Available online at www.sciencedirect.com
ScienceDirect
IFAC PapersOnLine 50-1 (2017) 225–232

Virtual Bidding: Equilibrium, Learning, and the Wisdom of Crowds ⋆

Jonathan Mather * **Eilyan Bitar** ** **Kameshwar Poolla** *

* *Department of Mechanical Engineering, UC Berkeley*
** *School of Electrical and Computer Engineering, Cornell University*

Abstract: We present a theoretical analysis of virtual bidding in a stylized model of a single bus, two-settlement electricity market. North-American ISOs typically take a conservative approach to uncertainty, scheduling supply myopically in day-ahead (DA) markets to meet expected demand, neglecting the subsequent cost of recourse required to correct imbalances in the real-time (RT) market. This can result in generation costs that far exceed the minimum expected cost of supply. We explore the idea that virtual bidding can mitigate this excess cost incurred by myopic scheduling on the part of the ISO. Adopting a game-theoretic model of virtual bidding, we show that as the number of virtual bidders increases, the equilibrium market outcome tends to the socially optimal DA schedule, and prices converge between the DA and RT markets. We additionally analyze the effects of virtual bidding on social welfare and the variance of the price spread. Finally, we establish a repeated game formulation of virtual bidding, and investigate simple learning strategies for virtual bidders that guarantee convergence to the Nash equilibrium.

Keywords: Electric Power Systems, Game Theory, Economics, Stochastic Programming, Learning Algorithms, Asymptotic Analysis

1. INTRODUCTION

In electricity markets, virtual bidding (VB) allows market participants to buy and sell electricity without the obligation to physically produce or consume it. This opens up market participation to financial entities or third parties without generation or load assets, allowing them to take advantage of arbitrage opportunities and promote market liquidity. VB is similar in nature to futures trading in more traditional commodity markets, where contracts are settled financially and no physical delivery takes place.

Deregulated electricity markets are typically characterized by centralized multi-settlement markets administered by an independent system operator (ISO). More specifically these markets have both a day-ahead (DA) and real-time (RT) market. In the DA market, the ISO collects demand bids and supply offers from participants and, based on the expected transmission network conditions, determines an economic unit commitment and dispatch with associated locational marginal prices (LMPs) for each hour of the next day. A similar economic dispatch procedure is conducted in the RT market, but in response to real-time system conditions, typically at five to fifteen minute intervals. The important distinction between the two markets is that cleared DA schedules are just financial contracts that can be settled at real-time prices, whereas the RT market represents physical delivery of energy *i.e.* no power flows in the DA market. It is this fact that enables the inclusion of VB that is not backed by physical assets in electricity markets.[1] A more complete discussion of these issues can be found in Hogan (2016).

A virtual bid in such a market structure is comprised of a buy (sell) bid in the DA market, matched by a sell (buy) offer in the RT market, such that any position taken up in the DA is completely liquidated in the RT market, with no obligation to physically produce or consume electricity. This allows virtual bidders to arbitrage the price difference between the DA and RT markets. This should in general cause the DA and RT prices to converge in expectation, as any price gap can be exploited by a risk neutral speculator. This is why VB is sometimes referred to as convergence bidding. It is also important here to highlight the difference between explicit and implicit VB. In the absence of an explicit VB mechanism, participants backed by physical assets can still make implicit virtual bids, for example bidding more capacity than they have available into the DA market and then purchasing the shortfall on the spot market in real time. Implicit VB can cause market power issues, and compromise the integrity of load and generation forecasts. Allowing a mechanism for explicit VB, as described above, goes some way to mitigating these issues. More broadly, whenever we discuss VB in this paper, we are referring to explicit VB.

The benefits of virtual bidding are discussed at length in Hogan (2016); Celebi et al. (2010); Isemonger (2006), and are generally characterized as: improved liquidity, mitigation of market power, improved market efficiency and price formation, reduced price volatility, and providing market participants with the ability to hedge price risk. A poten-

⋆ Supported in part by NSF grants ECCS-1351621, CNS-1239178, IIP-1632124, US DoE under the CERTS initiative.

[1] VB is implemented in the majority of North-American ISOs, including PJM, NYISO, ISO-NE, MISO, and CAISO.

10.1016/j.ifacol.2017.08.038

tial downside of virtual bidding highlighted in the above works is the incentive for a virtual bidder in possession of a bilateral or external position to influence the profitability of this position through virtual trades. This is of particular relevance to those traders in possession of financial transmission rights (FTRs), as described in Ledgerwood and Pfeifenberger (2013), to the effect that both ISO-NE and PJM enforce revenue capping when a participant makes a virtual bid which affects its own FTR revenue stream. Parsons et al. (2015) also suggest that virtual traders can exploit approximations in market designs to make profits without improving system operation, for example real-time ramping requirements that are not considered in the DA market. Some attempts have been made to quantify the efficiency effects of virtual bidding through empirical studies testing for the existence of profitable bidding strategies. See Saravia (2003); Borenstein et al. (2008); Li et al. (2015); Jha and Wolak (2015).

In this paper we focus on the ability of virtual bidding to improve outcomes in electricity markets with uncertainty. Hogan (2016) emphasizes this as one of the most valuable aspects of VB, yet also highlights the lack of rigorous work or analysis in this area, mainly due to the complexity involved. Modern electricity markets face increasing uncertainty in both supply and demand with a growing penetration of renewable and distributed generation. ISOs typically take a conservative approach to uncertainty, scheduling supply myopically in the DA market to meet expected demand, and neglecting the subsequent cost of recourse required to correct imbalances in the real-time (RT) market. They also hold significant reserve margins to manage large deviations or deal with contingencies. This deterministic approach to power markets provides reliable and secure system operation, but it can be costly. Recent advances in stochastic and robust optimization have shown that significant cost reductions can be achieved by more explicitly incorporating uncertainty into market clearing algorithms. See Bertsimas et al. (2013); Munoz-Alvarez et al. (2014); Hreinsson et al. (2015). Such approaches are tractable for real, large-scale, power systems; however, they face resistance from ISOs and system operators due to their perceived complexity, opacity, and reduction in system reliability.

We propose the novel thesis that, under certain assumptions, deterministic system operation with virtual bidding approximates the results of stochastic system operation, obviating the need for implementing new market algorithms. We demonstrate this result on a stylized model of a single bus, two-settlement electricity market. While a simple model, the results are instructive and point the way to models that more closely approximate the true operation of real power systems in future work. Our model is similar in nature to that proposed by Tang et al. (2016), although the equilibrium analysis, welfare analysis, and learning dynamics presented here are novel. All of these analyses are shown to depend on the accuracy of the aggregate beliefs of the population of virtual bidders. In short, the wisdom of the crowd. Our contributions are as follows:

- We characterize the unique, pure strategy Nash equilibrium of a population of profit-maximizing virtual bidders with heterogeneous beliefs about the market in which they participate.
- We show that as the number of virtual bidders increases, the DA ISO schedule approaches the socially optimal schedule, and prices converge in expectation between the DA and RT markets.
- We investigate simple learning strategies for individual speculators and characterize conditions under which they converge to the unique Nash equilibrium.

Organization: The remainder of the paper is organized as follows. In Section 2, we formulate a model of the two-settlement market and the virtual bidding mechanism. In Section 3 we characterize the pure Nash equilibrium among virtual bidders, and discuss its effect on social welfare. In Section 4 we propose simple learning dynamics under which virtual bidders reach the Nash equilibrium, and Section 5 concludes.

Notation: Denote by $\mathbb{R}$ and $\mathbb{R}_+$ the sets of real and nonnegative real numbers, respectively. Denote the transpose of a vector $x \in \mathbb{R}^n$ by $x^\top$. Let $x_{-i} = (x_1, .., x_{i-1}, x_{i+1}, .., x_n) \in \mathbb{R}^{n-1}$ be the vector including all but the i^{th} element of x. Denote by $\mathbf{1}$ the vector of all ones, and by $\mathbf{E} := \mathbf{1}\mathbf{1}^\top$ a square matrix of all ones. Denote by $\text{diag}(x_1, \ldots, x_n)$ the diagonal matrix with diagonal elements $\{x_i\}_{i=1}^n$.

2. MARKET MODEL

We consider a simplified model of a two-settlement electricity market administered by an independent system operator (ISO) for a copper plate power system.[2] The electricity market is cleared in two stages: day-ahead (DA) and real-time (RT). In the DA market, the ISO must determine an initial dispatch of supply subject to uncertainty in the eventual realization of demand, which we assume to be perfectly inelastic and denote by $D \in \mathbb{R}_+$. We describe uncertainty in the ISO's prior belief about demand by modeling D as a random variable with mean $\mu := \mathbb{E}[D]$ and variance $\sigma^2 := \text{Var}(D)$.

The ability to schedule supply in the DA market is essential, as certain generation resources (e.g., coal and nuclear) have limited ramping capability, and must therefore be scheduled well in advance of the required delivery time. We define the production cost in the DA market according to a convex quadratic function of the form

$$C_{\text{DA}}(x) := \frac{1}{2}\alpha x^2,$$

for all production levels $x \geq 0$. Here, $\alpha > 0$ is assumed to be fixed and known by the ISO.

In RT market, demand is realized, and any mismatch between supply scheduled in the DA market, say x, and the realized demand D must be compensated through an adjustment of supply in the amount of $D - x$. The subsequent balancing cost incurred in the RT market is assumed to be a convex quadratic function of the form

$$C_{\text{RT}}(D - x) := \frac{1}{2}\beta(D - x)^2 + \gamma(D - x),$$

[2] We use the term *copper plate* here to imply a lossless, unconstrained transmission system.

where $\beta > 0, \gamma$, are assumed fixed and known by the ISO.[3] The inclusion of the affine term in the RT cost is an approximation of the fact that in reality the DA and RT cost functions will be coupled. Fast-ramping generators that have not been dispatched in the DA market may bid their spare capacity into the RT market.γ may be interpreted as the minimum marginal cost of fast-ramping generators in the RT market, such that $\gamma = \alpha\overline{x}$, where $\overline{x}$ is the total capacity available in the DA market.

We define the total expected cost of supply incurred under a DA schedule $x \geq 0$ as

$$J(x) := C_{\mathrm{DA}}(x) + \mathbb{E}[C_{\mathrm{RT}}(D - x)]. \tag{1}$$

Finally, the price at which energy is traded in each of the DA and RT markets is set by the ISO according to the marginal cost of supply in each market. Accordingly, given a DA dispatch of supply in the amount of $x \geq 0$, the DA and RT prices of energy are determined according to

$$P_{\mathrm{DA}}(x) := \alpha x \quad \text{and} \quad P_{\mathrm{RT}}(D - x) := \beta(D - x) + \gamma,$$

respectively, and the RT-DA price spread is defined as

$$\Delta(x) := P_{\mathrm{RT}}(D - x) - P_{\mathrm{DA}}(x).$$

Naturally, a priori uncertainty in demand will manifest itself as uncertainty in the RT price.

Remark 1. Implicit in our assumption of quadratic cost functions, in both the DA and RT markets, is the assumption that the underlying aggregate supply function in each market is linear and affine respectively. This is a common assumption in the power system economics literature, see for example Baldick et al. (2004). Throughout the paper, we interpret these supply functions as representing the true marginal cost of generation in each market. The treatment of more sophisticated models, which capture the effect of generator strategic behavior on the determination of these supply functions (in combination with strategic virtual bidding) represents an interesting and open direction for future research.

2.1 Conventional Market Clearing

The approach to market clearing practiced by the majority of North-American ISOs today is inherently myopic in nature. That is to say, the ISO schedules supply in the DA market to minimize the immediate system cost based on a point estimate (forecast) of demand, which we denote by $\widehat{D}$. In doing so, the ISO neglects the subsequent cost of recourse required to compensate imbalances that might arise between supply scheduled in the DA market and realized demand. Needless to say, the cost incurred by a myopic approach to scheduling such as this may far exceed the minimum expected cost of supply, which we formally define as

$$J(x^{\star}) := \min\{J(x) : x \in \mathbb{R}_{+}\}.$$

A straightforward calculation shows the optimal DA schedule to satisfy[4]

$$x^{\star} := \arg\min\{J(x) : x \in \mathbb{R}_{+}\} = \frac{\beta\mu + \gamma}{\alpha + \beta}.$$

This optimal DA schedule results in an ex-ante no-arbitrage condition, such that

$$P_{\mathrm{DA}}(x^{\star}) = \mathbb{E}[P_{\mathrm{RT}}(D - x^{\star})].$$

This is equivalent to stating that the expected price spread is equal to zero, $\mathbb{E}[\Delta(x^{\star})] = 0$. Myopic scheduling on the part of the ISO will result in a non-zero price spread in expectation

$$\mathbb{E}[\Delta(\widehat{D})] = (\alpha + \beta)\left(x^{\star} - \widehat{D}\right),$$

which can be exploited by speculators for profit. In what follows, we investigate the extent to which the speculative behavior of *virtual bidders* might drive the procurement of supply in the DA market towards the optimal procurement level $x^{\star}$.

2.2 Virtual Bidding

Consider a two-settlement electricity market in which a set of virtual bidders, $\mathcal{N} = \{1, \ldots, N\}$, participate. We assume that each virtual bidder is risk-neutral and seeks to maximize the expected profit they derive through price arbitrage between the DA and RT markets. Moreover, we assume that all virtual bids are quantity bids[5], such that the total supply x scheduled by the ISO in the DA market takes the form

$$x = \widehat{D} + \sum_{i=1}^{N} v_i,$$

where $v_i \in \mathbb{R}$ denotes the quantity bid of the ith virtual bidder. We adopt the sign convention that $v_i > 0$ ($v_i < 0$) corresponds to a demand bid (supply offer) in the DA market. We denote by $v = (v_1, \ldots, v_N)$ the virtual bid profile, and by $V := \sum_{i=1}^{N} v_i$ the aggregate virtual bid. It follows that the DA and RT prices induced under a virtual bid profile v are given by

$$P_{\mathrm{DA}}(x) = \alpha(\widehat{D} + V),$$
$$P_{\mathrm{RT}}(D - x) = \beta(D - \widehat{D} - V) + \gamma,$$

respectively, and the RT-DA price spread is equal to

$$\Delta(\widehat{D} + V) = (\alpha + \beta)\left(\frac{\beta D + \gamma}{\alpha + \beta} - \left(\widehat{D} + V\right)\right)$$

These price functions are illustrated in Figure 1. In Figure 1a we see that the ISO schedules supply myopically to meet expected demand. In Figure 1b, the RT price is determined by the realization of demand. In Figure 1c the DA schedule is adjusted due to virtual bidding. In Figure 1d, we see that the RT price is still determined by the realization of demand, but is impacted by the virtual bids. In the model we consider, we allow for asymmetry in the beliefs held by individual virtual bidders regarding the market in which they participate. Namely, we assign to each virtual bidder $i \in \mathcal{N}$ a belief defined according to the tuple $(\alpha_i, \beta_i, \gamma_i, \mu_i)$, representing what virtual bidder i believes the DA and RT cost coefficients and mean value of demand to be.[6] The

[3] We make no assumption on the relative values of α and β, although generally in practice $\beta > \alpha$, reflecting the fact that it is more expensive to procure power in real-time than schedule it forward.

[4] Finding this solution in the more general network case with constraints amounts to solving a two-stage stochastic optimization problem.

[5] In practice, virtual bids allow for the specification of both price and quantity, thereby allowing virtual bidders to reveal their willingness to pay (accept) in addition to their quantity bid (offer).

[6] For now it is assumed that the ISO forecast of demand $\widehat{D}$ is common knowledge, although this will not be necessary for the learning dynamics presented in Section 4.

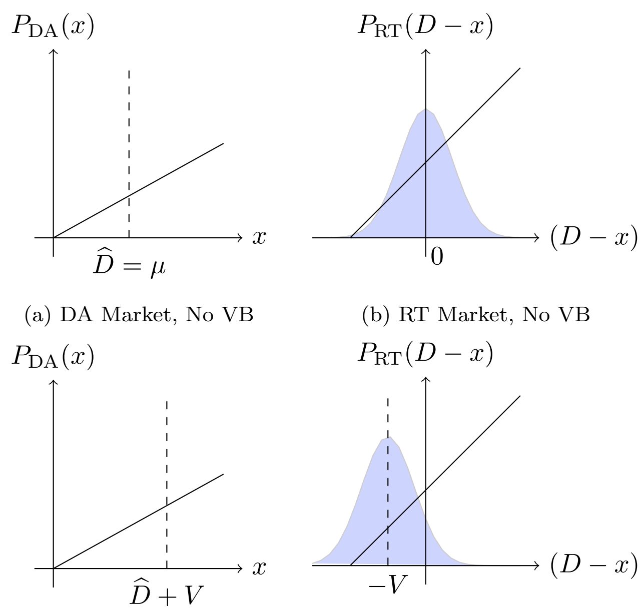

(a) DA Market, No VB (b) RT Market, No VB

(c) DA Market, With VB (d) RT Market, With VB

Fig. 1. DA and RT Markets, with and without virtual bidding

expected payoff to virtual bidder $i \in \mathcal{N}$, under a virtual bid profile v, is therefore defined according to

$$\begin{aligned} \pi_i(v_i, v_{-i}) &:= \mathbb{E}\left[\Delta_i(\widehat{D}+V)v_i\right], \\ &= (\alpha_i + \beta_i)\left(x_i^\star - \left(\widehat{D}+V\right)\right)v_i \end{aligned} \tag{2}$$

where $\Delta_i(\cdot)$ is the RT-DA price spread calculated using the beliefs of virtual bidder i, and $x_i^\star := (\mu_i\beta_i + \gamma_i)/(\alpha_i + \beta_i)$ represents the implicit estimate of the optimal DA schedule $x^\star$ by virtual bidder i. The collection of payoffs $\pi = (\pi_1, \ldots, \pi_N)$ together give rise to a normal-form (Cournot) game between the virtual bidders, which we refer to as the *virtual bidding game*. We define its equilibrium as follows.

Definition 1. (Nash equilibrium). The bid profile $v \in \mathbb{R}^N$ defines a *pure strategy Nash equilibrium* of the virtual bidding game if for each $i \in \mathcal{N}$, it holds that

$$\pi_i(v_i, v_{-i}) \geq \pi_i(\overline{v}_i, v_{-i}) \quad \text{for all} \quad \overline{v}_i \in \mathbb{R}.$$

3. EQUILIBRIUM ANALYSIS

We proceed with an explicit characterization and analysis of the equilibrium of the virtual bidding game defined in Section 2.2. Before proceeding, it will be convenient to measure the quality of the belief that each virtual bidder $i \in \mathcal{N}$ holds about the market in which they participate according to the quantity

$$\eta_i := x_i^\star / x^\star,$$

Naturally, the closer η_i is to one, the more accurate is the belief held by virtual bidder i. We say that virtual bidder i has perfect belief if $\eta_i = 1$. We define the market belief profile according to the vector $\eta := (\eta_1, \ldots, \eta_N)$. With this notation in hand, we present the following characterization of the equilibrium of the virtual bidding game.

Theorem 1. The virtual bidding game admits a unique pure strategy Nash equilibrium $v^\star \in \mathbb{R}^N$ satisfying

$$v_i^\star = \left(\eta_i - \frac{\sum_{j=1}^N \eta_j}{N+1}\right)x^\star - \left(\frac{1}{N+1}\right)\widehat{D}, \tag{3}$$

for each $i \in \mathcal{N}$.

The theorem is proved in Appendix A. It is immediate to see that under perfect beliefs (i.e., $\eta_i = 1$ for all $i \in \mathcal{N}$), this unique pure strategy Nash equilibrium $v^\star$ is symmetric, and reduces to

$$v_i^\star = \frac{x^\star - \widehat{D}}{N+1}, \tag{4}$$

for all $i \in \mathcal{N}$. All further discussion in this Section refers to the unique pure strategy Nash equilibrium under heterogeneous beliefs in (3). We see that the equilibrium action of each virtual bidder is to fill some fraction of the quantity gap between the optimal DA schedule and the myopic ISO schedule. This equilibrium action of each virtual bidder is a function of the quality of their own belief, and the average quality of the beliefs of all virtual bidders.

3.1 The Wisdom of Crowds

We wish to consider the effect on virtual bidding on the physical DA schedule at equilibrium, which is given by

$$x_N := \widehat{D} + \sum_{i=1}^N v_i^\star.$$

It is first useful to characterize the belief quality of the 'crowd' of virtual bidders. One way to model this is to assume that the individual beliefs of virtual bidders are drawn in an independent and identically distributed (IID) fashion from a common probability distribution. That is to say, we model the belief profile $\eta = (\eta_1, \ldots, \eta_N)$ as a collection of IID random variables having mean and variance

$$\mu_\eta := \mathbb{E}[\eta_i] \quad \text{and} \quad \sigma_\eta^2 := \mathrm{Var}(\eta_i),$$

for all $i \in \mathcal{N}$. In addition, we assume the belief profile η to be independent of the demand D.[7]

Definition 2. (Wisdom of the crowd). We define the crowd of virtual bidders to be *wise* if $\mu_\eta = 1$, i.e., their belief is correct on average.

It is not difficult to show that the DA schedule and price spread, which emerge at equilibrium, satisfy

$$x_N = \left(\frac{1}{N+1}\sum_{i=1}^N \eta_i\right)x^\star + \left(\frac{1}{N+1}\right)\widehat{D}$$

and

$$\Delta(x_N) = (\alpha+\beta)\left(\frac{\beta D + \gamma}{\alpha+\beta} - x_N\right).$$

We have the following Corollary to Theorem 1, which characterizes their asymptotic values as the the number of virtual bidders grows large.

Corollary 1. (Asympotic Market Efficiency). Assume that the virtual bidders collectively behave according to the

[7] This assumption may be strong, as it is not unreasonable to expect that the quality of private estimates and demand may be correlated in some fashion.

Nash equilibrium (3). As the number of virtual bidders participating in the market grows large, it holds that

$$\lim_{N\to\infty} x_N = \mu_\eta x^\star$$

and

$$\lim_{N\to\infty} \mathbb{E}\left[\Delta(x_N)\right] = (\alpha+\beta)(1-\mu_\eta)x^\star.$$

Namely, if the crowd is wise (i.e., $\mu_\eta = 1$), the DA schedule, which emerges at the Nash equilibrium, converges to the optimal DA schedule as the number of virtual bidders tends to infinity. As a result, the expected price spread between the RT and DA markets also converges to zero. Such asymptotic market behavior is to be expected, as a large number of virtual bidders will naturally compete away any ex-ante arbitrage opportunity.

We draw the following conclusions from this result, assuming that the desired outcome of implementing virtual bidding is improved market efficiency. First, it is important to have a crowd. The larger the number of virtual bidders, the smaller the expected arbitrage opportunity available to each bidder at equilibrium, and the closer one gets to the optimal DA schedule at equilibrium. In reality, the number of participants is likely to be determined by transaction costs associated with virtual bidding, and the risk premia that risk-seeking or risk-averse virtual bidders will demand or are willing to pay. It is not in the interests of an ISO to restrict access to virtual bidding markets in any way, for example through uplift payments associated with virtual bids as described in Hogan (2016).

Second, it is important that the crowd is wise. This is not something that can be prescribed per se, but is a phenomenon that has been observed in many contexts. From estimating the weight of an ox, Galton (1907), to modern day prediction markets, Surowiecki (2004), the crowd average generally outperforms individual estimates. One might also surmise that if participants have skin in the game, they are more likely to be invested in the quality of their own estimate, thus improving the crowd estimate.

Third, Corollary 1 holds for an arbitrary ISO forecast of demand $\widehat{D}$. Of course, the closer $\widehat{D}$ is to $x^\star$, the closer x_N will be to $x^\star$. It remains to be seen whether an aggregate crowd estimate of $x^\star$, could outperform one generated by a central ISO.

Given our distributional interpretation of beliefs, it is also possible to explicitly characterize the variance of the price spread, which results at the Nash equilibrium. Recalling that $\mathrm{Var}(D) = \sigma^2$, the spread variance in the absence of virtual bidding is easily calculated as

$$\mathrm{Var}(\Delta(\widehat{D})) = \beta^2\sigma^2.$$

In the presence of virtual bidders, we have

Corollary 2. The spread variance at the virtual bidding Nash equilibrium is equal to

$$\mathrm{Var}(\Delta(x_N)) = \beta^2\sigma^2 + (\alpha+\beta)\left(\kappa^2\frac{\sigma_\eta^2}{N}(x^\star)^2\right). \quad (5)$$

where $\kappa := \frac{N}{N+1}$ is a nondimensional parameter measuring the size of the crowd.

It follows directly that

$$\mathrm{Var}(\Delta(x_N)) \geq \mathrm{Var}(\Delta(\widehat{D}))$$

for any number of virtual bidders N. It can also be seen that as the number of virtual bidders grows large, it holds that

$$\lim_{N\to\infty} \mathrm{Var}(\Delta(x_N)) = \mathrm{Var}(\Delta(\widehat{D}))$$

This result is independent of the wisdom of the crowd, and states that at equilibrium the variance of the price spread under virtual bidding is lower bounded by the variance of the price spread under the myopic ISO schedule. Under these assumptions, the spread variance *never* decreases after the introduction of virtual bidding. This is due to the underlying variance in demand, that is not addressed at all by virtual bidding. Additionally, a large variance in the distribution of beliefs among virtual bidders only serves to worsen the variance of the spread, although this effect is mitigated as the number of virtual bidders increases.

This theoretical result would seem to be at odds with empirical results presented by Jha and Wolak (2015), which demonstrate that spread variances decreased after the introduction of virtual bidding in the CAISO market. However, it should be noted that they attribute this reduction in variance to the reduction of implicit virtual bidding by physical assets, and the fact that DA physical generation schedules should be closer to their real-time outputs under explicit virtual bidding, thereby reducing the need for costly purchases by the ISO to account for deviations in real-time. In our analysis we have not considered implicit virtual bidding by physical participants, but this would present an interesting avenue for further study. One could conjecture that under implicit virtual bidding the spread variance might increase due to both the false reporting of true physical production schedules, and heterogeneity of beliefs among implicit virtual bidders.

3.2 Welfare Analysis

We now investigate the social welfare properties of the virtual bidding Nash equilibrium. As demand is assumed to be inelastic, social welfare is naturally defined according to the expected cost of generation $J(x)$, which we previously defined in (1). To simplify the analysis we assume that a myopic ISO takes as its demand forecast $\widehat{D} = \mu$, although the results hold for arbitrary forecasts $\widehat{D}$.

We first see that under a myopic ISO dispatch, $\widehat{D} = \mu$, in the absence of virtual bidding, the generation cost takes the form

$$J(\mu) = \frac{1}{2}\left(\alpha\mu^2 + \beta\sigma^2\right).$$

If the ISO adopts the socially optimal dispatch $x^\star$, then the generation cost is

$$J(x^\star) = J(\mu) - \frac{1}{2}\frac{(\gamma-\alpha\mu)^2}{(\alpha+\beta)}.$$

As expected, we see that $J(x^\star) \leq J(\mu)$. We note the following identity

$$x^\star - \mu = \frac{\gamma-\alpha\mu}{\alpha+\beta},$$

such that

$$J(x^\star) = J(\mu) - \frac{1}{2}(\alpha+\beta)(x^\star-\mu)^2.$$

We now consider the generation cost at the equilibrium of virtual bidders. Since we assume that the individual beliefs of virtual bidders are drawn in an IID fashion from a common probability distribution we have that

$$J(x_N) = \mathbb{E}_\eta \left[C_{\mathrm{DA}}(x_N) + \mathbb{E}_D \left[C_{\mathrm{RT}}(D - x_N) \right] \right] \qquad (6)$$

where we must take expectations with respect to the random belief profile $\eta = (\eta_1, \dots, \eta_N)$. Assuming that the crowd is wise (i.e., $\mu_\eta = 1$), it can be shown that

$$J(x_N) = \frac{1}{2}\left(\alpha\mu^2 + \beta\sigma^2\right) + \frac{1}{2}(\alpha+\beta)\left(\kappa^2 \frac{\sigma_\eta^2}{N}(x^\star)^2 + \kappa(\kappa-2)(x^\star - \mu)^2\right),$$

We see that $J(x_{N=0}) = J(\mu)$, and that $J(x_{N\to\infty}) = J(x^\star)$, as expected. In general, however, virtual bidding may actually increase the generation cost due to the positive contribution from the variance of virtual bidders' beliefs. This can generally be characterized as occuring for a low number of virtual bidders, with high variance in beliefs. For a given variance, this positive term will be offset by a sufficiently large population of virtual bidders, since the final term is a strictly decreasing function of N, on the interval $N \in [0, \infty)$. We can in fact explicitly characterize the number of bidders after which the generation cost is strictly decreasing and less than or equal to the cost under a myopic ISO dispatch, denoted N_{dec}. For a fixed set of parameters $(\alpha, \beta, \gamma, \mu, \sigma_\eta)$, we have

$$N_{\mathrm{dec}} = \sigma_\eta^2 \frac{(x^\star)^2}{(x^\star - \mu)^2} - 2.$$

For the generation cost to be strictly decreasing for all $N \geq 1$, we require that

$$\sigma_\eta^2 \leq \frac{3(x^\star - \mu)^2}{(x^\star)^2}$$

At equilibrium, for $N \geq N_{\mathrm{dec}}$, virtual bidders never profit at the expense of loads. The expected cost of generation decreases in the presence of virtual bidders, and the marginal cost reduction associated with the addition of a new virtual bidder is always positive.

The exact impact of virtual bidders on social welfare will be a function of the specific market parameters, however these results highlight again the importance of a *crowd* of virtual bidders. Even the effect of a high variance in beliefs can be mitigated by the presence of a large number of virtual bidders.

4. REACHING EQUILIBRIUM

While the above results hold at the unique Nash equilibrium of the virtual bidders, actually reaching this equilibrium is a more subtle question. We consider simple learning dynamics for each virtual bidder, assuming that the two-settlement market is a repeated game in a homogeneous environment. In practice this might represent one hour of a day across many weeks, assuming similar patterns of weather and demand.

The best response of virtual bidder i is defined as

$$v_i^{BR} := \arg\max\{\pi_i(v_i, v_{-i}) : v_i \in \mathbb{R}\}$$

assuming that the actions of the other virtual bidders v_{-i} are given. It can be shown that this is equal to

$$v_i^{BR} = \frac{1}{2}\left(x_i^\star - \left(\widehat{D} + V_{-i}\right)\right)$$

where $V_{-i} = \sum_{j\neq i} v_j$. At equilibrium v_i^{BR} is equivalent to $v_i^\star$ in (3).

We consider a smoothed best-response learning dynamic, where at each iteration virtual bidder i plays a weighted sum of their previous action and their best response to the previous actions of all other bidders, with smoothing parameter θ_i, where $0 \leq \theta_i \leq 1$. See Fudenberg and Levine (1998); Hopkins (1999). The learning dynamic then takes the form

$$v_i(k+1) = \theta_i v_i(k) + (1-\theta_i) v_i^{BR}(k) \qquad (7)$$

$$= \theta_i v_i(k) + \frac{(1-\theta_i)}{2}\left(x_i^\star - \left(\widehat{D} + V_{-i}(k)\right)\right) \qquad (8)$$

where $v_i(k)$, $V_{-i}(k)$ represents the value of v_i, V_{-i}, respectively at the kth iteration, and $v_i^{BR}(k)$ represents the best response of player i to the actions of all other players at iteration k. We note that virtual bidder i will observe the quantity $(\widehat{D} + V_{-i}(k))$, assuming that the DA ISO dispatch $x(k)$ is published. We also note that $x_i^\star$ is only dependent on the beliefs of virtual bidder i. Thus (8) represents a valid learning dynamic, dependent only on the available information at each iteration. We also assume that the myopic ISO dispatch $\widehat{D}$ does not change, and that each virtual bidder does not change their beliefs $(\alpha_i, \beta_i, \gamma_i, \mu_i)$.

If $\theta_i = 0$, then this learning dynamic would constitute naive best response, where the virtual bidder plays their optimal action at each iteration assuming other virtual bidders do not change their actions. It is interesting to note that this same naive best response strategy is obtained if one attempts to solve (2) using gradient descent with exact line search, suggesting that virtual bidding could in fact be a form of gradient-descent algorithm that approximates the solution of the ISO problem.

Remark 2. An alternate interpretation of the learning dynamic in (8) is as the expected trajectory under a randomized update policy. At each iteration the virtual bidder i adopts their previous action with probability θ_i, and their best response to the previous actions of all other bidders with probability $(1-\theta_i)$. All the following results hold for this stochastic learning dynamic; however, the concept of asymptotic stability is replaced with convergence in expectation.

4.1 Stability and Convergence Analysis

We assume that all virtual bidders adopt the learning dynamic in (8). Considering the collective learning dynamics of all virtual bidders it can be shown that

$$(v(k+1) - v^\star) = A_\Theta (v(k) - v^\star)$$

where A_Θ is defined as

$$A_\Theta := \frac{1}{2}\left((\mathbf{I} + \Theta) - (\mathbf{I} - \Theta)\mathbf{E}\right)$$

where $\Theta = \mathrm{diag}(\theta_1, \dots, \theta_N)$.

We have the following result

Theorem 2. As $k \to \infty$, $v(k) \to v^\star$, under the learning dynamics in (8), if

$$\frac{N-3}{N+1} < \theta_i < 1, \quad \forall i = 1, \dots, N \qquad (9)$$

The theorem is proved in Appendix B. Theorem 2 is equivalent to stating that the unique pure Nash equilibrium is globally asymptotically stable under the learning dynamics in (8), if condition (9) is satisfied.

We note that naive best response, *i.e.* $\theta_i = 0$, is only asymptotically stable for $N < 3$. As N grows larger, the feasible range of θ_i, over which the learning dynamics are asymptotically stable, shrinks. As to whether real virtual bidders would adopt smoothing parameters which satisfy (9) is unclear.

We now consider the speed of convergence of the learning dynamics. We have that

$$\begin{aligned}(v(k+1) - v^\star) &= A_\Theta \left(v(k) - v^\star\right) \\ &= \left(A_\Theta\right)^k \left(v(0) - v^\star\right) \\ \|(v(k+1) - v^\star)\| &= \|\left(A_\Theta\right)^k \left(v(0) - v^\star\right)\| \\ &\le \|A_\Theta\|^k \|\left(v(0) - v^\star\right)\| \\ &= \rho\left(A_\Theta\right)^k \|\left(v(0) - v^\star\right)\|,\end{aligned}$$

where $\rho(A_\Theta)$ denotes the spectral radius of A_Θ. We see that the virtual bidders converge linearly in expectation to the Nash equilibrium at the rate of the spectral radius of A_Θ. It can be shown that

$$\rho\left(A_\Theta\right) = \max\left(\frac{1+\theta_{\max}}{2}, \frac{(1-\theta_{\min})}{2}N - \frac{1+\theta_{\min}}{2}\right)$$

For fast convergence we want $\rho\left(A_\Theta\right)$ as small as possible. This minimum is achieved if all virtual bidders adopt the same smoothing parameter $\theta_i = \frac{N-2}{N+2}$, $\forall i = 1, \ldots, N$. For large N, this convergence will be slow. More generally, the speed of convergence is limited by the fact that the only information each virtual bidder receives on the actions of the other players is the sum of their bids. This means that if we allowed virtual bidders to update and improve their estimates $(\alpha_i, \beta_i, \gamma_i\mu_i)$, at each iteration this would not necessarily improve the speed of convergence. In fact it would only serve to shift the Nash equilibrium towards the equilibrium under perfect beliefs. If we assume that the crowd is wise, $\mu_\eta = 1$, and remains wise as virtual bidders improve their private estimates, then the quality of the information received by each virtual bidder, namely the sum of the bids of other bidders, is not improved by updated private beliefs.

5. CONCLUSIONS

We have analysed a simple model of a two-settlement market under a myopic ISO dispatch, which provides insight into the equilibrium behavior of virtual bidders. The key results are as follows. At equilibrium, if the crowd of virtual bidders is wise, the DA schedule tends to the social optimum, and the expected price spread tends to zero, as the number of virtual bidders grows large. Additionally the variance of the price spread under virtual bidding is always greater than or equal to the variance in the case where there is no virtual bidding, explicit or implicit. We have also proposed simple learning dynamics, which have as their asymptotically stable equilibrium the Nash equilibrium of the virtual bidding game.

It is important to acknowledge the differences between this simplified model and real-world two-settlement power markets. One problem, highlighted by Parsons et al. (2015), is that typically the DA and RT clearing algorithms are run in different ways. Namely that the DA dispatch must consider commitment costs, and the RT dispatch must consider constraints such as ramping limits of generation. Another issue is that virtual bids are settled at DA hourly prices, but the RT market is typically run every 5-15 minutes. The 'RT price' that is used to settle virtual bids is often the average hourly RT price. These concerns distort the incentives of virtual bidders, and can lead to undesirable behavior. Furthermore, real power markets are run on networks, with generation and load varying from node to node, in addition to requiring DA schedules to satisfy transmission constraints and contingency scenarios. We hope to address this general network problem in future work.

Finally, the environment in which real virtual bidders are speculating and learning is far from homogeneous. It is in fact highly heterogeneous, with network conditions, generation costs, and parameter distributions changing from day to day. This makes learning very difficult, and it is questionable whether virtual bidders can ever reach the equilibrium solutions presented in this paper. An interesting piece of further work would be to understand how far away the actions of real bidders are from equilibrium and the effect that this has on social welfare and price convergence.

REFERENCES

Baldick, R., Grant, R., and Kahn, E. (2004). Theory and application of linear supply function equilibrium in electricity markets. *Journal of Regulatory Economics*, 25(2), 143–167. doi:10.1023/B:REGE.0000012287.80449.97.

Bertsimas, D., Litvinov, E., Sun, X.A., Zhao, J., and Zheng, T. (2013). Adaptive robust optimization for the security constrained unit commitment problem. *IEEE Trans. on Power Syst.*, 28(March), 1–8. doi: 10.1109/TPWRS.2012.2205021.

Borenstein, S., Bushnell, J., Knittel, C.R., and Wolfram, C. (2008). Inefficiencies and Market Power in Financial Arbitrage: A Study of California's Electricity Markets. *The Journal of Industrial Economics*, 56(2), 347–378. doi:10.1111/j.1467-6451.2008.00344.x.

Celebi, M., Hajos, A., and Hanser, P.Q. (2010). Virtual bidding: The good, the bad and the ugly. *Electricity Journal*, 23(5), 16–25. doi:10.1016/j.tej.2010.04.014.

Fudenberg, D. and Levine, D. (1998). Learning and Evolution: Where Do We Stand? Learning in games. *European Economic Review*, 42, 631–639. doi:10.1016/S0014-2921(98)00011-7.

Galton, F. (1907). Vox Populi. *Nature*, 75, 450–451.

Hogan, W.W. (2016). Virtual bidding and electricity market design. *Electricity Journal*, 29(5), 33–47. doi: 10.1016/j.tej.2016.05.009.

Hopkins, E. (1999). A Note on Best Response Dynamics. *Games and Economic Behavior*, 29(1-2), 138–150. doi:10.1006/game.1997.0636. URL `http://dx.doi.org/10.1006/game.1997.0636`.

Hreinsson, K., Vrakopoulou, M., and Andersson, G. (2015). Electrical Power and Energy Systems Stochastic security constrained unit commitment and non-spinning reserve allocation with performance guarantees. *International Journal of Electrical Power and Energy Systems*, 72, 1–7. doi:10.1016/j.ijepes.2015.02.017.

Isemonger, A.G. (2006). The Benefits and Risks of Virtual Bidding in Multi-Settlement Markets. *Electricity Journal*, 19(9), 26–36. doi:10.1016/j.tej.2006.09.010.

Jha, A. and Wolak, F. (2015). Testing for Market Efficiency with Transactions Costs : An Application to Convergence Bidding in Wholesale Electricity Markets.

Ledgerwood, S.D. and Pfeifenberger, J.P. (2013). Using virtual bids to manipulate the value of financial transmission rights. *Electricity Journal*, 26(9), 9–25. doi: 10.1016/j.tej.2013.09.014.

Li, R., Svoboda, A.J., and Oren, S.S. (2015). Efficiency impact of convergence bidding in the california electricity market. *Journal of Regulatory Economics*, 48(3), 245–284. doi:10.1007/s11149-015-9281-3.

Munoz-Alvarez, D., Bitar, E., Tong, L., and Wang, J. (2014). Piecewise Affine Dispatch Policies for Economic Dispatch under Uncertainty. 4–8.

Parsons, J.E., Colbert, C., Larrieu, J., Martin, T., and Mastrangelo, E. (2015). Financial Arbitrage and Efficient Dispatch in Wholesale Electricity Markets Erin Mastrangelo Financial Arbitrage and Efficient Dispatch in Wholesale Electricity Markets. (February).

Saravia, C. (2003). Speculative Trading and Market Performance : The Effect of Arbitrageurs on Efficiency and Market Power in the New York Electricity Market. (November).

Surowiecki, J. (2004). *The Wisdom of Crowds: Why the Many are Smarter Than the Few and How Collective Wisdom Shapes Business Economies Societies and Nations*. New York Doubleday.

Tang, W., Rajagopal, R., Poolla, K., and Varaiya, P. (2016). Model and Data Analysis of Two-Settlement Electricity Market with Virtual Bidding. *Proceedings of the IEEE Conference on Decision and Control*.

Appendix A. PROOF OF THEOREM 1.

Proof. Theorem 1. We begin by considering the payoff-maximizing action $v_i^\star$ of virtual bidder i, given the actions of all other virtual bidders v_{-i}.

$$v_i^\star = \arg\max\{\pi_i(v_i, v_{-i}) : v_i \in \mathbb{R}\}$$

Since $\pi_i(v_i, v_{-i})$ is a strongly concave function of v_i, we have that $\nabla\pi_i(v_i^\star, v_{-i}) = 0$ is a neccesary and sufficient condition for optimality. Solving we find that

$$v_i^\star = \left(x_i^\star - \left(\widehat{D} + \sum_{j\neq i}^{N} v_j + v_i^\star\right)\right) \tag{A.1}$$

We now assume that an equilibrium $v^\star = (v_1^\star, \dots, v_N^\star)$ exists, and will show that this is indeed the case. Summing over $v_j^\star$, we see that

$$\begin{aligned}\sum_{j=1}^{N} v_j^\star &= \left(x^\star \sum_{j=1}^{N} \eta_j - N\widehat{D} - N\sum_{i=1}^{N} v_j^\star\right) \\ &= x^\star \frac{\sum_{i=1}^{N} \eta_i}{N+1} - \frac{N}{N+1}\widehat{D}\end{aligned} \tag{A.2}$$

Substituting (A.2) into (A.1), we see that

$$v_i^\star = \left(\eta_i - \frac{\sum_{j=1}^{N}\eta_j}{N+1}\right)x^\star - \left(\frac{1}{N+1}\right)\widehat{D}$$

Since $\nabla\pi_i(v_i^\star, v_{-i}^\star) = 0$, $\forall i$, this is an equilibrium of the virtual bidding game, and is unique due to the strong concavity of the payoff function. □

Appendix B. PROOF OF THEOREM 2.

Proof. Theorem 2. We assume that all virtual bidders adopt the learning dynamic in (8), and denoting $\eta = [\eta_1, \dots, \eta_N]^\top$, and $\chi = \left(\eta x^\star - \widehat{D}\mathbf{1}\right)$, we have the full system update as

$$\begin{aligned}v(k+1) &= \Theta v(k) + \frac{(\mathbf{I}-\Theta)}{2}\left(\chi - (\mathbf{E}-\mathbf{I})v(k)\right) \\ &= \frac{(\mathbf{I}+\Theta) - (\mathbf{I}-\Theta)\mathbf{E}}{2}v(k) + \frac{(\mathbf{I}-\Theta)}{2}\chi\end{aligned}$$

It is straightforward to show that $v^\star$ in (3) is a unique fixed point of this iteration. It can also be shown that

$$(v(k+1) - v^\star) = A_\Theta(v(k) - v^\star)$$

where A_Θ and Θ are as defined in the text. To show asymptotic stability of the unique pure Nash equilibrium $v^\star$ under these learning dynamics, we require that $|\rho(A_\Theta)| < 1$. We cannot characterize $\lambda(A_\Theta)$ analytically, but we simply require that $\lambda_{\max}(A_\Theta) < 1$, and $\lambda_{\min}(A_\Theta) > -1$. We have that

$$\begin{aligned}\lambda_{\max}(A_\Theta) &= \lambda_{\max}\left(\frac{(\mathbf{I}+\Theta) + (\Theta-\mathbf{I})\mathbf{E}}{2}\right) \\ &\le \lambda_{\max}\left(\frac{\mathbf{I}+\Theta}{2}\right) + \lambda_{\max}\left(\frac{(\Theta-\mathbf{I})\mathbf{E}}{2}\right) \\ &= \frac{1+\theta_{\max}}{2} - \lambda_{\min}\left(\frac{(\mathbf{I}-\Theta)\mathbf{E}}{2}\right) \\ &\le \frac{1+\theta_{\max}}{2} - \lambda_{\min}\left(\frac{(\mathbf{I}-\Theta)}{2}\right)\lambda_{\min}(\mathbf{E}) \\ &= \frac{1+\theta_{\max}}{2}\end{aligned}$$

and using a similar analysis it can be shown that

$$\lambda_{\min}(A_\Theta) \ge \frac{1+\theta_{\min}}{2} - \frac{(1-\theta_{\min})}{2}N$$

Solving the following simple inequalities completes the proof.

$$\frac{1+\theta_{\max}}{2} < 1, \qquad \frac{1+\theta_{\min}}{2} - \frac{(1-\theta_{\min})}{2}N > -1 \qquad \square$$