Understanding Strategic Bidding in Multi-Unit Auctions: A Case Study of the Texas Electricity Spot Market

Ali Hortaçsu and Steven L. Puller*

June 2007

Abstract

We examine the bidding behavior of firms competing in the newly created spot market for electricity in Texas, where electricity generating firms submit hourly supply schedules to sell power. We characterize an equilibrium model of bidding into this market and use detailed firm-level data on bids and marginal costs of production to compare actual bidding behavior to theoretical benchmarks derived from our model. We find that firms with large stakes in the market performed close to the theoretical benchmark of static profit-maximization. However, several smaller firms utilized excessively steep bid schedules that significantly deviated from this benchmark. Our results suggest that payoff scale has an important effect on firms' willingness and ability to participate in complex, strategic market environments. We find that the bidding behavior of the smaller firms contributed significantly to productive inefficiency in this new market, although the smaller firms moved closer to the theoretical bidding benchmarks over time.

*University of Chicago, Texas A&M University. Email:hortacsu@uchicago.edu and puller@econmail.tamu.edu. We thank seminar participants various universities and conferences. We are grateful for assistance with data and institutional knowledge from Parviz Adib, Tony Grasso, and Danielle Jaussaud at the Public Utility Commission of Texas. The Editor Igal Hendel, two anonymous referees, Severin Borenstein, Jim Bushnell, Stephen Holland, Marc Ivaldi, Julie Holland Mortimer, Shmuel Oren, Peter Reiss, Steve Wiggins, Joaquim Winter and Frank Wolak provided very helpful comments. Hailing Zang, Anirban Sengupta, Jeremy Shapiro, and Joseph Wood provided capable research assistance. Hortaçsu was a visitor at Harvard University and the Northwestern University Center for the Study of Industrial Organization during the course of this research, and gratefully acknowledges both institutions' hospitality and financial support. Puller was a visitor at the University of California Energy Institute's Center for the Study of Energy Markets, for whose hospitality he is grateful. Hortaçsu acknowledges financial support from the National Science Foundation (SES-0449625) and Puller acknowledges support from the Texas Advanced Research Program (010366-0202).

1 Introduction

Many recent empirical analyses of oligopoly competition, including the analysis of bidding in auction markets, rely crucially on assumptions regarding the model of firm behavior. In a typical paper, a researcher has data on firms' prices or bids and seeks to estimate the underlying costs of production or valuation of the auctioned object. By assuming that firms behave according to a particular strategic equilibrium model of profit maximization, the researcher can map firms' observed pricing or bidding decisions into their unobserved costs or valuations.[1] The inferences drawn from such approaches rely on the assumed strategic behavior. In most instances, testing the validity of a particular equilibrium model is left to the laboratory, where the researcher assigns costs/valuations to subjects and compares the subject behavior to the behavior predicted by the equilibrium model of competition. Outside of the laboratory, it is difficult to assess equilibrium models because data usually are not available on bidder costs/valuations.

In this paper, we analyze the recently restructured electricity market in Texas, where we have the advantage of having very detailed bidding and marginal cost data on a rich cross-section of generation firms. Our study builds on the work of Wolfram (1999) and Sweeting (2005) on the electricity market of England and Wales, Wolak (2003a) on Australia, Borenstein, Bushnell, Wolak (2002) and Puller (2007) on California, who also use marginal cost data to investigate theories of oligopolistic firm behavior.[2] These data allow us to construct benchmarks for each firm's optimal bid functions and compare those to the actual bids. Thus, we "measure" the extent to which the different firms in our sample maximize expected profits and explore reasons for observed deviations from (static) profit maximization.

To construct profit maximization benchmarks, we need to account for institutional complexities of the Texas electricity market in our theoretical model. In this market, most of the electricity is traded through bilateral forward contracts between generators and users of electricity. To meet last-minute changes in aggregate electricity demand that fall beyond or below contracted quantities, generation firms submit bids to adjust their production into an hourly "balancing market" administered by ERCOT (Electricity Reliability Council of Texas). Firms participating in this market include large formerly regulated utilities, merchant generating firms, and small municipal utilities and power cooperatives. The hourly market clearing mechanism is a multi-unit, uniform-price auction – firms bid supply functions and winning sellers earn the price at which aggregate supply bids equal demand.

We model competition in the hourly balancing market using Wilson's (1979) "share auction" formulation.[3] In our model, firms choose bid functions to maximize expected profits under uncertainty coming from two sources. First, total demand for balancing power is determined by events such as weather shocks, so it is stochastic from the perspective of the bidder. Second, firms cannot predict the equilibrium bids of

[1] See Bresnahan (1989) on empirical models of oligopolistic markets and Hendricks and Porter (forthcoming) and Athey and Haile (forthcoming) for reviews of the recent empirical literature on auctions. Klemperer (2003) and Einav (2004) clarify the connections between oligopoly models and auction models.

[2] Several other papers use "outside" estimates of marginal cost to measure price-cost margins and test strategic oligopoly models. For example, Genesove and Mullin (1998) analyzes the sugar industry in early 20th century. Hendricks, Porter and Boudreau (1987) investigate the ex-post returns of bidders in Outer Continental Shelf oil and gas lease auctions using drilling and production data from auctioned tracts. Bajari and Hortaçsu (2003) use experimental data with assigned bidder valuations to gauge the performance of structural econometric models of auctions.

[3] Ausubel and Cramton (2002), Wang and Zender (2000), Hortaçsu (2002b), Viswanathan, Wang and Witelski (2002) develop this theoretical framework further. See Cramton (2003) for a particularly accessible account of these models.

their rivals with certainty because each firm possesses private information on their own forward contracts to supply power. These contract obligations determine the firms' net buy or net sell positions in the balancing market, and therefore affect bidding incentives.[4] Because they are private information, these obligations generate uncertainty from the perspective of other bidders. We characterize the Bayesian-Nash equilibrium of bidding into the balancing market. We show that when supply schedules are restricted to be additively separable in price and private information on contract quantities, equilibrium bid schedules are "ex-post optimal" and therefore are straightforward to compute, given information on firms' contract positions and their marginal costs of generation.

A similar benchmark, "best-response bidding," is utilized in Wolak (2003a) and Sweeting (2006) to analyze bidding behavior in the Australian and England & Wales electricity markets. Wolak (2003a) builds on the Klemperer and Meyer (1989) supply-function equilibrium (SFE) model to motivate this benchmark.[5] In the SFE model, which is nested by the Wilson (1979) model, the source of uncertainty is aggregate demand shifts, and firms do not possess private information regarding each others' marginal costs or contract positions. While the assumption that generation costs are common knowledge across bidders is realistic in electricity markets, where a lot of information is publicly available about each firm's generation technology and the spot price of fuel, it is less likely that firms have accurate information regarding each others' contract positions on a high-frequency basis. Our modelling framework allows for the presence of private information regarding contract positions, and provides conditions under which "best-response bidding" can be supported as an ex-post optimal (Bayesian-Nash) equilibrium outcome. The ex-post optimality feature of this benchmark allows us to avoid pooling data across auctions, and hence avoids potential measurement biases due to the presence of unobserved (to the econometrician) factors that vary from auction to auction. In Section 4.3.2, we also provide a test of the conditions needed for ex-post optimality. Along with the ex-post optimal bidding benchmark, we also test the ex-ante optimality of bidding. We assess whether one can outperform the bidders by constructing the ex-post optimal benchmark, conditioning on *past* realizations of the residual demand curve; which is available to the bidders.

An important requirement for constructing all of these bidding benchmarks is having hour-to-hour information on contract positions. However, this information is typically not available to economic researchers. Wolak (2003a) avoids this problem by utilizing proprietary information on contracts obtained from a generation company in the Australian market.[6] Since our empirical focus is to analyze the heterogeneity of bidding performance across a wide variety of firms operating in ERCOT, and obtaining information on contract quantities for this wide array of firms was not feasible, we develop a method to infer contract quantities using marginal cost data and the observed bid function. Our method, described in Section 3, relies on a behavioral assumption that is much weaker than profit-maximization. In essence, we merely require bidders to understand that they can end up being either net *buyers*; in which case, they should

[4]Wolak (2000), Bushnell, Mansur and Saravia (2006) also point out that forward contract positions are important determinants of bidding behavior.

[5]The SFE model has been influential in the modelling and analysis of electricity markets, as in Green and Newberry (1990), Green (1992), Rudkevich (1999), Baldick, Grant and Kahn (2004), Crawford, Crespo and Tauchen (forthcoming).

[6]Sweeting (2006), in contrast, makes the assumption that the generators have 80% contract cover, and tests robustness to alternative assumptions.

try to mark-*down* the market price, or net *sellers*; in which case they should mark-*up*. This practice was acknowledged by all of the firms that we interviewed during our research.[7]

The main empirical finding of the paper is that larger firms perform closer to our benchmark for (static) profit maximization. The smaller firms tend to submit bid functions that are "excessively steep" so that these firms are not called to supply much power to the balancing market even when it is ex-post profit-maximizing to do so. In Section 5, we argue that this finding is best explained by the presence of scale economies in setting up and maintaining a successful bidding operation – an intuition confirmed by our interviews with traders in the market. Thus, the observed patterns of bidding in this market can be "rationalized" given the fixed costs of establishing a sophisticated trading operation. We discuss this cost of "sophistication" in section 5.1. Finally, we find some evidence of learning by the small firms over our sample period. The learning rate is a 10% performance improvement per year.

The observed deviations from theoretical benchmarks are quantitatively important; we find that this behavior leads to significant efficiency losses. In section 6, we describe the two sources of efficiency losses. The first is the efficiency loss due to the (optimal) exercise of market power by profit maximizing firms. The second is the efficiency loss due to the "excessive steepness" of small firms' bid schedules that we can not reconcile with expected profit maximizing behavior. When we decompose the total efficiency losses into these two components, we find, somewhat surprisingly, that the latter source of inefficiency is larger. The inefficiency generated by the smaller firms suggests that market performance could be improved by the consolidation of small firm bidding operations or the use of a market mechanism with less strategic complexity.

Our contributions to the growing literature on characterizing strategic behavior on restructured electricity markets are both methodological and policy-oriented. By applying the Wilson (1979) model on share auctions, we extend the analysis of Wolak (2003a) and Sweeting (2006) to the case where firms possess private information regarding their contract positions. We also devise a method to estimate firms' (private information) contract positions based on a weak behavioral restriction. This allows us to conduct tests of ex-post and ex-ante profit maximization for a rich cross-section of firms on this market, for which information on contracts is not available.

We also believe that our empirical results have generalizable implications for settings outside of the electricity industry. Aside from controlled experimental settings, there is limited empirical evidence on the importance of payoff scale and learning in real-world strategic environments. We present clear evidence that both mechanisms are in effect on ERCOT. We also document, however, that deviations from profit maximizing behavior are economically significant. We believe that this latter finding has important implications for market design.

The outline of the rest of the paper is as follows: in section 2, we describe the institutional setting of the Texas electricity balancing market. In section 3, we model strategic bidding in this market as a uniform-price share auction. We discuss the empirical implications of our model. In section 4, we compare our theoretical benchmarks with the actual bids in the data. Section 5 discusses these results and explores

[7] Wolak (2003a) demonstrates how one can obtain contract quantities under the assumption of profit maximization. Since our goal is to test profit-maximization, the method proposed by Wolak is not feasible in our context.

the role of payoff scale in explaining deviations from ex-post optimal bidding. Section 6 calculates the efficiency losses and section 7 concludes.

2 How Bidding Occurs in ERCOT's Balancing Energy Market

We analyze electricity transactions that occur through spot market auctions. In the Texas wholesale electricity market, most trades occur via bilateral agreements. In addition to this bilateral market, ERCOT, the system operator, conducts an auction to balance supply and demand in real-time. Approximately 2-5% of energy is traded in this "spot market," called the Balancing Energy Services auction, and we analyze the bidding into this auction.

The mechanics of electricity transactions on this market can be summarized as follows.[8] One day before production and consumption occur, ERCOT accepts schedules of quantities of electricity to inject and withdraw at specific locations on the transmission grid. Firms' day-ahead schedules are fixed quantities that do not vary in price. The day-ahead schedules may differ from the firms' forward contract position. Those supply ("generation") and demand ("load") schedules also may differ from the actual production and consumption in real-time for a variety of reasons such as an unpredictably hot day or an outage at a powerplant. The balancing market operates in real-time to balance actual load and generation. Depending upon whether more or less power is needed than the day-ahead schedule, the balancing demand can be positive or negative. As the time of production and consumption nears, ERCOT estimates how much balancing energy is required. Because there are virtually no sources of demand that can respond to prices in real-time, balancing demand is perfectly inelastic.

Bidders offer to increase ("INC") and decrease ("DEC") the amount of power supplied relative to their day-ahead schedule. Firms submit *hourly* INC and DEC bid schedules that must be increasing monotonic step functions with up to 40 "elbow" points (20 INC and 20 DEC bids). These bids may be changed up until one hour prior to the operating hour. The bid schedules apply to each of the four 15-minute intervals of the hour. In addition, the bidder observes real-time information on its units' generation, the load it is obligated to serve, and its net short or long position in the balancing market.[9]

Procurement occurs using a uniform-price, multi-unit auction. ERCOT clears the balancing market four times every hour by intersecting the hourly aggregate bid function with the 15-minute perfectly inelastic demand function. A generator called to INC is *paid* the market clearing price for all INC sales (i.e. production beyond the day-ahead schedule). Likewise, a generator called to DEC *pays* the market clearing price for the quantity of output reduced. A generator that DECs reduces output and purchases power from ERCOT at the market clearing price to satisfy existing contract obligations.

Bidders appear to have a great deal of information on the competitive environment when they choose their bid functions. Our conversations with several market participants suggest that traders have good

[8] See Wilson (2002) and Joskow (2000) for more detailed discussions of transactions in restructured electricity markets. Further details on the ERCOT market are in Baldick and Niu (2005).

[9] Examples of both the bidding and operations interfaces can be found in a supplementary materials section at *http://econweb.tamu.edu/puller/HPsupplementary.htm*. We thank the real-time trading desk at Bryan Texas Utilities and Reliant/TexasGenco for allowing us to visit and observe their trading desk and allowing us to interview their traders. A variety of other market participants provided very helpful insights into trading in the balancing market.

information on their rivals' marginal costs. The powerplants in Texas have very similar production technologies, and there are publicly available data on the fuel efficiency of each generating unit. Traders appear to know the major generating units that are on and off-line at any point in time. In addition, some market participants purchase real-time data on the generation of large rival plants from an energy information company named Genscape that developed a technology measuring real-time output with remote sensors installed near the transmission lines out of a plant. This can be useful strategic information not only when initial bids are submitted but also if the trader wants to change the bids an hour before the market clears.

However, even if firms have a good idea of each others' marginal cost schedules, they seldom have information about competitors' contract obligations. These contracts are signed bilaterally in an over-the-counter market where it is difficult to monitor transactions, and they are seldom publicized. As pointed out by Wolak (2000, 2003a) and will become clear in the next section, these contract obligations significantly affect bidders' incentives to exercise market power, hence this constitutes a very important source of private information.

The information available to the bidders may allow them to accurately estimate the *distribution* of their residual demand curve in an upcoming auction. The residual demand is the perfectly inelastic total balancing demand minus bids by all other firms. Total demand is stochastic, but shocks to total demand (e.g. weather) only shift residual demand left and right in a parallel fashion. The distribution of rival bids can be inferred in two ways. A trader equipped with knowledge of rivals' marginal costs and the distribution of their contract positions can compute (as we do in section 3) the equilibrium mapping of costs and forward positions to bids. Alternatively, and perhaps more plausibly, every trader has access to the aggregate supply schedule with a two day lag. By knowing the recent aggregate supply curve as well as her own recent bids, the trader can infer the recent aggregate bids by all rivals. To the extent that rival bids several days before are similar to current rival bids, the trader can infer the shape of residual demand before placing her bids.

Congestion of the transmission grid poses a slight complication for our analysis. ERCOT is geographically divided into several zones. If transmission lines between zones are not congested, the balancing market is a single unified market across all Texas. However, when lines become congested, ERCOT divides the state into separate markets with different market clearing prices. During congested hours, bids by some firms are technically feasible while bids by others are not. Therefore, our analysis uses only uncongested hours, which represent 74% of the time intervals during our sample period.

We analyze weekdays from the beginning of the market in September 2001 to January 2003. Although we can analyze any period of the day, we focus on 6:00-6:15pm because the most flexible type of generators that can respond to balancing calls without large adjustment costs are likely to be online during this peak hour of the day. The average number of megawatts (MW) traded (both positive and negative) during this time interval is 915 MW.[10] Demand for balancing power is both negative and positive in many hours – the interquartile range is from -709MW to +615MW.

[10] To compare this with other intervals, the average MW transacted across all intervals varies only slightly over the day with an average volume of 890. Demand can vary substantially during the 15-minute intervals within a bidding hour. The average difference between the minimum and maximum balancing demand is 617 MW for the 6:00-7:00pm bidding hour. Because firms can submit only one bid function for each hour, optimal bids must cover a wide range of possible demand quantities.

The Texas market consists of a variety of investor-owned utilities, independent power producers, municipal utilities and power cooperatives. The two largest players are the two large former incumbent utilities: TXU and Reliant, owning 24% and 18% of capacity, respectively. The other major investor-owned utilities (and capacity shares) are Central Power & Light (7%) and West Texas Utilities (2%). Municipal utilities (e.g. City of San Antonio Public Service (8%) and City of Austin (6%)) and power cooperatives (e.g. Lower Colorado River Authority (4%)) also sell to the balancing market. Finally, there are a large number of merchant generation firms of various sizes (e.g. Calpine (5%)). The generation technology is primarily natural gas and coal with small amounts of nuclear, hydroelectric and wind generation.

Each firm has some level of contract obligation to directly supply retail customers or provide power to utilities that serve retail load. If the day-ahead generation schedule differs from the contract obligation, the firm is either long or short entering into the balancing market. This residual contract position affects bidding incentives in the auction (Wolak, 2003a). We propose a method to estimate the long or short position in section 3.

3 An Equilibrium Model of Bidding on the ERCOT Balancing Market

We now provide a model of strategic behavior in this market that incorporates the uncertainty faced by each firm when making their bidding decisions. To do this, we follow the uniform price share auction setup of Wilson (1979). For ease of exposition, assume that firms sign forward contracts and then bid *all* electricity through the auction, i.e. assume that there is no day-ahead scheduling. This allows us to simplify notation. After deriving the model, we explain how we map this model into a market with a day-ahead schedule followed by a balancing auction.

We index the costs of generation (at time t) of the N firms in this market by $\{C_{it}(q), i = 1, .., N\}$. We take total demand $\tilde{D}_t(p) = D_t(p) + \varepsilon_t$ to be the sum of a deterministic price-elastic component and a stochastic constant term.[11] Prior to the auction, each firm has signed contracts to deliver certain quantities of power each hour, given by QC_{it}, at a fixed price PC_{it}. We assumes that these contracts have been written a long enough time ago that they are taken as "sunk" decisions from the perspective of a bidder making real time decisions on the balancing market.[12]

In each time period t, each firm simultaneously submits a supply schedule, $S_{it}(p, QC_{it})$, which we restrict to be continuously differentiable, with bounded derivatives. Given the supply schedules of each firm, the auctioneer computes the market clearing price, p_t^c, which satisfies the below market clearing condition:

$$\sum_{i=1}^{N} S_{it}(p_t^c, QC_{it}) = \tilde{D}_t(p_t^c) \tag{1}$$

[11] In our application, we assume that demand is perfectly inelastic. However, our model generalizes to the case where $D'(p) \leq 0$.

[12] Although the analysis of the two-stage game with endogenous contract quantity choices adds a further and interesting strategic dimension, we found that traders who negotiate forward contracts and traders who bid in the balancing (real-time) market are distinct, and often belong to different trading desks. Thus our model can be thought of as applying to the real-time trader.

Each firm gets paid $S_{it}(p_t^c, QC_{it})p_t^c$ due to the uniform pricing rule. Hence, firm i's ex-post profit, upon the realization of market clearing price p_t^c is:

$$\pi_{it} = S_{it}(p_t^c, QC_{it})p_t^c - C_{it}(S_{it}(p_t^c)) - (p_t^c - PC_{it})QC_{it}$$

The firm's payoff from its contract position is $-(p_t^c - PC_{it})QC_{it}$, because it has to refund its customers any differential between the contract and market prices for the contracted sales. Wolak (2003a) has shown this is identical to a contract for differences.

The most important source of uncertainty in the profit equation above is p_t^c, the market clearing price at time t. In a strategic equilibrium, the uncertainty in p_t^c, from the perspective of firm i, is due to two factors: the uncertainty in market demand, $\tilde{D}_t$, and the unobserved components of i's competitors' profit maximization problems, i.e. the contract positions and prices of rival firms, $\{(QC_{jt}, PC_{jt}), j \in -i\}$.

Following the discussion in Section 2, each firm bidding into the market is assumed to know its rivals' total cost functions. This knowledge would be sufficient to calculate equilibrium bids by rival firms, except the firm does not know its rivals' contract positions. Thus, following Wilson (1979), we look for a Bayesian-Nash equilibrium characterization of the game, in which firms' strategies are of the form $S_{it}(p, QC_{it})$.

To characterize a Bayesian-Nash equilibrium, we define a probability measure over the realizations of the market clearing price, from the perspective of firm i, conditional on firm i's private information contract quantity, QC_{it}, and the fact that firm i submits the supply schedule, $\hat{S}_{it}(p)$, while his competitors are playing their equilibrium bidding strategies, $\{S_{jt}(p, QC_{jt}), j \in -i\}$:

$$H_{it}(p, \hat{S}_{it}(p); QC_{it}) \equiv \Pr(p_t^c \leq p | QC_{it}, \hat{S}_{it}(p))$$

Utilizing the definition of the market clearing price in equation (1), we can rewrite this probability distribution as:

$$\begin{aligned} H_{it}(p, \hat{S}_{it}(p); QC_{it}) &= \Pr(\sum_{j \epsilon -i}^{N} S_{jt}(p, QC_{jt}) + \hat{S}_{it}(p) \geq \tilde{D}_t(p) | QC_{it}, \hat{S}_{it}(p)) \\ &= \int_{QC_{-it} \times \varepsilon_t} 1\{\sum_{j \in -i} S_{jt}(p, QC_{jt}) + \hat{S}_{it}(p) \geq D_t(p) + \varepsilon_t\} dF(QC_{-it}, \varepsilon_t | QC_{it}) \end{aligned}$$

where the first line follows from the fact that the event "$p_t^c \leq p$" is equivalent to there being excess supply at price p. In the second line, $1\{.\}$ is the indicator function for the enclosed event, and $F(QC_{-it}, \varepsilon_t | QC_{it})$ denotes the joint distribution of the vector of contract quantities, $\{QC_{jt}, j \in -i\}$, and the demand noise, ε_t, conditional on the contract position of bidder i, QC_{it}.[13]

We now rewrite the bidder's expected utility maximization problem, where $U(\pi)$ is the utility enjoyed by the bidder from making π dollars of profit. This general utility formulation allows for both risk averse

[13] Observe that we have not imposed independence on this joint distribution; contract quantities and demand noise can be correlated. However, as will be evident from the bidder's objective function below, this is not a common value environment, as other bidders' contract quantities do not enter into the bidder's ex-post utility.

and risk neutral firms.

$$\max_{\hat{S}_{it}(p)} \int_{\underline{p}}^{\bar{p}} U\left(p\hat{S}_{it}(p) - C_{it}(\hat{S}_{it}(p)) - (p - PC_{it})QC_{it}\right) dH_{it}(p, \hat{S}_{it}(p); QC_{it})$$

where the expectation is taken over all possible realizations of the market clearing price, weighted by the probability density, $dH_{it}(p, \hat{S}_{it}(p); QC_{it})$. The Euler-Lagrange necessary condition for the (pointwise) optimality of the supply schedule $S^*_{it}(p)$ is given by:[14]

$$p - C'_{it}(S^*_{it}(p)) = (S^*_{it}(p) - QC_{it}) \frac{H_S(p, S^*_{it}(p); QC_{it})}{H_p(p, S^*_{it}(p); QC_{it})} \tag{2}$$

where

$$\begin{aligned} H_p(p, S^*_{it}(p); QC_{it}) &= \frac{\partial}{\partial p} \Pr(p^c_t \leq p | QC_{it}, S^*_{it}(p)) \\ H_S(p, S^*_{it}(p); QC_{it}) &= \frac{\partial}{\partial S} \Pr(p^c_t \leq p | QC_{it}, S^*_{it}(p)) \end{aligned}$$

$H_p(p, S^*_{it}(p); QC_{it})$ is the "density" of market clearing price when firm i bids $S^*_{it}(p)$. $H_S(p, S^*_{it}(p); QC_{it})$ can be interpreted as the "shift" in the probability distribution of the market clearing price, due to a change in $S^*_{it}(p)$, i.e. this is the term that captures the "market power" of firm i. Notice that this derivative is always non-negative, because an increase in supply weakly lowers the market clearing price, which weakly increases the probability that the market clearing price is lower than a given price p.

The above derivation relied on the assumption that supply schedules are continuously differentiable. In reality, however, firms are allowed to bid 40 price-quantity points; restricting supply schedules to (non-differentiable) step functions. A characterization of bidding behavior in uniform price auctions when the strategy space is restricted to a discrete number of steps is pursued in McAdams (2006) and Kastl (2006a,b). Kastl (2006b) obtains the result that as the number of steps available to bidders grows without bound, necessary conditions characterizing bidding behavior in the discrete strategy game converge to the necessary conditions for the game with differentiable supply schedules.[15]

First-order condition (2) can be seen as a "markup" expression, where the markup in price above the marginal cost depends on how much market power firm i can exercise by shifting the distribution of the market clearing price through its own supply function $S^*_{it}(p)$. As an intuitive consequence, observe that if $H_S \to 0$, i.e. there is no market power, price equals marginal cost. Also, note that where $S^*_{it}(p) < QC_{it}$, the firm is a net seller and bids below marginal cost.

Finally, observe that where $S^*_{it}(p) - QC_{it} = 0$, $p = C'_{it}(S^*_{it}(p))$. This allows us to infer the unobserved contract positions of the bidders:

Proposition 1 *If $C'_{it}(S^*_{it}(p))$ is observed, one can calculate the contract position QC_{it}, by finding the quantity where the supply function of the firm intersects its marginal cost function.*

The empirical implementation of (2) requires the estimation of $H_{it}(p, S^*_{it}(p); QC_{it})$ (and its partial derivatives), for each bidder i, in every period t. $H_{it}(p, S^*_{it}(p); QC_{it})$ is the equilibrium belief of bidder i

[14] The proof can be found in Hortacsu and Puller (2005).

[15] Although this is a comforting result from the perspective of assuming differentiable bid schedules, Kastl (2006a) finds that if bidding extra points is a costly activity, the constrained optimal bidding behavior may depart significantly from equation (2). We discuss some evidence regarding this possibility in Section 5.1, footnote 31.

regarding the distribution of the market clearing price in auction t, conditional on his bidding strategy, $S^*_{it}(p)$. Obtaining econometric estimates of this probability distribution might require strong parametric assumptions regarding the specification of bidder-specific beliefs, and especially the role played by economic unobservables entering into bidders' beliefs across different time periods.[16] The latter concern is potentially the most troublesome: if bidders condition their beliefs on factors unobservable to the economist, estimating $H_{it}(p, S^*_{it}(p); QC_{it})$ by pooling data from a series of auctions without taking these unobservable factors into account may lead to incorrect estimates.

Observe also that the set of first-order conditions (2) for each firm, when written as a system of equations, characterizes equilibrium strategies, $S_{it}(p_t, QC_{it})$ for given primitives of the game.[17] The computation of equilibrium strategies is not a trivial task, however, since $H_{it}(p, S^*_{it}(p); QC_{it})$ is determined endogenously through the market-clearing condition (1), and depends on the joint distribution of contract positions and the distribution of demand noise.

However, we now show that the characterization of equilibrium strategies (2) is greatly simplified when the functional form of the supply function strategies, $S_i(p, QC_i)$ is restricted to a class of strategies that are *additively separable in the private information* possessed by bidders:

Proposition 2 *If supply function strategies $S_i(p, QC_i)$ are restricted to the class of strategies: $S_i(p, QC_i) = \alpha_i(p) + \beta_i(QC_i)$, the markup relation (2) is given by the familiar "inverse-elasticity" markup rule: $p - C'_i(S_i(p, QC_i)) = \frac{S_i(p,QC_i)-QC_i}{-RD'_i(p)}$ where $RD'_i(p)$ is the price derivative of the ex-post realization of the residual demand curve faced by bidder i.*

Proof. See Appendix B. ■

The intuition underlying this result is straightforward. The additive separability restriction implies that, in equilibrium, the residual demand function faced by bidders is also additively separable in its random component – i.e. all uncertainty (from the perspective of bidder i) *shifts* the residual demand curve but does not *rotate* it. Given this, it is easily seen that (subject to the concavity of the profit function), the bid function $S_i(p)$ provides a *pointwise best-response* to every possible realization of the residual demand curve. This also means that the class of additively separable equilibrium strategies, when they exist, are *ex-post optimal*, in the sense that seeing other bidders' supply functions would not change bidder i's choice of supply function.[18]

As an important caveat, however, note that the additive separability restriction is an *a priori* restriction on bidding strategies. It is not necessarily true that every specification of marginal cost functions, $C'_i(q)$ and joint distribution of contract quantities QC_i, will lead to equilibrium strategies of this form.[19] Hortacsu and Puller (2005) works through an example in which firms possess linear marginal cost curves (these

[16]Some of these assumptions and possible estimation strategies based on such assumptions are discussed in the discriminatory (pay-as-bid) share auction context by Hortaçsu (2002).

[17]The primitives of this game are the set of firms who are participating, N, their cost curves $C_{it}(q), i = 1, .., N$, the joint distribution of contract quantities and the distribution of the uncertain demand component.

[18]The additive-separability restriction appears to be crucial. Note that without this separability restriction, we can not, in general, collapse the stochastic terms (from the perspective of bidder i) into a single scalar random variable. See Appendix B for further discussion of the case where private information leads to rotations in residual demand, as opposed to shifts.

[19]In particular, for certain specifications of marginal costs, a bidder's best-response to additively-separable bidding strategies by her opponents may not be additively separable.

can be asymmetric across firms), under which equilibrium strategies are analytically characterizable, and satisfy the additive separability restriction. Based on our inspection of actual marginal cost curves, linearity appears to be a reasonable approximation. Finally, we note that the additive separability restriction is testable and we provide several tests in section 4.3.2.

This result has an immediate practical application:

Proposition 3 *Suppose supply function strategies $S_i(p, QC_i)$ are restricted to the additively separable class of strategies: $S_i(p, QC_i) = \alpha_i(p) + \beta_i(QC_i)$. Then given data on the marginal cost function, one can compute the ex-post optimal supply curve $S_i^{xpo}(p)$, which is the ex-post best response to the observed realization of the residual demand curve.*

Observe that under the above restriction, a single realization of the residual demand curve, $RD_i(p, \varepsilon, QC_{-i})$ is enough to compute $\frac{d}{dp}RD_i(p, \varepsilon, QC_{-i}) = RD_i^{'}(p)$ for all realizations. Then, for a range of prices, $p \in [\underline{p}, \bar{p}]$, one can solve the equation for S, in terms of p and QC_i:

$$p - MC_i(S) = \frac{S - QC_i}{-RD_i^{'}(p)} \tag{3}$$

to trace out $S^{xpo}(p, QC_i)$, which constitutes an ex-post best-response to all possible realizations of residual demand.[20]

Thus, although additive separability is a restrictive assumption, its imposition aids us greatly in solving for optimal supply schedules. Perhaps more importantly, ex-post optimality aids us in dealing with economic unobservables. Observe once again that the ex-post optimal supply schedule corresponding to a given marginal cost schedule is derived using a *single*, ex-post observation of residual demand. The proposed empirical procedure does not pool data across auctions, and thus avoids the problem of making strong assumptions as to how to model the role of unobservables.

4 Analysis of Observed Bid Schedules

In our empirical application, we implement Proposition 3. We use data on bidders' marginal cost functions to calculate the ex-post optimal supply curve, which is also an equilibrium bid function under the restrictions imposed in Proposition 2. Then we compare ex-post optimal bid schedules to actual bids.

Now we explain how we map the model of section 3 into the actual market. As described in section 2, firms in ERCOT do not bid all electricity through the auction. Rather, the firms submit a fixed-quantity day-ahead schedule, and then compete in the balancing auction to increase or decrease supply from that day-ahead quantity. In order to implement Proposition 3 to model bidding in the balancing auction, we modify our marginal cost function and contract quantity. First, we account for the fact that the day-ahead schedule implies that some of the firm's generating capacity is already committed to produce the day-ahead quantity. Thus, we shift the total marginal cost function to the left by the day-ahead quantity. This balancing marginal cost function represents the marginal cost of increasing output and the marginal savings of reducing output relative to the day-ahead schedule. Second, the day-ahead quantity is used

[20]This could be extended to show that given QC_i, one can compute the entire marginal cost curve rationalizing a supply curve $S_i(p)$ observed in the data, using a *single* realization of the residual demand curve.

to satisfy some of the firm's forward contract positions. Any remaining contract position, the balancing contract quantity, is the QC_{it} that affects bidding into the balancing market.[21] Therefore, in the context of our notation in section 3, $S_{it}(.)$ represents supply bids into the balancing market, $C_{it}(q)$ represents the costs/savings of increasing/reducing output relative to the day-ahead schedule, and QC_{it} represents the quantity that the firm is long or short on its contracted sales after the day-ahead schedule and upon entering the balancing market.

Our empirical strategy is illustrated in Figure 1. In order to measure the contract position, QC_{it}, we use Proposition 1. QC_{it} is measured as the quantity at which the actual (balancing) bid schedule intersects the (balancing) marginal cost function (point A in Figure 1).[22] Note that Proposition 1 allows us to identify the contract position under certain forms of suboptimal or non-equilibrium bidding. This identification approach is valid as long as the firm is sufficiently sophisticated to bid above (below) marginal cost when it is a net seller (buyer), even if it errs in the size of the markup. More formally, we can identify QC_{it} in equation (2) even if firms have incorrectly calculated $\frac{H_S(p,S_{it}^*(p);QC_{it})}{H_p(p,S_{it}^*(p);QC_{it})}$. Our conversations with various market participants lead us to believe that traders clearly recognize the rationale for marking up bids for quantities greater than the contract position (and vice versa), but traders have different heuristics for choosing the size of the markup.

Each firm's residual demand $RD_i(p)$ is the realized total demand minus the bids by all rival firms. Suppose RD_1 in Figure 1 is the actual realization of residual demand for firm i. We calculate $RD_1^{'}(p)$ and find the ex-post optimal (price, quantity) bid to be Point B, where the marginal revenue curve corresponding to $RD_1(p)$, MR_1, intersects the marginal cost curve MC. Also, we can calculate the ex-post optimal bid under other possible realizations of uncertainty (i.e. other realizations of total balancing demand, ε, and rivals' private information, QC_{-i}). Because each form of uncertainty acts to shift residual demand in a parallel fashion, we can consider another possible realization of residual demand as RD_2 – the realized residual demand shifted parallel to the left. Under this realization of uncertainty, the optimal bidpoint is given by Point C, where $MR_2 = MC_i(q)$. We repeat this operation by adding parallel shifts to the actual realization of the residual demand curve to find the set of ex-post optimal points for various realizations of uncertainty. The ex-post optimal bid *function* is traced out by the set of ex-post optimal bid points to generate $S_i^{xpo}(p, QC_i)$. Note that our assumption that the slope of residual demand is independent of uncertainty is necessary for the set of ex-post optimal points to be on a monotonic bid function.

We should note that the residual demand function is a step function whose derivatives are either zero or infinity, which renders the literal evaluation of equation 3 impossible. We follow two methods to address this: First, we follow Wolak (2003a) to obtain a "smoothed" version of the residual demand function to calculate the marginal revenue curve.[23] Second, we perform a grid search on the "unsmoothed" residual

[21] Notice that the lumping of the day-ahead quantity and the balancing bids does not affect the strategic nature of the game because the bidders are not provided any information about each others' actions until the market clear.

[22] One could be concerned that there is a small set of distinct contract quantities, however the empirical distribution we recover using Proposition 1 suggests that our assumption of a continuous differentiable distribution of contract quantities is very reasonable.

[23] Let $\{(p_1, q_1)...., (p_K, q_K)\}$ represent the price and incremental quantities that form the residual demand curve seen in the data. The smoothed version of this function is $RD(p) = \sum_{k=1}^{K} q_k \kappa(\frac{p-p_k}{h})$ where $\kappa(.)$ is a kernel function. With this representation, the derivative of residual demand is $RD'(p) = \sum_{k=1}^{K} q_k \frac{1}{h} \kappa'(\frac{p-p_k}{h})$ We used a normal kernel and the smoothing parameter, $h = 10MW$ throughout our analysis.

demand function and find the ex-post profit maximizing point for each parallel shift in residual demand. In practice, we found that the ex-post profitability benchmark that is generated by the grid search departs negligibly from the ex-post profitability benchmark generated by the "smoothed" residual demand.[24]

4.1 Data

We use hourly data on balancing demand and firm-level bids and marginal costs. To construct each firm's marginal cost function, we utilize data on the generating units that are operating on a given hour and the hourly declared capacity of each unit. The marginal cost of operating units represents the variable costs – fuel, operating and maintenance, and SO2 permit costs – of coal and natural gas fired units. A growing body of literature has developed on measuring the marginal cost of electricity production (for example, see Wolfram (1999), Borenstein, Bushnell and Wolak (2002), Mansur (2004a), Puller (2007), Joskow and Kahn (2002), Bushnell and Saravia (2002), Bushnell, Mansur and Saravia (2006)), and we use the same approach. Details of the data and additional institutional considerations are discussed in Appendix A.

Our "marginal cost of balancing power" function is the costs of providing more power (INCing) or the cost savings of reducing production (DECing) from the day-ahead scheduled quantity. We construct this function by first calculating the total marginal cost function in a given hour, and then subtracting the quantity that has already been scheduled one-day ahead. Total marginal cost is the marginal production cost of units that are reported by ERCOT to be operating and available in period t. It is reasonable to assume that firms produce in a least cost manner, so we stack up the marginal cost of each generating unit from cheapest to most expensive and construct the total marginal cost function. We have data on how much generation has been scheduled day-ahead by the firm. Thus the marginal cost of supplying balancing power is the total marginal cost function with the origin re-centered at the day-ahead scheduled quantity. Certain types of generating units that cannot supply power on short notice are then excluded from this marginal cost stack. Natural gas-fired units, and to a lesser extent coal units, can be adjusted on relatively short notice to other production levels. Other types of units such as nuclear, wind, and hydroelectric typically cannot respond to balancing market calls. Therefore, we exclude nuclear, wind, and hydroelectric generating units from the marginal cost portfolio for providing balancing energy.[25]

We use bid schedules for each firm for the 6:00-7:00pm hour of each non-congested weekday. To determine the sales into the balancing market, we intersect the actual bid with the realization of residual demand to determine sales. We assume a step function which is a good approximation to the actual dispatch algorithm used. This simulation of the actual auction predicts prices with only a 5% error.

4.2 Comparison of Actual Bids to Ex-Post Optimal Bids

We compare each firm's actual bids to ex-post optimal bids for each auction. Examples of this comparison for specific auctions are shown in Figure 2 which displays representative actual and ex post optimal bid functions for three large suppliers – Reliant, TXU, and Calpine – and for one small seller, Guadalupe.

[24] As an illustration, Reliant's ex post profitability is 79% under "smoothed" residual demand and 70% under the grid search.

[25] Some hydroelectric units may be able to respond to balancing calls, however these units represent less than 1% of total capacity and are primarily owned by Lower Colorado River Authority.

Competitive bidding is offering the "MC curve" and optimal bidding is offering the "Ex-post optimal bid." The intersection of the actual bids and marginal cost schedules is the contract position. For quantities above (below) the contract position, the ex-post optimal bid function is above (below) marginal cost. Visually, Reliant's bids appear much closer to the optimal bids than to the marginal cost function.

TXU is close to the ex-post optimal bid function on the INC side (Balancing MW > 0), but bids below ex-post optimal prices on the DEC side. We see this tendency for TXU to offer DECs at only very low prices throughout our sample.

Calpine offers some DEC bids but does not offer to INC supply. Although Calpine does offer INC bids in some periods, much of Calpine's bids are to DEC supply. Those DEC offers are often at prices substantially below ex post optimal bids.

Guadalupe submits bids that are much steeper than ex-post optimal. Because it is a small seller, Guadalupe's residual demand function is relatively flat as compared to the residual demand of the larger players. Nevertheless, it has some potential to bid strategically and exercise market power. However, the actual bid functions are significantly above the optimal bids in INC periods and below optimal bids in DEC periods. This suggests that bids significantly different from marginal cost are not intended as a means to exercise market power, but rather to avoid being called upon to change production from day-ahead schedules. Many small sellers show similar bidding patterns. We discuss reasons underlying these patterns in section 5.

The visual inspection of the figures are suggestive of the closeness of actual bidding behavior to ex-post optimal behavior, but a more meaningful metric to evaluate bidders' performance is to measure how much profit they have foregone ex-post by deviating from the ex-post optimal bidding schedule.

To calculate the profit deviation, we calculate the difference of the producer surplus obtained at the actual submitted price/quantity point (point D in figure 1), and the surplus obtained at the ex-post optimal point (point B in figure 1). We calculate this difference each firm-auction for 20 simulations of residual demand which we construct by adding uniformly distributed noise (with support -200MW to 200MW) to the actual demand. These simulations allow us to evaluate the optimality of several points on the bid function that could determine output under other possible realizations of residual demand. The results are generally robust to the scale of the noise added.

We also calculate the producer surplus achieved relative to a benchmark of "suboptimal" behavior to compare how much distance is closed between the benchmark of suboptimal pricing and optimal pricing.[26] One possible benchmark is behaving non-strategically and bidding marginal cost. However, it appears that the "default" behavior is to bid to avoid being called to supply balancing power. As shown in the sample figures, smaller firms choose to bid only small quantities relative to both competitive and optimal bidding. Therefore, we measure performance as the fraction of (dollar) distance between "no bidding" and ex post optimal bidding that is realized by the actual bids. Producer surplus in the balancing market is:

$$\pi_{it} = S_{it}(p_t^c)p_t^c - C_{it}(S_{it}(p_t^c)) - (p_t^c - PC_{it})QC_{it}$$

[26] We cannot use a measure such as the fraction of possible profits achieved because some firms tend to be short on their contract positions entering the balancing market, and we would have to make an assumption about the contract price. The measure we construct differences out the contract price and avoids this complication.

We calculate π_{it} for three scenarios: (1) ex post optimal bidding ($S_{it} = S_{it}^{XPO}$), (2) actual bidding ($S_{it} = S_{it}^{O}$), and (3) bidding to avoid the balancing market (($S_{it} = 0$) i.e. not bidding at all and buying/selling at the market clearing price any net short or long position on contracts). Our primary measure of performance is:

$$PercentAchieved = \frac{\pi^{Actual} - \pi^{Avoid}}{\pi^{XPO} - \pi^{Avoid}}$$

which we calculate for each firm in the market.

One should keep in mind that *PercentAchieved* is only a metric of the generator's performance in the *balancing market*. It does not account for the profitability of the vast majority of output that is sold through bilateral transactions, and therefore may understate the overall profitability of electricity sales. In order to measure the profits of bilateral sales we would require data on contract prices, but such data are not available. However, we can construct an upper bound for overall profitability by assuming that each generator is maximizing profits in the bilateral market but may be sacrificing profits in the smaller balancing market:

$$\begin{aligned} Upper\,Bound\,Total\,Percent\,Profitability = \quad & PercentAchieved * \%Sales\,in\,Balancing \\ & + 100\% * \%Sales\,in\,Bilaterals \end{aligned}$$

4.2.1 Ex-post Profitability

We compute each firm's ex-post profitability in each of the 6-7pm auctions in our sample. Table 1 compares output and producer surplus in the balancing market under actual and ex-post optimal bidding. Column 6 shows the average quantity sales if firms submit ex-post optimal bids and column 5 shows the average producer surplus from bidding ex-post optimally rather than not bidding. For the largest firms with average sales of at least 250 MWh under ex-post optimal bidding, the potential profits average about $2300/hour, with Reliant having the most "money on the table" of $4,333 each hour. Firms that would sell less than 250 MWh under ex-post optimal bidding would average about $750/hour in potential profits ranging from $2,176/hour for Air Liquide America to $31/hour for Denton Municipal Electric.

The first column displays the percent of potential profits achieved under the actual bid schedules. Reliant achieves 79% of ex-post optimal profits, or $3,422/hour of the $4,333/hour of potential profits. The next two firms closest to ex-post optimal profits are two relatively small municipal utilities – Brownsville (50%) and Bryan (45%). TXU, the second largest incumbent utility, achieves 39% of potential profits while the largest independent power producers, Tenaska Gateway and Calpine, achieve 41% and 37%, respectively. The two other major incumbent utilties – West Texas Utilities and Central Power and Light

– capture only 8% of potential ex-post profits. The other large municipal utilities earn a moderate fraction of potential profits: Austin (30%) and San Antonio (23%). The largest electricity cooperatives earn lower profits: LCRA (25%), Brazos (15%), and South Texas (3%).[27]

Several characteristics of bid schedules appear to drive the foregone ex-post profits. We examined many sets of actual and ex-post optimal bid functions similar to Figure 2. The key characteristic of underperforming bidders is that the actual bid functions tend to be "too steep" relative to optimal bids. By bidding too high during INC hours and too low during DEC hours, firms sell less than under ex-post optimal bidding and forego individually profitable sales. As shown in columns 6 and 7 of Table 1, firms sell less than the ex-post optimal quantity on average. Interestingly, the cause of foregone profits is not that firms are bidding more competitively than individually optimal, but rather that bid prices are too high.

It is important to keep in mind that these profitability measures do not necessarily represent generators' overall performance in trading electricity. Many generators may focus their strategic efforts in the bilateral market where the vast majority of transactions occur. If firms focus strategic efforts on the bilateral markets, the overall performance is substantially higher. The upper bound on total profitability, shown in column 8 of Table 1, ranges from 41% to 98% with a mean of 80%. These metrics are more likely to reflect the firms' overall performance.

4.3 Additional Tests of Profit Maximization

It is important to note that our calculations of foregone profits rely on restrictions that we place on the economic environment. In particular, we assume that bidders' equilibrium perceptions of the uncertainty results in "parallel shifts" rather than "pivots" in residual demand. If this assumption is violated, it is possible that the set of ex-post optimal price-quantity points is a "cloud" of points that cannot be connected by an increasing supply function. If this were the case, a profit-maximizing firm's solution to the (ex ante) *expected* profit maximization problem may not be the set of *ex-post* optimal price-quantity points. That is, testing whether firms maximize "ex-post" profits as we did would not be informative about whether firms maximize "ex-ante" profits. Since "ex-ante" profitability is a more accurate measure of bidder performance, we employ several testing strategies that do not impose restrictions on the relationship between uncertainty and residual demand. All of these tests suggest that the "additively separable in private information" restriction does not drive our findings of foregone profits.

4.3.1 Best-Response to Previous Rival Bids

First, we test if firms could significantly increase profits by using a simple bidding rule that utilizes only information available to traders at the time of bidding. The bidding rule we employ is a naive best response to recent rival bidding.

As discussed in section 2, traders have access to the aggregate bid function with a two day lag. Using

[27] Note that two firms, Extex Laporte and Air Liquide, earn lower profits under actual bidding than bidding to "avoid the market." Both firms, which are infrequent participants in the balancing market, have positive contract positions and relatively high cost units available. Both bid so they are called to INC despite the fact that it would be more profitable to not participate and buy its contract position from the market at a price lower than marginal cost.

their own bids from the past, traders can calculate the aggregate bids by rivals in the recent past.[28] To construct the naive best response of firm i for upcoming day t, we:

1. Use aggregate bids and own bids for day $t-3$ and calculate aggregate rival bids on day $t-3$

2. Assume rivals use the $t-3$ bid schedule on upcoming day t

3. Calculate ex-post optimal bid function for various realizations of day t total balancing demand

This algorithm uses only information available to firms when bids are submitted. We view this bidding rule as fairly unsophisticated – it uses only a small fraction of the information available to traders and it would be simple to program as an add-in to the trading interface used by the generators. We calculate producer surplus under "naive best response" bidding and compare to producer surplus under actual bidding. If uncertainty causes the ex-ante expected profit maximizing bid to differ from ex-post optimal bids, the actual profitability should be much closer to the naive best response benchmark.

Results are shown in columns 2 and 4 of Table 1. Across all firms, naive best response profits are substantially higher than actual profits and very close to ex-post optimal profits. The performance of actual bidding is significantly below the naive best response benchmark for all firms except Reliant. In fact, most bidders' performance measure of "Percent Achieved" rises very little when compared against the naive best response benchmark. Naive best response profits are very close to ex-post optimal profits – the former average \$1,193 and the latter average \$1,204. This suggests that our findings of foregone profits do not arise from the additive separability restriction. Moreover, this is consistent with the additive separability restriction. The results in columns 4 and 5 reveal that a firm's conditioning on RD_{t-3} instead of RD_t leads to negligible profit-losses. Thus, under a profit metric, the shifts in RD are purely parallel.

We also tested the restriction of how uncertainty affects residual demand using a GMM based approach inspired by Wolak (2003a). We allow for a much more general relationship between uncertainty and residual demand, and derive first-order conditions for the choice of each bidpoint. These conditions yield moment conditions that are zero in expectation under the null of expected profit maximization. The results indicate that, except for Reliant, firms in this market violate the first-order optimality conditions that need to hold for expected profit maximization. Details are in Hortacsu and Puller (2005).

4.3.2 Testing Additive Separability

Since we can estimate bidders' (private information) contract positions, we also investigate whether the additive separability restriction holds in the data. According to this restriction, assuming that the firm's and its competitors' marginal costs are unaffected, exogenous shifts in a bidder's contract position (QC) should affect the intercept of the bid function, but not the slope.

To operationalize this test, we first fit a linear function to each day's bid function and calculate a slope term. The linear specification yielded excellent fit especially in the price range between \$0 and \$30. We

[28] Note that *individual firm* bid data are only available (to firms and econometricians) with a six-month lag, so firms are unable to use Proposition 1 in real-time to estimate rival firms' contract quantities and resolve some of their uncertainty. Some of rivals' uncertainty stems from variation in QC_{it}. We find that balancing contract positions by a firm varies across time; for example, the standard deviation of QC_{it} is 449, 161, and 7 for Reliant, Calpine, and Guadalupe, respectively.

also used a linear specification to calculate the slopes of bidders' daily marginal cost functions (intercepts change very little during the sample period).

Optimal bids can depend on (the parameters of) competitors' marginal cost schedules in complex ways. Unfortunately, incorporating every competitor's marginal cost parameters in a regression specification would exhaust degrees of freedom rapidly. Therefore, instead of using competitors' marginal cost slopes and intercepts as control variables, we use the slope of the (realized) residual demand curve seen by each bidder (we do not use the intercept, as this is the uncertain part of residual demand not seen by the bidder). Note that, under additive separability, this residual demand derivative can be seen as a "sufficient statistic" encoding changes in competitors' costs.

The first regression in Table 2 reports the panel regression of the bid function slope on the estimated contract quantity, controlling for residual demand and (own) marginal cost variation, along with a linear time trend and bidder fixed effects. Although the coefficient on contract quantity is statistically significant at the 5% level and positive, the economic significance of this correlation is not large, as the amount of variation in bid function slope that is explained by changes in contract quantity is small. Specifically, the average standard deviation of daily contract quantities (across firms) is 338, and multiplying this by the coefficient estimate 0.001 leads to only 4.2% of the average bid function slope in the sample (which is 8.0).

In the second column of Table 2, we repeat the regression in the first column, but also control for auction fixed effects, which can be interpreted as factors (common across bidders) that bidders take into account when formulating their bids, but not explicitly taken into account by our simplified econometric specification. This specification yields a weaker (both economically and statistically) relationship between contract quantities and bid function slope. Note that accounting for the auction fixed effect leads to a dramatic increase in the coefficient on residual demand slope; which, in theory, should be an important determinant of bid functions. This latter fact is suggestive of an omitted variables problem in the first regression, which is ameliorated by the fixed effect specification.

In Table 3, we report the results of the first regression in Table 2 estimated at the bidder level. We display the results for those bidders who submit their own bid schedules, and do not use an intermediary "qualified scheduling entity (QSE)."[29] Reliant, the most successful bidder (in terms of ex-post and ex-ante profit maximization), appears to conform to the additive separability restriction – changes in contract quantity do not have a statistically significant effect on the slope of Reliant's bid function. The next most successful bidder in terms of ex-post profit maximization, City of Bryan, also satisfies this restriction. Although TXU violated additive separability during its first month of bidding, we fail to reject additive separability in its subsequent bidding patterns. In contrast, additive separability appears to break down for Calpine and City of Austin; however, the amount of variation in bid function slope explainable by shifts in contract position is less than 20% in both of these cases. The violation of additive separability is strongest for LCRA (Lower Colorado River Authority); variation in contract quantities can explain 50% of the variation in bid function slope.

These tests indicate that the additive separability restriction holds on average across the bidders, and that there is heterogeneity across bidders in terms of how close they come to satisfying this restriction.

[29] We discuss this sample further in Section 5.1.

The heterogeneity pattern affirms that the theoretical restriction might be a good approximation to reality; bidders who appear to perform well both from the ex-post and ex-ante profit maximization benchmarks (especially Reliant, who performs best) come close to satisfying additive separability.

5 Explaining Deviations from Optimal Bidding

This section investigates explanations for the observed deviations from static profit maximization and the considerable heterogeneity across firms in terms of performance. We explore whether the observed deviations are driven by characteristics of the firms such as the firm size, the firm type (e.g. investor-owned utility vs. municipal utility), and the generation technology. We find that the most significant determinant of performance is the size of the stakes that each firm has in the balancing market, which suggests there are scale economies to participation in the balancing market auctions. Finally, we document evidence of a modest degree of learning by the small firms.

5.1 Participation Costs and Scale Economies in Bidder Performance

We consider the hypothesis that small firms might not have sufficiently large dollar stakes to justify the fixed cost of participating in the balancing market. To do this, we first have to clarify what we mean by "participation." We find that many firms forego profits because their bid functions have a large range of prices at which quantity offered is zero. For example, in Figure 2, Guadalupe is not offering to supply additional INC energy until the price reaches \$33 nor offering to reduce production until price is \$11 despite the fact that one of its units that can both INC and DEC has marginal cost of \$28. These bid patterns effectively price the bidders out of the market for plausible realizations of residual demand (in fact, as Table 1 shows, some firms at the bottom of this table are almost always priced out of the market). In separate calculations, we find that the six firms with the highest measures of $PercentAchieved$ are called to produce balancing power in 67% of the actions. In contrast, the other firms are called to produce in only 31% of auctions, despite the fact that it's ex post optimal to produce in 89% of the auctions.

An interpretation of these bid patterns is that it does not pay for the small firms to bid the optimal markup even if this optimal markup would allow them to profit from incrementing/decrementing their power generation, or "participating" in the balancing market much more often. This interpretation is plausible if it is costly for these firms to calculate the optimal markup conditional on participating.

The fixed and variable costs of running a trading operation are likely not to be trivial. One market participant suggested that even a simple bidding operation would require an upfront expenditure of \$3 million with annual operating costs of \$1 million, and that most sophisticated trading operations could be much more expensive. The magnitude of these costs is not trivial compared to the "money-left-on-table" figures reported in Table 1.[30] Another component of this "fixed cost of participation" is institutional. ERCOT only allows certified companies, called "qualified scheduling entities"(QSEs) to submit bids in the balancing market. All other firms have to route their bids through a QSE, or contract with a QSE

[30] An annual expenditure of \$1 million corresponds to \$114 per hour for one year of operations (365 days, 24 hours), without correcting for the fact that the hour we are analyzing is a peak hour.

to conduct their bidding operations. This suggests that only firms with greater dollar stakes may find it optimal to incur the fixed costs of becoming a QSE.

The presence of such fixed costs leads to substantial heterogeneities in bidding behavior – not just in outcomes, but also in the strategies that are being used. Many bidders do not make full use of the strategy space available to them, but rather use coarse-grained bidding strategies. The bid rules allowing 40 price-quantity points afford generators a large degree of flexibility in bidding. However, none of the bidders make full use of the 40 bidpoints that they can use to trace out their optimal bidding functions. Among the firms serving as their own QSE, the firm earning the greatest fraction of ex-post profits (Reliant) also uses the largest number of bidpoints, averaging 22.2 points per bid schedule. None of the other firms use more than 13 points on average. Apparently, traders choose not to formulate refined bid strategies with desired quantities for many potential realizations of the market clearing price.[31]

Even conditional on paying the fixed cost of becoming a QSE, scale economies still appear to matter. This is clearly seen in Figure 3, which displays the relationship between bidding performance and size for generation firms that act as their own QSE. Our measure of performance is the percent of ex-post optimal profits, calculated in the manner described in section 4.2. Our measure of stakes in the balancing market is the volume of sales under ex-post optimal bidding (using other size measures, such as actual volume of sales, or firms' total capacity, yield similar patterns). There is a positive relationship between "Percent Achieved" and optimal sales volume. The figure includes the fitted linear relationship which is positive and marginally significant when all firms are included. When Bryan is excluded, the relationship is even stronger and highly significant.

We now examine the extent to which stakes are correlated with performance for the broader sample of firms in a regression context, along with other firm-specific factors that we believe might affect performance. We regress each firm's measure of "Percent Achieved" on a measure of stakes in the balancing market – the volume of sales under ex-post optimal bidding (SIZE). Also included are firm-level covariates on firm type (independent/merchant power producer, municipal utility, and investor-owned utility) and whether the firm acts as its own bidder (OWNBIDDER). Finally, we include dummy variables for whether the firm's generation technology is at least 50% comprised of two technologies that are less flexible to quick changes in output (COAL and COMBINED-CYCLE).

Results are reported in Table 4. The baseline regression in column 1 yields a result that is consistent with the "scale hypothesis:" a 1000 MWh increase in sales is associated with a 52 percentage point increase in "Percent Achieved".

[31] Kastl (2006a) brings an interesting perspective to the "coarse-grained" bidding strategies on ERCOT. In Kastl's model of bidding in Czech treasury auctions, there is an explicit marginal cost of submitting a price-quantity point. Thus, "coarse" bids are constrained optimal, and can depart significantly from bids that can comprise of a larger number of points. Although the cost of adding bidpoints may explain a portion of the foregone ex post profits, it appears the majority of foregone profits results from not 'optimally' bidding. To see this, suppose that there was a cost to adding bidpoints that restricted firms to submitting the number bidpoints that we observe them using (rather than 40). For example, TXU uses an average of 12.6 bidpoints. We calculate Naive Best Response (NBR) profits using 12 equally spaced prices between $10 and 40$. Note that since the 12 price points are fixed rather than chosen optimally, we will understate NBR profits and thus overstate Percent Achieved. Even relative to this "constrained" benchmark, TXU's Percent Achieved is only 60%. Performing the same exercise for Calpine (using 7 points) yields a Percent Achieved of 45%. We also used a similar approach to get a lower bound on the implied cost of using an additional bid point. To do this, we calculated TXU's NBR profits using 12 vs. 13 bid points (we kept the 12 equally spaced points constant, and varied the 13th point to maximize TXU's profit). TXU's incremental profit gain from adding the 13th bid point was $1.59.

Column 2 suggests that a "corporate governance" based explanation is not borne in the data. Controlling for size and technology, the performance of municipal utilities appears, in fact, to be slightly (2.5 to 4.8%) better than that of investor-owned utilities, though the regression coefficient is not statistically significant at conventional levels. Moreover, merchant firms actually seem to be the weakest performers.

We also find that the technology mix of a firm does not appear to affect its performance on the balancing market. In column 3 we add measures of technology type and, find that owning a large fraction of coal and combined-cycle generation units does not negatively impact performance.

We can best measure scale effects if we focus on those firms that choose to establish their own bidding operation rather than those that outsource. In column 4, we control for whether the firm performs its own bidding and allow the effect of SIZE to vary by OWNBIDDER status. Larger stakes are associated with higher performance for firms that serve as their own bidders. For firms that perform their own bidding, a 1000 MWh increase in optimal sales volume is associated with an 86 percentage point increase in Percent Achieved.

Moreover, if one were to view the choice to serve as their own bidder as a revealed preference, then the threshold volume of sales where it becomes profitable to construct an in-house bidding operation rather than to outsource is 71 MWh (=-(-.071/.001)). The results are similar when we control for firm type and technology (column 5). A 1000 MWh increase in optimal sales volume is associated with an 97 percentage point increase in Percent Achieved and the threshold size for ownbidding is 163 MWh.

Explanations of deviations from static profit maximization, other than scale economies, do not have strong empirical support. Explanations based on risk-preferences are ruled out by the fact that the first-order condition of optimality (equation(2)) does not depend on the curvature of the utility function. Collusion appears to be unlikely because the heterogeneity in bid patterns is consistent with a collusive coalition that includes the small but not the large players, a possibility we believe to be unlikely. Engineering-based explanations such as unmeasured adjustment costs or transmission constraints also appear implausible due to the periods we choose to analyze; for a more detailed discussion, see our working paper (Hortacsu and Puller (2005)).

5.2 Learning

One might also expect bidders to gradually learn the rules of competition in this market, and improve their bidding behavior over time. To explore this hypothesis further, we examine whether there are any time trends in bidder performance.

Table 5 reports the results of a regression of individual bidder's daily performance (measured by the percentage of ex-post profit achieved for that day) on a time trend and controls for seasonality and the bidder's ex-post optimal generation in that period. The specification includes bidder fixed-effects to account for firm-level sources of variation. Notice that, for the entire sample of firms, the estimated coefficient for the time trend is positive, and indicates a 3 percentage point improvement in performance for every 100 days (or roughly 10% over a year). It is interesting to note that this time trend is not significant (though of the same estimated magnitude) for the top six bidders. For the rest of the firms, the time trend in performance is strongly significant.

Thus, our data support a learning hypothesis; though the rate of learning in this market strikes us as being rather slow (especially compared to rates of learning reported in laboratory experiments).[32] Firms in this market face a considerable amount of uncertainty, which may slow bidders' ability to infer optimal decision rules from their experiences. Moreover, several bidders have told us that they do not have the resources to perform detailed "ex-post" analyses that will enable them to assess how successful their bidding has been, and that they view our exercise as providing useful information in this regard.

6 Quantifying Efficiency Losses

Inefficiencies arise in electricity markets when the bids do not lead to least-cost production.[33] In a balancing market, firms bidding above marginal cost on the INC side may not be called upon to produce despite the fact they have low cost generators. Similarly, firms bidding below marginal cost on the DEC side may not be called to reduce production even if they have high cost plants operating.

We measure the cost of production in the balancing market implied by actual bids and compare those costs to the production costs under various counterfactual bidding behaviors. Our benchmark for efficiency is competitive bidding, which is an equilibrium under an alternative market design – the multi-unit Vickrey auction. We calculate the production costs if firms instead were to bid their marginal cost functions. These calculations suffer from one data limitation – the generation data are incomplete for a few of the smaller firms. This does not affect our ability to analyze bidding behavior for the remaining firms but prevents us from calculating efficient dispatch for many days in our sample. Therefore, these measures of efficiency loss should be interpreted with caution for they only represent a fraction of our sample period.

Results are shown in Table 6. The average hourly dispatch costs in the balancing market are $29,874 under actual bidding and $23,571 under marginal cost bidding, implying that actual production costs are 27% higher than least cost production.

We can decompose the welfare losses into the two major sources of inefficiency. The first source of inefficiency is the strategic exercise of market power – large firms face steeper residual demand functions and thus have incentives to bid steeper than marginal cost. This may result in inefficient production when low cost units are withheld from the market while higher cost units are called to generate.[34]

A second source of inefficiency, at least potentially, is the behavior of small generators. As we have noted, many of the smaller participants submit extremely steep bid functions – even though static profit maximization suggests that they should bid very close to marginal cost. As a result, they are often not called to produce despite the fact that it would be efficient to do so. This source of inefficient production can arise from a variety of sources as we discuss above (e.g. fixed cost of establishing a sophisticated trading

[32] It is difficult to make a direct comparison between "rounds" in the laboratory and "days" in electricity markets. However, if we are willing to equate "rounds" with "days," bidding behavior appears to converge (in percentage profit terms) to theoretical predictions quicker in many laboratory experiments. See Kagel (1995), Section F for examples from auction experiments.

[33] Because total demand is perfectly inelastic, prices higher or lower than competitive levels do not cause suboptimal levels of consumption. All inefficiencies are productive rather than allocative. For a more general discussion of the efficiency properties of uniform price multi-unit auctions, see Ausubel and Cramton (2002).

[34] Borenstein, Bushnell and Wolak (2002) find the actual production costs to be 14% higher than competitive levels in the California market in 2000 due to market power. Mansur (forthcoming) refines the methodology from Borenstein, Bushnell and Wolak to correct for production constraints such as start-up costs, and finds that actual production costs in the Pennsylvania, New Jersey, Maryland market exceed competitive costs by 3-8%. Note that our analysis only calculates productive inefficiencies in the *balancing* market for units that have already started and submitted day-ahead schedules.

operation). To the extent that the inefficiencies result from scale economies, the fixed cost of setting up trading operations or outsourcing to power marketers should be included in a complete welfare analysis.

To decompose the total amount of production inefficiency into these two components, we separate the firms into two groups: "strategic" bidders who exercise market power optimally, and"non-strategic" bidders who bid exceedingly steep schedules that effectively minimizes their participation in the market. We first compute counterfactual generation costs in this market under the assumption that "strategic" firms ignore their market power and bid competitively, i.e. we calculate what the total generation cost would have been if the "strategic" firms were bidding their marginal costs. In this counterfactual, we assume that the "non-strategic" firms continue to bid what they are observed to bid, i.e. that they do not respond to the (counterfactual) change in the behavior of their strategic competitors.[35]

We assume that the three largest firms – Reliant, TXU, and Calpine – as well as the other 3 firms in the top 6 in Table 1 (Brownsville, Bryan and Tenaska) are "strategic." We find counterfactual dispatch costs when these six firms bid marginal cost and all other firms submit their actual bid schedules. The difference in dispatch costs between this counterfactual bidding strategy and the actual bids can be interpreted as the "efficiency loss due to market power." The remaining efficiency loss is due to non-strategic firms that bid so as to not participate in the market.

Table 6 decomposes the efficiency losses. Again, the total efficiency loss due to misrepresentation of marginal costs is on average \$6,303 per hour, or 27% of the total cost of efficient generation in the balancing market. If the "strategic" firms were to bid their marginal costs, the total efficiency loss would have been only \$1203 less than the actual efficiency loss. This means that most (81%) of the observed efficiency loss is due to the steep bid schedules submitted by the "non-strategic" bidders.

We should note, once again, that the above calculation relies on relatively few intervals (62 out of a total of 220 periods – the latter periods are especially prone to the missing data problem). This calculation is largely based upon the first six months of the market's operation. However, this calculation points out that the observed deviations from static profit maximization are not without economic consequence. In fact, they cause larger efficiency losses than the "near-optimal" exercise of market power by the "strategic" firms. Submitting bid functions that are too steep not only sacrifices producer surplus, but also prevents technologically efficient firms from supplying energy to the balancing market.

One would expect that these efficiency losses due to the "non-strategic" firms' behavior would dissipate over time. Unfortunately, as we noted above, we face data limitations that prevent us from conducting a long time series analysis, especially in the later part of the sample. However, market forces are at play to reduce the inefficiencies. To the extent there are scale effects, the small firms may outsource bidding decisions or consolidate bidding activities across plants. It is noteworthy that there are a variety of power marketers and large energy trading firms that offer energy asset management services to generators in ERCOT. Such outsourcing can increase participation by "non-strategic" firms and reduce inefficiencies without substantially increasing the fixed cost of bidding expertise.

[35]Note that the "reverse" of this counterfactual, where we set "non-strategic" bids equal to marginal cost, and do not allow "strategic" bidders to respond, would be less realistic, since our results show that the "strategic" bidders do respond to changes in the residual demand they are facing. Counterfactuals involving the equilibrium response of "strategic" bidders are complicated by the fact that multi-unit auctions can have multiple equilibria.

7 Conclusion

Our analysis of the ERCOT balancing market yields the following conclusions. The first conclusion, we believe, is a comforting one for economic theory. In a marketplace that is marked by considerable uncertainty and institutional complexity, two factors that may pose analytical challenges for both the firms competing in the market and the economists who are observing (and, in some cases, advising) them, firms with large stakes in the market behave close to theoretical predictions of a strategic model of oligopolistic interaction. Indeed, as two empirical industrial organization researchers who have previously utilized such models to infer supply and demand parameters in other markets, we are swayed to interpret the behavioral pattern displayed by Reliant as good news for previous and future empirical work on oligopolistic markets. More specifically, the confirmation of the basic predictions of the uniform price share auction model is important, as one could use this model to forecast bidder behavior in restructured electricity markets that are being put into operation in many different parts of the world.

However, our results also suggest some amount of caution when analyzing and predicting the behavior of smaller players in newly restructured markets. Smaller firms submit bids that differ substantially from the benchmarks we construct for optimal bidding. This finding is not inconsistent with rational economic behavior by these bidders, however. As argued in section 5.1, "participation" in the balancing market may have nontrivial costs, and the behavioral pattern *across* firms appears to confirm this hypothesis.

Our third result is that small firms' deviations from optimal bidding is economically important. In section 6, we calculated that small firms' bidding patterns led to the major portion of losses in productive efficiency. This suggests interesting new avenues for market design that explicitly takes into account the strategic complexity, hence the participation costs imposed by proposed market mechanisms. Such a consideration may favor dominant strategy implementable mechanisms, such as the Vickrey-Clarke-Groves (VCG) mechanism, over others. However, as pointed out by Milgrom (2004) (chapter 2.5), the VCG mechanism suffers from serious pitfalls of its own. Nevertheless, we view theoretical research in this area to prove extremely fruitful for real-world applications.

APPENDIX A: DATA

Hourly bid schedules by each bidder, or qualified scheduling entity (QSE), are from ERCOT. QSEs occasionally bid for more than one firm. For example, in the South zone in 2001, the QSE named Reliant bid for both Reliant and the City of San Antonio. We match the bid functions to all units for which the QSE bids. So for all units owned by both Reliant and the City of San Antonio in the South in 2001, we match the bid function to the generation data. However, interpretation of the results becomes problematic when an observed bid function represents the bids by more than one firm. Because the results are some combination of two firms' behavior, we will not interpret results in such situations. We only interpret our results as firm-level behavior when at least 90% of all electricity generated by owners using that QSE can be attributed to a single owner. We make one exception to this 90% rule – TXU Generation which comprises 87% of the generation for TXU the QSE in North 2002.

We measure the variable costs of output using data on each unit's fuel costs and the rate at which the unit converts the fuel to electricity. For each 15-minute interval, we have data from ERCOT on whether a generating unit is operating, its day-ahead scheduled generation, and its hourly available generating capacity. We measure the marginal cost of units that burn natural gas and coal. For each unit, we have data on the fuel efficiency (i.e. average heat rate). Each unit is assumed to have constant marginal cost up to its hourly operating capacity, an assumption that is common in the literature. The ERCOT system is largely separated from other electricity grids in the country so there are virtually no imports.

Daily gas spot prices measure the opportunity cost of fuel for natural gas units. We use prices at the Agua Dulce, Katy, Waha, and Carthage hubs for units in the South, Houston, West, and North zones, respectively. We assume a gas distribution charge of \$0.10/mmBtu. Coal prices are monthly weighted average spot price of purchases of bituminous, sub-Bituminous, and lignite in Texas, reported in Form FERC-423. Coal-fired plants in Texas are required to possess federal emission permits for each ton of SO2 emissions. In order to measure average emission rates, we merge hourly net metered generation data from ERCOT with hourly emission data from EPA's CEMS to calculate each unit's average pounds of SO2 emissions per net MWh of electricity output. The emissions each hour are priced at the monthly average EPA permit price reported on the EPA website.

In order to deal with complications posed by transmission congestion, we restrict our sample to daily intervals 6:00-6:15pm during which there is no interzonal transmission congestion during the 6-7pm bidding hour. We believe *intra*zonal (or local) congestion is also likely to be rare during these intervals.

We do not account for the sales of operating reserves. We do not incorporate the possibility that some of the available capacity to INC in our data may be sold as reserves. However, the amount of operating reserves procured are small as a fraction of total demand.

We measure the marginal cost of INCing or DECing from the day-ahead schedule of output. We account for the fact that units cannot DEC down to zero output without incurring costs of startup and facing constraints on minimum downtime. It is unlikely that revenue from the balancing market would be sufficiently lucrative to compensate a unit for shutting down. Therefore, we assume that each operating unit cannot DEC to a level below 20% of its maximum generating capacity.

APPENDIX B: PROOF OF PROPOSITION 2

Given the additively separable form of the bidding strategies $S_i(p, QC_i) = \alpha_i(p) + \beta_i(QC_i)$, use the market clearing condition (1) above to represent the event $\{p_t^c \leq p | QC_i, S_i\}$, i.e. there is excess supply at p, conditional on firm i bidding S_i at this price:

$$\sum_{j \neq i} \beta_j(QC_j) - \varepsilon \geq D(p) - S_i - \sum_{j \neq i} \alpha_j(p)$$

The left hand side of this inequality can be labeled as a (bidder-specific) random variable, θ_i, that does not depend on p, and the right hand side is a deterministic function of price. Let $\Gamma_i(.)$ denote the cdf of θ_i and $\gamma_i(.)$ denote the pdf (both conditional on the bidder's contract quantity, QC_i). Given these:

$$\begin{aligned}
H_p(p, S_i; QC_i) &= \frac{\partial}{\partial p} \Pr(p_t^c \leq p | QC_i, S_i) \\
&= \frac{\partial}{\partial p} \Pr(\theta_i \geq D(p) - S_i - \sum_{j \neq i} \alpha_j(p)) \\
&= \frac{\partial}{\partial p} [1 - \Gamma_i(D(p) - S_i - \sum_{j \neq i} \alpha_j(p))] \\
&= -\gamma_i(D(p) - S_i - \sum_{j \neq i} \alpha_j(p)) \frac{\partial}{\partial p}(D(p) - S_i - \sum_{j \neq i} \alpha_j(p))
\end{aligned}$$

and

$$\begin{aligned}
H_S(p, S_i; QC_i) &= \frac{\partial}{\partial S_i} \Pr(p_t^c \leq p | QC_i, S_i) \\
&= -\gamma_i(D(p) - S_i - \sum_{j \neq i} \alpha_j(p)) \frac{\partial}{\partial S_i}(D(p) - S_i - \sum_{j \neq i} \alpha_j(p))
\end{aligned}$$

Evaluating the derivatives gives $\frac{H_p(p,S_i;QC_i)}{H_S(p,S_i;QC_i)} = -[D'(p) - \sum_{j \neq i} \alpha_j^{'}(p)]$. Now, observe that with the above restrictions, the residual demand function faced by firm i (for a given realization of the random variables $\{\varepsilon, QC_{-i}, i = 1, ..N\}$), is given by:

$$RD_i(p, \varepsilon, QC_{-i}) = D(p) + \varepsilon - \sum_{j \neq i} \alpha_j(p) - \sum_{j \neq i} \beta_j(QC_j) \tag{4}$$

with derivative:

$$RD_i^{'}(p) = D'(p) - \sum_{j \neq i} \alpha_j'(p)$$

which yields the result in the proposition.

Note that in one other case where we can collapse the multiple stochastic terms into a scalar random variable, we do not obtain the ex-post optimality result. This is the case when $S_i(p, QC_i) = \alpha_i(p) + \beta_i p QC_i$ and $\tilde{D}(p) = D(p) + p\varepsilon$, i.e. private information and aggregate uncertainty leads to pure *rotations* of the residual demand curve. In this case market clearing condition becomes $\theta_i = \sum_{j \neq i} \beta_j QC_j - \varepsilon \geq \frac{1}{p}\left[D(p) - S_i - \sum_{j \neq i} \alpha_j(p)\right]$, but $\frac{H_p(p,S_i)}{H_S(p,S_i)} \neq \frac{1}{RD_i^{'}(p,\varepsilon,QC_{-i})}$.

REFERENCES

S. C. Athey and P. A. Haile. Nonparametric Approaches to Auctions, *Handbook of Econometrics*, Vol 6., forthcoming.

L.M. Ausubel and P. Cramton. Demand Reduction and Inefficiency in Multi-Unit Auctions, working paper, University of Maryland , 2002.

P. Bajari, A. Hortaçsu. Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data, NBER working paper no. 9889, 2003.

R. Baldick, R. Grant, and E. Kahn. Theory and Application of Linear Supply Function Equilibrium in Electricity Markets, *Journal of Regulatory Economics*, 2004.

R. Baldick and H. Niu. Lessons Learned: The Texas Experience. *Electricity Deregulation: Choices and Challenges.* J. Griffin and S. Puller (eds.) University of Chicago Press, 2005.

S. Borenstein, J.B. Bushnell, and F.A. Wolak. Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market. *American Economic Review*, 92(5), 1376-1405, 2002.

T. Bresnahan. Empirical Studies of Industries with Market Power. *Handbook of Industrial Organization* v. 1. edited by R. Schmalensee and R. Willig. New York, NY: North Holland, 1989.

J.B. Bushnell and C. Saravia. An Empirical Assessment of the Competitiveness of the New England Electricity Market. Center for the Study of Energy Markets working paper CSEMWP-101, 2002.

J.B. Bushnell, E.T. Mansur, and C. Saravia. Vertical Arrangements, Market Structure and Competition: An Analysis of Restructured U.S. Electricity Markets, mimeo, Yale, 2006.

P. Cramton. Competitive Bidding Behavior in Uniform-Price Auction Markets, Report to Federal Energy Regulatory Commission, Dockets EL00-95-075 and EL00-98-063, 2004.

G. Crawford, J. Crespo, and H. Tauchen, Bidding Asymmetries in Multi-Unit Auctions: Implications of Bid Function Equilibria in the British Spot Market for Electricity, *International Journal of Industrial Organization*, forthcoming.

H. Demsetz and K. Lehn, The Structure of Corporate Ownership: Causes and Consequences. *Journal of Political Economy*, 93, 1155-1177.

L. Einav, A Note on The Analogies between Empirical Models of Auctions and of Differentiated Product Markets, working paper, Stanford University, 2004.

D. Genesove, W. Mullin. Testing Static Oligopoly Models: Conduct and Cost in the Sugar Industry, 1890-1914. *RAND Journal of Economics*, 29(2), 1998.

R. Green. Increasing Competition in the British Electricity Spot Market. *Journal of Industrial Economics*, 1992.

R. Green and D. Newbery. Competition in the British Electricity Spot Market. *Journal of Political Economy*, 100(5), 929-953, 1990.

L.P. Hansen. Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica*, 50, 1029-54, 1982.

K. Hendricks and R. H. Porter. An Empirical Perspective on Auctions. *Handbook of Industrial Organization*, Vol 3., forthcoming.

K. Hendricks, R. H. Porter and B. Boudreau. Information, Returns, and Bidding Behavior in OCS Auctions: 1954-1969. *Journal of Industrial Economics*, 35, 517-42, 1987.

A. Hortaçsu. Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market, working paper, University of Chicago, 2002a.

A. Hortaçsu. Bidding Behavior in Divisible Good Auctions: Theory and Evidence from the Turkish Treasury Auction Market, working paper, University of Chicago, 2002b.

A. Hortaçsu and S.L. Puller. Understanding Strategic Bidding in Restructured Electricity Markets: A Case Study of ERCOT, NBER Working Paper 11123, 2005.

P. Joskow. Deregulation and Regulatory Reform in the U.S. Electric Power Sector. In S. Peltzman and C. Winston, editors, Deregulation of Network Industries: The Next Steps. Brookings Press, 2000.

P. Joskow and E. Kahn. A Quantitative Analysis of Pricing Behavior In California's Wholesale Electricity Market During Summer 2000. *The Energy Journal*, 23(4), 1-35, 2002.

J. Kagel. Auctions: A Survey of Experimental Research. *Handbook of Experimental Economics* J. Kagel and A. Roth (eds.) Princeton University Press, 1995.

J. Kastl. Discrete Bids and Empirical Inference in Divisible Good Auctions, working paper, Stanford University, 2006a.

J. Kastl. On the Existence and Characterization of Equilibria in Private Value Divisible Good Auctions, working paper, Stanford University, 2006b.

P. Klemperer. Why Every Economist Should Learn Some Auction Theory. *Advances in Economics and Econometrics: Invited Lectures to 8th World Congress of the Econometric Society.* M. Dewatripont, L. Hansen and S. Turnovsky (eds.) Cambridge University Press, 2003.

P. Klemperer and Meyer, M. Supply Function Equilibria in Oligopoly under Uncertainty, *Econometrica*, 57(6), 1243-1277, 1989.

K. Kuhn and M. Machado, Bilateral Market Power and Vertical Integration in the Spanish Electricity Spot Market, CEPR Working Paper No. 4950, 2004.

E.T. Mansur. Upstream Competition and Vertical Integration in Electricity Markets, *Journal of Law and Economics*, 50(1), 125-156, 2007.

E.T. Mansur. Measuring Welfare in Restructured Electricity Markets, *Review of Economics and Statistics*, forthcoming.

D. McAdams. Partial Identification in Multi-Unit Auctions, MIT Sloan working paper, 2006.

P. Milgrom. *Putting Auction Theory To Work: The Churchill Lectures in Economics.* Cambridge University Press, 2004.

S.L. Puller. Pricing and Firm Conduct in California's Deregulated Electricity Market, *Review of Economics and Statistics*, 89(1), 2007.

A. Rudkevich. Supply Function Equilibrium in Power Markets: Learning All the Way. TCA Technical Report Number 1299-1702, Tabors Caramanis and Associates, 1999.

A. Sweeting. Market Power in the England and Wales Wholesale Electricity Market: 1995-2000, working paper, Northwestern University, 2005.

S. Viswanathan, J. Wang, and T. Witelski, Optimal Bidding In Multi-Unit Discriminatory Auctions, working paper, Duke University, 2002.

N.H. von der Fehr and D. Harbord. Spot Market Competition in the UK Electricity Industry. *Economic Journal*, 103, 531-546, 1993.

J.J.D. Wang and J.F. Zender. Auctioning Divisible Goods. *Economic Theory*, 19(4):673-705, 2002.

R. Wilson. Auctions of Shares. *Quarterly Journal of Economics*, 93, 675-689, 1979.

R. Wilson. Architecture of Power Markets. *Econometrica*, 70(4), 1299-1340, 2002.

F.A. Wolak. An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market. *International Journal of Economics*. 14(2), 2000, 1-39.

F.A. Wolak. "Identification and Estimation of Cost Functions Using Observed Bid Data: An Application to Electricity Markets," in M. Dewatripont, L.P. Hansen, and S.J. Turnovsky, eds., Advances in Economic and Econometrics: Theory and Applications, Eight World Congress, Volume II. New York: Cambridge University Press, 2003a, pp. 133-169.

F.A. Wolak. Measuring Unilateral Market Power in Wholesale Electricity Markets: The California Market 1998 to 2000. *American Economic Review*, 93(2), 425-430, 2003b.

C.D. Wolfram. Measuring Duopoly Power in the British Electricity Spot Market. *American Economic Review*, 1999.

Table 1: Outcomes Under Actual, Ex-Post Optimal and Naive Best Response Bidding

Firm	Percent Achieved Relative To		Producer Surplus ($/hour)			Quantity Sales (MWh)		Upper Bound Total Prof
	XP Optimal (1)	Naive BR (2)	Actual (3)	Naive BR (4)	XP Optimal (5)	XP Optimal (6)	Actual (7)	(8)
Reliant	79%	80%	3,422	4,268	4,333	507	431	98%
Brownsville PUB	50%	50%	173	343	343	42	17	88%
City of Bryan	45%	45%	221	488	488	56	30	85%
Tenaska Gateway Partners	41%	41%	456	1,111	1,115	182	72	88%
TXU	39%	41%	1,243	3,056	3,159	441	133	97%
Calpine Corp	37%	38%	820	2,168	2,214	408	102	91%
Denton Municipal Electric	35%	35%	11	31	31	3	1	98%
Ingleside Cogeneration	31%	31%	171	541	541	81	25	79%
City of Austin	30%	31%	581	1,889	1,907	271	48	84%
Rio Nogales LP	28%	28%	109	393	393	60	10	93%
Lower Colorado River Auth	25%	25%	367	1,471	1,488	274	28	88%
City of San Antonio	23%	24%	290	1,221	1,241	266	49	90%
Gregory Power Partners	20%	20%	143	720	722	96	14	82%
Midlothian Energy	17%	17%	171	1,016	1,024	175	15	86%
Cogen Lyondell Inc	16%	16%	408	2,523	2,523	269	34	67%
Tractebel Power Inc	16%	16%	127	795	795	97	15	79%
Brazos Electric Power Coop	15%	15%	101	676	677	82	6	79%
Lamar Power Partners	15%	15%	266	1,800	1,808	272	30	79%
Mirant Wichita Falls	14%	14%	16	114	114	18	2	83%
BP Energy	14%	14%	134	993	994	135	17	80%
City of Garland	13%	13%	128	1,018	1,019	115	5	80%
Hays Energy	8%	8%	64	775	777	111	8	82%
West Texas Utilities	8%	8%	132	1,635	1,635	224	11	82%
Central Power & Light	8%	8%	185	2,375	2,407	352	35	80%
Guadalupe Power Partners	6%	6%	140	2,356	2,380	396	12	77%
Tenaska Frontier Partners	5%	5%	52	1,044	1,051	144	7	80%
South Texas Electric Coop	3%	3%	8	298	298	44	1	81%
Sweeney Cogeneration	2%	2%	10	409	409	46	2	85%
Brazos Valley Energy LP	0%	0%	1	134	134	12	0	68%
AES Deepwater	0%	0%	1	969	969	92	0	60%
Frontera General LP	0%	0%	0	984	1,004	197	0	62%
TGC	0%	0%	0	405	405	70	0	81%
South Houston Green Power	0%	0%	0	60	60	7	0	70%
Air Liquide America	-8%	-8%	-181	2,174	2,176	103	9	55%
Extex Laporte LP	-81%	-81%	-1,230	1,497	1,497	92	116	41%

Producer surplus measures are all relative to profitability under not bidding into the balancing market and are measured in average dollars for each 6-7pm auction in our sample.

Percent Achieved is the ratio of producer surplus under actual bidding to producer surplus under the benchmark for optimal bidding (either ex-post optimal or Naive Best Response).

Actual is the actual bids submitted by the firms.

XP Optimal is ex-post optimal bidding constructed as we describe in section 3.

Naive BR is the naive best-response bid to rivals' lagged bids as we describe in section **??**.

Quantity of sales are total volumes and, therefore, may sum to more than the average net demand for balancing energy.

In column 8, the Upper Bound for Total Profitability = (Percent Achieved XP Optimal∗%Sales in Balancing)+(100% ∗ %Sales in Bilateral) as defined in section 4.2.

Table 2: Testing Additive Separability: Panel Regressions

	Bid Fn Slope	Bid Fn Slope
Contract Quantity	0.0010	0.00049
	(0.0003)*	(0.00031)
Slope of Residual Demand	-0.002	0.38
	(0.004)	(0.01)**
Slope of Marginal Cost	0.129**	0.075**
	(0.017)	(0.011)
Time Trend	0.000044	
	(0.000033)	
Bidder FE	Y	Y
Auction FE	N	Y
Nobs	2254	2254
R-squared	0.70	0.79

Note: Dependent variable is the slope of each bidder-auction's bid function using a linear fit. Sample includes all bidders in 2001. The slope and intercept of the bid function, residual demand function, and marginal cost are based on a linear fit of the bid points between $0 and $30.

* significant at 5%

** significant at 1%

Table 3: Testing Additive Separability: Individual Bidders

	% Achieved	Bid Fn Slope
Reliant	79%	-0.00081
		(0.0029)
City of Bryan	45%	-0.0025
		(0.0014)
TXU (first month)		0.021
		(0.013)
TXU (after first month)	39%	0.0025
		(0.0013)
Calpine	37%	0.0044
		(0.0018)*
City of Austin	30%	0.0055
		(0.0015)**
LCRA	25%	0.0136
		(0.0014)**
City of Garland	13%	0.0011
		(0.0007)
South TX	3%	0.0014
		(0.0006)*

The reported coefficients are the coefficients of each firm's contract quantity. Unreported control variables include the slope of the marginal cost and residual demand functions, and a time trend.

* significant at 5%

** significant at 1%

Table 4: Relationship Between Profitability and Firm Characteristics

	(1)	(2)	(3)	(4)	(5)
Size (MWh)	0.00052	0.00047	0.00035	-0.00015	-0.00035
	(0.00031)*	(0.00026)*	(0.00028)	(0.00025)	(0.00024)
Size*OwnBidder				0.00101	0.00132
				(0.00044)**	(0.00042)***
OwnBidder				-0.071	-0.215
				(0.111)	(0.109)*
Merchant Firm		-0.098	-0.167		-0.170
		(0.135)	(0.146)		(0.098)*
Municipal		0.048	0.025		0.047
		(0.135)	(0.142)		(0.089)
Coal			-0.021		0.043
			(0.113)		(0.120)
Combined-Cycle			0.082		0.127
			(0.058)		(0.071)*
Constant	0.100	0.158	0.200	0.158	0.262
	(0.054)*	(0.142)	(0.152)	(0.058)**	(0.098)**
Observations	34	34	34	34	34
R-squared	0.15	0.28	0.31	0.34	0.50

Notes: Dependent variable is the Percent Achieved for firm i over the entire sample as defined in section 4. Parameters are estimated by OLS. White standard errors are in parentheses. The excluded category for firm type is Investor-Owned utililty.
* significant at 10%
** significant at 5%
*** significant at 1%

Table 5: Evolution of Performance and Learning

	All Firms	Top 6 Firms	Non Top 6 Firms
Days Since Market Began	0.00031	0.00033	0.00028
	(0.00010)**	(0.00029)	(0.00009)**
Volume Optimal Output (GWh)	0.16	0.25	0.09
	(0.06)**	(0.10)*	(0.06)
Off-Peak Season	0.01361	-0.04437	0.03164
	(0.01892)	(0.03799)	(0.01885)
Constant	0.02330	0.21053	-0.01441
	(0.03030)	(0.07503)**	(0.02886)
Observations	9765	2103	7662
R-squared	0.25	0.18	0.11

Notes: Model includes firm fixed effects. The dependent variable is the *PercentAchieved* for firm i on day t. Parameters are estimated by OLS. White standard errors in parentheses.
* significant at 5%
** significant at 1%

Table 6: Decomposition of Efficiency Losses from Observed Bidding Behavior

Bidding Counterfactual	**Average Production Cost**
Actual Bids for all firms	\$29,874
"Strategic" firms Bid MC, Others Bid Actual	\$28,671
All Firms Bid MC (Vickrey auction)	\$23,571
Total Efficiency Losses	\$6,303
"Strategic Bidders"	\$1,203
"Non-Strategic Bidders"	\$5,100

"Strategic" firms are Reliant, TXU, Calpine, Brownsville, Bryan and Tenaska Gateway Partners. The average production cost is the hourly average of the absolute value of the cost of dispatch (either costs in INC hours or cost savings in DEC hours) during uncongested weekdays for the 6-7pm bidding hour.

Figure 1: Empirical Strategy

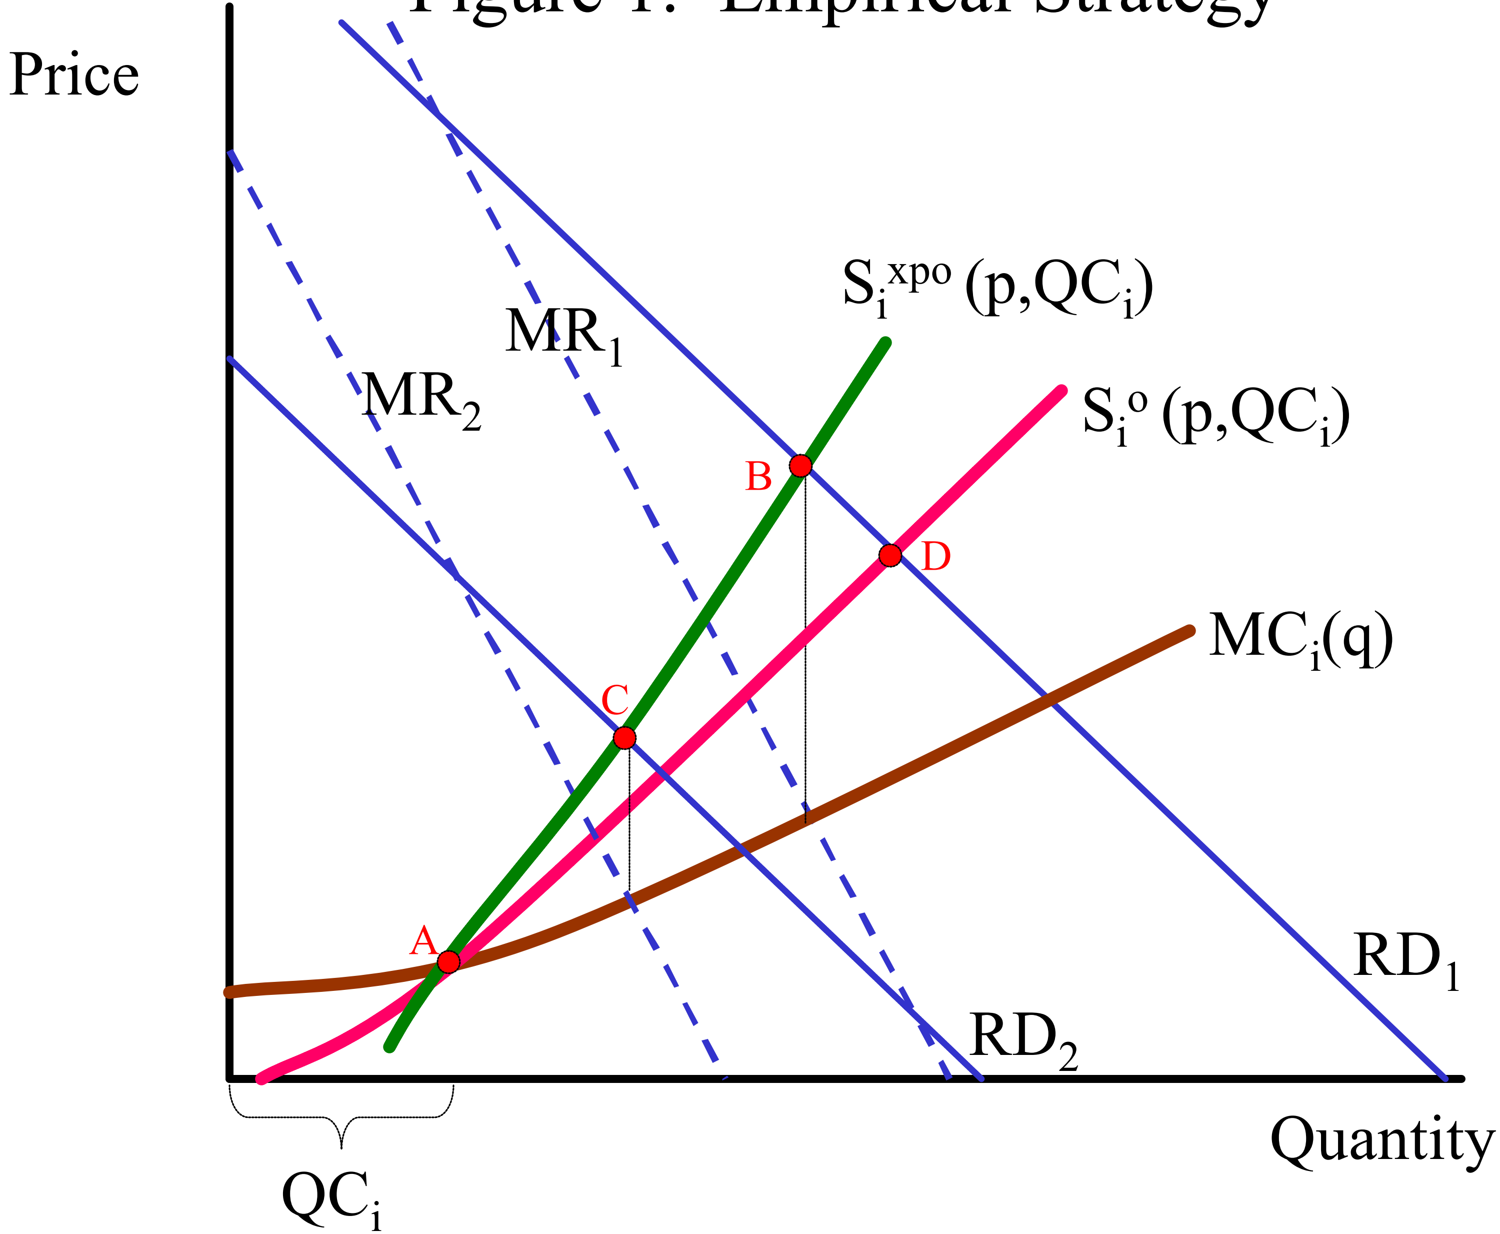

Figure 2: Examples of Actual and Optimal Bid Functions

Reliant

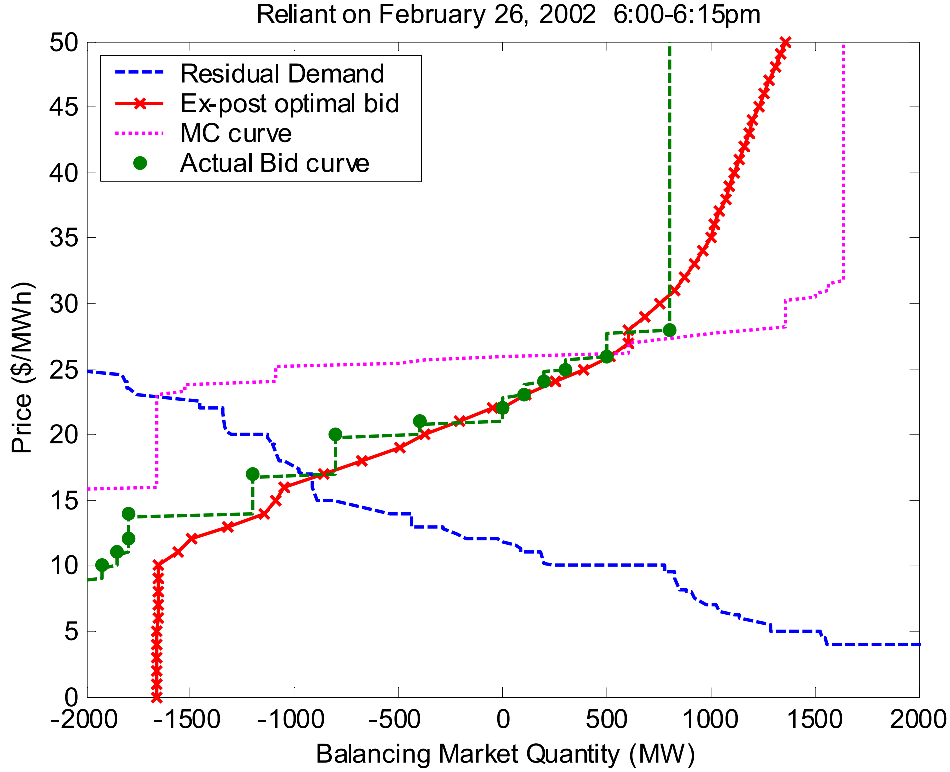

TXU

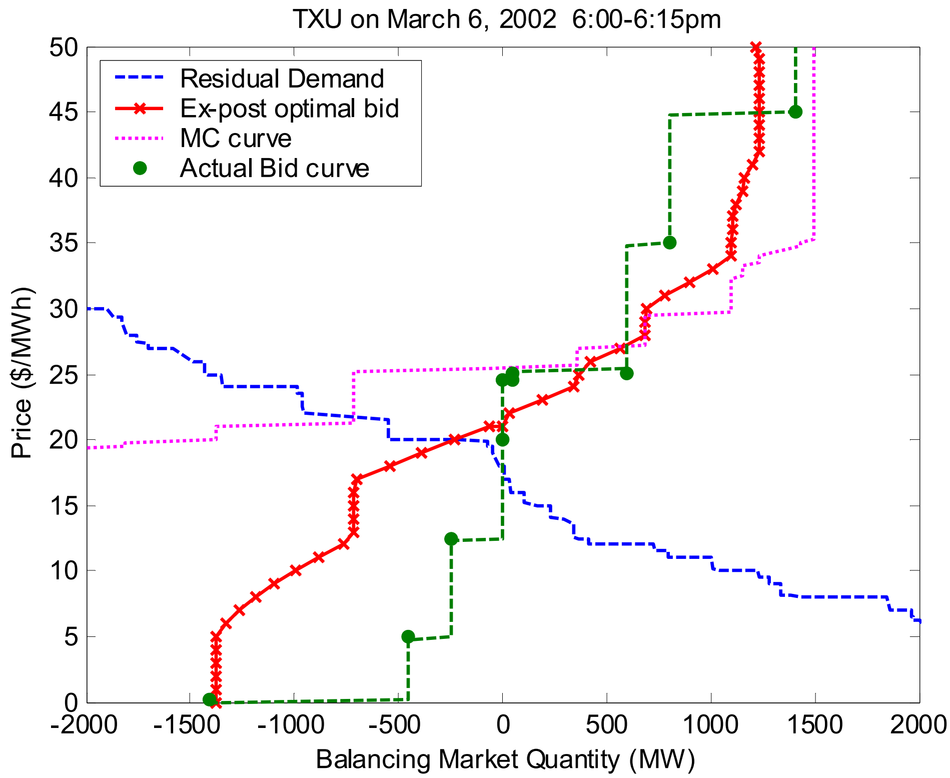

Figure 2 (contd)

Calpine

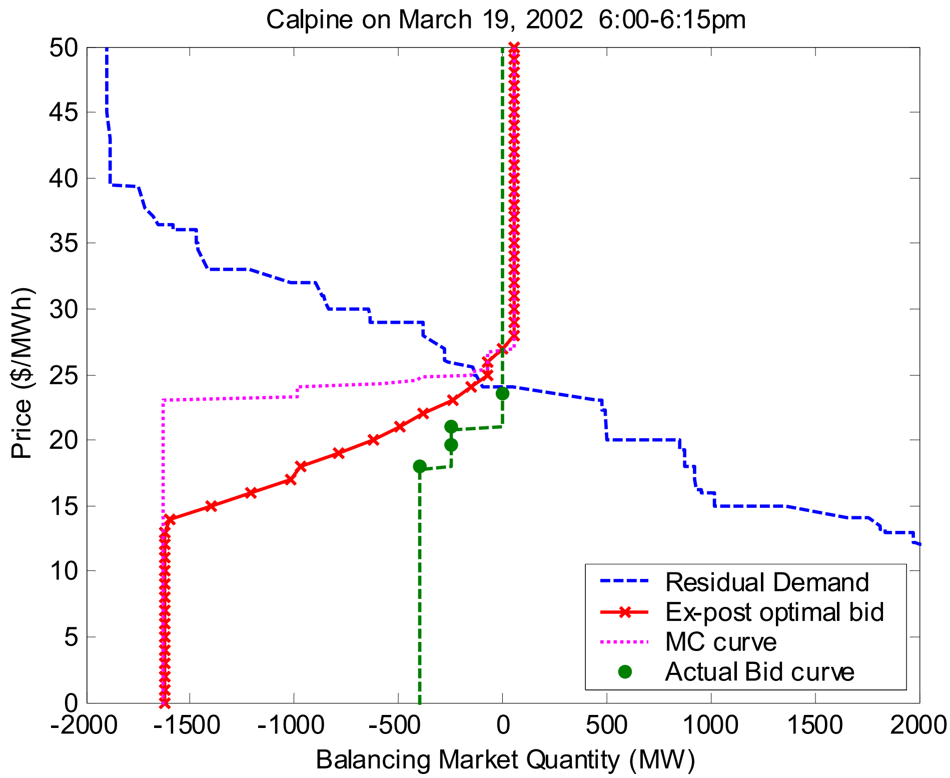

Guadalupe Power Partners

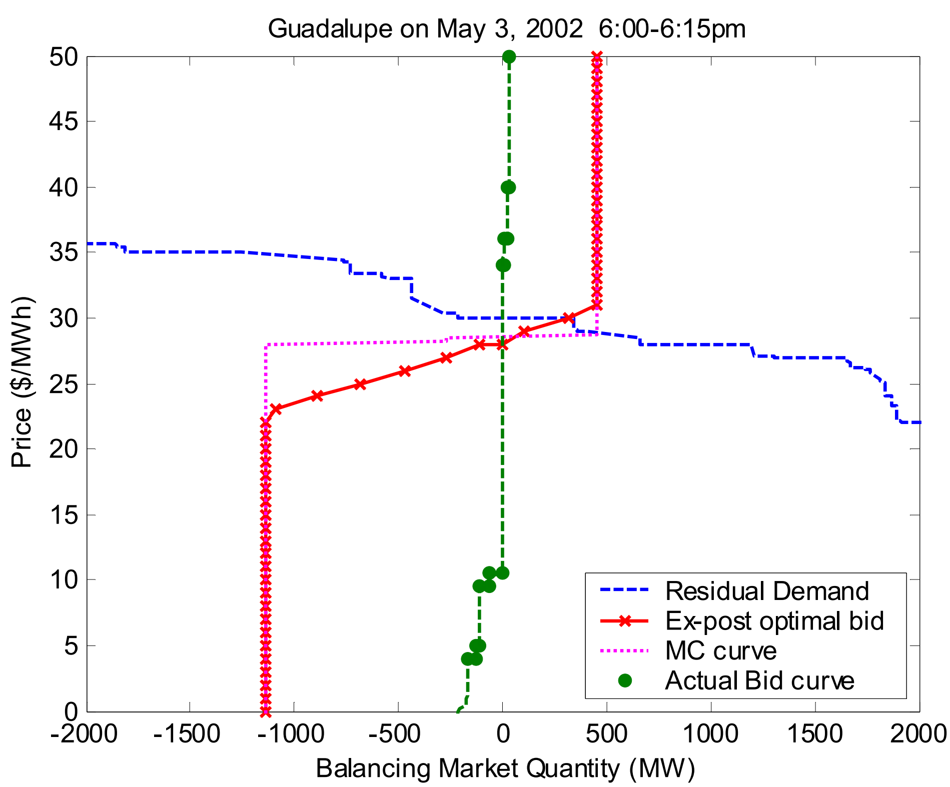

Figure 3

Performance vs. Stakes in Balancing Market
For Firms Serving as Their Own Bidding Entities

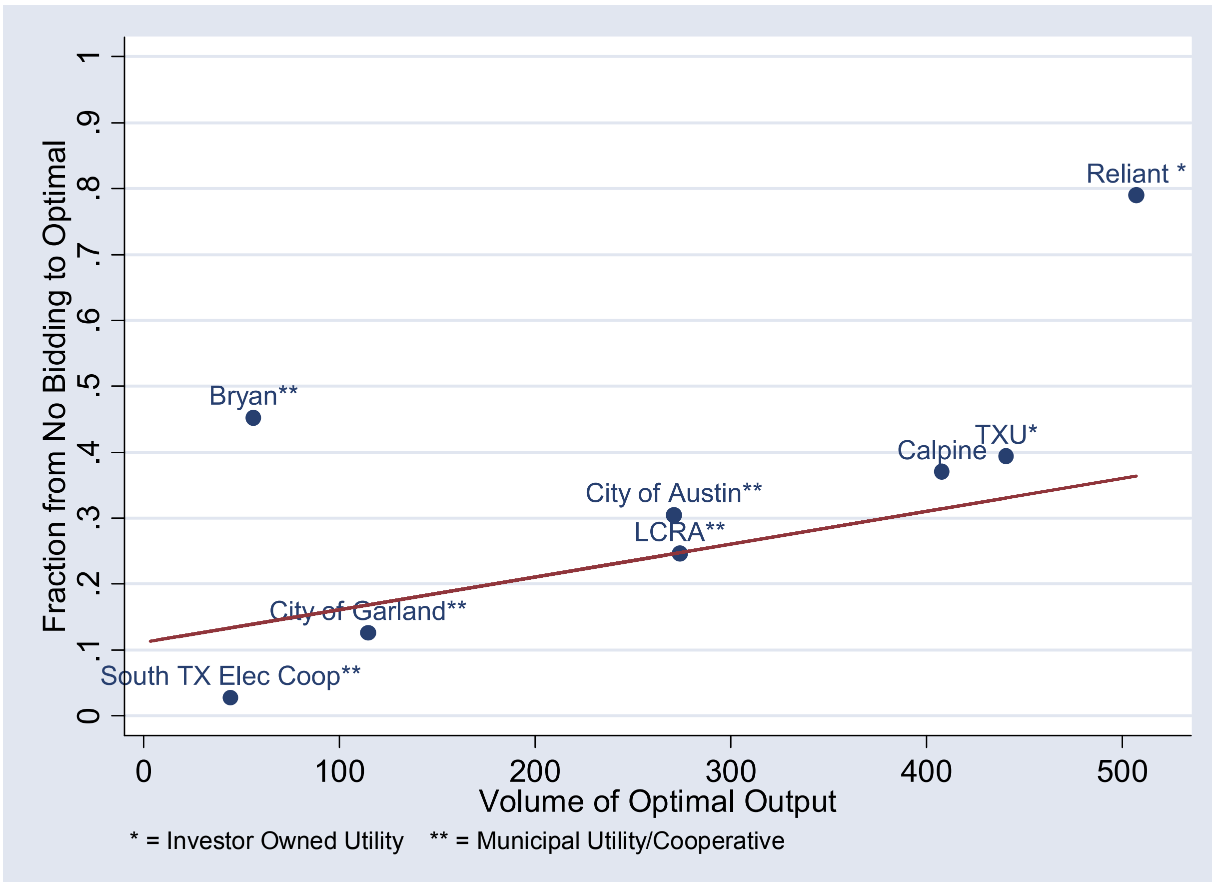

Percent Achieved is the percent of possible producer surplus achieved relative to not bidding, as defined in section 4.2 and reported in column 1 of Table 1. Volume of optimal output is the mean of the absolute value of quantity sales under ex-post optimal bidding, as reported in column 6 of Table 1. We restrict this analysis to firms that serve as their own bidding entities because we can best measure "stakes" of a bidding entity by focusing on firms that serve as their own bidders.