Arrow Lecture: IIAS, Jerusalem, July 2018

Ariel Pakes, Harvard Unviersity

June 25, 2018

Just Starting Out : Learning and Price Competition in a New Market

U. Doraszelski, G. Lewis, and A. Pakes
(Wharton, Microsoft Research, Harvard)

March 9, 2017

Goal: Empirical start at the analysis of market responses to changes in its environment.

- A goal of Industrial Organization is to predict the response of markets to policy &/or environmental changes.
- Counterfactual outcomes are typically analyzed by looking for a rest point where each firm is doing the best it can given the actions of other firms.
- This ignores both the transition period and the problem of selecting one of the many possible rest points.
- We empirically analyze the transition from a market allocation determined by governmental fiat to an auction. Inital uncertainty about demand and rival behavior led to bids which we analyze with models which allow for learning about rivals' play and the response of quantities to bids.
- We show that over time prices stabilized, converging to a rest point that is consistent with equilibrium play and that the learning models lead to notably better bid predictions.

Case Study

- FR market UK: FR is a product that is used to keep the frequency of electricity in the wires within pre-specified bounds. It is needed for the system to keep running.
- Quantities that must be held are set by regulation.
- Prior to November 1 2005, FR was obtained by fiat. The owner operator of the grid (National Grid, or NG) has governor controls which can take over a generator.
- Generators were paid an administered price.

Deregulation.

- November 1: Market is deregulated.
- Firms compete on price through bids.
- NG now "buys" the right to "governor control" a generator.
- We follow the market for the next six years.

Bids.

- Separate bids for each generator. Typical station has several generators. Some firms own multiple stations.
- Bids are submitted monthly and consist of a bid and a max quantity (tied to "operating position" of generator).
- Bids submitted by the 20th of the prior month.
- Accepted bids gives NG ability to instruct a generator into FR mode whenever it choses during the month.

- Providing FR is costly: a BM unit in FR mode which is called incurs additional wear and tear and a diminished "heat rate" as it may have to make rapid adjustments to its energy production in response to supply and demand shocks.

Market Clearing

- NG pays bid $\times$ time BM unit put in FR mode in MWh (plus +/- compensation for change in costs of energy supplied to balancing market).
- A supercomputer with a proprietary program does the quantity allocation in real time to cost minimize subject to constraints: transmission, reliability, position in the balancing market (it needs "head-room" and "foot-room").
- We were told that the computer program was unchanged over our time period.

Data

- Bids (NG: monthly, by generating unit).
- Quantities (NG: daily, by generating unit).
- Main market positions (Elexon: half-hourly, by unit).
- Ownership structure, location and type of generators.

Structure of market.

- 130 BM units in 61 stations owned by 29 firms.
- We focus on the bids of 10 big players (table 1) whose combined market share is 80-85%.
- All of these players were active pre-deregulation ($\Rightarrow$ experience with demand and cost conditions).
- Demand is highly variable over days and is seasonal, but pretty stable (figure 3) across years .
- Aggregate quantity regulated = f(quantities of electricity demanded).

Table 1: Firms with the largest frequency response revenues

{revenue_ta

Rank	Firm name	Num Units Owned	Total Revenue	Revenue Share (%)	Cumulative Share (%)
1	Drax Power Ltd.	6	99.4	23.8	23.8
2	E.ON UK plc	20	67	16	39.9
3	RWE plc	23	48.4	11.6	51.6
4	Eggborough Power Ltd	4	29.8	7.1	58.7
5	Keadby Generation Ltd	9	24.2	5.8	64.5
6	Barking Power Ltd	2	17.8	4.2	68.8
7	SSE Generation Ltd	4	15.2	3.6	72.5
8	Jade Power Generation Ltd	4	15	3.6	76.1
9	Centrica plc	8	14.7	3.5	79.6
10	Seabank Power Ltd	2	14	3.3	83

Inflation-adjusted revenue in millions of british pounds (base period is October 2011). There is information on 72 months in the data. The number of units owned is the maximum ever owned by that firm during the sample period.

over the sample period, or about £1,400,000 per month. Drax is a single-station firm, while the next two largest firms, E.ON and RWE, are multi-station firms. Anecdotally, Drax's disproportionate share is attributable to having a relatively new plant, with accurate governor controls, making it attractive for providing FR. The smallest firm, Seabank, still makes around £200,000 per month. This suggests that the FR market was big enough that firms may have been willing to devote an employee's time to actively managing their bidding strategy, at least when the profitability of the market became apparent. In 2006 Drax indeed hired a trader to specifically deal with the FR market.[6] Within a year, Drax's revenue from the FR market increased more than threefold.

Physical environment. The FR market operates within a relatively stable environment. Starting with the demand for FR, the left panel of Figure 3 plots the monthly quantity of FR. Though this series is clearly volatile, it is no more volatile at the beginning than at the end of the period we study (and as we show below, the bids are). The right panel of Figure 3 shows some evidence of modest seasonality.

As further evidence of the stability of the demand for FR, the left panel of Figure 4 plots the daily total amount of electricity demand over the sample period. The response requirement

[6]Source: private discussion with Ian Foy, Head of Energy Management at Drax.

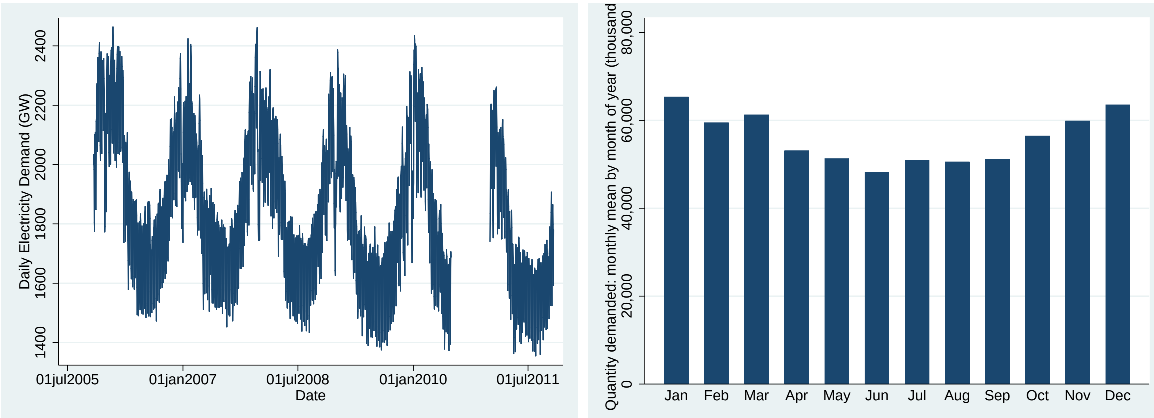

Figure 4: Total amount of electricity demanded by day (data missing from *** to *** GREG, CAN YOU PLEASE FILL IN? *** left panel) and on average by month-of-year (right panel). {main_marke

In addition to the mandatory frequency response (MFR) that is the focus of this paper, NG uses long-term contracts with BM units to procure FR services. This is known as firm frequency response (FFR). Figure 5 plots the monthly quantity of FFR and, for comparison purposes, that of MFR (see also the left panel of Figure 3). The quantity of FFR remains relatively stable over our sample period up until July 2010, when it almost doubles and thereafter remains stable at the new level.

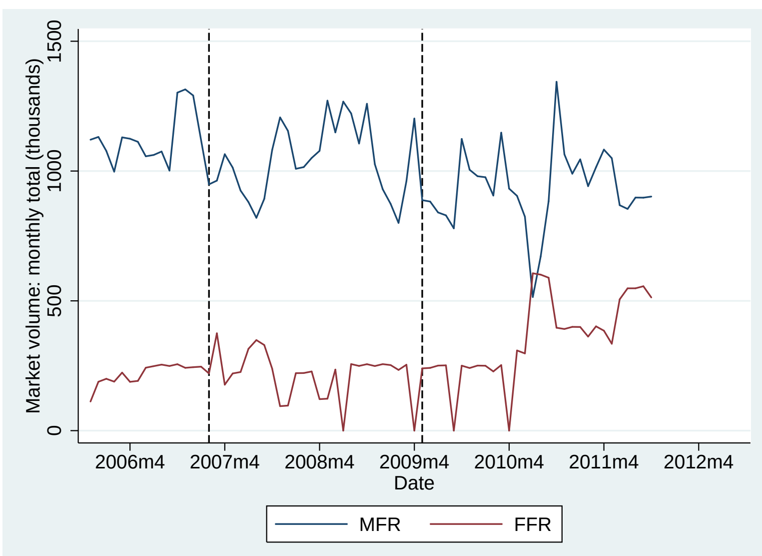

Figure 5: MFR and FFR quantities by month. {ffr_quanti

Changes in Market Environment Over Time.
(Divide six years into three periods: see below.)

• Differences in physical environment across periods: need to distinguish their effects on bids from learning induced changes.

- NG can also use long-term contracts with BM units to procure FR services (called FFR). This is fairly stable proportion of total FR until the middle of third period (figure 4).
- Fuel prices (figure 4): Also fairly stable until the later period.
- Number of active stations and quantity share of active stations constant until 2010 (third period, figure 5). ("Active": A station can absent itself from the market by submitting an unreasonably high bid).

Conclude. Bidding environment (cost, demand, participants) fairly stable until last period.

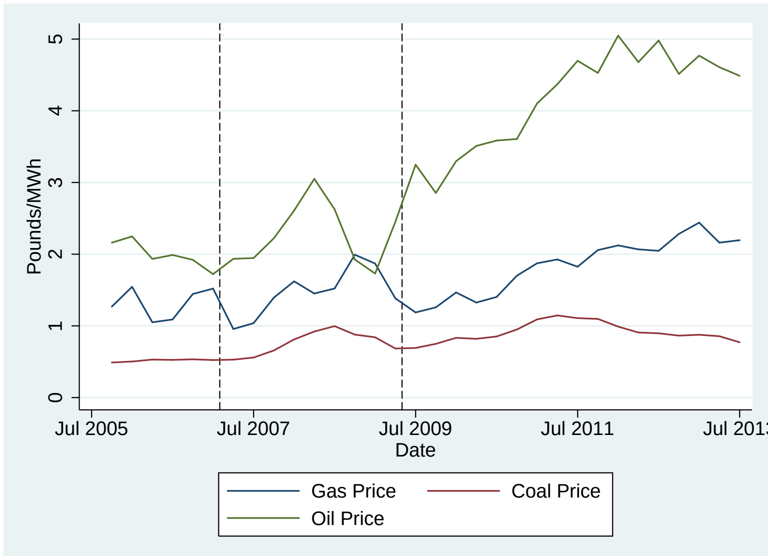

Figure 7: Fuel prices by quarter. {main_marke

3 Evolution of the FR market {evolutio

Our discussion divides the evolution of the FR market into three phases that differ noticeably in bidding behavior. Figure 8 shows the average monthly price of FR, computed as the quantity-weighted average of bids, with vertical lines separating the three phases. For comparison purposes, Figure 8 also shows the unweighted average of bids.

During the early phase from November 2005 to February 2007, the price exhibits a noticeable upward trend, moving from an initial price of £3.1/MWh to a final price of £7.2/MWh. The upward trend culminates in a "price bubble." During the middle phase from March 2007 to May 2009, this trend reverses itself and the price falls back down to £4.8/MWh. From June 2009 to the end of our study period in October 2011 there is no obvious trend at all. While there are fluctuations during this late phase, they are smaller, and the price stays in the range of £4.3/MWh to £5.1/MWh. The sharper movements in one direction are relatively (to the prior periods) quickly corrected by movements in the opposite direction.

The movements in the price of FR in the early phases in Figure 8 occurred despite the relative stability of the demand and supply conditions (see Section 2), and are too persistent to be driven by seasonality in the demand for FR. Although there are some changes in FFR and an upward trend in the number of active power stations as well as in the oil and gas

Turning from the demand to the supply of FR, a BM unit can opt out of the FR market by submitting an unreasonably high bid. The left panel of Figure 6 plots the number of "active" power stations over time, where we define a station as active if one of its BM units submits a competitive bid of less than or equal to £23/MWh (see Appendix A.2 for details). The number of active stations fluctuates a bit, ranging from 53 to 61 over the sample period. In the first four years of the FR market, the fluctuations are relatively small and none of the stations who become active or inactive is particularly large. The right panel of Figure 6 shows that the share of stations that are always active is steady at around 95%. There are some larger fluctuations in last two years of the FR market.

Figure 6: Number of active power stations by month (left panel) and market share of always-active power stations (right panel). {active_sta

Finally, Figure 7 plots quarterly fuel prices paid by power stations in the UK over time. Fuel prices may matter for the FR market in that they change the "merit order" in the main market. For example, when gas is relatively expensive, gas-powered BM units may be part-loaded and therefore available for FR, whereas coal-powered BM units may be operating at full capacity and thus require repositioning in the BM in preparation for providing FR. Though there are some upward trends in oil and — to a lesser extent — gas prices, they are largely confined to the end of the sample period.

In sum, until the middle of 2009, the physical environment and demand and supply conditions are stable. After that date, FFR plays a larger role and the number of active power stations rises, as do oil and gas prices. Thus, at least prior to the middle of 2009 any volatility in bids is unlikely to be caused by changes in demand or supply conditions.

Market Outcomes

- Table 1 again: the stakes vary from
  - Drax, $\approx$ 16.6 million 2011 $\mathcal{L}$/year, to
  - Seabank, $\approx$ 2.3 million $\mathcal{L}$/year.
  - Enough to hire a person, even for Seabank.
- Share weighted average bid (figure 6).
  - Initial rising price period (from 3.1 to 7.2 $\mathcal{L}$/Mwh).
  - Period 2: Price move down (from 7.2 to 4.8 $\mathcal{L}$/Mwh).
  - Period 3: prices "cycle" (between 5.1 and 4.3 $\mathcal{L}$/Mwh).
- Graph differences in bids across periods in
  - probability of bid change (figure 7),
  - magnitude of changes (figure 8),
  - variance of bids across BM units (figure 9),
  - breakdown of variance: within and between firm (figure 10),
  - Bids of largest 8 firms by period (figure 11).

Table 1: Firms with the largest frequency response revenues

{revenue_ta

Rank	Firm name	Num Units Owned	Total Revenue	Revenue Share (%)	Cumulative Share (%)
1	Drax Power Ltd.	6	99.4	23.8	23.8
2	E.ON UK plc	20	67	16	39.9
3	RWE plc	23	48.4	11.6	51.6
4	Eggborough Power Ltd	4	29.8	7.1	58.7
5	Keadby Generation Ltd	9	24.2	5.8	64.5
6	Barking Power Ltd	2	17.8	4.2	68.8
7	SSE Generation Ltd	4	15.2	3.6	72.5
8	Jade Power Generation Ltd	4	15	3.6	76.1
9	Centrica plc	8	14.7	3.5	79.6
10	Seabank Power Ltd	2	14	3.3	83

Inflation-adjusted revenue in millions of british pounds (base period is October 2011). There is information on 72 months in the data. The number of units owned is the maximum ever owned by that firm during the sample period.

over the sample period, or about £1,400,000 per month. Drax is a single-station firm, while the next two largest firms, E.ON and RWE, are multi-station firms. Anecdotally, Drax's disproportionate share is attributable to having a relatively new plant, with accurate governor controls, making it attractive for providing FR. The smallest firm, Seabank, still makes around £200,000 per month. This suggests that the FR market was big enough that firms may have been willing to devote an employee's time to actively managing their bidding strategy, at least when the profitability of the market became apparent. In 2006 Drax indeed hired a trader to specifically deal with the FR market.[6] Within a year, Drax's revenue from the FR market increased more than threefold.

Physical environment. The FR market operates within a relatively stable environment. Starting with the demand for FR, the left panel of Figure 3 plots the monthly quantity of FR. Though this series is clearly volatile, it is no more volatile at the beginning than at the end of the period we study (and as we show below, the bids are). The right panel of Figure 3 shows some evidence of modest seasonality.

As further evidence of the stability of the demand for FR, the left panel of Figure 4 plots the daily total amount of electricity demand over the sample period. The response requirement

[6]Source: private discussion with Ian Foy, Head of Energy Management at Drax.

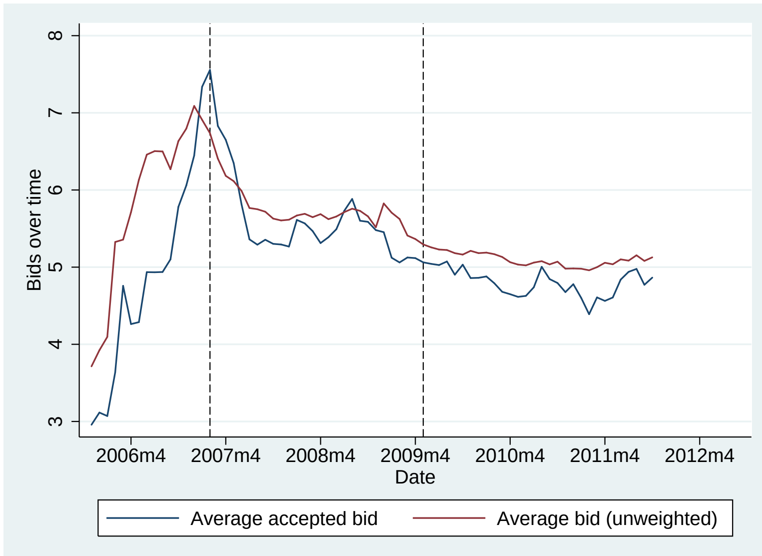

Figure 8: Quantity-weighted and unweighted FR price by month. Weights are based in month t.

{monthly_pr

prices, most of that action occurs towards the end of the sample period, when the price of FR has become quite stable. We therefore look for an alternative explanation for the changes in bidding behavior over time. In particular since none of the participants in this market had any experience bidding into it, it seems unlikely that they had strong priors about either how their competitors would bid, or about how their allocation of FR would vary with their bid conditional on how their competitors would bid.

We begin with a summary of how bidding behavior changed from one phase to the next. After providing the overview, we look more closely at the role of individual power stations.

Early or rising-price phase (November 2005 – February 2007). In the early or rising-price phase, firms change the bids of their BM units more often and by larger amounts (in absolute value) than in the middle and late phases. On average, the bids of 4 out of 10 BM units change each month by between £1/MWh and £3/MWh (conditional on changing). This is illustrated in Figures 9 and 10.

In addition to changing their bids more often and by noticeably larger amounts, firms tender very different bids in the early phase. Figure 11 shows that the range of bids as measured by the variance of bids across BM units is an order of magnitude larger than in the middle

Initial Period: Diverse Strategies.

- Dominant change is upward (price "bubble").
- 4 out of 10 bid changes/month. Changes are between 1-3 $\mathcal{L}$.
- Drax (1/4 of revenues) does little for a year, then hires someone who triples their profits (Figure 11).
- Bids differ widely: between 1.5 and 10 $\mathcal{L}$/Mwh. Large variance across generators (Figure 9).
- Some intra-firm variance in bids which goes away in middle period (Figure 10). Experimentation?

and late phases.

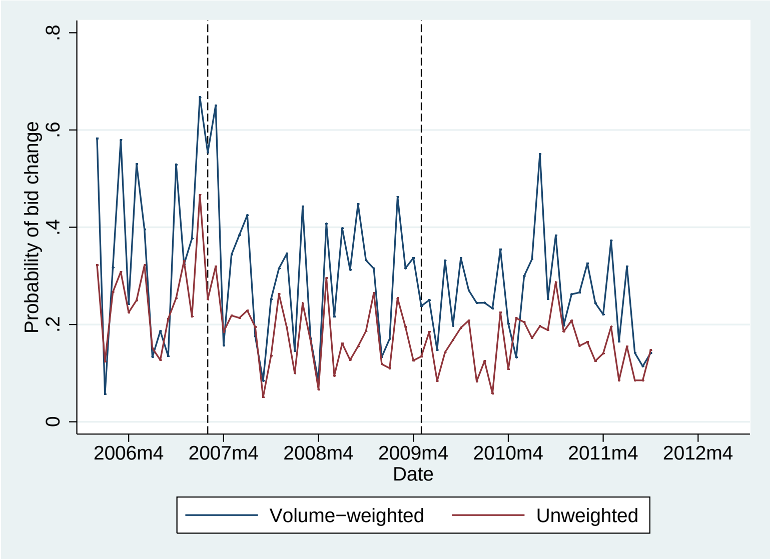

Figure 9: Quantity-weighted and unweighted probability of a bid change between month t and $t-1$. Weights are based in month $t-1$.

{changesbid

Comparing the left and right panels of Figure 12 shows that most of the variance stems from differences in bids between firms (across-firm variance, right panel) rather than from differences between BM units within firm (within-firm variance, left panel). What within-firm variance there is, is highest in the early phase and then declines, suggesting that firms initially experimented by submitting different bids for their BM units, and that such experimentation became less prevalent over time.

Figure 13 shows the monthly bids of the eight largest power stations by revenue in the FR market. The top left panel provides a more detailed look at the early phase. In line with the wide range of bids documented in Figure 11 and the right panel of Figure 12, the levels and trends of the bids are quite different across stations. Firms seem to experiment with different bids during the early phase of the FR market. Barking, Peterhead and Seabank bid very high early on — pricing themselves out of the market — and then drift back down into contention. The remaining stations start low and then gradually ramp up. The big increase in bids by Drax during late 2006 and early 2007 leads to the "price bubble" in Figure 8.

Middle or falling-price phase (March 2007 – May 2009). In the middle or falling-price phase, firms change the bids of their BM units less often and by much smaller amounts

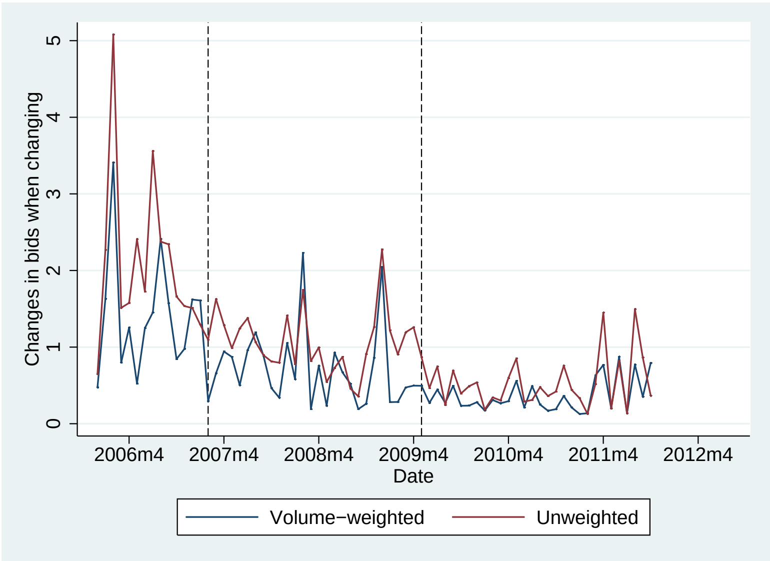

Figure 10: Quantity-weighted and unweighted absolute value of bid change conditional on changing between month t and $t-1$. Weights are based on month $t-1$ and are zero if the BM unit's bid did not change.

{bidchanges

(in absolute value) than in the early phase. As Figures 9 and 10 illustrate on average the bids of 3 out of 10 BM units change each month by around £1/MWh (conditional on changing). Figure 11 shows that the range of bids is much narrower than in the early phase.

The top right panel of Figure 13 provides more detail. The "price bubble" bursts when Seabank and Barking sharply decrease their bids and steal significant market share from Drax. Drax follows Seabank and Barking down, and this inaugurates intense competition and the noticeable downward trend in the price of FR in Figure 8. Experiments with increased bids are not successful. Drax, for example, increased its bid at the end of 2007 for exactly two months, giving its rivals an opportunity to see its increased bid and followed suit. When no one did, Drax decreased its bid.[7]

The dominant trend in the top right panel of Figure 13 is for the bids of the different power stations to move toward one another. The way this happens is that the stations that entered the middle phase with relatively high bids decreased their bids while the firms that entered the phase with relatively low bids maintained those bids. This intense competition generated

[7]Also, toward the middle of 2008, Eggborough increased its bid. While this was followed by lesser increases by Drax and Connah's Quay, Eggborough soon undercut Drax and Connah's Quay and was subsequently followed down by them.

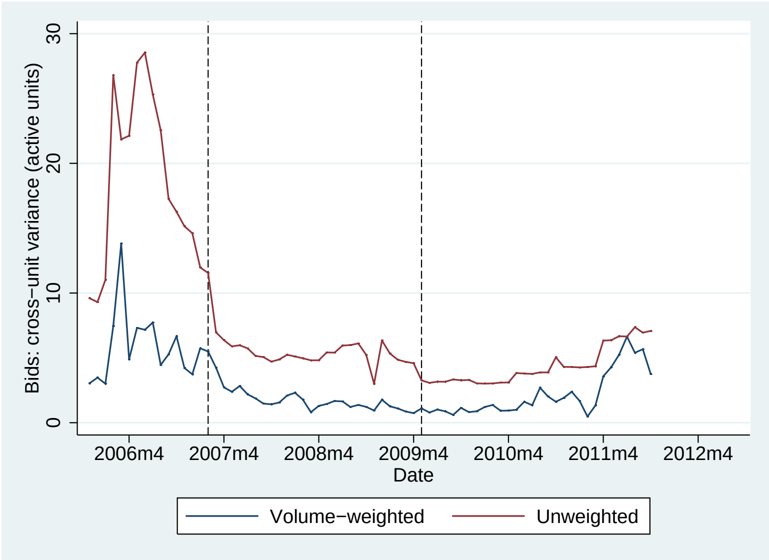

Figure 11: Quantity-weighted and unweighted variance in bids across BM units by month. Weights are based in month t.

{bidvarianc

the marked decrease in the range of bids in Figure 11.

Late or stable-price phase (June 2009 – October 2011). In the late or stable-price phase, firms change the bids for their BM units as often as in the middle or falling-price phase, but by much smaller amounts (in absolute value). As Figures 9 and 10 illustrate, on average, the bids of 3 out of 10 BM units change each month by around £0.5/MWh (conditional on changing). Figure 11 shows that the range of bids is again much narrower than in either of the earlier phases. The bottom panel of Figure 13 provides more detail. While bids at some power stations continue to fall (Rats and Cottam), others are more erratic or rise (Drax and Eggborough), and others are almost completely flat (Peterhead). Overall, however, the bids of the different stations are noticeably closer to one another in this phase. By the time the FR market has entered its late phase, the impression prevails that it has reached a "rest point" that is periodically perturbed by small changes in the physical environment.

Summary. The early phase of the FR market is characterized by heterogeneous bidding behavior and frequent and sizeable adjustments of bids. During the middle and late phases, bids grow closer and the frequency and size of adjustments to bids falls.

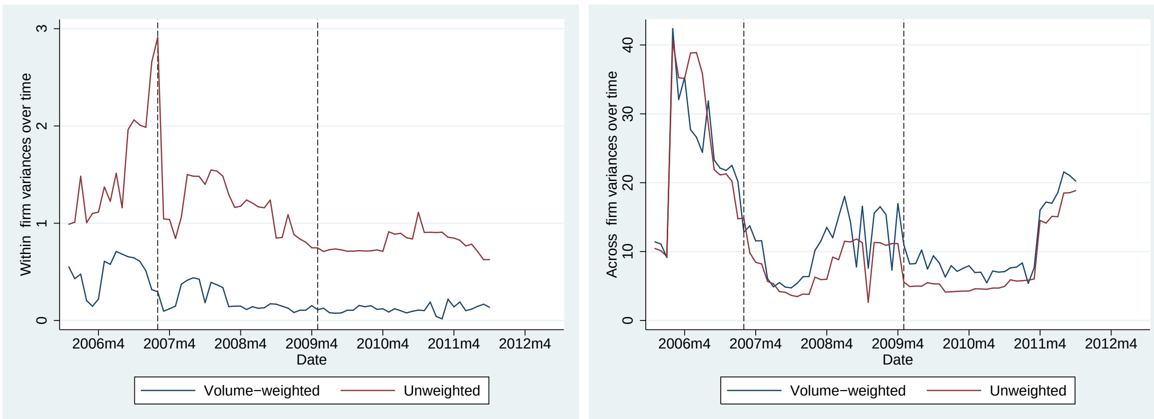

Figure 12: Quantity-weighted and unweighted variance in bids within a firm (left panel) and across firms (right panel). The right panel shows quantity-weighted variance across firms in the quantity-weighted mean firm bids and the unweighted variance across firms in the unweighted mean firm bids. {bidvarianc

In the early phase firms had no prior experience of bidding in this market. One may therefore expect the firms who think the market is a profit opportunity to experiment with their bids. This view is consistent with a comment by Ian Foy whose e-mail states: "The initial rush by market participants to test the waters having no history to rely upon; to some extent it was guess work, follow the price of others and try to figure out whether you have a competitive edge." Apparently the different firms pursue different strategies with at least some firms responding to rivals' experiments. As a result a model able to explain bidding behavior in this period is likely to have to allow firm to consider the gains from alternative experiments in a competitive environment; a task beyond the scope of this paper.

We view the middle or falling-price phase as a period of firms learning about how best to maximize current profits. That is, we treat the middle phase as a period dominated by firms bidding to "exploit" perceived profit opportunities rather than to experiment. Section 5 analyzes this phase by integrating some familiar learning models.

Finally, we view the late or stable-price phase as the FR market having reached an understanding of the behavior of competitors, the resulting allocation of FR, and the likely impact of changes in the physical environment. As a result, firms are able to adjust with quick small changes to the perturbations which occurred in the late phase.

Figure 13: Quantity-weighted average bids of the largest power stations by month. November 2005 – February 2007 (top left panel), March 2007 – May 2009 (top right panel), and June 2009 – October 2011 (bottom panel). Stations ranked by revenue in the FR market during early and middle phases. Bids are censored above at £10/MWh to improve visual presentation.

{top8_1}

Middle Period: Process of Convergence.

- "Bubble bursts". Barking and Peterhead undercut and gain share from Drax who then counters and prices head down.
- Attempts to coordinate on high price fail (Drax, Barking), & bids converge to those of low bidding firms.
- Big 8 bids begin $\in [4, 9]$ ends $\in [4, 5.5]$ (figure 11).
- **Dramatic** decrease in inter-firm variance (figure 9), and intra-firm variance almost goes to zero (figure 10).
- Less bid changes/month (3 vs 4 out of 10 per month) and smaller changes than in prior period ($\approx$ 1 vs 2 $\mathcal{L}$/mwh) .

Final Period

- This is the period with environmental changes.
- Bid changes still 3 out of 10/month (as in period 2).
- But bid changes are smaller $\approx .5\mathcal{L}$/Mwh, and are no longer primarily in one direction.
- All bids stay $\in [4, 5.5]$.
- **Possible takeaway:** environmental changes' impacts small relative to those generated by learning about competitors' play and/or demand responses to bid changes.

General: Three Periods

- Probability of a bid change:
  Period 1 $\gg$ Period 2 $\approx$ Period 3.
- Magnitude of bid change:
  Period 1 $>$ Period 2 $>$ Period 3.
- Period 1: bids increase with high inter-firm variance.
- Period 2: high bids fall causing dramatic fall in this variance.
- Period 3: bids cycle in a relatively narrow range.

Why an analytic framework?

- Analyze relationship between bids and profits over time.
- To construct profits from alternative bids we need an FR specific
    - demand system, and
    - cost function.

Demand.

- Use all generators of top 10 firms.
- Other generators in "outside good". Monthly dummies pickup "its bid".

Model

- RHS observables in addition to bid (our variable of interest):
  (i) generating unit and month FEs,
  (ii) fraction of time fully or part loaded in BM market,
  (iii) providing positive contracted FFR.
- Use logits (explain shares, no need for market size).
- Allow AR(1) disturbance. Innovation assumed not known when bid is made.

$$s_{j,t} = \frac{e_{j,t} exp[\alpha lnb_{j,t} + \beta x_{j,t} + \gamma_j + \mu_t + \zeta_{j,t}]}{1 + \sum_k e_{k,t} exp[\alpha lnb_{k,t} + \beta x_{k,t} + \gamma_k + \mu_t + \zeta_{k,t}]}.$$

Zero Shares (the $e_{j,t}$)

- About 20% bid high enough never to be called.
- Logit can not approximate zero shares so we have a separate "in-market" equation: bid above highest bid (almost) ever accepted.
- Assume error in "eligibility" uncorrelated with innovation. in demand. Appendix generalizes using a joint normal specification; but it doesn't change estimates of interest.

Demand Estimation Results (tables 2 & 3)

- Estimated using a quasi first difference (NLLS).
- Significant coefficients of expected sign (in addition to negative bid; part loaded $>$ fully loaded $>$ 0, FFR $<$ 0).
- AR(1): highly significant $\rho = .41$.
- Predicted shares (integrating out empirical distribution of $\zeta_{j,t}$) vs actual share; $R^2 = .67$.
- Figure 12: Fit over time (shows that we are doing better than just using FE).

Table 3: Demand System Estimates

{demandsyst

	Market Share		Eligibility
	OLS	NLLS	ML
Log bid	-1.648***	-1.614***	
	(0.132)	(0.119)	
Fully loaded	1.666***	1.949***	2.501***
	(0.220)	(0.182)	(0.355)
Part loaded	2.111***	2.234***	2.168***
	(0.156)	(0.139)	(0.335)
Positive FFR volume	-0.794***	-0.587**	-0.500
	(0.200)	(0.245)	(0.461)
Unit and Month FE	yes	yes	yes
ρ	–	0.41	–
s.e. ρ	–	0.03	–
R^2 (in shares)	0.49	0.66	–
N	3831	3509	5175

In the first two columns, the dependent variable is the log ratio of the share to the outside good share. In the last column it is an indicator for eligibility. The second market share specification allows for an AR(1) process in the error term, and we estimate the quasi-first-differenced equation by non-linear least squares (we provide an estimate of the autocorrelation coefficient ρ and the standard error of that estimate). The R^2 measure reported is for the fit of predicted versus actual shares (again omitting zero-share observations). Standard errors are clustered by bmunit. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

operating range. The coefficient on positive FFR volume in $x_{j,t}$ is negative and significant, indicating that a BM unit has a smaller share of the MFR market if it is already under contract with NG, also as expected. Finally, the NLLS estimates from equation (4) in the second column of Table 3 provide evidence of persistence in the unobservable characteristics $\xi_{j,t}$ as the $AR(1)$ coefficient ρ is positive and significant.

The third column of Table 3 shows ML estimates from equation (2). They are in line with our logit model for market shares. In particular, the coefficients on fully loaded and part-loaded are positive and significant, indicating that a BM unit is more likely to be eligible for providing FR services if it is up and running.

To assess goodness of fit, we predict the market share of BM unit j in month t conditional on $s_{j,t} > 0$. To do so, we sample independently and uniformly from the empirical distribution of residuals $\hat{\xi}_{j,t}$ for the OLS specification in equation (3) and from the empirical distribution

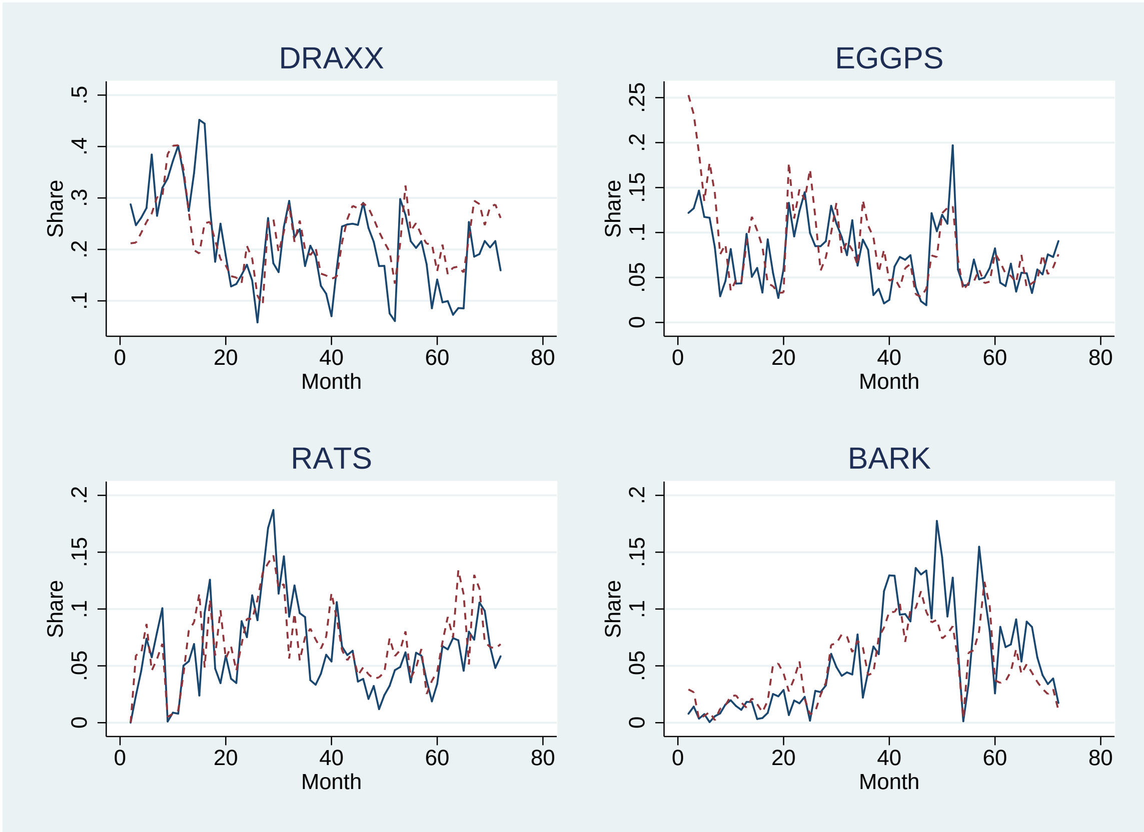

Figure 14: Goodness of fit. Realized (blue, solid) and predicted (red, dashed) market share by month for the four largest power stations Drax (top left panel), Eggborough (top right panel), Ratcliffe (bottom left panel), and Barking (bottom right panel). {fit_pic}

where M_t is market size in month t and our notation emphasizes that the market share of BM unit j in month t depends on the bids b_t, characteristics x_t and ξ_t, eligibilities e_t of all BM units, as well as on the true parameters θ_0 of the demand system. In contrast to market share, market size M_t is independent of bids b_t because the response requirement NG is obligated to satisfy is exogenously determined by government regulation as a function of the demand for electricity.

We estimate the marginal cost $c_i = (c_j)_{j \in \mathcal{J}_i}$ for the BM units that are owned by firm i from the bidding behavior of the firm in the late or stable-price phase of the FR market from June 2009 to October 2011. We maintain that a firm's bidding behavior stems from the firm "doing its best" in the sense of choosing its bid to maximize its expected profit conditional on the information available to it. More formally, the bids $b_{i,t}$ of firm i in month $t \geq 44$

Table 3: Demand System Estimates

{demandsyst

	Market Share		Eligibility
	OLS	NLLS	ML
Log bid	-1.648***	-1.614***	
	(0.132)	(0.119)	
Fully loaded	1.666***	1.949***	2.501***
	(0.220)	(0.182)	(0.355)
Part loaded	2.111***	2.234***	2.168***
	(0.156)	(0.139)	(0.335)
Positive FFR volume	-0.794***	-0.587**	-0.500
	(0.200)	(0.245)	(0.461)
Unit and Month FE	yes	yes	yes
ρ	–	0.41	–
s.e. ρ	–	0.03	–
R^2 (in shares)	0.49	0.66	–
N	3831	3509	5175

In the first two columns, the dependent variable is the log ratio of the share to the outside good share. In the last column it is an indicator for eligibility. The second market share specification allows for an AR(1) process in the error term, and we estimate the quasi-first-differenced equation by non-linear least squares (we provide an estimate of the autocorrelation coefficient ρ and the standard error of that estimate). The R^2 measure reported is for the fit of predicted versus actual shares (again omitting zero-share observations). Standard errors are clustered by bmunit. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

operating range. The coefficient on positive FFR volume in $x_{j,t}$ is negative and significant, indicating that a BM unit has a smaller share of the MFR market if it is already under contract with NG, also as expected. Finally, the NLLS estimates from equation (4) in the second column of Table 3 provide evidence of persistence in the unobservable characteristics $\xi_{j,t}$ as the $AR(1)$ coefficient ρ is positive and significant.

The third column of Table 3 shows ML estimates from equation (2). They are in line with our logit model for market shares. In particular, the coefficients on fully loaded and part-loaded are positive and significant, indicating that a BM unit is more likely to be eligible for providing FR services if it is up and running.

To assess goodness of fit, we predict the market share of BM unit j in month t conditional on $s_{j,t} > 0$. To do so, we sample independently and uniformly from the empirical distribution of residuals $\hat{\xi}_{j,t}$ for the OLS specification in equation (3) and from the empirical distribution

Cost Estimation: Assumptions

- Main cost is wear and tear of machines and heat rate loss. We assume a constant mc (do not observe the amount of times in a day that it generators are called).
- Recall that the bids in the last period are more stable, despite environmental changes.
- Accordingly we assume that in this period firms chose bids to maximize their perceptions of expected profits.

This assumes away

- experimentation (seems consistent with the data in this period), and
- most collusive models.

Collusion? Figure 14 which provides the distribution of the within and between firm correlations in:
(i) generators bids and in
(ii) the direction of changes in those bids.
They are both centered close to zero (and we expect some positive correlation due to demand shocks).

Online appendix shows that if they were maximizing joint profits, or the top three firms were maximizing their joint profits, estimates of mc are negative.

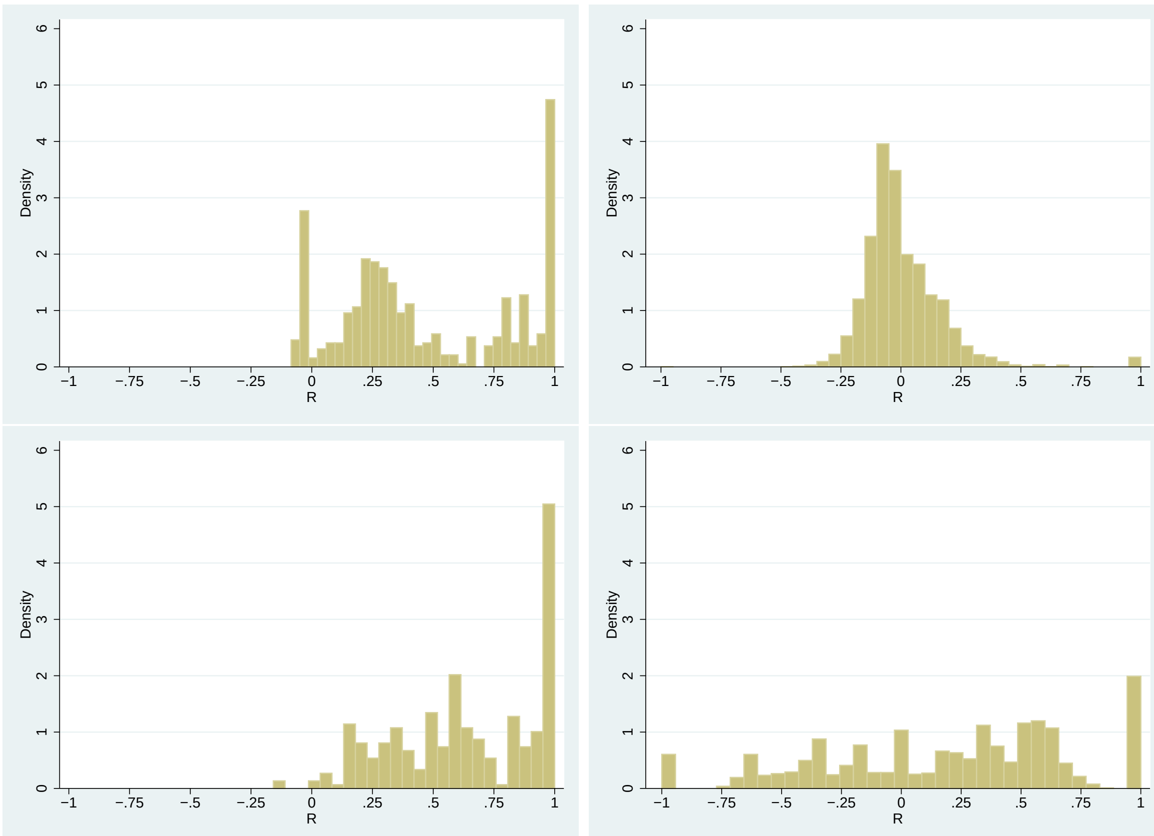

Figure 14: Top left is within-firm correlation in bid changes; top right is across-firm correlation in bid changes; bottom left is within-firm correlation in direction of change (conditional on both changing); bottom right is across-firm correlation in directions.

its bid (conditional on both BM units in the pair changing their bids). In Figure 14 we plot the distribution of correlation coefficients separately for BM units owned by the same firm ("within firm", left panels) and for BM units owned by different firms ("across firms", right panels). Note that we expect some across-firm correlation in both the timing and the direction of bid changes due to common shocks to demand.

The within-firm correlations for the timing and direction of bid changes in the left panels are positive and substantial. This reinforces our contention that decisions are centralized at the level of the firm rather than made at the level of the BM unit. The right panels show correlations pretty much evenly distributed around zero, consistent with independent decision making across firms.

Our second approach is more direct: we assume particular collusive arrangements and infer

Formal Model.

If $\mathcal{E}_{(\cdot)}(\cdot|\Omega_{i,t-1})$ indicates agent's perceptions, our assumption is that in the last period firms chose their bids to

$$\max_{b_{i,t}} \mathcal{E}_{b_{-i,t},\xi_t,e_t,\theta}\left[\sum_{j\in\mathcal{J}_i}(b_{j,t}-c_j)\,M_t s_j(b_t,x_t,\xi_t,e_t;\theta)\middle|\Omega_{i,t-1}\right],$$

Sources of uncertainty that a firm faces:

- Uncertainty about its rivals' bids $b_{-i,t}$, and
- Demand uncertainty generated by the fact that
  - the parameters θ of demand may not be known (quantity response to bids), &
  - the realizations of ξ_t and e_t (disturbances in demand and acceptance probability).

Behavioral Assumptions for Estimating Costs.

To maximize need to endow firms with perceptions on costs and the sources of demand uncertainty.

- $c_i \in \Omega_{i,t-1}$ (but not necessarily c_{-i}).
- Our estimated $\hat{\theta} \rightarrow_{a.s.} \theta_0$.
- By the final period the firms' perceptions are on average correct (do not need to specify precisely what those perceptions are – the topic of the next section).

These assumptions, standard identification conditions, and a ULLN imply that $\forall k \in \mathcal{J}_i$ and $\forall i$

$$p \lim_{T\to\infty} \left(\frac{1}{T-44} \sum_{t=44}^{T} \left[M_t s_k(\cdot,\hat{\theta}) + \sum_{j\in\mathcal{J}_i} (b_{j,t} - c_j)\, M_t \frac{\partial s_j(\cdot,\hat{\theta})}{\partial b_{k,t}}. \right] = 0 \right).$$

- Then consistent estimates of marginal costs can be backed out of the average of the f.o.c.'s for bids in the final period.
- This requires only a simple matrix inversion.
- Note: though we did not need to specify how perceptions are formed (the subject of the next section), we do need to assume that the average of the f.o.c.'s after $t = 44$ converges to zero; the paper gives primitive condition which insures this happens.

Cost Estimates: Results: Table 4.

- Ave.=1.41 pounds/MwH. Range of values=[1.04,1.6].
- Standard deviation across BM units .66.
- Within station the generators are almost always of the same type and vintage.
- Within station standard deviation .01 to .06.
- Standard errors of estimates small (about .04).
- "Cost reflective" price pre-deregulation, 1.7, and they presumably had to give them a markup over costs.

Table 4: Cost estimates for the top 8 stations (by total revenue)

{costtable}

Station	# Units	Fuel	Vintage	Mean	Std. Dev. (within station)
Barking	2	CCGT	1994	1.2	.01
Connah's Quay	4	CCGT	1996	1.04	.03
Cottam	4	Coal	1969	1.35	.04
Drax	6	Coal	1974	1.06	.04
Eggborough	4	Coal	1968	1.53	.06
Peterhead	1	CCGT	2000	1.54	0
Ratcliffe	4	Coal	1968	1.33	.06
Seabank	2	CCGT	1998	1.59	.01

Summary statistics on the unit-specific cost estimates derived from solving the firm first order condition arising from the demand system, reported as the within-station average cost and standard deviation in costs.

nor do we place any restrictions on the learning process (e.g. that players are Bayesians, and update their perceptions accordingly). Still, which learning processes have this convergence property is a complicated problem, and one we leave for future research.

Results: estimates. The average of the marginal cost c_j that we estimate for the $J = 72$ BM units owned by the ten largest firms is £1.40/MWh, with a standard deviation of £0.66/MWh across BM units.[14] The estimates are reasonably precise, with an average standard error of £0.04/MWh. By comparison, pre CAP047 the "cost reflective" administered price was around £1.7/MWh.[15] Since we expect some markup to be built into the administered price, the marginal cost we recover is in the right ballpark.

Table 4 shows the average marginal cost for the BM units belonging to the eight largest power stations. They are quite reasonable and vary between £1.04/MWh and £1.6/MWh across stations. The standard deviation of marginal cost within a station is very small, on the same order as the standard error of the estimates. Most of the variation in marginal cost is therefore across stations.

Table 5 shows the result of regressing marginal cost on the characteristics of the BM unit.

[14]Because one BM unit has zero share during the late phase, we impute its marginal cost with that of the other BM unit in the same power station.

[15]We have two sources: Figure 2 and a document prepared just prior to CAP047 by NG for Ofgem, the government regulator (`www.ofgem.gov.uk/ofgem-publications/62273/8407-21104ngc.pdf`). It states in paragraph 5.3 that the holding payment is "of the order of £5/MWh" for the bundle of primary, secondary, and high response, implying an average of £1.67/MWh per type of FR.

Residuals from FOC by period: Figure 15

- Note the difference in axis for the two graphs (scale changed by a factor of 5).
- Average in three periods: .002, -.0004, ≈ 0 (by construction).
- Standard deviations across generators also falls: .006, .002,.001.
- Table 6: R^2 between (estimated) residual in t and $t-1$ drops over time.
- Bids do not change every period, and even in the last period this generates serial correlation.
- When bid changes; serial correlation ≈ 0 in last period,

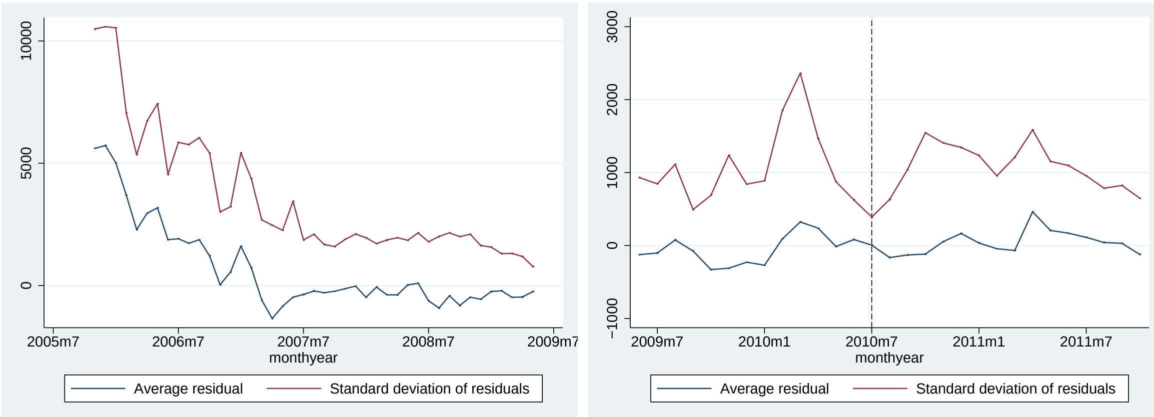

Figure 15: Average and standard deviation of residuals during the early and middle phases (left panel) and late phase (right panel). In the right panel, the dotted line indicates July 2010, when FFR volume nearly doubles.

{residual_s

magnitude smaller than in the early phase of the FR market. Apparently in the late phase firms adjust to changes in their environment quite quickly. This is quite different than the behavior in the early phases, although at that time firms had little experience in predicting their competitors bids and may not have known certain parameters of the demand process.

We also examine whether the residuals are autocorrelated. The first three columns of Table 6 display the coefficients from separate regressions of the residual on its lagged value for each of the three phases of the FR market, including BM-unit fixed effects in all regression. In the last three columns we further restrict attention to observations in which the BM unit's bid changed between months.

We find significant autocorrelation in all regressions but the last. Assuming our specification and cost estimates are correct, this indicates that some firms are making systematic mistakes. This may reflect persistent differences between a firm's perceptions of its expected profits and the reality. This makes particular sense in the early and middle phases of the FR market where firms had little experience and behaved quite differently.

At the same time, it is striking how the R^2 falls over the three phases of the FR market, indicating that the lagged value explains progressively less of the variation in the residual. In the third phase, we find that we find significant autocorrelation using all observations (column 3), but that it essentially disappears when we restrict attention to observations in which the BM unit's bid changed between months (column 6). Our interpretation of this

Table 6: Autocorrelation in residuals

{autocorrel

	Early	Middle	Late	Early	Middle	Late
		All			Bid changes only	
Lagged residual	0.542***	0.343***	0.445***	0.389***	0.126***	0.029
	(0.087)	(0.050)	(0.063)	(0.059)	(0.042)	(0.080)
R^2	0.63	0.48	0.20	0.74	0.38	0.08
N	1080	1931	2088	355	449	401

The dependent variable is the residual in the FOC at the estimated costs. Controls are the lagged residual and unit fixed effects. The regressions with bid changes only include only observations in which the unit's bid was different from its bid in the previous period. Standard errors are clustered by unit. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

is that by the end firms have reasonably accurate perceptions, and when they choose to update their bids they do so in a way that accounts for the information contained in the lagged residual.

Collusion? We now briefly examine the possibility of collusion between the firms in our data. To begin, we use first-order conditions that reflect different collusive arrangements to re-estimate the marginal cost c_j for the $J = 72$ BM units owned by the ten largest firms during the stable-price phase. If the top 10 firms colluded and therefore maximized the combined profits of all their BM units, then the implied cost is £-9.8/MWh on average across BM units. If the top 3 firms colluded, then the implied cost is £-0.25/MWh. We estimate cost to be negative because demand is relatively inelastic, and so rationalizing their bids in the face of increased market power requires firms to have low costs.

Next we look for coordination in the timing and direction of bid changes across BM units, as this could be a sign of collusion being established or breaking down. To capture timing, we define a dummy for BM unit j changing its bid between months $t-1$ and t and, to capture direction, another dummy for the BM unit increasing its bid. We compute all pairwise correlations between BM units in the dummy for a BM unit changing its bid and in the dummy for a BM unit increasing its bid (conditional on both BM units in the pair changing their bids). In Figure 16 we plot the distribution of correlation coefficients separately for BM units owned by the same firm ("within firm", left panels) and for BM units owned by different firms ("across firms", right panels). Note that we expect some across-firm correlation in both the timing and the direction of bid changes due to common shocks to demand.

Table 7: "Lost Profits" and "Bidding Errors".

Upper bound to lost profits assumes that firms maximize knowing
(i) what their competitors' bid and
(ii) the parameters of the demand function.
Computation does not require a learning model.

- Table 7: "Lost profits" are small in percentage terms; for 6 out of 10 firms $\leq 3\%$, always $\leq 8.2\%$ (and that was for Seabank whose profits were lowest).
- Actual vs optimal bid differences noticeably larger; 4 firms $\geq 20\%$, and the smallest is 8%.
- Though profits may not change much as a result of misperceptions during the learning process it is clear that bids do. Flat return function at optima making it difficult to learn?
- Reminiscent of Ackerloff and Yellen (1985): small individual deviations can lead to large changes in aggregate behavior.

Table 7: Profit Statistics

Rank	Firm name	Months Changed	Matched Direction (%)	Total Profit	Ex-Post Lost Profit (%)	Ex-Post Bid Diff.	Ex-Post Bid Diff. (%)
1	Drax Power Ltd.	23	80	68.4	2.03	1.1	20
2	E.ON UK plc	52	67	44.4	2.01	.76	14
3	RWE plc	15	86	25.1	3.06	.93	23
4	Eggborough Power Ltd	18	58	18.2	1.43	.84	17
5	Keadby Generation Ltd	17	80	14.9	5.12	1.03	19
6	Barking Power Ltd	57	49	11.8	6.17	.65	15
7	Jade Power Generation Ltd	15	54	10	4.74	1.24	24
8	SSE Generation Ltd	17	62	9.4	.47	.42	8
9	Seabank Power Ltd	9	89	9.1	8.21	1.14	20
10	Centrica plc	42	73	8.7	2.59	.92	18

Months changed indicates the total number of months in which a firm changed the bid of one or more of their units. Matched direction indicates the share-weighted percent of bid changes that are in the same direction as our estimated ex-post-optimal bid. Ex-post lost profit indicates the share-weighted percent increase in profit by taking the estimated ex-post-optimal bid instead of the actual bid. Ex-post bid difference indicates the share-weighted average absolute difference between the ex-post-optimal and actual bids, in absolute and relative terms.

necessarily optimal adjustments when they do. This combination of rational inattention (Sims 2003) and "satisficing" behavior (Simon 1955) may have been optimal along the observed path of play given the firms' various human resource and institutional constraints.

Yet, as Akerlof and Yellen (1985) have noted, even small departures from perfect rationality may lead to aggregate behavior that is quite different from equilibrium and thus hard to predict with equilibrium models. To investigate bidding incentives during the disequilibrium phase of the observed play, we need the learning models that we outline in Section 5 below.

Collusion? Before doing so, we briefly examine the possibility of collusion between the firms in our data. We try two different approaches. The first is to look for coordination in the timing and direction of bid changes across BM units, as this could be a sign of collusion being established or breaking down. To capture timing, we define a dummy for BM unit j changing its bid between months $t-1$ and t and, to capture direction, another dummy for the BM unit increasing its bid. We compute all pairwise correlations between BM units in the dummy for a BM unit changing its bid and in the dummy for a BM unit increasing

Learning.

Time period for analysis.

- Period 1. Heterogenous behavior, experimentation (and we don't have a model for either).
- Look at Periods 2 & 3. No evidence of collusion (Figure 14).

Assume they know from pre-deregulation experience.

- c_i (but possibly not c_{-i}),
- the $AR(1)$ process generating ζ_t,
- the objective p.d.f of e_t, &
- the BM-unit F.E.'s (the γ_j or NG's preferences for BM units).

Unlikely to know.

- Other firms' bids
- share response to bid (our "α") and possibly response to time varying characteristics (β) and month fixed effects (μ_t) .

Learning Models.

About competitors' play.

- Use variants of fictitious play: The distribution of b_{-i} is geometric weighted average over observed *vectors* of competitors' prior play.
- Weight for $t - \tau \propto \delta^{\tau}$. Consider models $F(\delta)$ for
  - $\delta = 1$ (equal weights for past periods),
  - $\delta = .5$,
  - $\delta = 0$ (best response).

About parameters of demand system.

- Statistical (or "adaptive") learning. NLLS estimates using only data available when bid was made.
  - $A(z)$; statistical learning about z, all other parameters known: $z \in \Big\{\alpha, (\alpha, \mu), (\alpha, \beta), (\alpha, \mu, \beta)\Big\}$.
  - $A(\emptyset)$; all parameters known.

- Models tried: $\Big(equilib., F(\delta)\Big) \otimes \Big(A(z), A(\emptyset)\Big)$.

Comparison Model: Complete Information Nash Equilibrium.

$\forall (i,t)$ bids solve the system of J_i equations

$$0 = E_{\xi_t, e_t}\Big[M_t s_k(b_t^*, x_t, \xi_t, e_t; \theta_0) +$$

$$\sum_{j \in \mathcal{J}_i} \left(b_{j,t}^* - c_j \right) M_t \frac{\partial s_j(b_t^*, x_t, \xi_t, e_t; \theta_0)}{\partial b_{k,t}} | M_t, x_t, \xi_{t-1}, \theta_0, c \Big] \quad \forall$$

• More stringent assumptions then used to estimate cost. $E_{\xi_t, e_t}[\cdot|\cdot]$ in contrast to prior $\mathcal{E}_{b_{-i,t}, \xi_t, e_t, \theta_t}[\cdot|\cdot]$ indicates that the expectation uses; the objective probability distribution of ξ_t conditional on ξ_{t-1}, the objective distribution of e_t, and θ_0.

• To make estimation practical we replace ξ_{t-1}, θ_0, and c by our estimates of those objects, evaluate the expectation operator using Monte Carlo integration, and numerically solve the system. Always only found one equilibria (but do not know of a uniqueness proof).

Measures of predictive accuracy.

- **Predict NG's costs.** The average (over time) of the share weighted absolute difference between the predicted and actual bids (predicted weights for models, actual weights for data).

$$CE^{(\delta,y)} = \frac{1}{T}\sum_t CE_t^{(\delta,y)} = \frac{1}{T}\sum_t \left| \sum_j \left(b_{j,t}^{(\delta,y)} - b_{j,t} \right) \frac{s_{j,t}}{\sum_j s_{j,t}} \right| .$$

- The mean square prediction error (traditional fit measure).

$$MSE^{(\delta,y)} = \frac{\sum_t \sum_j \left(b_{j,t}^{(\delta,y)} - b_{j,t} \right)^2 s_{j,t}}{\sum_t \sum_j s_{j,t}} .$$

Prediction horizons.

- Single period predictions; firms use $(b_\tau)_{\tau \le t-1}$, $(s_\tau)_{\tau \le t-2}$.
- Multi-period predictions: (i) when sampling competitors' bids in subsequent periods, we sample the predicted bids, and (ii) when estimating $\hat{\theta}_t$ the data set now includes predicted bids and market shares.

Sequential Estimates of Demand Parameters.

Figure 14. Over time they converge to the full sample estimates (by construction).

- Price sensitivity parameter (α).
  - When we estimate it alone it is within 2% of the final number at the start of period 2.
  - α estimates "deteriorate" when also estimating other parameters.
- Other parameters
  - Fully loaded and part loaded gradually approach, and
  - coefficient of FFR stabilizes after they start using more FFR.

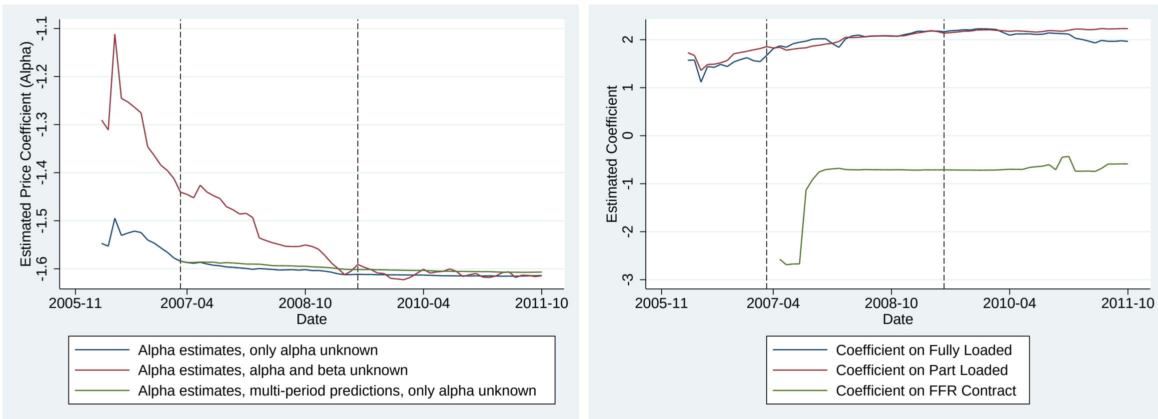

Figure 14: Sequential estimates $\hat{\alpha}^{(t)}$ under $A(\alpha)$ and $A(\alpha, \beta)$ (left panel) and $\hat{\beta}^{(t)}$ under $A(\alpha, \beta)$ (right panel) by month.

5. $A(\emptyset)$: Fix all parameters θ at the full-sample estimates in Table 3.

Figure 14 illustrates the sequential estimates. The left panel shows $\hat{\alpha}^{(t)}$ obtained alternatively under $A(\alpha)$ and $A(\alpha, \beta)$ and the right panel shows $\hat{\beta}^{(t)}$ obtained under $A(\alpha, \beta)$. Over time the sequential estimates by construction approach the full-sample estimates obtained in Section 4.1. The sequential estimates for the price sensitivity parameter α start out small in absolute value and gradually decrease. Until the late phase of the FR market, $\hat{\alpha}^{(t)}$ is considerably smaller under $A(\alpha)$ than under $A(\alpha, \beta)$; in fact the sequential estimates produced by $A(\alpha)$ is within 2% of the full sample estimate by the start of the second phase. The sequential estimate $\hat{\beta}^{(t)}$ on positive FFR volume is far from its full-sample estimate until month 25 (because there is little FFR volume in the early part of the data), but stabilizes thereafter. The sequential estimates $\hat{\beta}^{(t)}$ on part-loaded and fully loaded are much less volatile and gradually trend towards their full-sample estimates.

Predictions. We combine fictitious play with adaptive learning to make predictions. We measure fit by comparing our predictions to the observed bids. We make two kinds of predictions. The first is one-period predictions: for each month t during the second and third phases of the FR market, we take the data available to the firms at the time they bid (which includes bids $(b_r)_{r \leq t-1}$ and market shares $(s_r)_{r \leq t-2}$) and predict their bids. This corresponds to the thought experiment of predicting the next move of a player in a game and is analogous to the one-step-ahead predictions used to assess predictive accuracy in the experimental and

Predictions for Middle (falling-price) Period.

Single period prediction results (table 9).

Results on $F(\delta)$.

- The preferred model for perceptions about rivals' bids is either $F(0)$ (best response) or $F(0.5)$ [either accuracy measure].
- We searched over $\delta \in [0, 1]$. Smallest MSE at $\delta = 0.3$, but little difference for $\delta \in [0, 0.6]$.
- Firms weight competitors' more recent observations disproportionately when forming perceptions b_{-i}.

Results on $A(z)$

- The minimum MSE model is either $A(\alpha)$ or $A(\emptyset)$; but $A(\alpha)$ does noticeably better on predicting the average bid.
- Extending adaptive learning to additional parameters noticeably worsen both accuracy measures.

Table 9: Middle Phase: Prediction Error in NG's Monthly Cost ($CE^{(\delta,y)}$) *

	Single Period Prediction					Multi-Period Prediction	
	$A(\emptyset)$	$A(\alpha)$	$A(\alpha,\mu)$	$A(\alpha,\beta)$	$A(\alpha,\beta,\mu)$	$A(\emptyset)$	$A(\alpha)$
F(0)	0.28	0.23	1.31	0.57	0.95	0.44	0.31
F(0.5)	0.28	0.24	1.33	0.59	0.97	0.41	0.29
F(1)	0.46	0.39	1.06	0.39	0.69	0.54	0.41
Eq.	0.46	0.37	1.91	0.73	1.33	-	-

* $F(x)$ denotes fictitious play with $\delta = x$. $A(x)$ denotes adaptive learning where all parameters except x are known and x is sequentially estimated. Single period predictions take the data from period $t-1$ and predict bids in period t. Multi period predictions take the bid data from the period before each phase, and make a series of bid predictions for the entire phase, taking the path of all other covariates as known.

or the weighted average of the square prediction error $\left(b_{j,t}^{(\delta,y)} - b_{j,t}\right)^2$, where the weight $\frac{s_{j,t}}{\sum_t \sum_j s_{j,t}}$ is the share of BM unit j in month t in the total share of the inside goods across months. We again use MSE^* to denote the corresponding measure for the complete information Nash equilibrium.

5.4 Results

We begin with discussing how well different learning models fit the data from the middle phase of the FR market and then move on to the late phase.

Middle or falling-price phase. Table 9 provides the average absolute prediction error ($CE^{(\delta,y)}$ and CE^*) for the middle phase and Table 10 provides the mean square prediction error ($MSE^{(\delta,y)}$ and MSE^*). The entries in the table are arranged to reflect our two-way classification of learning models: the first three rows specify the fictitious play model used to form perceptions of rivals' bids, followed by the complete information Nash equilibrium, and the columns specify the adaptive learning model used to form perceptions of the demand parameters. The five columns on the left pertain to single-period predictions and the two columns on the right to multi-period predictions.

Starting with the single-period predictions, the 2×2 sub-matrix in the top left corner of Tables 9 and 10 have the smallest values. Regardless of whether we are concerned with $CE^{(\delta,y)}$ or $MSE^{(\delta,y)}$, the preferred model for perceptions about rivals' bids is either $F(0)$ (full decay beyond the immediate past, or adaptive best response) or $F(0.5)$ (intermediate decay).

Multi Period Prediction Results (table 10).

- As expected fit worsens, but similar qualitative results to single period: i.e. $\left(A(\alpha), F(0) \text{ or } F(.5)\right)$ do best.
- The multi-period predictions accentuate the importance of allowing for adaptive learning (small differences get magnified).

Table 10: Middle Phase: Mean Squared Error of Bid Predictions ($MSE^{(\delta,y)}$) *

	Single Period Prediction					Multi-Period Prediction	
	$A(\emptyset)$	$A(\alpha)$	$A(\alpha,\mu)$	$A(\alpha,\beta)$	$A(\alpha,\beta,\mu)$	$A(\emptyset)$	$A(\alpha)$
F(0)	1.22	1.20	3.81	1.95	2.78	1.31	1.22
F(0.5)	1.22	1.20	3.91	1.99	2.85	1.26	1.19
F(1)	1.41	1.33	2.80	1.52	2.03	1.46	1.33
Eq.	1.36	1.26	6.70	2.29	4.13	-	-

* $F(x)$ denotes fictitious play with $\delta = x$. $A(x)$ denotes adaptive learning where all parameters except x are known and x is sequentially estimated. Single period predictions take the data from period $t-1$ and predict bids in period t. Multi period predictions take the bid data from the period before each phase, and make a series of bid predictions for the entire phase, taking the path of all other covariates as known.

The preferred model for perceptions about demand is either $A(\emptyset)$ (no demand uncertainty) or $A(\alpha)$ (uncertainty about the price sensitivity parameter α). Both measures of fit noticeably deteriorate under the adaptive learning models $A(\alpha,\beta)$, $A(\alpha,\mu)$ and $A(\alpha,\beta,\mu)$ that presume that firms are uncertain about additional demand parameters.

When we searched for a point estimate of the decay parameter δ in the fictitious play model, we obtained the smallest mean square prediction error at $\delta = 0.3$. However, there was very little difference in the mean square prediction error when δ took on values between 0 and 0.6. So we conclude only that the data from the middle phase of the FR market appears to favor a fictitious play model in which firms rely disproportionately on more recent observations to form beliefs about rivals' bids.

The single-period predictions in the five columns on the left of Tables 9 and 10 seem to give only a slight edge to $A(\alpha)$ over $A(\emptyset)$ (as is suggested by Figure 14). This edge, however, is accentuated in the multi-period predictions in the two columns on the right. The multi-period predictions compound differences between the sequential estimates $\hat{\alpha}^{(t)}$ of the price sensitivity parameter α that underlie $A(\alpha)$ and the full-sample estimate that underlies $A(\emptyset)$. More surprising is how well the multi-period predictions using $A(\alpha)$ do. They have essentially the same mean square prediction error as the single-period predictions using $A(\alpha)$ and do just 25% to 35% worse in terms of the average absolute prediction error. In contrast, the multi-period predictions using $A(\emptyset)$ do 5% to 10% worse than the single-period predictions using $A(\emptyset)$ in terms of the mean square prediction error and 50% worse in terms of the average absolute prediction error.

Perhaps the most striking feature of Tables 9 and 10 is that the fictitious play models $F(0)$

Compare to Nash Equilibrium.

Table 11 regressions.

- Learning models, do noticeably (and statistically significantly) better.
- The difference between the learning model falls over time.

Figure 15: Compares predictions for costs.

- As expected single period do better than multi-period. However even multi-period prediction do better that the equilibrium model; especially early on.
- Reason: the equilibrium bids go almost immediately go to something that looks like our rest point, and the data does not.

Conclude. Learning models do better in that they mimic the observed slow movement in the data towards equilibrium.

Table 11: Middle Phase: Differences Between the Learning Model with $(\delta, y) = (0, \emptyset)$ and Nash Equilibrium

	Dependent Variable			
	$CE_t^{(0,\emptyset)} - CE_t^*$		$\left(b_{j,t}^{(0,\emptyset)} - b_{j,t}\right)^2 - \left(b_{j,t}^* - b_{j,t}\right)^2$	
Constant	-0.182	-0.357	-0.144	-0.528
	(0.037)	(0.066)	(0.079)	(0.200)
Time		0.013		0.027
		(0.004)		(0.011)
N	27	27	1470	1470

An individual observation for the CE_t regressions is a month that is formed from aggregating across BM units with weights as specified in section 5.3. For the third and fourth column, an individual observation is a BM unit in a given month; these regressions are estimated via share-weighted least squares and the standard errors are clustered at the month level. The regressors are a constant and a linear time trend; time=1 for the first month of the phase.

and $F(0.5)$ outperform the complete information Nash equilibrium. This is more pronounced for the average absolute prediction error, but clearly noticeable for both measures of fit. Table 11 further illustrates just how different the predictions from the learning and equilibrium models are. In this table, we specifically compare $A(\emptyset)$ and $F(0)$ (i.e. adaptive best response with the full sample estimate) and Nash equilibrium. To do this, we run regressions of $CE_t^{(0,\emptyset)} - CE_t^*$ on a constant, and then a constant and time trend (columns 1 and 2). Columns 3 and 4 show the corresponding regressions when the dependent variable is instead $\left(b_{j,t}^{(0,\emptyset)} - b_{j,t}\right)^2 - \left(b_{j,t}^* - b_{j,t}\right)^2$, the difference in squared prediction error. In those regressions, we share-weight the observations to give greater importance to precise bid predictions for units with large market share.

We find statistically and economically significant differences between the learning model and equilibrium predictions, both with and without a time trend. Moreover, the regressions that include a time trend indicate that the difference starts out large and declines over the middle phase of the FR market, with both the initial difference and the subsequent decrease clearly significant. We find quantitatively similar statistically and economically significant differences between the predictions of the other leading learning models (i.e. $(\delta, y) = (\emptyset, 0.5), (\alpha, 0), (\alpha, 0.5)$) and equilibrium.

Figure 15 shows the time path of the price of FR paid by NG as predicted by alternative learning models and the complete information Nash equilibrium as well as the time path of

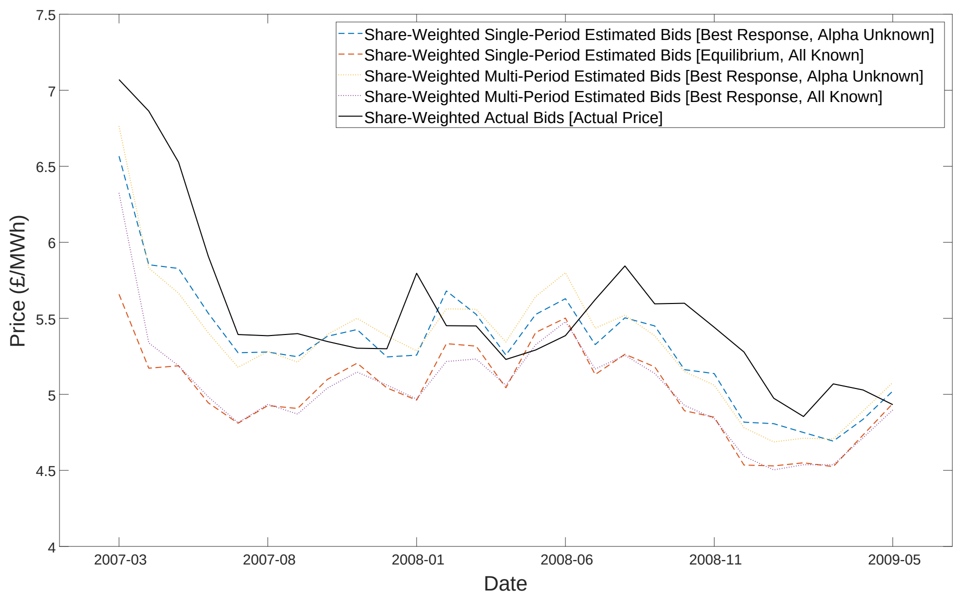

Figure 15: Predicted and actual FR price by month. Share-weighted average computed from bids predicted by fictitious play $F(0)$ with adaptive learning $A(\alpha)$ in both single and multi-period simulations, complete information Nash equilibrium without demand uncertainty $A(\emptyset)$, and actual bids. Middle phase.

the purple and yellow line is so notable indicates that allowing for learning about the price sensitivity parameter α is important for the multi-period predictions to come close to the actual price of FR.[34]

Late or stable-price phase. Tables 11 and 12 provide the average absolute prediction error $CE^{(\delta,y)}$ and CE^{*} and the mean square prediction error $MSE^{(\delta,y)}$ and MSE^{*} for the late phase of the FR market; they are analog to Tables 8 and 9 for the middle phase. Both measures of fit improve in the late phase for the alternative learning models and for the complete information Nash equilibrium. In particular, the mean square prediction error is a third or less of its value in the middle phase.

Recall that we estimate the marginal cost c by setting the time-average of the first-order conditions over the late phase to zero (see equation (8)). It may therefore not be a surprise

[34]Our discussion is about relative fit; all models underestimate NG's cost of procuring FR services during the middle phase.

Predictions for late (stable-price) period.

Table 13 for NG's costs, and 14 for MSE.

- MSE's: 1/3rd or less of those in the middle phase **despite** the environmental changes in this period.
- Best response still does better than equilibrium but the difference is now very small.
- In terms of MSE, it is difficult to tell the difference between any of the adaptive learning models (as expected, as they are all based on similar information).
- However, for cost estimates there is a preference for adaptive learning with all parameters unknown.
- When the environment is changing. it may well be that the best approximation to the demand system is also. Then an adaptive learning model may be needed.

Table 13: Late Phase: Prediction Error in NG's Monthly Costs ($CE^{(\delta,y)}$)*

	Single Period Prediction					Multi-Period Prediction	
	$A(\emptyset)$	$A(\alpha)$	$A(\alpha,\mu)$	$A(\alpha,\beta)$	$A(\alpha,\beta,\mu)$	$A(\emptyset)$	$A(\alpha)$
F(0)	0.21	0.21	0.25	0.25	0.18	0.24	0.31
F(0.5)	0.22	0.22	0.25	0.26	0.19	0.23	0.30
F(1)	0.33	0.33	0.36	0.36	0.27	0.24	0.32
Eq.	0.25	0.25	0.30	0.30	0.21	-	-

* $F(x)$ denotes fictitious play with $\delta = x$. $A(x)$ denotes adaptive learning where all parameters except x are known and x is sequentially estimated. Single period predictions take the data from period $t-1$ and predict bids in period t. Multi period predictions take the bid data from the period before each phase, and make a series of bid predictions for the entire phase, taking the path of all other covariates as known.

Table 14: Late Phase: Mean Squared Error of Bid Predictions ($MSE^{(\delta,y)}$) *

	Single Period Prediction					Multi-Period Prediction	
	$A(\emptyset)$	$A(\alpha)$	$A(\alpha,\mu)$	$A(\alpha,\beta)$	$A(\alpha,\beta,\mu)$	$A(\emptyset)$	$A(\alpha)$
F(0)	0.40	0.40	0.44	0.48	0.43	0.40	0.46
F(0.5)	0.41	0.41	0.45	0.49	0.44	0.40	0.45
F(1)	0.52	0.53	0.58	0.64	0.56	0.42	0.49
Eq.	0.42	0.42	0.49	0.52	0.42	-	-

* $F(x)$ denotes fictitious play with $\delta = x$. $A(x)$ denotes adaptive learning where all parameters except x are known and x is sequentially estimated. Single period predictions take the data from period $t-1$ and predict bids in period t. Multi period predictions take the bid data from the period before each phase, and make a series of bid predictions for the entire phase, taking the path of all other covariates as known.

from the model $(F(0), A(\emptyset))$. The fact that the difference between the red and the blue line is so notable indicates that allowing for learning about the price sensitivity parameter α is important for the multi-period predictions to come close to the actual price of FR.

Late or stable-price phase. Tables 13 and 14 provide the average absolute prediction error $CE^{(\delta,y)}$ and CE^{*} and the mean square prediction error $MSE^{(\delta,y)}$ and MSE^{*} for the late phase of the FR market; they are analog to Tables 9 and 10 for the middle phase. Both measures of fit improve in the late phase for the alternative learning models and for the complete information Nash equilibrium. In particular, the mean square prediction error is a third or less of its value in the middle phase.

Recall that we estimate the marginal cost c by setting the time-average of the first-order

Late Phase Comparison to Nash Predictions.

- Table 15. MSE differences are no longer statistically significant, and if we put in time dummies, neither are the cost differences.
- Figure 16, Cost predictions. In this last part of the period best response and equilibrium seem to be predicting each other (though best response does marginally better).

Does the fact that equilibrium does so much better in the last than the middle period Illustrate the difficulty of learning relative to that of adaptation to environmental change?

Table 13: Late phase: Comparison of learning model ($F(0)$,$A(\emptyset)$) and complete information Nash equilibrium

	$CE_t^{(0,\emptyset)} - CE_t^*$		$\left(b_{j,t}^{(0,\emptyset)} - b_{j,t}\right)^2 - \left(b_{j,t}^* - b_{j,t}\right)^2$	
Constant	-0.023	-0.007	-0.002	0.026
	(0.007)	(0.014)	(0.001)	(0.012)
Time trend		-0.002		-0.002
		(0.001)		(0.001)
N obs	29	29	1529	1529

In the first two regressions with $CE_t^{(0,\emptyset)} - CE_t^*$ as the dependent variable, the unit of observation is a month. In the last two regressions with $\left(b_{j,t}^{(0,\emptyset)} - b_{j,t}\right)^2 - \left(b_{j,t}^* - b_{j,t}\right)^2$ as the dependent variable, the unit of observation is a BM unit-month. In the latter, observations are weighted by $\frac{s_{j,t}}{\sum_t \sum_j s_{j,t}}$ as described in Section 5.3 and standard errors are clustered by month.

error of about 0.02 (compared to about 0.19 in the middle phase), and a precisely estimated average difference in mean square prediction error of about 0.002 (compared to about 0.19 in the middle phase).

The learning models continue to do slightly better than the equilibrium model, at least in terms of average absolute prediction error. One can see why in Figure 16, which is the analog of Figure 15 for the middle phase, and shows the predicted time paths of the cost of FR paid by NG from alternative models. In the late phase, the equilibrium predictions are slightly above the single- and multi-period predictions from the learning models which, in turn, are slightly above the data. However, these differences are much smaller than those in the middle phase, and the learning and equilibrium models move in very similar ways; indeed they seem to mimic each other.

Our findings support the presumption that the learning models are "well behaved" in the sense that they seem to converge to a complete information Nash equilibrium or something close to it. Convergence is not generally guaranteed under fictitious play (see Shapley (1964) and Fudenberg and Levine (1998)), so it is encouraging that we find empirical support for it. To the extent that learning models are used in counterfactual analysis (see Lee and Pakes (2009) and Wollman (2016)), this suggests that counterfactual outcomes may be close to equilibrium outcomes that could plausibly be sustained over time. On the other hand, we also see that there is still some room left between our best predictions and the actual bids.

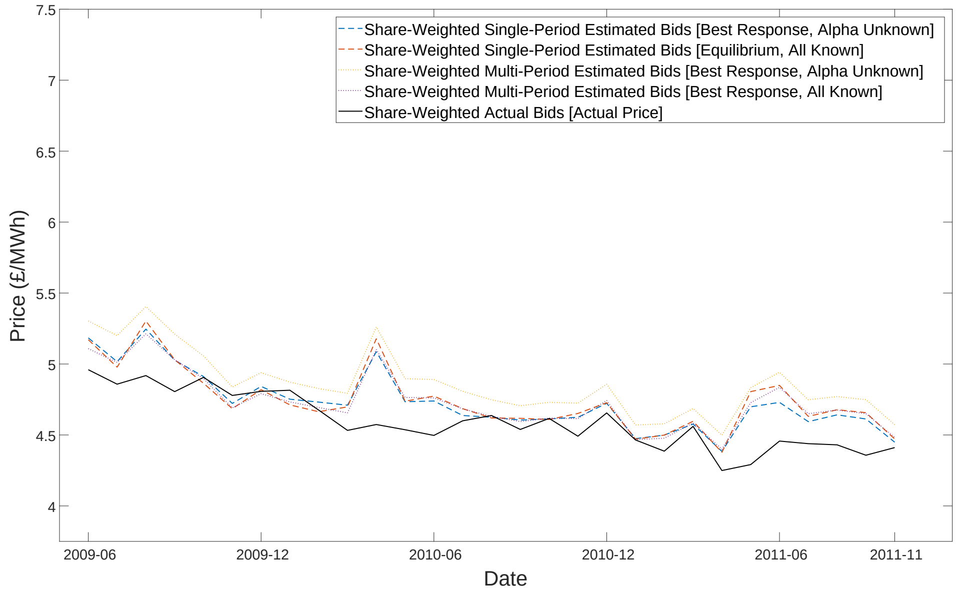

Figure 16: Predicted and actual FR price by month. Share-weighted average computed from bids predicted by fictitious play $F(0)$ with adaptive learning $A(\alpha)$ in both single and multi-period simulations, complete information Nash equilibrium without demand uncertainty $A(\emptyset)$, and actual bids. Late phase.

While the learning and equilibrium models converge to about the same place, there remains the question whether some of these models are better able to adjust to the environmental changes that occurred during the late phase. Here we have less to say. When we regress the actual bids on BM-unit fixed effects and the predicted bids from each model separately, we obtain highly significant positive coefficients on the predicted bids. However, the coefficients obtained from the different models are virtually identical, in line with the fact that the different models are close to indistinguishable according to mean square prediction error.

Efficiency consequences. An interesting question is whether the out-of-equilibrium behavior in the middle phase of the FR market had important efficiency consequences, or merely amounted to a transfer from NG to firms.[35] As FR demand is perfectly inelastic, efficiency hinges on the social cost of meeting that demand. To further assess this, we first compute the average cost per MWh of FR provided by the $J = 72$ BM units owned by the

[35] We thank a referee for suggesting that we investigate this topic.

Overview of Empirical Results.

- There was an initial phase with heterogenous responses and experimentation by some, that we have not tried to model.
- This was followed by a second phase in which there was a distinct movement towards a set of policies which abided by all the characteristics of a Nash equilibrium.
- The **transition (or middle) period** lasted 2.5 years (30 periods), and in it fictitious play models that weighted more recent observations on competitors' play more heavily and adaptive expectations models which estimated the price response parameters do distinctly better than equilibrium predictions.

- **Final period** Despite many more environmental changes, the fit of both the learning & the equilibrium model improves dramatically. This despite the fact that this is the period when the environment is changing in meaningful ways.
- Suggests that once agents have a better understanding of the play of their competitors, they respond rather quickly and thoroughly to changes in their environment with very little experimentation before adjusting to a new equilibrium?