# Equilibrium in Wholesale Electricity Markets

Nathan Larson and David Salant*

First version: November, 2002
This revision: November 13, 2003

**Abstract**

We develop game theoretic models to evaluate strategic behavior in deregulated electricity markets, with particular attention given to the market rules in place in California through the summer of 2000. We prove existence of a Nash equilibrium under two particular sets of market rules used by the CALPX and CAISO respectively. Next we derive a lower bound (strictly above marginal cost) on average equilibrium prices when there is a positive probability that at least one generator is capacity-constrained. Finally, we compare two competing methods for modelling competition in power markets: supply function equilibrium and discrete, multi-unit auctions and illustrate shortcomings of both approaches.

# 1 Introduction

In this paper, we develop a series of models and simulations to illuminate the factors underlying price formation in *supply schedule markets*. In the markets we consider, supply and demand clear in a centralized marketplace according to precisely specified rules. The signature feature of these markets is that each seller is required to offer its supply to the market in the form of a supply schedule. Individual sellers' supply schedules are aggregated to form a market supply curve which the market administrator uses to clear demand and set the market-clearing price. Supply schedule markets are commonly used to procure power generation and transmission and have also been used in corporate procurement. Although the stakes in these markets are quite high, and despite renewed academic scrutiny in the wake of the volatile energy prices of recent years, outcomes in supply schedule markets are not well understood.

Although there are undoubtedly many reasons why modelling these markets is difficult, we would like to highlight two prominent ones. The first is their complexity: choosing an optimal supply schedule involves formulating detailed expectations about the bids

*Affiliations and contact information: Nathan Larson, Robert H. Smith School of Business, University of Maryland, n.larson@rhsmith.umd.edu. David Salant, National Economic Research Associates, David.Salant@nera.com

of rival sellers and the elasticity of demand and then performing a high dimensional optimization. Market equilibrium is then a situation in which all of these expectations and optimizations are consistent with each other. The second complication is repeated interaction – a supply schedule market is often repeated periodically, sometimes on an hourly or daily basis, so the same set of market players may find themselves competing not just once, but many times over. Some studies have taken a "simplify and ignore" approach to these two issues: that is, approximate the true market environment with a simpler form of competition which the sellers (and the analyst) can solve, and exclude past and future competition from the analysis in order to focus on a single market. Studying a simpler market environment is appropriate when there are details that are extraneous and can be safely ignored; unfortunately, as we shall show, outcomes in supply schedule markets can be exquisitely sensitive to fine market details. Furthermore, as a practical matter, when faced with a highly complex decision problem individual sellers are likely to bring past experience to bear in assessing how their rivals are likely to bid and hence how profitable different strategies are likely to be. Long run behavior in the market will depend on how – and whether – learning nudges the sellers toward a stable equilibrium, but in order to study this we need a model that explicitly recognizes the dynamic nature of the market. We build a model of learning in supply schedule markets and study it through simulations. Intriguingly, we find that the model is characterized by significant and persistent volatility – a feature that is characteristic of real-world market data and that static models cannot replicate.

We start off in Section 2 with a brief overview of various notions of equilibrium in a static supply schedule game. In addition to reviewing the existing literature, one objective of this overview is to highlight certain features that are frequently present in real-world markets but often omitted from analysis – notably the "lumpiness" of supply schedule bids. We show that this lumpiness creates strong incentives for price undercutting that can make the existence of an equilibrium tenuous. The intuition behind the fragility of equilibrium in the static situation will be helpful in understanding the volatility that is generated in the dynamic model with learning.

Section 3 introduces our learning dynamics and applies them to study competition in a market with two suppliers. The model is kept deliberately simple – each supplier submits a step function supply schedule with at most two steps – in order to illustrate the price dynamics most effectively. Although the static game has a "competitive" equilibrium with price equal to the cost of the marginal unit of unsold supply, average market clearing prices in the dynamic setting are generally higher than this – sometimes substantially so. Furthermore, prices tend to fluctuate cyclically in a way that is reminiscent of Edgeworth cycles.

Section 4 extends this approach to a more elaborate model in which the sellers submit schedules with as many as thirty steps. As this is on the order of the number of steps permitted in the power markets of many US states, the qualitative results may be taken as indicative of what might be expected in those markets.[1] Furthermore, it has been conjectured that ignoring the lumpiness of supply schedules may of little consequence when the true supply schedules are close to smooth functions. We show that there is one sense in which this conjecture is true – over time, the average market price can be close to the price predicted in a static supply function equilibrium. However, this average masks significant fluctuations in the market price from one day to the next; these fluctuations show no signs of dying out as the supply schedules grow closer and closer to smooth curves. Section 5 concludes with a discussion of these results and suggestions for the direction of future work.

# 2 Equilibrium in Supply Schedule Markets

In order to fix ideas, we will briefly describe two of the sub-markets that were charged with procuring power for the state of California during the summer of 2000. In addition to serving as concrete examples for the sometimes abstract discussion that follows, these markets are interesting in their own right, as the role of market design in California's 2000 energy crisis is still much debated. With these examples in mind, we will sketch the major modelling approaches that have been employed to identify equilibrium outcomes in markets such as these.

## 2.1 Example: The California Power Markets

While there were a number of different options for trading energy in California in 2000, including bilateral trading, a block-forward market, a market for ancillary services, and other markets, we are specifically interested in the day-ahead market operated by the CALPX and the hour-ahead imbalance energy market operated by the CAISO. The former market was intended to approximately clear supply and demand on a day-ahead basis, allowing the California Independent System Operator (ISO) to develop a provisional dispatch schedule, while the latter was supposed to resolve the inevitable deviations from the day-ahead plan due to plant outages, real-time fluctuations in demand, and the like.

**The CAISO Hour-Ahead Market**

The CAISO hour-ahead market was held twenty-four times per day and was principally intended for reserve energy needed to meet real-time imbalances. Because generation units differ in terms of response time, an important characteristic for reserve energy, this market was divided into several sub-markets to distinguish different types of reserves. In each of these sub-markets, a bid by a generator was required to be a step function supply schedule with up to ten steps. That is, a bid could be characterized as a sequence

[1] Note, however, that we have made no attempt to calibrate the model to mimic any particular market, so direct quantitative comparisons are inappropriate.

of ten pairs $(p_i, q_i)$ for $i = 1, 2, ..., 10$. The interpretation is that $p_1$ is the lowest price at which the generator would offer its first $q_1$ megawatts of energy, $p_2$ is the lowest price at which it would offer between $q_1$ and $q_2$ megawatts, and so on. The ISO would assemble the bids of all of the generators into an aggregate supply curve. Then on a real-time basis (ten times per hour) it would assess the imbalance between actual load (demand) and previously scheduled supply – this difference can be thought of as the real time market demand curve. By intersecting this demand curve with the aggregate supply curve, the ISO would determine a market-clearing price. All units of generation offered at or below this price would be dispatched and paid the market-clearing price.[2]

**The CALPX Day-Ahead Market**

The bulk of supply and demand was intended to be scheduled in one of two daily markets run by the CALPX, the day-ahead and day-of markets. Although the two markets had similar bidding rules, in practice there was virtually no trading volume in the day-of market, so we focus on bidding in the day-ahead market. In this market, generators submitted independent supply schedules for each hour of the following day. Each supply schedule was permitted to consist of up to sixteen price-quantity pairs. Successive pairs were connected by straight lines to form a continuous, piece-wise linear supply curve with up to fifteen line segments. Furthermore, these supply curves were required to be strictly upward sloping – in contrast with the requirement of step function bids in the CAISO hour-ahead market. As in the CAISO market, the CALPX market was cleared by aggregating these supply curves across all bidders and equating supply with demand.[3]

### Variations on a Theme

The CALPX and CAISO markets illustrate two common templates for supply schedule markets, but there are some variations that are worth keeping in mind. One important question is whether supply bids are meant to be binding across a broad range of demand conditions or not. For example, the England and Wales Power Pool operates a market similar to the CALPX market, but where generators in the CALPX market could submit twenty-four different schedules each day – one for each hour of the day – generators in England and Wales submit a single daily bid which is binding for every hour.

[2] This is not the full story. In conjunction with their supply schedule bids, generators also placed a single capacity payment bid. The ISO would first procure a certain amount of reserve capacity by selecting generation units in order of their capacity bids. Only generators with winning capacity bids would be considered for real-time dispatch as described in the text. The capacity payment can be thought of as compensation for keeping a unit available in case it is needed for dispatch, while the energy payment (as described in the text) is payment for the actual power delivered.

[3] The framework also included the possibility of decremental bids that could be called upon to reduce excess generation and mitigate congestion; we will abstract from this.

## 2.2 Modelling Approaches

Most analysis of supply schedule markets has set aside the complications of repeated markets to focus on predicting the outcome in a single market in isolation. There are still a number of complicating details to consider. Among other issues, one must think about whether pricing is uniform or pay-as-bid, whether bids binding for a narrow or a broad range of demand conditions, and what the specific format in which bids must be submitted is. Furthermore, one must worry about the elasticity of demand, whether there are significant non-convexities in bidders' costs (in power markets, this might mean the fixed costs associated with starting up each marginal generation unit), and whether bidders have access to information or communication channels that might help them to collude. This paper will focus in particular on one issue that has proved divisive for modellers: how (and whether) to account for discreteness in bidders' supply schedules.

Existing analysis can be split loosely into two groups that differ in the approach taken to lumpiness in supply schedules. One line of work, following Klemperer and Meyer (1989), has treated supply schedules as smooth continuous functions. With this modelling assumption, there are established techniques for calculating a *supply function equilibrium* (SFE) for the market.[4] A second line of work, following von der Fehr and Harbord (1993) emphasizes supply lumpiness by considering a seller's supply schedule to consist of a collection of bids – one for each of a number of discrete units of supply. The underlying model is then one of a multi-unit auction (MUA), and the goal of the analysis is to identify the Nash equilibria of this auction. Between these two branches, the SFE approach has been used substantially more often in subsequent work, largely because it can be quite tractable if suitable assumptions on the shape of cost and demand curves are adopted. Nonetheless, there are some compelling reasons to think that a discrete approach to modelling supply schedules might be appropriate, particularly in the power market context. The first is technological: each generator's capacity is divided among a finite number of plants, and the fixed cost associated with starting up a plant can be substantial relative to the marginal cost of supplying an additional unit of power once the plant is up and running. Thus, it may be natural for a generator to develop its pricing strategy on a plant-by-plant basis. The second reason has to do with the market rules: in practice, all operational markets restrict bidders to submitting a finite number of price-quantity points rather than a continuous curve. SFE proponents would counter that when the number of discrete steps available to bidders is relatively large, when lumpiness in bids is aggregated over a number of suppliers, and when there is some uncertainty about supply (due to unplanned outages, for example) and demand, then the expected residual demand faced by bidders will be relatively smooth, and their optimal supply schedules will be closely approximated by continuous functions. While this argument has intuitive appeal, we know of no rigorous results to support it.

[4] The name can be a bit confusing. A supply function equilbrium is simply an application of the standard concept of Nash equilibrium to a particular, specialized type of model – not an alternative way to define an equilibrium.

The distinction between the two approaches is not just academic; they predict dramatically different market outcomes. In the multi-unit auction framework, pure strategy equilibria (when they exist) tend to exhibit Bertrand-like pricing – that is, prices at or near marginal cost. The driving force for this result is familiar from basic models of Bertrand competition: by undercutting a competitor's price by an infinitessimal amount, a generator can win a non-infinitessimal increase in its quantity dispatched. In contrast, with continuous, upward-sloping supply curves, undercutting a competitor only allows a generator to capture an infinitessimal increase in quantity, so prices in a SFE are typically not driven down to marginal cost. In fact, because there is a sort of strategic complementarity among generators – the less elastic your supply curve is, the less elastic my residual demand curve will be, and hence, the steeper my optimal supply curve – SFE models generally admit a range of different equilibria, with peak prices ranging between the competitive and Cournot levels. Depending on parameters, the latter may entail markups of perhaps *several times* marginal cost. Clearly, if this sort of modelling is intended to eventually inform market design and competition policy, the question of how two reconcile these two approaches, or at least establish conditions under which one is more appropriate than the other, seems pressing.

If we take the discreteness of the market seriously, as in the MUA approach, one disturbing implication is that pure strategy equilibria may not exist. To see why this is so, consider that the quantity allocated to any one bidder is not a continuous function of its bid – it jumps by a discrete positive amount when the bid on a marginal unit crosses below the market-clearing price. As long as the market-clearing price is above the cost of its marginal unit, a bidder can achieve a jump in its profits by undercutting the (former) market-clearing price with its bid on that unit. By itself, this implies that prices higher than the cost of the first losing unit are not stable. However, there is also an overcutting, or demand-reduction incentive. When the market-clearing price is close to marginal cost, the owner of the last winning unit makes little profit on that unit. By taking it off the market (or offering it at a high price) this owner may be able to drive up the market-clearing price and substantially increase profits on his infra-marginal units. Thus, prices near marginal cost may not be stable either.

As a consequence, if an equilibrium exists, it may require each firm to randomize, that is, to avoid being too predictable. Predictable bidding is costly in that it results in being undercut. Furthermore, in some cases, an equilibrium may fail to exist altogether. When prices are constrained to a discrete grid, the underlying game played by the bidders is finite, and Nash's classic existence result applies. However, when payoff function discontinuities are too irregular in a sense defined more precisely below (and by Dasgupta and Maskin (1986)), then this existence result need not carry over to the limiting case of a continuous price grid. Although the arguments involved are technical, they carry an important practical message. If the continuous limit of discrete price grid equilibria exists and is an equilibrium, then we can have some confidence that incentives and market outcomes adjust relatively smoothly as supply schedule lumpiness is smoothed out. On the other hand, if this limit fails to converge, then we can conjecture that incentives and outcomes will be quite sensitive to the precise specification of the grid. This sensitivity is in addition to any variability in outcomes due to the fact that bidders are using mixed strategies. The question of whether equilibrium exists with continuous prices is still open, but for one special case we have the following result.

**Proposition 1** *A two player step function auction, with strategies restricted to a single step, has an equilibrium in mixed strategies when both firms have the same capacities and zero marginal costs.*

The proof is sketched in the appendix. In closing, it is worth mentioning one additional point. One might be tempted to think of the problems introduced by lumpy supply bids as artificial, in the sense that if bidders were not constrained – either by market rules, or by practical limits on formulating a strategy – to use discrete, unit-by-unit bids, these problems would go away. However, this would be too optimistic. Even if bidders are able to formulate continous supply function bids, if the underlying cost cost curves are not smooth, then the supply functions freely chosen by bidders typically will not be smooth either. In this case, as Anderson and Philpott (2002) have shown, the existence of a continuous function that is a best response to rivals' supply functions is not guaranteed.

## 2.3 Why Simulations?

In light of the issues raised above, there is a compelling case for investigating supply schedule markets through simulations. First, if one accepts that bidders will arrive, either *via* introspection or through learning, at an equilibrium, one must still ask which is the appropriate model – one that emphasizes continuous or discrete supply schedules? Second, if one subscribes to the SFE approach, the question of *which* SFE (as there are typically many) arises. Or if the MUA approach is preferred, the question of how to compute an equilibrium (which will often be in analytically intractable mixed strategies) cannot be avoided. Finally, as any of these equilibrium approaches presumes that bidders hold correct expectations about the (highly complex) strategies pursued by their rivals and adopt responses that are optimal across many dimensions, one must ask how they arrive at these beliefs and optimal responses. A simulation model in which bidders adjust their strategies as they learn about the market environment is well-positioned to address all of these questions.

Our computational investigation of bidding by power generators contributes to a growing literature on computing equilibria in supply function games. Green and Newberry (1992) pioneered the approach of numerically computing supply function equilibria. Their approach is based on assuming continuity of the supply schedules and quadratic costs, and numerically solving the differential equations that characterize different equilibria. Their method is agnostic about which SFE, out of the continuum of possible equilibria, will actually be observed. More recently, there has been interest in explicitly simulating the process by which generators adjust their bids from one day to the next;

the hope is that this will not only provide a natural way to compute equilibria but will also shed light on which equilibria are likely to arise. Day and Bunn (2001) is a notable contribution in this vein. Finally, of substantial interest for our work is Baldick and Hogan (2002), who show that when generators' supply bids are subject to a certain type of noise (and with quadratic costs and continuous bid functions), only the linear SFE is stable. Their simulations generally confirm that under these conditions, bidding near the linear SFE is to be expected. Our results present a more nuanced view. We show that when discreteness is an integral feature of supply bids – even with a relatively fine grid – price instability and cycles are robust market outcomes.

# 3 Learning in a Simple Supply Schedule Auction

This section develops a model of learning in repeated supply schedule games and applies it to a simple, discrete unit market game. Our analysis of this model shows how bidder behavior can diverge from the predictions of both the MUA and SFE models. The game that we study has a unique pure strategy outcome in which the price is equal to the marginal cost of the most efficient unit not to be dispatched. However, we demonstrate analytically that this equilibrium will not be realized under our learning process. Simulations of the model dynamics reveal cylicality similar to that of Edgeworth cycles and average prices that are substantially above marginal cost. We start by introducing the market game and performing a traditional equilibrium analysis. Then we introduce the learning model and present the results of the simulations.

## 3.1 The Benchmark Game

There are two generators, $A$ and $B$, each of which has two discrete, equally sized blocks of capacity to offer to the market. For each firm, the marginal cost of power supplied by its first unit is 0 (across its entire capacity), and the marginal cost of power supplied by its second unit is 1. Demand is inelastic and certain: it is equal to 2 units up to a cap price of 3. The firms compete by bidding a single price for power from each of their units; that is, each firm submits a pair of positive real numbers $(p_1, p_2)$. The market is cleared by ranking all of the bids in ascending order to form an aggregate supply curve and selecting the total quantity to clear supply and demand, so as long as there are at least two units offered at a price of 3 or below, the units with the lowest and second-lowest prices are dispatched. Each dispatched unit is paid the system price which is equal to the bid of the marginal dispatched unit (i.e., the second-lowest bid). Ties are broken by dispatching the lower cost unit first, or by lottery if the tied units have the same cost.[5]

**Proposition 2** *In every pure strategy Nash equilibrium in weakly undominated strategies of the game described above:*

*i) The equilibrium price is* 1*.*

[5] Under other tie-breaking rules, an equilibrium may fail to exist.

*ii) Each generator dispatches one unit and earns a profit of* 1*.*
*iii) Each generator bids its second unit at its marginal cost of* 1*.*

**Proof.** First, observe that bidding its second unit below cost is a weakly dominated strategy for each firm. Next, suppose there were an equilibrium with a system price $p$ strictly greater than 1. Total profits in such an equilibrium must be less than $2p$, so at least one firm, say A, must be earning less than or equal to $p$. But then A could offer both of its units at $p' = p - \varepsilon$, thereby earning $2p - 1 - 2\varepsilon > p$ for $\varepsilon$ small enough, so this cannot be an equilibrium. Alternatively, suppose there were an equilibrium with $p < 1$. Because we have ruled out bidding below cost, each firm must be selling its first unit in this equilibrium and earning $p$. But then either firm could raise the price of its first unit to $1 - \varepsilon > p$ and continue to sell one unit but at a higher price, so this cannot be an equilibrium either. Therefore, the price in equilibrium must be equal to 1, so each generator earns profits of at most 1. Now suppose that some generator were to bid its second unit at a price $p_2 > 1$. Then the other could boost the system price to $p_2 - \varepsilon$ by raising the price of its first unit, thus earning a profit greater than 1. Thus, each generator must bid its second unit at a price of 1.

Finally, note that bidding both units at a price of 1 is a symmetric equilibrium. Neither firm can do better by lowering the price of its first unit or raising the price of its second unit, as neither of these actions changes the equilibrium price or allocation. If a firm were to raise its first unit price, it would fail to sell, losing a potential profit of 1. Conversely, if a firm were to lower its second unit price, it would be forced to operate that unit at a loss. There are also asymmetric equilibria in which one firm bids $(p_1, 1)$ and the other bids $(1, 1)$, with $p_1 < 1$. ■

In passing, we also note that there is also a continuum of equilibria in weakly dominated strategies in which both firms bid both units at $p < 1$ for any $p \in [0, 1)$. In these equilibria, each firm is indifferent to bidding its second unit below cost because it knows that unit will never be dispatched.

This example exhibits the standard intuition associated with Bertrand-like pricing models: the system price is bid down to the marginal cost of the most efficient losing unit. If the price were any higher, the owner of this most efficient losing unit would have an incentive to undercut slightly in order to get this unit dispatched at a profit. Notice, as well, that in equilibrium the price at which a firm offers its less efficient unit is irrelevant for its own profits. One might try to motivate the decision to price this unit at marginal cost with a story along the following lines:

1. Suppose A were to set $p_2^A > 1$. Then,

2. B would respond by raising the price on its first unit to just below $p_2^A$, driving up the system price.

3. But then A would have an incentive to reduce $p_2^A$ to undercut B's response, thereby dispatching his second unit at a profit.

4. But then B would respond by undercutting A again, and so on.

Presumably this process could end with prices converging back to 1. However, there is another possibility. At Stage 2, B must raise the price on its second unit by more than it raises the price of its first unit. (Otherwise it would end up dispatching the more expensive unit but not the less expensive one.) But B is more or less indifferent to *how much* more it raises the second unit price, since this unit is not expected to be dispatched in any case. (After all, A's first unit is still priced at 1.) If B raises its second unit price enough, then at Stage 3, A may find it more profitable to respond by raising $p_1^A$ to just below $p_2^B$, thus selling only one unit, but at a high price, rather than reducing $p_2^A$ in order to sell two units at a lower price. In this second story, the eventual convergence of all prices back to 1 seems less obvious, and the prospect of cycling or perhaps convergence to a mixed strategy equilibrium seems like a possibility. In the next section, we develop a model of a dynamic price adjustment process to replace the informal stories described above.

## 3.2 Learning: The Dynamics of Price Adjustment

Our price adjustment model is a version of stochastic best reply dynamics, and is based loosely on the idea that each generator tends to shift its pricing strategy toward bids that would be good responses to the historical pattern of prices chosen by its opponent. Let $f_t^i$ be the probability density function from which generator $i$ chooses its bid at time $t$. We assume that $f_0^i > 0$ on $\{(p_1, p_2) : 0 \leq p_1 \leq p_2 \leq 3\}$, so that initially every feasible price pair has some possibility of being selected. Let $\pi(\mathbf{p}\,, f)$ be the expected profit earned by bidding the price vector $\mathbf{p} = (p_1, p_2)$ against an opponent playing according to the distribution $f$. We assume that the generators adjust their bidding according to the following system of equations:

$$f_t^i = (1-\alpha) f_{t-1}^i + \alpha g_t^i \qquad i = A, B \tag{1}$$

That is, each generator's distribution over bids is a weighted average of its distribution last period and an update term $g_t^i$. This update term places greater weight on bids that would have earned relatively higher profits against the distribution used by the opponent in the last period:

$$g_t^i(\mathbf{p}) = e^{\kappa\pi(\mathbf{p}, f_{t-1}^{-i})} / D_t^i \tag{2}$$

where $D_t^i = \int e^{\kappa\pi(\mathbf{p}', f_{t-1}^{-i})} d\mathbf{p}'$ is just a normalization constant.

One can motivate these equations of motion as follows. After every period, a generator reviews the distribution of bids by its opponent in that period. (Think of a period as comprising a number of opportunities to bid – hourly markets within a day, for example – and that generators reassess their own strategies at the end of the day, after observing their opponent's bids over the course of the day.) The generator attempts to calculate a best response to that opponent distribution, but does not do so perfectly – the update distribution $g_t^i$ it comes up with places greater, but not infinite, weight on more profitable responses. This could be because analyzing the market is costly and time-consuming or perhaps because of small unmodelled cost shocks that shift the optimal response. Finally, there is inertia in the generator's pricing strategy: it merges its update over its optimal strategies with the distribution it used last period to determine its probability of making various bids today.

This adjustment process is more flexible than it may at first appear. Notice that equation (2) is just the familiar logit equation from discrete choice modelling. By setting $\kappa$ very high, one can model generators whose updates are arbitrarily close to placing full weight only on optimal responses. Furthermore, by additionally setting $\alpha$ to one, one can model generators as responding instantly and myopically to the last action chosen by their opponents as in the motivational stories from the previous section. On the other hand, if one believes that generators react gradually to changes in strategy by their opponents, then this can be incorporated by setting $\alpha$ lower.

Speaking loosely, we will call a pair of strategy distributions stable if the adjustment process converges to them from an open set of initial conditions. We have not tried to prove whether (1) converges, although our simulations will shed some light on this. However, it is clear that if stable strategy distributions exist, they must satisfy

$$f_t^i = e^{\kappa\pi(\mathbf{p}, f_{t-1}^{-i})}/D_t^i \tag{3}$$

Equation (3) is a continuous version of the logit equilibrium defined by McKelvey and Palfrey (1995) for discrete strategy spaces. They show that solutions to (3) converge to Nash equilibria in the limit as $\kappa$ grows large, so one can view the adjustment process as a way to predict which (approximate) Nash equilibrium is likely to be played. We will call a Nash equilibrium stable if it is the limit of a sequence of stable pairs of strategy distributions as $\kappa$ goes to infinity. Our first result is the following.

**Proposition 3** *The game has no stable pure strategy equilibria.*

**Sketch of Proof**

Suppose there were a sequence of logit equilibria converging to the symmetric equilibrium in which both generators bid $(1, 1)$. In order for second unit bids to converge to 1, it must be true that second unit bids near 1 sometimes win (at each of these logit equilibria), making it strictly more profitable to bid near 1 than at higher prices. But this means that first unit bids must sometimes lose, so first unit bids cannot converge to 1 at a faster rate than second unit bids. However, the rate at which each unit's bid converges to 1 is driven by the loss in profits that would be incurred if that unit were bid too high. This loss in profits is roughly constant at 1 for the first unit but goes to 0 for the second unit, so the distribution of first unit bids *must* converge at a faster rate, contradicting the previous claim.

The driving force behind this result is the assumption that a generator will choose a bid relatively arbitrarily when it is perceived to be irrelevant to its payoff. Later we will investigate outcomes when there is demand uncertainty, so that every bid is payoff-relevant with positive probability. But first, we investigate which outcomes, if any, are stable under our price adjustment process.

## 3.3 Computation of the Stable Equilibrium

In order to further analyze the behavior of the dynamics specified by (1), we have conducted numerical simulations for a range of parameter values. This section reviews the results. To simplify the computation, we start the simulations with symmetric initial strategy distributions, allowing us to work with a single difference equation. We conjecture that the results would not change substantially with asymmetric initial distributions, but we have not yet checked this.

### 3.3.1 The Low Inertia Case

In this scenario, $\alpha$ was set to 0.3, so each generator places a weight of 30% on approximately optimal responses to the most recent pattern of bidding by its opponent. Various values of $\kappa$ were tested; the results presented are for $\kappa = 100$, so a bidding strategy becomes about 2.7 times less likely to be used for each decrease of 0.01 in its expected profits. (For this game, a 0.01 decrease in absolute profits is roughly equivalent to 1% of total profits.) The initial strategy distributions used were uniform over the set of feasible strategies (but the results do not appear to be very sensitive to the initial conditions). Average market-clearing prices along various other measures are displayed in Figure 1. (The transient initial dynamics are omitted.) For these parameters, the market settles into a cycle with a length of about twenty periods. At the trough of this cycle, generators are bidding both units in at prices near 1.06 with high probability. However, at these prices, second unit profits are both rare and low, so there is little incentive for generators to bid their second units competitively. Second unit prices begin to drift away from marginal cost. But this means that generators can get away with boosting prices on their first units - as they realize this, first unit prices begin to drift upward as well, as does the market-clearing price. But as first unit prices begin to rise, it becomes more attractive to bring second unit prices back down in an attempt to supply the entire market. This effect reins in first unit prices as well, and so prices gradually drift back down to around 1.06. At the peak of the cycle, the average market-clearing price is about 1.3 - 30% above the "competitive" price level. Even at the trough of the cycle, average prices are bounded well away from marginal cost. The average market price over the entire cycle is 1.16.

### 3.3.2 The High Inertia Case

In this scenario, $\alpha$ was set to 0.01, so generators update their strategies relatively sluggishly. This might reflect various factors, such as bureaucracy and costs associated with evaluating and changing bids, or an internal preference for price stability, or perhaps a desire to stay below the regulatory radar by avoiding dramatic bid swings. The other parameters were set as for the previous case. Results of these simulations are presented in Figure 2. One can observe that the cycles persist, although their period is naturally much longer. The magnitude of the price swings has also diminished, particularly on the top end - the highest average price observed over the course of the cycle is now only about 16% above the competitive level. However the average market price over the cycle is approximately 1.12, lower than in the low inertia case, but not substantially so.

### 3.3.3 Low Inertia with Imperfect Profit Maximization

In the two cases presented above, the generators were assumed to be capable of quite precise profit maximization. Because optimal responses generally involve undercutting, this tends to destabilize the market and encourage price cycles. If generators make mistakes more frequently, then undercutting will be more difficult, and prices may be more stable. Here, we set $\kappa = 20$, so that a 0.02 drop in the expected payoff of a strategy now means it is about 1.49 times less likely to be used. We assume relatively rapid strategy updating of $\alpha = 0.3$. Results are presented in Figure 3. In contrast with the earlier figures, in this case we include the transient initial dynamics to provide a flavor for the speed of convergence. As anticipated, price fluctuations are now virtually negligible relative to the earlier cases. Furthermore, average prices are substantially higher than in either of the previous cases at roughly 1.27.

### 3.3.4 Reducing the Price Cap

In this scenario, we explore the effect of reducing the price cap from 3 to 1.5. The other parameters are as in the low inertia case ($\alpha = 0.3, \kappa = 100$). Results are in Figure 4. Here the outcome is quite different: the market-clearing price converges to the competitive level of 1 (actually, just below it). It is not entirely obvious why the reduction in the price cap should have such a substantial effect, since the new cap would not appear to constrain market-clearing prices, which were always less than 1.5 under the old cap. However, the price cap does help to constrain second unit prices from drifting too high when they are payoff irrelevant. This in turn makes it less attractive to gamble by raising one's first unit price - the chance of failing to sell is higher.

### 3.3.5 Demand Uncertainty

In this simulation run, we relax the assumption that demand is known to be equal to two units with certainty. Demand is assumed to be uniformly distributed on the interval $[1, 3]$. Generators are assumed to be capable of dispatching fractional amounts of each units capacity as necessary in order to clear demand. We set $\alpha = 0.1$, but otherwise the parameters are as in the baseline cases. The results are presented in Figure 5, and in a slightly different format from the other figures. The lower graph shows the average net profit (not market-clearing price, in this case) of a generator over time. The upper graphs are contour plots in $p_1 - p_2$ space that indicate the most frequently used strategies at selected points in time. At the trough of the price/profit cycle (near Period 60), second unit bids are clustered around 1.5 and first unit bids are fairly evenly distributed between 0 and 1.5. Observe that the first unit bids only come into play if demand falls between 1 and 2. At this point in the cycle, generators are relatively indifferent between setting a low first unit price in order to be sure to dispatch all of the first unit's capacity - with the risk of earning a low price if the competing firm does likewise - and setting a high first unit price but dispatching less than the unit's full capacity in expectation. Second unit bids are now directly payoff relevant whenever demand falls between 2 and 3. In these cases, raising or lowering one's second unit bid involves a standard tradeoff between increasing the chance of dispatching that unit at a profit and reducing the price at which the first (inframarginal) unit is dispatched. After prices have fallen far enough, driving down profits on the second unit when it is dispatched, this tradeoff begins to favor supply reduction. Second unit prices shoot up to between 2 and 2.5 (Periods 80 and 10). This in turn briefly induces the generators to boost their first unit prices to around 2 before the gradual process of price decline begins again.

Clearly, because capacity constraints are sometimes binding under demand uncertainty, marginal cost pricing would not be expected in this game. However, because equilibria in games with capacity constraints are notoriously difficult to compute, it would not be clear *a priori* how large the average markup should be expected to be. For example, the largest payoff that a generator could earn by simply offering all of its capacity at the cap price of 3 is just 0.75 (with probability 0.5, demand is between 2 and 3 in which case the generator dispatches 0.5 units on average). In contrast, the profit earned by each generator when both offer their full supply at a price of 1 (although this is no longer an equilibrium) is higher, at 0.875. From this, one might surmise that even if $(p_1 = 1, p_2 = 1)$ is no longer an equilibrium, prices in equilibrium might be close to this level on average in spite of the presence of capacity constraints. However, in fact, market-clearing prices average around 1.42 over the course of the price cycle, substantially above marginal cost.

### 3.3.6 Discussion

Our analysis of bidding behavior under step function bidding is admittedly quite limited. We have restricted attention to a very simple, stylized example and have analyzed one particular price adjustment dynamic. Nonetheless, the results are suggestive of effects that would persist in more sophisticated models. In particular, we have demonstrated that when profit-maximization is a bit noisy, Bertrand-like equilibria with prices at or near marginal cost need not be stable - even when they are essentially unique in the set of pure strategy Nash equilibria. Furthermore, this is true even in the limit as profit-maximization becomes perfect. We believe this notion that firms compute only approximate best responses to be an intuitive and appealing one. One could alternatively posit that firms sometimes experiment with alternative strategies. Because in practice it is often only the market-clearing price (and not competitors' bids) that is observed, both optimization errors and experimentation seem to be plausible features of generator bids.

One might infer that the stable cycles we observe represent the time average of mixed strategy equilibria of the games analyzed. Strictly speaking, however, this cannot be true for the case in which demand is certain. Because the strategy space is compact, the support of any mixed strategy must be bounded above. In particular, if there were a symmetric mixed strategy equilibrium, there would be an upper bound $\bar{p}_2$ on the support of second unit bids. Any generator bidding this upper bound would earn zero profit on its second unit, as it would never be dispatched. However, our stable outcomes have the property that generators dispatch their second units at a profit with positive probability, even in the noiseless limit. But this means that any generator bidding $\bar{p}_2$ could do strictly better with some lower bid, contradicting its inclusion in the mixed strategy. (This argument is a bit loose, but could be formalized and extended to asymmetric equilibria.) Therefore, our outcomes cannot be viewed literally as mixed strategy equilibria, although in a sense they may represent "approximate" mixed strategy equilibria.

# 4 Learning with Many-Step Supply Schedules

The fact that the temptation to undercut one's rival can lead to instability and cycles is familiar from the theory of capacity constrained price competition. The preceding simulations help to illustrate how instability and cycles persist when we replace the capacity constraints with an increasing unit cost and relax the constraint that a firm must offer all of its units at the same price. However, because the model used is still quite stylized, its implications for real-world power markets are at best suggestive. In this section, we develop and test a simulation model that more closely approximates real markets. In this setting, we can examine the extent to which activity in the simulated market can be represented by a supply function equilibrium. To preview the results, we find that in some cases the *time-average* behavior in the market is fairly well described by an affine SFE, but this average generally masks substantial and persistent cyclicality in prices. Moreover, market outcomes can be sensitive to the rule that firms use to update their supply schedules from one day to the next. This indicates a need for caution – without knowing more about the rules of thumb that real firms actually employ, market predictions based on theory and simulations should be treated tentatively.

## 4.1 The Model

Our simulations will model competition in supply functions between two generators for a single day. (The framework could be extended without much trouble to handle larger numbers of firms.) Each generator controls a number of discrete units. We assume that generation costs are quadratic, in the sense that a generator's total cost of providing $q$

units of power is $\frac{1}{2}c_i q^2$, where $c_i$ is a cost parameter for generator $i$. We also assume that there is a price grid, so that each generator's supply bid consists of the number of units it is willing to supply at each point on the price grid.

Supply and demand for the day are cleared in 24 hourly markets. Each generator submits a single supply function at the beginning of the day which is binding over all of the hourly markets. (This aspect of the model most closely mirrors the rules in the UK and Australia and contrasts with the CALPX rules, where generators could submit separate bids for each hour.) The demand curve for each hour is linear with variation in load over the course of the day captured by different intercepts for each hour. Both the slope of the demand curve and its deterministic variation from hour to hour are public knowledge. There is a uniform price for all units dispatched in each hour.

Unplanned outages constitute the only exogenous source of uncertainty in the model. Each unit has a small, independent probability (typically 10%) of being out of service for the day. Before market-clearing, each generator's supply bid is amended to account for out of service units. For example, if a generator has two units out of service, one which was offered at a price of 10 and another which was offered at a price of 15, its supply schedule is reduced by two units at price 15 and above and by one unit for prices between 10 and 15.

## 4.2 The Learning Dynamics

As earlier, each generator will aim to adjust its supply schedule toward a best response to the residual demand it faced yesterday. However, in this case it is impractical to compute hypothetical payoffs for an entire distribution of prospective bids (the space of possible bids is too large), so we model strategy adjustments in the direction of a single best response. Below, we discuss more precisely how a generator computes its approximate best response and the speed of its adjustment.

First, note that in computing a generator's best response, we will implicitly assume that it can observe the residual demand curve it faced yesterday. Since this amounts to an assumption that it can infer its rival's bids, this is a fairly innocuous assumption for markets in which bids are public; however, if the only information available to a generator consists of its own bid and hourly prices and quantities, then this assumption may be quite strong. A generator begins by calculating the profit-maximizing price and quantity on each hourly residual demand curve. If these price-quantity pairs trace out an upward sloping supply curve, then this is the generator's profit-maximizing response to yesterday's residual demand. More generally, when the price and quantity grids are relatively coarse, this set of optimal points may include pairs such as $(p_1, q_1)$ and $(p_2, q_2)$ where $p_2 > p_1$ but $q_2 < q_1$. In these cases, it is impossible to construct an upward sloping supply function that optimizes pointwise against each hourly residual demand curve. We assume that the generator constructs an approximate best response function that attempts to hit all of the hourly optima but deviates where necessary to avoid bending backward.

In describing a generator's updated strategy, we will distinguish between the supply function $s_i^t$ that it would like to submit if it could offer fractions of a unit at different prices and its bid $b_i^t$ which has a whole number of units at each price. A generator's updated strategy is composed of a weighted average of its desired bid yesterday and its best response:

$$\begin{aligned}\Delta s_i^t &= s_i^t - s_i^{t-1} \\ &= \beta(BR_i(b_{-i}^{t-1}) - s_i^{t-1}) \\ b_i^t &= \text{round}(s_i^t)\end{aligned}$$

The parameter $\beta$ can be thought of as reflecting the degree to which a generator thinks its rival's most recent behavior, as opposed to its longer term behavior, is the best indicator of how it will bid today. Alternatively, $\beta < 1$ can be taken as a stylization of the idea that generators do not update their bids on a daily basis. The need to distinguish between $s_i^t$ and $b_i^t$ arises when $\beta$ is small and individual units are relatively large. In this case, an updating rule that relies only on $b_i^t$ can get stuck: the term $\beta(BR_i(b_{-i}^{t-1}) - b_i^{t-1})$ may round to zero units even when $BR_i(b_{-i}^{t-1}) - b_i^{t-1}$ is consistently non-zero. Using $s_i^t$ in the updating rule means that the generator will eventually adopt even small strategy changes if the gain to doing so is persistent.

## 4.3 Simulations

This section presents the results of preliminary simulations of the model. We should note that the parameters have been chosen for illustrative value and are not intended to correspond to any real market. In the benchmark scenario considered, the cost parameters of the two generators are $c_1 = .2$ and $c_2 = .4$ respectively, and we assume that neither generator faces a binding capacity constraint. Prices lie on a 100 point grid between 0 and 20, and the hourly demand curve $D_h(p)$ is given by

$$D_h(p) = 10 + z(h) - p$$

for $h \in \{1, 2, ..., 24\}$. Load is assumed to vary linearly over the course of the day: $z(h) = 2h$. Figure 6 shows the maximum and minimum hourly demand, as well as the aggregate marginal cost curve for the industry.

In the sequel, we will often make reference to the affine supply function equilibrium of this model. One advantage of working with a quadratic cost specification is the existence of an easily calculated SFE in linear strategies which in this case is given (approximately) by $q_1 = 1.56\,p$, $q_2 = 1.26\,p$. In addition to possessing the virtue of convenience, the affine SFE is of interest because Baldick and Hogan (2001) have demonstrated that it is the only stable equilibrium under a particular form of perturbation.

Figure 7 presents snapshots of the generators' bid functions over time, starting from the affine SFE on Day 1. For reference, dotted lines show the affine SFE. While the supply schedules remain in the general neighborhood of the affine SFE, they do not converge to it. Instead, they cycle between phases with relatively more (Day 82) and less (Day 400) competitive pricing. Furthermore, if the market starts further from the affine SFE, these cycles tend to be more pronounced. Figure 8 shows the evolution of the market starting from a different set of initial conditions.[6] The variance of the average daily price (over the first 500 days) is 0.174, close to three times as great as when the market starts at equilibrium (0.064). Only a fraction (0.035) of the daily price variation can be attributed to the exogenous plant outages; the rest is endogenous. This suggests that any initial price volatility in the market tends to be preserved by the strategic adjustments of the generators. Volatility is also exacerbated when the generators are more responsive to recent market conditions. Figure 9 shows daily average prices when generators place a relatively high weight ($\beta = 0.5$) on responding to the previous day's residual demand. For comparison, the price series generated by affine SFE bidding (including outage shocks) is overlaid. Both the additional volatility and the cyclicality induced by the best response dynamics are pronounced.

The simulations allow us to investigate the impact of unit start-up costs and other non-convexities. As an example, suppose a "plant" consists of a group of five generation units. A generator incurs a fixed startup cost of 20 in each hour that it dispatches any of its first five units, $2 \times 20$ in each hour that it dispatches units one through five, plus any of its second five units, and so on.[7] Figure 10 shows the price series with start-up costs (for $\beta = .5$), with the no start-up cost series from Figure 9 for comparison. The generators recover some, but not all, of these costs through higher prices – average profits fall to 168 and 131 (versus 184 and 144 without start-up costs). And perhaps more surprisingly, price volatility is substantially dampened – the variance drops from 0.166 to 0.045.

### 4.4 Discussion

Our model of bidder adjustment through learning can be seen as complementary to a growing body of work on stability in supply schedule markets. Many contributions in this area begin with the premise that bidders can submit continuous supply functions, and hence look at the question of supply function equilibrium stability. Since there is generally a continuum of SFEs, ranging from those with shallow supply curves and low prices to ones with steep supply curves and high prices, stability analysis may be a promising way to give the SFE approach more predictive power. Because this type of stability analysis is often motivated by a heuristic description of a dynamic process of strategy adjustment, much of the work on stability has a flavor that is similar to our learning-based simulation approach.

[6] Actually, these initial conditions are the result of running the system for a few days with $\beta = .5$, starting from the affine SFE.

[7] This is a very simplistic portrayal of start-up costs – a more complete model would account for the fact that the number of times a plant must be started depends on whether or not the hours during which it operates are contiguous.

Loosely, an equilibrium is stable if bidders converge back to it after their strategies are perturbed slightly. An obvious difficulty with an analytical approach to SFE stability is that because strategies are functions, the space of possible perturbations to consider is very large. Most studies to date have circumvented this problem by examining stability with respect to a limited class of perturbations. One common approach is to assume that supply functions are linear. An early example in this vein is Rudkevich (1998) who exploits the fact that when bidders have quadratic costs and face linear residual demand, their best responses are also linear. Using this fact, he shows that if suppliers start off bidding their (linear) marginal cost curves, then a best response process will converge to the unique affine SFE. However, this analysis does not provide any insight into whether a linear equilibrium will be stable if the initial supply curves are not linear.

Baldick and Hogan (2002) broaden the scope of the analysis by considering initial conditions with non-affine supply curves; this is important, as it permits them to study the robustness of non-affine SFEs. Their notion of stability involves perturbing the initial conditions and then studying whether bidders' responses to this perturbation bring strategies closer to the initial conditions or carry them further away. To keep the analysis tractable, they restrict attention to a particular type of perturbation (and response) which replaces the rightmost portion of a supply curve with a linear segment. In this setting, they show that all SFEs except the affine one will be unstable. It is probably best to take this result in two pieces. First, it is probably safe to conclude that no single non-affine SFE is stable, as all are vulnerable to linear deviations. Second, if one is willing to assume that bidders take a (more or less) linear-quadratic view of the world, then only linear deviations need to be considered – in this case, there is only one SFE and it is stable. However, if we relax the assumption that bidders default to linear schedules, then there is nothing to preclude cycling in the *set* of non-affine SFEs. In fact, our simulations suggest that this may be exactly what bidders are likely to do.

The assumption of continuous supply functions limits these studies somewhat, as all real-world markets involve some form of a discrete schedule. Two papers that look at outcomes under specific formats for the supply schedule are Anderson and Xu (2002) and Day and Bunn (2001). The first of these examines an Australian-style power market with a two stage bidding process. First, each bidder announces a set of prices. Then a bidder specifies a supply schedule for each hour of the day – using only the prices it specified at the first stage. Anderson and Xu study the second stage of this process, which has similarities to a quantity-setting game, under a Cournot-like best response dynamic. They find that equilibria involving low markups above marginal cost are stable, but the status of other equilibria is unclear; they offer an example in which the dynamics settle into a stable cycle. Day and Bunn study a similar setting in which an essential difference is that a bidder's quantity offers at the pre-set prices are connected by straight lines to form a piecewise linear supply curve. A second difference is that quantities are restricted to discrete amounts (each taken to represent a single plant). Under a boundedly rational "better reply" dynamic (a bidder moves one plant per period to a different price bin in order to achieve the greatest possible improvement in profits), they present simulations in which bidding settles into a neighborhood around the "Cournot SFE", that is, the SFE with the highest possible prices and profits. While finding a useful synthesis of this diverse array of different techniques and results is a challenging task, one broad conclusion that can be drawn is that we should expect qualitatively different outcomes when bidders select prices and quantities for their supply schedules sequentially rather than simultaneously.

Perhaps our main contribution to this discourse on is to emphasize the importance of step function bidding in contributing to marketplace instability. Because of the self-sustaining nature of supply function equilibria – the steeper your supply curve is, the steeper my optimal response tends to be – small strategy adjustments can end up being dramatically amplified. Step function bids always present opportunities for profitable undercutting or demand reduction, even when the steps are potentially quite small, and these small adjustments can act as seeds for substantial volatility.

# 5 Conclusion

It is commonplace in the burgeoning literature on competition in deregulated power markets to make relatively broad, and sometimes *ad hoc*, simplifying assumptions about how firms will behave. Indeed, given the complexity of some of these markets, some such assumptions are essential if any headway is to be made. However, on occasion, simplification can be misleading; the tendency to focus on a taxonomy of competition that includes only Cournot and Bertrand (and increasingly on supply function equilibria as well) is a good example of this. Quite often, power markets operate under relatively idiosyncratic rules that can generate a rich assortment of different competitive outcomes. In order to develop a better understanding for the spectrum of these possible outcomes, basic game theoretic modeling of particular market rules can be quite useful. This paper develops such models for market rules resembling those in the California power markets, and analyzes them from several angles of attack. Our analysis is preliminary, and could be extended in a number of directions. On equilibrium existence, for example, further work is needed on the interlinkages between markets and the potential impact of non-convexities (like start-up costs). It would also be useful, although difficult, to develop more complete analytical characterizations of equilibrium bidding behavior to complement our simulation results. Such a characterization would be helpful in assessing the robustness of modeling tools like supply function equilibria.

# 6 Appendix

**Proposition 1: Sketch of proof:**

One form or continuity which ensures convergence of the equilibria of a sequence of finite games to equilibrium of the continuous game is upper-semicontinuity. Unfortunately, it turns out that these payoffs need not be upper nor lower semi-continuous. But, for at least the two-player symmetric game, we can appeal to a result of Dasgupta and Maskin to show existence when the strategy set is limited to setting one price for the entire capacity. What is required is that in some sense the game is that at any jump discontinuities, the two players' payoffs jump in opposite directions. Note, that strategies can be characterized as prices, $p_j, j = a, b$, with $0 \leq p_j \leq p^c$, where $p^c$ is the maximum price for which demand is non-negative. In what follows, we assume that demand is allocated evenly when two firms quote the same price for one step, and that overall demand is continuous. We let $\Pi_j(p_A, p_B)$ denote the profit of firm $j$ at $(p_A, p_B)$. We assume, for sake of simplicity that there are no costs. In what follows, we can show that an equilibrium exists in the two person game when $N = 1$ and the two firms have equal capacities.

The profits are continuous except where $p_i = p_j$ for some two prices offered by the two different firms.

To show existence, it suffices, by Theorem 5b of Dasgupta, Maskin (1986a), to show that for each $p \in [0, P]$,

(i) $lim_{pA \to p^-, p_B \to p^+} \Pi_i(p_A, p_B) \geq \Pi_i(p, p) \geq lim_{pA \to p^+, p_B \to p^-} \Pi_(p_A, p_B)$ and

(ii) $lim_{pA \to p^-, p_B \to p^+} \Pi_j(p_A, p_B) \leq \Pi_i(p, p) \leq lim_{pA \to p^+, p_B \to p^-} \Pi_j(p_A, p_B), ji.$

where the left (right) inequality is strict in the first expression if and only if the right (left) inequality is strict in the second expression.

The discontinuities in profits occur when $p_A = p_B = p$. At such a point, if $p_A$ approaches $p$ from below, then (ii) is satisfied for firm A and (i) is satisfied for firm B.

Therefore, by Theorem 5b of Dasgupta-Maskin (1986a), the game has an equilibrium in mixed strategies.

Note that in an N-person auction, one firm's gains are not solely at the expense of one rival. Moreover, the sum of the payoff functions need not be upper semi-continuous. For this reason, Theorem 5 and its corollaries in Dasgupta and Maskin (1986a) do not apply directly. Dasgupta and Maskin (1986b) considered another game, of insurance contract competition, which failed to satisfy the upper-semicontinuity of the sum of payoffs as well. The payoff functions here, as in the model of Section 4 of Dasgupta-Maskin (1986b), are continuous except on a set of measure 0 that satisfy (A1) of Dasgupta and Maskin (1986a), i.e., $U_j = U_k$ for some $j \neq k$. As in Theorem 5 of Dasgupta and Maskin (1986b), we can define modified payoffs which have an upper-semicontinuous sum, and also satisfy a version of weak lower semi-continuity, labeled Property $\alpha^*$ in Dasgupta-Maskin (1986a). Given this is the case, then an M firm game with N steps will have an equilibrium in mixed strategies, with the modified payoffs. Property á* essentially ensures an atomless equilibrium distribution for the modified game. And if the distribution is atomless, the same strategies will be an equilibrium for the original game, as the payoffs can only differ on sets of measure zero.

# References

Anderson, E., and A. Philpott (2002): "Using Supply Functions for Offering Generation into an Electricity Market," *Operations Research*, 50(3), 477–89, AGSM Working Paper 02-002.

Anderson, E., and H. Xu (2002): "Nash Equilibria in Electricity Markets with Discrete Prices," Australian Graduate School of Management Working Paper 02-002.

Baldick, R., and W. Hogan (2002): "Capacity Constrained Supply Function Equilibrium Models of Electricity Markets: Stability, Non-decreasing Constraints, and Function Space Iterations," UCEI POWER Working Paper.

Dasgupta, P., and E. Maskin (1986a): "The Existence of Equilibrium in Discontinuous Economic Games I: Theory," *Review of Economic Studies*, 53(1), 1–26.

——— (1986b): "The Existence of Equilibrium in Discontinuous Economic Games II: Applications," *Review of Economic Studies*, 53(1), 27–41.

Day, C., and D. Bunn (2001): "Divestiture of Generation Assets in the Electricity Pool of England and Wales: A Computational Approach to Analyzing Market Power," *Journal of Regulatory Economics*, 19(2), 123–41.

Green, R., and D. Newberry (1992): "Competition in the British Electricity Spot Market," *Journal of Political Economy*, 100(5), 929–53.

Klemperer, P., and M. Meyer (1989): "Supply Function Equilibria in Oligopoly Under Uncertainty," *Econometrica*, 57(6), 1243–77.

McKelvey, R., and T. Palfrey (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10(1), 6–38.

Rudkevich, A. (1999): "Supply Function Equilibrium in Power Markets: Learning All the Way," TCA Technical Paper 1299-1702.

von der Fehr, N.-H., and D. Harbord (1993): "Spot Market Competition in the UK Electricity Industry," *Economic Journal*, 103(418), 531–46.

,

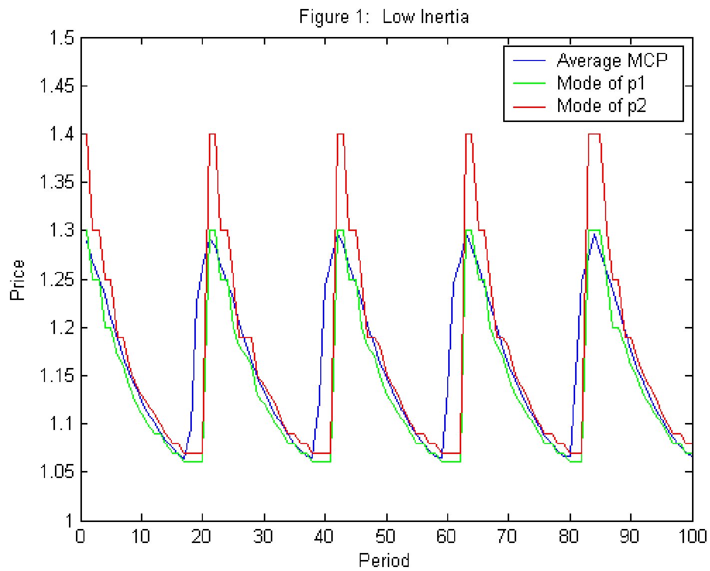


Figure 1:

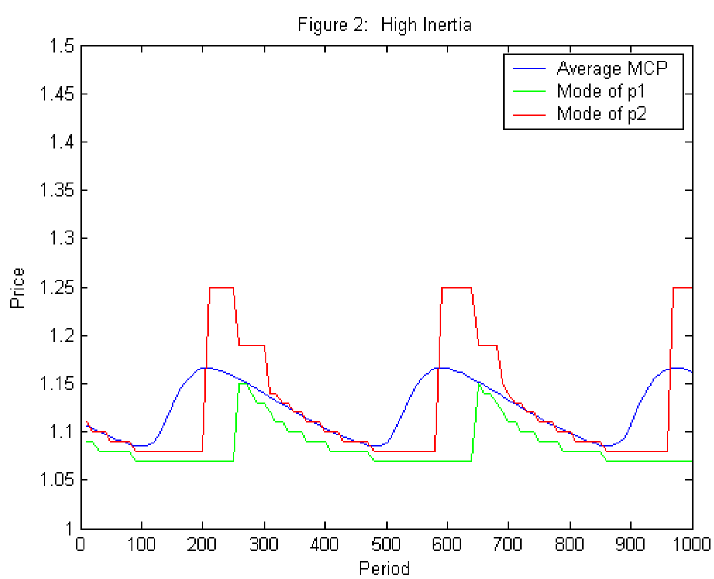


Figure 2:

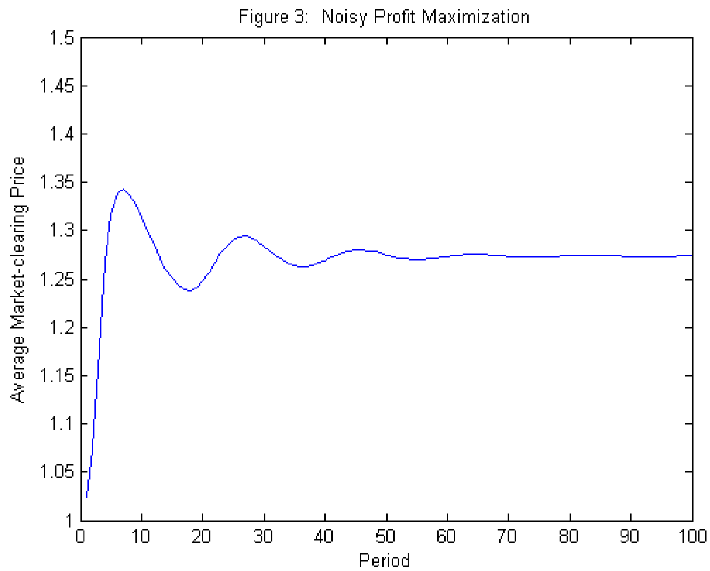


Figure 3:

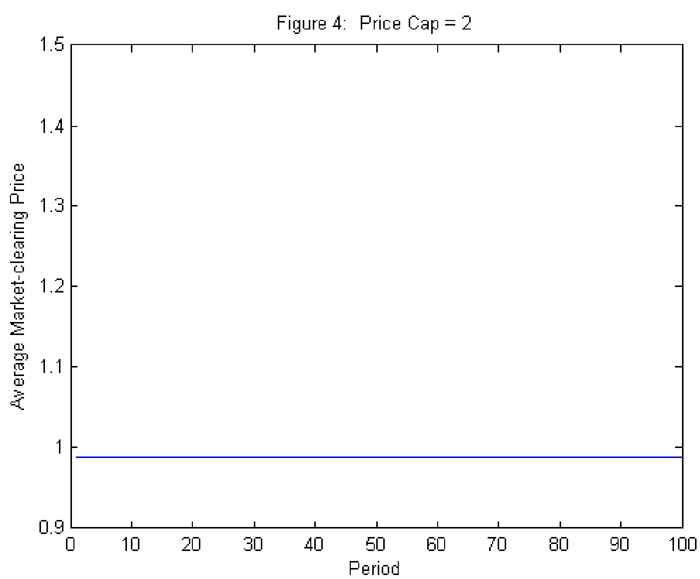


Figure 4:

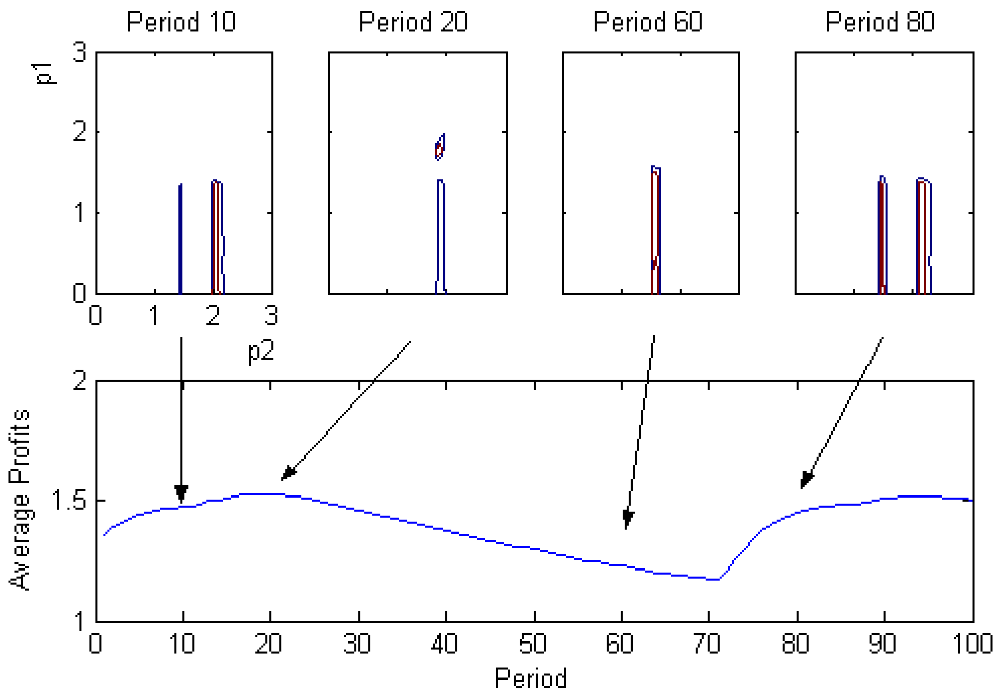


Figure 5:

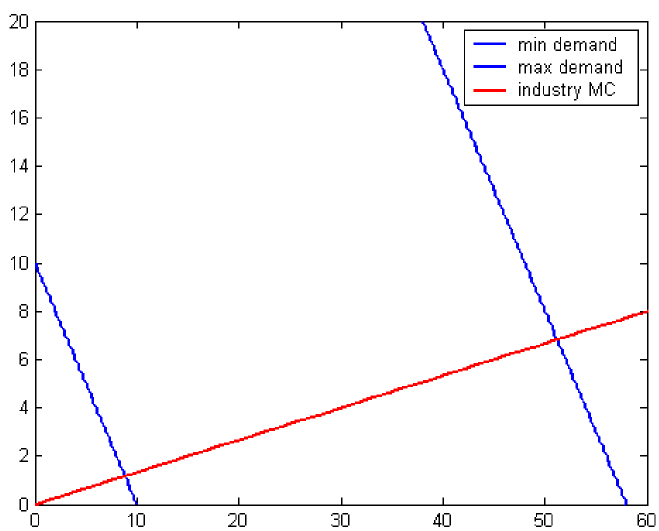


Figure 6:

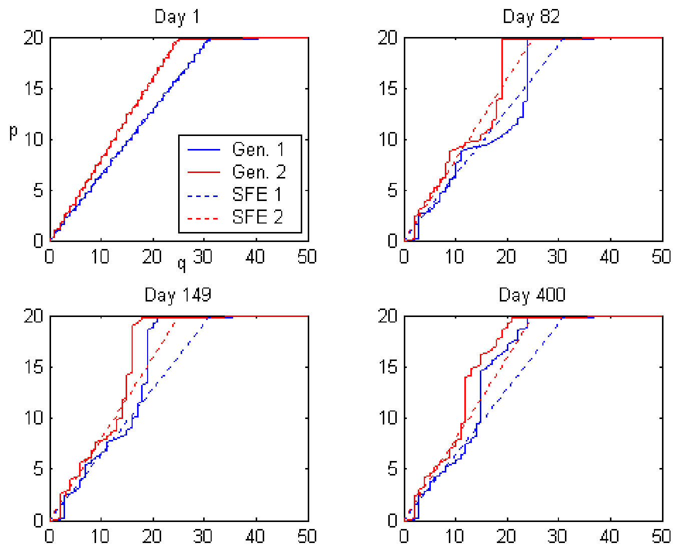


Figure 7:

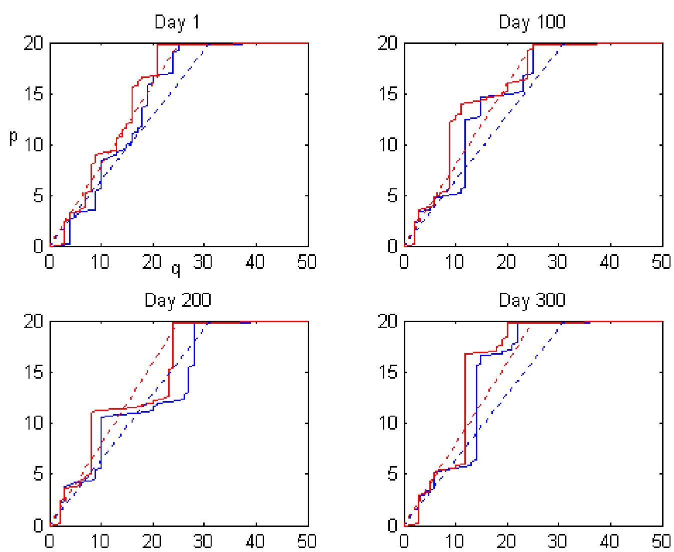


Figure 8:

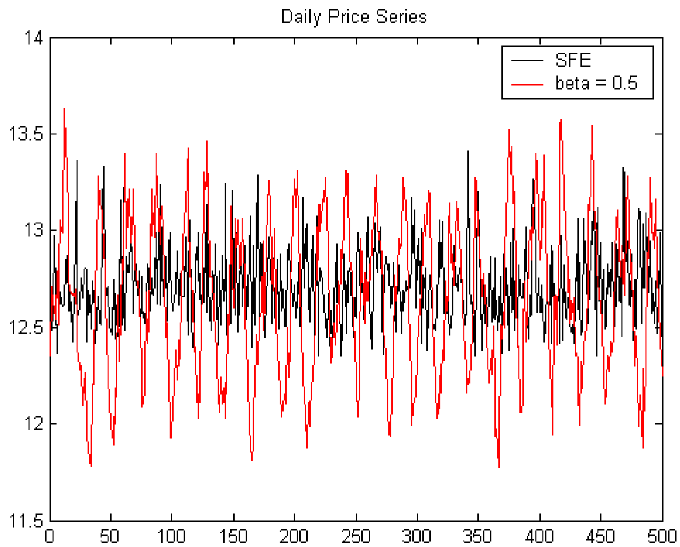


Figure 9:

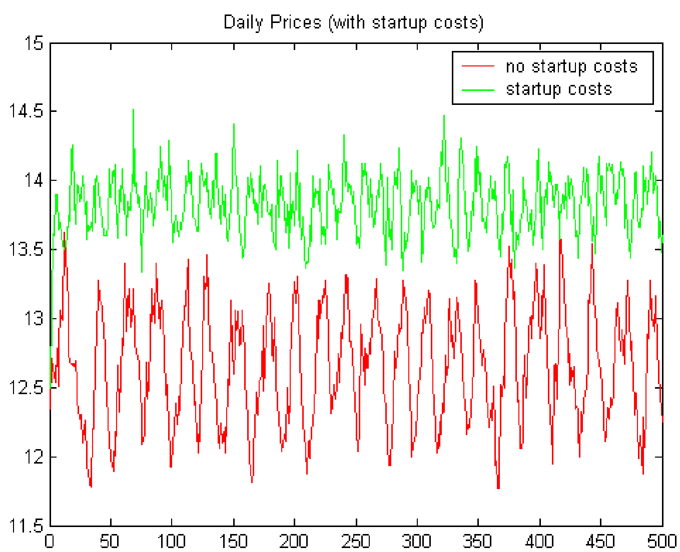


Figure 10: