Virtual Bidding and Electricity Market Design

William W. Hogan

May 25. 2016

Table of Contents

Virtual Bidding and Electricity Market Design

William W. Hogan

Summary

Efficient electricity day-ahead market designs include virtual transactions. These are financial contracts awarded at day-ahead prices and settled at real-time prices. In PJM these virtual transactions include incremental offers (INCs) that are like generation offers, decremental bids (DECs) that are like demand bids, and up-to-congestion bids (UTCs) that are like transmission price spread bids. Virtual transactions offer potential benefits to improve the efficiency of electricity markets, mitigate market power, enhance price formation, hedge real-time market risks, and price those risk hedging benefits.

The role and performance of virtual transactions has been a subject of controversy. A report by PJM addresses some of these controversies, identifies possible problems in the present implementation of virtual transactions with the associated settlement rules, and makes recommendations for changes in the treatment of virtual transactions. The PJM report is generally supportive of the contribution of virtual transactions as improving overall market performance. Illustrative examples in the report highlight these contributions and add to the general understanding of the benefits and some of the problems with its current rules for treating virtual transaction.

Although these examples help in explaining the mechanics of virtual transactions, and the interactions with the underlying physical market, the examples do not provide a framework for evaluating the overall cost and benefits of virtual transactions. The PJM analysis is not alone in this regard, because the evaluation task is not easy. There is no readily available template waiting to be applied to the PJM case. The limited available analyses from other regions indicate that the benefits are material and outweigh the costs, but no available studies cover all the relevant issues. However, going beyond examples of particular outcomes to consider, the broader context is important. Looking to the broader framework can change both the diagnosis of the symptoms and the prescriptions for the cures.

Under the current PJM market rules, there is an asymmetry in the settlement treatment of different types of virtual transactions, applying residual uplift charges to INCs and DECs but not to UTCs. One of the PJM recommendations is to eliminate this asymmetry by extending the same uplift treatment to UTCs. The argument is based on allocation of uplift costs according to the deviations between real-time quantity and day-ahead schedules. This approach is particularly problematic for virtual transactions, which by design involve a 100% deviation.

There is no simple connection between deviations, uplift costs and market efficiency. Under a broader equilibrium analysis there can be conditions where there is no relationship between any of these components. Furthermore, the allocation of properly defined residual costs according to a cost causation argument can in itself be a contradiction. More importantly, the focus on uplift cost causation is misplaced. The important question is the aggregate net benefit of virtual transactions, not the residual cost. If virtual transactions increase the net benefits in the market, then there is no incentive-based reason to assign additional costs to virtual transactions. The

criterion for assigning residual costs would then turn to doing the least damage to the performance of the market.

A better symmetric solution is to avoid any uplift allocation to virtual transactions. The residual cost allocation would then apply to real load; liquidity and entry in financial day-ahead virtual transactions would be enhanced; market power would be reduced; accurate price formation would be supported; and the efficiency of the overall PJM electricity market system should be improved. This reversal of the conventional wisdom follows from a broader framework than that applied by PJM for consideration of the costs and benefits of virtual transactions.

This broader framework builds directly on the basic principles of efficient electricity market design. Stepping back to consider first principles makes it easier to see the connections among the components of market design, in order to consider the function and benefits of virtual transactions from the perspective of aggregate market performance. PJM's own analysis provides many examples of the contributions and effects of virtual bidding, but does not connect the examples to the broader framework of electricity market design principles. Furthermore, going beyond the uplift allocations, the PJM recommendations restricting the use of virtual transactions do not follow necessarily even from a narrower evaluation perspective. The principal problem PJM identifies with virtual transactions is a computational burden that would be only indirectly affected by uplift allocations, and could be addressed through other means with fewer negative consequences for the broader market design, such as by continuation of bidding budgets that allowed flexibility in the choice of virtual transactions.

Restricting explicit virtual bidding, as PJM proposes, creates market power for those who can make implicit virtual bids. Explicit virtual bidding mitigates or eliminates this market power, provides liquidity, improves price formation, allows hedging, connects naturally with longer term financial transmission rights, helps reveal defects in market design, and on average should improve system operations.

The PJM report appears in a context where virtual bidding is under attack. While a complete cost benefit analysis is not available, the PJM analysis can be expanded to enhance both the understanding of the role of virtual bidding and the policies that support overall electricity market efficiency.

Virtual Bidding and Electricity Market Design

William W. Hogan[i]

May 25, 2016

Introduction

Efficient electricity day-ahead market designs include virtual transactions. In particular, in markets with a multi-settlement system including a day-ahead and a real-time market, day-ahead transactions can clear bids and offers that are strictly financial and are not intended for physical fulfillment in real-time. These day-ahead financial transactions are settled against real-time prices in the same manner as other day-ahead market transactions. A recent PJM study reviewed the operation of its energy market, discussed the role of virtual transactions, and offered recommendations on proposed rule changes that would affect the scope and treatment of virtual transaction participation in their day-ahead market. (PJM, 2015). The purpose of the present paper is to comment on this PJM analysis and recommendations.

A full analysis of the impacts of virtual bids must immediately consider and model outcomes in an electricity market with uncertainty. Assessing the costs and benefits of virtual transaction on electricity market outcomes with even an approximation of the complications induced by realistic unit commitment and dispatch, e.g., taking into account uncertainty about the level of real-time load and resource availability, would be difficult. The PJM report does not attempt such an analysis, but argues primarily from examples that pertain to a context without uncertainty. There are important features of the implicit assumptions in the PJM analysis that affect the conclusions about the costs and benefits of virtual bidding, and introduction of reasonable and realistic changes to these assumptions would lead in a different direction than PJM in specifying recommendations.

PJM Analysis Overview

The PJM analysis presents background context and provides discussion and analysis of many of the issues surrounding virtual bidding. The context can be alarming, as in the consideration of possibilities of market manipulation, where PJM raises the specter of "…perhaps going so far as to eliminate outright virtual trading in RTO markets" (PJM, 2015, p. 9). However, although PJM notes the concerns about market manipulation, the PJM analysis neither takes a position on this matter nor pursues explicitly the how and the where of possible market manipulation. Essentially, market manipulation is treated as a separate topic and while presented as context is not afforded material discussion or analysis.

The focus of the PJM analysis is narrower and addresses the efficiency effects and benefits of virtual trading assuming that market participants are simply responding to market signals without an attempt to manipulate those signals. This aspect of the PJM analysis is generally supportive of the impacts of virtual bidding. The discussion and examples in the analysis illuminate the issues and are instructive in expanding our understanding of the many dimensions of the benefits, and costs, of virtual bidding.

Broadening the discussion of virtual bidding requires more background about the context of electricity market design. The review below summarizes the critical and relevant engineering-economic elements of efficient electricity market design, short-term electricity system operations, and long-term contracts, and their relationship to the use of virtual transactions in electricity markets.[1] The interconnections among these topics have implications for the evaluation of recommended changes in the treatment of virtual transactions. The main conclusions are that the PJM recommendations are in certain cases inconsistent with the broader principles of efficient market design, and in other cases there is no direct connection between the PJM analysis and its recommendations. Both the scope of PJM's analysis and the breadth of its recommendations should be expanded.

Electricity Market Design

The special characteristics of the electrical transmission network create strong instantaneous interactions in how power flows between and among generators injecting energy into the grid and loads withdrawing energy. Dealing with these interactions induces related interactions in the elements of electricity market design. Although the structure and interconnections may be familiar, it is often helpful to go back to the basics to understand how the pieces fit together. Forgetting the details of the larger context linking the market design economics to engineering principles can result in analyses and recommendations that can neglect the requirements of efficient electricity market design and recreate problems already solved. A relevant case in point appears in recommendations for undoing financial transmission rights in PJM, see (Monitoring Analytics, 2016, sec. 13), therein ignoring the long history of the fundamental transmission problem they were intended to solve (Hogan, 1992) (Hogan, 2002a) (Pope, 2016).

The central idea of efficient electricity market design is to recognize the critical engineering characteristics of the power system, operate that power system efficiently, and utilize prices and associated incentives that are consistent with and motivate efficient operation.

The distinctive critical characteristics of the power system are the lack of adequate storage, meaning that most power must be generated contemporaneously with its use, and limits on whether and how system operators can adjust which transmission lines power flows on as it moves through the grid (parallel flows). Due to the lack of adequate storage, the speed of power flows, and response of other engineering elements of the system, system operators need to maintain essentially instantaneous balance, i.e., equivalence, of generation and load. This balance between generation and load occurs as power flows on the grid along every parallel path between supply sources and load sinks in quantities determined by the engineering ratings of each specific transmission line, among other things, rather than through a system where the pattern of the flows can be controlled by valves and pipes. In effect, therefore, the use of the transmission grid, in terms of power flows on transmission lines is determined by the distribution of load across the system and the dispatch of supply at different locations by the system operator. In every interconnected grid, a system operator is required to control the dispatch in order to control the flows on the grid within security limits.

These are not new challenges, and are familiar to power engineers who have well-developed techniques and tools for economic dispatch to coordinate and control flows on the transmission

---

[1]   This is an updated and expanded version of (Hogan, 2012).

2

system to maintain reliability. In choosing the generation dispatch to supply load within the limits of power flow constraints, there is still a great deal of flexibility, so some criterion needs to be applied. The natural approach for choosing among alternative feasible dispatches is to minimize the costs or maximize the net benefits of the electricity system operation. The term of art is to choose the "economic dispatch" to meet the load at the least cost subject to the security (i.e., reliability) constraints of the electricity system.

An efficient design for real-time markets should address the special challenges of electricity system operation and support the intended economic outcomes by providing a spot market basis for development of and reliance on forward contracts. The essence of the successful electricity market design in PJM and elsewhere was to organize the real-time spot market around the principles of bid-based, security-constrained, economic dispatch with the associated locational prices (Hogan, 1992). Under this market system, market participants are able to buy, sell and trade electricity through a non-discriminatory organized spot market. Settlement prices are the real-time locational prices. Charges for transmission service between locations are settled at the difference in the locational prices for the injection and withdrawal. The real-time locational prices can be volatile, but forward contracts allow market participants to hedge the real-time prices.

Applying security-constrained economic dispatch is a well-developed practice in power systems. It developed using engineering estimates of the operating costs of generation. The adaptation to markets was to replace the engineering cost estimates with the bids and offers of the market participants. With this change in the inputs, the form of the economic dispatch remained otherwise unchanged.

The second innovation of markets was to apply consistent prices to the purchases and sales determined in the economic dispatch. A by-product when determining the economic dispatch is the calculation of the marginal costs of incremental power at each location. Following the usual definition of competitive markets, these marginal costs define the market-clearing prices associated with the economic dispatch. Under reasonable simplifying assumptions about the nature of the dispatch, taking these prices as given the generators and loads would have no incentive to deviate from the dispatch.[2] These spot prices are known in the PJM system as Locational Marginal Prices (LMP) (Schweppe, Caramanis, Tabors, & Bohn, 1988).

Using any other materially different pricing system would by construction create a fundamental inconsistency with the market quantities determined in the economic dispatch. Because of this inconsistency, implementation of a pricing system other than LMP would require surrendering the benefits of efficient dispatch, restricting open access, or abandoning the principle of non-discrimination, or all of the above. There is no other pricing system that is compatible with economic dispatch, open access and non-discrimination. Therefore, the centerpiece of successful market design is bid-based, security-constrained, economic dispatch with locational marginal prices.

---

[2] The principal simplifying assumption employed is convexity of the cost function. More generally, the market clearing prices depend on the absence of a duality gap (Gribik et al., 2007).

Electricity Market Design and Forward Contracts

Electricity production is capital intensive. Furthermore, the cost structure implies that short-run system marginal costs and the associated prices will be volatile. This creates an interest in forward contracts to allocate the ubiquitous risks in the industry. Both customers and producers see forward contracts as inherently useful. Customers are interested in forward contracts of a variety of forms to manage the risks associated with future purchases of energy. Similarly, producers are interested in forward contracts to manage the complementary risks created by high investment in generating assets to be repaid through an otherwise uncertain stream of revenue.

In the traditional regulated industry, most of the management of risk between consumers and producers was handled through vertical integration of regulated utilities. Contracts played an important role, particularly for smaller companies and municipalities that were not able to diversify their supply through vertical integration, but forward contracts were a supplement rather than the main instruments for risk allocation.

Electricity restructuring, especially that which involved separation of generation, transmission and load, changed the business environment and created a much greater emphasis on the importance of forward contracts as the means for addressing the risk allocation that had been inherent in vertical integration. In addition, electricity restructuring was expected to expand the richness and diversity of contract forms and agents, giving rise to a much greater role for intermediaries that would trade in forward contracts that were essentially financial instruments that need not be created or held to match any particular physical transaction (Hogan, 2002a).

Early discussion of electricity restructuring in the 1980s emphasized the importance of forward contracting and markets without describing a coherent framework to meet this need (Joskow & Schmalensee, 1983). The argument was that markets and contracts could play a critical role, but the challenge was to provide a workable framework for access to the transmission system in order to support these markets.

Contracts for Difference

An important feature of successful electricity market design is the necessity to separate the financial role of contracts used to allocate risk and the physical operation of the system. It is difficult to impossible to rely on forward contracts to govern short-term physical operations and electricity dispatch because of the strong interactions in physical power flows on the transmission grid. An essential feature of successful market design is to enable separation of the dispatch from the contracts, meaning that the system operator does not need to know about the contracts or use any information embedded in those contracts to operate the electricity system. As a result, all forward contracts in electricity systems with organized markets are essentially financial contracts within the organized markets used to reallocate the financial risk of physical system operations. Scheduled transactions between organized markets can require decisions with lead times of an hour, after which the essential financial nature of the forward contract is replaced by a short term schedule with a closer connection to physical dispatch.

The basic forward contracts in electricity markets include a number of simple building blocks. For example, a simplified contract for difference at a location would specify a fixed quantity for receipt by a load and delivery by a generator, say for 100 MW in real-time at the location of the load. The actual physical transaction would be through the real-time spot market. The load would purchase 100 MW through the spot market. The generator would sell 100 MW through the spot market. The contract for difference would require payment of the difference between

the contract price and the real-time spot price by the load to the generator (or paid by the generator to the load if the difference is negative). The net result is for delivery of 100 MW from the generator to the load at the contract price. The contract for difference formulation provides a building block facilitating many additional types of transactions. For example, if the load or the generator deviates from the 100 MW contract quantity, the quantity difference is in effect sold to or purchased from the spot market. The system operator is not involved, nor even aware of the imbalance on the contract, which is simply a financial arrangement.

A variant of the contract for difference is the option, where the spot payment is made only in the event of the choice by the party that has acquired the option right. In combination, contracts for difference and options provide the raw material for market participants to craft an almost unlimited range of alternatives for structuring the allocation of short term energy price risk through forward contracts. Using the real-time spot market prices as the basis for settlements, it is possible to apply many different contract forms found in other markets and for other commodities to the case of electricity.

Financial Transmission Rights

Economic dispatch addresses the strong interactions of power flows in an electrical network. For equilibrium to hold at the efficient economic dispatch, the difference in market-clearing locational spot prices must equal the opportunity cost of transmission between the locations. Therefore, the consistent spot price for transmission is this difference in locational prices. If it were not for the strong interactions in the network, it would be possible to define a set of physical rights for transmission, and these physical rights could be traded to produce an efficient use of the network. In equilibrium, the spot price of these tradable physical transmission rights would be equal to the spot price of transmission under economic dispatch. But a workable system of such physical rights is not available.

The replacement for the unworkable physical transmission right is the financial transmission right (FTR) to collect the difference in the locational prices (Hogan, 1992). Spot electricity prices are volatile, and this transmission spot price is even more volatile. The FTR is the right to collect the difference in the locational prices.[3] In effect, the FTR is paid the equivalent of a physical right sold at the real-time spot price without the necessity of actually trading the physical right. This provides a hedge for real-time physical transactions between locations. A physical transaction incurs a spot transmission charge equal to the difference in the locational prices between the source and sink; this charge is the net of the spot price value of the generation injection at its source and the spot price cost of the load withdrawal at its sink. The payment to the owner of an FTR is this same difference in locational prices. If the physical transaction and the FTR are exactly matched, then the net payments cancel as though the schedule had used a physical transmission right. From the perspective of the physical schedule, the cost of connecting the source to the destination is the cost of acquiring the FTR. Acquisition of an FTR in advance, at a known price provides a hedge for the possibly volatile difference in real-time locational prices between a source and a sink.

---

[3] Here we address the usual FTRS defined for congestion differences, ignoring losses. The definition of FTRs includes possible treatment as obligations or options. In the case of obligations the holder of the FTR receives payment when the difference in congestion costs is positive and makes a payment when the difference in negative. Under the FTR option, the case of negative difference in congestion costs does not require a payment. The mix of options and obligations affects the simultaneous feasibility of FTRs. But for the present discussion the differences are not important and the focus of discussion is on the treatment of FTR obligations.

The revenue to fund payments to the owners of the FTRs arises from the residual remaining after paying suppliers and charging loads for their economic dispatch quantities at spot prices. The total payments by load exceed the payments to generators, reflecting the congestion value arising from the use of transmission to move lower cost generation to serve load in higher cost locations, and the marginal pricing of transmission losses; together these comprise the spot market surplus. If the full set of outstanding FTRs is simultaneously feasible for the grid conditions used in the economic dispatch, then under certain regularity conditions, the net of the spot market payments in the physical market will be sufficient to fund payments for the FTRs (Hogan, 2002b).

This use of transmission payments in the form of congestion rents to support FTRs is a consequence of the FTR design, but not a fundamental purpose. The history is clear as to the need for some form of fundamental transmission rights, and that physical transmission rights are impossible to design under the principles of open access and non-discrimination. The contrary view of the PJM market monitor illustrates how isolated recommendations, separated from the history of development of the basic principles of efficient market design, can lead to recommendations that would fundamentally undermine the entire structure. In particular, the advice to ignore the hedging function of FTRs and simply find ways to refund congestion rents without the use of FTRs, (Monitoring Analytics, 2016, sec. 13), would unravel a key component of long-term contracting.

The FTR provides a critical piece in the elements of a workable and efficient electricity market design under the principles of open access and non-discrimination. The core contribution is in providing a substitute for the congestion hedges of the unavailable physical transmission rights. The existence of FTRs creates the opportunity to replicate many other features of efficient markets with an array of forward contracts and hedging instruments.

Day-Ahead Markets and Multi-Settlement Systems

The description above of economic dispatch and FTRs applies to the real-time market. The contracts for differences and FTRs are relatively long-term arrangements that typically apply across many real-time periods. In principle and in practice, the system operator need not consider these contracts when conducting the dispatch. For a variety of reasons, the design of electricity markets includes or soon gravitates to include formal integration of forward markets that allow for advance notice for unit commitments, physical schedules, and financial hedges. This typically takes the form of a day-ahead market scheduling 24 consecutive hours, which allows for more alternatives and flexibility in sequencing the commitment of units and dealing with the complex dynamics of generation ramping to meet forecast changes in the level of load. The market design could even include a multi-settlement system with hour-ahead, day-ahead, and longer forward markets. The details are slightly different in each actual ISO/RTO implementation, due to differences in desired or mandated scheduling lead times and mechanisms.

In essence, the day-ahead market operates analogously to the real-time market. Participants make demand bids and supply offers. The system operator combines these offers and bids with a description of the expected transmission network conditions and determines an economic unit commitment and dispatch with associated LMPs for each hour of the next day. The system operator settles the market by charging or paying market participants for purchases and sales in

the day-ahead market at the day-ahead prices. Bilateral transmission schedules pay for their schedule from source to sink at the difference in the day-ahead prices.

There are other relevant details about commitment costs and related limitations on generation. But for the present discussion an important feature of this day-ahead, market-clearing, economic dispatch is that the schedules are effectively all day-ahead financial contracts that will be cleared either physically or financially in real-time. No power flows in the day-ahead market. In real-time, when power actually flows, there will be another economic dispatch based on real-time offers and bids. Although the mechanics of accounting and settlements may be different in each ISO/RTO, the net settlement result for each market participant is the same as if all the day-ahead schedules were liquidated at the real-time locational spot price, and all real-time physical transactions were settled at the respective locational spot prices. This formulation of the accounting is sometimes referred to as a "gross pool." An alternative and financially equivalent accounting process is to settle the differences between the day-ahead and real-time quantities for a participant at a location at the real-time price. This is sometimes referred to as a "net pool." Note that the settlement system characterization does not affect the dispatch or the ultimate net payments by the market participants. The economics and the aggregate financial outcome and, hence, economic incentives are unchanged.

This recognition that all cleared schedules in the day-ahead market are financial contracts that can be settled at real-time prices rather than through physical delivery of energy, immediately opens the possibility for entities that are not "physical loads" or "physical generators" to participate in the day-ahead market. Given the uncertainty in real-time, e.g., concerning weather conditions affecting load and generation levels, the day-ahead contracts provide a hedge against the volatility of real-time prices. To the extent that expected real-time prices differ from the day-ahead prices, there is an arbitrage opportunity. This arbitrage possibility creates an incentive for purely financial participants to make bids and offers for financial contracts in the day-ahead market that will be settled at real-time prices.

Allowing these "virtual" bids and offers, such as incremental bids (INCs) and decremental offers (DECs) in PJM, to participate in the day-ahead market promotes entry and other benefits of competition that come with increased liquidity. Virtual bidding promotes price discovery, accurate price formation, market-based redistribution of risk, and an opportunity to price risk in the electricity market.

The existence of a multi-settlement system affects the settlement of FTRs. Any economic dispatch coordinated by the system operator inherently involves using the transmission grid to support the power flows from generation to load. Embedded in the decisions about the dispatch is an assignment of the use of the grid. Hypothetically, if there were physical transmission rights, the holders of these rights would sell them day-ahead to be reconfigured and reassigned to those market participants using the grid under the economic dispatch. In the absence of physical transmission rights, the same reassignment must apply under an FTR paradigm. Inherent in the day-ahead market, therefore, must be the payment for the FTRs at the day-ahead price and charges at the day-ahead price to those receiving day-ahead schedules for the grid, where these schedules are a reassignment of the liquidated and reconfigured rights of the FTR holders.

FTRs in a multi-settlement system are hedges against the prices in the first of a sequence of coordinated dispatches and settlements. In the case of the day-ahead market, this means that FTRs hedge the volatility in day-ahead prices, not real-time prices. In order to provide a

complete forward hedge of the locational differences in real-time congestion costs, a market participant would need to have FTRs and convert these FTRs into a day-ahead financial contract (i.e., schedule) in order to settle in real-time. For example, a holder of an FTR obligation could introduce a day-ahead bilateral schedule for the same amount of transmission between source and destination. Treated as a virtual contract that would be settled at real-time, the day-ahead schedule would hedge the locational difference of real-time LMPs.

This link between virtual schedules in day-ahead, FTRs, and real-time prices is necessary and inherent in the design of electricity markets. The physical analogy would be to conduct a reconfiguration auction for physical transmission rights in the day-ahead. Market participants would sell the long-term transmission rights day-ahead and purchase short-term rights for use in the real-time dispatch. In the absence of a workable system of physical transmission rights the combination of FTRs and day-ahead virtual transactions addresses the same need. The FTR is the long-term right that is settled each day in the day-ahead market. The day-ahead virtual transaction is the short-term right that can be used to continue this hedge to real-time to hedge the difference in real-time locational prices. The intimate connection between FTRs and day-ahead market transactions is an essential part of efficient and workable electricity market design.

Virtual Transaction Types

The PJM paper discusses many purposes and advantages of virtual transactions in the day-ahead market (PJM, 2015). Explicit virtual transactions in the PJM day-ahead market include incremental offers (INCs) that are like a generation offer, decremental bids (DECs) that are like demand bids, and Up-To Congestion (UTCs) contracts that are similar to FTRs.[4]

The inherent connection between FTRs and virtual transactions is most pronounced in the UTC product found in PJM. A similar financial contract can be found in the organized Texas market in ERCOT.[5] The UTC product has many characteristics similar to FTRs. The basic idea is to arrange a virtual day-ahead schedule between two locations with a bid for the maximum payment expressed as the difference in the LMPs between source and destination. This is a generalization in financial terms of a simple price taking physical bilateral schedule. Clearing the pure virtual bilateral schedule would not depend on the price differential in the day-ahead market. By contrast, the UTC product allows a bid for the maximum price at which the transaction should clear. This bid allows the market participant to limit the cost that will be incurred to obtain the real-time hedge.

The PJM report reviews the complicated history of the development of UTC virtual transactions, constraints on their use, and the allocation of uplift charges (PJM, 2015, pp. 16–17).

The bid for a UTC is directly analogous to the bids for FTRs in the forward FTR auctions for financial transmission rights. In the forward auction, bids to purchase FTRs are determined by the market participant, and the bid limits the amount the participant is willing to pay to hedge the difference in day-ahead congestion prices between the source and the destination. The UTC

---

[4]   There are many possible variants of virtual transactions. For instance, a "congestion and loss" product would be the functional equivalent of a UTC sourced at the reference bus. The essence of the main policy issues would not change materially for a wider array of virtual products.

[5]   Day-Ahead   Market   Point-To-Point   (DAM   PTP)   Obligations   and   Options. http://www.ercot.com/mktinfo/dam/.

applies to the real-time congestion and loss prices between a source and destination. The bid on the UTC product limits the amount the market participant would pay for the relevant hedge in essentially the same way as the bid for the FTR.

Like an FTR, a UTC hedges the difference in locational prices. The FTR hedges the locational difference in day-ahead LMPs. The UTC hedges the locational difference in real-time LMPs. The FTR covers marginal congestion costs. The UTC, despite its name, covers both marginal congestion costs and the marginal cost of losses. This UTC formulation is sometimes known as a "spread bid."

A simple price-taking bilateral schedule for transmission between two locations as allowed in the day-ahead market, would also hedge real-time differences in LMPs for both losses and congestion. The advantage of the UTC is that it includes a bid that permits a limit on the price differential paid day-ahead to obtain the real-time protection. In effect, a pricing taking bilateral schedule between two locations is equivalent to a UTC between those same locations with a bid price so high (infinite) that it always clears. It is not evident why the appropriate implicit bid on the bilateral is effectively infinite, while the current PJM market rules limit a bid on a UTC to between zero and ± 50.

The ability to roll FTRs forward from the day-ahead market to real-time to hedge real-time spot prices is materially facilitated by allowing a flexible UTC product. This would improve the ability to arbitrage locational differences and support price convergence between day-ahead and real-time, resulting in more accurate price formation.

Markets and Virtual Transactions

As explained in the PJM report, virtual transactions provide a means to improve convergence between day-ahead and real-time prices. Virtual transactions add to the liquidity in the day-ahead market to facilitate settlement of longer term forward contracts arranged outside the organized dispatch. This same increased liquidity helps to moderate or eliminate the ability to exercise market power. Virtual transactions provide market participants with hedges to reduce price variation for real-time settlements. PJM provides examples of uses of virtual transactions to provide a mixed portfolio of day-ahead and real-time transactions, as in mitigating the risks for generators that might trip offline in real-time.

The PJM perspective credits a range of benefits to virtual transactions including reducing the cost of commitment and dispatch and the risk hedging benefits of forward contracts. By comparison, the analysis of (Parsons, Colbert, Martin, & Mastrangelo, 2015) takes the view that the effects on dispatch are the only relevant benefits of virtual bids and offers. From this perspective, without changing the dispatch, virtual bidding creates "parasitic profits" that extract money from the market. If there were no risk-averse participants in electricity market, this parasitic-profits conclusion might appear more relevant under the assumption of no beneficial changes in the dispatch. But with this same risk-neutral perspective, the need for forward hedging contracts would also vanish. In the real electricity system, risk management is a material concern, and there are real benefits in mitigating risks through forward contracting (Bessembinder & Lemmon, 2002). Although it may be difficult to quantify the risk hedging benefits to compare with the expected cost of commitment and dispatch, market participants can express their willingness to pay to achieve both types of efficiency. Hence, unlike the

assumption in (Parsons et al., 2015), the evaluation of the benefits of virtual bidding involves risk hedging and price formation, as well as the impact on the commitment and dispatch.

An important impact on commitment and dispatch is the effect day-ahead virtual transactions can have on the incentives for physical generation and load to participate in the day-ahead market. For some physical market participants, there may be no incentive to participate in the day-ahead market. For instance, renewable energy suppliers that are paid a fixed contract price for their actual real-time production may be operating under a contract where day-ahead commitments actually increase their risks. Importantly, PJM elaborates on the role of virtual bids and offers to compensate for physical generator and load quantities missing from the day-ahead market.

> "On average, fixed demand bids in the Day-Ahead Market account for about 95 percent of the load forecast for the next operating day. On a peak load day where the real-time load is about 150,000 MW, five percent of the load is 7,500 MW which is equivalent to about seven nuclear plants. On a percentage basis it is small but in terms of real megawatts it is substantial. Without some form of virtual trading, this amount of load could go un-procured in the Day-Ahead Market leading to discounted prices and inadequate resource commitments" (PJM, 2015, p. 23).

The PJM analysis emphasizes a connection between virtual offers and the resulting physical commitment and dispatch, placing this as the primary, or perhaps only, purpose and test of benefits for virtual transactions.

> "Driving the day-ahead commitment closer to what is needed in real time to maintain system reliability is an important function provided by virtual transactions. In order for virtual transactions to accomplish this function, they must impact the scheduling and commitment of the physical resources on the system. When this is done in a direction leading to a day-ahead resource commitment that more closely aligns with real-time needs, market clearing prices will reflect this and the transaction will be profitable. When the opposite occurs, the transaction will not be profitable and, therefore, the market participant is incentivized not to submit the same transaction again.

> If there are persistent scenarios found where virtual transactions drive physical unit commitments in the Day-Ahead Market that are different than real time and yet the transactions are still profitable, or, in cases where virtual transactions are cleared that are not meaningfully impacting the day-ahead resource commitment yet are extracting profits from the market, not only is there no value added, but the transactions are actually detrimental to efficient market operation" (PJM, 2015, p. 18).

If the reference to "value added" refers narrowly to efficient dispatch then the latter statement is true by definition. If the argument is that dispatch efficiency is the only value added, then the PJM position repeats the error of ignoring all the risk hedging and price formation benefits of virtual transactions.

Although this connection to the real-time physical dispatch seems a plausible criterion on its face for evaluating the merits of virtual transactions, the actual situation is more nuanced. The form and purposes of the day-ahead market and the use of virtual transactions to affect "scheduling

and commitment" invoke some issues – especially the impact of uncertainty – that are passed over in the PJM report but would be important in evaluating the contributions of virtual transactions.

The design of the real-time market as a bid-based, security-constrained, economic dispatch with locational prices follows directly from the operational practices of system operators and the requirements for providing open access and non-discrimination in the use of the transmission system. If there were no uncertainty, the design of the day-ahead market would follow naturally to simply replicate the same structure and design as in real-time to maintain consistency between the two markets. With no uncertainty, a reasonable test of the performance of the day-ahead market, including the use of virtual transactions, would be the extent that the "scheduling and commitment" "aligns with real-time needs."

In the presence of uncertainty about real-time needs, however, the design and criteria for evaluating day-ahead market outcomes are not so straightforward and may not include the expectation of outcomes identical to the real-time market. For instance, there is an important distinction between "scheduling" and "commitment" in the day-ahead market. It is the day-ahead commitment of resources that have long lead-times, e.g., units that require time to ramp up to its minimum load level of output, that has a physical impact that can affect the outcomes in real-time. By contrast, day-ahead energy schedules, including dispatch of all available units, can change after the day-ahead and is not necessarily affected by day-ahead choices. In the stylized framework, the "commitment" is fixed under uncertain conditions (e.g., forecasts of real-time load) in the day-ahead, but the different eventual outcomes of this uncertainty in real-time can and will be accommodated in the real-time dispatch as the uncertainty is resolved. This distinction is important but it is not maintained in the PJM analysis, where the discussion refers to "commitment" but the examples all pertain to energy dispatch. Hence, the illustrative examples may or may not indicate problems with the day-ahead market.

It might be natural to formulate the day-ahead market optimization problem so as to minimize the total expected cost of the actual commitment and real-time dispatch of supply to meet forecasted load. This expected-value framework appears to be implicit in the PJM discussion and criteria for evaluating the contribution of virtual transactions. However, this framework is quite different from the actual approach taken in present designs of day-ahead markets. The expected-value approach would require the day-ahead analysis to describe and incorporate the probability distribution for possible realizations of the real-time market, and decision variables would be the commitments for all long lead-time actions that would affect the real-time dispatch costs. In formal terms, this would be a sequential stochastic optimization problem, not the security-constrained deterministic model PJM employs. Formulating and solving such stochastic optimization problems is an active area of research, but it is not what the current electricity market designs attempt or could achieve. The information and computational challenges are daunting, and the day-ahead results would be quite different (Bjørndal, Bjørndal, & Midthun, 2016).

The design of the PJM day-ahead market does not explicitly account for the uncertainty affecting the real-time dispatch. Given the bids and offers, including virtuals, the PJM day-ahead market finds the corresponding security-constrained least cost unit commitment and dispatch. Some market participants are restricted to cost-based offers. Others, especially including virtuals, can submit market-based bids and offers. The system operator solves for the optimal unit commitment, dispatch solution and associated locational prices, along with any necessary uplift

payments. If the uncertainty impacting the real-time dispatch is small, then the resulting dispatch solution for the day-ahead market may be close to the eventual real-time dispatch.

Without any virtual bids and offers, the resulting day-ahead commitment and dispatch could, however, deviate significantly from the requirements in real-time. For example, if only bids and offers backed by physical facilities were allowed, and some of the load did not participate in the day-ahead market, as illustrated in the PJM analysis, this could result in an inadequate commitment and dispatch to meet real time load, just as illustrated and documented in the PJM analysis. Virtuals are an important part of resolving this deficiency of sole reliance on the physical offers to reach a day-ahead commitment that reflects market expectations of real-time conditions.

In the presence of uncertainty the distinction between physical and virtual offers is not sharp; i.e., if a day-ahead schedule differs from the real-time quantity it will not always be clear whether this has arisen because of uncertainty at the time of the day-ahead scheduling process, or because of the intentional submission of a day-ahead schedule inclusive of some virtual quantity intended to be settled financially in real-time. As a practical matter, offers from physical generators and bids from physical loads can deviate significantly from the actual real-time outcomes. Differences between real-time conditions and what was expected day-ahead caused by uncertain outcomes appear identical to virtual transactions in the ISO settlement accounting. The principal distinction of formal virtual transactions is that they can be clearly recognized in the protocols for the bids and offers presented for the day-ahead market. The deviations between day-ahead schedules and real-time outcomes are known in advance for explicit virtual transactions. This is unlike the uncertainty associated with physical forecasts where the deviations cannot be known in advance.

Under uncertainty, and even with the best efforts of all the physical and virtual offers, the connection between the day-ahead and real-time dispatch can be complicated. For example, as discussed below, even the simplest case of day-ahead markets can create outcomes that deviate from the expected value of the real-time dispatch. Setting aside the hedging activities, one role for virtual bids and offers is in improving price formation. With no market power, no transaction costs, full information, and risk neutrality, virtual bids and offers should match the expected locational prices in real-time. The standard analysis indicates that day-ahead prices at a particular location will be driven by expectations about real-time prices at that the same location. Any sustained deviation from this condition would produce an arbitrage opportunity for profitable entry through appropriate financial contracts.   If day-ahead prices are above expected real-time prices, it would be profitable to offer a financial contract for sale of energy in the day-ahead that would be balanced by the profitable purchase of that same energy in the real-time market.   Similarly, if the day-ahead price is lower than the expected real-time price, then it would be profitable to purchase energy day-ahead through a financial contract that would be in turn balanced by a profitable sale in real-time. The actual realized price in real-time may differ substantially due to the uncertain real-time outcomes. But the day-ahead price should on long-run average be close to the real-time price. Of course, with the high volatility of electricity prices, the long-run could be quite a long period of time over which to measure convergence of prices between the two markets, thereby reflecting accurate price formation. Nonetheless, there is a close connection between the expected prices in the two markets, day-ahead and real-time.
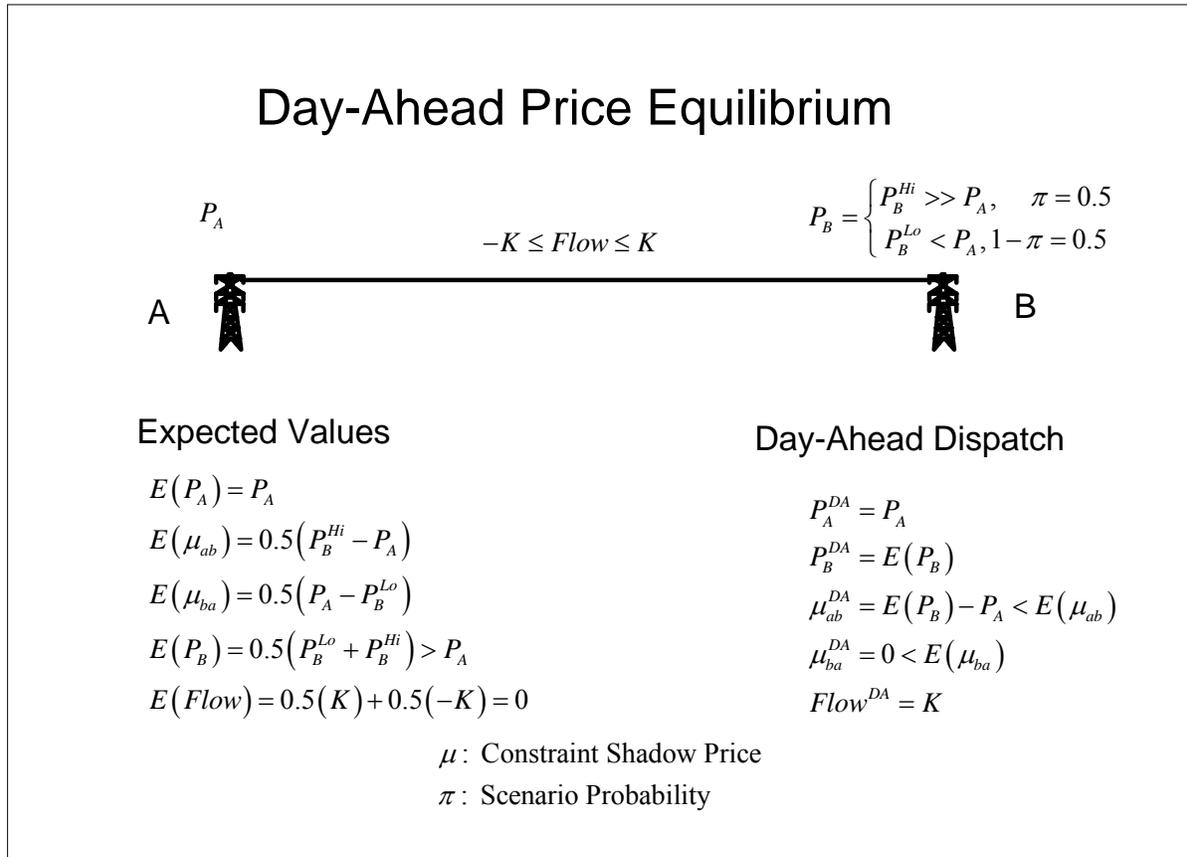
The connection, however, is neither obvious nor inherently easy to analyze in the general case. In particular, deviations between real-time and the day-ahead dispatch solution can arise for

many reasons. And deviations may not go hand-in-hand with added uplift costs. This situation is especially important for virtual transactions which always produce a deviation between day-ahead and real-time quantities.

Deviations Without Uplift or Efficiency Effects

The equilibration of day-ahead prices and expected real-time prices does not mean that expected dispatch in real-time will be the same as the dispatch day-ahead, nor does it imply that the same transmission constraints will be binding or have the same congestion costs.

A simplified example illustrates the general point. Consider the case in the accompanying Figure for "Day-Ahead Price Equilibrium." The model here abstracts from the full details to focus on the key parameters. There is no commitment decision; all the implied generation is available for the real-time dispatch at no added cost. Hence, there is no question about ultimate dispatch efficiency, and no need for uplift payments.

Day-Ahead Price Equilibrium

P_A

$$-K \leq Flow \leq K$$

$$P_B = \begin{cases} P_B^{Hi} >> P_A, & \pi = 0.5 \\ P_B^{Lo} < P_A, 1-\pi = 0.5 \end{cases}$$

A   B

Expected Values

$$E(P_A) = P_A$$
$$E(\mu_{ab}) = 0.5(P_B^{Hi} - P_A)$$
$$E(\mu_{ba}) = 0.5(P_A - P_B^{Lo})$$
$$E(P_B) = 0.5(P_B^{Lo} + P_B^{Hi}) > P_A$$
$$E(Flow) = 0.5(K) + 0.5(-K) = 0$$

Day-Ahead Dispatch

$$P_A^{DA} = P_A$$
$$P_B^{DA} = E(P_B)$$
$$\mu_{ab}^{DA} = E(P_B) - P_A < E(\mu_{ab})$$
$$\mu_{ba}^{DA} = 0 < E(\mu_{ba})$$
$$Flow^{DA} = K$$

μ : Constraint Shadow Price
π : Scenario Probability

The transmission system is the same day ahead and in real time. The conditions at node A imply a net supply-and-demand condition that yields a constant price, P_A under all real-time conditions.

The conditions that will appear at node B in real-time are uncertain with two possible outcomes with equal probability. The load and generation conditions could produce a high real-time price,

P_B^{Hi} at B. Or the conditions could produce a low price, P_B^{Lo}. We assume that $P_B^{Hi} >> P_A$, so in this case the real-time dispatch will involve transmission flow from A to B at the upper limit of the constraint on the transmission line, K, because under these assumptions cheaper supply is available at A than at B. In the other real-time case, where $P_B^{Lo} < P_A$, the transmission flow will be from B to A and the active constraint will be the lower limit on the line.

At the time of the day-ahead dispatch, assume that there are sufficient virtual bids and offers to ensure that the day-ahead locational prices equal the expected values for the real-time. Everything works perfectly as designed in terms of bidding and pricing. But the results also show that the day-ahead quantity dispatch does not equal the expected value of the real-time dispatch quantities. For example, because of the particular asymmetry in prices employed for the example, the day-ahead transmission flow is K megawatts, from A to B, even though the expected real-time transmission flow is zero.   The associated constraint shadow prices also do not equal the expected value of the constraint prices from real-time. The day-ahead shadow price on the transmission constraint from B to A is $\mu_{ba} = 0$, but the expected value of the real-time shadow price is $E(\mu_{ba}) = 0.5(P_A - P_B^{Lo}) > 0$. Adding further details about the loads and generation, we would also find a difference between the day-ahead load and generation dispatch quantities versus the corresponding expected real-time quantities. Apparently accurate price formation to achieve price convergence between day-ahead and expectations of real-time is not the same thing as convergence of the day-ahead dispatch quantities to equal the expected real-time dispatch quantities.

The reason this occurs, as discussed by the PJM report, is that the problem in the day-ahead model can be quite different than in real-time. With enough uncertainty about the real-time, essentially any and all constraints could have non-zero expected shadow prices as anticipated day-ahead. There is an upper bound (i.e., defined by the number of nodes minus one) on the number of constraints that can be binding in the day-ahead dispatch, but the number of binding constraints could be quite large in the day-ahead solution compared to the relative handful of actual binding constraints in any particular real-time dispatch. The result is that the quantity dispatch in the day-ahead will be different than the real-time dispatch, even in expected value terms.

Hence, quantity deviations between day-ahead and real-time are inevitable. These deviations apply by definition to virtual transactions and by necessity to physical schedules. As shown in this example, the deviations can occur without having any dispatch efficiency effects and without any uplift requirements. The basic principles illustrated through this simplest of cases extends to a more general model with many network interactions and security constraints.

Uplift Without Deviations or Efficiency Effects
Uplift can arise for many reasons, and may not result from day-ahead versus real-time deviations, and may not have any efficiency effects. The simplest case appears when the day-ahead dispatch involves unit commitment costs. Even without any uncertainty, the lumpiness of the unit commitments costs can give rise to a situation where there is no market clearing price in the day-ahead market that fully supports the solution (Gribik, Hogan, & Pope, 2007). The unit commitment and economic dispatch found cannot be maintained without some side payments to compensate bidders that otherwise would have an incentive not to follow the dispatch at the market prices. This approach preserves efficiency but creates an uplift requirement.

Introducing uncertainty would not change the underlying premise that an uplift payment is required without considering deviations between day-ahead and real-time quantity dispatch. Including virtual transactions in the analysis would have the effect of smoothing the day-ahead commitment and dispatch problem and reducing the required uplift payments.

Deviations With Uplift But No Efficiency Effects

The elements of a case with both deviations and uplift, but no efficiency effects, can be seen through a slight modification of the equilibrium analysis in the simple case of the Figure "Day-Ahead Price Equilibrium." Suppose now we assume that the high-priced generator at node B comes with a unit commitment cost and this generator is needed to meet load in the high-priced real-time case.

Since this generator's variable cost is above the expected real-time price, this high-priced generator would never be committed in the day-ahead dispatch. In the practice in PJM and other RTOs, this generator would have to be committed in a reliability pass to make sure it would be available in the event of the high demand case in real time. Since the high-priced generator would not be included in the day-ahead dispatch, it would require an uplift payment to cover the cost of commitment.

Given the reliability constraint, the real-time dispatch would always be efficient. Any virtual transactions and some physical transactions would differ from the day-ahead schedule, so there would be deviations. Hence, in this equilibrium analysis, we would have deviations with uplift and no efficiency effect.

Equilibrium Analysis

The analysis of costs and benefits of virtual transactions cannot be formulated or found from looking at individual dispatch outcomes as in the PJM analysis. And there is no single connection between deviations, uplift and efficiency effects. As illustrated by the simple examples, equilibrium analysis is required for a number of reasons. First, many proffered examples of apparent anomalies between day-ahead and real-time show isolated effects that could not persist in equilibrium. To evaluate the effect of virtual transactions it is necessary to find anomalous examples that could exist in equilibrium rather than as an isolated surprise. Second, in the presence of uncertainty, the equilibrium condition should compare the day-ahead dispatch with the expected values of the real-time costs, rather than the day-ahead dispatch with the actual real-time results. It is understood that individual real-time outcomes will deviate substantially from the day-ahead prediction, where the variance of real-time prices is much larger than the variance of day-ahead expectations. Third, the equality between the day-ahead locational price and expected real-time price at the same location does not extend to all the other metrics of the dispatch. In the end, settlements are for energy transactions and these are the prices that equilibrate through day-ahead and real-time convergence. The day-ahead schedule can affect the day-ahead commitment, and day-ahead trading can help improve this commitment. In the absence of uncertainty and with a well behaved problem without lumpy constraints, the quantity dispatch solutions -- i.e., day-ahead and expected real-time -- also would match perfectly. But with uncertainty and with the actual lumpy start-up and minimum load conditions any comparison of the expected real-time and day-ahead dispatches results would require more analysis than has been done.

In this sense, the connection between the analysis and the conclusions in the PJM report is missing. The PJM analysis focuses on examples of day-ahead decisions and particular real-time outcomes. These examples are valuable and instructive in understanding how the electricity market system works. But they are not directed at the essence of the interactions between day-ahead and real-time under uncertainty, and for this reason do not provide a test of the performance of virtual transactions. It remains as a challenge to develop a complete model of all the benefits and costs of virtual transactions.

The most obvious test of the impact of virtual transactions on prices would be to measure the degree of price convergence between day-ahead market solutions and real-time prices over a long period with and without the inclusion of virtual bids, to assess whether the inclusion of the virtual bids results in more accurate price formation. This would be complicated by any concurrent change in transactions costs and uplift allocations, but the standard would be a small difference in prices. The annual reports on the state of the market summarize such analyses as a regular feature and find very close convergence between prices. "Markets fluctuate continuously and substantially from positive to negative. The difference between the [locational] average day-ahead and [locational average] real-time prices was -$0.93 per MWh in 2014 and -$0.73 per MWh in 2015. The difference between average day-ahead and real-time prices, by itself, is not a measure of the competitiveness or effectiveness of the Day- Ahead Energy Market" (Monitoring Analytics, 2016, p. 17). The emphasis here should be on "by itself," but it is clear that this price convergence is an essential part of the story.

It is a greater challenge to separate out the independent effect of explicit virtual bidding in supporting this price convergence. However, some attempts to exploit natural experiments to analyze the impact of virtual bidding reached the conclusion that introduction of explicit virtual bidding is an important contributor to improved price convergence (Jha & Wolak, 2015) (Li, Svoboda, & Oren, 2015).

A related focus of the evaluation of virtual bidding would address the contribution to the many other objectives and benefits outlined by the PJM report. Some of these features are necessary as a theoretical matter. For example, something like the UTC virtual transaction is necessary for a market participant that wants to settle an FTR against real-time prices. This was and is an important part of the basic market design. Others, such as the value of liquidity and hedging, present added difficulties in measuring the benefits, but the PJM analysis suggests why these benefits could be substantial.

The impact on the "commitment and schedule" is important, but the evaluation would require more to identify the costs and benefits with and without virtual trading. The examples that emphasize these dispatch impacts tend to focus on cases where the underlying problems are "modeling discrepancies" (PJM, 2015). See also (Parsons et al., 2015). It is to be expected that when market design features create unintended disconnects from reality, real-time dispatch and day-ahead dispatch would create unintended differences in prices that would be a magnet for virtual trading. The probable effect of the virtual trading would be to induce more accurate price formation, but this may not eliminate the impact of the underlying market defect. Whether this is a problem with virtual trading or with market design is subject to debate. The first choice would be to fix the market design. If the market design defect is unavoidable, then the least disruptive solution should be pursued. If this requires limits on virtual trading, there should be some identifiable connection between the proposed cure and the underlying disease.

Going further, to do a full-blown efficiency analysis of the impact of virtual trading, would be an interesting research project and a challenge. However, as demonstrated by the simple example above, posing the question in the proper light would not be easy. Hence, while PJM is correct that "[v]irtual transactions that impact prices in the Day-Ahead Market but that do not result in physical resource commitments that more closely reflect what are actually needed in real time do not result in more efficient market operation," the available studies have not gone very far down the path of analyzing this criterion (PJM, 2015, p. 17). How would we describe the efficiency of the outcome with and without virtual trading? It is not as simple as showing the dispatch would be the same, even on average. How would we quantify the effect of different unit commitment results under uncertainty, as separate from the dispatch? Answering these questions would be an important part of any full analysis of the empirical costs and benefits of virtual trading.

As of yet, empirical cost benefit questions have not even been posed succinctly to deal with all of the commitment and dispatch costs and related benefits. The work of (Jha & Wolak, 2015) uses a regression approach to conclude that the introduction of virtual bidding in California improved price convergence, and reduced fuel cost slightly, but there is no attempt to address other costs and benefits. See also (Woo, Zarnikau, Cutter, Ho, & Leung, 2015) on price convergences with virtual bidding and increased wind penetration in California. The trading model of (Li et al., 2015) shows improved price convergence, but focusses only on the marginal incentives so as to avoid the complication of addressing changes in commitment and dispatch.

An empirical analysis using the actual dispatch model for the Independent System Operator of New England (ISONE) addresses both price convergence and dispatch performance by comparing the day-ahead and real-time markets with and without explicit virtual bids, assuming all other bids stay the same. Although this does not account for different implicit virtual bidding strategies, the counterfactual modeling approach comes closer to addressing the total cost issues. The authors conclude "that the financial entity participation not only results in reduced DAM-RTM price deviations but also leads to DAM dispatch results that are 'closer' to those of RTMs. Therefore, financial player participation improves the ability of the RTO to ensure system security" (Güler, Gross, Litvinov, & Coutu, 2010, p. 294). On balance, therefore, this study concludes that virtual bidding improves both price formation and physical system operations.

Dispatch Interactions

In principle, FTRs, UTCs, and all types of virtual transactions could be constructed through private arrangements outside the organized market administered by the RTO. Any consenting parties could write a contract that settled against PJM's LMPs, whether for day-ahead or real-time markets. These derivative contracts would be subject to market oversight, but they would not be seen by the system operator. These financial contracts would not be considered in the commitment and dispatch.

The difficulty with leaving derivative contracting to the private market arises again from the absence of physical transmission rights and the inability of a bilateral market for transmission to address the strong interactions in the flow of power on the grid. While it is true that anyone can write a contract that looks like an FTR, UTC, or virtual bid, only the system operator can support a set of contracts that fully respects and utilizes the limited capacity of the grid. The need for coordination through the system operator is most obvious in the case of FTRs, but the same principles apply to any financial products that depend on prices that reflect the actual flow of

power on the grid. In economic terms, the transaction costs of organizing an efficient commitment, dispatch and hedging configuration are relatively small when conducted through the system operator, and much larger or even prohibitive when left to the bilateral market.

The advantages of including these financial transactions in forward auctions or forward markets are clear. The clearing and scheduling of accepted transactions can reflect the real limits of the grid, with all of its strong and complex interactions. This expands the set of feasible transactions and should both increase efficiency and reduce risk. The flip side of this inclusion in coordinated auctions by the system operator is that the financial contract bids and offers can affect the commitment and dispatch choices of the system operator, at least to some degree, and capture the benefits found by (Güler et al., 2010).

The degree of interaction with the physical commitment and dispatch depends on how the virtual and other financial transactions are represented in the forward markets. For example, FTRs that are only for congestion cost create no direct impact on losses or forward energy contracts. The award of FTRs may have some indirect effects on the development and availability of long lead-time generation facilities and loads, but the impact on the day-ahead or real-time dispatch would be de minimis. Likewise, in the real-time dispatch, financial bids and contracts are no longer part of the solution and it is only real physical conditions that determine the final dispatch and prices. In the day-ahead, the issue is more complicated and the degree of interaction depends on how the system operator models UTCs, virtual bids and other financial contracts.

The possibility of interaction between financial contracts and the market dispatch raises the concern that the financial bids could be used to manipulate the market (Haas, 2009). A difficulty arises in considering the proper test of manipulation. The appropriate counterfactual in testing manipulation would be an equilibrium solution without the financial bids. Consider the simplifying assumption of complete information, with a common probability distribution characterizing the uncertainty of real-time prices, and risk neutral financial participants. Then the day-ahead financial participants would produce financial bids at the common expected real-time price. With no restrictions on entry, any financial bids that deviated from the common expected real-time price would either lose money or would not clear. Under this condition, the strictly financial bidders, who could not affect the real-time price, could not affect the day-ahead price. Hence, the ability to affect day-ahead prices must depend on some combination of circumstances violating these simplifying assumptions, such as restrictions on entry, external limits on participation by risk neutral financial traders, transactions costs including collateral costs, analysis costs, or a more complicated information setting (Lo Prete & Hogan, 2014).

From this perspective, and as discussed in greater depth below, limitations on virtual bids and financial transactions work in the wrong direction. Expanded liquidity and ease of entry would improve the operation of the market and create a closer approximation of the idealized competitive day-ahead market. Given the benefits of coordinated day-ahead and real-time markets and expanded opportunities for hedging this provides a strong case that limitations on financial bids should be avoided or at least face a strong burden of justification.

Settlements and Cost Allocation

The principles of cost causation focus on achieving the benefits of efficient price signals. Prices should be set to reflect costs on the margin. If the costs are well behaved, the balance induced by equating price and marginal cost supports the efficient economic outcome. With the simplest

representation of supply and demand, including increasing marginal costs and decreasing marginal benefits, the efficient solution establishes a welfare maximizing market equilibrium. This is what motivates economic dispatch as the core feature of the electricity market model. And this idealized model would not create residual costs, with or without virtual transactions.

Energy Uplift for Residual Costs

In the presence of fixed costs, joint products, and other departures from the assumptions of a simple competitive market structure, prices that equate marginal benefits and marginal costs may not cover the full costs of production. This raises the question of how to allocate costs using principles that go beyond the simplest application marginal cost pricing to accomplish a cost allocation conforming to the cost causation principle. In the electricity market, for example, the problem arises when there are startup costs, minimum run times and other constraints on generators which imply that the efficient, least-cost solution may not be compatible with any given set of prices for outputs of the facility. In other words, there is no unique set of internally consistent prices that would send price signals that would lead independently profit maximizing (or cost minimizing) market participants to reach, through the "invisible hand," the least cost dispatch solution. To incent market participants to produce or consume quantities corresponding to the least cost solution, there is a need to allocate the costs that cannot be covered at the prices determined by the variable costs.

These residual costs often go under the name of energy uplift. The uplift charges can vary significantly. For example, in PJM "[t]otal energy uplift charges decreased by \$646.3 million, or 67.3 percent, in 2015 compared to 2014, from \$960.5 million to \$314.2 million" (Monitoring Analytics, 2016, p. 20). The allocation of these costs can be problematic and is a subject of continuing review in PJM.[6]

In the case of PJM, an example is the Balancing Operating Reserve (BOR) charge that covers a variety of startup and related costs that might exceed the revenues obtained at spot prices for the output schedule for a generator in the day-ahead market. The BOR costs reductions were the major source of the reduction of overall energy uplift charges in PJM in 2015. However, they remain the largest component of the total uplift cost. The current PJM method for allocating these costs is based on real-time deviations from day-ahead schedules. "PJM calculates for each Operating Day the Balancing Operating Reserve charges to deviations by allocating the total cost of Balancing Operating Reserve for deviations on a regional basis to each customer account based on their daily share of the sum of the total hourly deviations in each region (RTO, East, and West)" (PJM, 2016, p. 37).

The energy uplift charge allocation question, like similar applications such as the Revenue Sufficiency Guarantee (RSG) in the Midwest System Operator, produces an often confusing and circular conversation (Hogan, 2008). Although this BOR allocation assigns different costs to different actions, the allocation of these costs based on deviations from day-ahead schedules does not arise from any fundamental model. As shown above in the simple equilibrium analysis, deviation from day-ahead schedules should be expected in any model with significant uncertainty, and deviations do not necessarily give rise to uplift costs.

---

[6]   The March 2016 meeting of the PJM Energy Market Uplift Senior Task Force reviewed the PJM report on virtual bidding; http://www.pjm.com/committees-and-groups/task-forces/emustf.aspx. See also (FERC, 2014).

The lack of an explicit model to deduce cost causation implies that the allocation method is more an administrative compromise than the product of a principled analysis. The difficulty is fundamental. The underlying assumptions behind the cost causation argument, with the link between marginal costs, prices and incentives, do not apply to all possible costs. The existence of discontinuities in the generator costs structure—the generator is on or off; the startup cost is incurred or it is not incurred, without any intermediate possibility—means that some of the costs may not be connected to marginal changes in output. A marginal change in output may have little or no impact on uplift costs, even though the total costs of starting and maintaining the active generator may be large. If prices equal to marginal costs do not cover the full costs, then the residual costs need to be covered by appeal to some other principle (Gribik et al., 2007).

In other words, appeals to cost causation principles to allocate properly defined residual costs are self-contradictory. The very definition of the residual costs is for that part of the total costs that is not amenable to attribution at the margin. If the cost could be connected at the margin to changes in load or generation, it should be part of the energy price. But, by definition, there is no cost causation allocation available for the residual costs left over after accounting for cost causation. Although there are many appeals to cost causation claims for uplift allocation, closer examination of the claims leads to a recognition that cost causation argument is at least difficult to apply in practice and, more importantly, may be flawed in principle (ISO_New_England, 2016, p. 2).

In the allocation of joint or residual fixed costs, without the connection to cost causation at the margin, there is an inherent arbitrariness to the allocation. If there were no consequences of this cost allocation decision in terms of choices of market participants in the market that could affect market efficiency, an administrative compromise would present no policy problem. With multipart tariffs, the costs may be included in separate charges that differ from and are in addition to the marginal cost charges for output. The basic principle would be to allocate the costs in a way that would have the least adverse impact on the choices made in the market. For example, allocating the uplift costs to network connection charges would be better than adding to a so-called "uplift" charge on load billed per megawatt-hour. If an uplift charge is necessary, it should be allocated to the least price-responsive loads. If a non-discriminatory uplift charge is required, it should be spread across the widest possible base of loads that cannot bypass or avoid the charge.

Efforts to avoid this logic for cost allocation, by finding a cost causation connection for the residual costs left over from cost causation allocations, can only mislead. This is especially true for charges like the BOR, to the extent that they are residuals after already taking into account settlements given the spot prices of output. The magnitude of the BOR is not only a function of generation total cost; it is also a function of the spot pricing rule. For example, the PJM market monitor recommends correcting spot pricing defects that contribute to uplift charges (Monitoring Analytics, 2016, p. 21). By definition, small changes in output provide no guidance for the cost allocation, and examples of large changes in the market big enough to create a correspondingly large change in costs inherently require arbitrary decision about joint effects, not independent marginal decisions. In effect, the cost allocation problem for residual costs inherits the problems of lumpiness and joint costs that give rise to the existence of residual costs in the first instance. It is easy to fall into the trap of seeking a cost causation allocation, but it is the wrong path to follow.

The problem of uplift allocation is especially important in dealing with transactions like FTRs, virtual INCs, DECs, and UTC transactions, financial contracts that do not imply or produce physical delivery or load in the real-time electricity market. By design and construction, these financial contracts will be settled at prices determined in the spot market, but the observed quantity will always be zero in the real-time physical flows. The underlying economics of the financial contract are driven by the expected value of the real-time price that will apply to the financial settlement of the contract. By design, the quantity deviation between day-ahead and real-time for the financial contract is the full quantity, and for a competitive bidder there is no connection between this deviation and the appropriate economic analysis of how much (in dollars) to bid. In other words, the system operator knows in advance that the schedule of the individual bidder will "deviate" between the day-ahead market and real-time market because this "deviation" is an inherent characteristic of the virtual transaction. Further, there is not necessarily any deviation between the aggregate day-ahead market schedules and real-time schedules because the virtual bid may proxy for supply or demand or power flows that are present in real-time but otherwise would not be represented in the day-ahead market Hence, allocating costs to these virtual contracts based on deviations of the individual bidder does not have a foundation in the economics of a competitive bid and creates perverse incentives to avoid virtual transactions. Any added charge to the cost of settling the virtual contract creates a wedge between the expected real-time price and the day-ahead price, reducing the incentive and the ability to promote convergence of day-ahead and expected real-time prices, resulting in less accurate price formation. Uplift allocation to any virtual contracts has material consequences that work at cross purposes to good electricity market design.

A purpose of these contracts is to hedge or arbitrage in the face of uncertainty about prices. With no uncertainty there would be no demand for these hedges. And without risk aversion that gives rise to hedging, there would be no need for these contracts. In the real world, with uncertainty and risk aversion, these financial contracts improve the operation of markets in the many ways described in the PJM report. However, the fundamentals dictate that the supply and demand for the financial contracts would be very sensitive to transactions costs, including any assignment of the residual uplift costs.

Virtual Trading and Financial Contracts
To avoid complicated simulations or examples of calculation of uplift charges, it helps to step back and think about the market conditions that give rise to the residual costs in the first instance, and the role of financial contracts in these markets. For example, much of the intuition about cost causation comes from an implicit connection to simple real-time markets for electricity in the so called "day one" structure that includes a real time spot market with one-part offers and bids for supply and demand. The offers and bids describe textbook supply and demand curves, and economic dispatch produces an efficient equilibrium with market clearing prices. Under these simplifying assumptions, we do not have lumpy decisions, the market clearing prices would cover the costs, and no residual costs remain to allocate to make suppliers whole for their bid costs. Financial contracts could be arranged ahead of time in the bilateral market, but the actual dispatch would depend only on the final offers and bids provided to the system operator. The spot deliveries would differ from the financial contracts, so there would be substantial deviations for these bilateral contracts. But the deviations would be accounted for through the bilateral transactions and not known to the electricity system operator. In this simplified world, there would be no energy uplift cost to allocate.

Setting aside the problem of the lead time for starting up, we could modify this simple market to include multipart offers and bids to reflect startup costs, minimum run times and the other complications of the actual structure of costs for physical generation supply. Assume for the sake of discussion that bilateral financial contracts are not organized through the system operator, but are strictly in the external financial market. With the introduction of multipart bids the lumpiness of the economic dispatch solution would create the problem of residual costs that might not be covered by spot prices. Hence, cost allocations, such as for the energy uplift charge, would be necessary. It would appear that the uplift charges arose not because of deviation from day-ahead schedules, but because of the fixed costs, minimum run levels and other constraints of generation. Allocation of the costs to the generators that caused the costs would be circular, because it is the under-recovery by these same generators that we seek to cover with uplift payments to incent them to follow system operator commitment and dispatch instructions. In this hypothetical case the financial contracts, which are still bilateral and would not even be known to the system operator, would not be the cause of the uplift charges, and deviations of the financial contracts would not be used for cost allocation.

The lead time associated with startup and related schedules drive these "day one" markets towards the "day two" structure that includes a day-ahead market operated by the electricity system operator. The day-ahead market and associated schedules facilitate operations by allowing market participants to exploit a wider range of commitment and dispatch decisions coordinated and sequenced over a 24 hour period. The expanded opportunity set, coupled with the a closer connection of the day-ahead market schedules to the actual physical characteristics of load response and physical generation operation, creates the potential for a more efficient unit commitment coordinated by the ISO that would reduce the total cost of meeting load. In principle, this day-ahead market could operate without including purely financial contracts at all. With no virtual bidding, including exclusion of implicit virtual bids, accurate price formation will become a more complicated issue. But set that aside for the moment. The resulting commitment and dispatch, reflecting the lumpiness of the unit commitment decisions, would give rise to the residual cost allocation requirement such as for the uplift charges. There could still be bilateral financial contracts, but there would be no requirement that the system operator even know about these contracts, which could be settled separately against the published prices. From the perspective of the system operator, there would be no observed deviations for the financial contracts and no cost allocation to the financial contracts. The residual costs would exist, but would not be caused by the financial contracts.

Once we have the day-ahead market structure in place, it becomes clear that there would be substantial advantages to the market as a whole to include financial contracts under the purview of the organized electricity markets. This is a main theme in the PJM report. Virtual bidding provides the advantages of additional entrants in the day-ahead market, better price convergence, better price formation, increased liquidity for hedgers, and a natural way to resolve long-term financial transmission rights that address the locational differences in prices. This integrated market would be impossible to fully replicate through a strictly bilateral financial market outside the organized market.

Viewed in this way, the focus is on overall system performance, not on residual costs alone. The residual costs arise independently of the financial contracts. The assembly of FTRs, INCs, DECs, and UTCs included in economic commitment and dispatch bring an added benefit to the market interaction, improving efficiency and lowering overall costs. Movement of financial

contracts into the organized market run by the system operator can affect operations in ways that are beneficial to the system. From this macro perspective, the presumption is that including explicit virtual bidding in organized markets will improve aggregate performance. At a minimum, any alternative presumption should bear the burden of proof.

However, the movement of financial contracts from the bilateral market into the organized market also makes the deviations of the financial contracts from the real-time market visible. Thus follows the conundrum. The residual costs in energy uplift charges arise because of the lumpiness inherent in the multipart offers and bids for generation and load. There may be some interaction between the financial contracts and the commitment decisions, but these interactions are intended to reduce total costs, not add to the total costs. Furthermore, it is the total costs, commitment and real-time dispatch, that should be the focus of any cost analysis of the impact of financial transaction on the day-ahead market and not the organization of these costs in different accounts such as energy uplift.

The change in commitment decisions may have some impact on dispatch and prices, and these changes may increase (or decrease) the allocation of costs into the residual category.[7] But these changes to energy uplift charges at their root are not caused by the deviation of individual market participant real-time schedules from their day-ahead contracts. In this important sense, financial contracts do not cause residual costs. In equilibrium, and on average, including financial contracts should improve the aggregate efficiency of the system. Virtual bidding cannot overcome all defects in market design, but virtual bidding offers real benefits in the aggregate.

Consequences of Uplift Allocations
These benefits of coordinating financial contracts would be threatened by any increase in transaction costs or allocation of residual costs to the financial contracts. In the first instance, the parties always have the option to move back into the external bilateral market where the transactions costs are higher but the deviations are not available as an indicator for residual cost allocations (Deng & Oren, 2006). Even worse, financial participants might withdraw from the market altogether and eliminate the efficiency gains of the more transparent and liquid forward market they engender, operated with explicit recognition of day-ahead conditions and transmission interactions.

This high level perspective provides a view that would be lost by trying to do *ceteris paribus* simulations of changes in financial offers and bids, to calculate the impact on energy uplift charges. This simulation approach would mislead. The better perspective is that the residual costs arise because of the lumpiness of the technology and the need for multipart bids for day-ahead and real-time dispatch. In the absence of financial transactions, there would still be energy uplift charges. The right perspective is that the financial transactions reduce overall costs and provide better incentives for efficient markets, irrespective of the effect on residual costs.

Allocation of residual costs to financial transactions cannot be supported by cost causation arguments. The incentive effects of such allocations are perverse, because even small increases in these transactions costs can have a material effect on the activity of financial participants. By contrast, real load, in real-time, has few alternatives to consumption. Financial participants have

---

[7] See Table 2 in (PJM, 2012b), where the "changes" in day-ahead cost are negative for the inclusion of UTC transactions and positive for inclusion of other virtual transactions (sign conventions explained through an inquiry to PJM). The argument is even more pronounced for the reliability unit commitment phase which allows analysis without the explicit virtual transactions.

many more options. Once we move past the cost causation argument, allocation of costs to the highly price-sensitive financial contract segment with the most options, in order to lessen the residual cost allocation to real load segments which will not change behavior, works in the wrong direction and reduces the overall benefits of organized markets.

In the PJM system, there is a separation of the allocation of uplift according to a rule that involves a judgment about the impacts of particular types of bids (PJM, 2016). The current PJM policy is to allocate some uplift costs to deviations from day-ahead schedules, with no separate treatment of financial transactions, and the remainder of the residual costs to load. The general arguments above about virtual transactions are particularly applicable to the case of UTCs. In the practical implementation of markets there are always approximations inherent in the implementation of security constrained economic commitment and dispatch. In the case argued by PJM, these approximations mean that UTCs do not affect the commitment organized through the multipart bids handled by the system operator (PJM, 2012a). Since they do not affect the day-ahead commitment, UTCs cannot affect the real-time dispatch and costs. Hence, currently no deviation charges are allocated to UTCs. Likewise, since FTRs are strictly financial contracts that are established before the day-ahead commitment and dispatch, the judgment is that there are no residual cost allocations to FTRs in this category. There is a separate issue of FTR allocation of "balancing congestion," related to over allocation of real-time transmission capability, but that is a different story of a design implementation mistake.

An objection to this logic arises because the current uplift cost allocation scheme in PJM exempts UTCs but includes allocations to deviation in other virtual transactions (Monitoring Analytics, 2016, p. 22) (PJM, 2015). The essence of the argument appears to flow not from a concern based on first principles of efficiency or cost causation, but from the asymmetry of the treatment between other virtual transactions and UTCs. Since the UTC is equivalent to a pair of virtual transactions with a linking constraint, why should it be treated differently? Is the justification for exempting UTCs nothing more than the particular simplification chosen for the PJM commitment decision? If there is no cost causation argument and the allocation is intended as a compromise to achieve rough justice, does the cost allocation have any effect on the market?

From the main argument in this paper, it follows that virtual transactions should help improve the overall efficiency of the market. Requiring financial contracts to settle for the difference in the electricity price of the transaction between the forward market and real-time market sets up the right incentives to affect day-ahead prices by improving convergence with expected real-time prices reflecting accurate price formation. But allocation of residual costs to the financial transactions does not follow from a coherent application of the principles of cost causation, or at least not by any analysis yet advanced. And any cost allocation to virtual transactions creates perverse incentives that have material consequences on the efficiency of the market.

The critique is correct that there is an asymmetry between the treatment of some virtual transactions and UTCs. The solution, however, need not be to adopt the flawed residual cost energy uplift allocation to INC and DEC virtual transactions and extend it to UTCs. The solution is to preserve the exemption for UTCs and FTRs, and then extend the same status to all virtual transactions. With the residual cost allocation applying only to real load, liquidity and entry in financial day-ahead virtual transactions would be enhanced, and the efficiency of the overall system should be improved.

PJM Virtual Bidding Recommendations

The PJM analyses cover many interesting examples, particularly in describing how virtual transactions can help with hedging and risk allocation. Anyone could benefit from reading and working through these examples. We could go further and extend these examples to include explicit analysis of equilibrium under uncertainty. This would be a useful direction for further work, and would likely reinforce and elaborate the understanding of the theoretical benefits of virtual trading.

PJM Virtual Bidding Examples

The existing PJM analysis could be separated into three related issues: (i) need to support efficient operations; (ii) potential for unintended outcomes due to market design defects or "modeling discrepancies"; (iii) computational challenges in managing a large volume of transactions in the day-ahead market.

Support of efficient operations is important, but the analysis does not yet speak to the essential elements of the efficient operation motivation. As illustrated above by the simple equilibrium analysis with day-ahead trading, the implicit assumptions about differences between day-ahead and real-time energy dispatch are not supported. Furthermore, the focus in analyzing the impact of virtual trading should be on assessing its impact on unit commitment decisions rather than assumed differences between day-ahead and real-time dispatch choices or transmission flows. The PJM analysis refers to the importance of commitment decisions throughout the report, but does no explicit analysis of those commitment decisions. This analysis would not be easy. But it is important to recognize that the absence of the analysis undermines the PJM conclusions because these do not generalize from the case of no uncertainty and no commitment decisions.

The examples of unintended outcomes appear almost exclusively to refer to cases of modeling discrepancies that could be considered to be market design defects. For example, the inconsistent treatment of a "dead bus," which is completely disconnected from the system due to the topology surrounding that bus, in the day-ahead and real-time can create price convergence problems. (p. 41) This seems true but the immediate question would be to consider the options for changing the treatment of dead buses, which can affect physical transactions, implicit virtual transactions, and explicit virtual transactions. There is nothing to suggest that the "dead" bus issue is inherent and has no solution other than limiting virtual bids or applying uplift to UTCs.

The discussion of the "load zone" settlements (pp. 42-45) is another analysis that is suggestive but incomplete. The cases shown are not equilibrium examples, so it is not clear what conclusions about virtual trading would follow. But to the extent there is a problem, the analysis does not point yet to the impacts of explicit virtual trading. A more natural question would be to revisit the reasons for price averaging across load zones and question why loads bid zonally instead of nodally. The market efficiency challenges of zonal models are well known, and were the reason for largely abandoning the zonal model for generation in the first instance (Hogan, 2002a). If virtual bidding presents problems for zonal aggregation, but not a fully nodal system, then the analysis would point to fixing the market defects of zonal aggregation.

These examples of market defects do highlight the possibility that virtual bidding will reveal underlying problems in market design. The market design will always have some unexpected consequences, and virtual bidding could largely fix the price convergence problems without correcting the inefficiencies of the underlying design. The treatment of persistent or unavoidable

market defects is a larger question, where virtual bidding can help but other policy solutions would be required (Hogan, 2014).

The PJM example of the "Small Positions on Low-Risk Paths" (p. 38) is an exception to the paper's concentration on market design flaws. However, it is not clear why this is a problem. The PJM analysis reports that these are transactions with a small expected payoff, but a high variance. These UTC transactions lose a small amount of money for many periods, but then see a very high payoff on a few occasions. This would seem to demonstrate that there are underlying risks that parties might want to hedge against. There is no claim that this is on average distorting the market. The only objection appears to be that transactions would have little direct efficiency effects on the dispatch. But there is no consideration of the risk allocation effects, and this objection would move the concern to the efficiency analysis that has not been done.

A large number of day-ahead virtual bids and offers could present a computational challenge for the unit commitment and day-ahead dispatch. The nature of security-constrained dispatch creates a very large number of potential constraints. In real-time, relatively few of these constraints will be active, and there is a well-established and manageable procedure for addressing these constraints. However, as illustrated by the simple equilibrium model above, many more of these constraints would have positive expected shadow prices day-ahead. Not all of these constraints would be active in the day-ahead solution, but it is logical that many more constraints might be active in the day-ahead solution. PJM reports that this larger number of binding constraints day-ahead is their empirical experience (p. 37).

This computational issue has to do with both the total number and the types of bids. Currently PJM imposes a "soft cap" budget on the number of individual virtual transactions. This virtual bid budget rule is an example of connecting the solution more directly to the problem analysis. Other approaches related to the value of the bids and offers might be considered. However, changing the flexibility and efficiency of the market design should be the last resort in dealing with computational problems, not the first choice of convenience. For example, reducing the number of bidding locations is likely to reduce the number of "unique" virtual bids, but could also produce a substantial reduction in flexibility. By contrast, limiting the budget for unique bids, but allowing bidders the flexibility to choose the locations, might be less disruptive to the market and yet address PJM's computational concerns.

PJM Recommendations
The PJM recommendations are (a) to align the locations for INCs and DECs at the locations with actual physical settlements; (b) to limit UTC at generation buses to be UTC sources only; and (c) to allocate uplift consistently across the different types of virtual transactions.

As a group, the PJM recommendations for restrictions on virtual bidding do not follow from its analysis. For example, the recommendation to limit virtual transactions to locations where there are physical settlements would need to be expanded at least to include all the nodes that are used as sources or destinations for monthly FTRs. Otherwise, the restriction would undo the intended market design element that allows FTRs to hedge either day-ahead or real-time prices. This is part of a related discussion that would address the FTR model in PJM.

In all the examples that deal with the treatment of "modeling discrepancies," the first consideration of analysis and recommendations should be to address the discrepancies. This is not included in the list of recommendations, but it should be the priority choice if we are to

promote market efficiency. The experience has often been that failure to address the underlying market design structure, and its compromises, often leads to unanticipated problematic results. It may be that fixing the market design is impossible, or politically difficult, but this assumption should not be the default choice or left unchallenged (Hogan, 2002a) (Hogan, 2013).

The limitation of UTCs to be only a source at generation buses ignores both the value of counterflow and the uncertainty about the direction of congestion. There is no analysis that connects any of the PJM examples to this recommendation. If the system consisted of simple radial connections or congestion was always in the same direction, then the intuition might apply. But this is not the case. In effect, this recommendation is the functional equivalent of prohibiting DEC bids at generation nodes. PJM has not made such a recommendation, but there is no analysis or explanation of the choice. The limitation on the direction of UTCs appears to be no more than another ad hoc way to limit the number of UTC transactions.

The number and types of virtual transactions can pose a computational problem. A more direct approach would address the computational problem by considering further what types and how many virtual bids could overload the system. Then the analysis would illuminate how the budget on total offers does or does not solve the computational problem. The PJM report offers no theoretical argument about why the computational problem dictates a need to control the source or destination of some virtual bids. As it stands, all that the rule for UTCs at generation buses might accomplish would be to substitute riskier INC and DEC combinations as the alternative, which could harm rather than help overall market performance.

Finally, the lack of symmetry in the uplift allocation points to a deeper problem. The main argument in the present paper is that the cause and effect of energy uplift allocation is a more complicated matter. There is a prima facie attraction to making the allocation rules consistent across different types of virtual transactions. A complete analysis of the costs and benefits of any uplift allocation to virtual transactions has not been done. As explained above, the best solution may be simply to eliminate the allocation of uplift to virtual transactions altogether. The cost causation argument for uplift allocation is inherently flawed. Absent cost causation, uplift allocation to the most elastic sources is exactly the wrong policy. This uplift allocation to virtual transactions has dramatic effects and is perhaps more harmful than anything else in reducing overall market liquidity and efficiency.

Conclusion

Virtual bidding is a necessary part of an efficient multi-settlement system. Implicit virtual bidding is effectively unavoidable, and explicit virtual bidding provides a critical component of efficient electricity market design. Restricting explicit virtual bidding creates market power for those who can make implicit virtual bids. Explicit virtual bidding mitigates or eliminates this market power, provides liquidity, improves price formation, allows hedging, connects naturally with longer term financial transmission rights, helps reveal defects in market design, and on average should improve system operations. The PJM report provides many illustrative examples and useful discussions of the benefits of virtual bidding. However, the discussion arises in a context where virtual bidding is under attack. By its very nature, the efficiency benefit of virtual bidding arises from dealing with uncertainty. A full treatment of uncertainty has not been done, but even a partial analysis indicates that recommendations for reforms of virtual bidding and associated cost allocations would go in an entirely different direction from that appearing in the

PJM report. Maintaining a close connection with the principles of efficient electricity market design would be necessary to avoid expensive and unwanted unintended consequences.

References

Bessembinder, H., & Lemmon, M. L. (2002). Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets. *The Journal of Finance*, *57*(3), 1347–1382. doi:10.1111/1540-6261.00463

Bjørndal, E., Bjørndal, M., & Midthun, K. (2016). *Congestion Management in a Stochastic Dispatch Model for Electricity Markets*. Retrieved from http://helper.ipam.ucla.edu/publications/enec2016/enec2016_12751.pdf

Deng, S. J., & Oren, S. S. (2006). Electricity derivatives and risk management. *Energy*, *31*(6-7), 940–953. doi:10.1016/j.energy.2005.02.015

FERC, F. E. R. C. (2014). *Uplift in RTO and ISO Markets*. Retrieved from https://www.pjm.com/~/media/committees-groups/task-forces/emustf/20140911/20140911-item-03-ferc-uplift-report-081314.ashx

Gribik, P. R., Hogan, W. W., & Pope, S. L. (2007). Market-Clearing Electricity Prices and Energy Uplift. Retrieved from http://www.hks.harvard.edu/fs/whogan/Gribik_Hogan_Pope_Price_Uplift_123107.pdf

Güler, T., Gross, G., Litvinov, E., & Coutu, R. (2010). On the economics of power system security in multi-settlement electricity markets. *IEEE Transactions on Power Systems*, *25*(1), 284–295. doi:10.1109/TPWRS.2009.2030398

Haas, H. J. (2009). Spread Bidding : MA Concerns and Mitigation Outline. Monitoring Analytics, September 10. Retrieved from http://www.monitoringanalytics.com/reports/Presentations/2009/Spread_Bidding_MA_Concerns_and_Mitigation_Outline_20090910.pdf

Hogan, W. W. (1992). Contract networks for electric power transmission. *Journal of Regulatory Economics*, *4*(3), 211–242. Retrieved from http://ezp-prod1.hul.harvard.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=bth&AN=16580807&site=ehost-live&scope=site

Hogan, W. W. (2002a). Electricity market restructuring: reforms of reforms. *Journal of Regulatory Economics*, *21*(1), 103–132. Retrieved from http://ezp-prod1.hul.harvard.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=bth&AN=17768581&site=ehost-live&scope=site

Hogan, W. W. (2002b). Financial transmission right formulations. Retrieved from http://www.hks.harvard.edu/fs/whogan/FTR_Formulations_033102.pdf

Hogan, W. W. (2008). Revenue Sufficiency Guarantees, Cost Causation, And Cost Allocation. FERC, Docket No. EL07-86-000. Retrieved from http://www.hks.harvard.edu/fs/whogan/Hogan_RSG_100908.pdf

Hogan, W. W. (2012). *Electricity Market Design Financial Transmission Rights, Up To Congestion Transactions and Multi-Settlement Systems*. Retrieved from http://www.hks.harvard.edu/fs/whogan/Hogan_UTC_071612.pdf

Hogan, W. W. (2013). Electricity Scarcity Pricing Through Operating Reserves. *Economics of Energy & Environmental Policy*, 2(2), 65–86. Retrieved from http://www.pserc.cornell.edu/empire/2_2_a04.pdf

Hogan, W. W. (2014). *Electricity Market Design Flaws and Market Manipulation*. Retrieved from http://www.hks.harvard.edu/fs/whogan/Hogan_MDFMM_02_03_14.pdf

ISO_New_England. (2016). *Report of ISONE: Price Formation in Energy and Ancillary Services Markets Operated by Regional Transmission Organizations and Independent System Operators, FERC Docket No. AD14-14-000*. Retrieved from https://elibrary.ferc.gov/IDMWS/common/opennat.asp?fileID=14163250

Jha, A., & Wolak, F. A. (2015). *Testing for Market Efficiency with Transactions Costs: An Application to Convergence Bidding in Wholesale Electricity Markets*. Retrieved from http://web.stanford.edu/group/fwolak/cgi-bin/sites/default/files/CAISO_VB_draft_VNBER_final.pdf

Joskow, P. L., & Schmalensee, R. (1983). *Markets for power: an analysis of electrical utility deregulation. MIT Press Books*. MIT Press Book. Retrieved from http://ideas.repec.org/b/mtp/titles/0262600188.html

Li, R., Svoboda, A. J., & Oren, S. S. (2015). Efficiency impact of convergence bidding in the california electricity market. *Journal of Regulatory Economics*. doi:10.1007/s11149-015-9281-3

Lo Prete, C., & Hogan, W. W. (2014). Manipulation of Day-ahead Electricity Prices through Virtual Bidding in the U. S. *IAEE Energy Forum*, 31–32. Retrieved from http://www.iaee.org/en/publications/newsletterdl.aspx?id=244

Monitoring Analytics. (2016). *2015 State of the Market Report for PJM* (Vol. 2). Retrieved from http://www.monitoringanalytics.com/reports/PJM_State_of_the_Market/2015/2015-som-pjm-volume2.pdf

Parsons, J. E., Colbert, C., Martin, T., & Mastrangelo, E. (2015). *Financial Arbitrage and Efficient Dispatch in Wholesale Electricity Markets* (No. WP 2015-002). Cambridge, MA. Retrieved from http://www.mit.edu/~jparsons/publications/20150300_Financial_Arbitrage_and_Efficient_Dispatch.pdf

PJM. (2012a). FTR Revenue Stakeholder Report. PJM, April 30. Retrieved from http://www.pjm.com/documents/~/media/documents/reports/20120430-ftr-revenue-stakeholder-report.ashx

PJM. (2012b). PJM's Response to the 2011 State of the Market Report. PJM Interconnection, May 11. Retrieved from http://www.pjm.com/documents/~/media/documents/reports/20120511-pjms-responce-to-the2011-state-of-the-market-report.ashx

PJM. (2015). *Virtual Transactions in the PJM Energy Markets*. Retrieved from http://www.pjm.com/~/media/committees-groups/committees/mrc/20151022/20151022-item-01-virtual-transactions-in-the-pjm-energy-markets-whitepaper.ashx

PJM. (2016). Operating Agreement Accounting. Retrieved from http://www.pjm.com/~/media/documents/manuals/m28.ashx

Pope, S. L. (2016). Affidavit of Dr. Susan L. Pope on Behalf of Elliott Bay Energy Trading, LLC, FERC Docket Nos. ER16-121-000, EL16-6-000, EL16-6-001. Retrieved from http://elibrary.ferc.gov/idmws/common/opennat.asp?fileID=14182961

Schweppe, F. C., Caramanis, M. C., Tabors, R. D., & Bohn, R. E. (1988). *Spot pricing of electricity.* Kluwer Academic Publishers. Retrieved from http://books.google.com/books?id=Sg5zRPWrZ_gC&pg=PA265&lpg=PA265&dq=spot+pr icing+of+electricity+schweppe&source=bl&ots=1MIUfKBjBk&sig=FXe_GSyf_V_fcIuTm UtH7mKO_PM&hl=en&ei=Ovg7Tt66DO2x0AH50aGNCg&sa=X&oi=book_result&ct=re sult&resnum=3&ved=0CDYQ6AEwAg#v=onep

Woo, C. K., Zarnikau, J., Cutter, E., Ho, S. T., & Leung, H. Y. (2015). Virtual Bidding, Wind Generation and California's Day-Ahead Electricity Forward Premium. *Electricity Journal, 28*(1), 29–48. doi:10.1016/j.tej.2014.12.006

---

[i] William W. Hogan is the Raymond Plank Professor of Global Energy Policy, John F. Kennedy School of Government, Harvard University. This paper draws on research for the Harvard Electricity Policy Group and for the Harvard-Japan Project on Energy and the Environment. Preparation of this paper was supported by the Financial Marketers Coalition. The author is or has been a consultant on electric market reform and transmission issues for Allegheny Electric Global Market, American Electric Power, American National Power, Aquila, Atlantic Wind Connection, Australian Gas Light Company, Avista Corporation, Avista Utilities, Avista Energy, Barclays Bank PLC, Brazil Power Exchange Administrator (ASMAE), British National Grid Company, California Independent Energy Producers Association, California Independent System Operator, California Suppliers Group, Calpine Corporation, CAM Energy, Canadian Imperial Bank of Commerce, Centerpoint Energy, Central Maine Power Company, Chubu Electric Power Company, Citigroup, City Power Marketing LLC, Cobalt Capital Management LLC, Comision Reguladora De Energia (CRE, Mexico), Commonwealth Edison Company, COMPETE Coalition, Conectiv, Constellation Energy, Constellation Energy Commodities Group, Constellation Power Source, Coral Power, Credit First Suisse Boston, DC Energy, Detroit Edison Company, Deutsche Bank, Deutsche Bank Energy Trading LLC, Duquesne Light Company, Dyon LLC, Dynegy, Edison Electric Institute, Edison Mission Energy, Electricity Corporation of New Zealand, Electric Power Supply Association, El Paso Electric, Energy Endeavors LP, Exelon, Financial Marketers Coalition, FirstEnergy Corporation, FTI Consulting, GenOn Energy, GPU Inc. (and the Supporting Companies of PJM), GPU PowerNet Pty Ltd., GDF SUEZ Energy Resources NA, Great Bay Energy LLC, GWF Energy, Independent Energy Producers Assn, ISO New England, Koch Energy Trading, Inc., JP Morgan, LECG LLC, Luz del Sur, Maine Public Advocate, Maine Public Utilities Commission, Merrill Lynch, Midwest ISO, Mirant Corporation, MIT Grid Study, Monterey Enterprises LLC, MPS Merchant Services, Inc. (f/k/a Aquila Power Corporation), JP Morgan Ventures Energy Corp., Morgan Stanley Capital Group, Morrison & Foerster LLP, National Independent Energy Producers, New England Power Company, New York Independent System Operator, New York Power Pool, New York Utilities Collaborative, Niagara Mohawk Corporation, NRG Energy, Inc., Ontario Attorney General, Ontario IMO, Ontario Ministries of Energy and Infrastructure, Pepco, Pinpoint Power, PJM Office of Interconnection, PJM Power Provider (P3) Group, Powerex Corp., Powhatan Energy Fund LLC, PPL Corporation, PPL Montana LLC, PPL EnergyPlus LLC, Public Service Company of Colorado, Public Service Electric & Gas Company, Public Service New Mexico, PSEG Companies, Red Wolf Energy Trading, Reliant Energy, Rhode Island Public Utilities Commission, Round Rock Energy LP, San Diego Gas & Electric Company, Secretaría de Energía (SENER, Mexico), Sempra Energy, SESCO LLC, Shell Energy North America (U.S.) L.P., SPP, Texas Genco, Texas Utilities Co, Tokyo Electric Power Company, Toronto Dominion Bank, Transalta, TransAlta Energy Marketing (California), TransAlta Energy Marketing (U.S.) Inc., Transcanada, TransCanada Energy LTD., TransÉnergie, Transpower of New Zealand, Tucson Electric Power, Twin Cities Power