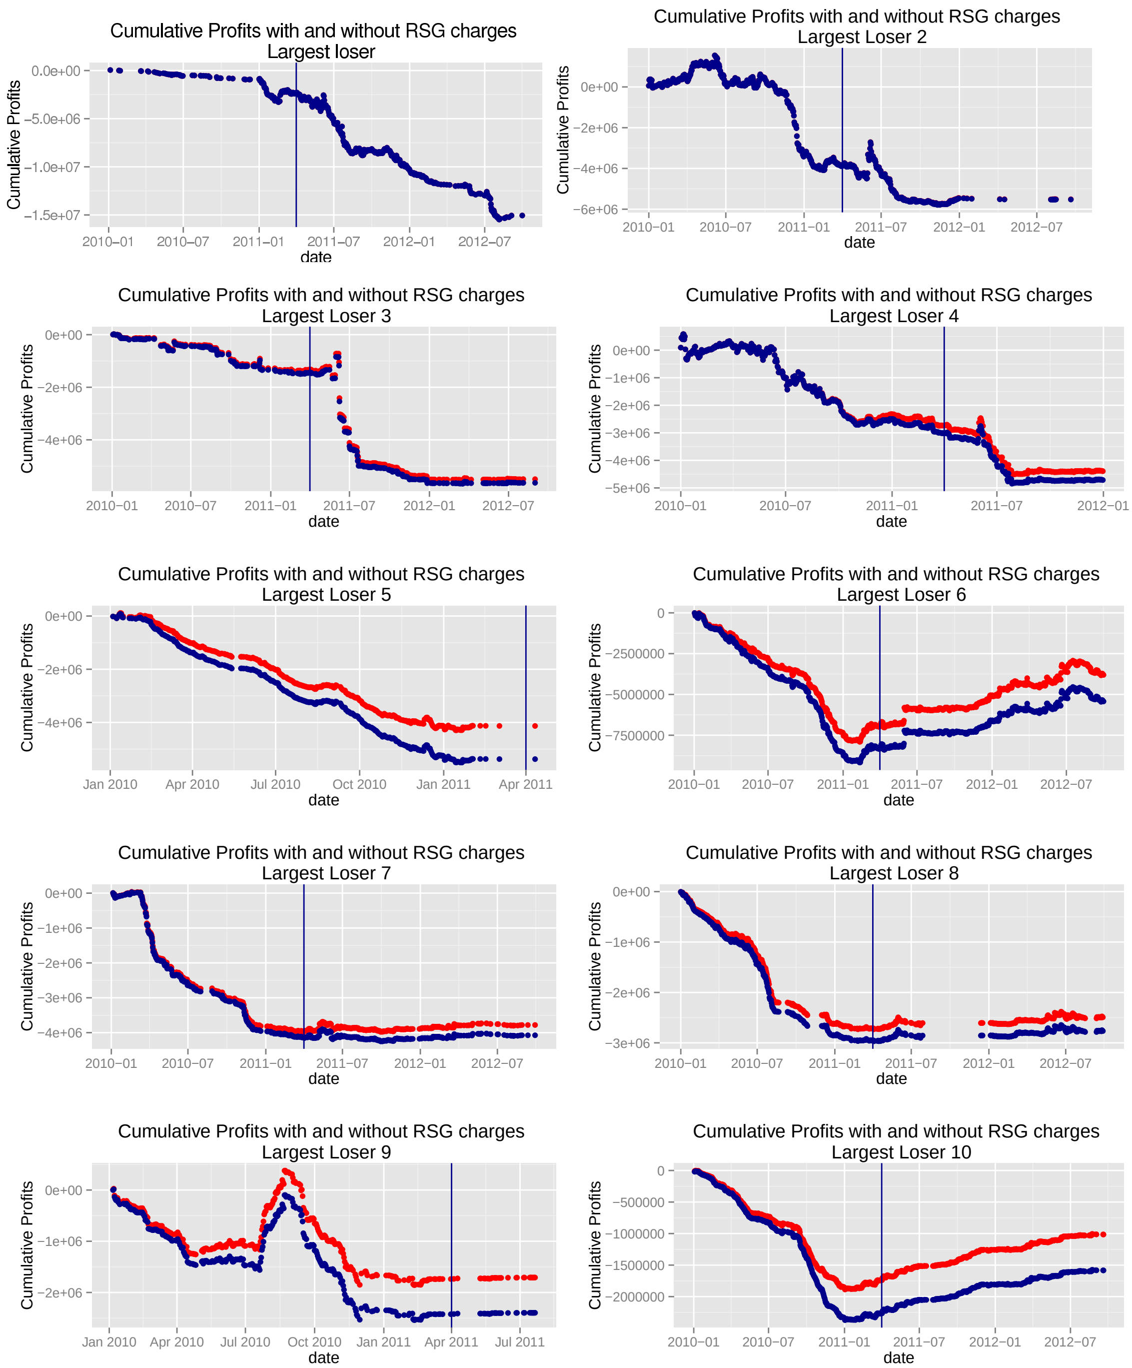

Figure 11: Cumulative Profits for the largest losers. The red dots do not include RSG charges.

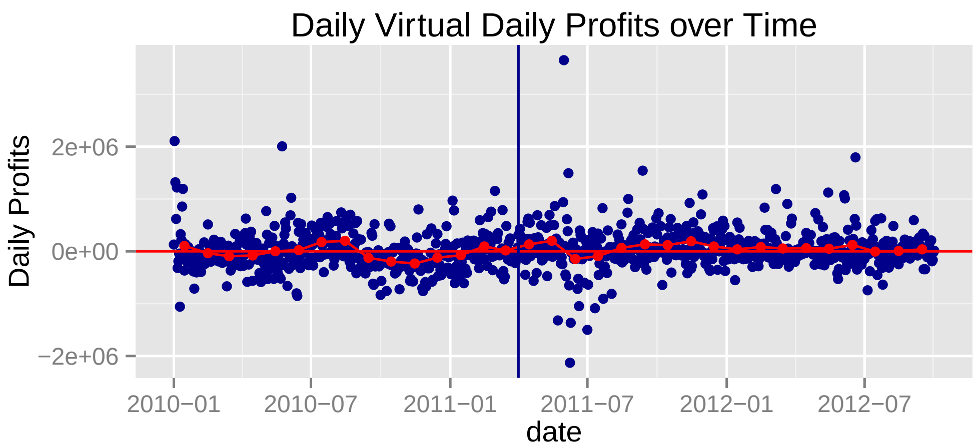

Figure 12: Daily Virtual Profits over time (This includes RSG charges, true profits

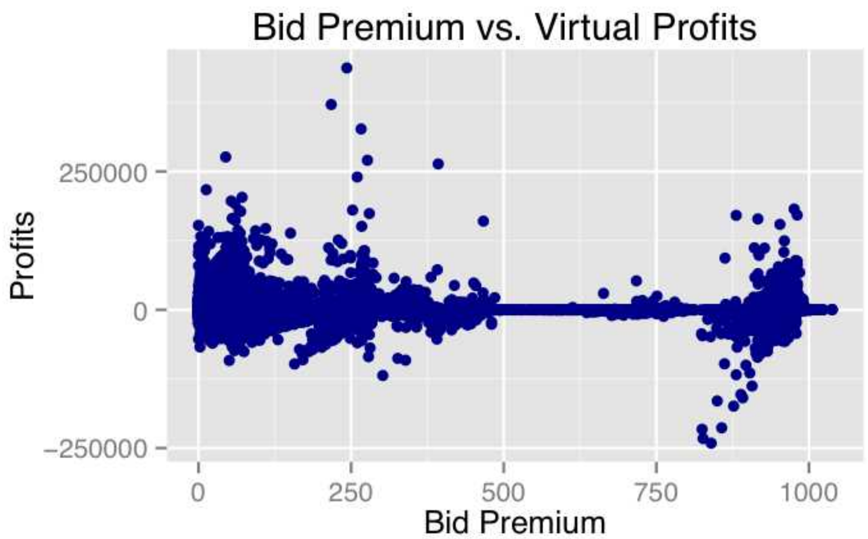

Figure 13: Two Types of bidders. Bid Premium and profits for price-sensitive bids only.

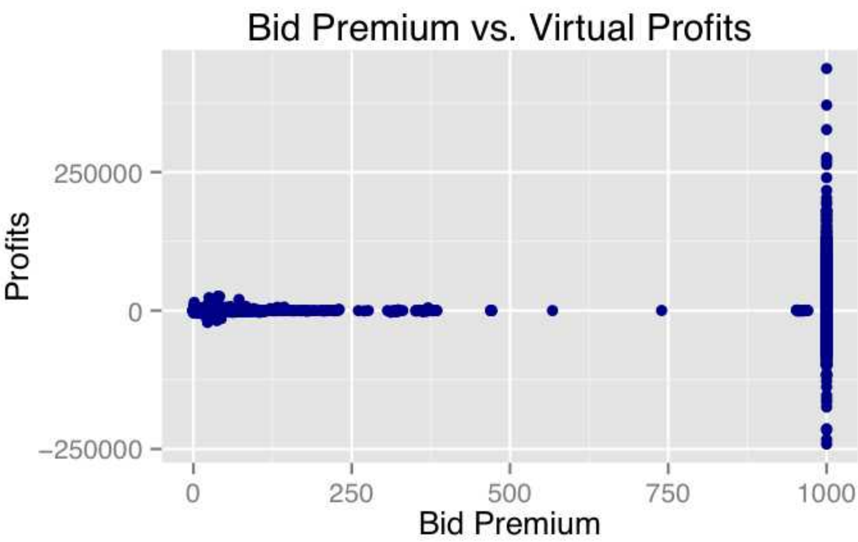

Figure 14: Bid premium and profits including price-takers.

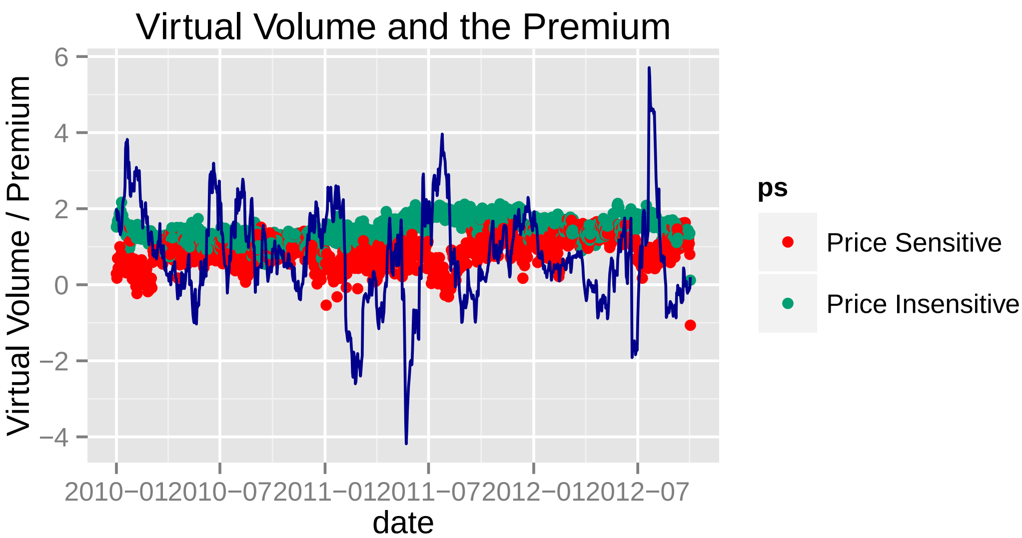

Figure 15: Price sensitive and price taking virtual volume

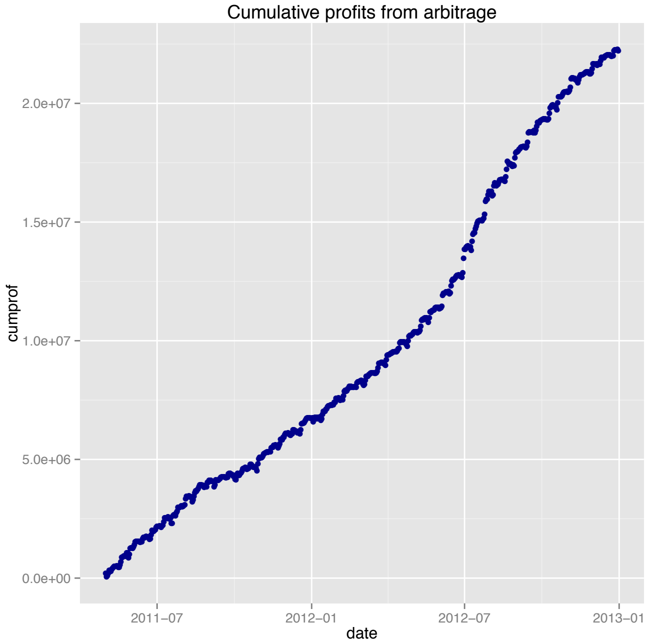

Figure 16: Cumulative profits from using monthly rolling regressions to arbitrage.

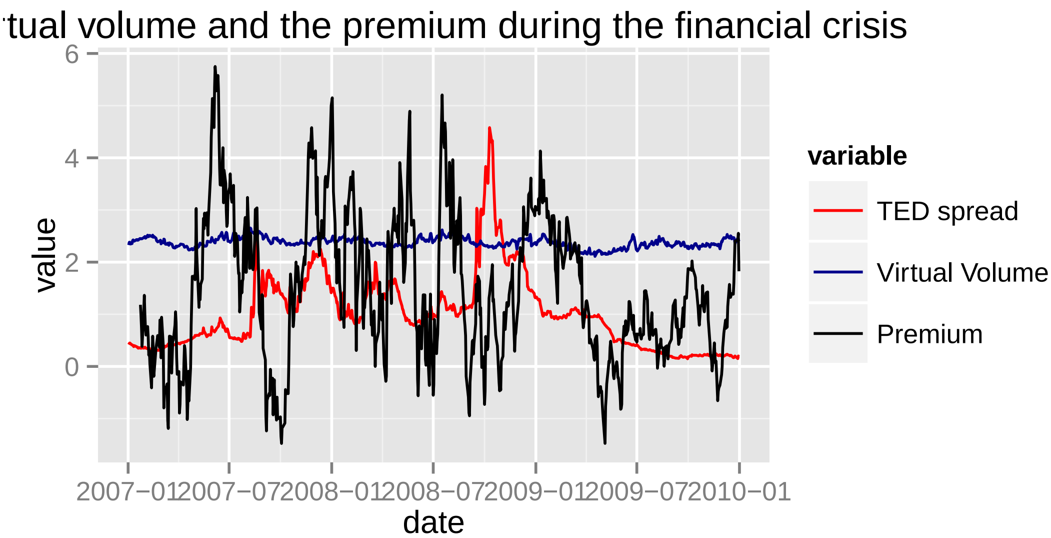

Figure 17: TED spread and virtual volume traded.

References

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.

Akshaya Jha and Frank A Wolak. Testing for market efficiency with transactions costs: An application to convergence bidding in wholesale electricity markets. 2013.

S.D. Ledgerwood and J.P. Pfeifenberger. Using virtual bids to manipulate the value of financial transmission rights. 2012.

Henry Louie and Kai Strunz. Locational marginal pricing in north american power systems. *ETG-Fachbericht-Netzregelung und Systemführung*, 2008.

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`.