

How good are supply function equilibrium models: an empirical analysis of the ERCOT balancing market

Ramteen Sioshansi · Shmuel Oren

Published online: 2 December 2006
© Springer Science+Business Media, LLC 2006

Abstract We present an empirical analysis of a supply function equilibrium model in the Texas spot electricity market. We derive conditions for optimal bidding behavior in a spot market with ex ante bilaterally contracted sales. By estimating costs, we are able to derive a set of ex post-and ex ante-optimal supply functions and use a non-parametric behavioral model to compare our theoretically optimal supply functions to actual offers made. Our results show that with the exception of the largest generators, firms make offers with mark-ups and markdowns far in excess of what a model of profit-maximizing behavior suggests.

Keywords Electricity markets · Supply function equilibrium · ERCOT · Market power

1 Introduction

With the recent move towards liberalized electricity markets, there has been a need for economic theory to predict behavior in and performance of restructured markets. Industrial organization gives a wide variety of equilibrium models based on varying behavioral assumptions, which have been used to this end. Proponents claim this type of modeling can provide market participants a means of evaluating their past bidding behavior, market monitors a means of

R. Sioshansi (✉) · S. Oren
Department of Industrial Engineering and Operations Research,
University of California at Berkeley, Berkeley, CA 94720 USA
e-mail: ramteen@ieor.berkeley.edu

S. Oren
e-mail: oren@ieor.berkeley.edu

examining the efficiency of their markets, and regulators a means of comparing different market designs. Klemperer and Meyer's (1989) supply function equilibrium (SFE) model is often touted as a good model of spot electricity markets because it encapsulates the underlying structure of the market well. This can be seen in there being many applications of the model to predict market performance, Green and Newbery (1992), Newbery (1998), and Green (1996) being some of the seminal studies in this area.

In spite of the myriad applications of the SFE model, there has been limited empirical analysis showing the soundness of the model in characterizing actual firm behavior. For example, Wolfram (1999) and Kim and Knittel (2004) attempt to provide this analysis for conjectural variations type models, showing them to generally be uninformative. In this paper we use historical offer data from years 2002 to 2003 in the spot Balancing Electricity Service (BES) market administered by the Electricity Reliability Council of Texas (ERCOT) to test the behavioral predictions of an SFE model. Our work is related to that of Hortaçsu and Puller (2005) and Niu (2005). Hortaçsu and Puller conduct an empirical analysis of the BES market by applying Wilson's (1979) share-auction model and assuming each firm's contract price and position is private information. Their share-auction model is more robust than a standard SFE due to its allowing firms to have this private information. To make their model analyzable, however, they make an assumption that each firm's optimal supply function will be *additively separable and linear in the private information* (AS-LPI). This assumption essentially amounts to rivals' private contract information entering a firm's profit function as a horizontal shift in its residual demand curve, and their model becomes a standard SFE, yielding the same inverse-elasticity markup rule. In their analysis they calculate each firm's ex post-optimal supply function (EOSF), which is the firm's optimal response to the actual offers of its rivals, and test for consistency of the model by comparing actual to potentially achievable profits. Their results show the large incumbent utilities to perform moderately well while most of the smaller power generating companies (PGCs) submit supply functions which are too 'steep.' They then show the efficiency losses from this observed behavior—both from the large PGCs exercising their market power and the small PGCs withholding their generation from the market. Niu analyzes a linear SFE model by comparing the actual market-clearing price for energy (MCPE) to that which would result from her theoretically profit-maximizing benchmark. Her results show actual prices match her theoretical predictions well when the BES clears for incremental energy, but there is a large gap when it clears for decremental service. These two studies of the BES focus their analyses primarily on market outcomes. We take a different approach, which is to compare the entire range of the actual and optimal offer curves. Moreover, we conduct our analysis by comparing both EOSFs as well as a set of Nash equilibrium supply functions for what we term the strategic bidders in the market. Our results turn out to be similar to the aforementioned studies, showing the two large PGCs (TXU and Reliant Energy) to submit offer curves which are somewhat consistent with our derived EOSFs. The next largest PGC, Calpine, offers with markups which are far in

excess of what our ex post-optimal analysis predicts. By comparing Calpine's offers across different time periods, we provide evidence suggesting that these high markups are largely due to a period of learning during the first quarter of 2002. As for the other PGCs, we find that they submit offer curves which are too 'steep,' with markups and markdowns far in excess of our model's predictions. We provide evidence suggesting that these excessively high offers are meant to economically withhold their generation from the balancing market. Finally, by restricting our analysis to the three large PGCs and calculating Nash equilibrium supply functions for these firms alone, we show that the SFE model does a relatively good job of characterizing the bidding behavior of what we term 'strategic bidders' for incremental balancing energy offers.

The remainder of this paper proceeds as follows: Sect. 2 describes the ERCOT markets and specifically the BES spot market. Sect. 3 discusses our supply function model of the BES market, our assumptions underlying the model, and the methodology for deriving our EOSFs. In Sect. 4 we present our econometric analysis comparing actual supply functions to our calculated EOSFs, and discuss some findings of learning and economic withholding by various PGCs. Sect. 5 presents our methodology for deriving Nash equilibrium sets of ex ante-optimal supply functions, and compares these to the actual offers of what we term the strategic bidders in the market. Sect. 6 concludes our analysis, discusses some implications for the efficient design of spot balancing markets, and highlights some of the issues underlying the use of supply function models by regulators in market monitoring and design.

2 The ERCOT electricity markets

ERCOT acts as the system operator for the NERC region by the same name, which covers most of southern and central Texas.¹ Restructuring efforts in ERCOT began to take hold in 2001, with market-based trading and dispatch beginning in late 2001. Wholesale electricity trading, procurement of ancillary services, and reliability are achieved both through bilateral action on the part of market participants and a number of centrally operated energy markets, with the bulk of wholesale electricity traded bilaterally between parties. Prior to each day, market participants submit resource and obligation schedules to ERCOT through Qualified Scheduling Entities (QSEs). The QSEs are meant to act as intermediaries between stakeholders and the system operator, and they also act as aggregators for smaller generators and utilities—by allowing multiple firms to submit schedules and bids through a single QSE. For example, in 2001 the same QSE submitted schedules and bids for both Reliant Energy and City Public Service of San Antonio.

Invariably, suppliers and consumers of electricity would need a market through which to buy and sell excess energy, since load forecasts are never

¹ A small portion of western Texas is part of the WECC, northern Texas is part of SPP, and the region east of Houston is part of SERC.

exactly correct and to account for real-time contingencies such as line or generator outages. The BES is meant to serve as this spot market in which PGCs—through their QSEs—can submit bids to increment (inc) and decrement (dec) their generation. Up until mid-2002, QSEs were required to submit balanced day-ahead schedules. This balanced schedule requirement was intended to discourage the use of the BES as a market to procure baseload resources, and help ensure that the PGCs and load-serving entities (LSEs) faithfully use it only as a market of last resort in which to procure balancing energy. Since the balanced schedule requirement has been dropped the BES now averages slightly higher balancing sales than it used to, although it is still used mainly as a balancing market, with typically only 3–5% of total sales traded in the market. The market itself operates much like a commodity spot market. For each hour, PGCs submit price/quantity offers specifying the amount of energy they are willing to inc or dec at a given price, subject to a \$1,000 price cap.² Absent transmission constraints, the market clears as a single power pool—ERCOT aggregates the offers into a supply curve, and for each 15-minute interval intersects an essentially price-inelastic demand for balancing energy³ with the supply curve to determine the least-cost dispatch and a uniform MCPE. Participation in the BES market is voluntary, with the exception of a regulatory rule imposed by the Public Utility Commission of Texas (PUCT) that all QSEs are required to offer to decrement at least 15% of their scheduled energy at any price within the price caps. The PUCT's rationale behind this requirement was to ensure that adequate decremental energy is available due to a fear that the QSEs would overschedule resources day-ahead.

Given the basic characteristics of the BES, an SFE-type model should be a theoretically sound representation. SFE assumes firms commit themselves to supply functions—which is the equivalent of submitting quantity/price offers. The market clearing mechanism intersects the aggregate supply of the firms with the market demand function, which need not be price-elastic. In reality, generators will have excellent information regarding their competitors. Operating costs are relatively easy to estimate using engineering techniques; the marginal generating cost of a fossil-fuel driven plant can be estimated from its heat rate (a measure of its thermal efficiency), which can be determined by combining institutional knowledge and a variety of commercial sources of heat rate information. Given that we as academicians were able to estimate these costs, it is no leap of the imagination to assume that generators can do it as well. As for the actual operating status of a rival's plants, a PGC will see a brief fluctuation in the power grid's voltage and frequency if a large generator is taken offline. Moreover, there are firms that monitor the operational status of plants and sell this information on a real-time basis. Thus, a PGC should be able to predict which of its rival's plants are operating at any given time. Finally, generators

² Note that negative offer prices are allowed, primarily for decremental energy, with a price floor of $-\$1000$.

³ Although ERCOT does allow demand-side bids, so few are submitted at such high prices that balancing load is for all intents and purposes price-inelastic.

interact in the BES market on an hourly basis everyday. This repeated interaction essentially makes this an infinite-horizon repeated game. The SFE model assumes that firms play Nash equilibrium strategies, which is often an unrealistic assumption in single-shot games due to bounded rationality of players, difficulty in predicting rivals' behavior, and other cognitive and behavioral limitations of the parties involved. Due to the repeated nature of the BES market, it is possible that even if generators would not play Nash equilibrium strategies in a single-shot game, they may be able to converge towards a Nash outcome through the repeated interaction and its associated learning effects.⁴

3 A supply function model of the ERCOT BES

Based on the inherent characteristics of the ERCOT BES market it is a common belief that an SFE-type model should well describe the behavior of the firms involved. The specific model we use is an SFE which takes into account the contracted supply position of each generating firm, while allowing for uncertainty in demand for balancing energy. Our derivation shows that contractual obligations affect optimal bidding only through the quantity contracted and not the contract price.⁵

3.1 Derivation of supply function model

To derive generator i 's optimal offer curve, we solve its profit-maximization problem for any realization of system load. We define the notation $s_j(p)$ to be firm j 's supply function—specifying quantity supplied at each price— $c_j(q_j)$ to denote firm j 's total cost function, Q_j^{DA} to be the quantity that firm j has contractually obligated itself to supply at the contracted price, p_j^C ,⁶ and $D(p, \epsilon)$ to be the stochastic market demand for balancing energy. We assume that this market demand has the separable form $D(p, \epsilon) = D(p) + \epsilon$, in which $D(p)$ is a deterministic function of price, and ϵ is a random shock with support $\epsilon \in [\epsilon_{\min}, \epsilon_{\max}]$. Firm i 's objective is to maximize its profits:

⁴ It is worth noting that the repeated interaction in the BES does also allow for supergame equilibria, especially a multitude of cooperative or collusive equilibria in which suppliers raise the MCPE above what would result from repeated stage-game Nash behavior. Based on our analysis of bidding behavior, however, we do not believe this to be the case in the BES.

⁵ Specifically, our derivation shows optimal offers in the BES are independent of the contract price so long as that price is not a function of the MCPE. Because bilateral contracts are meant to hedge against spot price volatility, contracts generally exhibit this type of price independence.

⁶ Because only the contract position and not the price affects the profit-maximizing behavior of a firm in the BES market, we assume a single contract price for notational ease. One could reformulate the problem with multiple contract prices, although the result would remain the same.

$$\begin{aligned}
\max_p \Pi_i(p) &= p \cdot \left(D(p) + \epsilon - \sum_{j \neq i} (s_j(p) + Q_j^{DA}) \right) \\
&\quad - c_i \left(D(p) + \epsilon - \sum_{j \neq i} (s_j(p) + Q_j^{DA}) \right) - (p - p_i^C) Q_i^{DA} \\
&= p \cdot RD_i(p, \epsilon) - c_i(RD_i(p, \epsilon)) - (p - p_i^C) Q_i^{DA}, \tag{1}
\end{aligned}$$

where $RD_i(p, \epsilon) = D(p) + \epsilon - \sum_{j \neq i} (s_j(p) + Q_j^{DA})$ is firm i 's residual demand function, for any possible realization of ϵ . Note that we have specified forward contracts to settle as contracts for differences based on the prevailing MCPE in the BES market. Firm i sells all its generation through the BES market and reimburses (or is reimbursed by) its contract counterparties for the difference between the ex post realization of the BES MCPE and the contracted price. Differentiating Eq. 1 and setting the result equal to zero gives the first-order necessary condition (FONC) for firm i 's profit-maximization problem:

$$p - c'_i(RD_i(p, \epsilon)) = - \frac{RD_i(p, \epsilon) - Q_i^{DA}}{\frac{\partial}{\partial p} RD_i(p, \epsilon)}. \tag{2}$$

Given the offers of its rivals, the FONC in (2) is a differential equation characterizing firm i 's optimal choice of p for each possible ϵ . Moreover, our specification of the profit function as a contract for differences implicitly defines a boundary condition on this differential equation. If $RD_i(p, \epsilon) = Q_i^{DA}$ (firm i has zero dispatch in the BES), then $p = c'_i(RD_i(p, \epsilon))$. Furthermore, if firm i 's rivals bid non-decreasing supply functions (which they must in the BES), we will have $\frac{\partial}{\partial p} RD_i(p, \epsilon) < 0$. This then gives us $RD_i(p, \epsilon) > Q_i^{DA} \implies p > c'_i(RD_i(p, \epsilon))$ and $RD_i(p, \epsilon) < Q_i^{DA} \implies p < c'_i(RD_i(p, \epsilon))$, i.e. each firm will markup its inc offers above marginal cost, and markdown its dec offers below marginal cost. When a PGC is dispatched to inc generation, it is paid the MCPE and incurs the marginal generating cost of increasing output. Symmetrically, if a PGC is dispatched to dec, it foregoes the cost of generation but must pay ERCOT the MCPE (it essentially 'buys' its scheduled generation back from ERCOT). Thus, the markup and markdown rules implied by Eq. 2 make intuitive sense.

3.2 Assumptions and data

Derivation of optimal supply functions requires data on generation costs of the firms. Implicit in our derivation is that firms decide their bidding at a firm-wide level—as opposed to generation plants or units making individual offer decisions. Our analysis is confounded by an identification problem due to the QSE relationships between firms. A number of PGCs interact with ERCOT through a QSE which is used by other PGCs. In most cases, however, the schedules of individual PGCs can be identified because their generation assets and bids are

Table 1 Individually-identified generating firms

PGC	% Gen. Cap.	PGC	% Gen. Cap.
TXU	22	BP Energy	< 1
Reliant Energy	17	Bryan Texas Utilities	< 1
Calpine	8	City of Garland	< 1
Central Power and Light	6	Rio Nogales Power Project	< 1
City of San Antonio Public Service	6	Tenaska Gateway Partners	< 1
City of Austin	4	Cogeneration Lyondell	< 1
Lower Colorado River Authority	3	Bastrop Energy Partners	< 1
West Texas Utilities	2	Mirant Wichita Falls Management	< 1
Midlothian Energy	2	South Texas Electric Cooperative	< 1
Guadalupe Power Partners	2	Brownsville Public Utility Board	< 1
Lamar Power Partners	2	AES Deepwater	< 1
Brazos Electric Power Cooperative	1	Gregory Power Partners	< 1
Sweeny Cogeneration General	< 1	Extex Laporte	< 1
Hays Energy	< 1	Denton Municipal Electric	< 1
Tractabel Power	< 1	Air Liquide	< 1
Ingleside	< 1		

located within a congestion zone where no other PGC sharing the same QSE is present. Thus, although two PGCs may use a single QSE, if their assets are in different zones, then the QSE schedules and offers from the two firms can be distinguished. In some cases, however, multiple PGCs sharing congestion zones submit offers through a single QSE—in these cases individual offers cannot be distinguished. We assume that if a single PGC represents more than 70% of the actual electricity generated within a congestion zone for its QSE, then all offers and schedules from that QSE within that congestion zone are for that single firm. As Table 1 shows, we are able to cover the major PGCs and the vast majority of the bidding assets in ERCOT. Optimal offers for firms which cannot be identified are not derived, but their *actual* offers are used in conjunction with actual system load in deriving the deterministic portion of market demand, $D(p)$.

To derive each firm's cost function, we assume that plants in its generation portfolio are dispatched in economic merit order. Thus, a firm that is generating 5,000 MW will generate its 5,000 cheapest megawatts available. PGCs dispatching resources out of merit order due to operational or other constraints is not captured in our analysis. Because of their significant ramping constraints, we assume nuclear units do not bid strategically and are instead run at 100% of available capacity. Finally, because of the difficulty in estimating resource availability, we exclude hydroelectric, wind, and solar plants from the generation portfolios. We feel justified in making this simplifying assumption since these renewables constitute less than 5% of ERCOT's installed generation capacity.

For fossil fuel-driven plants, we assume that each unit has a constant marginal cost. The fuel cost is the product of its heat rate and fuel price. We impute an average heat rate for each month using heat produced and net generation for each plant as reported in EIA Form 906. For months in which that data is unavailable, the tested heat rate is used instead. We realize that there is some endogeneity from using this average heat rate, as it is affected by bidding behavior through the actual dispatch of a given unit, but since the bulk of generation is traded bilaterally we believe this effect to be minimal. Fuel prices for natural gas is estimated using the Henry Hub spot price plus \$0.10/mmBTU for transportation. Although it is common practice for PGCs to contract for fuel and pay a price different from the spot price, it nonetheless represents the opportunity cost of burning the fuel. The cost of other fuels are estimated using the heat content-weighted average of that fuel procured in each given month, as reported in EIA Form 423. In addition to fuel costs, generators are also subject to emission fees from both the US Environmental Protection Agency (EPA) and the Texas Commission on Environmental Quality (TCEQ). The TCEQ charges each polluting plant the greater of a fixed fee to administer its monitoring program and a charge based on actual emissions of pollutants. For plants paying the fixed fee, we assume no marginal emission cost. For plants subject to the variable emission-based charges, we estimate the cost per megawatt by dividing the total charge for the year by the plant's net generation for the year. In addition to the TCEQ's emission program, the EPA charges for SO₂ emissions as part of its acid rain program. The program is administered through an emissions trading program, whereby a polluter must obtain emission credits for each ton of SO₂ emitted. Using TCEQ data, we are able to estimate the average SO₂ output per megawatt-hour generated and multiply that by the cost of an emission permit as reported by Cantor-Fitzgerald Environmental Trading Brokerage. Finally, we add an estimated variable operations and maintenance cost for each plant based on its generating technology.

In determining a firm's cost of providing inc and dec service, we assume that all units which had not experienced an outage (as recorded in ERCOT's outage scheduler) were available to ramp generation. Thus, we ignore ramping constraints and any intertemporal constraints on a unit's on or off time. Taking account of these constraints would require detailed operational data not available to us.

Our derivation of the supply function model showed that a PGC's forward contract position will affect its profit-maximizing offer function only through the quantity contracted, Q_i^{DA} , and not the contract price. One difficulty in accounting for this impact is that, in general, a PGC's contract position is private information which is not available to us. In their analysis of the BES, Hortaçsu and Puller (2005) estimate each PGC's contract position by counterfactually applying the boundary condition of Eq. 2, which states that a PGC's optimal markup is zero if its imbalance quantity is zero. That is, letting p_i^{DA} be the price at which PGC i offers zero imbalance sales, if PGC i is submitting offers in accordance with the supply function model then its contracted quantity should satisfy $c'_i(Q_i^{DA}) = p_i^{DA}$. Although this method of finding contracted sales is

correct when PGCs are acting in accordance with the supply function model, our aim is to analyze the extent to which firm behavior is consistent with the model and as such would prefer to use a forward contract estimate which does not require us to assume that PGCs behave as the model dictates. We therefore follow the approach seen in Niu (2005), which is to use a PGC's day-ahead schedule as a proxy for the contract position. Although a generator may wish to deviate from its contracted sales, we implicitly assume any such deviation can be implemented through the BES and as such a PGC would schedule its contracted position.

3.3 Derivation of ex post-optimal supply functions

We begin our analysis by first deriving for each firm a set of EOSFs. That is, in each bidding period, for each firm, we find a supply function satisfying Eq. 2 for every possible ϵ , given the *actual realized* offers of its rivals. We then compare the actual offers of each firm to our theoretically EOSFs. One reason for conducting this ex post analysis is to screen-out PGCs whose behavior is far from optimal in our subsequent evaluation of the Nash equilibrium SFE model. Qualitative analysis of the offers shows that most PGCs, especially smaller municipalities, cooperatives, and cogenerators, are reluctant to deviate from their scheduled generation by participating in the BES. These firms will typically economically withhold their generation by making offers with substantial markups and markdowns to minimize the odds of being dispatched except when there is extremely high demand for BES energy. Because the EOSFs are the optimal reaction to the actual offers of each PGC's rivals, it captures the degree to which a firm is able to anticipate its rivals' behavior and optimally react, even when rivals are not Nash players. By excluding these obvious non-Nash players in our subsequent Nash calculation, we will effectively recognize that such persistent behavior is accounted for by the strategic PGCs, whose behavior we try to model with our Nash equilibrium model.

One complication which arises in the supply function analysis is the format of offers into the BES. The differential equations governing the optimal behavior of the firms assumes continuously differentiable supply functions. Offers into the BES, however, are price/quantity pairs defining a step function. ERCOT limits each QSE to submitting 20 steps for each of the inc and dec side of its supply function (giving a total of 40 steps) in each bidding period. von der Fehr and Harbord (1993) raise the issue of step functions confounding an SFE-type analysis, concluding that only mixed-strategy equilibria will exist. Baldick and Hogan (2002) argue to the contrary saying that with enough offer points a step function can closely resemble a continuously differentiable supply function. While ERCOT allows a total of 40 steps in each offer curve, PGCs actually use around 5–10 steps, making the argument somewhat tenuous. Nonetheless, in comparing offer curves, we 'flatten' our optimal offer curve into a step function to make it conform to the actual offer curve submitted. In doing so, we assume the offer quantities to be fixed, based on those quantities actually used by the firm. For instance, in the hour ending 3:00am on 4 August, 2002 Tenaska

Gateway Partners' offer curve consisted of 10 steps—5 for inc and 5 for dec service. The inc offers had steps at 625, 626, 650, 651, and 768 MW. In deriving optimal supply functions we assume that the PGC uses the same offer quantities, and we flatten the supply function by using the optimal prices at those quantities in each 'flat.' Furthermore, in constructing each firm's residual demand curve, the step function format of the offer curves will give a stepped residual demand function. Since it would be overly zealous to assume that PGCs anticipate the exact location of these steps, we smooth-out the residual demand curve using a kernel function as seen, for instance, in Wolak (2003). We estimate the derivative of the residual demand function using a finite difference method, as direct differentiation of the kernel function proved highly sensitive with respect to the choice of the smoothing parameter.

Our period of study is 2002 and 2003, with all bid periods (24 per day) included. However, assuming the BES clears as a simple power pool requires that there be no binding transmission constraints. Thus, any hour with a congested transmission line would have to be removed from our sample, since the BES would clear with multiple MCPEs for each congestion zone. In accounting for the effects of congestion, we removed any day in which there was *any* interzonal congestion. The rationale for removing days with *any* congestion from the sample is that any anticipation of congestion on the part of PGCs could likely lead to substantially different behavior on the part of any PGCs with locational market power. We felt that excluding days with any interzonal congestion (even in a few of the 15-min-long clearing periods) would, to some extent, control for such a distortion of the bids.⁷ This restriction does not, however, allow us to control for periods in which transmission lines were nearly congested. Again, any anticipation of congestion in these instances could likely affect a PGC's optimal offer decision. After excluding congested days, our sample was reduced to 340 days, giving a total of 8,160 bidding periods.

4 Comparison of actual offers to ex post-optimal supply functions

In comparing the actual to optimal offer curves, our approach is to study the nature of the entire range of the supply function as opposed to using pointwise tests of optimality, such as those seen in Hortaçsu and Puller or Niu's study of the BES. A simple qualitative comparison of the actual and theoretical offer curves suggests that the supply function equilibrium model is a rather poor representation of the behavior of most firms. We find that almost all firms offer dec service with substantial markdowns from marginal cost, which are much greater than can be explained by our model. Furthermore, most firms overbid their inc service, with the exception of the two largest power producers: TXU and Reliant Energy. The third major player in the market, Calpine, submits

⁷ Our sample selection criterion differs from that in Hortaçsu and Puller who use the same single hour in each day and exclude hours during which congestion occurred. Thus if congestion occurred in periods adjacent to the sample period but not during the sample period the sample was not excluded.




Fig. 1 Actual and optimal offers example

inc offers which are on average about one order of magnitude greater than ex post-optimal, although we later show that this is due to substantially overpriced offers early on in the infancy of the BES market. Figure 1 shows a sample actual offer curve and EOSF for City of San Antonio. Note that in the particular bid period shown, inc offers (quantities greater than zero) match rather closely, whereas the dec offers differ substantially with the offers in the middle of the curve being marked down more than the optimal and the markdown being less than optimal for extreme decs.

4.1 Non-parametric model of firm behavior

In order to compare the offers quantitatively we estimate an econometric model of firm behavior which posits that generators, in making their offers, choose markups over marginal cost to be some multiple of the theoretically optimal markup. Defining $b_{i,t}(q)$, $b_{i,t}^*(q)$, and $c'_{i,t}(q)$ to be firm i 's actual offers, optimal offers, and marginal cost function (respectively) in bidding period t , we can define firm i 's actual and optimal markup in bidding period t as $MU_{i,t}(q) = b_{i,t}(q) - c'_{i,t}(q)$ and $MU_{i,t}^*(q) = b_{i,t}^*(q) - c'_{i,t}(q)$, respectively. Our model is then:

$$MU_{i,t}(q) = \phi_i \cdot MU_{i,t}^*(q),$$

where the multiplier, ϕ_i , can be thought of as a measure of conduct. A value of ϕ_i close to zero yields zero markup or perfectly competitive behavior. Higher values of ϕ_i would be indicative of more rational profit-maximizing behavior, with $\phi_i = 1$ being perfect rationality. Values of $\phi_i > 1$ would be indicative of a firm trying to exclude itself from the market, overly zealous exercise of market

power, collusion, or other behavior which is inconsistent with or not accounted for by our model.

One possible estimation approach would be to assume the ϕ_i parameter fixed along the range of the supply function, in which case the model could be estimated by standard parametric techniques such as least squares. We believe this specification would be overly restrictive on firm behavior, by assuming constant ‘conduct’ in bidding throughout their offer curves. The shape of actual offer curves suggests that some firms opt to ‘hockey-stick’ their supply functions—offering most of their generation at reasonable prices and a small quantity at a very high price, giving their offer curve a hockey-stick shape—even more than theoretically optimal. Furthermore, we find that most PGCs submit dec offers with extremely high markdowns (suggesting a high value of ϕ_i), but some offer inc bids with more reasonable markups (suggesting a lower value of ϕ_i). In view of such anecdotal evidence against a fixed multiplier, we opt for a model with a variable conduct multiplier that varies with quantity:

$$MU_{i,t}(q) = \phi_i(q) \cdot MU_{i,t}^*(q), \quad (3)$$

where $\phi_i(q)$ is an unspecified smooth function of the offer quantity, q . In order to estimate this model, we divide Eq. (3) through by the optimal markup, $MU_{i,t}^*(q)$, to yield a standard non-parametric model:

$$\frac{MU_{i,t}(q)}{MU_{i,t}^*(q)} = \phi_i(q) + \eta_{i,t}(q), \quad (4)$$

where the error term, $\eta_{i,t}(q)$, with $\mathbb{E}[\eta_{i,t}(q)] = 0$ and $\text{Var}(\eta_{i,t}(q)) < +\infty$, allows for the fact that a firm may misestimate its rivals’ offers or miscalculate its own optimal reaction. We estimate the model in Eq. 4 using a Nadaraya (1964)—Watson (1964) kernel estimator, in which $\phi_i(q)$ is estimated as:

$$\hat{\phi}(q) = \frac{\sum_{n=1}^N K\left(\frac{q-q^n}{h}\right)\phi^n}{\sum_{n=1}^N K\left(\frac{q-q^n}{h}\right)},$$

where ϕ^n and q^n are observed markup ratios and corresponding imbalance quantity. The function, $K(\cdot)$, is called a kernel or smoothing function, and we opt for the standard of using the normal density function. The parameter, h , which is called the window width or bandwidth, controls the weight put on neighboring observations in estimating $\hat{\phi}(q)$. We use an optimal bandwidth, $h^* = \mathcal{O}(N^{-1/4})$, which ensures the estimator is consistent.

In order to make offers between different periods comparable we normalize the offer quantities, q , to be the fraction of the total offer amount that a firm makes in each given bid period, meaning q is restricted to $q \in [-1, 1]$. For example, in Fig. 1, San Antonio made 4 inc and 5 dec offers for that hour. Its inc quantities were for 15, 10, 25, and 25 MW, giving a total offer of 75 MW. Those bid points would correspond to $q = \{1/5, 1/3, 2/3, 1\}$. Symmetrically, the dec

bids would correspond to values of $q \in [-1, 0)$. We make these normalizations because firms may offer varying absolute quantities in different bid periods, yet our model attempts to capture how the markup multiplier, $\phi_{i,t}(q)$, varies along a firm's supply function.

4.2 Non-parametric estimates

By using the offer data from each individual firm, we estimate the model in Eq. 4, comparing its actual offers to a supply function which is an optimal response to the realized offers of its rivals. In making this comparison we seek to determine the extent to which each PGC can conjecture the actual behavior of its rivals and optimally react to that conjecture by submitting offers that maximizes its profits ex post. In general, this could prove to be a complicated task as a firm would have to predict the behavior of each rival and take account of that in making its own offer decision. Due to the repeated nature of the interaction, though, we surmise that over time a rational generator may be able to anticipate its rivals' behavior and adjust its pattern of offers to that behavior. As Fig. 1 shows, City of San Antonio has been able to make such an adjustment to some extent in submitting inc offers which closely match its ex post-optimal supply function.

Figures 2–4 show the estimated $\hat{\phi}_{NP}(q)$ for the three largest PGCs participating in the BES, along with an asymptotic pointwise 95% confidence interval, and a $\phi = 1$ line which would correspond to ex post profit-maximizing behavior. In addition, Table 2 summarizes the estimated values and gives upper- and lower-confidence interval bounds of $\hat{\phi}_{NP}(q)$ for the three major PGCs. Table 3 gives summary estimates for the conduct curves of all PGCs. Our estimates show that with the exception of TXU and Reliant Energy, most PGCs' markups are several orders of magnitude above those predicted by our EOSF model. This result is consistent with anecdotal evidence as well as other analyses of the BES market.

We note that many of the $\hat{\phi}_{NP}$ estimators have a peak in the neighborhood of $q \approx 0$ with a corresponding wide confidence interval band. Recall that offers with $q \approx 0$ correspond to net balancing sales $RD_i(p, \epsilon) - Q_i^{DA} \approx 0$, which our derived optimality condition implies should have a markup or markdown close to zero. Thus, the ratio $\phi_i = MU_{i,t}(q)/MU_{i,t}^*(q)$ will be very sensitive to any error a firm makes in calculating its offers in this range or in our estimating a firm's marginal cost (which is subject to some uncertainty), giving the peak and the corresponding wide variance band due to the heteroskedasticity of the errors at $q \approx 0$.

The estimates of TXU's conduct curve show their actual markups to be within a reasonable multiple of our derived EOSFs, at between 10% and 120% of optimal. This is still somewhat surprising as TXU is a major holder of generating assets within ERCOT (approximately 22% of nameplate capacity) and often a pivotal supplier in the BES, which affords it much market power. One possible explanation for this apparent restraint is the fear of regulatory action. After




Fig. 2 TXU conduct function estimates

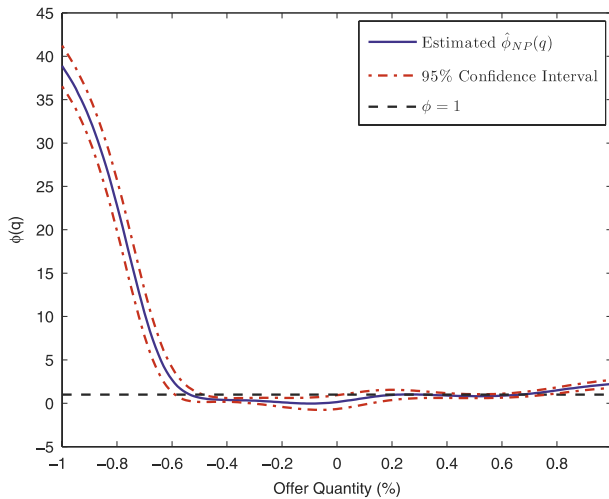


Fig. 3 Reliant energy conduct function estimates

the inception of the BES market, the PUCT has revised the ERCOT tariffs several times, adding various market mitigation rules such as: a ‘shame-cap,’ which immediately reveals the identity of any QSE submitting an inc or dec bid into the BES above or below a certain price; and a ‘hockey-stick’ curtailment rule, which Hurlbut, Rogas, and Oren (2004) designed to help mitigate price-gouging in bid periods where the BES offer stack is exhausted. Moreover, the PUCT has been under political pressure from advocacy groups to impose further market-mitigation rules due to the perceived persistence of unjustifiably high MCPs. Many of these market-mitigation procedures have been opposed by the PGCs, including TXU. Part of TXU’s restraint in its bidding behavior




Fig. 4 Calpine conduct function estimates

Table 2 Summary statistics of conduct function estimators

q	-1	-0.8	-0.6	-0.4	-0.2	0	0.2	0.4	0.6	0.8	1
TXU											
$\hat{\phi}_{NP}(q)$	0.410	0.298	0.155	0.272	0.712	1.247	0.235	0.089	0.109	0.167	0.226
$\hat{U}_{NP}(q)$	0.446	0.354	0.185	0.360	1.443	2.705	0.758	0.141	0.141	0.192	0.241
$\hat{L}_{NP}(q)$	0.373	0.242	0.125	0.184	-0.020	-0.211	-0.288	0.037	0.077	0.141	0.210
Reliant Energy											
$\hat{\phi}_{NP}(q)$	38.888	22.816	2.702	0.386	0.108	0.141	0.962	0.896	0.875	1.486	2.225
$\hat{U}_{NP}(q)$	41.24	25.758	4.120	0.606	0.630	0.936	1.548	1.169	1.107	1.843	2.697
$\hat{L}_{NP}(q)$	36.536	19.873	1.284	0.166	-0.413	-0.654	0.377	0.623	0.644	1.129	1.752
Calpine											
$\hat{\phi}_{NP}(q)$	37.851	22.703	13.476	19.051	36.385	33.531	3.637	0.202	0.947	9.775	21.785
$\hat{U}_{NP}(q)$	65.276	42.197	20.046	24.974	62.567	88.293	23.714	3.239	7.254	32.741	52.15
$\hat{L}_{NP}(q)$	10.425	3.208	6.907	13.127	10.203	-21.231	-16.441	-2.835	-5.360	-13.192	-8.580

may reflect a form of self-imposed mitigation to keep the MCPE sufficiently low to try and ward off further regulatory action. Due to TXU’s position as the dominant PGC in the market, it was a pivotal supplier in many of the bidding periods in 2002 and 2003, in which case it could have easily set the MCPE at the price cap of \$1,000. Such behavior, however, would have triggered an investigation and potentially regulatory intervention. A second plausible rationale for TXU’s bidding behavior is the native load served by TXU Energy—a subsidiary of the TXU Corporation which is one of the largest LSEs in Texas. During the study period, TXU Energy’s retail rates were frozen by the PUCT, and given TXU Energy’s large customer base of 2.9 million, an excessively high MCPE may have actually reduced the holding company’s total profits as it could not recover the full cost of energy procurement through retail rates. The effect of

Table 3 Firm conduct estimates for ex post-optimal bid curves

q	-1	-0.8	-0.6	-0.4	-0.2	0	0.2	0.4	0.6	0.8	1
TXU											
$\hat{\phi}_{NP}(q)$	0.410	0.298	0.155	0.272	0.712	1.247	0.235	0.090	0.109	0.167	0.226
$\hat{U}_{NP}(q)$	0.446	0.354	0.185	0.360	1.443	2.705	0.758	0.141	0.141	0.192	0.241
$\hat{L}_{NP}(q)$	0.373	0.242	0.125	0.184	-0.020	-0.211	-0.288	0.037	0.077	0.141	0.210
Reliant Energy											
$\hat{\phi}_{NP}(q)$	38.888	22.816	2.702	0.386	0.108	0.141	0.962	0.896	0.875	1.486	2.225
$\hat{U}_{NP}(q)$	41.24	25.758	4.120	0.606	0.630	0.936	1.548	1.169	1.107	1.843	2.697
$\hat{L}_{NP}(q)$	36.536	19.873	1.284	0.166	-0.413	-0.654	0.377	0.623	0.644	1.129	1.752
Calpine											
$\hat{\phi}_{NP}(q)$	37.851	22.703	13.476	19.051	36.385	33.531	3.637	0.202	0.947	9.775	21.785
$\hat{U}_{NP}(q)$	65.276	42.197	20.046	24.974	62.567	88.293	23.714	3.239	7.254	32.741	52.15
$\hat{L}_{NP}(q)$	10.425	3.208	6.907	13.127	10.203	-21.231	-16.441	-2.835	-5.360	-13.192	-8.580
Central Power and Light											
$\hat{\phi}_{NP}(q)$	6.482	7.086	8.252	11.3	17.665	19.195	-17.738	-14.751	-10.479	-6.581	-4.454
$\hat{U}_{NP}(q)$	9.566	10.218	10.349	13.145	21.818	53.54	0.142	-8.226	-7.324	-4.197	-1.919
$\hat{L}_{NP}(q)$	3.398	3.954	6.155	9.455	13.511	-15.149	-35.618	-21.277	-13.634	-8.965	-6.990
City of San Antonio Public Service											
$\hat{\phi}_{NP}(q)$	2.921	2.481	2.855	4.327	6.882	6.523	0.625	2.498	5.795	10.692	15.167
$\hat{U}_{NP}(q)$	3.158	2.665	2.996	4.535	7.245	8.634	2.273	3.700	6.896	12.342	16.916
$\hat{L}_{NP}(q)$	2.684	2.298	2.715	4.118	6.518	4.413	-1.023	1.297	4.693	9.042	13.419
City of Austin											
$\hat{\phi}_{NP}(q)$	12.192	6.335	3.536	4.263	7.896	20.122	30.647	25.85	23.646	27.333	33.879
$\hat{U}_{NP}(q)$	12.919	6.935	3.914	4.858	10.346	36.898	56.852	46.089	40.105	40.209	44.359
$\hat{L}_{NP}(q)$	11.466	5.734	3.158	3.668	5.445	3.347	4.442	5.611	7.188	14.457	23.4
Lower Colorado River Authority											
$\hat{\phi}_{NP}(q)$	4.908	4.987	5.450	7.100	9.099	7.459	-0.807	-0.949	-1.022	1.868	3.490
$\hat{U}_{NP}(q)$	5.680	5.756	5.994	7.749	10.118	9.776	1.595	0.487	0.534	5.689	6.993
$\hat{L}_{NP}(q)$	4.135	4.218	4.906	6.451	8.081	5.142	-3.208	-2.386	-2.578	-1.954	-0.013
West Texas Utilities											
$\hat{\phi}_{NP}(q)$	8.363	8.122	10.655	19.319	31.167	19.712	-42.041	-30.326	-22.213	-14.949	-10.192
$\hat{U}_{NP}(q)$	9.878	9.188	12.177	21.456	34.289	63.716	11.601	1.605	-1.858	-1.648	-1.222
$\hat{L}_{NP}(q)$	6.849	7.056	9.132	17.182	28.046	-24.293	-95.683	-62.256	-42.568	-28.249	-19.161
Midlothian Energy											
$\hat{\phi}_{NP}(q)$	86.959	69.883	50.416	430.98	1821	2415.9	1946.8	399.65	8.887	-0.184	-0.685
$\hat{U}_{NP}(q)$	112.54	108.07	123.5	772.12	2383.4	2839.3	2656.3	837.64	73.59	3.909	0.600
$\hat{L}_{NP}(q)$	61.375	31.693	-22.669	89.847	1258.6	1992.4	1237.4	-38.328	-55.815	-4.276	-1.971
Guadalupe Power Partners											
$\hat{\phi}_{NP}(q)$	137.69	105.76	70.338	47.232	202.87	846.16	216.07	16.876	3.274	27.155	116.1
$\hat{U}_{NP}(q)$	158.29	134.64	94.787	75.534	275.21	1014	363.75	59.314	16.801	44.643	142.09
$\hat{L}_{NP}(q)$	117.1	76.871	45.889	18.929	130.52	678.32	68.399	-25.562	-10.254	9.666	90.106
Lamar Power Partners											
$\hat{\phi}_{NP}(q)$	3299.7	3054.2	1192.5	32.262	4.591	0.309	-5.148	-2.923	-1.416	-2.572	-3.773
$\hat{U}_{NP}(q)$	3982.3	3860.8	1998.1	155.34	17.65	12.728	2.464	2.190	16.376	24.903	29.265
$\hat{L}_{NP}(q)$	2617	2247.7	386.89	-90.815	-8.469	-12.111	-12.759	-8.036	-19.208	-30.047	-36.811
Brazos Electric Power Cooperative											
$\hat{\phi}_{NP}(q)$	240.14	153.06	58.654	48.951	63.833	89.879	112.28	42.427	7.619	48.723	63.835
$\hat{U}_{NP}(q)$	258.63	181.44	73.436	60.235	79.14	125.89	396.89	135.58	98.954	254.29	134.97
$\hat{L}_{NP}(q)$	221.65	124.69	43.871	37.667	48.527	53.872	-172.32	-50.732	-83.716	-156.85	-7.299

Table 3 continued

q	-1	-0.8	-0.6	-0.4	-0.2	0	0.2	0.4	0.6	0.8	1
Sweeny Cogeneration General											
$\hat{\phi}_{NP}(q)$	72.132	65.248	80.056	120.92	150.36	-51.798	-165.05	-92.266	-60.122	-46.99	-35.291
$\hat{U}_{NP}(q)$	79.296	68.959	82.617	125.59	155.95	-1.311	-147.21	-85.501	-56.175	-42.795	-33.075
$\hat{L}_{NP}(q)$	64.968	61.537	77.496	116.25	144.77	-102.28	-182.88	-99.03	-64.07	-51.184	-37.506
Hays Energy											
$\hat{\phi}_{NP}(q)$	74.401	66.108	146.38	1533.8	3752.1	4255	3844.9	1901.2	145.79	-0.411	-1.911
$\hat{U}_{NP}(q)$	87.246	95.471	429.08	2459.3	4700.6	4853.1	4925.3	3353.2	611.45	35.956	1.182
$\hat{L}_{NP}(q)$	61.555	36.745	-136.32	608.33	2803.5	3656.9	2764.5	449.21	-319.88	-36.777	-5.004
Tractabel Power											
$\hat{\phi}_{NP}(q)$	408.7	371.92	375.67	2284	6403.9	7618.7	4249.3	482.14	-2.531	0.828	15.847
$\hat{U}_{NP}(q)$	435.56	436.93	857.9	3659.5	7767.1	8616.1	5486	925.11	51.921	9.2552	25.917
$\hat{L}_{NP}(q)$	381.84	306.92	-106.55	908.48	5040.7	6621.2	3012.6	39.171	-56.983	-7.598	5.777
Ingleside											
$\hat{\phi}_{NP}(q)$	659.66	150.31	23.786	13.834	16.168	0.215	-10.786	-6.120	-3.263	-3.647	-7.976
$\hat{U}_{NP}(q)$	832.64	229.6	46.04	20.423	17.632	6.556	-8.541	-5.28	-2.520	-1.771	-4.856
$\hat{L}_{NP}(q)$	486.69	71.031	1.531	7.245	14.705	-6.125	-13.032	-6.959	-4.006	-5.522	-11.097
BP Energy											
$\hat{\phi}_{NP}(q)$	1139.1	738.91	161.76	555.18	2267.4	2420.9	1036.4	135.6	21.412	201.92	396.49
$\hat{U}_{NP}(q)$	1253.4	973.98	286.45	1024.2	3054	2862.9	1426.5	306.18	93.832	488.14	558.66
$\hat{L}_{NP}(q)$	1024.7	503.83	37.068	86.209	1480.9	1978.9	646.35	-34.976	-51.008	-84.299	234.33
Bryan Texas Utilities											
$\hat{\phi}_{NP}(q)$	67.502	62.421	73.056	91.267	104.45	102.79	-29.111	-59.537	-40.866	-55.148	-73.723
$\hat{U}_{NP}(q)$	73.779	71.937	84.456	105.14	122.1	152.76	87.021	-13.49	-5.718	-24.767	-60.958
$\hat{L}_{NP}(q)$	61.224	52.906	61.656	77.391	86.809	52.814	-145.24	-105.58	-76.014	-85.529	-86.488
City of Garland											
$\hat{\phi}_{NP}(q)$	31.018	37.131	45.934	60.541	72.503	44.056	-49.748	-38.466	-26.87	-21.353	-20.896
$\hat{U}_{NP}(q)$	32.217	38.807	47.866	63.218	77.552	75.353	-35.201	-31.337	-20.741	-14.073	-16.523
$\hat{L}_{NP}(q)$	29.818	35.455	44.002	57.865	67.454	12.76	-64.295	-45.595	-32.998	-28.634	-25.269
Rio Nogales Power Project											
$\hat{\phi}_{NP}(q)$	130.25	140.26	184.3	272.66	361.84	327.53	59.127	-7.509	-8.773	-6.187	-5.346
$\hat{U}_{NP}(q)$	161.52	173.09	230.04	341.29	469.88	527.04	153.7	15.574	-0.317	-0.750	-0.327
$\hat{L}_{NP}(q)$	98.985	107.43	138.56	204.03	253.79	128.02	-35.448	-30.593	-17.229	-11.624	-10.364
Tenaska Gateway Partners											
$\hat{\phi}_{NP}(q)$	23.603	10.066	6.707	7.232	10.2	1.885	-12.027	-5.519	-2.289	-0.755	-2.377
$\hat{U}_{NP}(q)$	26.747	12.04	7.684	8.035	11.051	4.790	-10.427	-3.987	-0.457	2.369	-1.174
$\hat{L}_{NP}(q)$	20.459	8.093	5.729	6.429	9.349	-1.021	-13.627	-7.051	-4.121	-3.879	-3.579
Cogeneration Lyondell											
$\hat{\phi}_{NP}(q)$	33.909	41.682	37.625	57.343	165.75	238.03	36.914	-8.322	1.543	156.63	231.86
$\hat{U}_{NP}(q)$	348.4	376.69	117.34	73.837	210.66	324.32	87.83	-2.745	17.484	223.97	294.08
$\hat{L}_{NP}(q)$	-280.58	-293.32	-42.089	40.848	120.84	151.74	-14.002	-13.899	-14.399	89.291	169.64
Bastrop Energy Partners											
$\hat{\phi}_{NP}(q)$	7107.2	6968.4	5388	719.98	8.163	-2.980	11.014	12.234	4.586	-0.562	-1.916
$\hat{U}_{NP}(q)$	8177.7	8103.8	6670.1	1401.9	135.12	15.845	21.962	36.881	71.695	97.758	107.94
$\hat{L}_{NP}(q)$	6036.6	5833	4105.9	38.038	-118.79	-21.804	0.067	-12.413	-62.524	-98.883	-111.78
Mirant Wichita Falls Management											
$\hat{\phi}_{NP}(q)$	1900.4	1652	1272.6	902.76	646.16	485.1	362.01	193.97	76.32	73.529	84.1
$\hat{U}_{NP}(q)$	2366.5	2418.3	2093.3	1580.4	1191.5	1103.7	1320.8	591.04	256.6	191.43	147.37
$\hat{L}_{NP}(q)$	1434.3	885.67	451.88	225.14	100.81	-133.51	-596.79	-203.11	-103.96	-44.377	20.826

Table 3 continued

q	-1	-0.8	-0.6	-0.4	-0.2	0	0.2	0.4	0.6	0.8	1
South Texas Electric Cooperative											
$\hat{\phi}_{NP}(q)$	246.67	223.42	261.99	384.41	516.49	581.31	119.56	-46.411	-46.736	-70.477	-75.031
$\hat{U}_{NP}(q)$	280.9	266.41	284.11	409.09	554.71	711.57	306.58	-9.351	-11.376	-46.834	-62.865
$\hat{L}_{NP}(q)$	212.44	180.44	239.87	359.74	478.26	451.05	-67.468	-83.472	-82.095	-94.121	-87.197
Brownsville Public Utility Board											
$\hat{\phi}_{NP}(q)$	1291.8	554.13	74.277	7.9159	5.0634	4.4561	0.61217	-1.313	-1.961	-8.626	-17.122
$\hat{U}_{NP}(q)$	1699.7	851.84	178.92	28.581	11.426	7.055	1.767	-0.237	1.075	0.720	-6.346
$\hat{L}_{NP}(q)$	883.83	256.42	-30.363	-12.749	-1.299	1.857	-0.543	-2.389	-4.997	-17.972	-27.898
AES Deepwater											
$\hat{\phi}_{NP}(q)$	48.884	48.884	48.884	48.885	49.378	316.22	585.6	586.1	586.1	586.1	586.1
$\hat{U}_{NP}(q)$	72.919	74.573	80.243	92.525	117.53	587.56	791.82	719.64	682.07	664.71	659.65
$\hat{L}_{NP}(q)$	24.849	23.196	17.525	5.244	-18.773	44.893	379.37	452.56	490.14	507.49	512.55
Gregory Power Partners											
$\hat{\phi}_{NP}(q)$	71.079	73.507	114.15	140.83	126.88	24.593	-33.78	-18.166	-5.669	2.490	2.893
$\hat{U}_{NP}(q)$	78.674	84.024	137.63	166.58	160.98	79.681	-1.816	-6.3	6.169	19.022	18.114
$\hat{L}_{NP}(q)$	63.485	62.991	90.674	115.08	92.787	-30.495	-65.744	-30.032	-17.506	-14.041	-12.329
Extex Laporte											
$\hat{\phi}_{NP}(q)$	253.88	235.31	151.71	16.094	1.523	1.564	2.230	3.197	3.982	3.996	3.792
$\hat{U}_{NP}(q)$	343.13	575.13	600.84	187.94	45.69	7.981	3.309	3.932	4.482	4.545	4.186
$\hat{L}_{NP}(q)$	164.62	-104.51	-297.41	-155.75	-42.643	-4.853	1.151	2.462	3.482	3.447	3.398
Denton Municipal Electric											
$\hat{\phi}_{NP}(q)$	85.735	63.219	42.33	53.15	89.583	153.19	197.96	509.23	3021.7	4885.7	5285.9
$\hat{U}_{NP}(q)$	92.863	82.649	59.818	89.571	153.93	434.32	713.13	1365.7	4441.7	6735.6	5592.3
$\hat{L}_{NP}(q)$	78.606	43.79	24.842	16.728	25.239	-127.94	-317.21	-347.29	1601.7	3035.9	4979.6
Air Liquide											
$\hat{\phi}_{NP}(q)$	644.11	613.1	652.72	921.92	1362.3	651.69	42.235	-18.553	-10.483	-16.994	-26.711
$\hat{U}_{NP}(q)$	709.21	700.35	793.14	1179.2	1863.3	1047.2	167.52	8.163	-4.304	-10.769	-19.871
$\hat{L}_{NP}(q)$	579.01	525.85	512.29	664.67	861.4	256.19	-83.046	-45.269	-16.663	-23.22	-33.552

this restraint may eventually wear down, however, once the regulated ‘price to beat’ lapses in the ERCOT market.

Similar to TXU, Reliant’s actual inc bids seem to match our theoretical optima rather closely, with the $\hat{\phi}_{NP}$ estimator being relatively close to 1. In contrast to TXU, though, Reliant’s dec bids are on average marked down far below what our model predicts. This pattern of high markdowns on dec bids, which is seen with most PGCs has also been observed by others who have studied the BES. The behavior is attributed to both a reluctance on the part of PGCs to ramp down their generation (especially combined-cycle gas turbines) due to heat rate and ramping considerations, as well as other costs not accounted for in standard engineering estimates such as higher maintenance costs and gas imbalance charges. Hortaçsu and Puller analyze the ‘bid-ask spread’ between the lowest inc and highest dec bid and find that most PGCs have a spread which is wider than can be explained by the estimated cost of adjusting output.

The most surprising behavior we find amongst the large generators is that of Calpine. Calpine is an independent power producer with no native load obligation, less regulatory oversight than the incumbent investor-owned utilities

(IOUs) such as TXU and Reliant, and it has been increasing its share of asset holdings in Texas (both through purchases and investments in new generation) to around 8% of nameplate capacity. In spite of its good position in the market and potential for exercise of market power, our estimates show Calpine's bid markups to average an order of magnitude greater than theoretically optimal, effectively pricing its generation out of the BES market. We find that Calpine's seemingly irrational behavior took place mainly in the first quarter of 2002, and we are able to show that there is a statistically significant difference in its bidding behavior during and after the first quarter of 2002. This suggests that Calpine was either initially reluctant to participate in the BES market or that there were some learning effects associated with participation in the BES market. Since the BES market began operation only in late 2001, by 2002 the PGCs had only been bidding in it for a few months, which suggests that Calpine may simply not have known how to bid.

The remainder of our estimates, which are summarized in table 3, show the other PGCs to fairly consistently submit markups and markdowns far in excess of what is ex post-optimal. One possible explanation of these bidding patterns is that they may be due to some form of anticompetitive or collusive behavior on the part of the PGCs. In the context of the BES, collusion would manifest itself with firms submitting high inc and low dec offers in an effort to raise the MCPE for incs and lower the MCPE for decs from what would result in a non-cooperative equilibrium. At first glance, the steep bids of the small PGCs may seem indicative of collusive behavior not accounted for in our model. Our findings and other analyses of the BES, however, point against this conclusion. As we noted earlier, the bulk of generating assets in the ERCOT control area are held by the three large PGCs. Furthermore, because the vast majority of electricity sales are contracted forward of the BES market, the small PGCs hold a relatively small percentage of balancing resources. Moreover, we find that on average the three large PGCs offer the bulk of balancing resources into the BES—often over 60%. These facts, coupled with TXU and Reliant submitting offers close to their EOSFs implies that a small PGC overpricing its incs and underpricing its decs would have little effect on the MCPE since the PGC would essentially price itself out of the BES market. In fact, the excessive markups and markdowns of the small PGCs is commonly explained by them wanting to exclude themselves from the market, which we study further in the following section. In her study of the BES, Niu shows that actual inc prices and those which would result from a linear SFE are on average within 3.9% of one another, indicating that even if the smaller PGCs are colluding to drive up the price of incremental balancing energy, these attempts have met with little to no success. Her analysis of dec prices, however, show the actual price to average twice her theoretical calculations. If the excessively high dec markdowns represent ramping costs or constraints, gas imbalance charges, or other factors not accounted for in these analyses, then the discrepancy in dec prices would represent a true cost of service. Otherwise, this could indicate some form of anticompetitive behavior for decremental energy.

The results of our ex post analysis are fairly consistent with Hortaçsu and Puller's findings. One of their metrics of rational bidding is to compare a firm's profits under optimal bidding to that under their actual offers, and calculate the percent of potential profits a firm foregoes with its actual offers. Their results show that amongst the three largest PGCs, Reliant performs best, realizing 79% of its potential profits in the BES. TXU and Calpine, by contrast, perform more poorly yielding only 39 and 37% of their potential profits. Their estimates show the smaller PGCs to consistently perform poorly, generally realizing only a small fraction of their potential BES profits. It is, however, worth noting that even in the best cases of Reliant and TXU, all our statistical tests reject the null hypothesis that $\hat{\phi}_{NP}(q) = 1$, meaning none of the generators are behaving in accordance with ex post-optimal behavior.

4.3 Market participation

Our estimates of the smaller PGCs' markup ratios show them to submit offer curves which are far steeper than our model predicts. The most common explanation for this observation is that smaller power producers, cooperatives, and cogenerators prefer to generate according to their contracted schedules. To help ensure this they overbid their generation into the balancing market with markups and markdowns significantly higher (in some extreme cases up to three orders of magnitude greater) than predicted by our model. Such behavior would act to keep them out of the balancing market, except when demand for balancing resources is sufficiently high to guarantee a high profit for being dispatched in the BES market. One reason behind this desire may be a real or perceived cognitive cost of solving a sophisticated optimization problem to make a bidding decision. Smaller generators may have so little 'money on the table' from participating in the BES, that there is little incentive to do so. Likewise, cogenerators are often only tangentially involved in the electricity market and may focus their efforts on more-lucrative bilateral sales as opposed to the BES. Furthermore, since cogeneration is a byproduct of their primary production process, they may have a strong disincentive to adjust their output in order to garner slim margins in the BES as doing so may affect their primary production. Similarly, municipalities may be primarily interested in generating and procuring resources for their native load, and as such may have less of a profit-motive to actively participate in the BES market.

This conjectured behavior is supported to some extent by the inc and dec patterns of the various PGCs. Table 4 shows that many of the smaller PGCs often opt to exclude themselves from the BES by submitting only dec offers, which they are required to, without any inc offers. Furthermore, as Table 5 shows, four of the small generators submitted only a single dec offer at the price floor of $-\$1,000$ in a number of bid periods. These observed patterns of bidding are indicative of many of the small PGCs participating in the BES only in so far as they are required to by the PUCT's dec offer requirement.

To concretely explore this relationship between dec offer patterns and the 'size' of a PGC, we estimate a binary response model. We define $y_{i,t}$ to be an

Table 4 Percentage of bidding periods with only DEC bids submitted

PGC	% DEC bid Only	PGC	% DEC bid only
Brazos Electric Power Cooperative	97.3	City of Garland	15.4
Mirant Wichita Falls Management	86.8	Air Liquide	11.9
Hays Energy	84.3	City of San Antonio Public Service	9.1
Midlothian Energy	71.6	BP Energy	7.6
Bryan Texas Utilities	71.4	Guadalupe Power Partners	7.5
Bastrop Energy Partners	53.5	Central Power and Light	7
Lamar Power Partners	41.7	Cogeneration Lyondell	2.3
Gregory Power Partners	39.8	Denton Municipal Electric	2.1
Rio Nogales Power Project	39.4	Ingleside	1.1
Tractabel Power	37.5	AES Deepwater	0.9
Brownsville Public Utility Board	33.3	City of Austin	0.7
South Texas Electric Cooperative	32	Lower Colorado River Authority	0.4
Sweeny Cogeneration General	26.4	Reliant Energy	0.2
West Texas Utilities	17.2	Extex Laporte	0.1
Tenaska Gateway Partners	17.1	TXU	0
Calpine	16.1		

Table 5 Percentage of bidding periods with only a single DEC bid at -\$1000

PGC	% -\$1000 DEC bid only
Bastrop Energy Partners	24.5
Brownsville Public Utility Board	23.7
Lamar Power Partners	16
Sweeny Cogeneration General	9.3

indicator variable, with $y_{i,t} = 1$ if PGC i submits only dec offers in period t and $y_{i,t} = 0$ otherwise. We then estimate the limited dependent variable model:

$$\text{Prob}\{y_{i,t} = 1|\mathbf{x}_{i,t}\} = F(\beta_0^\top \mathbf{x}_{i,t}), \tag{5}$$

where $\mathbf{x}_{i,t}$ is a vector of regressors, β_0 is the vector of parameters to be estimated, and $F(\cdot)$ is a cumulative distribution function. One of our regressors, profPGC_i , is an indicator variable for whether PGC i is what we designate a “profit-driven PGC.” These PGCs include the incumbent IOUs: TXU, Reliant, West Texas Utilities, and Central Power and Light; as well as the large independent power producer, Calpine. We include a regressor, $\%CAP_{i,t}$, the percent of available capacity that a PGC has committed day-ahead to account for capacity effects on bid patterns. We also include a set of regressors, $\%\Delta\text{SCHED}_{i,t}^\tau$, $\tau = \{-2, -1, 1, 2\}$, giving the percent change in a PGC’s scheduled production from bidding hour t to bidding period $t + \tau$, which are meant to account for binding ramping constraints preventing a PGC from offering inc service. We finally

Table 6 Estimates of LDV model

Regressor	$\hat{\beta}_{MLE}$	$\mathbb{E}_N[\hat{\beta}_{MLE} f(\hat{\beta}_{MLE}^T x)]$
WEEKDAY	-0.070676	-0.01314
ON-PEAK	0.14715	0.027359
%CAP	-1.6112	-0.29956
% Δ SCHED ₋₂	0.0095387	0.0017735
% Δ SCHED ₋₁	0.18526	0.034444
% Δ SCHED ₁	0.0067523	0.0012554
% Δ SCHED ₂	0.0056042	0.001042
profPGC	-0.19935	-0.037063

include indicator variables ON-PEAK and WEEKDAY to account for any difference in dec offer patterns between different time periods. We specify the distribution function in the LDV model in Eq. 5 to be logistic, and estimate it by maximum likelihood. Using a Wald test, we find all the estimated coefficients to be statistically significantly non-zero at the 1% level. Our estimates, which are summarized in Table 6, show that a profit-driven PGC is on average less likely to submit only dec offers. We further find that changes in scheduled production deter a PGC from offering inc offers, with the change from the previous hour having the highest impact. This suggests ramping constraints, which are ignored in our offer curve derivation, may play a role in explaining some bidding behavior. Finally, we find that having capacity committed reduces a PGC's probability to offer only dec offers. While seemingly counterintuitive, this suggests that a PGC with a higher proportion of its capacity scheduled to generate, has greater flexibility in ramping generation among a wider portfolio of operating units.

4.4 Learning effects

A question to ask in a complicated market such as the BES is whether participants can gradually learn to better bid against their opponents through their repeated interactions in the market. In many empirical and experimental settings, players have demonstrated the ability to 'converge' towards playing equilibrium strategies even when equilibria are rarely seen in a one-shot variant of the game. Similarly, one may expect that over time PGCs may learn to better conjecture their rivals' behavior and react to those beliefs. Indeed, this result bears itself out quite strikingly with the behavior of Calpine, which our estimates showed to be submitting offers with markups and markdowns up to 37-times its ex post-optimal response. The data shows that this is due in large part to exorbitantly high markups in the first quarter of 2002. Since by that point the BES market had only been in operation for a few months, this could indicate a period of learning on the part of Calpine. It has alternately been suggested that Calpine may have initially been reluctant to participate in the BES and was only willing to 'test the waters' in a limited fashion, and only after observing the market for a few months was willing to earnestly 'jump in.'

We demonstrate this learning phenomenon by estimating a partial-linear semiparametric variant of the behavioral model in Eq. 4:




Fig. 5 Semiparametric estimates of Calpine’s conduct function

$$\frac{MU_{i,t}(q)}{MU_{i,t}^*(q)} = \beta_i^\top \mathbf{Z}_{i,t} + \phi_i(q) + \eta_{i,t}(q), \tag{6}$$

where $\mathbf{Z}_{i,t}$ consists of indicator variables for the different time periods we compare bidding patterns across. As usual, we assume $\mathbb{E}[\eta_{i,t}|\mathbf{Z}, q] = 0$.

In analyzing Calpine’s learning effects, the matrix of linear regressors, \mathbf{Z} , consists of an indicator variable, Q1-Y2002, for bids submitted in quarter 1 of 2002, and another indicator variable, Q1, for bids submitted in quarter 1 of 2002 or 2003. The inclusion of an indicator for quarter 1 of either year is meant to control for any seasonal variation in bidding behavior. We estimate the model using the Speckman (1998) conditional moment method. Figure 5 shows the semiparametric estimates and gives the estimated values of $\hat{\beta}_{SP}$, both of which are shown to be statistically significantly non-zero at the 1% level by a standard Wald test. As can be seen, the first quarter of 2002 showed exceedingly high markups, followed by bidding patterns which more closely match our theoretical model thereafter. It is also evident that once controlling for high markups and markdowns during the first quarter of 2002, Calpine’s offers exhibit a pattern somewhat similar to that of Reliant with excessively marked-down dec offers and more reasonably priced inc offers.

5 Comparison of actual offers to Nash equilibrium supply functions

Finally, we wish to address the question of whether bids into the BES conform to a Nash equilibrium set of supply functions. While related to the analysis of the EOSFs, this allows us to compare bids to a full equilibrium as opposed to

the partial equilibrium analysis we have conducted thus far. Assuming that an SFE model truly describes the behavior of the generating firms, PGCs should submit supply functions with the Nash property that no generator can profitably unilaterally deviate.

Besides the computational complexity of the problem, one of the difficulties in finding Nash equilibrium supply functions is the multiplicity of equilibria. This multiplicity arises because if firm i 's rivals are submitting elastic supply functions then firm i should also submit a more elastic supply function, otherwise it would price itself out of the market. Similarly, if firm i 's rivals are submitting inelastic supply functions then firm i should submit a more inelastic supply function. As such, there has been an extensive theoretical literature trying to overcome the non-uniqueness of supply function equilibria by imposing further assumptions and structure on the model.

5.1 Methods of finding unique supply function equilibria

Rudkevich (1999) and Baldick, Grant and Kahn (2000, 2004) study linear supply functions. Rudkevich derives optimality conditions giving a Nash equilibrium for a supply function when marginal costs and market demand are linear, firms are restricted to bidding linear supply functions, and the support of the demand shock has at least two distinct points. He shows these conditions to yield a unique solution and proposes using an iterative myopic best response algorithm to solve for the equilibrium. Baldick et al. give conditions under which this myopic best response algorithm is a contraction mapping—implying the technique will converge from any starting point and that the unique equilibrium is stable. They also show that while there do exist a multitude of non-linear equilibria (which could arise if firms are not restricted to submitting linear supply functions), if a firm's rivals all bid linear supply functions it is then optimal for the firm to bid a linear supply function. Although the linear SFE model has these attractive properties, we find that in our case it falls short of capturing some complexities of the market. For one, most of our marginal cost estimates exhibit non-linearities, which we would like to be able to account for in our derivation. Second, the linear supply function model assumes that the generating firms have no capacity constraints, meaning that the model cannot take account of the relative 'size' of the PGCs. From running some sample instances we found the linear SFE to be very close to competitive marginal-cost bidding—often with markups of less than \$1 over marginal cost. The reason the linear model yields such competitive outcomes is exactly due to the lack of capacity constraints; because the model treats the multitude of smaller generators as having unbounded capacity, the equilibrium solution turns out to be very competitive. Niu overcomes this shortcoming of the linear model by only considering the four largest PGCs and 'lumping' the remaining the PGCs into a fifth 'fringe' bidder. Thus the linear model she uses sees only five capacity-unconstrained PGCs and yields bids with more sensible markups. The final shortcoming of the linear model is that it does not allow PGCs to 'hockey-stick' their bids. We

mentioned that this behavior is prevalent in the offer curves of most all the PGCs, and we would like to capture such non-linear supply functions in our model.

In their original paper, Klemperer and Meyer show that if the demand shock, ϵ , has an unbounded support, then there will be a unique set of supply functions solving the governing differential equations and satisfying the second-order optimality conditions. Although this assumption has the attractive advantage of allowing for a unique equilibrium with general-form supply functions, the unbounded support assumption is rather tenuous in a balancing market which normally dispatches less than 5% of total electricity traded in ERCOT.

A more recent line of theoretical research in SFE has explored the use of capacity constraints and price caps, which give an additional set of boundary conditions and yield unique equilibria. In a series of papers, Holmberg (2004, 2005a,b) explores this technique with symmetric and asymmetric firms. His method essentially amounts to assuming that the capacity constraints of the generators will be binding at the price cap and solves for an equilibrium by integrating an expanding set of coupled differential equations backwards. This requires an assumption that the support of the demand shock is sufficiently high so as to exhaust the generating capacities of the bidders with some positive probability. It also implicitly assumes that the bidders will want to 'hockey-stick' their bids so as to reach the price cap when their capacity constraints are binding. Although we find the assumptions of this model to fit the realities of the BES market well, the solution technique of solving a series of stiff differential equations backwards is too computationally complex for us to feasibly run the model on our sample of 8,160 bidding periods. Anderson and Hu (2005) develop a unique SFE along the same lines, in which all but one firm reaches its capacity limit prior to the price cap. Moreover, their model allows for general marginal cost and demand functions, and they provide a tractable means of finding this equilibrium by formulating the problem as a complementarity problem. Due to the soundness of their model's assumptions and the relative simplicity of finding equilibria, we use Anderson and Hu's technique.

Although these supply function models can be applied to derive both Nash equilibrium inc and dec offers, we opt to analyze only the inc part. The reason we make this restriction is because of our earlier finding in analyzing the EOSFs that most PGCs (including Reliant and Calpine) bid their dec service with markdowns which are excessively high and cannot be explained by our optimal response model. Furthermore, we restrict our equilibrium analysis to the three largest PGCs: TXU, Reliant, and Calpine. We make this restriction both to make the problem tractable and because of our earlier evidence suggesting that the smaller PGCs are not rational strategic players in the BES market. The comparison of actual bids to EOSFs showed the smaller PGCs to be submitting bids far in excess of what our model predicts, which acts to price their generation out of the BES market. If we include these firms as strategic players, our model would essentially conjecture that these firms are making rational offers which would make each firm's residual demand more elastic, thereby making the theoretical equilibria more competitive. Because our analysis shows these smaller

firms not to be acting in accordance with rational profit-maximizing behavior, we opt to hold their bids as fixed at their actual offer curves in constructing the market demand function, $D(p)$, and the residual demand functions for the three strategic PGCs. In making these assumptions, we are essentially conjecturing that the strategic PGCs account for the non-strategic behavior of the smaller PGCs in estimating their residual demand function and calculating their supply function response to their strategic and non-strategic counterparts.

5.2 Derivation of capacity-constrained Nash equilibrium supply functions⁸

The basic model has n capacity-constrained firms. Let $c_i(q_i)$ be firm i 's total cost function, which we assume to be convex, differentiable, and non-negative, and let \bar{q}_i be firm i 's generating capacity. We will assume that each firm has a different initial marginal cost (which is true in our data set) and that they have been numbered such that $c'_1(0) < c'_2(0) < \dots < c'_n(0)$. Again, we assume that demand is given by the function $D(p, \epsilon) = D(p) + \epsilon$, in which $D(p)$ is a deterministic function which is strictly decreasing, differentiable, and concave. The stochastic portion of the demand, ϵ , is again assumed to have support $\epsilon \in [\epsilon_{\min}, \epsilon_{\max}]$, and we suppose that ϵ is distributed according to density function, $f(\epsilon)$, which is strictly positive everywhere on this support. Finally, we assume that there is a price cap on the market, \bar{p} , and that prices are always non-negative (which they will be for inc energy). Although Anderson and Hu's model does not directly account for scheduled sales, we can incorporate them into the model by defining each firm's generation quantity, q_i , to be $q_i = q_i^{\text{tot}} - Q_i^{DA}$, where q_i^{tot} is firm i 's total generation. Thus, q_i simply measures the amount of incremental energy sold above any sales contracted outside the BES, and the contracted quantity, Q_i^{DA} , simply shifts the 'zero point' of the marginal cost and supply functions.

We now state Anderson and Hu's main uniqueness result, the details and proof of which are in their paper.

Theorem 5.1 *If*

$$-D(c'_1(0)) < \epsilon_{\min} < -D(c'_2(0)),$$

then any supply function equilibrium is part of an ordered family, and only the lowest (smallest offered quantity at any given price) can have the property that all but one of the firms reach their capacity limits prior to the maximum price.

The condition of the theorem simply means that at price $c'_1(0)$ there is always some demand even when $\epsilon = \epsilon_{\min}$, but when the price reaches $c'_2(0)$ there may not be. Another interpretation of the assumption is that at the minimum shock there is a single economic supplier (*i.e.* the supplier with the lowest initial marginal cost).

⁸ Because Anderson and Hu's paper details their model, we only mention the underlying assumptions, state their main uniqueness result, and discuss the technique used to solve for the equilibrium. Interested readers should consult their paper for complete details.

Anderson and Hu’s technique to solve for an equilibrium is based on discretizing the demand shock and approximating each firm’s supply function as being piecewise linear. Given a fixed positive integer, K , they assume the demand shock can now take one of K discrete values $\epsilon \in \{\epsilon_1, \epsilon_2, \dots, \epsilon_K\}$, with $\epsilon_{min} \leq \epsilon_1 \leq \epsilon_2 \leq \dots \leq \epsilon_K \leq \epsilon_{max}$. They assume that when $\epsilon = \epsilon_k$ the MCPE will be p_k , where $\sum_i s_i(p_k) = D(p_k) + \epsilon_k$, each firm’s generation will be $q_{i,k} = s_i(p_k)$, and its supply function will have slope $\beta_{i,k} = s'_i(p_k) \geq 0$. Next, they define for each firm a set of points $\tilde{p}_{i,1}, \tilde{p}_{i,2}, \dots, \tilde{p}_{i,K-1}$ at which its piecewise linear supply function kinks. Thus, we can write each firm’s supply function as:

$$s_i(p) = \begin{cases} q_{i,1} + \beta_{i,1}(p - p_{i,1}), & 0 \leq p \leq \tilde{p}_{i,1} \\ q_{i,k} + \beta_{i,k}(p - p_{i,k}), & \tilde{p}_{i,k-1} \leq p \leq \tilde{p}_{i,k}, \quad k = 2, \dots, K - 1 \\ q_{i,K} + \beta_{i,K}(p - p_{i,K}), & \tilde{p}_{i,K-1} \leq p \leq \bar{p}; \end{cases}$$

where the $p_{i,k}$ is the price that firm i conjectures for when $\epsilon = \epsilon_k$. With this piecewise linear form, we can now write firm i ’s profit maximization problem as:

$$\begin{aligned} \max_{p_{i,k}} & \left[D(p_{i,k}) + \epsilon_k - \sum_{j \neq i} s_j(p_{i,k}) \right] p_{i,k} - c_i(D(p_{i,k}) + \epsilon_k - \sum_{j \neq i} s_j(p_{i,k})) \\ \text{s.t.} & \quad 0 \leq p_{i,k} \leq p \\ & \quad 0 \leq D(p_{i,k}) + \epsilon_k - \sum_{j \neq i} s_j(p_{i,k}) \leq \bar{q}_i, \end{aligned}$$

which gives a set of FONC for an optimum. When the FONC from all the firms’ profit-maximization problems are assembled, and we impose the equilibrium condition, $p_{1,k} = p_{2,k} = \dots = p_{n,k} = p_k \forall k$, that the firms all conjecture the same price under each demand shock realization, we arrive at the following set of equilibrium conditions:

$$\begin{aligned} q_{i,k} - (p_k - c'_i(q_{i,k}))(\sum_{j \neq i} \beta_{j,k} - D'(p_k)) + \lambda_{i,k} - \mu_{i,k} &= 0 & \forall i, k \\ \sum_i q_{i,k} &= D(p_k) + \epsilon_k & \forall k \\ 0 \leq p_k &\leq \bar{p} & \forall k \\ q_{i,k+1} - q_{i,k} + \beta_{i,k+1}(\tilde{p}_{i,k} - p_{k+1}) + \beta_{i,k}(p_k - \tilde{p}_{i,k}) &= 0 & \forall i, k=1, \dots, K-1 \\ p_k < \tilde{p}_{i,k} < p_{k+1} & & \forall i, k \\ \beta_{i,k} &\geq 0 & \forall i, k \\ 0 \leq q_{i,k} \perp \mu_{i,k} &\geq 0 & \forall i, k \\ q_{i,k} \leq \bar{q}_i \perp \lambda_{i,k} &\geq 0 & \forall i, k, \end{aligned} \tag{7}$$

where $\mu_{i,k}$ and $\lambda_{i,k}$ are Lagrange multipliers on the lower- and upper-bound capacity constraints, respectively. Anderson and Hu then prove the following theorem, which we restate, showing that for K sufficiently large the piecewise linear functions will well approximate the actual equilibrium.

Theorem 5.2 *Let $\{s_i^*(p)\}_{i=1}^n$ be a supply function equilibrium on $[0, \bar{p}]$. Then, for K large enough, there exists a solution $\epsilon_k, p_k, q_{i,k}, \tilde{p}_{i,k}, \lambda_{i,k}, \mu_{i,k}$ to the equilibrium*

conditions (7), such that $D(p_k) + \epsilon_k = \sum_i s_i^*(p_k)$, and $q_{i,k} = s_i(p_k)$. Moreover $\beta_{i,k} = s_i^*(p_k)$ for all but a finite number of k .

In order to solve for an equilibrium, Anderson and Hu formulate a mathematical program with equilibrium constraints (MPEC) with the equilibrium conditions (7) as constraints. The difficulty in solving such an MPEC is that the complementarity conditions between the capacity constraints and their Lagrange multipliers do not satisfy constraint qualification conditions, making the problem difficult for standard solvers. They suggest overcoming this issue by choosing $\rho \geq 0$ and relaxing the complementarity conditions so that:

$$\begin{aligned} q_{i,k} \mu_{i,k} &\leq \rho \\ (\bar{q}_i - q_{i,k}) \lambda_{i,k} &\leq \rho. \end{aligned}$$

One could then solve the MPEC by first solving the problem with a large starting value for ρ , and iterating by reducing ρ and resolving the relaxed MPEC at each step until reaching a sufficiently small final ρ .

Our implementation follows the Anderson and Hu method. For each bidding period we fit quadratic marginal cost and demand functions to the actual data by method of least-squares. The capacity constraints, \bar{q}_i , are set based on the actual quantity each PGC offers into the BES in that bidding period, as opposed to using the total nameplate capacity of the generating units available to that PGC. We determine generating capacities from the actual bids because our nameplate estimates do not account for resources being held for self-scheduled reserves, ramping constraints on the amount of energy available, and other physical limitations. Moreover, if a PGC has excess capacity available there is no rationale for physically withholding those resources from the BES, since it can easily economically withhold the generation by submitting it with an excessively high offer. We set the lower-bound of the support of the demand shock, ϵ_{\min} , so as to satisfy the assumptions of Theorem 5.1. Although Anderson and Hu's simulations worked well using the CONOPT optimization package, the parameters of our fitted functions were very poorly scaled and CONOPT could almost never find an initial feasible solution, even with our large starting value of ρ . As such, we opted to use the filterSQP solver [details of the algorithm are available in Filters and Leyffer (1999)] on the NEOS optimization server [details of which are available in Czyzy, Mesnier and More (1998)], which performed much better, and attempted to solve the MPEC with a final ρ less than 10^{-9} . Though the equilibrium problem for each bidding period is feasible, the solver had difficulty with some instances in which the function parameters and capacity constraints were particularly poorly scaled. Out of a total of 8,160 bidding periods, 4,297 converged with $K = 10$, another 1,781 with $K = 5$, and 1,518 with $K = 3$, leaving 564 instances which could not converge. In all, 7,596 instances or 93% of the total sample solved. Our econometric analysis includes only those 7,596 instances which converged. Qualitative inspection of the cases we could not solve do not reveal any pattern that is likely to bias our results.

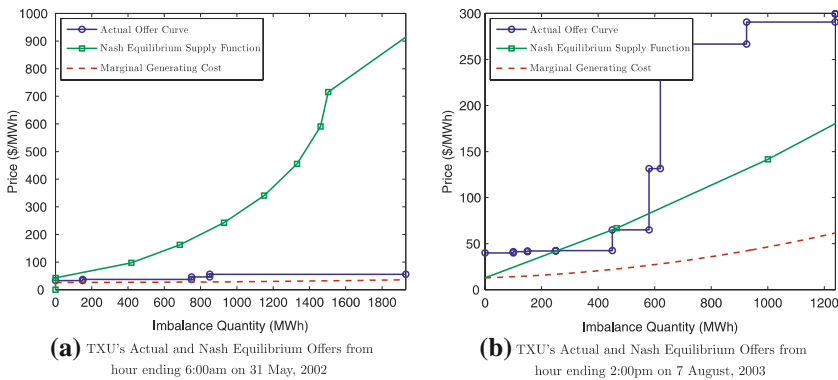


Fig. 6 Actual and Nash equilibrium offers example

5.3 Comparison of actual offers to nash equilibrium supply functions

In order to compare the actual offers to the Nash equilibrium supply functions, we will again estimate the same behavioral model used in analyzing the EOSFs, with specific reference to Eq. 4. Because we now wish to test the extent to which the offers of the three strategic PGCs match our derived theoretical optima, we will estimate our behavioral model both using the entire cross-section of offers from the three PGCs and for each PGC individually. Figure. 6 shows two sample actual and Nash equilibrium offer curves for TXU from two different bidding periods. In the left pane, we see that TXU’s actual offers are relatively close to its marginal generating costs, whereas its Nash equilibrium supply function requires higher offer prices. The derived equilibrium also shows the non-linear shape of the Nash equilibrium supply function, which we would not be able to capture in a linear SFE model. In the right pane, the actual and equilibrium offers match closely for imbalance quantities less than 600MW, but TXU prices its inc offers above 600 MW at higher than the equilibrium calls for. Figures 7–10 show our estimated $\hat{\phi}_{NP}(q)$ and the 95% confidence interval. The second pane in the figures shows the same plot with the axes truncated to show the estimates with greater granularity. Moreover, Table 7 summarizes the estimate and gives upper- and lower-confidence interval bounds. Again, the $\hat{\phi}_{NP}(q)$ estimator spikes at $q \approx 0$ because of the extreme sensitivity of the ϕ multiplier to errors when the imbalance quantity is close to zero.

As the figures show, while the $\hat{\phi}_{NP}(q)$ estimator varies considerably across the supply stack, the $\phi = 1$ line lies within the confidence interval bounds in all four estimates. As such, we are unable to reject the null hypothesis that $\hat{\phi}_{NP}(q) = 1$. In fact the plots with the truncated axes show that while the $\hat{\phi}_{NP}(q)$ estimator is not identically one, it is nonetheless within a close neighborhood of one. Moreover, in comparison to the estimates of $\hat{\phi}_{NP}(q)$ from the analysis of the EOSFs of the smaller PGCs, we see that the behavior of the three major PGCs are better predicted by a Nash supply function equilibrium than by a EOSF. Given the large magnitude of the $\hat{\phi}_{NP}(q)$ estimator for low values of

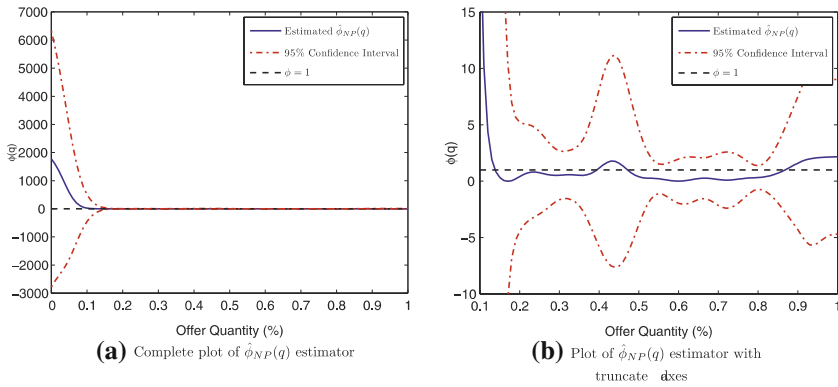


Fig. 7 Estimated industry conduct function

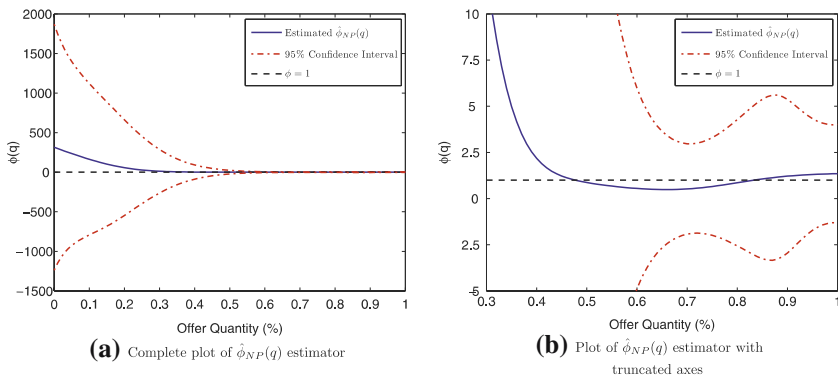


Fig. 8 TXU's estimated conduct function

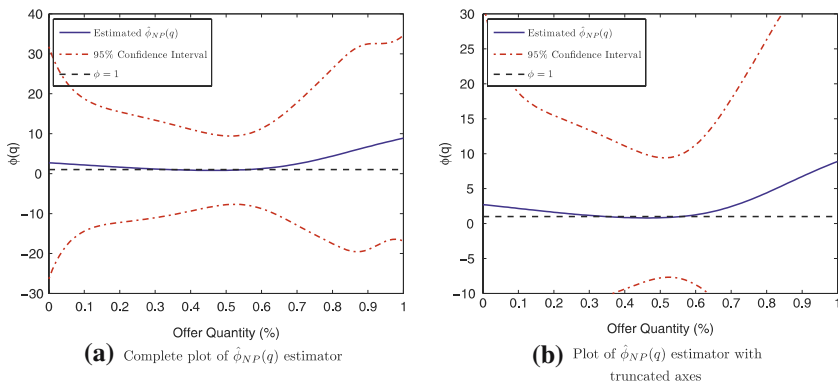


Fig. 9 Reliant's estimated conduct function

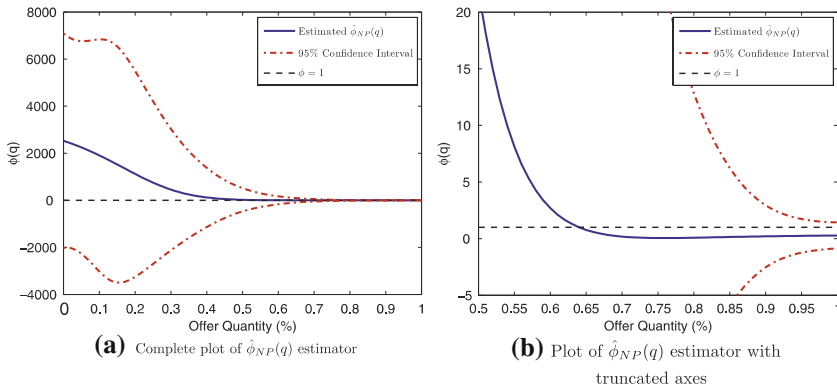


Fig. 10 Calpine's estimated conduct function

Table 7 Summary statistics of conduct functions

q	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
Overall											
$\hat{\phi}_{NP}(q)$	1775.9	23.858	0.402	0.525	1.155	0.389	0.000	0.146	0.324	1.594	2.166
$\hat{U}_{NP}(q)$	6333.9	499.4	5.228	2.691	8.129	4.630	1.962	2.441	1.370	7.49	9.034
$\hat{L}_{NP}(q)$	-2782.1	-451.68	-4.425	-1.640	-5.819	-3.851	-1.961	-2.149	-0.722	-4.302	-4.702
TXU											
$\hat{\phi}_{NP}(q)$	315.58	163.95	57.755	12.313	2.178	0.871	0.549	0.524	0.859	1.216	1.352
$\hat{U}_{NP}(q)$	1870.8	1120.2	661.96	287.19	91.994	23.956	5.981	2.971	4.279	5.374	4.014
$\hat{L}_{NP}(q)$	-1239.6	-792.32	-546.45	-262.56	-87.638	-22.214	-4.882	-1.924	-2.562	-2.943	-1.310
Reliant											
$\hat{\phi}_{NP}(q)$	2.714	2.167	1.623	1.164	0.869	0.833	1.259	2.435	4.400	6.728	8.885
$\hat{U}_{NP}(q)$	31.757	18.738	15.465	13.375	11.102	9.442	11.273	17.796	26.321	32.48	34.647
$\hat{L}_{NP}(q)$	-26.329	-14.404	-12.219	-11.047	-9.363	-7.776	-8.756	-12.927	-17.521	-19.024	-16.877
Calpine											
$\hat{\phi}_{NP}(q)$	2531.7	1908.4	1122.3	455.01	121.47	21.997	2.696	0.180	0.084	0.204	0.277
$\hat{U}_{NP}(q)$	7086.9	6840.5	5498.8	3032.9	1377.7	516.3	171.41	49.213	12.85	2.931	1.429
$\hat{L}_{NP}(q)$	-2023.5	-3023.6	-3254.3	-2122.9	-1134.8	-472.3	-166.02	-48.852	-12.681	-2.522	-0.875

q and the corresponding wide variance bands, however, the predictions of the SFE model are highly sensitive to miscalculations on the part of the PGCs or the economist applying the theory. Nonetheless, other spot market analyses, such as Niu's, suggest that precisely pinpointing the bidding behavior of the firms is not crucial in determining price outcomes of the market. Although there are efficiency and profit consequences from irrational behavior of the smaller PGCs, if one's primary concern is the price of balancing energy, then the SFE model performs admirably with regard to incremental energy – with Niu's estimates showing actual inc prices to on average be within 3.9% of linear SFE predictions. Moreover, the fact that she restricted herself to a linear SFE model suggests that non-linear supply functions may not even be necessary to predict market price outcomes. The reason a linear model may be sufficient is that the

BES will rarely exhaust the supply of the PGCs, as such the ‘hockey-sticked’ portion of the supply stacks rarely set the MCPE. Although there tends to be a large disparity between predicted and actual dec prices, if the excessively low dec bids are due to operating costs and constraints not captured in our data, then this is not necessarily a failing of the model.

6 Conclusion

The results of our analysis show that most bidders participating in the ERCOT BES do not act in accordance with what is predicted by an optimal response model. This is characterized by them making offers with markups and mark-downs which are far in excess of that implied by our theoretically ex post-optimal supply functions. With smaller power producers, municipalities, cooperatives, and cogenerators, we have explained this as a general reluctance to participate in the BES. Amongst the large PGCs, TXU and Reliant’s behavior matched our theoretical optima most closely. Calpine, the other large PGC, was found to be bidding with markups far greater than our model predicted, although we have shown this to be due to exceedingly high markups in the first quarter of 2002—when we control for the excessively high offers in this period, Calpine’s offer behavior more closely matches that of Reliant.

When we incorporated capacity constraints into the model and restricted attention to the three strategic PGCs, the actual offers were within a ballpark of the Nash equilibrium supply functions. Moreover our estimates showed that we could not reject the null hypothesis that the actual observed behavior was not statistically significantly different from Nash. Thus our results suggest that an SFE model can well describe offer behavior in electricity spot markets, although it is highly dependent on judicious application to firms we believe to be bidding strategically. Though the model does not exactly predict the behavior of bidders, the fact that the three dominant players tend to set the margin and bid in accordance with the model gives it a fair amount of power in predicting price outcomes. However, SFE is considered an attractive model of spot electricity markets in part because it assumes a strategy space and firm behavior which is reminiscent of the actual price/quantity offers submitted by generators. Yet in a sense, our analysis suggests the one attractive point of SFE—its behavioral assumptions and predictions—tends to not be in tune with the individual behavior patterns of a large segment of the market. This begs the question whether complicated general-form SFE models, such as the one we employed, may simply be an overly complex model of the market, which gives no better behavioral or price predictions than simpler models. For instance, linear SFE models have been shown to give fairly accurate price predictions even though they cannot capture complicated hockey-stick type bidding. Even more conventional Cournot-type models have been applied to studying markets. Moreover, many market power studies and models for assessing the economic impacts of transmission constraints and other nuances of power systems rely on simpler Cournot modeling. Although Cournot models cannot incorporate the

institutional details which SFE models can, this study suggests that at least with respect to determining equilibrium prices, the results of Nash equilibrium models are quite robust and it may be sufficient to use Cournot models which allow inclusion of transmission congestion and other details of power systems. Cournot models do have the difficulty, however, that electricity balancing markets have little or no demand elasticity. As such, Cournot models usually require estimating a residual demand by assuming there is a competitive fringe. Our analysis suggests, however, that the fringe suppliers in the market behave far from competitively, and that their behavior can be the most difficult to predict.

Our analysis also highlights some of the issues arising in the oversight of a market by regulators and market monitors. With the behavior of the smaller PGCs, we observed these firms submitting offers with markups and markdowns far in excess of what our supply function model suggests would be profit-maximizing in a non-cooperative equilibrium. Observation of such behavior may lead us to conclude these firms are exercising market and acting to raise the MCPE, however a more thorough analysis shows these firms' residual demand is far too elastic and as such these bids merely price the firms out of the market. This demonstrates that in undertaking market mitigation, regulators and market monitors should not be focused solely on behavior (for example, observing high markups), but also on outcomes (for example, whether high markups drive up the MCPE). This type of restraint applies not only to simple 'eyeballing' of market data for high offers, but also to use of sophisticated equilibrium models such as the one we studied. We can also this type of analysis to draw conclusions about how market mitigation measures should be targeted. A 'shame cap' which is meant to punish any firm with an offer above a certain threshold may be targeting the wrong culprit if the high offer is being used by the firm to ensure it does not have to deviate from its schedule. As such, it may be more sensible to only target high bids which set (or are close to setting) the margin. However, depending on the impact the observed economic withholding has on efficiency of the market, regulators may wish to take steps to facilitate entry into the market by these firms, which may be acting to exclude themselves. Hortaçsu and Puller's analysis of the BES shows these efficiency losses to be sizeable, suggesting a need by the regulator to adjust the market design. By contrast, the restraint showed by TXU in its BES offers suggests that giving PGCs native load obligations with a fixed or capped retail rate may help keep their behavior in the balancing market in check. We may even go further, concluding that LSEs should be encouraged to retain at least a portion of their generating assets, and that excessive forced divestitures may increase PGCs' incentives to exercise market power. On the flip side, we were able to show the model's predictions are fairly in line with the behavior of the larger PGCs, suggesting the type of equilibrium analysis we've undertaken has some value in analyzing policy decisions. For example, in a market with a handful of large strategic firms and a small fringe, these types of models may prove to be one useful tool in analysis of mergers, divestitures, or alternate market designs. It is important to bear in mind, however, that these models can result in vastly different conclusions depending on nuanced modeling assumptions.

Finally, we draw some simple market design conclusions. We showed the smaller PGCs submit non-optimal offers, instead relying primarily on the bilateral markets, and larger PGCs submit non-optimal dec offers. If the behavior of the smaller PGCs is due to some actual or perceived cost of transacting in the balancing market, for instance from having to determine offer prices in a complicated optimization problem, this suggests efficient trade may benefit from creating voluntary liquid day-ahead and real-time markets such as those in PJM which facilitate trading and reduce transaction costs. The data in our study is taken from the current ERCOT market, which by design tries to rely primarily on bilateral markets. It is hoped that the new nodal design which includes a voluntary two-settlement market will improve liquidity and facilitate optimal behavior by participants.

Acknowledgements The work described in this paper was supported in part by NSF Grant EEC0119301 and the Consortium For Electric Reliability Technology Solutions (CERTS) through a grant from the Department of Energy. In addition, the first author would like to thank the US Department of Energy's Energy Information Administration, which provided invaluable support through the EIA Research Fellows Program, conducted in conjunction with the American Statistical Association. The Research Program is designed to produce theoretical and empirical statistical and economic analyses that serve to improve the statistical programs of the EIA. The opinions and conclusions expressed in this paper are those of the authors and do not necessarily represent those of the EIA. In addition to the generous support, we would like to thank Ross Baldick, Richard Gilbert, Steve Puller, Armin Sorooshian, Ruey-Pyng Lu, and seminar participants at UCEI for discussion and suggestions. We would also like to thank Tony Grasso, Ted Hailu, and Stan Kaplan, for providing invaluable assistance with data. Any errors are entirely our own.

References

- Anderson, E. J., & Hu, X. (2005). finding supply function equilibria with asymmetric firms. Working Paper.
- Baldick, R., Grant, R., & Kahn, E. (2000). Linear supply function equilibrium: generalizations, application, and limitations. UCEI POWER Working Paper PWP-078.
- Baldick, R., Grant, R., & Kahn, E. (2004). Theory and application of linear supply function equilibrium in electricity markets. *Journal of Regulatory Economics*, 25, 143–167.
- Baldick, R., & Hogan, W. (2002). capacity constrained supply function equilibrium models for electricity markets: stability, non-decreasing constraints, and function space iterations. UCEI POWER Working Paper PWP-089.
- Czyzyk, J., Mesnier, M. P., & More, J. J. (1998). The NEOS server. *IEEE Journal on Computational Science and Engineering*, 5, 68–75.
- Fletcher, R., & Leyffer, S. (1999). User Manual for filterSQP. University of Dundee.
- Green, R. (1996). Increasing competition in the british electricity spot market. *The Journal of Industrial Economics*, 44, 205–216.
- Green, R., & Newbery, D. (1992). Competition in the British electricity spot market. *The Journal of Political Economy*, 100, 929–953.
- Holmberg, P. (2004). Unique supply function equilibrium with capacity constraints. Uppsala Universitet Department of Economics Working Paper 2004:20.
- Holmberg, P. (2005a). Asymmetric supply function equilibrium with constant marginal costs. Preliminary Version.
- Holmberg, P. (2005b). Numerical calculation of an asymmetric supply function equilibrium with capacity constraints. Uppsala Universitet Department of Economics Working Paper 2005:12.
- Hortaçsu, A., & Puller, S. L. (2005). testing models of strategic bidding in auctions: a case study of the texas electricity spot market. UCEI CSEM Working Paper 125.

- Hurlbut, D., Rogas K., & Oren, S. S. (2004). Protecting the market from Hockey Stick pricing: how the public utility commission of texas is dealing with potential price gouging. *The Electricity Journal*, 17, 26–33.
- Kim, D.-W., & Knittel, C. (2004). Biases in static oligopoly models? Evidence from the California electricity market. UCEI CSEM Working Paper 131.
- Klemperer, P. D., & Meyer, M. A. (1989). supply function equilibria in oligopoly under uncertainty. *Econometrica*, 56, 1243–1277.
- Nadaraya, E. A. (1964). On estimating regression. *Theory of Probability and Its Applications*, 9, 141–142.
- Newbery, D. (1998). Competition, contracts, and entry in the electricity spot market, *The RAND Journal of Economics*, 29, 726–749.
- Niu, H. (2005). Models for electricity market efficiency and bidding strategy analysis. Ph.D. dissertation. Electrical and Computer Engineering Department, The University of Austin at Texas.
- Rudkevich, A. (1999). Supply function equilibrium in power markets: Learning all the way. TCA Technical Report Number 1299-1702, Tabors Caramanis and Associates.
- Speckman, P. E. (1998). Kernel smoothing in partial linear models. *Journal of the Royal Statistical Society, Series B (Methodological)*, 50, 413–436.
- von der Fehr, N.H. M., & Harbord, D. (1993). Spot market competition in the uk electricity industry. *The Economic Journal*, 103, 531–546.
- Watson, G. S. (1964). Smooth regression analysis. *Sankhya—The Indian Journal of Statistics, Series A*, 26, 359–372.
- Wilson, R. (1979). Auctions of shares. *Quarterly Journal of Economics*, 93, 675–689.
- Wolak, F. A. (2003). Measuring unilateral market power in wholesale electricity markets: The California Market 1998 to 2000. *American Economic Review*, 93, 426–430.
- Wolfram, C. (1999). Measuring duopoly power in the british electricity spot market. *American Economic Review*, 89, 805–826.