Online Appendix: Not For Publication

Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets

Ali Hortaçsu Fernando Luco
Steven L. Puller Dongni Zhu

A Profitability of Actual Bids

In this appendix, we report metrics of 'money left-on-the-table' across firms during the first 1.5 years of the market, as in Hortacsu and Puller (2008). Given our data on costs and bids, for each firm in each auction we calculate producer surplus under two scenarios: (1) best-response bidding and (2) bidding vertically at the contract position which is essentially not participating in the balancing market except to meet contract obligations. Then we calculate realized producer surplus under actual bidding, and we compute the fraction of potential profits relative to non-participation that were achieved by actual bidding. Table A.1 reports the average percent of potential profits realized by different firms across the first 1.5 years of the market.[1] A large firm Reliant – which is the large firm depicted in panel (a) of Figure 4 (a) – realized 79% of realized profits. However, all of the other firms realized less than one half of potential profits. The firm-level profits 'on-the-table' average between $1000 and $4000 each hour.

[1]In this table we report profitability only for the 12 firms that we will model in section V.

Table A.1: Firm-Level Profitability in First 1.5 Years of Market

Firm	Percent of Potential Profits Achieved
Reliant	79%
City of Bryan	45%
Tenaska Gateway Partners	41%
TXU	39%
Calpine Corp	37%
Cogen Lyondell Inc	16%
Lamar Power Partners	15%
City of Garland	13%
West Texas Utilities	8%
Central Power and Light	8%
Guadalupe Power Partners	6%
Tenaska Frontier Partners	5%

This reports the percent of potential profits that are achieved with actual bids relative to a benchmark where firms do not participate in the balancing market, i.e. bid vertically at the contract position. The figures represent the firm-level profitability averaged over auctions in the first 1.5 years of the market. Source: Hortacsu and Puller (2008).

B Additional Evidence on Learning

In this appendix, we show descriptive evidence that firm patterns of offering small quantities into the balancing market via bids that are 'too steep' is a phenomenon equally prevalent in the first and second year of the market. We test whether firms offer more generation capacity into the market in the second versus the first year of the market, or whether bidding that is 'too steep' is equally present in both years. Specifically, for each firm-auction, we calculate the amount of generation capacity that the firm offers relative to the contract position at the market-clearing price. We call this variable *Participation Quantity*. We define *Participation Quantity*$_{it}$=$|(S_{it}(p_t^{mcp}) - QC_{it})|$, using absolute value to capture bidding behavior for quantities above and below the contract position. A firm bidding vertically at the contract position is measured as *Participation Quantity*=0, but firms bidding with more elasticity have positive measures of *Participation Quantity*. We test whether firms offer additional generation into the market in the second year. Results are shown in Table B.1. In all specifications we include firm fixed effects so that we can test if firms participate more in the second year of the market relative to their participation in the first year. Column (1) shows that firms offer *less* generation in the second year, however the point estimate of *Year 2* (-35 MW) is neither economically nor statistically significant. Column (2) conditions on whether balancing demand is positive, and the point estimate is even smaller and not statistically different from zero. Column (3) conditions on the day of week and yields nearly an identical estimate. Finally, column (4) estimates the relationship for only the small firms and finds that these firms offer a very small amount of additional capacity in the second year – 1.52 MW – and this is not statistically different from zero.

This persistence of small quantities offered into the market suggests that learning is slow in this market. Formal tests of learning are reported in section V.

Table B.1: Offered Quantities into Market in Year 2 vs Year 1

	All Firms (1)	All Firms (2)	All Firms (3)	Small Firms (4)
Year 2	-34.76	-15.85	-16.15	1.52
	(42.42)	(34.24)	(34.70)	(2.90)
Firm Fixed Effects	Yes	Yes	Yes	Yes
INC Fixed Effects	No	Yes	Yes	Yes
Day of Week Fixed Effects	No	No	Yes	Yes
Observations	2264	2264	2264	1029
R^2	0.01	0.03	0.04	0.09

The dependent variable *Participation Quantity*$_{it}$ is the megawatt quantity of output bid at the market-clearing price relative to the firm's contract position in auction t, i.e. $|S_{it}(p^{mcp}) - QC_{it}|$. The sample period is the first 1.5 years of the market and *Year 2* is a dummy variable for the second year. Standard errors clustered at the firm-level are reported in parentheses.

C Evidence that Bidding Rules Do Not Bias Best-Response Bids as a Benchmark for Expected Profit Maximization

In this appendix, we test if firms would increase profits by following a simple trading rule that follows all bidding rules and places no restrictions on how uncertainty affects residual demand. The trading rule takes advantage of an institutional feature of the Texas market – the grid operator publicly released the aggregate bid schedule with a 2 day lag; therefore firms can learn their rivals' aggregate bid function with a 2 day lag. Suppose firms were to use the lagged bid data to create best-response bid functions to rivals' bids from 3 days prior to each auction, and submit these bids to the current auction. We compute lagged best-response bids and call these "naive best response" bids. Then we use the naive best-response bids and clear the market with the actual (step function) residual demand for the current auction. We find that this simple trading rule firm significantly outperforms the actual realized profits for all but the largest firm.

The results of this test are shown in Figure C.1. For example, TXU's actual bids yield 39.3% of the profits that would have been realized under our best-response benchmark. However, if TXU had used the simple trading rule, it would have earned 96.7% of best-response profits, which indicates that there is strong persistence in the shape of residual demand across auctions. Similarly, all firms except Reliant would have significantly increased profits by following this simple trading rule.

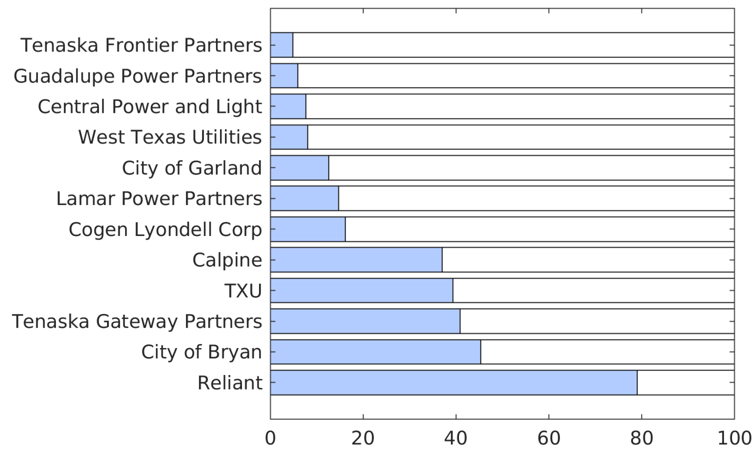

(a) Realized profits

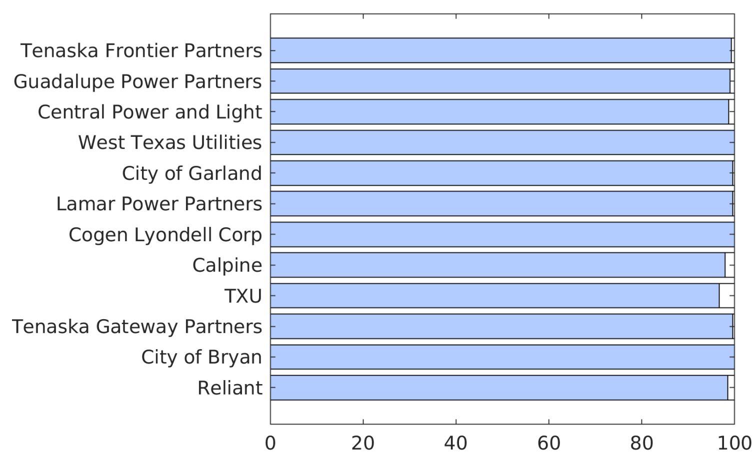

(b) Profits best-responding to lagged bids

FIGURE C.1: Realized profits and predicted profits best-responding to lagged bids, as percentage of potential profits

D Evidence that Firms Do Not Misrepresent Capacity

We rule out the possibility that our measure of capacity – the firm's self-declared capacity for each day – overstates the actual capacity. We compare each firm's stated capacity to the highest amount of production that we observe during our sample. All firms are observed to use at least 75% of stated capacity and on average to use 90% of stated capacity (see Table D.1). This suggests that our finding that firms do not bid significant capacity into the balancing market is not driven by overstating capacity. Moreover, the concern of overstated capacity does not apply to periods when firms *decrease* production, or 'dec', and we observe deviations from best-response in 'dec' intervals as well.

TABLE D.1: Capacity utilization relative to self-declared capacity

Firm	Maximum capacity utilization (%)
Reliant	81.72
City of Bryan	76.59
Tenaska Gateway Partners	125.88
TXU	97.13
Calpine Corp	83.84
Cogen Lyondell Inc	81.12
Lamar Power Partners	76.19
City of Garland	93.57
West Texas Utilities	92.92
Central Power and Light	98.82
Guadalupe Power Partners	74.69
Tenaska Frontier Partners	93.40

Note: The table reports maximum capacity utilization relative to self-declared capacity for each day, for the firms that we consider in the Cognitive Hierarchy.

E Examining the Impact of Selection of Firms to Include in the Cognitive Hierarchy

In this appendix, we explore how choosing a different set of firms to include in the Cognitive Hierarchy affects our baseline estimates. Our strategy is to narrow the set of firms to incorporate into the CH to those that would result in no more than 15 percent of auctions being lost. This set consists of five firms. We denote this set $\mathscr{F}$. We then separately incorporate each each firm $j \in \mathscr{F}$ to the CH. However, because our objective is to examine how sample selection impacts our findings, we want to keep 12 firms in the hierarchy throughout the exercise. For this reason, for each firm $j \in \mathscr{F}$ that we include in the hierarchy, we sequentially loop over the set of original firms in the hierarchy, and we drop one each time. This leads to having 12 firm samples for each $j \in \mathscr{F}$. We re-estimate the baseline specification of our model for each of these samples using 10 random initial points (using Halton sequences), and recover the estimated parameters for each of these estimation routines.

To examine the impact of our sample selection, we use the estimated parameters of the exercise described above and compute, for each exercise, the expected type associated with the probability distribution implied by those estimates. We report the cdf of the expected type, for the largest and smallest firm in our data, in Figure E.1, together with the cdf associated with our main specification. The figure shows that varying the sample of firms included in the CH has little impact on the estimated types.

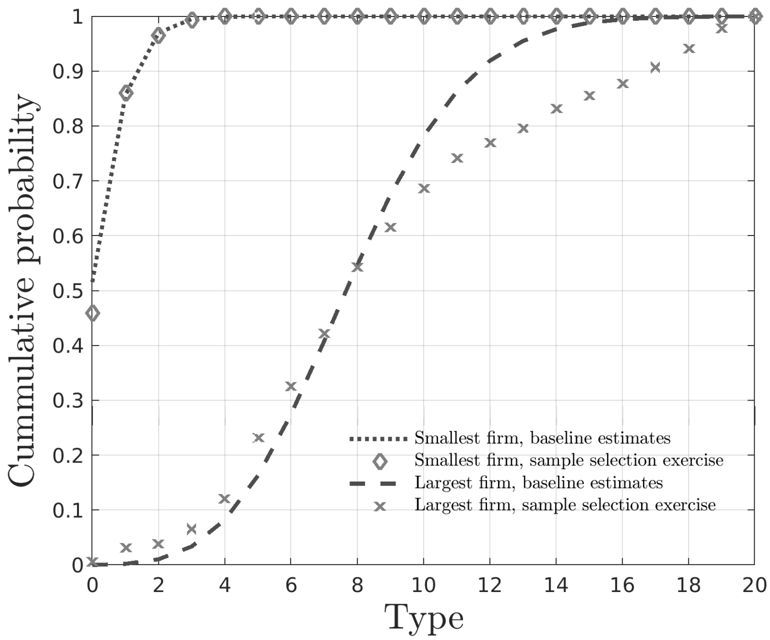

FIGURE E.1: Robustness to Sample Selection of Firms Entering Cognitive Hierarchy: Estimated expected types and probability distributions over types for the smallest and largest firms

F Proof that CH Bids for Level-k Players are Additively Separable in Price and Contract Position

In this appendix, we show that bids for level-k players, for $k > 0$, are additively separable in price and private information on contract quantity, so bid functions take the form: $S^1_{it}(p, QC_{it}) = \alpha^1_{it}(p) + \beta^1_{it}(QC_{it})$.

To see this, consider first the case of bidders type-1. In this case, bids $S^1_{it}(p)$ can be calculated from equation (5), which can be rewritten as

$$\begin{aligned}
S^1_{it}(p) &= \left[(p - C'_{it}\left(S^1_{it}(p)\right)\right] \frac{H_p\left(p, S^1_{it}(p); k_i, QC_{it}\right)}{H_s\left(p, S^1_{it}(p); k_i, QC_{it}\right)} + QC_{it} \\
&= \left[(p - C'_{it}\left(S^1_{it}(p)\right)\right] \frac{-\int_{l_{-i}} \gamma\left(D_t(p) - \hat{S}^1_{it}(p)\right) D'_t(p)\Delta(l_{-i})}{\int_{l_{-i}} \gamma\left(D_t(p) - \hat{S}^1_{it}(p)\right) \cdot \Delta(l_{-i})} + QC_{it} \\
&= \alpha^1_{it}(p) + QC_{it}
\end{aligned}$$

because the argument $[(p - C'_{it}(S^1_{it}(p))] \frac{-\int_{l_{-i}} \gamma\left(D_t(p) - \hat{S}^1_{it}(p)\right) D'_t(p)\Delta(l_{-i})}{\int_{l_{-i}} \gamma\left(D_t(p) - \hat{S}^1_{it}(p)\right) \cdot \Delta(l_{-i})}$ is a function of price p.

Therefore, bids of type-1 bidders are additively separable and can be represented by $S^1_{it}(p) = \alpha^1_{it}(p) + QC_{it}$, where $\alpha^1_{it}(p) = [(p - C'_{it}(S^1_{it}(p))] \frac{-\int_{l_{-i}} \gamma\left(D_t(p) - \hat{S}^1_{it}(p)\right) D'_t(p)\Delta(l_{-i})}{\int_{l_{-i}} \gamma\left(D_t(p) - \hat{S}^1_{it}(p)\right) \cdot \Delta(l_{-i})}$.

Similarly, for a bidder type-k, bids $S^k_{it}(p)$ are given by

$$\begin{aligned}
S^k_{it}(p) &= \left[(p - C'_{it}\left(S^k_{it}(p)\right)\right] \frac{H_p\left(p, S^k_{it}(p); k_i, QC_{it}\right)}{H_s\left(p, S^1_{it}(p); k_i, QC_{it}\right)} + QC_{it} \\
&= \left[(p - C'_{it}\left(S^1_{it}(p)\right)\right] \frac{-\int_{l_{-i}} \gamma\left(D_t(p) - \sum_{j \neq i} \alpha^{l_j}_{jt}(p) - \hat{S}^1_{it}(p)\right) \tilde{D}'_t(p)\Delta(l_{-i})}{\int_{l_{-i}} \gamma\left(D_t(p) - \hat{S}^1_{it}(p)\right) \cdot \Delta(l_{-i})} + QC_{it} \\
&= \alpha^k_{it}(p) + QC_{it},
\end{aligned}$$

where $\alpha^k_{it}(p) = [(p - C'_{it}(S^1_{it}(p))] \frac{-\int_{l_{-i}} \gamma\left(D_t(p) - \sum_{j \neq i} \alpha^{l_j}_{jt}(p) - \hat{S}^1_{it}(p)\right) \tilde{D}'_t(p)\Delta(l_{-i})}{\int_{l_{-i}} \gamma\left(D_t(p) - \hat{S}^1_{it}(p)\right) \cdot \Delta(l_{-i})}$ and $\tilde{D}'_t(p) = D'_t(p) - \sum_{j \neq i} \alpha'^{l_j}_{jt}(p)$.

G Examining the Impact of Assumption 3

In this section, we examine the extent to which Assumption 3 (i.e., Γ being Uniform) may impact our analysis. As described in the text, this assumption is made due to computational complexity. Therefore, to examine how Assumption 3 may impact our findings, we need to make some simplifications. Instead of significantly reducing the number of firms and types that we can allow for in estimation, we use the estimated parameters that we report in section V and predict bidding behavior assuming that Γ_i is Uniform and that Γ_i is Normal, and compare the distribution of predicted bids. Importantly, in this exercise we include the two largest firms in our data, because if assumptions about Γ have any impact on our findings, this impact will show up when considering relatively high-type firms. This is so because our assumption about how level-0 bidders behave implies that Γ does not affect bidding behavior of bidders type 0 and 1 (note that this is true for any distributional assumption regarding Γ). We present our findings in Figure G.1 and Figure G.2. The figures show a number of important insights.

First, as we described above, Γ does not enter the first-order condition of bidders type 0 and 1. This is shown in the figures and implies that, for low-type bidders, assuming that Γ is Uniform (or any other distribution) is irrelevant.

Second, the difference between distributions at higher types is small, suggesting that the Uniform distribution allows us to take into account the same degree of uncertainty as an alternative distribution, with considerable savings on computational complexity.

Finally, the small differences between distributions, in particular at low types, is appealing. This is so because most of the inefficiencies that we estimate are caused by low-type firms (see section VIII). For this reason, assuming that Γ is Uniform not only does not appear to have (material) impact on the predicted bid functions, but also suggests that it does not affect our estimation of efficiency gains.

Because the exercise just described leads us to conclude that there are no (material) advantages in using a different distribution for Γ, relative to the Uniform distribution, and the cost of departing from the Uniform distribution is significant (we need to restrict the number of types and firms in the CH, which also restrict the degree of heterogeneity in firms attributes that we can consider), we believe that our decision to assume that Γ is Uniform is the best one as it allows us to include more firms (and heterogeneity) and types than alternative assumptions.

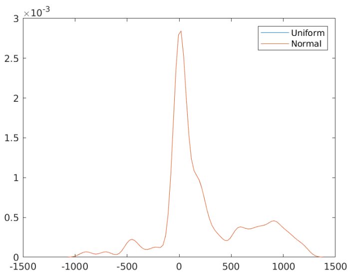

(a) $k = 0$

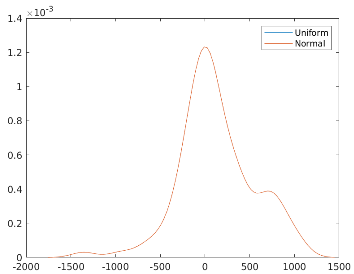

(b) $k = 1$

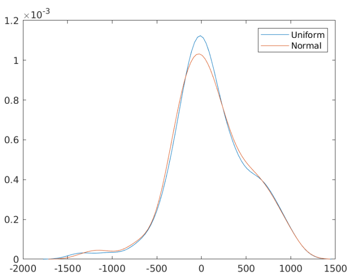

(c) $k = 2$

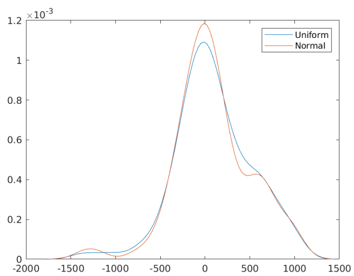

(d) $k = 3$

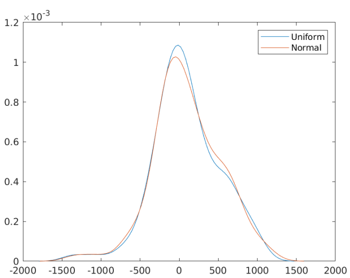

(e) $k = 4$

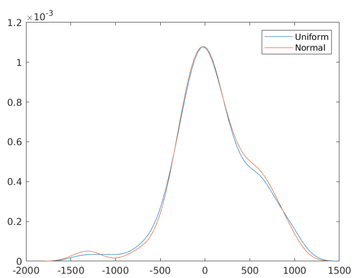

(f) $k = 5$

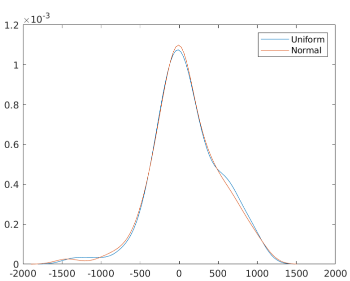

(g) $k = 6$

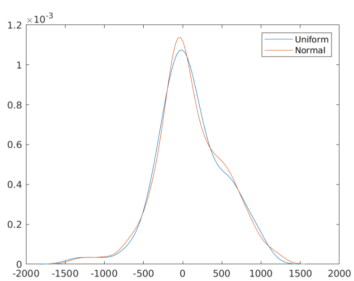

(h) $k = 7$

FIGURE G.1: Largest firm

Notes: The figures report the distribution of (quantity) bids at the market clearing price, across all auctions, for the largest firm in the balancing market.

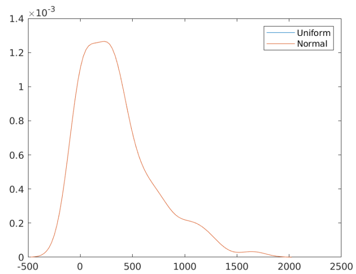

(a) $k = 0$

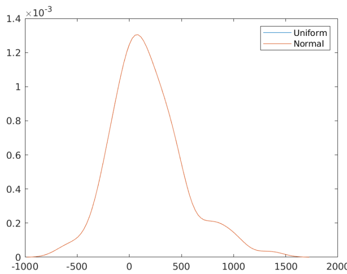

(b) $k = 1$

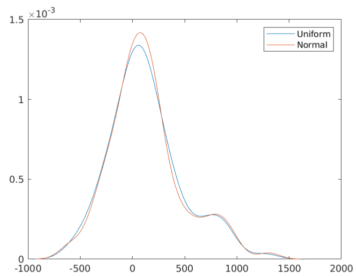

(c) $k = 2$

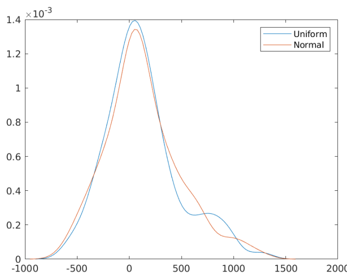

(d) $k = 3$

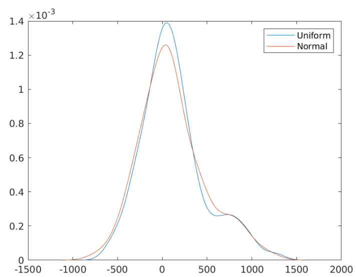

(e) $k = 4$

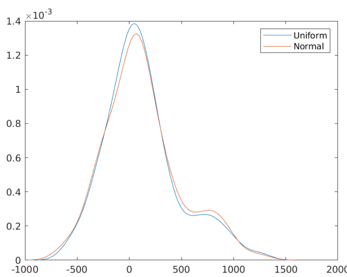

(f) $k = 5$

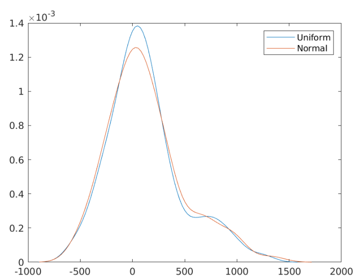

(g) $k = 6$

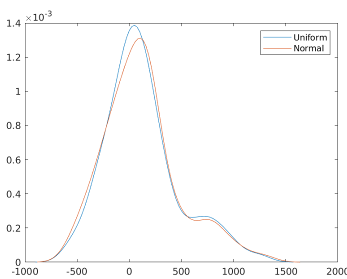

(h) $k = 7$

Figure G.2: Second largest firm

Notes: The figures report the distribution of (quantity) bids at the market clearing price, across all auctions, for the second largest firm in the balancing market.

H Evidence There is Not Economically Significant Uncertainty in the Slope of Supply by Unmodeled Firms

In this appendix, we show that the slope of supply by the firms that are not included in the Cognitive Hierarchy is not economically significant from the perspective of introducing uncertainty to residual demand slope by the modeled firms.

Recall that in our implementation, the $D(p)$ in the theoretical model corresponds to balancing demand (which is perfectly inelastic) *net* of (elastic) supply by the unmodeled firms. In this context, our model would imply that there is no uncertainty about the *slope* of the supply of the unmodeled firms.

While it is not literally true that there is no uncertainty, to a first-order, there is not economically significant slope uncertainty. Across the 99 auctions, there is variation in the slope of unmodeled firm supply, but some of this is predictable by the 12 firms, from, say changes in fuel costs or other seasonal factors. It is therefore necessary to assess if there is economically meaningful variation in the *uncertainty* of slope. In order to do so, we run a regression with 99 observations where the dependent variable is the slope of the aggregate supply bids by the unmodeled firms (specifically, it is the $\frac{dq}{dp}$ where we linearize using the aggregate bid quantity plus and minus \$10 around the mean market-clearing price). We regress this slope on known cost factors – daily gas prices, total system load, and day-of-week and week-of-year fixed effects. Then we take the residuals, interpreting the residuals as the slope uncertainty that cannot be predicted. When we do this, we find that the standard deviation of the residuals is 7.10 (i.e. the quantity sensitivity to a one dollar change in price has a standard deviation of 7.10 MW).

The scale of this uncertainty is economically very small. To demonstrate this, we illustrate the impact of this uncertainty on the largest firm, which is the one for which uncertainty is most important because it is a high-type firm. (The low-type firms will be less sensitive to slope uncertainty.) Specifically, we calculate how much profits would change if the large firm were "incorrect" in the slope of residual demand by 2 standard deviations of this level of uncertainty. We calculate the largest firm's profits of submitting two different bids: (1) the best-response to the actual residual demand slope when that slope is realized, and (2) the best-response when the firm *thought* residual demand had a slope that was bigger by two standard deviations of this uncertainty (i.e. had a $\frac{dq}{dp}$ that was $2*7 = 14$ MW larger) but the *actual* residual demand is realized. The profits are essentially identical — across the 99 auctions, the median difference in profit is zero.

We conclude that the amount uncertainty in residual demand *slope* in our 12 firm implementation of our model is very small, thus justifying our modeling assumption.

I Evidence that Firms Do Not Bid Marginal Cost

In this appendix, we provide evidence that firms do not bid marginal cost. If firms submit bids that deviate from best response in the direction of marginal cost bidding, then the actual quantity sold into the auction should be *greater* than predicted sales under best-response bidding. However, this is not the case in our setting – Table I.1 shows that firms systematically sell less output that the best-response level of output, which supports our assumption that firms deviate by submitting bids that are steeper than best-response.

Table I.1: Comparison of Actual Output to Best-Response Output

Firm	**Actual Output**	**Best-Response Output**
Reliant	431	507
TXU	133	441
Calpine	102	408
Guadalupe	12	396
Central Power & Light	35	352
Lamar Power Partners	30	272
Cogen Lyondell	34	269
West Texas Utilities	11	224
Tenaska Gateway	72	182
Tenaska Frontier	7	144
Garland	5	115
Bryan Texas Utilities	30	56

Note: This table reports average output under actual bids and best-response bids from the opening of the market until January 31, 2003. The numbers are the average of the absolute value of sales so that output during both INC and DEC intervals has the property that marginal cost bidding will yield actual output greater than best-response output and that bidding steeper than best-response will yield actual output less than best-response output. Source: Hortacsu and Puller (2008).

J Robustness to Another Time Period of the Day

In this appendix, we show that our estimated relationship between firm size and type is not driven by the time period of the day that we select. We estimate the baseline model (column 1 of Table 1) using data from the time period 7–8pm. The lines in Figure J.1 depict the CDFs of the estimated types for firms of various sizes – the smallest, fourth from the smallest, ninth from the smallest, and the largest. The solid lines are the estimated type distributions reported in the paper for our 6–7pm sample. The grey lines show the corresponding estimates from the 7–8pm sample. The CDFs of the smaller sets of firms essentially overlap. There are slight differences in the CDFs for the large firms, and these differences are driven by the estimated constant of the 7–8pm sample being slightly smaller (in absolute value) than for our primary sample period.

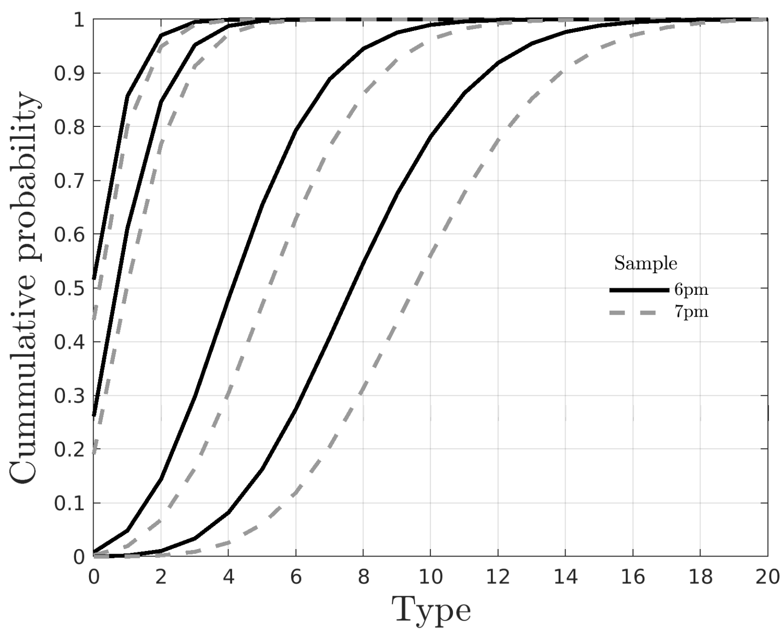

FIGURE J.1: Estimated CDFs for two samples. From left to right the CDFs correspond to those of larger firms, as reported in Figure 6 in the main text

K Estimated types over time

In this appendix, we show that learning is, in practice, minimal. We re-estimate our main specification including a linear trend, which allows us to capture whether the estimated types change over time. Here we present the estimated probability distributions over types for the smallest and largest firms in our data, for the first and last week in our sample. The figure shows that for both firms the implied probability distributions over types overlap, which suggest that types did not change significantly during our sample period.

FIGURE K.1: Estimated CDFs over types for the smallest and largest firm, for the first and last week of the sample (*Size* Specification)

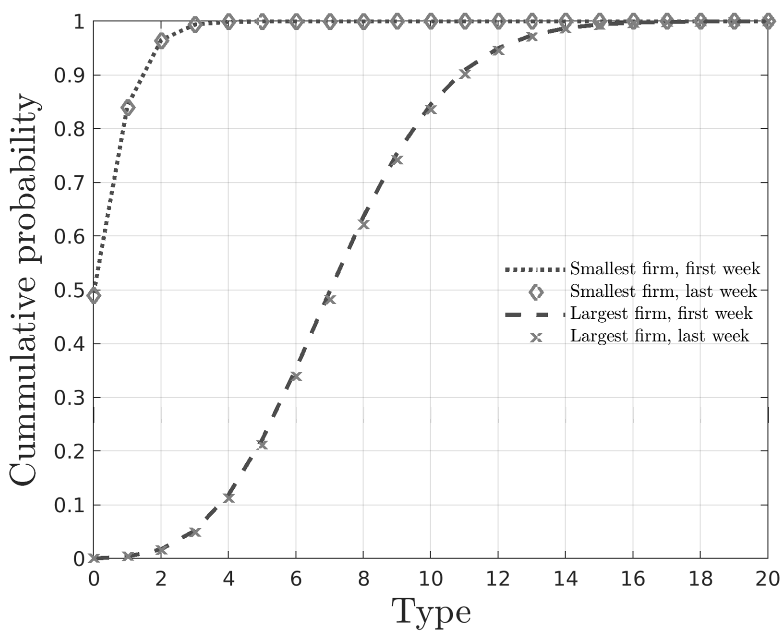

L Diminishing Returns to Increasing Sophistication

In this appendix, we investigate whether the private returns to increasing sophistication are decreasing or increasing. We do this for two firms, a small one (the one with the highest probability of being type 0) and a medium-size one. Because types are parameterized by size, we sequentially increase firm size until the firm reaches the same capacity as the largest firm in the market. As before, we do this for auctions that clear on the INC and DEC side separately. The results are reported in Figure L.1, that reports incremental returns relative to the status quo of each firm. The figure confirms that there are decreasing marginal returns to increasing sophistication.

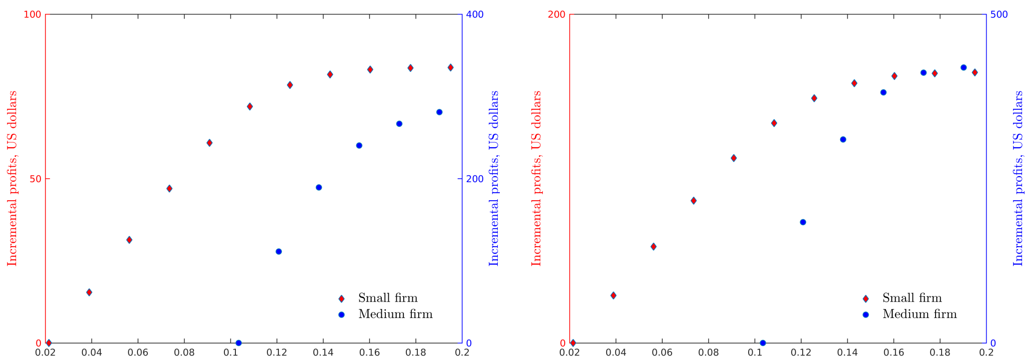

(a) INC side

(b) DEC side

FIGURE L.1: Marginal returns to increasing sophistication

References

Hortacsu, Ali, and Steven L. Puller. 2008. "Understanding Strategic Bidding in Multi-Unit Auctions: A Case Study of the Texas Electricity Spot Market." *RAND Journal of Economics*, 39(1): 86–114.