MIT CEEPR

MIT Center for
Energy and Environmental
Policy Research

Financial Arbitrage and Efficient Dispatch in Wholesale Electricity Markets

John E. Parsons, Cathleen Colbert,
Jeremy Larrieu, Taylor Martin and
Erin Mastrangelo

February 2015

CEEPR WP 2015-002

A Joint Center of the Department of Economics, MIT Energy Initiative and MIT Sloan School of Management.

Financial Arbitrage and Efficient Dispatch in Wholesale Electricity Markets

John E. Parsons[*], Cathleen Colbert[†], Jeremy Larrieu[†], Taylor Martin[†] and Erin Mastrangelo[†‡]

February 10, 2015

Abstract

Virtual bidding is a type of transaction introduced into wholesale electricity markets to improve competition and pricing. This paper analyzes the theory behind virtual bidding and describes circumstances under which it does not work as advertised. The case for virtual bidding is predicated on an oversimplified model of the multi-settlement market design. The complexity of the unit commitment and optimal power flow problems forces the actual market algorithms to make compromises with the theoretical model. These compromises create situations in which virtual bidders can profit without improving system performance. Indeed, in these situations, virtual bidding can add real costs to system operation. The paper illustrates this with a specific case study of virtual bidding in California, and with a matching numerical illustration. The paper explains the general nature of the problem with experiences in other regions and other situations. The fault with virtual bidding identified in this paper needs to be incorporated into any assessment of the costs and benefits of virtual bidding.

[*]MIT Sloan School of Management and MIT Center for Energy and Environmental Policy Research (CEEPR).

[†]Federal Energy Regulatory Commission (FERC), Division of Analytics and Surveillance in the Office of Enforcement.

[‡]The opinions expressed in this paper are the authors' opinions only, and do not necessarily represent the views of the Commission nor of any of the Commissioners. John Parsons' work on this research was made possible by appointment as a Visiting Scholar in FERC's Division of Analytics and Surveillance in the Office of Enforcement from June 2013 through September 2014.

1 Introduction

Virtual bidding is a type of transaction introduced into wholesale electricity markets to improve competition and pricing. Under certain circumstances, virtual bidding works as advertised. But not always. Situations can arise in which the profits from virtual bidding are a purely parasitic transfer from electricity producers and consumers. Indeed, in these situations, virtual bidding can add real costs to system operation. This paper analyzes how virtual bidding functions and details the situations under which it malfunctions. These are illustrated with specific examples.

The theory behind virtual bidding relies on a very strong assumption that the different stages of organized wholesale markets operate identically. It attributes any spread between Day-Ahead and Real-Time prices to one factor: a deficiency or surplus in either demand or supply. As a remedy, financial traders are recruited and offered a bounty equal to that spread. By augmenting either demand or supply with virtual bids, these traders narrow the spread. The promise is that doing so improves system performance and lowers system costs. Unfortunately, this theory overlooks the many real world complexities that break the assumed identity between the stages of the market. A spread can arise between Day-Ahead and Real-Time prices for many reasons other than the simple deficiency or surplus in demand or supply. In these cases, added virtual demand or virtual supply is unlikely to re-establish the right equilibrium, although it may sometimes tighten the spread. In these cases, the bounty can be an expense that buys no return. Worse still, the added virtual demand or virtual supply, since it is not counterbalancing any deficiency or surplus in real demand or supply, can add to system costs.

The next section of the paper provides some essential background on virtual bidding and its place in wholesale electricity market design. Section 3 then presents the problem. It includes a case study of convergence problems and the impact of virtual bidding in California, and a related numerical example to illustrate how the profits to virtual bidding can be parasitic and how virtual bidding can add to system costs. It also generalizes this case and example by identifying diverse examples driven by the same underlying fault. Finally, it presents a critical assessment of the empirical literature on the impact of virtual bidding. Section 4 concludes.

2 Background on Virtual Bidding

Virtual bidding is a subsidiary element of a larger wholesale market design, and so to understand virtual bidding it is necessary to step-back and appreciate the standard design of organized wholesale markets in the U.S.

2.1 Multi-Settlement Markets

A key element of wholesale market design is the balance crafted between two objectives. On the one hand, the electricity system must be managed as an integrated whole in order to assure stability and performance. On the other hand, great economic benefit can be had in exploiting decentralized decision making among competing participants. The standard design employed in the U.S. accomplishes this by encouraging freely negotiated bilateral contracts between generators and load, and then requiring that all trade ultimately be fed into a centrally organized dispatch and scheduling system administered by an independent authority, often called an independent system operator (ISO). The operator constructs the lowest cost schedule from bids and offers supplied by market participants while simultaneously respecting a wide array of system wide reliability and security constraints. The centrally managed system is called "bid-based, security constrained, economic dispatch."

For this centralized system, U.S. markets have adopted a structure known as "multi-settlement markets" in which bidding and dispatch is managed in a set of successive runs. One run, performed a day in advance, is appropriately called the Day-Ahead energy market. Generators offer supply at various terms. Load-serving-entities bid demand at various prices. Bilateral transaction schedules can be submitted as well. Using these, the operator commits certain units for the next day and establishes a generation and load schedule for each of the day's 24 hours. This produces 24 hourly clearing prices at each node of the network which are called the Day-Ahead locational marginal prices (LMPs). Later, as each individual hour approaches, a Real-Time energy market is administered based on revised generation offers and an updated forecast of load, producing adjustments to the generation schedule. This produces 24 new hourly clearing prices at each node which are called the Real-Time LMPs.[1]

It is common to think of the Day-Ahead and Real-Time market runs as successive versions of the same bidding and auction process. While the two markets are clearly successive, they

[1] In addition to the Day-Ahead and Real-Time energy markets, the multi-settlement market design involves other products and decisions. For example, the system sources various ancillary services, including reserves. Sometimes these decisions are also managed on a market basis, and sometimes that market is integrated with the two energy markets. See O'Neill et al. (2011) for a broad overview of the diversity of implementations.

are not identical. Unit commitment decisions made in the Day-Ahead market cannot be easily revised in the Real-Time market. Some units have a minimum start time and if they have not been scheduled ahead of time, they cannot be quickly called upon. A unit's available ramp rate may be conditional on the current level of generation, so the range of generation available in the Real Time market is determined by the previous hour's level. Some units have multiple configurations and are costly to reconfigure, so that, for all practical purposes, the configuration choices made in the Day-Ahead market determine their availability in the Real-Time market.

The two markets are also different in how they are implemented. A key issue is the sheer complexity involved with finding the lowest cost commitment decisions and dispatch schedule subject to system wide constraints. Colloquially, the problem is presented as simple task of assembling a supply stack: generations units are arranged according to their marginal cost from lowest to highest, together with their capacity, and the required load is sourced first from the lowest cost unit in the stack, then from the next lowest and so on until generation exactly matches load. The true unit commitment problem has to confront many fixed costs and discrete choices which raise the computational complexity enormously. The true dispatch schedule problem, known as the optimal power flow problem, is also more involved since the generation units are located across a network and the selected generation schedule needs to respect an array of complex power flow constraints such as thermal limits on the network cables and voltage limits. The many non-linearities in the system make it extremely difficult to solve. In reviewing the state of research on the task, staff at the Federal Energy Regulatory Commission (FERC) wrote that the complexities of the problem are so daunting that

> *Even 50 years after the problem was first formulated, we still lack a fast and robust solution technique for the full [alternating current optimal power flow problem]. We use approximations, decompositions and engineering judgment to obtain reasonably acceptable solutions to this problem.*[2]

For example, certain implementations reduce the complexity of the transmission problem by using a linear representation of the flow of active power and ignoring the flow of reactive power. The physical relationship between the power flow and the voltages and angles in the buses of the network is difficult to fully incorporate in the modeling, as are the full suite of operational constraints on certain generating units. Proxy constraints are imposed or other adaptations are devised to coax the model to a good solution. These often work well so long as the system is operating within a familiar range, but they need to be constantly re-tuned

[2] Cain et al. (2012), p. 4.

as conditions change.

Importantly for the issues at hand in this paper, the approximations, decompositions and engineering judgements employed are different across the algorithms used to solve the Day-Ahead and Real-Time schedules. The extra complexity of the unit commitment problem solved in the Day-Ahead algorithm demands that the problem be simplified elsewhere. For example, the Day-Ahead algorithm usually employs a simplified representation of the transmission network, and the Day-Ahead algorithm solves for unit commitment and dispatch decisions in hourly blocks, without concerning itself with the minute-by-minute details of the dispatch. The Real-Time market cannot get by with the simplified representation of the transmission system often used in the Day-Ahead market, so a more complete representation is used. The Real-Time market also develops a schedule for generation at a much more granular time scale, often to 5-minute intervals. All of this extra detail, however, demands that the Real-Time algorithm be simplified in other ways. For example, by being myopic and optimizing the dispatch in a given hour without taking into account the needs of the system much beyond that hour, and by using the Day-Ahead generation schedule as a starting point and optimizing locally around it.

Additionally, the Day-Ahead market is used to give some participants strategic flexibility. Most load, for example, is ultimately price insensitive. This is reflected in the Real-Time market by the fact load serving entities are not allowed to submit bids. Instead, an updated forecast of load is used as a price insensitive demand curve to match with the updated supply curve in order to clear the Real-Time market. So, what the load serving entities are doing when they bid into the Day-Ahead market is allocating the portion of their load that they will buy at the Day-Ahead's clearing price, and the residual portion that will be bought at the Real-Time price. Similarly, some renewable generation that is non-dispatchable enters the Real-Time market as price insensitive, while the generator can bid it into the Day-Ahead market. The rules for what can be bid into which market, and how, varies across the regional markets. Providing this strategic option may improve competition, and may also help elicit valuable information about demand and renewable supply.

When the multi-settlement system is functioning well, the Day-Ahead solution is broadly consistent with the Real-Time solution. Obviously, changed conditions, such as the sudden outage of a generator or transmission line, will necessitate a Real-Time solution that differs from the Day-Ahead solution. Hence, the Real-Time market is called a balancing market. But, absent changed conditions, the solutions should match one another. Even in the presence of changed conditions, the prices should not be too far from one another, or at least not often. The Day-Ahead solution should anticipate regular contingencies and produce a dispatch schedule that is robust to them. Market monitors watch for important inconsistencies between the two solutions which may signal that some adjustments to the implementation need to be made.

Certain events in the PJM market in 2012 provide a nice illustration of the complications involved in committing units, of the associated approximations and judgments made differently in the Day-Ahead and Real-Time markets, and of how these approximations and judgments are retuned over time. A number of combustion turbine units in PJM received a surprisingly large amount in Lost Opportunity Cost (LOC) credits that year—$138 million. These payments were made necessary because the units had been committed and received substantial generation awards in the Day-Ahead market, but in Real-Time were not dispatched often enough to cover their costs. Why was this happening? Normally, the Day-Ahead algorithm would only commit these units if their expected hours of generation would earn sufficient revenue to cover their costs. Investigation revealed that the problem was due to an important consideration left out of the algorithm: the need for black start capacity and for voltage support.[3]

Before 2012, these two services had often been supplied by coal fired generation units located in the eastern portion of PJM's territory. These units are qualified to provide black start, and their location closer to important load sinks provides the right voltage support. Since these units were the low bidders in the Day-Ahead energy market, they were committed and dispatched. Although the Day-Ahead market algorithm was missing criteria requiring provision of these services, the services were provided as a costless by-product of the energy supplied by these units. The top panel of Figure 1 shows this situation.

This state of affairs has been disturbed in recent years. The falling price of natural gas and tightening environmental constraints that are especially costly to coal units, among other things, have re-ordered the economic competition between coal-fired and gas-fired generation units, so that the gas-fired units have won an increasing volume of Day-Ahead energy awards at the expense of the coal-fired units. The middle panel of Figure 1 shows this situation. This new ordering of the units exposes the missing criteria in the Day-Ahead algorithm. These gas-fired units are not qualified to provide black start, and their location creates problems for voltage support. In PJM, load is weighted to the eastern region, while gas-fired generation is weighted to the western region. The prevailing flow of power is west-to-east, and the shift of dispatch towards the west created problems satisfying the reactive transfer constraint. After

[3] See PJM (2013), pp. 119-127 and PJM (2014), pp. 143-144.

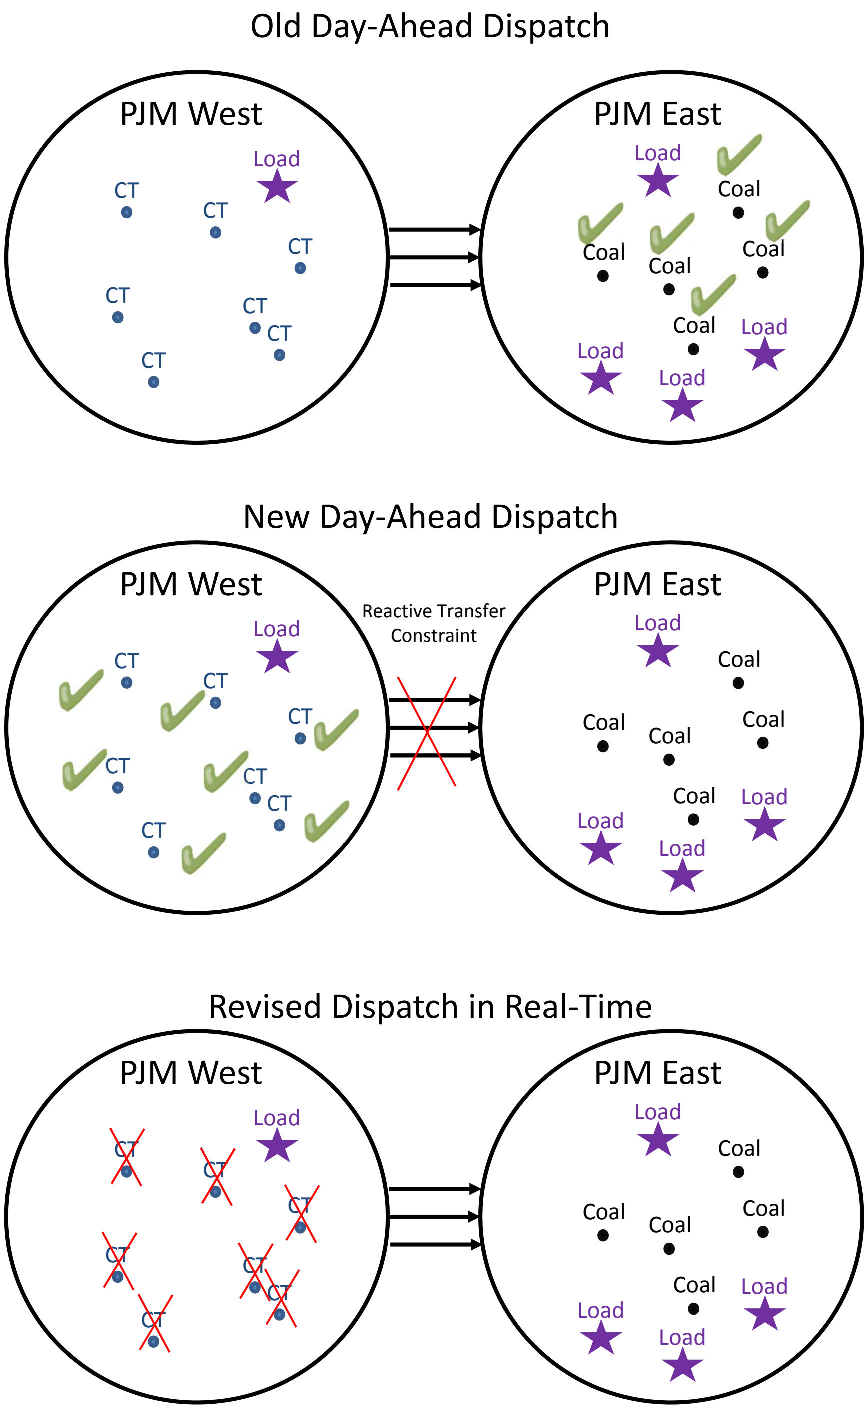

Figure 1: Illustration of Voltage Support Problem in PJM in 2012

the conclusion of the Day-Ahead market, when the system operator examined the dispatch plan and identified the failure to provide these two important needs, it would back down gas units and dispatch coal units in order to assure black start capability and voltage support. The bottom panel of Figure 1 shows this situation. This led to the large LOC payments to the gas units.

Having seen this happen, and having identified the cause, the system operator revised the Day-Ahead algorithm to incorporate the need for black start and voltage support which reestablished consistency between the Day-Ahead and the Real-Time solutions.

This situation illustrates the fallacy in thinking of electricity as a simple, undifferentiated commodity. In fact, the coal generation units provided multiple commodities, or a multi-attribute commodity: energy, but also black start capability and voltage support. Therefore, the simple, single supply curve for energy does not reflect the reality. The situation also illustrates the fact that the Real-Time market is not simply a later version of the Day-Ahead market, but a different implementation that can produce different outcomes even when the underlying supply and demand conditions are the same.

2.2 Convergence

One aspect of consistency between the Day-Ahead and Real-Time solutions is convergence between the Day-Ahead and Real-Time prices. This is measured using the DA/RT spread, i.e., the difference between the Day-Ahead LMP at a location and hour and the Real-Time LMP at the same location and hour. The spread can be measured at an individual node, or at a larger aggregate for which the LMPs are calculated such as a zone or hub. Since problems may manifest themselves in either positive or negative spreads, operators often focus on the absolute value of the DA/RT spread. Markets are said to converge when they produce a small DA/RT spread.

Annual market monitor reports for each of the regional wholesale markets regularly produce summary statistics on convergence, usually reporting both average DA/RT spreads and average absolute DA/RT spreads. Table 1 shows some of the statistics reported by the Market Monitoring Unit for the NYISO in its 2013 State of the Market Report. These numerous reports also include detailed analyses of the episodic and diverse situations in which persistently large DA/RT spreads, whether positive or negative or both, seem to be symptomatic of underlying problems in market operations.

In contrast, the large academic literature focuses almost exclusively on the average DA/RT spread. This literature has analyzed average DA/RT spreads across several regional

Table 1: Price Convergence between Day-Ahead and Real-Time Markets in Select Zones of the NYISO, 2011-2013: annual average DA/RT spreads and annual average absolute DA/RT spreads, in percent. Copy of Table 3 from NYISO (2014), p. 29.

	Avg. Diff % (DA - RT)			Avg. Absolute Diff %		
Zone	**2011**	**2012**	**2013**	**2011**	**2012**	**2013**
West	1.4%	0.0%	-1.9%	24.0%	26.4%	36.3%
Central	1.1%	0.6%	1.3%	25.7%	25.5%	29.5%
Capital	2.6%	2.9%	4.5%	28.1%	27.0%	33.1%
Hudson Valley	0.9%	0.9%	-0.8%	30.0%	29.9%	33.9%
New York City	1.8%	0.8%	-1.4%	32.4%	31.4%	35.0%
Long Island	0.9%	1.7%	-6.5%	35.5%	42.1%	46.5%

markets and over different time windows, consistently finding spreads that are statistically different from zero.[4] Conditional on a number of variables, these average spreads are sometimes positive and sometimes negative, and the range is large. Positive average spreads are more prevalent than negative. The conditional variation in the spreads is important to keep in mind since it means that averages across different conditions may obscure sizable, but offsetting spreads. The easiest form of conditioning is to sort the data according to season or month of the year and according to peak- or off-peak hours or by each hour of the day. Positive average spreads are more likely during peak hours and summer months. Other research attempts to identify the underlying state variables driving the conditional variation,

[4] Longstaff and Wang (2004) study DA/RT spreads in PJM's Eastern Hub in the period between June 2000 and November 2002. Douglas and Popova (2008) extend the dataset to December 2004 (and also backwards to January 2000), Ullrich (2007) extends the dataset to May 2007, and Haugom and Ullrich (2012) extend the dataset to December 2010. Pirrong and Jermakyan (2008) analyze average DA/RT spreads in PJM's Western Hub between 2000 and 2003. Saravia (2003) analyzes average DA/RT spreads in two regions of the NYISO from November 1999 to late-2003. Nakano (2007) looks across all zones in NYISO from November 1999 through November 2004. Hadsell and Shawky (2006) analyze peak hour average DA/RT spreads in all zones in NYISO from January 2001 to June 2004, and Hadsell and Shawkey (2007) look at DA/RT spreads in all hours in two zones using data extended to March 2005. Hadsell (2008) analyzes average DA/RT spreeads in eight zones of ISO-NE from January 2004 to December 2007, and Hadsell (2011) does the same for the period of March 2003 to February 2007. Werner (2014) analyzes average DA/RT spreads in the Southeast Massachusetts zone of ISO-NE between March 2005 and June 2011. Bowden et al. (2009) analyzes average DA/RT spreads in five hubs in MISO from September 2005 to December 2007. Birge et al. (2013) analyze average DA/RT spreads in MISO from January 2010 to September 2012 at nodal detail. Jha and Wolak (2013) analyze average DA/RT spreads at CIASO's three main load aggregation points—for PG&E, SCE and SDG&E—from April 2009 to December 31, 2012 and do some tests on spreads at other nodes. Their paper does some analysis of the variance of spreads. Borenstein et al. (2008) study 1-day forward v. spot spreads in the two main regions of the California market between April 1998 and November 2000, when an earlier market design was in operation.

such as the level or volatility of demand or a variety of operational constraints, such as plant outages, low natural gas storage, or reduced transmission capacity. Positive average spreads are more often observed on high load days, high volatility periods, and when other factors tighten.

As a benchmark, the major hubs of the PJM market often see an average DA/RT spread of 2-5% depending upon the hours and seasons over which the average is taken. In some hours and seasons the average premium may be negative and in some it may be even more positive. Larger average DA/RT spreads, whether positive or negative, arise at specific nodes.

Evaluating convergence is more difficult than it sounds at first blush because the right benchmark is hard to pin down. Many analysts assume that the expected DA/RT spread should be zero. However, this ignores the problem of risk and the appropriate risk premium. Real-Time prices are more volatile than Day-Ahead prices, and some participants may be willing to pay a premium in the Day-Ahead market to avoid the risk of the Real-Time price. Risk premia have been repeatedly documented across a wide range of commodities, as for many financial securities, and we should expect the same to be true for electricity. Whether the market risk premium is positive or negative or zero whether it is sellers who pay buyers, or buyers who pay sellers, or a wash is a complicated issue determined by industry factors as well as by the place of the industry within larger financial markets and the macroeconomy. Indeed, the right premium can be positive in some hours and negative in others, as the model of Bessembinder and Lemmon (2002) predicts, or may be conditional on other factors that vary over time. There is plentiful documentation of large risk premia in longer term electricity forward or futures prices. See, for example, Pirrong and Jermakyan (2008) who analyze the market risk premium in PJM's Western Hub between 1997 and 2005 for terms extending out several months, Bessembinder and Lemmon (2002) who analyzed the market risk premium in California between April 1998 and July 2000 using month-ahead forward contracts versus CALPX spot prices, Bunn and Chen (2013) who analyze both 1-day and 1-month forward premia in the UK electricity market from February 2007 to February 2010.[5]

[5]Bunn and Chen (2013) cite additional papers:

> *Hadsell and Shawky (2006), Diko et al. (2006), and Gjolberg and Johnsen (2001), Weron (2008) as well as Daskalakis and Markellos (2009) find significant premia in the NYISO, APX, Powernext and Nord Pool long-term electricity markets respectively, whilst Bierbauer et al. (2007), Kolos and Ronn (2008), Benth et al. (2008), Daskalakis and Markellos (2009) and Redl et al. (2009) and Kolos and Ronn (2008) find a negative forward premium for monthly, quarterly and yearly contracts at the EEX (German) market. Regarding the latter, Benth et al.(2008) relate the term structure of the forward premium to the net hedging demand of consumers and producers, producing a model that yields decreasing absolute values of forward*

This literature accords with a larger literature finding risk premia embedded in futures prices across a wide array of commodities. Nevertheless, both Borenstein et al. (2008) and Jha and Wolak argue for a zero risk premium and argue that any DA/RT spread reflects inefficiencies.

2.3 Virtual Bidding

Virtual bids are a special class of bids authorized in the Day-Ahead market designed to arbitrage DA/RT spreads. A virtual supply offer—also known as an increment offer, or INC—clears when the offered price is less than the resulting Day-Ahead price. A virtual demand bid—also known as a decrement bid, or DEC—clears when the bid price is greater than the resulting Day-Ahead price. However, virtual bids never result in an obligation to supply or take power. Instead, they earn a cash payoff that is primarily a function of the DA/RT spread. The gross payoff to a cleared virtual supply offer is the DA/RT spread multiplied times the quantity cleared:

$$\pi_S = (DA - RT)Q, \tag{1}$$

where DA is the Day-Ahead LMP at a particular location and particular hour, and RT is the Real-Time LMP for the same location and hour. The gross payoff to a cleared virtual demand bid is the negative DA/RT spread multiplied times the quantity cleared:

$$\pi_D = (RT - DA)Q. \tag{2}$$

Since virtual bids payoff in cash, the bidder can be a financial entity that neither owns generation nor serves load, although generators and load serving entities can also submit virtual bids. Since virtual bidding is advocated as a way to improve convergence it is also known as convergence bidding.

In addition to this gross payoff, cleared virtual bids sometimes incur other charges which can be material to the net profitability of a virtual bidding strategy. Some of these charges depend upon the outcome of the Day-Ahead and Real-Time markets. The rules for these vary by ISO, and have varied through time as well. Since the main points of this paper hold with or without these charges, we do not go into them in any more detail.

Virtual bids directly impact the unit commitment and dispatch of generators as well as payments between generators and load. A rough sense of the impact can be conveyed using

premia (eventually getting negative) when time to maturity or delivery period length increases.

the simplified presentation of the Day-Ahead market as a double auction, with generator offers compiled into a supply stack and load-serving entity bids compiled into a demand stack and the dispatch schedule and prices determined by the intersection of supply and demand. In constructing the supply stack, virtual supply offers are treated just like generators' supply offers—they shift the supply stack out to the right. In constructing the demand stack, virtual demand bids are treated just like load-serving entities' demand bids—they shift the demand stack out to the right. The supply and demand stacks augmented with the virtual bids determine the clearing Day-Ahead price and which bids and offers clear. Other things equal, cleared virtual supply offers reduce the total quantity of physical generation scheduled in the Day-Ahead market and lower the Day-Ahead price. Correspondingly, cleared virtual demand bids add to total quantity of physical generation scheduled in the Day-Ahead market and raise the Day-Ahead price.[6]

Virtual bidding is a regular feature of wholesale market design in regions across the U.S. that operate according to the standard multi-settlement design. The largest of these regional markets, the mid-Atlantic's PJM, began operating its multi-settlement market in June of 2000, and a version of virtual bidding was a part of its operation from the start. New York's multi-settlement market, operated by the NYISO, began operation in November 1999, and virtual bidding was incorporated two years later, in November 2001. In 2002, FERC incorporated virtual bidding into its proposed Standard Market Design.[7] The New England region's multi-settlement market, operated by ISO-NE, included virtual bidding from its beginning in March 2003. The midwest region's multi-settlement market, operated by MISO, included virtual bidding from its beginning in April 2005. The Texas wholesale market, operated by ERCOT, shifted to a nodal design with a centralized multi-settlement structure in December 2010, and virtual bidding was included as a part of this shift. California's mutli-settlement market, operated by CAISO, began operation in April 2009, and virtual bidding was incorporated in February 2011. Most recently, the central plains region's mutli-settlement market, operated by SPP, included virtual bidding from its beginning in March 2014. While details of virtual bidding vary across regional markets, and have varied over time within each regional market, the basic design has been constant.

Table 2 shows the volume of cleared virtual bids in 2013 by ISO—virtual demand, virtual

[6]This is only a first approximation of the impact of virtuals on dispatch. For example, after the Day-Ahead auction is complete, the ISO may examine the results and determine that additional generation needs to be dispatched in order to assure reliability. If the auction had settled with net virtual supply, then less generation was dispatched due to virtuals. But the ISO's reliability choices can reverse this outcome.

[7]FERC (2002).

supply, and spread contracts which are a combination of one virtual demand and one virtual supply at different locations. The volume shown is across all hours and across all nodes. As a benchmark against which to measure the volume, the table shows the total load in 2013 by ISO. In measuring the significance of virtuals, one needs to consider alternative metrics, and the table provides two crude metrics: a total and a net. The total equals virtual demand plus virtual supply plus two times virtual spreads. The volume of spread contracts is doubled since each contract is equivalent to one virtual demand and one virtual supply bid at different locations. This raw total is not the complete picture. Virtual demand and virtual supply bid into the same hour and the same location cancel one another out, so totaling them together exaggerates the impact of virtual bidding. Of course, virtual demand and virtual supply do not have to be placed at exactly the same location for their effects to cancel out or to substantially offset one another. Since a network consists of many different load nodes and generation nodes, operationalizing the definition of 'same node' is technically quite complicated. For simplicity, the table shows an extreme choice by netting all demand and all supply within the system at the same hour. This measures the scale of virtuals with respect to their potential impact on the system energy price, while disregarding their scale with respect to potential impact on the price of congestion or losses at different nodes or zones. Since some hours will have net demand and other hours will have net supply, the net figure is broken down into two pieces: one piece is the sum across all hours in which there is net demand and the other piece is the sum across all hours when there is net supply. Measured on a total basis, PJM is by far the most active virtual market, due overwhelmingly to the volume of spread trades. Absent the volume of spreads, NYISO is the most active, with total virtuals equal to nearly 19% of load. Measured on a net basis, NYISO is the most active virtual market, with net virtuals equal to more than 8% of load.

The market monitor reports for the various regional markets generally indicate that the majority of these virtual trades are placed by financial participants—i.e., those having no load serving obligations, no physical generation and no marketing services in the region.

2.4 The Demand/Supply Shift Theory

Virtual bidding is a tool for enhancing competition in the Day-Ahead market by opening up the auction to traders without physical generation or load. Advocates have cited several distinct reasons to believe that without virtuals, the auction may not be sufficiently competitive.

First, market power is a longstanding problem in electricity markets. If a generator has

Table 2: Total Cleared Virtual Bids by ISO - 2013. (000 MWh)

ISO	Demand	Supply	Spread	Total	Net	Net D	Net S	Load
PJM	63,142	44,989	454,792	1,017,714	18,617	18,385	232	791,184
MISO	30,666	24,510		55,176	9,421	1,632	7,788	516,132
CAISO	16,353	19,661		36,014	6,292	1,492	4,800	234,356
NYISO	8,386	21,991		30,378	13,889	142	13,747	163,514
ISO-NE	2,012	1,797		3,809	1,857	1,036	821	127,206

Notes:

The volume in the column 'Spread' is for PJM's Up-to-Congestion bids, which are a type of virtual bid representing virtual demand at one location and virtual supply at another location.

Total = Demand+Supply+2*Spread

Net system energy, demand = sum(max{VD-VS,0}) across all hours.

Net system energy, supply = sum(max{VS-VD,0}) across all hours.

market power, it can raise the Day-Ahead price by withholding supply. Alternatively, if it is a load-serving entity that has market power, then it can lower the Day-Ahead price by withholding demand. These actions open up a spread between the Day-Ahead price and the expected Real-Time price: withholding supply creates a positive expected DA/RT spread, while withholding demand creates a negative expected DA/RT spread. Through the Day-Ahead unit commitment process, these actions may also reduce the availability of supply responsiveness in the Real-Time market which has additional negative consequences. There is evidence that the large California utility, PG&E, withheld demand this way in 2000—see Borenstein et al. (2008). There is evidence that some large owners of generation in certain regions of the NYISO exercised market power this way when it first opened its Day-Ahead market in 1999—see Saravia (2003) as well as Chaves and Perez (2010).

Second, an array of institutional features of this industry create differential bidding incentives across the two markets. An obvious example is when the price cap imposed on one market is higher or lower than the cap imposed on the other market. A subtler example is the regulation of load-serving entities which sometimes provides for a pass-through of Day-Ahead prices to the retail tariff, but does not allow a full pass-through of Real-Time prices, penalizing the company for paying high Real-Time prices. Moreover, load-serving entities are often prohibited from speculating, and certain strategies to profit off of the difference between Day-Ahead and expected Real-Time prices may be designated as speculation, so that company management is very conservative in its bidding strategy.

Finally, even when there are no explicit obstacles to competitive bidding by generation and load, there may be a value to expanding the pool of available bidders beyond those participants. Virtual bidders may bring to the market expertise in forecasting the complicated patterns of load and the complicated dynamics of the transmission system. Enabling that expertise to work its way into the auction may improve dispatch and pricing.

Any expected discrepency between the Day-Ahead and Real-Time prices is a profit opportunity for a virtual bidder. Virtual bids placed to capture this profit shift the supply or demand curve which changes the Day-Ahead clearing price and reduces the DA/RT spread. For example, suppose that a load serving entity withholds demand in the Day-Ahead market, lowering the Day-Ahead clearing price below the expected Real-Time clearing price. This makes a virtual demand bid profitable. The virtual demand bid shifts the demand curve to the right, increasing the Day-Ahead clearing price back towards the expected Real-Time clearing price.

The main acknowledged danger with virtual bidding is the possibility for market manipulation. FERC's Office of Enforcement investigated Constellation Energy's Commodity Group for this type of manipulation in the NYISO and nearby markets during 2007 and 2008, ultimately negotiating a consent agreement—see FERC (2012) for more details. This danger has long been acknowledged in the literature—see Isemonger (2006) for example. Recently Ledgerwood and Pfeifenberger (2013) give an exposition of how this type of manipulation works, and Birge et al. (2013) find evidence that virtuals were used to manipulate prices in MISO. Because of the danger of manipulation, a number of ISOs have provisions restricting virtual trading when a trader has certain complementary positions that could benefit from manipulation, and/or have clawback provisions to prevent traders from profiting on these complementary positions due to virtual trades.

3 A Fault with Virtual Bidding

3.1 Overlooking Essential Complexity

This paper's focus is on an as yet unacknowledged problem with virtual bidding that arises from the inherent implementation differences between the Day-Ahead and Real-Time markets. The case for virtual bidding implicitly assumes that electricity is a simple commodity and that generators can be stacked simply according to costs. The source of any discrepency between the prices in the Day-Ahead and Real-Time markets is either a simple deficiency in demand or supply bid into the Day-Ahead market. Virtual bids are able to correct this by adding to net demand or net supply. Gone are the multiple attributes that differentiate specific generators along multiple dimensions besides the marginal cost of the next unit of energy. Gone are the important discrete choices, non-convexities and other complications that make the unit commitment and optimal power flow problems so difficult to solve in practice. Gone are the many approximations, decompositions and engineering judgments that play an important role in determining the outcome, and that are applied differently across the Day-Ahead and Real-Time markets.

Because the real problem is so much more complex than intersecting a pair of simple supply and demand curves, and because the Day-Ahead and Real-Time markets employ algorithms with different approximations, decompositions and judgments, a DA/RT spread can arise even when there is no simple deficiency of supply or demand bid into the Day-Ahead market. Since the problem is not caused by a simple deficiency of supply and demand, virtual bidding may not help to converge the prices. Worse still, virtual bidding may help converge the prices, but convergence may not correspond to improved system performance. In these cases, the profits on virtual bids can be a purely parasitic transfer from electricity consumers. Moreover, since the underlying problem is not a simple deficiency in supply or demand, virtual bids that add to simple supply or demand can have a harmful impact on the Day-Ahead unit commitment and dispatch, adding real costs to the system.

To flesh out these points, the next subsection relates the case of DA/RT spreads in the first few years of CAISO's operation and the effect of introducing virtual bidding. The subsequent subsection details a numerical illustration of non-economic virtual bidding designed to tie back to the CAISO experience. The profits to virtual bidding in the illustration are purely parasitic, and virtual bidding is likely to add real costs to the system. The last subsection generalizes from the CAISO case study and related numerical illustration.

3.2 Case Study of Virtual Bidding in California

CAISO, California's ISO, began operating its new multi-settlement markets in April 2009. From its very beginning, this new market design persistently produced two troubling price patterns. The first was infrequent but very severe price spikes in the Real-Time market. The second was Hour-Ahead prices persistently below both the earlier Day-ahead price and the later Real-Time price. Where most multi-settlement markets have two main stages, California's added a third, the Hour-Ahead market. The Hour-Ahead market balances the imports and exports, while the Real-Time market balances internal resources. These were not the only patterns suggesting problems, but they were the most significant and lasting. As we shall see, these two troubling patterns turned out to be related in some respects, but in many respects they reflected distinct problems.

3.2.1 Ramping Requirements and Real-Time Price Spikes

Table 3 shows the problem created by infrequent, but severe price spikes in the Real-Time market. Hours of operation have been divided into two categories: those in which the Real-Time system marginal energy cost, a component of the price, spikes above $250/MWh, and those when it did not. The table shows the percentage of hours belonging to each category, and the average DA/RT spread for each category, as well as the average DA/RT spread across all hours. For 2009, the price spiked in fewer than 1% of the hours. In those hours, the average DA/RT spread was -$349.06/MWh. In the remaining 99% of the hours, the Day-Ahead price was above the Real-Time price, with an average DA/RT spread of $1.18/MWh. Nevertheless, the few hours with a Real-Time price spike were so extreme that across all hours the average Real-Time price was above the average Day-Ahead price, with an average DA/RT spread of -$2.21/MWh.

Most of these spikes had a common origin that was quickly identified: short intervals, often less than 5-minutes, when the pattern of load required a very fast ramp up in generation that exceeded the capacity of the units scheduled to generate that hour by the Day-Ahead algorithm.[8] The Day-Ahead market is generally resolved only at an hourly level. Bids and offers are made for hourly quantities which cannot inform the system about intra-hour details such as the 5-minute ramp rate. The Day-Ahead algorithm solves for the least-cost hourly generation awards. While the algorithm solution respects generation unit operating constraints such as ramp rate limits and minimum run rates, these constraints are often defined on an hourly granularity.

The algorithm only deals with constraints at a finer granularity insofar as the operator imposes them. Doing so is difficult and presses against computational limits. Consequently, the solution to the Day-Ahead algorithm may not provide adequate fast ramp capacity when and where it is needed. The Real-Time market, however, operates on the finer granularity of 5-minute intervals and the fast ramp requirements impose themselves. If the units already operating do not have sufficient ramping capacity, the system operator is forced to resort to very expensive alternatives. Thus the very short intervals with very large price spikes. And,

[8]The first quarterly market monitor report noted that "Many of the price spikes occurring in the ISO's 5-minute dispatch market (RTG) are due to shortages of ramping energy " CAISO (2009a), p. 21, fn 9. See also the next two quarterly reports: CAISO (2009b), CAISO (2010a). The annual report reads that More frequently, real-time price spikes resulted from ramping limitations and other constraints. CAISO (2010b), p. 3.21.

Table 3: CAISO DA/RT Spreads in Hours With and Without a Real-Time Price Spike

		Spike	No Spike	All Hours
2009	count of hours	62	6,538	6,600
	% of hours	0.94%	99.06%	100.00%
	DA/RT spread	$ (349.06)	$ 1.18	$ (2.21)
2010	count of hours	87	8,673	8,760
	% of hours	0.99%	99.01%	100.00%
	DA/RT spread	$ (343.68)	$ (0.65)	$ (4.06)
2011	count of hours	65	8,695	8,760
	% of hours	0.74%	99.26%	100.00%
	DA/RT spread	$ (341.80)	$ 2.32	$ (0.23)
2012	count of hours	70	8,714	8,784
	% of hours	0.80%	99.20%	100.00%
	DA/RT spread	$ (365.55)	$ 2.11	$ (0.82)
2013	count of hours	47	8,713	8,760
	% of hours	0.54%	99.46%	100.00%
	DA/RT spread	$ (335.67)	$ 3.98	$ 2.16
2009-2013	count of hours	331	41,333	41,664
	% of hours	0.79%	99.21%	100.00%
	DA/RT spread	$ (347.81)	$ 1.82	$ (0.97)

Source: Data are from OASIS. Calculations are author's.

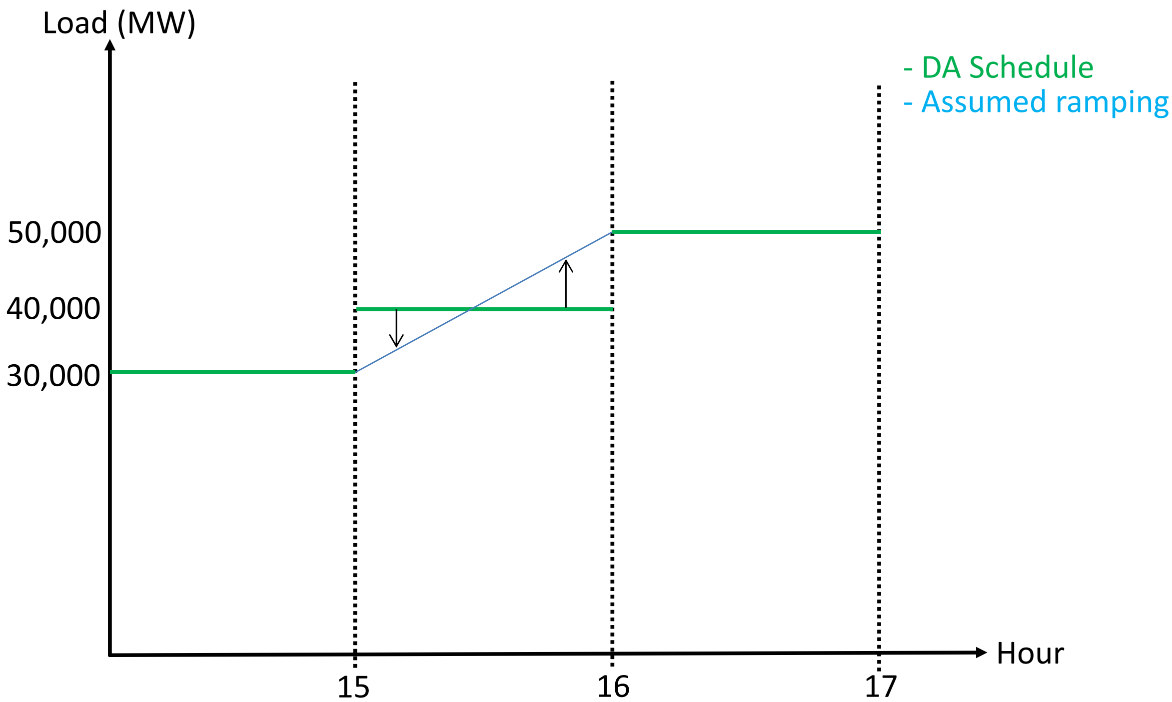

Figure 2: Illustration of the Hourly Granularity of the Day-Ahead Generation Schedule

since these spikes were not anticipated in the Day-Ahead dispatch, the average Day-Ahead price is less than the realized average Real-Time price.[9]

Figures 2 and 3 are a stylized illustration of the ramping assumption in the Day-Ahead market and the actual ramping in the Real-Time market. Figure 2 shows the Day-Ahead cleared generation increasing in hourly increments for hours ending 15, 16 and 17. Overlayed on top of these fixed hourly increments, the figure also shows a simple assumptions of linear ramping within hour 16. Figure 3 shows the actual Real-Time ramping within hour 16, and highlights a specific 5-minute interval when the ramp rate is especially severe. This is the interval when a shortage of ramping capacity is likely and Real-Time price spike occurs.

The problem can be ameliorated with better forecasting that anticipates the need for fast ramping, and with timely sourcing of cheaper ramping capacity.[10] Ideally, both actions

[9] While the focus here will be on the problem of upward ramping, the same incongruity between the hourly granularity of the Day-Ahead market and the 5-minute granularity of the Real-Time market can produce a downward ramping problem at other hours of the day. This has happened, too. However, the net impact on average prices was smaller in CAISO during these years. In order to keep the exposition simple, I have chosen to focus the narrative on the upward ramping problem.

[10] The two actions of forecasting and sourcing are sometimes not distinct from one another, because some ramping events arise when other system constraints are imposed, tracked and managed.

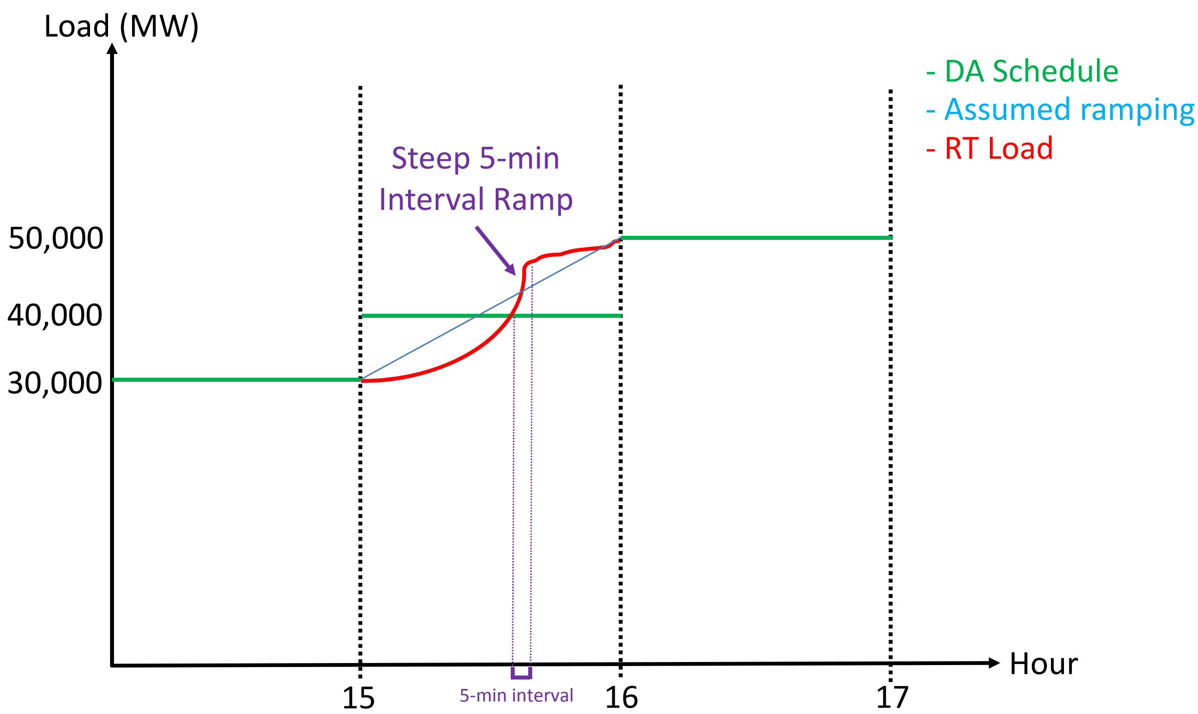

Figure 3: Illustration of an Intra-Hour Real-Time Load Ramp Exceeding the Day-Ahead Assumption

could be incorporated into the Day-Ahead algorithm, although doing so imposes complexity and may press against computational limits. More realistically, some actions taken after the Day-Ahead market but sufficiently before the Real-Time market may prove adequate, so long as there remains sufficient time and flexibility to identify lower cost ways of meeting the fast ramp requirement.

The difficulty and challenges in successfully resolving the problem are apparent in CAISO's operation over the next few years.

In 2009, the ISO quickly took several actions to reduce the frequency and severity of these Real-Time price spikes. In June, the ISO modified how certain transmission constraints were modeled. In August, the ISO adjusted how regulating reserve was managed. The ISO also made changes to how it exercised certain biases to manage load and transmission constraints. It is worthwhile quoting extensively on these changes because it gives the flavor of the many real world approximations and calibrations contained in the algorithms used to solve the optimal power flow problem, and shows how these must be tuned and re-tuned:

> *In Q3, the performance of the ISO's real-time market (RTM) for energy improved significantly as a result of a variety of steps taken toward the end of Q2 and beginning of Q3 that decreased the frequency and magnitude of price spikes not reflective of fundamental market conditions. Three of these changes that appear to have had very significant impacts include the following:*
>
> • *In early June, the pricing run of the RTD software was modified to allow transmission constraints to be exceeded by 5 MW instead of the previous threshold of .1 MW during the first 5-minute interval of the RTD optimization. This modification allows extra slack on a constraint that may not be fully resolved in a single 5-minute interval, but would otherwise have a significant impact on prices due to ramping constraints enforced in the RTD software.*
>
> • *Starting August 1, the RTD software was modified to represent how regulating reserve is used to balance short-term high-frequency load fluctuations. This modification allows limited relaxation of the power balance constraint through a lower scheduling run penalty price. These modifications would account for the effect of regulation ramping capability that will naturally be provided by resources providing regulation via Automated Generation Control (AGC).*
>
> • *In Q3, the ISO also implemented a tool that allowed phasing in bias (of load and transmission limits) across several market intervals rather than in one interval. This allows the market to adjust to new targets and limits more gradually (generally over a 15 minute period) and reduces the frequency of extreme prices and their impact on price convergence that otherwise would occur due to sudden*

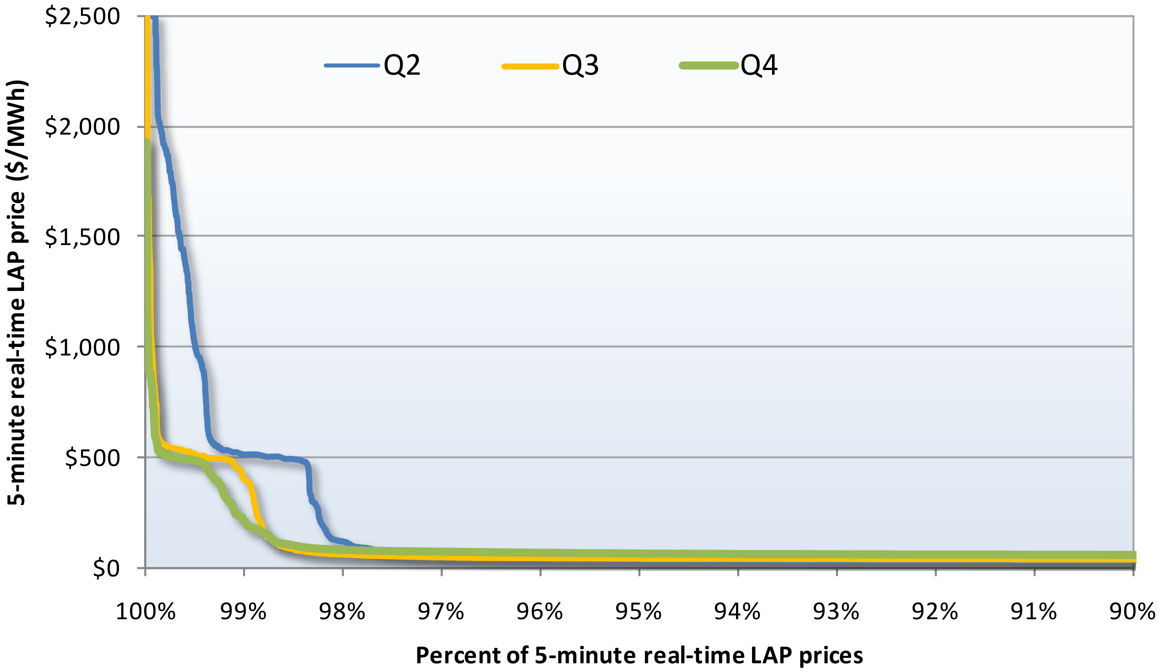

Figure 4: Top 10th Percentile of the Quarterly Real-Time Price Duration Curves for CAISO, 2009. Copy of Figure 3.14 from CAISO (2010), p. 3-20.

shocks.[11]

Finally, the ISO also made adjustments to exceptional dispatch procedures to address the problem.[12] These and similar measures seemed to quickly moderate the severity of the problem. Figure 4 shows the extreme tail of the Real-Time price duration curve for the first 3 quarters of the market's operation, which were Q2, Q3 and Q4 of 2009. There was a marked reduction in the number of the largest spikes from the first quarter of operation (Q2 of 2009) to the second and third quarters (Q3 and Q4 of 2009).

Despite this quick improvement, a troubling number of large price spikes persisted past the first few months of operation. Indeed, some of the improved performance in late 2009 may have been more a result of luck than of remedy. The second quarter of 2010 saw a recurrence of frequent, large spikes, despite an array of further actions taken by the ISO.[13] The numbers for 2010 in Table 3 reveal how the problem persisted. One contributing factor was the increase of the bid cap from $500/MWh to $750/MWh which increased the size of some spikes so that remedying the problem was a bit of a moving target. The cap was

[11] CAISO (2009b), pp. 7-8.
[12] CAISO (2010b), pp. 3.27-3.28.
[13] See CAISO (2011a).

raised again in April 2011, to $1,000/MWh. In 2011, the ISO took further action and the frequency of spikes decreased as the year progressed. In December 2011, the ISO began implementation of a flexible ramping constraint that had been in the works for more than a year. It continued to modify and fine tune this requirement over time, and had real success in assuring system-wide ramping capacity.[14]

From 2009 through 2011, the problem had been one of insufficient ramping capacity system-wide. Congestion had not been a major contributor to these spikes. However, that changed in 2012. More than half of the flexible ramping capacity produced was in the northern part of the ISO, and when congestion occurred, this ramping capacity was not always available where it was needed. Consequently, in 2012 the spikes were primarily driven by congestion and a shortage of ramping capacity in specific locations.[15] The ISO continued to increase the flexible ramping constraint in 2013 during peak hours, especially during the ramping hours. At the same time, the ISO also aligned transmission limits between the Day-Ahead and Real-Time markets, which reduced congestion.[16] The frequency of spikes fell again in 2013.

3.2.2 Low Hour-Ahead Prices

The second problem of low Hour-Ahead prices is manifest in Figure 5. Hour-Ahead prices were regularly lower than the Day-Ahead price that preceded them, as well as lower than the Real-Time price that followed. The Hour-Ahead market routinely decreased net imports from the level awarded in the Day-Ahead market, and then the Real-Time market often ended up having to re-source the power from expensive internal resources. Therefore, this anomaly had the effect of selling low and buying high, which obviously added to system costs.

[14] The CAISO market monitor's annual report for 2010 looked forward: The flexible ramping constraint will require that the software optimization results include a pre-specified amount of additional ramping capacity (beyond levels needed to simply meet the energy forecast). This new constraint is designed to ensure that sufficient upward and downward ramping capability from 5-minute dispatchable resources is committed and available to balance loads and supply on a 5-minute basis, taking into account the potential variability in actual system conditions. CAISO (2011b). The monitor then reported on the implementation in the annual report for 2012—see CAISO (2013), esp section 3.3. See also Abdul-Rahman et al. (2012). Further detailed information on the flexible ramping constraint implementation and related activities can be found on the CAISO website.

[15] Just over half of the flexible ramping capacity was in the northern part of the ISO system. When congestion occurs in the southern part of the system, this capacity can be stranded or unavailable for dispatch to help relieve congestion and meet system energy requirements in Southern California. CAISO (2013), p. 84. See also p.68.

[16] CAISO (2014), p. 76.

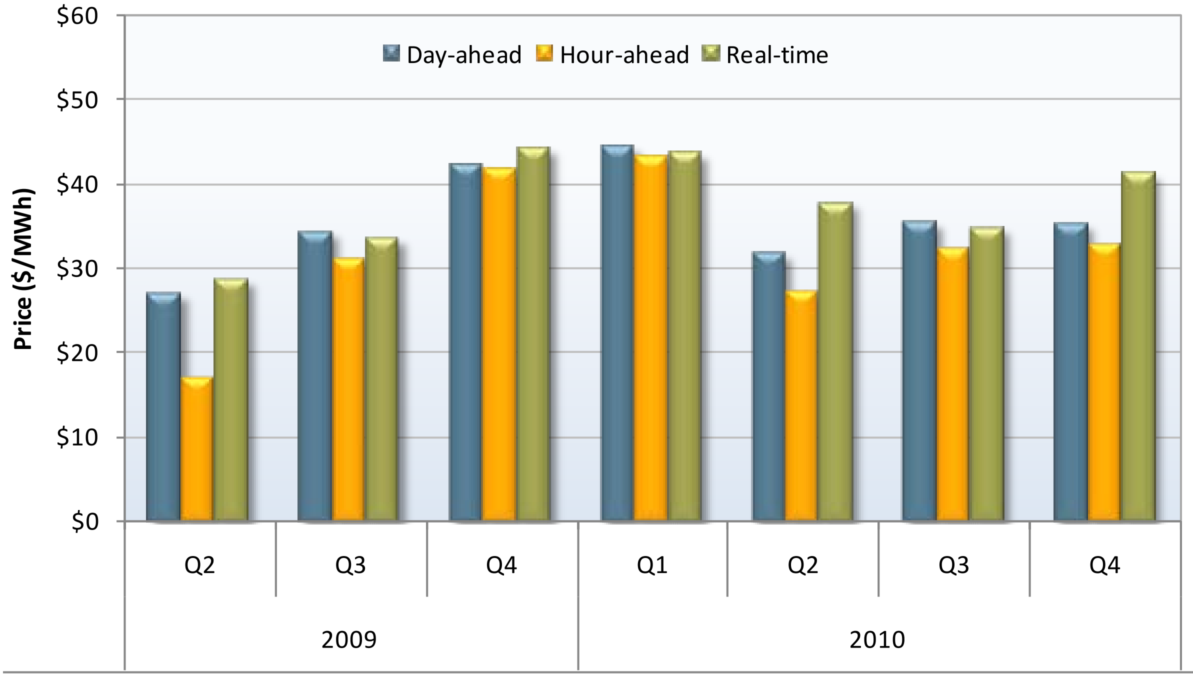

Figure 5: Quarterly Average Prices in CAISO's 3 Multi-Settlement Markets, 2009 & 2010. Prices are for the PG&E LAP, all hours. Copy of Figure 3.7 from CAISO (2011), p. 65.

It is not within the scope of this paper to address the problem of low Hour-Ahead prices in detail. However, since the history of virtual bidding against the Real-Time price spikes is so deeply intertwined with the history of virtual bidding against the low Hour-Ahead prices, it is necessary to present the basic outlines of the issue. The factors driving the Hour-Ahead anomaly are complicated, and resolving the problem has proven difficult. A contributing factor to the ramping problems driving the Real-Time price spikes was the different pattern of ramping assumed for imports in the Hour-Ahead market and the actual pattern that the Real-Time market had to accommodate. This is one of the many factors that was eventually fixed and that reduced the volume of Real-Time price spikes.[17] But it was not the only factor: after all, there is a difference between Real-Time prices being too high and Hour-Ahead prices being too low. As detailed in various Market Monitor reports and other CAISO documents, the disparity between Day-Ahead and Hour-Ahead prices reflected other factors, such as different modeling of power flow constraints in the Day-Ahead market and the Hour-Ahead market. Only as these other factors were diagnosed and resolved could the disparity between the Hour-Ahead and the Day-Ahead prices be diminished. As it happens, this disparity lasted into 2013, when modeling adjustments finally seemed to have eliminated

[17]CAISO (2009b), section 1.3.4, and CAISO (2011a) pp. 25-26.

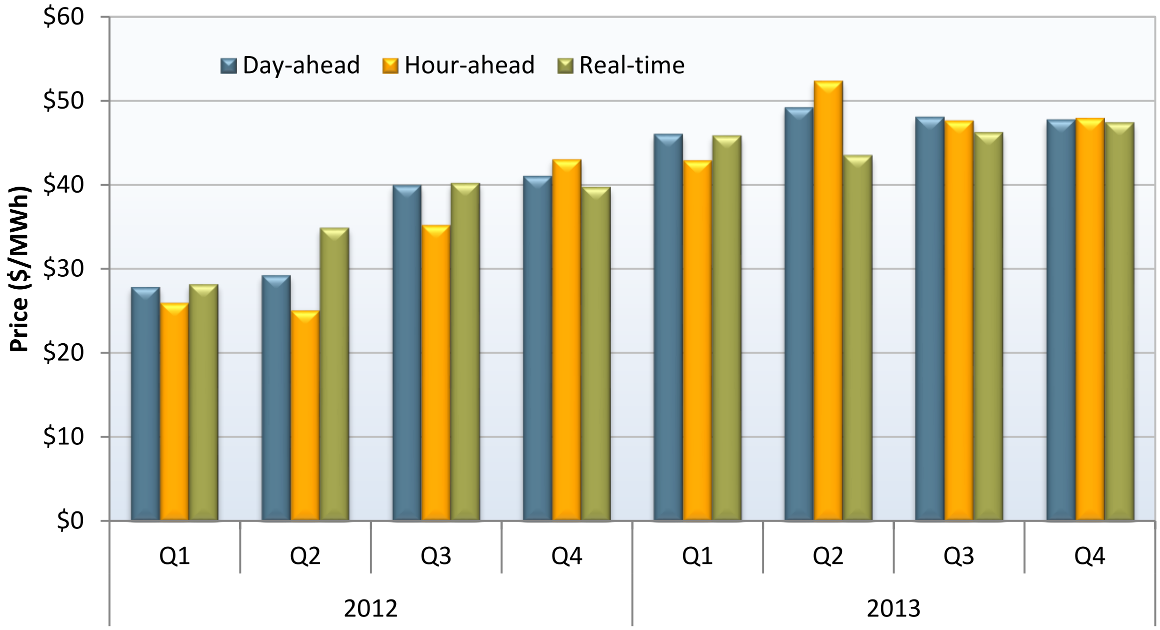

Figure 6: Quarterly Average Prices in CAISO's 3 Multi-Settlement Markets, 2012 & 2013. Prices are for System Energy, peak hours. Copy of Figure 2.9 from CAISO (2014), p. 73.

it, at least that year. Contrast Figure 6 with Figure 5 above.

3.2.3 The Role of Virtuals in CAISO

Virtual bidding was introduced to CAISO in February 2011. Both pricing problems had had already been diagnosed and publicly discussed, but neither had been fully resolved at that time. The problem of Real-Time price spikes could be exploited with a virtual demand bid placed at an internal location. The gross payoff to this virtual demand bid would be:

$$\pi_{VD,i} = (RT_i - DA_i) * Q. \tag{3}$$

The problem of low Hour-Ahead prices could be exploited with a virtual supply bid placed at an intertie. The gross payoff to this virtual supply bid would be:

$$\pi_{VS,i} = (DA_j - HA_j) * Q. \tag{4}$$

As can be seen in Figure 7, right from the start, the overwhelming volume of cleared virtual bids consisted in these two strategies, and the total cleared volume of these two types of bids increased through the first three months.

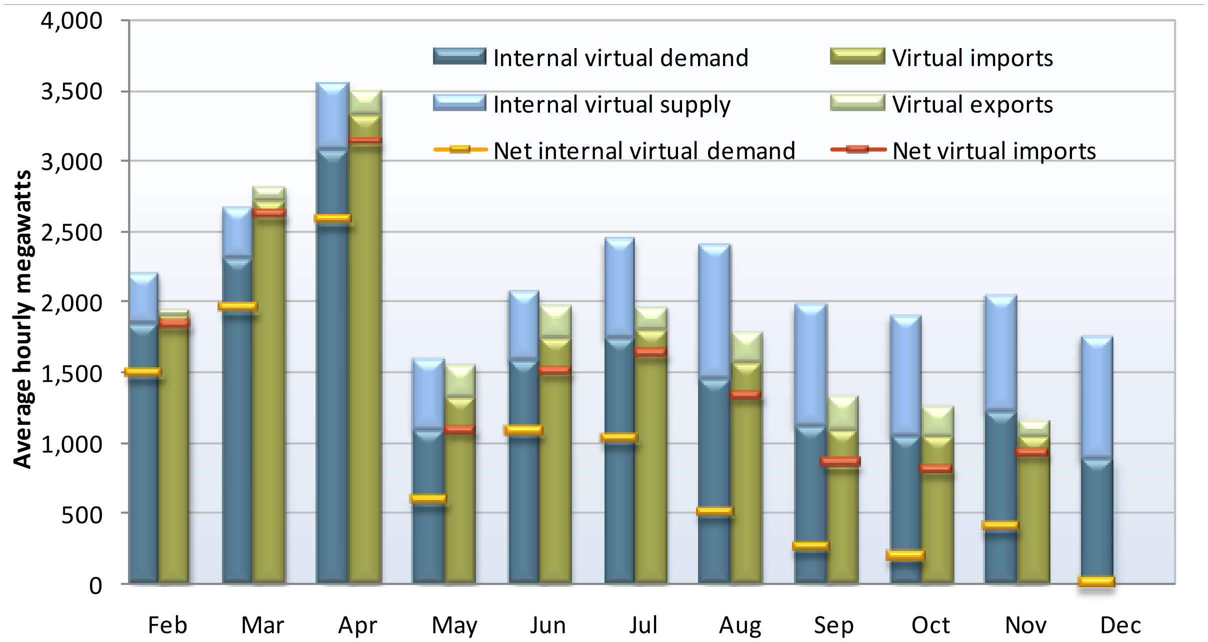

Figure 7: Average Monthly Cleared Virtual Bids at Inter-ties and Internal Locations. Copy of Figure 4.4 from CAISO (2012), p. 83.

The 1% of hours with Real-Time price spikes created by the ramping problem accounted for all of the profit made by virtual demand bids in 2011. During the other 99% of hours, virtual demand bids were highly unprofitable. However, the profit from the spikes more than compensated for the losses in the other hours. In December, when the frequency of price spikes decreased significantly, the profitability of virtual demand bids decreased to about zero.[18] In 2012, when Real-Time price spikes arose again due to congestion and insufficient ramping capacity in certain locations, again almost all profits from virtual demand bids resulted from these infrequent short intervals when there was insufficient ramping capacity.[19]

Many traders also combined the two strategies, creating matched pairs of virtual demand bid at an internal location and virtual supply bid at an intertie for the same hour. This produced the combined profit of both strategies with lower risk overall. The gross payoff to the matched pair would be:

$$\pi_{VD,i+VS,j} = (RT_i - DA_i) * Q + (DA_j - HA_j) * Q = (RT_i - HA_j) * Q + (DA_j - DA_i) * Q. \quad (5)$$

[18] CAISO (2012), p. 87.
[19] CAISO (2013), p. 110.

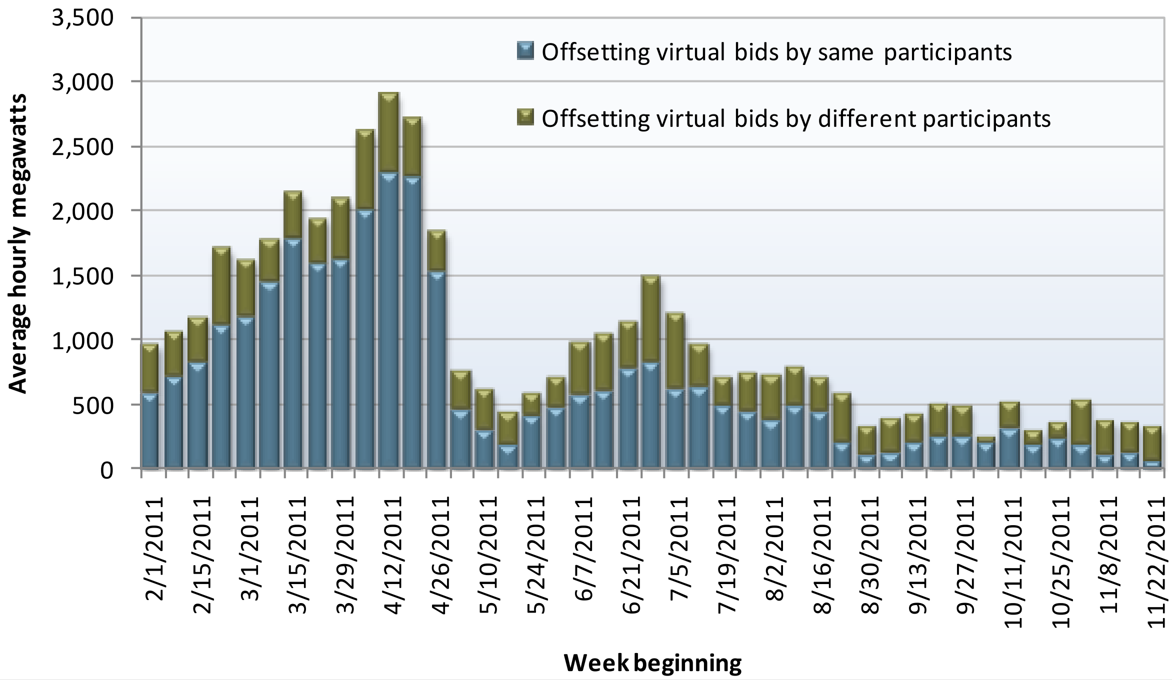

Figure 8: Offsetting Virtual Demand at an Internal Location and Virtual Supply at an Intertie. Copy of Figure 4.5 from CAISO (2012), p. 84.

Although the locations for the two virtual bids are different, so that the two Day-Ahead prices might be different, in practice this is not a problem. Locations can be chosen so that the two Day-Ahead prices are almost certain to be the same. Therefore, the gross payoff to the matched pair strategy is more succinctly written:

$$\pi_{VD,i+VS,j} \cong (RT_i - HA_j) * Q. \tag{6}$$

When the Real-Time price spiked, it would add profit to this strategy, and if the Hour-Ahead price were low, it would add profit, too. Risk was reduced because the uncertain relationship between the Day-Ahead and the two balancing prices was removed. Any updates to the load forecast, to generation plant outages or transmission constraints after the Day-Ahead market would enter both prices.

Offsetting positions by the same trader were a significant and increasing fraction of the cleared virtual trades through the first months of trading, as shown in Figure 8.

Virtual bidding could not remedy the underlying ramping problem. This is most transparent in this matched pair strategy. In the Day-Ahead market, the matching virtual demand and virtual supply bids effectively cancel one another out, producing a net zero impact on total system energy required. There is no change in the total amount of energy scheduled Day-Ahead. More importantly, no additional ramping capacity is put on to ameliorate the Real-Time price spikes, and no change is made to the Hour-Ahead scheduling algorithm so that the Hour-Ahead price is still too low and net imports are still reduced.

Nevertheless, the matched pair strategy generated significant profits to the virtual bidders which raised the cost paid by load.

The transparent uselessness of this virtual bidding strategy, and its obvious cost, immediately caught the attention of the ISO and many stakeholders. In April of 2011, the ISO initiated a stakeholder process to address the volume of offsetting virtual demand and import bids. At the same time, the actions being taken by the ISO to address the underlying causes of the price discrepancies—detailed above—were reducing the profitability of these strategies. According to the Market Monitor, these two factors may help explain the sharp reduction in volume of the offsetting cleared virtual bids seen in Figure 7, above.[20] Nevertheless, some use of this matched pair strategy continued, as did unmatched virtual supply offered at intertie locations. In November, the ISO suspended virtual bidding at the interties.

While it is straightforward to see that the matched pair virtual bidding strategy is not a remedy to the underlying scheduling issues behind the two pricing problems, the same is true about each of the two parts of the strategy taken individually. Virtual demand bids could not remedy the fast ramp issue leading to Real-Time price spikes, and virtual supply offers could not remedy the scheduling mistakes leading to low Hour-Ahead prices.

Cleared virtual demand does add to generation awards made in the Day-Ahead market, and also raises the Day-Ahead clearing price. Consequently, convergence is improved since the average DA/RT spread becomes less negative. However, this improved convergence does not correspond to improved operation and lower costs. The added generation awards need not materially increase the system's fast ramp capability—in fact, because of the complex interaction between the level of generation and the capacity to ramp up or down, added generation awards can paradoxically decrease ramp capability. These added Day-Ahead awards are for an hourly lump of power, and are not specifically for extra power in the short 5-minute intervals when it is needed. Therefore, the Real-Time price still spikes just as often and the size of the spike is just as high. In fact, since extra generation units receive Day-Ahead awards, but this extra generation is not really useful, virtual demand actually adds to total system costs.

[20]CAISO (2012), p. 84.

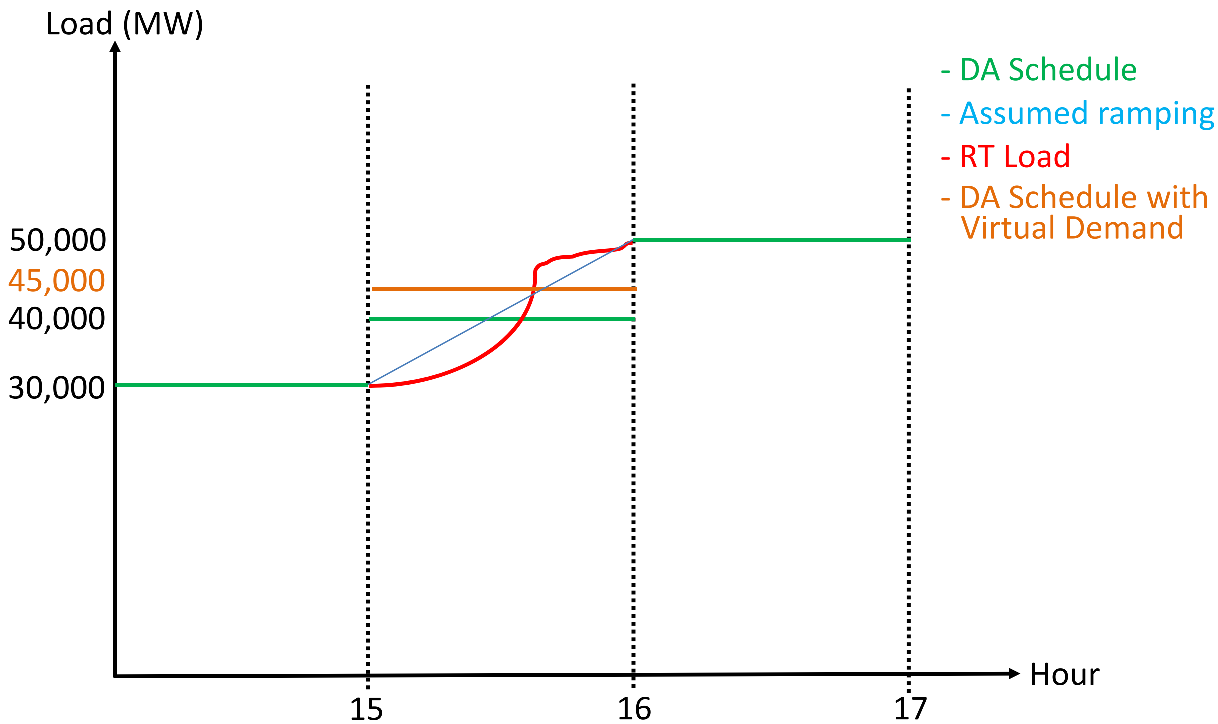

Figure 9: Illustration of the Impotency of Virtual Demand to Resolve Intra-Hour Real-Time Load Ramp Problem

Figure 9, extends the illustration begun in Figures 2 and 3, above, showing the impact of the extra generation award produced by virtuals and its irrelevance to the fast ramp problem.

Another way to appreciate the problem is to think of what virtual bidding does in the two different categories of hours—when there turns out to be a shortage of fast ramp capacity and therefore a Real-Time price spike, and when there does not. In the former case, the extra generation capacity is the wrong kind. The Day-Ahead price is increased, but without ameliorating the problem. The extra generation awards Day-Ahead may add to total system costs depending upon the broader picture of start-up and unit commitments across multiple hours. In the latter case, the average Day-Ahead price is already above the average Real-Time price, so that what is really wanted is fewer generation awards Day-Ahead, not more. So, in 99% of the hours, the virtual demand bid is producing exactly the wrong result, while in 1% of the hours it is not helping.

In cases like this, the increased convergence in the average DA/RT spread, therefore, gives a false impression that the system's operation has been improved.

Cleared virtual supply at the interties also fails as a remedy for the low Hour-Ahead price.

The problem is not one of too much generation scheduled Day-Ahead, which is what cleared virtual supply is targeted to removing. The problem is that the Hour-Ahead algorithm sells some of this generation, forcing the system to buy it back in the Real-Time market. At best, cleared virtual supply moves this sale to an earlier point in time—the Day-Ahead market instead of the Real-Time market—but leaves the system still needing to source power at a higher price in the Real-Time market.

The ISO has documented at various points how these various virtual bidding strategies added costs.[21]

3.3 A Numerical Illustration

3.3.1 Introduction

This is a numerical example to illustrate the mechanics of how virtual bidding interacts with the fast ramping problem as experienced in CAISO. The example is carefully constructed to be as simple as possible, while successfully demonstrating the key mechanics.

The example details the operation of the market over a single hour, broken down into twelve 5-minute intervals. It presents two alternative scenarios for the rate at which load ramps up over the hour. In Scenario A, load ramps up at a constant rate through the hour, while in Scenario B all of the ramp is concentrated into two 5-minute intervals in the middle of the hour.

The example assumes there are two types of generation units. One is a fleet of NGCC units which are the cheapest option for energy, but which have a limited capability for ramping. The other is a fleet of fast ramp units, which are very expensive to operate but which can serve the highest ramp rate demanded by the system in any interval. To keep the example simple, there are no unit start-up costs, no minimum start times, and no minimum operating levels for individual units. The only constraint is the aggregate rate of ramping the fleet of NGCC units. Therefore, the example is presented in terms of the total power produced by either type of generation, since there is no need to keep track of individual plants. The marginal cost of generation is a function of the aggregate energy produced by either type of generation in any 5-minute interval.

The optimal dispatch is straightforward: NGCC units are used up to their limit for ramping, and the fast ramp units fill in where necessary. The fast ramp units are never required in Scenario A, but are briefly required in Scenario B.

[21] CAISO (2012), p. 6, and CAISO (2013), pp. 8-9.

The example calculates Real-Time and Day-Ahead market outcomes, but without any explicit modeling of bids and offers into an auction. Instead, the price is set to match demand and supply, where demand is always a given level of load, and the supply curve is defined by the marginal cost of energy from the generation fleet. When the NGCC units are able to satisfy demand, it is the marginal cost of NGCC generation that sets the price, and when the fast ramp units are required, it is the marginal cost of fast ramp units that sets the price. The outcome in the Real-Time market depends upon the load scenario. For each scenario, the example shows a price for each 5-minute interval as well as the hourly Real-Time price calculated as the average of the 5-minute interval prices. Matching the CAISO experience, the price spikes very high in Scenario B, but only for a short period. Nevertheless, this short spike is enough to dramatically change the Real-Time price.

In anticipation of the Real-Time market, but with uncertainty about which Scenario will be realized, there is a Day-Ahead market. To match the feature of CAISO's price spike problem the example assumes that Scenario B is very rare, only occurring 1% of the time. Scenario A is the usual case, occurring 99% of the time. For each scenario, the example then calculates all of the payments made by load and to each type of generation, including Day-Ahead payments as well as Real-Time balancing payments and any uplift charges necessary to net the total system to zero. The example also calculates the expected payments, averaging across the two scenarios.

We solve the example for three different cases reflecting alternative assumptions about virtual bidding. First, in Case #1, we assume no virtual bidding. The Day-Ahead market is cleared by matching physical demand and physical supply. The result is that NGCC units receive generation awards equal to the expected load. The Day-Ahead price is not equal to the expected Real-Time price: there is a negative expected DA/RT spread. The next two cases allow virtual bidding, which seeks to profit off the DA/RT spread. With virtual bidding cleared physical generation need not exactly match expected physical load. The cleared net virtual bids will be the difference. In Case #2 there is enough volume of cleared virtual bids so as to narrow the negative expected DA/RT spread, while still leaving enough of a spread so that the virtual bids are profitable. Case #3 is the limiting case in which there is so much volume in virtual bids that the expected DA/RT spread is driven to zero, which means that the virtual bids have zero expected profit.

Virtual bidding has consequences, but it does NOT lower system costs, nor improve system performance in any other way. The example shows that it is possible to get improved convergence, as measured by the average DA/RT spread, without getting any improved performance. Profits paid to virtual bidders are a pure parasitic drain on the system: they are made at the expense of load. Virtual bidding also has the effect of increasing total payments from load to generation, without there being any additional generation service provided.

Although the example does not incorporate many of the frictions that make the Day-Ahead market useful—such as start-up costs, minimum start times, and minimum operating levels for individual units—the set-up makes clear that virtual bidding might actually make things even worse if these frictions were incorporated into the problem.

Having sketched the basic outline, we can now turn to presenting the numerical details of the example. First, we specify the features of generation, and then the two scenarios for load. We then describe the optimal dispatch in each scenario. Next, we detail the Real-Time market prices. Then we describe the three cases for the Day-Ahead market, with and without virtual bidding. Finally, we show the calculation of the payments from both the Day-Ahead and the Real-Time markets contingent on the two scenarios. We do this for each of the three cases. We use the comparison of payment results across the three cases to explain the effect of virtual bidding.

3.3.2 Supply—the Cost of Generation

The market has a set of NGCC power plants which can supply electricity according to the following marginal cost function:

$$C_N(x) = A_1 + A_2 * x + A_3 * x^2, \tag{7}$$

where x is the rate of generation in MW/h, A_1, A_2, and A_3 are given parameters, and C_N is the rate of marginal cash flow in \$/h. Table 4 shows the three parameter values used in the example, as well as the values at the three points along the supply curve.

In our illustration, there are no start-up costs and no minimum run rates, nor any other discrete constraints except for an aggregate ramp rate constraint. We assume that this suite of NGCC plants has a constant maximum rate of ramping capability equal to 30,000 MW/h. This translates to a ramp of 2,500 MW within a single 5-minute interval. In our illustration we will only have a need for upward ramp capacity. There is no change in this cost curve between the Day-Ahead market and the Real-Time market.

In addition to the suite of low cost NGCC power plants, the system has a set of expensive fast ramp plants. These have a simple, constant marginal cost function:

Table 4: Parameters of the Supply Curve of NGCC Generation

Function Parameters			Function Value at Sampled Output Levels	
			Output	Marginal Cost
A1	31.00		x	$C_N(x)$
A2	0.000100		MW/h	\$/MWh
A3	6E-09		40,000	44.60
			45,000	47.65
			50,000	51.00

$$C_F(x) = B, \tag{8}$$

where x is the rate of generation in MW/h, B is a given parameter, and C_F is the rate of cash flow in \$/h. We set $B = 2000$, which makes the fast ramp plants much more expensive than the NGCC plants at any range of generation relevant in this illustration. These units are only used when the NGCC plants are constrained in their ramp rate. For these units, too, there are no start-up costs and no minimum run rates. They are able to ramp at any rate, both up and down.

3.3.3 Demand Structure of Load

The single hour's load profile is broken down into twelve 5-minute intervals. Total load through the hour is 45,000 MW. Load immediately before the start of the hour is 40,000 MW, and load immediately after the end of the hour is 50,000 MW.

Scenario A Even, Slow Ramp

In scenario A, load ramps evenly throughout the hour as shown here in Table 5:

The 5-minute ramp in this scenario is within the capacity of the suite of NGCC plants, so the fast ramp units will not be needed.

Scenario B Sudden, Fast Ramp

In scenario B, load ramps suddenly in two 5-minute intervals at the middle of the hour as shown here in Table 6:

The 5-minute ramp in this scenario exceeds the capacity of the suite of NGCC plants, so the fast ramp units will be needed. They will ramp up through the two intervals when load is ramping up. Even once load stops ramping up, the NGCC plants will take two more 5-minute intervals to catch up, so the fast ramp units will operate during those two intervals,

Table 5: Scenario A: Real-Time Load in 5-minute Intervals

Interval	Start	Ramp	End	Avg
1	40,000	833	40,833	40,417
2	40,833	833	41,667	41,250
3	41,667	833	42,500	42,083
4	42,500	833	43,333	42,917
5	43,333	833	44,167	43,750
6	44,167	833	45,000	44,583
7	45,000	833	45,833	45,417
8	45,833	833	46,667	46,250
9	46,667	833	47,500	47,083
10	47,500	833	48,333	47,917
11	48,333	833	49,167	48,750
12	49,167	833	50,000	49,583
Average				45,000

Table 6: Scenario B: Real-Time Load in 5-minute Intervals

Interval	Start	Ramp	End	Avg
1	40,000	0	40,000	40,000
2	40,000	0	40,000	40,000
3	40,000	0	40,000	40,000
4	40,000	0	40,000	40,000
5	40,000	0	40,000	40,000
6	40,000	5,000	45,000	42,500
7	45,000	5,000	50,000	47,500
8	50,000	0	50,000	50,000
9	50,000	0	50,000	50,000
10	50,000	0	50,000	50,000
11	50,000	0	50,000	50,000
12	50,000	0	50,000	50,000
Average				45,000

ramping down from their peak.

3.3.4 The Real-Time Market

Table 7 shows the outcomes of the Real-Time market in each of the two scenarios. For each 5-minute interval, it shows the average load over the interval as well as the Real-Time dispatch instructions for the two types of generation.[22] It also shows the price in each 5-minute interval. When the load in an interval can be supplied exclusively by the NGCC power plants, the price for the interval is calculated from the NGCC supply curve:

$$P_i = C_N(L_i), \tag{9}$$

where i indexes the 5-minute intervals, with $i = 1, 2, 12$, and L_i is the average load during interval i. When the load in an interval requires generations by the fast ramp plants, the price for the interval is calculated from the fast ramp supply curve:

$$P_i = C_F(L_i) = B, \tag{10}$$

The hour's Real-Time price is the average interval price across all twelve 5-minute intervals:

$$R_T = \sum_{i=1}^{12} P_i/12, \tag{11}$$

<u>**Scenario A**</u>

In scenario A, the Real-Time price is:

$$R_T = \$47.70/MWh. \tag{12}$$

<u>**Scenario B**</u>

In scenario B, the Real-Time price is:

$$R_T = \$698.00/MWh. \tag{13}$$

[22] The table shows average load and generation within each interval. Across any two successive intervals with generation ramping at the same rate, the difference between the average generation will equal the ramp rate. However, across any two successive intervals with generation ramping at different rates, the difference between the average generation will be less than the ramp rate. Therefore, in the first interval when the NGCC power plants ramp at a rate of 2,500MW, the average generation is only incremented by 1,250. In the next interval the average generation is incremented by 2,500

Table 7: Real-Time Dispatch and Market Prices for the 2 Scenarios

	Scenario A				Scenario B			
Interval	Avg Load	NGCC Gen	Fast Ramp Gen	Price	Avg Load			Price
1	40,417	40,417	0	44.84	40,000	40,000	0	44.60
2	41,250	41,250	0	45.33	40,000	40,000	0	44.60
3	42,083	42,083	0	45.83	40,000	40,000	0	44.60
4	42,917	42,917	0	46.34	40,000	40,000	0	44.60
5	43,750	43,750	0	46.86	40,000	40,000	0	44.60
6	44,583	44,583	0	47.38	42,500	41,250	1,250	2,000.00
7	45,417	45,417	0	47.92	47,500	43,750	3,750	2,000.00
8	46,250	46,250	0	48.46	50,000	46,250	3,750	2,000.00
9	47,083	47,083	0	49.01	50,000	48,750	1,250	2,000.00
10	47,917	47,917	0	49.57	50,000	50,000	0	51.00
11	48,750	48,750	0	50.13	50,000	50,000	0	51.00
12	49,583	49,583	0	50.71	50,000	50,000	0	51.00
Average	45,000	45,000	0	47.70	45,000	44,167	833	698.00
Probability of Each Scenario				99%				1%
Expected Hourly Real-Time Price:					54.20			

Table 8: Day-Ahead Dispatch and Market Prices for Alternative Cases of Virtual Bidding

	DA Price	Cleared Bids and Offers: Load	NGCC Gen	Net Virtual Demand
Case #1: No Virtual Bidding	47.65	45,000	45,000	0
Case #2: Virtual Bidding, standard case	50.93	45,000	49,895	4,895
Case #3: Virtual Bidding, limiting case	54.20	45,000	54,409	9,409

This is 14 times the Real-Time price of scenario A. The higher cost is due to the need to turn to the very expensive fast ramp units. When scenario B is realized we say that a price spike occurs. Table 7 also shows the expected value of the Real-Time price, which is \$54.20/MWh due to the small probability of the very high price spike associated with the fast ramp required in Scenario B.

3.3.5 The Day-Ahead Market

The Day-Ahead market operates before it is known which scenario for load obtains, and must be based on the expectations. Moreover, the Day-Ahead market accepts bids and offers for the entire hour, and makes generation awards for the entire hour. The Day-Ahead market does not operate at the detail of the 5-minute intervals within the hour. The illustration assumes that load always bids its expected value for the hour, which is 45,000 MWh. It assumes that generation bids its marginal cost, which gives the supply curve described above. The Day-Ahead price clears the market by equating demand with supply. The outcome varies depending upon assumptions about virtual bidding. Table 8 shows the Day-Ahead market outcome under three different assumptions.

Case #1. No virtual bidding. The cleared physical load is 45,000 MWh, which means the cleared physical generation must be 45,000 MWh. The cheapest source for that is the NGCC units, and the supply curve for the NGCC units requires a price of \$47.65 to produce that amount. Note that in the absence of virtual bidding, the Day-Ahead price is less than the expected Real-Time price. The DA/RT spread is -\$6.55/MWh. In 99% of the hours, the Day-Ahead price is greater than the Real-Time price, but the 1% of the hours when the Real-Time price spikes very high creates the negative DA/RT spread across all hours.

The negative average DA/RT spread is enticing to virtual bidders, representing an expected profit of \$6.55/MWh for a cleared virtual demand bid. Virtual demand will make losses in 99% of the hours, but the large profit in the 1% will more than compensate. Of course, if virtual traders aggressively bid virtual demand, that will bid up the Day-Ahead price. Table A.5 shows two cases with virtual bidding. Case #2. Virtual bidding-the standard case. This case assumes that the Day-Ahead price is bid up half-way towards the expected Real-Time price, which leaves a smaller negative expected DA/RT spread of -$3.28/MWh. The cleared volume of net virtual demand that produces this move is 4,895MWh. This adds to the total cleared load in the Day-Ahead market, increasing the Day-Ahead generation awards to 49,895MWh. Case #3. Virtual bidding-the limiting case. This case assumes that the Day-Ahead price is bid up all the way to exactly equal the expected Real-Time price. The expected DA/RT spread is $0.00/MWh. The cleared volume of net virtual demand that produces this move is 9,409MWh. This adds to the total cleared load in the Day-Ahead market, increasing the Day-Ahead generation awards to 54,409MWh. In both cases, virtual bidding increases the amount of physical generation awards in the Day-Ahead market. However, in neither case will any of this extra generation be used. In Scenario A, the extra generation is simply not needed. In Scenario B, what is needed is fast ramp generation, but this is not what the virtual bidding elicits. Virtual bids have no way to signal the need for generation in any specific 5-minute interval: these bids are for total energy across the full hour. Consequently, the virtual bids do not improve the performance of the system. Therefore, any profit captured by virtual bidders is purely parasitic. It is a drain on the system that must ultimately be paid for by load. We can see this more clearly in the next section where we detail the payments made from load to generation and to virtual bidders.

Before turning to the payments, one other final observation about the Day-Ahead market is in order. Virtual bidding has improved convergence as measured by the average DA/RT spread. However, there has been no improvement in performance. This highlights the fact that average DA/RT spreads are a preliminary diagnostic tool with which to assess performance, and this single diagnostic needs to be blended together with other detailed information before a reliable assessment can be made about performance. One should be cautious about equating convergence with performance. One should be equally, if not more cautious about equating the profitability of virtuals with improved performance.

3.3.6 Payments from Load to Generation and Virtuals

Tables 9, 10 and 11, show the payments from virtuals to generation and virtuals for each of the three cases, respectively.

Case #1, with no virtuals, is shown in Table 9.

Table 9: Case #1, No Virtuals: Payments from Load to Generators

Scenario A

	to NGCC Generators			to Fast Ramp Generators			to Virtuals	from Load			
interval	DA	RT	Total	DA	RT	Total		DA	RT	uplift	Total
1	2,144,250	-205,529	1,938,721	0	0	0	0	2,144,250	-205,529	0	1,938,721
2	2,144,250	-170,004	1,974,246	0	0	0	0	2,144,250	-170,004	0	1,974,246
3	2,144,250	-133,684	2,010,566	0	0	0	0	2,144,250	-133,684	0	2,010,566
4	2,144,250	-96,547	2,047,703	0	0	0	0	2,144,250	-96,547	0	2,047,703
5	2,144,250	-58,574	2,085,676	0	0	0	0	2,144,250	-58,574	0	2,085,676
6	2,144,250	-19,743	2,124,507	0	0	0	0	2,144,250	-19,743	0	2,124,507
7	2,144,250	19,966	2,164,216	0	0	0	0	2,144,250	19,966	0	2,164,216
8	2,144,250	60,574	2,204,824	0	0	0	0	2,144,250	60,574	0	2,204,824
9	2,144,250	102,103	2,246,353	0	0	0	0	2,144,250	102,103	0	2,246,353
10	2,144,250	144,572	2,288,822	0	0	0	0	2,144,250	144,572	0	2,288,822
11	2,144,250	188,004	2,332,254	0	0	0	0	2,144,250	188,004	0	2,332,254
12	2,144,250	232,418	2,376,668	0	0	0	0	2,144,250	232,418	0	2,376,668
Average	2,144,250	5,296	2,149,546	0	0	0	0	2,144,250	5,296	0	2,149,546

Scenario B

	to NGCC Generators			to Fast Ramp Generators			to Virtuals	from Load			
interval	DA	RT	Total	DA	RT	Total		DA	RT	uplift	Total
1	2,144,250	-223,000	1,921,250	0	0	0	0	2,144,250	-223,000	0	1,921,250
2	2,144,250	-223,000	1,921,250	0	0	0	0	2,144,250	-223,000	0	1,921,250
3	2,144,250	-223,000	1,921,250	0	0	0	0	2,144,250	-223,000	0	1,921,250
4	2,144,250	-223,000	1,921,250	0	0	0	0	2,144,250	-223,000	0	1,921,250
5	2,144,250	-223,000	1,921,250	0	0	0	0	2,144,250	-223,000	0	1,921,250
6	2,144,250	-7,500,000	-5,355,750	0	2,500,000	2,500,000	0	2,144,250	-5,000,000	0	-2,855,750
7	2,144,250	-2,500,000	-355,750	0	7,500,000	7,500,000	0	2,144,250	5,000,000	0	7,144,250
8	2,144,250	2,500,000	4,644,250	0	7,500,000	7,500,000	0	2,144,250	10,000,000	0	12,144,250
9	2,144,250	7,500,000	9,644,250	0	2,500,000	2,500,000	0	2,144,250	10,000,000	0	12,144,250
10	2,144,250	255,000	2,399,250	0	0	0	0	2,144,250	255,000	0	2,399,250
11	2,144,250	255,000	2,399,250	0	0	0	0	2,144,250	255,000	0	2,399,250
12	2,144,250	255,000	2,399,250	0	0	0	0	2,144,250	255,000	0	2,399,250
Average	2,144,250	-29,167	2,115,083	0	1,666,667	1,666,667	0	2,144,250	1,637,500	0	3,781,750
Expected	2,144,250	4,952	2,149,202	0	16,667	16,667	0	2,144,250	21,618	0	2,165,868

The top panel of the table shows the payments made in Scenario A, while the bottom panel shows the payments made in Scenario B. At the bottom, the Table shows the expected payments across the two scenarios. In the top panel, payments are calculated for each of the twelve 5-minute intervals, and then the total hourly payment is shown at the bottom of the panel, in the row labeled 'Average.' The first set of columns show the aggregate payments to the NGCC generators. These are broken down into the payments for the generation awards received in the Day-Ahead market and the payments for the balancing amounts awarded in the Real-Time market. The next set of columns shows the payments to the fast ramp generators. The next column is the payments to the virtual bidders, which in this case are zero by definition. The final set of columns are the payment from load.

In Scenario A, payments to the NGCC units for the Day-Ahead award are constant through the hour, by definition. Early in the hour, as load is ramping slowly up, the total generation needed is smaller than the Day-Ahead award, so the NGCC units buy-back power at the clearing price in the Real-Time market. Therefore the payments are negative. Later in the hour, as load continues to ramp and is above the Day-Ahead award, the NGCC units sell the additional power at the clearing price in the Real-Time market. Therefore the payments are positive. Fast ramp generators. Since they are not needed, the payments are zero. Payments from load are the mirror image of payments to the NGCC generators.

In Scenario B, the fast ramp generators are needed during the middle of the hour, setting a very high Real-Time price. All of the power from the fast ramp generators is purchased in the Real-Time market. During the first couple of these intervals, power purchased from the fast ramp generators displaces Day-Ahead awards made to NGCC generators, so the NGCC units buy-back power at very high prices. In the later couple of these intervals, the NGCC generators are selling extra power in the Real-Time market at very high prices. Payments from load are the mirror image of the total payments made to the NGCC and the fast ramp generators.

Case #2, with virtuals, is shown in Table 10.

The NGCC generators receive larger payments than in Case #1 since the Day-Ahead price was bid-up higher. However, some of these payments are reversed in the Real-Time market. This happens because a significant portion of the Day-Ahead awards are bought back in the Real-Time market. The Day-Ahead market cleared with net virtual demand, which increased the volume of Day-Ahead awards made to physical generation in excess of the volume of load that actually arrives in either Scenario A or B. Virtual demand earns losses in Scenario A, but earns profits in Scenario B. Overall, virtual demand earns an expected

Table 10: Case #2, Virtual bidding-the standard case: Payments from Load to Generators and Virtuals

Scenario A

	to NGCC Generators			to Fast Ramp Generators			to Virtuals	from Load			
interval	DA	RT	Total	DA	RT	Total		DA	RT	uplift	Total
1	2,540,952	-425,019	2,115,933	0	0	0	-15,793	2,291,685	-205,529	13,984	2,100,139
2	2,540,952	-391,900	2,149,052	0	0	0	-15,793	2,291,685	-170,004	11,577	2,133,258
3	2,540,952	-358,027	2,182,925	0	0	0	-15,793	2,291,685	-133,684	9,130	2,167,131
4	2,540,952	-323,379	2,217,573	0	0	0	-15,793	2,291,685	-96,547	6,642	2,201,779
5	2,540,952	-287,935	2,253,017	0	0	0	-15,793	2,291,685	-58,574	4,113	2,237,223
6	2,540,952	-251,674	2,289,278	0	0	0	-15,793	2,291,685	-19,743	1,543	2,273,484
7	2,540,952	-214,575	2,326,377	0	0	0	-15,793	2,291,685	19,966	-1,067	2,310,583
8	2,540,952	-176,618	2,364,334	0	0	0	-15,793	2,291,685	60,574	-3,719	2,348,540
9	2,540,952	-137,781	2,403,170	0	0	0	-15,793	2,291,685	102,103	-6,411	2,387,377
10	2,540,952	-98,045	2,442,907	0	0	0	-15,793	2,291,685	144,572	-9,143	2,427,114
11	2,540,952	-57,387	2,483,565	0	0	0	-15,793	2,291,685	188,004	-11,917	2,467,772
12	2,540,952	-15,787	2,525,165	0	0	0	-15,793	2,291,685	232,418	-14,732	2,509,371
Average	2,540,952	-228,177	2,312,775	0	0	0	-15,793	2,291,685	5,296	0	2,296,981

Scenario B

	to NGCC Generators			to Fast Ramp Generators			to Virtuals	from Load			
interval	DA	RT	Total	DA	RT	Total		DA	RT	uplift	Total
1	2,540,952	-441,302	2,099,650	0	0	0	3,167,205	2,291,685	-223,000	3,198,170	5,266,855
2	2,540,952	-441,302	2,099,650	0	0	0	3,167,205	2,291,685	-223,000	3,198,170	5,266,855
3	2,540,952	-441,302	2,099,650	0	0	0	3,167,205	2,291,685	-223,000	3,198,170	5,266,855
4	2,540,952	-441,302	2,099,650	0	0	0	3,167,205	2,291,685	-223,000	3,198,170	5,266,855
5	2,540,952	-441,302	2,099,650	0	0	0	3,167,205	2,291,685	-223,000	3,198,170	5,266,855
6	2,540,952	-17,289,319	-14,748,367	0	2,500,000	2,500,000	3,167,205	2,291,685	-5,000,000	-6,372,846	-9,081,162
7	2,540,952	-12,289,319	-9,748,367	0	7,500,000	7,500,000	3,167,205	2,291,685	5,000,000	-6,372,846	918,838
8	2,540,952	-7,289,319	-4,748,367	0	7,500,000	7,500,000	3,167,205	2,291,685	10,000,000	-6,372,846	5,918,838
9	2,540,952	-2,289,319	251,633	0	2,500,000	2,500,000	3,167,205	2,291,685	10,000,000	-6,372,846	5,918,838
10	2,540,952	5,372	2,546,324	0	0	0	3,167,205	2,291,685	255,000	3,166,845	5,713,529
11	2,540,952	5,372	2,546,324	0	0	0	3,167,205	2,291,685	255,000	3,166,845	5,713,529
12	2,540,952	5,372	2,546,324	0	0	0	3,167,205	2,291,685	255,000	3,166,845	5,713,529
Average	2,540,952	-3,445,639	-904,687	0	1,666,667	1,666,667	3,167,205	2,291,685	1,637,500	0	3,929,185
Expected	2,540,952	-260,352	2,280,600	0	16,667	16,667	16,037	2,291,685	21,618	0	2,313,303

profit, which reflects the fact that the Day-Ahead price is still less than the expected Real-Time price. Payments from load are the mirror image of the total payments made to the NGCC, the fast ramp generators, and virtual bidders. The column labeled uplift trues up the payments for generation which are calculated for each separate 5-minute interval and payments to the virtuals which are fixed across the full hour. In each Scenario, the total uplift nets to zero across the full hour.

Case #3, also with virtuals, is shown in Table 11.

The NGCC generators receive larger payments than in Case #1 and larger than in Case #2 since the Day-Ahead price was bid-up still higher. Some of these payments are reversed in the Real-Time market when the extra power is bought back. Once again, virtual demand earns losses in Scenario A, but earns profits in Scenario B. However, in this case, virtual demand earns zero expected profit, which reflects the fact that the Day-Ahead price exactly equals the expected Real-Time price. Payments from load are the mirror image of the total payments made to the NGCC, the fast ramp generators, and virtual bidders.

A key item to examine across the three cases is the total expected payment from load.

In Case #1, the total expected payment from load is $2.166 million. In Case #2, the total is $2.313 million, an increase of $0.147 million. Although load pays more, load does not receive any additional power—in fact there is absolutely no difference in the source of power or in the schedule of power. A portion of this increased cost is accounted for by the $0.016 million profits captured by virtuals. The remainder is due to the higher payments made to the NGCC generators—note that payments made to the fast ramp generators never change across any of the cases, so the remainder must all be due to payments made to the NGCC generators. These higher payments to the NGCC generators are an indirect consequence of the virtual bidders, which have driven up the Day-Ahead price as well as the volume of awards to the NGCC generators in the Day-Ahead market. These higher payments made to NGCC generators have do not reflect any extra delivery of energy from the NGCC generators. In Case #3, the total expected payment from load is higher still, $2.461 million. All of this extra payments flow to the NGCC generators: indeed, the total expected payments to virtuals drops to zero, so that the NGCC generators now capture all of the differential in payments between Case #3 and Case #1. This is all due to the indirect effect of the virtuals in increasing the Day-Ahead price as well as the volume of awards to the NGCC generators.

Table 11: Case #3, Virtual bidding-the limiting case: Payments from Load to Generators and Virtuals

Scenario A

	to NGCC Generators			to Fast Ramp Generators			to Virtuals	from Load			
interval	DA	RT	Total	DA	RT	Total		DA	RT	uplift	Total
1	2,949,092	-627,437	2,321,655	0	0	0	-61,184	2,439,120	-205,529	26,880	2,260,470
2	2,949,092	-596,538	2,352,554	0	0	0	-61,184	2,439,120	-170,004	22,254	2,291,370
3	2,949,092	-564,922	2,384,170	0	0	0	-61,184	2,439,120	-133,684	17,550	2,322,986
4	2,949,092	-532,569	2,416,524	0	0	0	-61,184	2,439,120	-96,547	12,767	2,355,339
5	2,949,092	-499,457	2,449,636	0	0	0	-61,184	2,439,120	-58,574	7,906	2,388,451
6	2,949,092	-465,565	2,483,527	0	0	0	-61,184	2,439,120	-19,743	2,966	2,422,342
7	2,949,092	-430,874	2,518,218	0	0	0	-61,184	2,439,120	19,966	-2,052	2,457,034
8	2,949,092	-395,362	2,553,730	0	0	0	-61,184	2,439,120	60,574	-7,148	2,492,546
9	2,949,092	-359,008	2,590,084	0	0	0	-61,184	2,439,120	102,103	-12,323	2,528,900
10	2,949,092	-321,792	2,627,301	0	0	0	-61,184	2,439,120	144,572	-17,576	2,566,116
11	2,949,092	-283,692	2,665,400	0	0	0	-61,184	2,439,120	188,004	-22,907	2,604,216
12	2,949,092	-244,688	2,704,404	0	0	0	-61,184	2,439,120	232,418	-28,317	2,643,220
Average	2,949,092	-443,492	2,505,600	0	0	0	-61,184	2,439,120	5,296	0	2,444,416

Scenario B

	to NGCC Generators			to Fast Ramp Generators			to Virtuals	from Load			
interval	DA	RT	Total	DA	RT	Total		DA	RT	uplift	Total
1	2,949,092	-642,625	2,306,467	0	0	0	6,057,250	2,439,120	-223,000	6,147,598	8,363,718
2	2,949,092	-642,625	2,306,467	0	0	0	6,057,250	2,439,120	-223,000	6,147,598	8,363,718
3	2,949,092	-642,625	2,306,467	0	0	0	6,057,250	2,439,120	-223,000	6,147,598	8,363,718
4	2,949,092	-642,625	2,306,467	0	0	0	6,057,250	2,439,120	-223,000	6,147,598	8,363,718
5	2,949,092	-642,625	2,306,467	0	0	0	6,057,250	2,439,120	-223,000	6,147,598	8,363,718
6	2,949,092	-26,317,258	-23,368,166	0	2,500,000	2,500,000	6,057,250	2,439,120	-5,000,000	-12,250,035	-14,810,915
7	2,949,092	-21,317,258	-18,368,166	0	7,500,000	7,500,000	6,057,250	2,439,120	5,000,000	-12,250,035	-4,810,915
8	2,949,092	-16,317,258	-13,368,166	0	7,500,000	7,500,000	6,057,250	2,439,120	10,000,000	-12,250,035	189,085
9	2,949,092	-11,317,258	-8,368,166	0	2,500,000	2,500,000	6,057,250	2,439,120	10,000,000	-12,250,035	189,085
10	2,949,092	-224,840	2,724,252	0	0	0	6,057,250	2,439,120	255,000	6,087,383	8,781,502
11	2,949,092	-224,840	2,724,252	0	0	0	6,057,250	2,439,120	255,000	6,087,383	8,781,502
12	2,949,092	-224,840	2,724,252	0	0	0	6,057,250	2,439,120	255,000	6,087,383	8,781,502
Average	2,949,092	-6,596,390	-3,647,297	0	1,666,667	1,666,667	6,057,250	2,439,120	1,637,500	0	4,076,620
Expected	2,949,092	-505,021	2,444,071	0	16,667	16,667	0	2,439,120	21,618	0	2,460,738

3.3.7 Lessons of the Illustration

The point of this illustration is to provide a concrete detailed example of exactly the type of price spikes CAISO experienced in 2011 and 2012 which were caused by insufficient ramping capabilities. The illustration also helps illuminate the role of virtual bidding what it does, but more importantly, what it does not do. In Scenario #2, virtual bidding raises the cleared Day-Ahead price up to the expected Real-Time price, which is very high due to the small probability of a very expensive fast ramp requirement. However, this isn't actually doing anything productive in terms of improving the efficiency of Day-Ahead scheduling. The extra generation that is receiving Day-Ahead award is generation that cannot actually provide any fast ramp capability. So this extra generation is not going to reduce any costs to the system.

In point of fact, the situation in CAISO was much worse. The illustration has minimized an important cost. It assumes there are no start-up costs or other frictions associated with the unnecessarily large dispatch awards made in the Day-Ahead market. Virtual bidding is the reason for this unnecessary dispatch. Extra generation is ordered Day-Ahead from the NGCC power plants, only to be reversed in the Real-Time market. In actual practice, that unnecessary dispatch drives up costs. This is an additional deadweight cost adding still more charges to load.

3.4 A General Problem: Episodic, Disparate and Transient

The CAISO experience related above is a specific example of the fault with virtual bidding that can arise as a result of the complexities associated with the optimal unit commitment and optimal power flow problems and the approximations, decompositions and judgments employed differently across the Day-Ahead and Real-Time algorithms. The same fault has manifested itself in a variety of cases–different situations in different ISOs in different years. Each time, the details of the case have been identified and diagnosed by market monitors and other staff at the ISO and FERC, and each time a specific remedy has been devised. However, the common character of the cases has not been appreciated. Presentations about virtual trading do not yet acknowledge this fault. Moreover, the episodic, disparate and transient nature of the cases disguises the common underlying fault. In order to help establish the common underlying fault, we now quickly describe a couple of other examples using the framework already established.

3.4.1 Locational Prices and Marginal Loss Factor Estimations

Losses are a cost that must be taken into account in calculating the optimal dispatch. Power flowing on a transmission network incurs losses due to heating of transmission lines. The losses mean that the total power delivered to customers is less than the power generated. These losses are a cost that needs to be included in the price. The size of the losses depends on many factors, including the equipment used, the voltage, distance traveled, and the amount of power flowing on the line. The total losses on the system depend on how the power flows across the system. Two generators located at different points on the network will face different losses in delivering power to the same customer, and these losses need to be taken into account in calculating the optimal dispatch.

The Day-Ahead algorithm uses estimated marginal loss factors. These are very difficult to determine accurately, and system operators are constantly trying to improve their estimation. Errors in estimating losses sometimes translate into differences between the Day-Ahead and Real-Time prices at various locations, occasionally producing predictable DA/RT spreads. This produces an opportunity for virtual bidders to exploit.

As the 2011 market monitor report for MISO relates,

> *While some [virtual bidders] appeared to take positions across constrained paths to arbitrage differences in day-ahead and real-time congestion, a few participants employed price-insensitive transactions to exploit sustained locational price differences due to marginal loss factor divergence between the day-ahead and real-time markets. One participant who appeared to be arbitraging significant differences in marginal loss factors between the day-ahead and real-time markets ceased this activity after MISO modified its methodology to eliminate large transitory differences.*[23]

And, in ISONE we see a similar thing:

> *The substantial drop in virtual transactions at the nodal level occurred in May 2010 when the ISO deployed a software solution to address an inconsistency in loss modeling at certain locations. This modeling inconsistency, which we first detected in late 2008, motivated a significant quantity of virtual trading at the affected locations because they produced low levels of consistent virtual profits (due to predictable differences between day-ahead and real-time LMPs). Hence, when this inconsistency was remedied, the associated virtual trading at those nodes ceased.*[24]

[23] MISO (2012), pp. 21-22.

[24] missing

As these two examples suggest, virtual bidders often know that what they are targeting is a discrepancy between the system operator's methodology for estimating marginal loss factors.

The case of loss factors is an example where the contribution of virtual bidders to system performance may be either positive or negative. Taking unit commitment as given, virtual bidding changes generation awards and prices in the Day-Ahead market, but the eventual dispatch in the Real-Time market is unaffected. The virtual bidder's profit is an additional cost to the system. If we do not take unit commitment as given, but recognize that Day-Ahead prices impact unit commitment, we see that the improved convergence produced by virtual bidding may improve unit commitment and lower system costs. Whether the benefits–in lower system costs–are worth the profits paid to the virtual bidder, is an empirical question. There is no foundation for believing that the benefits must be worth the cost.

In addition to the benefit of improved unit commitment, virtual bidding against the mis-estimation of loss factors may also indirectly improve system operation. As the quotes above suggest, the persistent profitability of virtual bidding at specific locations can call the system operator's attention to where improvements in its loss factor estimation are needed. This is a very different argument in favor of virtual bidding than is usually given, and there are many reasons to be circumspect about it. Like other arms races, this one does not lead to an optimal investment of time and modeling resources. Instead, time and modeling resources are shifted to a particular set of problems susceptible to gaming by virtual bidders. It is also hard to believe that profits paid out of the pockets of consumers translate efficiently into incentives for system operators. Finally, it is not always a quick and easy matter to fix the problem of mis-estimated loss factors. Nevertheless, this indirect dynamic is a real one that needs to be appreciated and taken into account in weighing the impact of virtual bidding.

3.4.2 Reserves

Uncertainty multiplies the economic complexity of the optimal dispatch and energy pricing problem. Even the theoretically complete ACOPF problem is a compromise with the full stochastic optimization problem that uncertainty entails. In the standard ACOPF problem, uncertainty is handled by specifying defined resource requirements of various sorts on top of the energy dispatch plan. Most of these are known as reserve requirements of one sort or another, and the precise set of requirements and the nomenclature varies. These requirements are then integrated into the Day-Ahead and Real-Time software, which presents constant challenges of its own due to the special economics of reserves. These challenges sometimes show themselves in price dynamics that virtual bidding exploits and sometimes worsens. Two examples are highlighted here.

First, the Day-Ahead optimization does not always commit sufficient reserves of one kind or another. System operators therefore sometimes commit additional units after the Day-Ahead run is complete. Come the Real-Time market, the extra capacity on-line drives down the Real-Time price below what might have been expected in the Day-Ahead calculations. The full cost of committing these units is not included in the Real-Time bid, since start-up and minimum load costs are backstopped by uplift charges.

This has been a recurrent problem in a number of ISOs, notably ISONE:

> *Non-market-based commitment and dispatch tends to depress real-time prices. This creates a premium in the day-ahead market and participants will naturally act on these economic incentives to reduce their day-ahead schedules. This can take the form of reduced schedules by LSEs in the area, reduced virtual loads, or increased virtual supply (all of which reduce the net load scheduled in the area). This under-scheduling pattern is self-reinforcing to some extent because it increases the need for supplemental commitment, which tends to reduce real-time prices and increases the incentive to under-schedule. The most effective way to address this problem is to reduce the need for supplemental commitment and out-of-merit dispatch over time by improving the representation of contingency requirements in the market software. The ISONew England has made strides in reducing supplemental commitment in Connecticut, and this is reflected by increased convergence between day-ahead and real-time scheduling there.*[25]

The same issue shows up in ISONE reports for 2007, 2008, 2009, 2010 and 2012. Although the report on 2004 makes it sound like the problem is being resolved, it is the nature of these sorts of fixes, that the problem, even once seemingly resolved, recurs or occurs anew in a new region.

Virtual bidding does not help the situation: on the contrary, it exacerbates it. Because the additional dispatch lowers the Real-Time price, it incentivizes additional virtual supply into the Day-Ahead market. This lowers the dispatch of physical resources in the Day-Ahead market, which only increases the system operator's measure of reserve resources required. Virtual bidding creates a spiral in the wrong direction.

The second example is the specific economics of fast start peaking resources. These units are often committed in the Real-Time market. However, they often have very high minimum run levels which do not match the marginal cost methodology generally employed in the Real-Time algorithm. As a result, when these units are dispatched, their costs are

[25] ISONE (2005), p. 63.

not used to set the Real-Time price. Once again, this creates the situation in which the Real-Time price is less than the true cost of the Real-Time generation. It also puts the expected Real-Time price below the Day-Ahead price, assuming that the Day-Ahead market is accurately reflecting expected load and generation costs. This mismatch creates incentives for physical load to bid less into the Day-Ahead market and shift demand into the Real-Time market.

The system in ISONE is also a good exemplar of this problem:

> *After the day-ahead market, the ISO may need to commit additional generators with high commitment costs to meet local and system-level reliability requirements. Once the commitment costs have been incurred, these generators may be inexpensive providers of energy and reserves. Because these commitment costs are not reflected in the market prices, the real-time LMPs frequently do not reflect the full value of online and fast-start capacity when generators are committed for reliability. Like any other forward financial market, the day-ahead market LMPs tend to converge with the real-time LMPs. Hence, day-ahead LMPs also do not reflect the full value of online and fast-start capacity, which reinforces the tendency of the day-ahead market-based commitment to not satisfy reliability requirements.*[26]

This is a good example of how incomplete is the framework of a single simple supply curve for adequately representing the true economics of electricity generation.

Virtual bidding does not help this situation, either. It also incentivizes additional virtual supply into the Day-Ahead market. This lowers the dispatch of physical resources in the Day-Ahead market, which increases the reliance on fast start generators. Virtual bidding once again only contributes to a spiral in the wrong direction.

3.5 Empirical Evidence on the Impact of Virtuals

Although virtual bidding is a regular feature of wholesale market design, with many advocates, the actual evidence that it improves system performance is surprisingly scant and weak. The major support comes from studies of DA/RT spreads before and after the implementation of virtual bidding in the New York and California markets.

Two often cited studies that find evidence that virtual bidding contributed to convergence in NYISO are NYISO (2003) and Saravia (2003): the former looks at both the average absolute spread as well as the average level of the spread, while the latter focuses on the average level of the spread. Two key findings are: (i) system wide, the absolute value

[26] ISONE (2013), p. 93.

of the DA/RT spread was higher in the first two years of the market's operation, before virtual bidding was permitted, than in the subsequent two years, after virtual bidding was permitted; and (ii) the average DA/RT spread, which was positive throughout the roughly four years, was larger in the two years before virtual bidding and smaller in the two years after. These two facts certainly suggest that virtual bidding helps converge the market in both senses: reducing the absolute size of the variances, and centering the average premium closer to zero. Saravia (2003) goes beyond this raw comparison, developing a more complete model of imperfect competition and the predicted pricing patterns with and without virtual bidding, and finds empirical support for virtual bidding reducing a positive DA/RT spread along with other reflections of market power.

Unfortunately, neither study controls for the various other factors changing at the same time as virtual bidding was introduced–the NYISO study cautiously recognizes this problem and hedges any claims that virtual bidding caused the change in DA/RT spreads. Hadsell and Shawky (2007) point to shakeout issues in the initial year that the multi-settlement market operated, including an event of sham transactions targeted to manipulating prices. Chaves and Perez (2010) attempt to control for the changing market power mitigation regimes used through 2005, and find a siginificant residual effect of virtual bidding in reducing DA/RT spreads. Nakano (2007) focuses on the implementation of demand response programs in the summer of 2001, immediately before the start of virtual bidding, and finds that the drop in DA/RT spreads coincided with this rather than the virtual bidding.

Jha and Wolak (2014) study DA/RT spreads before and after virtual bidding was introduced in CAISO in February 2011. They find that in several of the hours (19 of 24), the DA/RT spread is closer to zero after virtual bidding. Whereas previously the spread was often very negative, now it is dramatically less negative or slightly positive. They also find that the volatility of the spreads has declined as have the volatility of Real-Time prices, and they find a differential change in DA/RT spreads across generation and non-generation nodes which they argue is consistent with virtual bidding reducing monopsony power.

Surprisingly, their study never addresses any of the underlying issues discussed by the market monitor. The study described CAISO's multi-settlement markets without mentioning the Hour Ahead market that drove such a share of the virtual bids actually placed. The study never mentions the many adjustments to the operation of the Day-Ahead market. In particular, it never mentions the changing requirements for ramping capacity and how that may have impacted DA/RT spreads and accounted for the changes documented.

The recent study by Larrieu (2014) also examines virtual bidding in CAISO from April

2009, when the market began, through February 2011, when virtual bidding was implemented, and then to March 2014. He focuses on the relationship between the absolute value of the DA/RT spread and the volume of virtual demand and the volume of virtual supply, controlling for other variables. While virtual supply is associated with a reduction in the average absolute DA/RT spread, virtual demand is associated with an increase. This is true at the system energy level, and the result is even stronger at the two aggregation points, NP-15 and NP-16. This increase could be construed as a consistent with the type of problems described in this paper.

In addition to these academic studies, each regional market produces an annual report by a market monitor, and these reports provide rich analyses of many issues, including the impact of virtual bidding. Since each report is designed as a review of the past year, the scope is limited. On the other hand, the monitors have deep knowledge of the structure of the industry in the region, of the institutional history, and of the specific details of bidding and other underlying factors, which is very valuable. While some of the authors of the reports are clearly advocates for virtual bidding, the nature of the reports and the authors' deep familiarity with system operations means that the variety of factors driving DA/RT spreads are prominently discussed. Occasionally, the impact of virtual bidding is scrutinized, too. Unfortunately, some market monitors simply invoke the assertion that virtual bidding improves market functioning, without having any empirical foundation for that assertion. For example, the Market Monitoring Unit for the New York ISO writes in its State of the Market Report for 2012 that:

> *Virtual trading activity helps align day-ahead prices with real-time prices, particularly when modeling and other differences between the day-ahead and real-time markets would otherwise lead to inconsistent prices. Overall, virtual traders were profitable in 2012, indicating that they generally improved convergence between day-ahead and real-time prices, which facilitates an efficient commitment of generating resources.*[27]

This assertion is not accompanied with any statistical evidence or citation, and conversation with its author confirm that there is no specific foundation aside from the general theory that virtual bidding is advantageous. Comparable conclusory comments sprinkle many market monitoring reports. The only way to determine whether or not virtual bidding has contributed to system performance is to examine the specific facts and circumstances. Contrast the idealistically insistent assertion above with this very pragmatic and empirical approach taken by the CAISO Market Monitoring Unit:

[27]NYISO (2013), p. iii.

> *Convergence bidders profit by taking advantage of differences between day-ahead, hour-ahead and real-time prices. In theory, if participants successfully profit from virtual bidding, this activity should drive day-ahead, hour-ahead and real-time prices closer. However, this theoretical impact of virtual bidding may not occur because of a market feature that makes the California market design different from most other ISOs. The financial settlement of inter-tie convergence bids based on hour-ahead prices has led to additional uplifts, known as imbalance offset costs, which can occur when prices diverge between the hour-ahead and real-time markets. Virtual bidding on inter-ties increased these imbalance offset costs by increasing the volume of transactions clearing at these different market prices.*[28]

Only a specific investigation can validate whether the profit to virtuals are a sign of their contribution or their costs. That investigation should be open to the complex causes for spreads and the potential costs of virtual bidding.

4 Conclusion

This paper has analyzed the theory and practice behind virtual bidding. It has focused attention on the often overlooked differences between the implementation of Day-Ahead and Real-Time markets which sometimes give rise to DA/RT spreads and which are assumed away by advocates for virtual bidding. The paper has shown that these differences can undermine the case for virtual bidding. The paper has used cases from various ISOs and across different years to illustrate the different manifestations of this common problem. The paper has established that:

- spreads between the Day-Ahead and the Real-Time price will often arise due to the many necessary approximations differently employed in the Day-Ahead and Real-Time algorithms;
- while virtual bidders can profit off of these spreads, oftentimes they cannot help resolve the underlying problem;
- in these cases, profits earned by virtual bidders can be a purely parasitic drain on the system, adding to the costs paid by load;
- in addition, virtual bidders may add to system costs;

[28] CAISO (2013), p. 36.

- convergence—a narrowing DA/RT spread—is an imperfect metric for evaluating system performance and the contribution of virtual bidders; virtuals may cause the average DA/RT spread to move closer to zero, and nevertheless all virtual profits are a purely parasitic drain, and, in addition, virtual trading has increased system costs.

Virtual bidding is a peculiar beast. It is a form of financial trading that has not developed organically out of trade among industry participants. It has instead been carefully grafted onto the design for wholesale trading of physical electricity. Virtual bidders are wanted because they will help drive out any and all disparities between the Day-Ahead and the Real-Time price—to drive DA/RT spreads to zero. The introduction of virtual bidders is analogous to the classical form of biological control in which an animal is introduced into an ecosystem to help manage pests. Biological control can be an effective tool, and there are many cases in which it has been successfully managed. But there are also cases of spectacular failure, especially where it has been implemented naively and the many possibilities for bad consequences have been overlooked. Unfortunately, many advocates of virtual bidding have been willing to overlook important cases in which it has proven costly to consumers and hurt system performance, driving up costs. Understanding when virtuals contribute to system performance and when they are parasitic and also hurt system performance is a difficult, empirical challenge.

REFERENCES

Abdul-Rahman, Khaled H., Hani Alarian, Mark Rothleder, Petar Ristanovic, Bogdan Vesovic and Bo Lu. 2012. Enhanced System Reliability Using Flexible Ramp Constraint in CAISO Market. IEEE Power and Energy Society General Meeting, July, pp. 1-6.

Bessembinder, Hendrik, and Michael L. Lemmon. 2002. "Equilibrium pricing and optimal hedging in electricity forward markets." the Journal of Finance 57.3: 1347-1382.

Birge, John, Ali Hortaçsu, Ignacia Mercadal and Michael Pavlin. 2013. The Role of Financial Players in Electricity Markets: An Empirical Analysis of MISO, Working Paper, October 22.

Borenstein, Severin, James Bushnell, Christopher Knittel and Catherine Wolfram. 2008. Inefficiencies And Market Power In Financial Arbitrage: A Study Of California's Electricity Markets, Journal of Industrial Economics 56(2): 347-378.

Bowden, Nicholas, Su Hu and James Payne. 2009. Day-Ahead Premiums on the Midwest ISO, Electricity Journal 22(2):64-73, March.

Bunn, Derek W., and Dipeng Chen. 2013. The Forward Premium in Electricity Futures, Journal of Empirical Finance 23:173-186.

Cain, Mary B., Richard P. O'Neill and Anya Castillo. 2013. History of Optimal Power Flow and Formulations. Optimal Power Flow Paper 1. FERC Staff Technical Paper, August.

CAISO. 2009a. Quarterly Report on Market Issues and Performance. July 30.

CAISO. 2009b. Quarterly Report on Market Issues and Performance. December 23.

CAISO. 2010a. Quarterly Report on Market Issues and Performance. February 1.

CAISO. 2010b. 2009 Annual Report on Market Issues and Performance. April.

CAISO. 2011a. Quarterly Report on Market Issues and Performance. February 8.

CAISO. 2011b. 2010 Annual Report on Market Issues and Performance. April.

CAISO. 2012. 2011 Annual Report on Market Issues and Performance. April.

CAISO. 2013. 2012 Annual Report on Market Issues and Performance. April.

CAISO. 2014. 2013 Annual Report on Market Issues and Performance. April.

Chaves, José Pablo and Yannick Perez. 2010. Virtual Bidding: A Mechanism to Mitigate Market Power in Electricity Markets: Some Evidence from New York Market, Working Paper.

Douglas, Stratford, and Julia Popova. 2008. Storage and the electricity forward premium, Energy Economics 30:1712-1727.

FERC. 2002. Notice of Public Rulemaking, July 31, 2002, 18 CFR Part 35, [Docket No. RM01–12–000], Remedying Undue Discrimination Through Open Access Transmission Service and Standard Electricity Market Design, Federal Register Vol. 67, No. 168, Thursday, August 29.

FERC. 2012. Order Approving Stipulation and Consent Agreement. Constellation Energy Commodities Group, Inc. Docket No. IN12-7-000. March 9.

Hadsell, Lester. 2008. Day-Ahead Forward Premiums on the New England ISO, Electricity Journal 21(4):51-57.

Hadsell, Lester. 2011. Inefficiency in Deregulated Wholesale Electricity Markets: the Case of the New England ISO, Applied Economics 43:515-525.

Hadsell, Lester, and Hany A. Shawky. 2006. Electricity Price Volatility and the Marginal Cost of Congestion: An Empirical Study of Peak Hours on the NYISO Market, 2001-2004, Energy Journal 27(2):157-179.

Hadsell, Lester, and Hany A. Shawky. 2007. One-Day Forward Premiums and the Impact of Virtual Bidding on the New York Wholesale Electricity Market Using Hourly Data, Journal of Futures Markets 27(11):1107-1125.

Haugom, Erik, and Carl J. Ullrich. 2012. Market Efficiency and Risk Premia in Short-Term Forward Prices, Energy Economics 34:1931-1941.

Isemonger, Alan G. 2006. The benefits and risks of virtual bidding in multi-settlement markets. The Electricity Journal, 19:26-36.

ISONE. 2005. 2004 Assessment of the Electricity Markets in New England, Potomac Economics, Independent Market Monitoring Unit, ISO New England, June.

ISONE. 2013. 2012 Assessment of the ISO New England Electricity Markets, Potomac Economics, External Market Monitor for ISO New England, May.

Jha, Akshaya, and Frank Wolak. 2014. Testing for Market Efficiency with Transactions Costs: An Application to Convergence Bidding in Wholesale Electricity Markets. Working paper, March.

Larrieu, Jeremy. 2014. Impact of Virtual Bidding on the Forward Premium in the California ISO. George Washington University, Working Paper, August.

Ledgerwood, Shaun, and Johannes Pfeifenberger. 2013. Using Virtual Bids to Manipulate the Value of Financial Transmission Rights, Electricity Journal 26(9):9-25.

Longstaff, Francis A., and Ashley W. Wang. 2004. "Electricity forward prices: a high-frequency empirical analysis." The Journal of Finance 59(4): 1877-1900.

MISO. 2012. 2011 State of the Market Report for the MISO Electricity Markets, Potomac Economics, Independent Market Monitor for MISO, June.

Nakano, Shoko. 2007. Impact Mitigation for Emergency Events: Their Effects on Day-Ahead and Real-Time Market Locational Based Marginal Pricing at the New York ISO, Hitotsubashi Journal of Economics 48:185-197.

NYISO. 2003. Report on Impact Of Virtual Trading on the Summer 2003 New York Electricity Markets. Market Monitoring and Performance Department, Market Analysis Unit, November.

NYISO. 2013. 2012 State of the Market Report for the New York ISO Markets, Market Monitoring Unit for the New York ISO, April.

NYISO. 2014. 2013 State of the Market Report for the New York ISO Markets, Market Monitoring Unit for the New York ISO, May.

Pirrong, Craig, and Martin Jermakyan. 2008. The Price of Power: The Valuation of Power and Weather Derivatives, Journal of Banking and Finance 32:2520-2529.

PJM. 2013. State of the Market Report for PJM. 2012. Volume 2: Detailed Analysis. Monitoring Analytics, Independent Market Monitor for PJM, March.

PJM. 2014. State of the Market Report for PJM. 2013. Volume 2: Detailed Analysis. Monitoring Analytics, Independent Market Monitor for PJM, March.

O’Neill, Richard P., Thomas Dautel and Eric Krall. 2011. Recent ISO Software Enhancements and Future Software and Modeling Plans, FERC Staff Report, November.

Saravia, Celeste. 2003. Speculative Trading and Market Performance: The Effect of Arbitrageurs on Efficiency and Market Power in the New York Electricity Market, Center for the Study of Energy Markets (CSEM) Working Paper, University of California, Berkeley.

Ullrich, Carl J. 2007. Constrained Capacity and Equilibrium Forward Premia in Electricity Markets, Working Paper.

Werner, Dan. 2014. Electricity Market Price Volatility: The Importance of Ramping Costs, Working Paper.