Just starting out: Learning and price competition in a new market *

Ulrich Doraszelski
University of Pennsylvania

Gregory Lewis and Ariel Pakes
Harvard University and NBER

January 12, 2014
(Preliminary and incomplete, please do not cite)

Abstract

Deregulation of the frequency response market in the UK allowed electricity firms to compete on price in an otherwise stable environment. While initial activity in the new market was volatile, we show that after three and a half years, price changes became relatively rare and small, consistent with convergence to a static Nash equilibrium in prices. We examine how well theoretical models of strategic learning do in predicting play during the period prior to convergence. Our analysis suggests that firms discount the past behavior of their opponents and pay relatively more attention to recent actions. There is also evidence of statistical learning, whereby firms estimate demand for their own products from the publicly available data on the past prices and sales volumes of all products. A model that combines simple best response dynamics with statistical learning fits well: it is able to explain 37% of the share-weighted variation in prices.

*We are grateful to Paul Auckland and Graham Hathaway of National Grid, and Ian Foy of Drax Power for useful conversations about this project. We have also benefited from discussions with Joseph Cullen and Mar Reguant, and from the comments of seminar participants at Boston College, Cornell, Kellogg and the NBER Productivity Lunch. Rebecca Diamond, Duncan Gilchrist, Daniel Pollmann, Sean Smith and Amanda Starc have all provided excellent research assistance.

1 Introduction

Entirely new markets arise frequently, often as a result of innovation or deregulation. When multiple firms enter such markets simultaneously, they may face considerable uncertainty about the demand for their products. This is partly because overall demand is typically uncertain, and partly because the demand for their product depends on the pricing decisions of their competitors. This combination of demand and strategic uncertainty makes pricing their own product difficult.

Understanding how firms price products in new markets — or indeed in existing markets following a sudden change in conditions — is an important open problem. A large focus of the empirical industrial organization literature has been the simulation of new equilibria following a counterfactual policy change of some form (e.g. a merger). But convergence to a new equilibrium is not guaranteed to be swift (or indeed certain), and so having a reliable model of out-of-equilibrium pricing dynamics would be extremely helpful. Moreover, such a model would help narrowing down the set of possibilities in cases where there are multiple counterfactual equilibria.

The pricing of differentiated products is generally modeled as a game, and learning in games has been well-studied in the theory literature. The focus there has been on the repeated play of static normal-form games i.e. on stable and known environments in which players are trying to maximize static payoffs. One organizing question in that literature is which learning models are guaranteed to converge to equilibrium, and under what conditions this occurs. Experimental economists have pushed this literature further, using lab data to work out which learning models best describe how people actually learn, and proposing new models as a result.

A second, distinct, theoretical literature offers equilibrium models with uncertainty about the state of nature. There is a long literature in applied mathematics and statistics analyzing bandit problems, in which forward-looking agents trade off "exploration" versus "exploitation". Economists have contributed to this literature by analyzing what happens when multiple agents compete in such an environment, noting informational free-riding incentives (e.g. Bolton and Harris (1999)) and incentives to "signal jam" (e.g. Mirman, Samuelson, and Urbano (1993)). But there has been little empirical work on learning in environments with either form of uncertainty, and consequently little field evidence to discriminate among the many models that have been offered by the theorists and experimentalists.

This paper attempts to fill this gap, with an analysis of the frequency response market in the UK. Broadly speaking, frequency response is a product required by the system operator in electricity markets to keep the system running smoothly. It is bought from electricity generating firms. Prior to November 2005, generators in the UK were required to provide frequency response to the system operator at a fixed system-wide price. Then the market was deregulated, and generators were allowed set their own prices, setting the stage for price competition. We study what happened next.

This particular market is well-suited for analyzing learning, as the environment is similar in many respects to the experimental environments designed in the lab. Prices are set simultaneously each month, as in a standard normal form game. The system operator is required to buy frequency response according to a public and fixed set of regulations, which means that demand is relatively stable. Because building a power plant is costly, there is little observed entry during the sample period, and the set of major players doesn't change. And while the market is somewhat concentrated, there are still ten big firms, so that tacit collusion is difficult and repeated game considerations may be less relevant. Having said that, there is the added element of demand uncertainty, which means that from the firms' point of view, the environment may evolve as they learn.

Our paper attacks three main questions: Do firms learn at all? Does the market settle down to an equilibrium? And which model of learning best explains the dynamic process that leads to this rest point?

We offer descriptive evidence that the market reaches a rest-point after three and a half years. Generators set substantially different prices in the early stages of the market, and update their prices regularly. But after three and a half years of repeated interaction, bid changes became infrequent and small, and the cross-sectional variance in bids shrinks dramatically.

The rest of the paper brings econometrics to bear in trying to assess whether this rest point is indeed a static Nash equilibrium, and how learning occurred. We estimate a flexible demand system using all the data, and then back out marginal costs under the assumption that firms play a static Nash equilibrium in the latter part of the data. Our estimated costs are in line with the "cost-reflective" fixed price prior to deregulation, which is evidence that the rest point constitutes a static Nash equilibrium.

Next, we consider a series of learning models. Recall that the firms may be uncertain about the behavior of their rivals, and in addition may be uncertain about the demand for their

products, given rival behavior. In our first set of models we assume away the second source of uncertainty to focus on strategic uncertainty. We allow firms to place differential weight on the past play of their competitors (in the theory literature, this is "fictitious play"). In our second set of models we allow firms to be uncertain about the demand parameters as well as rival strategies. We model them as engaging in statistical learning, essentially running regressions to determine the demand parameters, given the set of data available to them at each point in time. This approach originates from the macro literature on rational expectations (e.g. Townsend (1983)).

Our criterion for comparing the different models is their ability to make "one-step-ahead" predictions: we compare the predicted bid for time t given data from time $t-1$ to the actual bid at time t, where the predictions vary across models. None of the models in which firms know the demand system from the start fit the data particularly well. We are able to reject a fictitious play model in which firms believe that their rivals are equally likely to take any of their past actions. Instead we find that firms discount the past behavior of their opponents and pay relatively more attention to recent actions (for example, a dynamic best response model fits the data as well as a fictitious play model with geometric discounting of past actions).

By contrast, when we allow for statistical learning, we obtain a much better fit. A model in which firms re-estimate demand each period and best respond to their rival's actions given their current demand estimates fits the data quite well indeed: the (volume-weighted) R^2 is 0.37, which is impressive considering that this model has no estimated parameters (our simulated bid is just the optimal bid given the model). This suggests that it is important to consider other sources of uncertainty, in addition to strategic uncertainty, when projecting the evolution of a new market.

Related literature. There is a large theory literature on learning in static normal-form games. The organizing feature of this literature has been on deriving conditions under which the canonical models of fictitious play (Brown 1951) and reinforcement learning imply convergence to equilibrium (e.g. Milgrom and Roberts (1991), Fudenberg and Kreps (1993), Börgers and Sarin (1997), Hart and Mas-Colell (2000)). Around the same time, there were a number of meta-studies, combining data from multiple lab experiments: Boylan and El-Gamal (1993) finds that a fictitious play model performs better than adaptive best response; Cheung and Friedman (1997) find evidence of heterogeneity across players in the discount

rate in fictitious play models; Erev and Roth (1998) show that a simple one-parameter reinforcement learning model can fit their data quite well. Subsequent work has led to more general models, such as experience-weighted attractive learning (Camerer and Ho 1999); or models with sophisticated learners, who try to influence how other players learn (Camerer, Ho, and Chong 2002).

There has been no prominent empirical work on either of these topics (i.e. convergence to equilibrium, or estimating learning models), but there are some papers on related topics. Hortaçsu and Puller (2008) show that in the newly created spot market for electricity in Texas (the "main market" in our terminology), big firms made bids that were equilibrium best responses, but small firms failed to optimize fully (although their behavior improved over time). This suggests that the market converged to equilibrium, at least for players with large stakes. Goldfarb and Xiao (2011) show that managers with different experience and education levels make different entry decisions in local US telephone markets after deregulation. This is evidence in favor of heterogeneity in decision-making ability.

Surprisingly, there is also little empirical work on how agents learn about demand for their product from the observation of their and their rivals' sales. In the IO and marketing literature, a number of papers have examined how agents may learn their demand for experience goods from their own experimentation (Erdem and Keane 1996, Ackerberg 2003, Dickstein 2013). There is also a small empirical literature on observational learning, where agents see the choice of others but not outcomes (e.g. Zhang (2010) on patient decisions to accept a kidney offer, and Newberry (2013) on music downloads). Social learning has been more widely studied in other contexts (see e.g. Griliches (1957) on hybrid corn, Conley and Udry (2010) on fertilizer in Ghana, and Covert (2013) on fracking in the Bakken Shale).

Our contribution is to simultaneously attack both of these topics, in an environment that has much of the clean structure of the lab experiments, and at much higher stakes. The main limitation is that we have evidence on a particular game — a pricing game with differentiated products — played in a particular market, which may caution against general inference.

Paper Structure. The paper proceeds as follows. In the next two sections we describe the frequency response market, our data on it, and some descriptive evidence on how it evolved over time. The following section outlines our estimation strategy for recovering the supply and demand primitives. We then estimate a number of learning models and compare their comparative fit, before concluding. Additional information on the construction of the data

and the estimation procedures are to be found in the appendix.

2 Overview of the British electricity market

The British electricity market is a network of generators and distributors, connected by a transmission grid. This grid is owned and operated by a company called National Grid (NG). NG is responsible for the transmission of electricity from the generators to the distributors, as well as the balancing of supply and demand in real time.

The unit of exchange in this market is a given amount of power supplied for a half-hour (measured in megawatt hours (MWh)). About 98% of electricity is sold through some form of forward contracting. Bilateral forward contracts between generators and distributors account for the vast majority of these sales, and they can be formed months or even years in advance. There are also shorter term contracts (both day ahead and day of) which are often traded on power exchanges. One hour prior to the settlement period, both distributors and generators must submit their contracted positions to NG, who then holds an auction to equate supply and demand as expected over the settlement period. This multi-unit discriminatory auction is called the balancing mechanism (BM), and it accounts for the remaining 2% of electricity sales. The generators bidding in the BM and are called BM units. A power station typically consists of multiple BM units, and multiple power stations may be owned by a single firm. The market for electricity — referred to as the "main market" in the rest of the paper — is summarized in figure 1.

Frequency response. NG is obligated by government regulation to maintain a system frequency within a one-percent band of 50Hz. System frequency is determined in real-time by imbalances between the supply and demand of electricity. The higher demand is relative to supply, the lower the system frequency is, and vice versa. Imbalances occur due to shocks that cannot be corrected in advance through the BM. To balance the system, NG instructs one or more BM units into frequency response (FR) mode. Once in this mode, NG can rapidly adjust the energy production of the BM unit using so-called governor controls.

NG is required by government regulation to hold an certain amount of FR capacity at all times.[1] This response requirement is based on risk-response curves that assess the likelihood

[1]There are in fact three different kinds of frequency response: primary, secondary and high. Primary

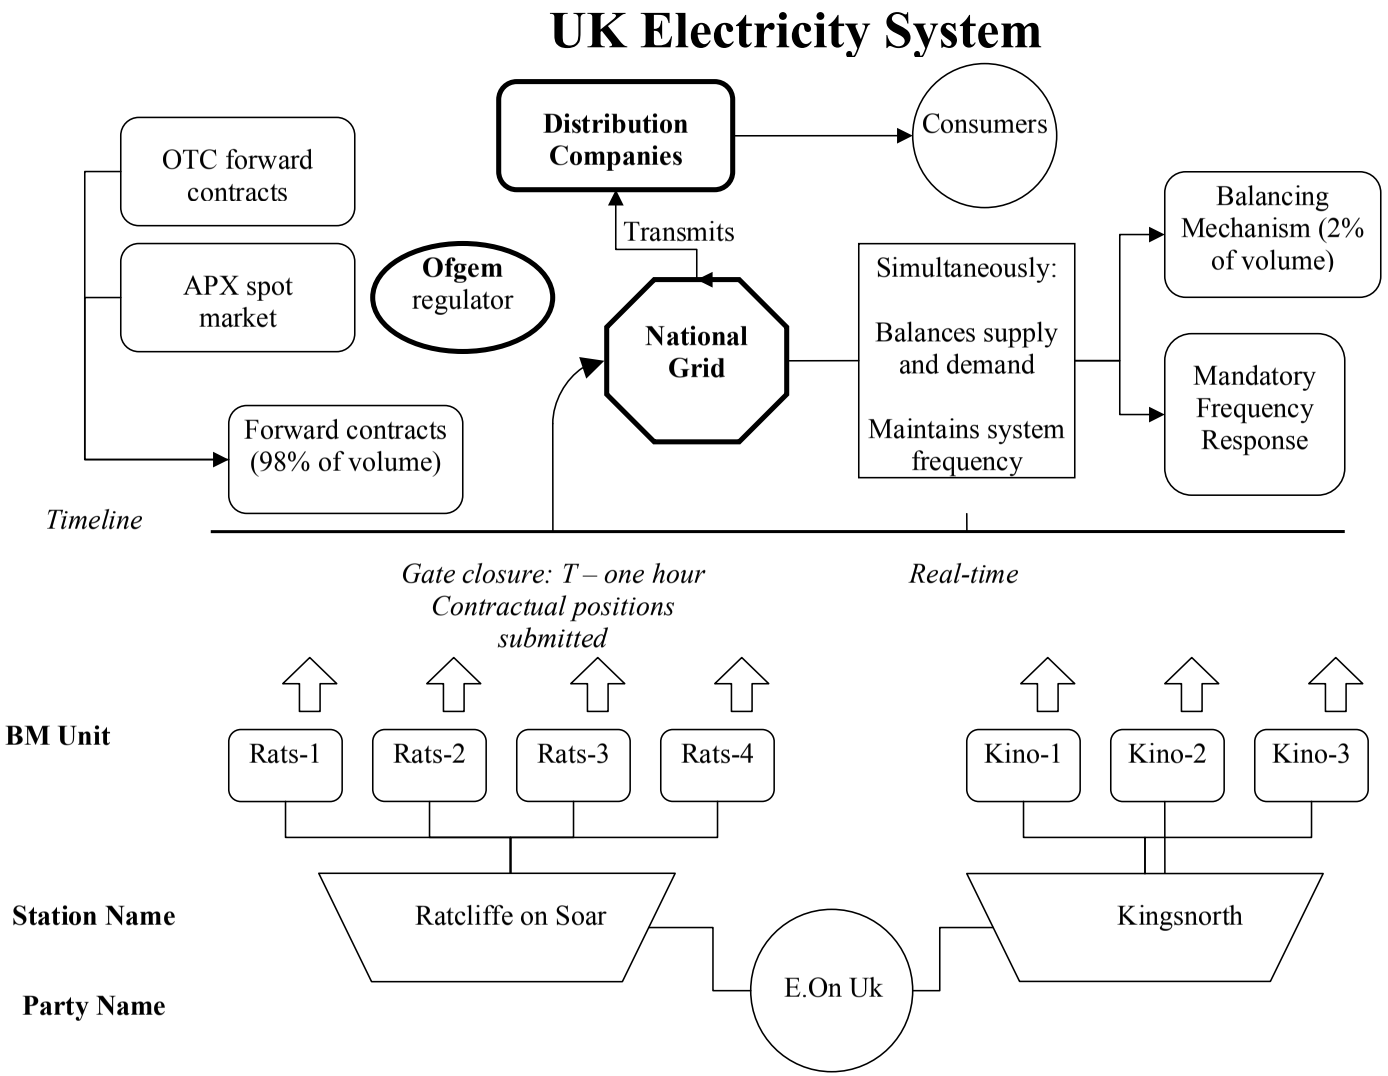


Figure 1: Overview of the British electricity market.

and magnitude of possible shocks given the total amount of electricity demanded. As the total amount of electricity demanded evolves, NG instructs BM units in and out of FR mode in real time to satisfy its response requirement. To the best of our knowledge, the response requirement was held fixed over the sample period.[2]

FR capacity is thus a second product, distinct from electricity, that BM units can sell to NG, and the FR market is distinct from the main market. Offering FR capacity is costly: a unit in FR mode incurs additional wear and tear as it may have to make rapid adjustments to its energy production in response to supply and demand shocks.[3] It is compensated by NG by a holding payment and an energy response payment. The holding payment is per unit of FR capacity and paid for the time that it is called into FR mode regardless of whether the BM unit has to adjust its energy production in response to demand or supply shocks. The energy response payment compensates the BM unit for actual adjustments to its energy production.[4] These energy response payments remained constant pre-and-post deregulation, and are considered by industry insiders to be a relatively small source of profits. For this reason, our focus on this paper will be on the holding payments.

Deregulation. Our interest in the FR market stems from a change in way holding payments were calculated, which occurred with the enactment of an amendment to the Connection and Use of System Code called CAP047. This change "went live" on November 1, 2005. Pre CAP047, providing FR was mandatory, and the holding payments were at an administered price (per unit of FR capacity) which had been fairly constant over time (see figure 2). CAP047 replaced the mandatory provision of FR with a market for FR. In this market, holding payments are determined by an auction.

response is additional energy from a BM unit that is available ten seconds after an event and can be sustained for a further twenty seconds; secondary response is additional energy that is available within thirty seconds for up to thirty minutes. High response is a reduction of energy within thirty seconds. These responses are technologically constrained and correspond to dilating the steam valve (primary), increasing the supply of fuel (secondary), and decreasing the supply of fuel (high). For historical reasons, BM units are instructed into FR mode in only two possible combinations: primary-high and primary-secondary-high. To simplify the presentation and analysis, we aggregate over the three types of FR (primary, secondary, and high); see the Data Appendix for details.

[2]We have checked the publicly available minutes of all meetings of the frequency response working group (comprising representatives from the generating firms and NG), and at no point was there a discussion of a change in the overall response requirements.

[3]It also runs less efficiently, with a degraded heat rate.

[4]If the BM unit produces more energy than it was initially contracted to in the BM, NG pays it 125% of the current market price per additional unit of energy; if the BM unit produces less energy, it pays NG 75% of the current market price.

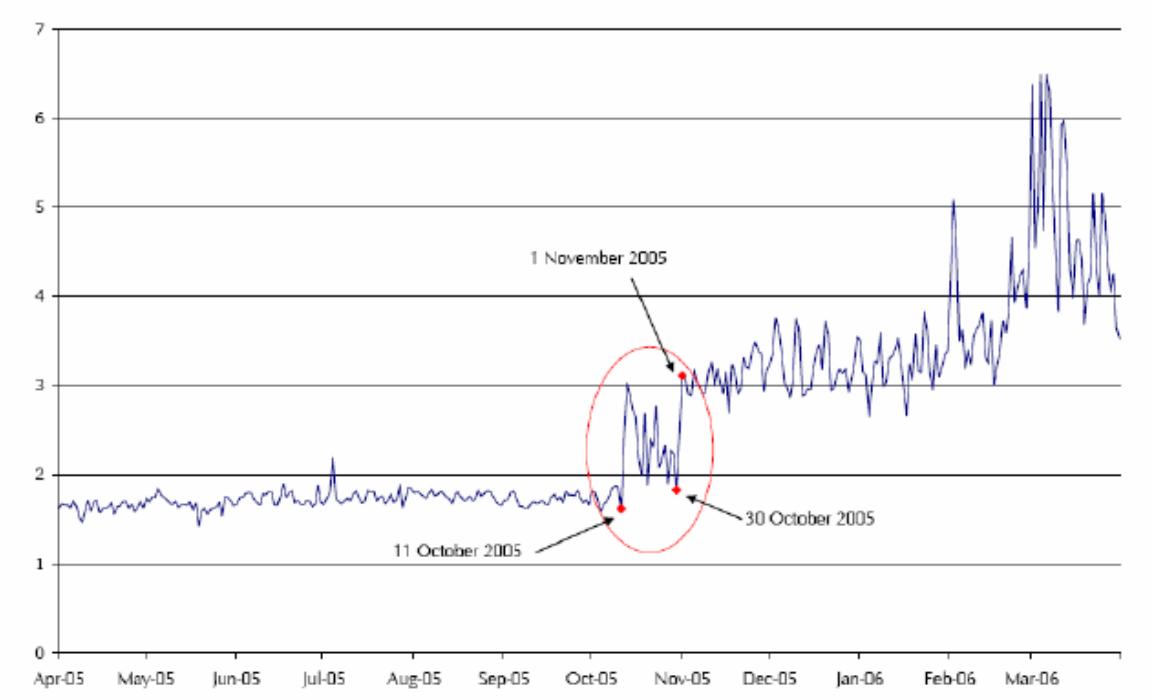


Figure 2: **Holding payments for high frequency response pre and post CAP047.** Source: National Grid.

Post CAP047, a BM unit tenders a bid each month for FR. The bid for the next month is submitted before the 20th of the current month, well in advance of electricity production, and consists of a price per unit of FR capacity. If called upon by NG, the BM unit is paid a holding payment equal to its bid per unit of FR capacity for time spent in FR mode (i.e. they get "paid-as-they-bid"). A bid commits the BM unit to offer FR at a fixed price over the next month. The quantity that a unit delivers if instructed into FR mode varies with its operating position and current system deviation according to a unit-specific contract between the generator and NG that is generally fixed over the sample period.[5]

NG can choose to call upon any unit at any time, and often does not choose the lowest bidding units to provide frequency response. Instead, it simultaneously accepts bids in the

[5]This contract takes the form of a 5×3 matrix for each type of frequency response (see footnote 1) that specifies the quantity delivered at 5 deload points (operating positions) and 3 system deviations (0.2Hz, 0.5Hz and 0.8Hz away from 50Hz). At other deload points and deviations, the quantity is found by linear interpolation. The matrices are proprietary information, but selected entries are published by NG in the capability data (see the data appendix). Generators can re-bid the matrix as well as the price each month, but for over 80% of the units, the observed entries don't change at all during the sample period.

BM and issues FR instructions to equate supply and demand and maintain the mandated amount of FR capacity in the most cost-effective way. In practice, this cost minimization problem is solved by a computer program running a linear program.

The market for FR was proposed by one of the largest electricity firms in the UK market, RWE. This proposal was bitterly opposed by NG, who argued that since their demand for FR is regulated and thus inelastic, generators would be able to exploit their market power and the price of FR would rise. The government regulator dismissed these concerns, and on November 1, 2005 introduced CAP047. Figure 2 shows that NG had every reason to worry about CAP047, as average daily holding payments for FR doubled within the year after the introduction of the market.

From the pre-CAP047 period, firms likely had an understanding of NG's requirements and their BM units' relative desirability given their main market positions, as well as the cost of providing FR. However, firms faced uncertainty as to the demand for their frequency response services because they didn't know how their rivals would bid, nor how price sensitive NG was. Our goal is to understand how firms learned to bid in the presence of this initial uncertainty, and how this contributed to the trends in holding payments over time.

Data. We collected most of our data from two public sources. Our data on the FR market comes from NG. For the post-CAP047 period we have the bids submitted by each BM unit at a monthly level and the quantities provided of each type of FR (in MWh) by each BM unit at a daily level. The combination of bid and volume data allows us to calculate the daily holding payment received by each BM unit.

The second data source is Elexon Ltd. They are contracted by the government regulator to manage measurement and financial settlement in the BM. For every BM unit we collected detailed data on the bids and acceptances in the BM every half-hour. This allows us to assess the operating position of the BM unit. Finally, we worked out which firms own which BM units through data from the Central Registration Agency.

Market participants. There are 132 BM units grouped into 63 power stations. The BM units belonging to the same power station tend to be identical. The 132 BM units are owned by 40 firms. Table 1 provides summary statistics on revenue in the FR market for the ten largest recipients over the first seven years of the market's existence (until October

Table 1: Firms with the largest frequency response revenues

Rank	Firm name	Total Revenue	Revenue Share (%)	Cumulative Share (%)	Months revenue > 0
1	Drax Power Ltd.	99.5	23.8	23.8	72
2	E.ON UK plc	67	16	39.9	72
3	RWE plc	48.5	11.6	51.5	72
4	Eggborough Power Ltd	29.8	7.1	58.7	72
5	Keadby Generation Ltd	24.2	5.8	64.5	72
6	Barking Power Ltd	17.8	4.2	68.8	72
7	SSE Generation Ltd	15.2	3.6	72.5	55
8	Jade Power Generation Ltd	15	3.6	76.1	69
9	Centrica plc	14.8	3.5	79.6	72
10	Seabank Power Ltd	14	3.3	83	72

Inflation-adjusted revenue in millions of british pounds (base period is November 2012). There is information on 72 months in the data.

31, 2011).[6] The ten firm concentration ratio is just over 80%, and the HHI is 76.5, so the FR market is mildly concentrated.

The participant that earned the most revenue, Drax, had about 20% of the FR market and earned over £90,000,000 over the sample period, or about £1,100,000 per month. Drax is a single-station firm, while the next two largest participants, E.ON and RWE, are multi-station firms. Anecdotally, Drax's disproportionate revenue share is attributable to having a relatively new plant, with accurate governor controls, making them attractive for providing FR. The smallest of these firms, Seabank, still makes around £200,000 per month. All of this suggests that the 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. For example, in 2006 Drax hired a trader to specifically deal with the FR market.[7] Within a year, this increased Drax's annual revenue from the FR market more than threefold.

[6]Most firms have positive revenue in every month; the exceptions tend to be single-station firms who may shut down occasionally for maintenance.

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

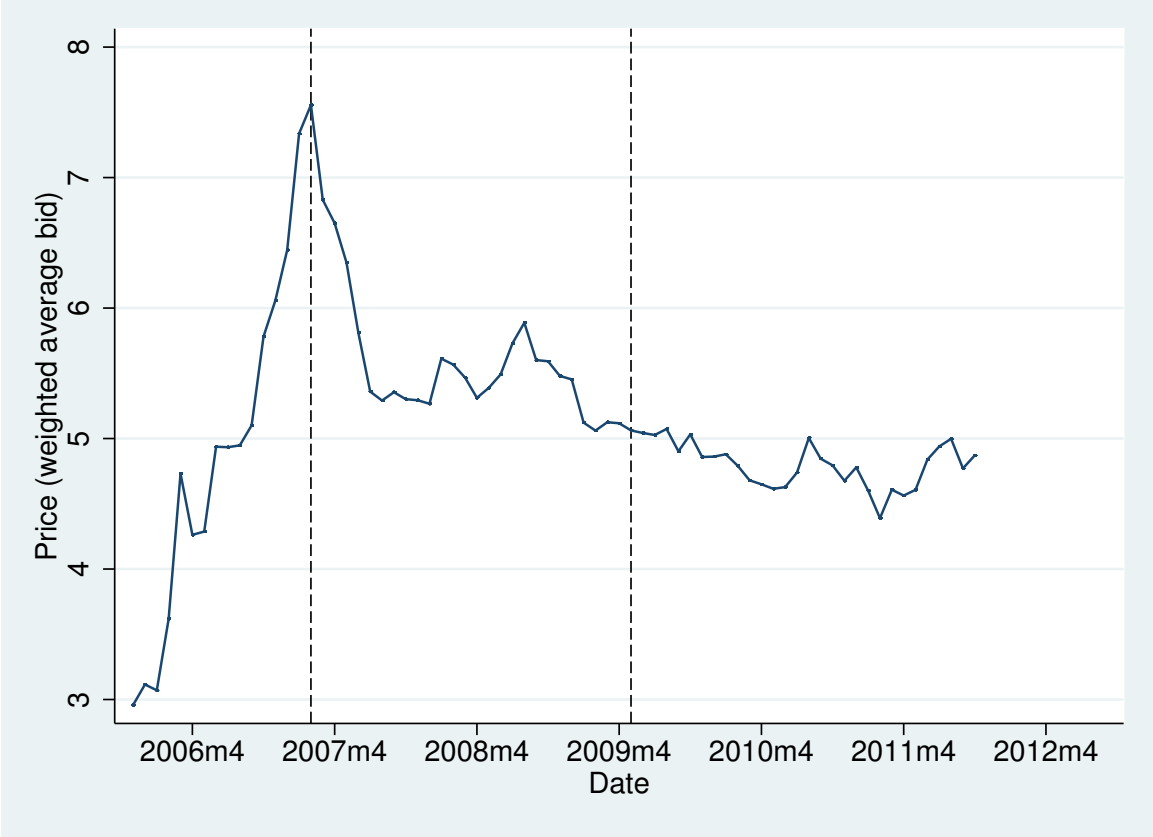


Figure 3: **Monthly prices over time.** Prices are weighted averages of bids across units, where the weights for each period are the volumes received by each unit.

3 The evolution of the frequency response market

We divide our analysis of the evolution of the FR market into three periods or phases. Figure 3 shows the prices over the entire period with the three periods delineated by vertical lines. During the first phase, from November 2005 until February 2007, the average holding price exhibits a noticeable upward trend, moving from an initial price of £3.1/MWh to a final price of £7.2. From March 2007 to May 2009 this trend reverses itself and prices fall back down to £4.8/MWh. From June 2009 to the end of our study period there is no obvious trend at all. While there are fluctuations, they are smaller, and prices stayed in the range of £4.3/MWh to £5.1/MWh. Moreover, during this third phase, the sharper movements in one direction are followed by corrections in the opposite direction.

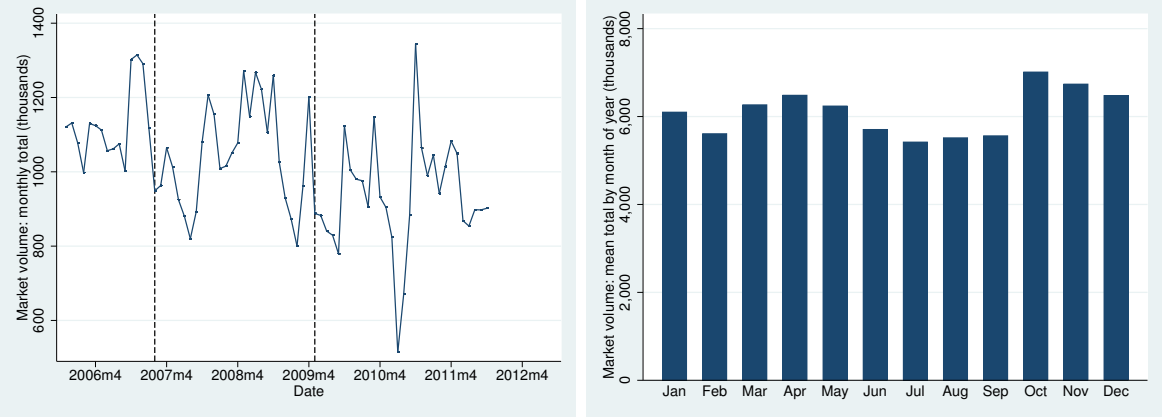


Figure 4: **Volumes over time and by season.** The left panel shows the total volume of frequency response for each month in the sample period, while the right panel shows the total volume by month-of-year (4 observations per month-of-year).

Environment. Our analysis assumes that the environment in which the FR market operates, i.e., the demand and supply conditions, is relatively stable. We now provide evidence that this is largely true. The left panel of figure 4 plots monthly data on the volume of FR. Though these are clearly volatile, they are no more volatile at the beginning than at the end of the period we study. There is a bit of evidence of seasonality in volumes: see the right panel of figure 4. While this may explain some of the fluctuations in the final period, the price movements in the first two periods are too persistent to be driven by seasonality. There is also a slight downward trend in volumes over the sample period, which we attribute to NG's signing of long-term contracts with individual firms to fulfill part of their frequency response requirements. Again, this is small.

As further evidence of stability, in the left panel of figure 5 we plot the realized daily electricity demand in the UK over the sample period. Though it is clearly seasonal, there are no obvious trends in electricity consumption. Since NG's response requirement derives from risk-response curves that did not change over the sample period, the stability of the main market implies the stability of frequency response demand that is documented above.

Turning from the demand for FR to the supply of FR, the right panel of Figure 5 plots the number of "active" stations over time, where a station is active if one of its units submits a competitive bid (bids ≤ 23). While the number of active stations does fluctuate a bit, ranging from 53–61 over the sample period, the movements are relatively small and none of the stations who become active or inactive are particularly large stations.

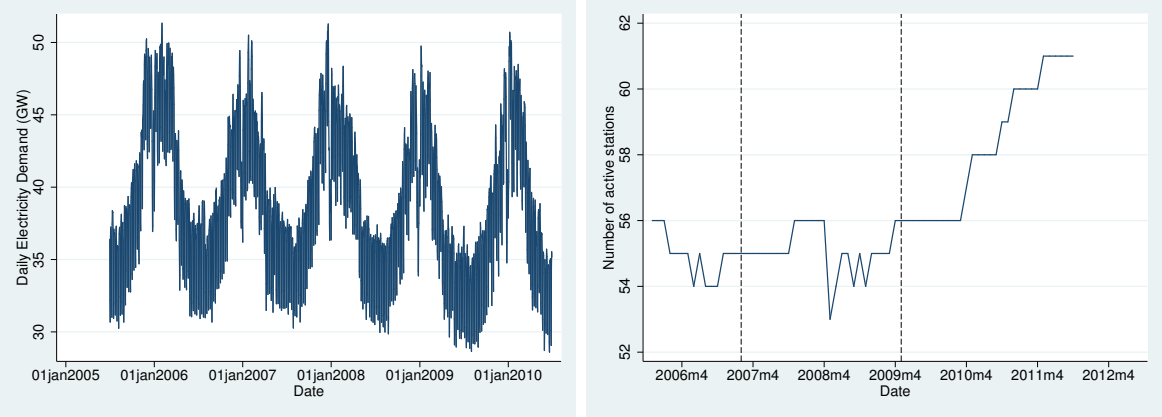


Figure 5: **Main market volume and number of active stations over time.** The left panel plots the total daily realized electricity demand in the UK over the sample period. This plots the number of active stations over time, where a station is "active" if one of its units submits a competitive bid (bids ≤ 23).

We conclude that the changes in the holding payments in Figure 3 occurred despite the relative stability of the demand and supply conditions. As a result we look to an explanation that is rooted in changes in bidding behavior over time. For each phase, we begin with an overview of how bidding behavior changed from period to period. After providing the overview we come back to look more closely at the role of individual power stations.

Early period (November 2005 – February 2007). Firms are most likely to experiment with different bids, during the first, or rising price, phase of our data. In this phase they adjust their bids more frequently as well as by larger amounts (in absolute value) than in any of our other periods. This is illustrated in Figures 6 and 7. The first graphs a weighted average of indicator variables for a BM unit changing its bid, and the second graphs the absolute value of the change conditional on changing. The weights are the fraction of FR volume in the base (lagged) period. The other fact that is worth noting is the cross-sectional variance in bids is an order of magnitude higher in the first period than in other periods, see Figure 8. We show below that during this period virtually all bidding changes by firms were bid increases. What the large cross-sectional variance in bids thus indicates is that the bid increases were not coordinated across firms in time.

Figure 9 shows the bids of the biggest 8 stations in the opening period.[8] The levels and

[8]We define "biggest" by revenue rank during the pre-equilibrium period, as we want to include all the

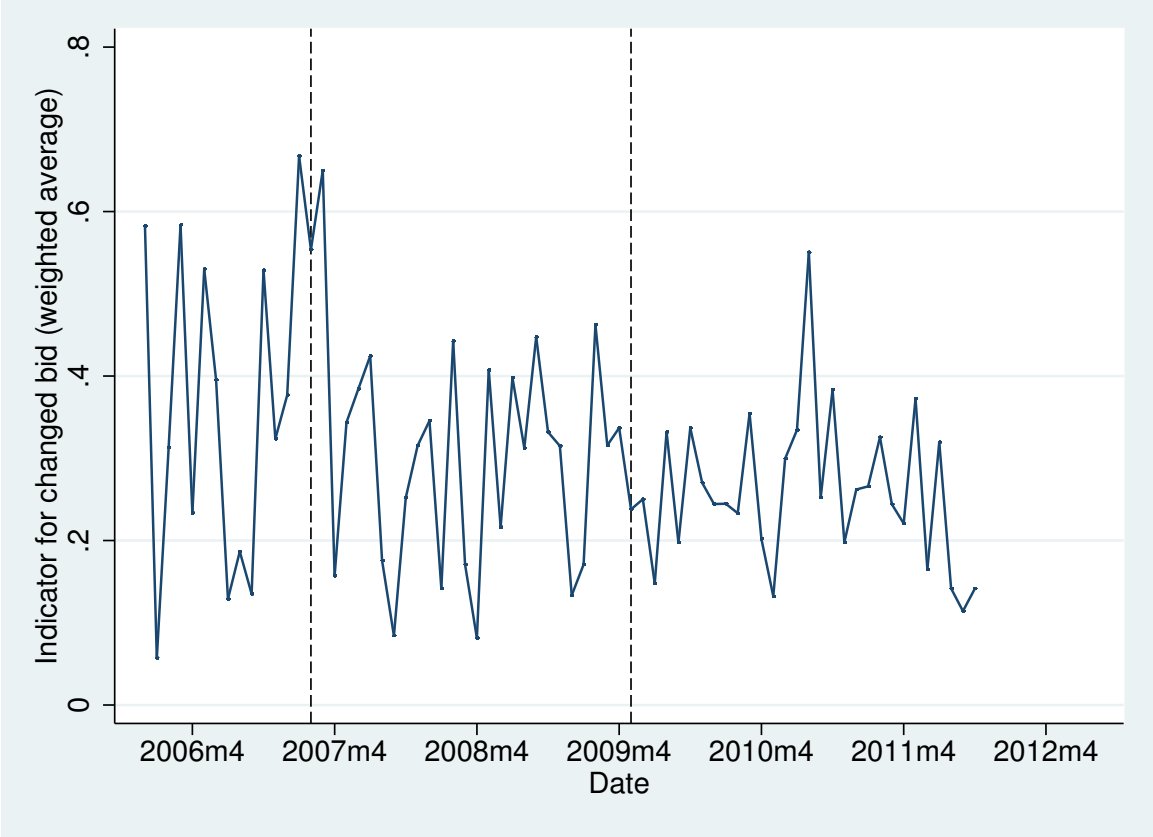


Figure 6: **Probability of a bid change over time.** This is a weighted average across units of an indicator for a change in their bid in period t, using quantity-weights from period $t-1$.

trends of the bids across stations are quite different. Peterhead and Seabank bid very high early on — pricing themselves out of the market — and then drift back down into contention. The remaining firms 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" we see in Figure 3.

At least to some extent, this heterogeneity in bidding strategies and accompanying price volatility is to be expected. Firms had no experience in this market and had to learn the likely profitability of different bidding strategies. This view is consistent with a conversation we had with Ian Foy of Drax Power, the largest provider of FR. He explained that initially they found it difficult to project revenue for each of their BM units. Though they were aware of the cost of providing FR in terms of both efficiency and wear and tear, it was hard to predict how others would bid and therefore the volume of FR they would attract. In each

big players during this period.

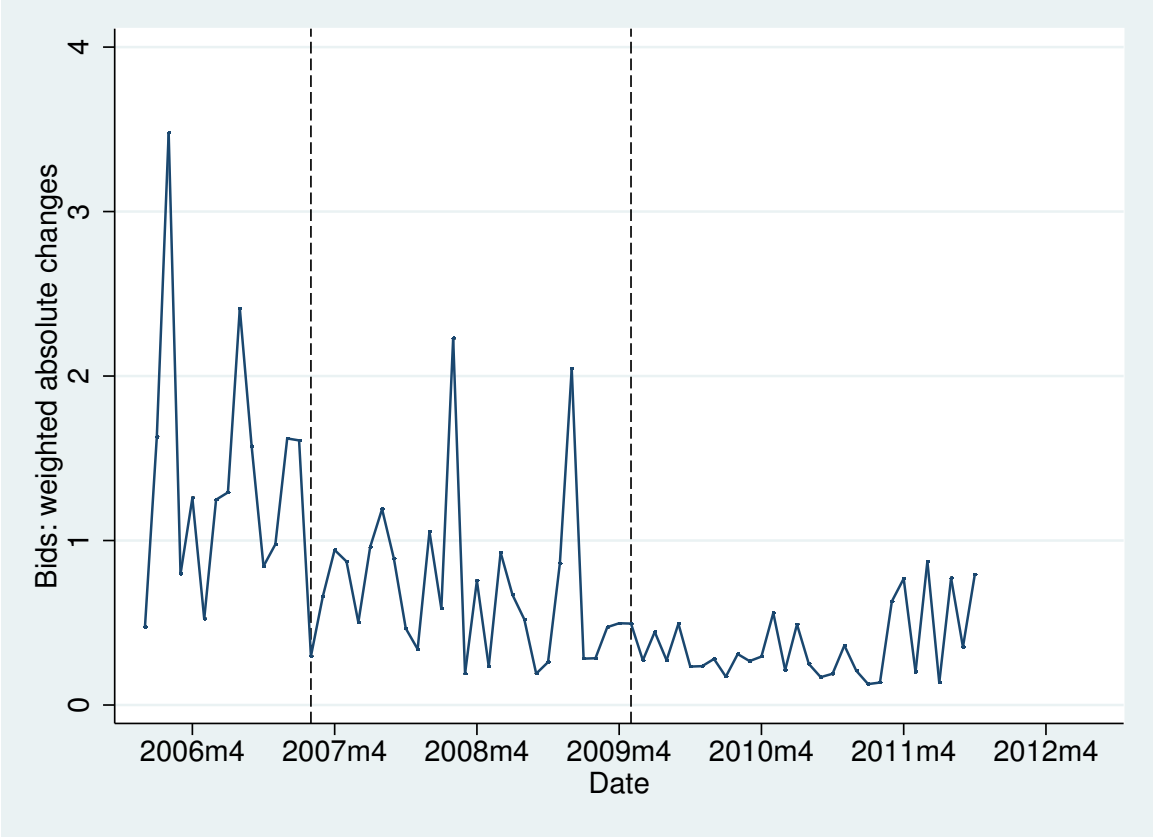


Figure 7: **Absolute bid changes over time conditional on changing.** This is a weighted average across units of the absolute change in their bid in period t, using quantity-weights from period $t-1$ and assigning a weight of zero if their bid did not change.

month they tried to predict the demand for their services at various prices, and changed their bid from the previous month if they thought they could make more money. After some period of time they no longer perceived any opportunities for increasing their profit.

Middle period (March 2007 – May 2009). There are still some sharp increases in bids in the second, or falling price, period. However in this period the upward bidding experiments are short lived. The dominant trend is for the bids of the different firms to move toward one another. The way this happens is that the firms that entered the period bidding higher decreased their bids while the firms that entered the period with relatively low bids maintained those bids. This combination generated the marked decrease in the cross-sectional variance in bids seen clearly in Figure 8. Indeed the overall impression is

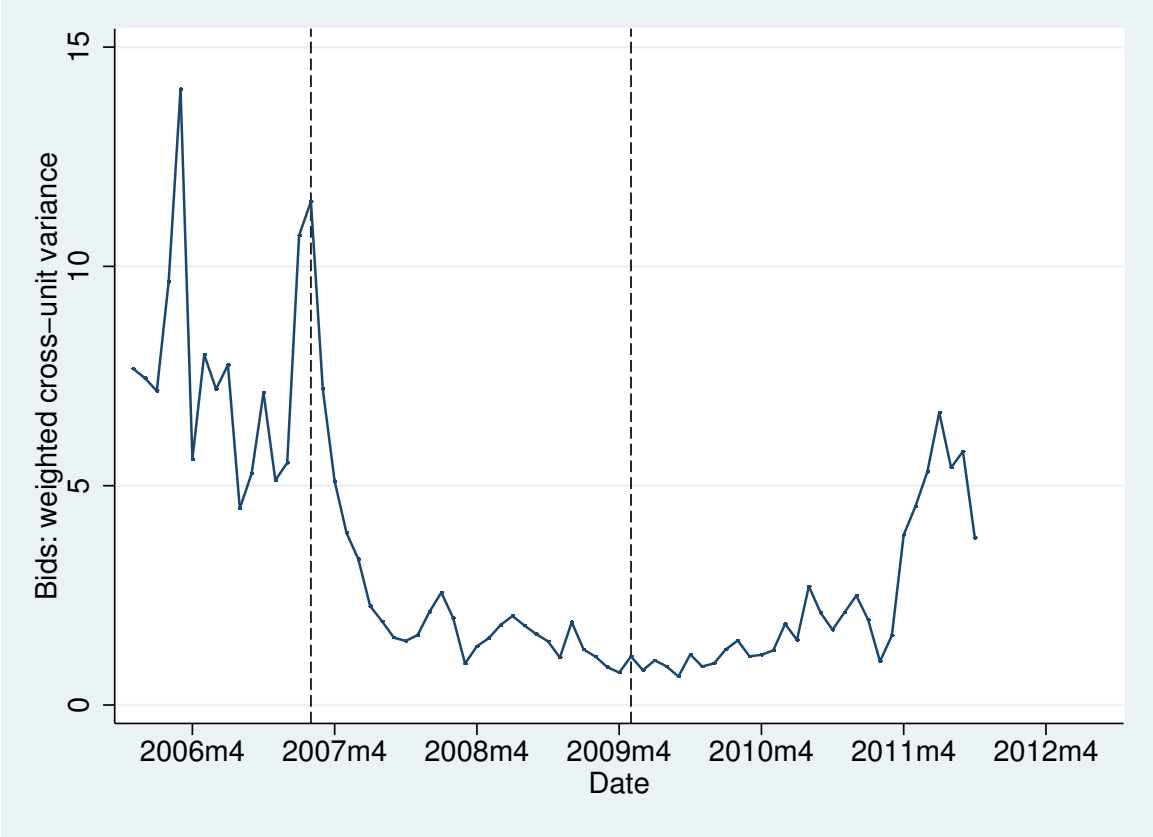


Figure 8: **Variance in bids across units.** Variance in the cross-section of bids at each point in time, where bids are quantity-weighted using contemporaneous quantities.

that this was a period of intense competition where the changes in bids that occurred were designed to undercut competitors' bids.

To give more detail, the "price bubble" bursts when Seabank sharply decreases their bid and steals significant market share from Drax. Drax follows Seabank down and this inaugurates the period of intense competition that we see in Figure 10. There are attempts to establish a higher price (notably by Drax) but they were not successful and bid cuts and falling prices are the dominating feature of this period.[9]

[9]Drax increased its bid at the end of 2007 for *exactly* two periods, giving its rivals an opportunity to see its bid increase and respond, presumably with the aim of some sort of tacit collusion. When no one followed suit, they decreased it.

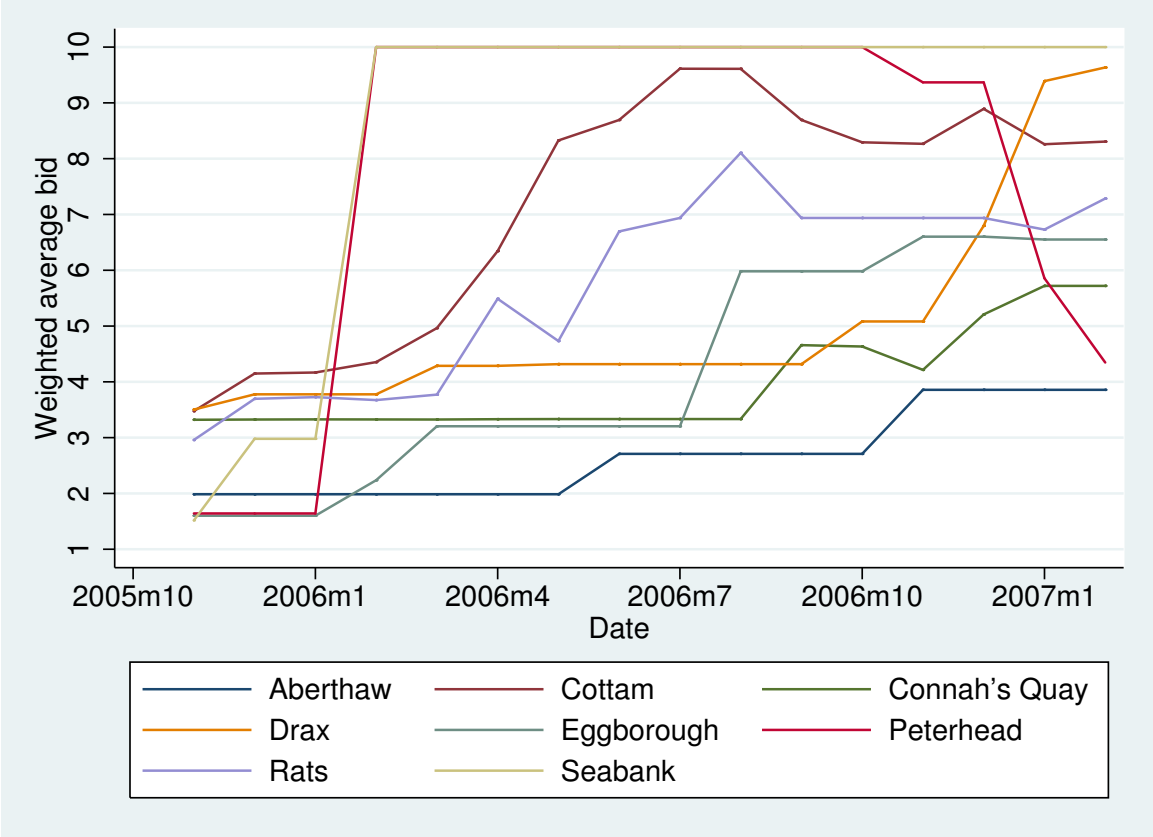


Figure 9: **Bids of the top 8 biggest stations (November 2005 – February 2007)** Stations are ranked by total revenues in the pre-equilibrium period. Bids are censored above at 10 to compress the axes and improve the visual presentation.

Final period (June 2009 – October 2011). There are fewer and smaller changes in bids during the final period, see again Figures 6 and 7. The smaller bid changes are accompanied by changes in average prices. However, each time prices move in one direction there is a mitigating effect in subsequent periods, bringing prices back to a level of about £4.8/MWh, the price at the beginning of the period, see Figure 3. By the time we reach this period it looks like prices have converged to an "equilibrium", or a rest point.

Figure 11 makes it clear that the range of bids in this period is much smaller than in the prior periods. While bids at some stations continue to fall (Rats and Cottam), others are more erratic or rise (Drax and Eggborough), and others are completely flat (Aberthaw). The fact that the bids are relatively stable and the resultant prices are relatively flat during this period leaves the impression that the market has reached some sort of "rest point" that is,

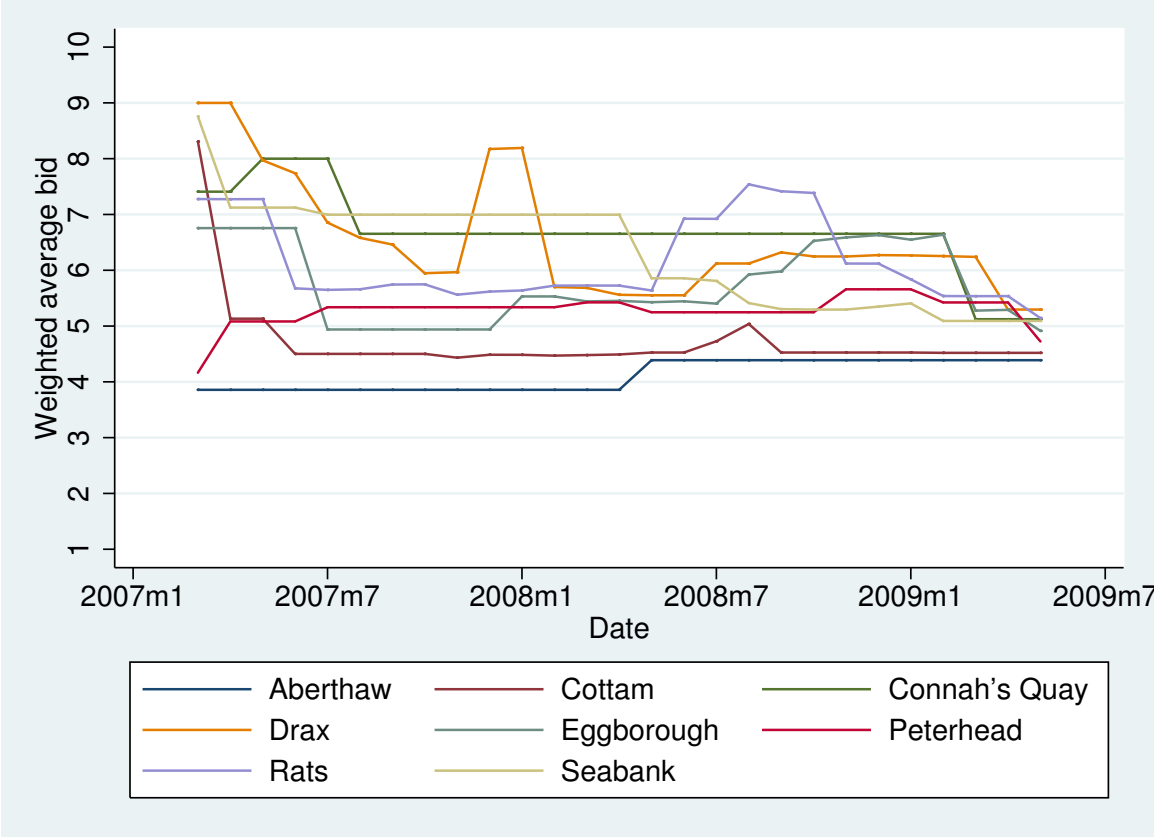


Figure 10: **Bids of the top 8 biggest stations (March 2007 – May 2009)** Stations are ranked by total revenues in the pre-equilibrium period

perhaps, periodically shocked by the volume changes we see in Figure 3.

4 Estimating demand and cost primitives

Our goal is to understand whether the changes in bidding behavior over time can be rationalized by common learning models. To do this, we would like to compare the observed bidding behavior with the predictions from a model. To keep things simple, we will restrict our attention to the bids of the ten biggest firms. We focus on models in which bidders maximize their expected static profit given their model of the world and their current information set. So to predict behavior, we need to estimate the profit function, which in turn requires estimates of the demand for frequency response and the costs of providing it.

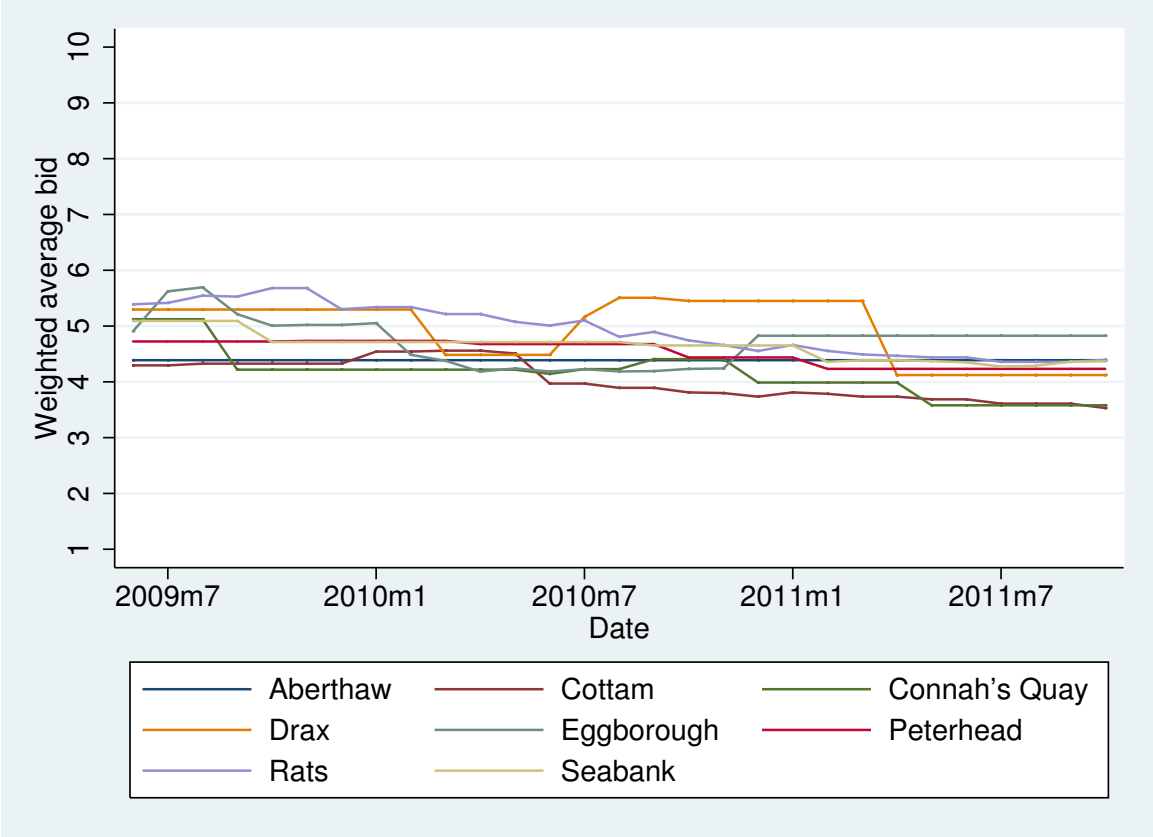


Figure 11: **Bids of the top 8 biggest stations (June 2009 – October 2011)** Stations are ranked by total revenues in the pre-equilibrium period

Before getting into specific estimation strategies, we formalize the above ideas. Let i index firms, j index BM units, and t index months. Let $\mathcal{J}_i$ denote the indices of the BM units that are owned by firm i and $b_{i,t} = (b_{j,t})_{j \in \mathcal{J}_i}$ the bids for these BM units. We adopt the usual convention to denote the bids for all BM units in month t by $b_t = (b_{i,t}, b_{-i,t}) = (b_{j,t}, b_{-j,t})$.

We assume that each BM unit has a marginal cost of providing frequency response $c_{j,t}$. As this cost mainly reflects wear-and-tear on the machine, we think it is likely to be stable over time, but we allow it to be time-varying for reasons we will explain in section 4.2 below. The realized profit for firm i in period t is then equal to:

$$\pi_{i,t} = \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_{j,t}) q_{j,t} \tag{1}$$

We assume that NG allocates frequency response volumes to unit j according to a residual

demand curve $q_j \equiv q_j(b_t, x_t, \xi_t; \theta)$ that is parameterized by θ. The vector x_t stands in for a set of unit characteristics that are mutually known to all firms at the time of bidding, while ξ_t is a set of mutually unknown and time-varying characteristics. We allow for the possibility that firms do not know the demand parameters θ, and are thus uncertain about the form of the demand system. They also may not know how their opponents will bid, and so $b_{-i,t}$ is unknown when choosing $b_{i,t}$, resulting in strategic uncertainty.

We summarize both forms of uncertainty by endowing each firm with a time-varying expectation operator $\mathcal{E}_{i,t}$ with which they assess their expected profits to taking any action given their information set $\Omega_{i,t}$ (their expectations must be measurable with respect to their information set). To keep things simple, we assume that firms maximize static profits in each period. This is consistent with the theory and lab experiments on strategic learning, but is in contrast to the literature on optimal experimentation and signal jamming under demand uncertainty, in which firms sacrifice current period profits to gain information or mislead a rival. There is some evidence that some firms experiment in the opening phase of the data, but this does not appear to be an important feature of the middle and later periods which will be the focus of our later analysis.

Thus we assume that bids $b_{i,t}$ solve the following static maximization problem:

$$\max_{b_{i,t}} \mathcal{E}_{i,t} \left[\sum_{j \in \mathcal{J}_i} (b_{j,t} - c_{j,t})\, q_j(b_t, x_t, \xi_t; \theta) \middle| \Omega_{i,t} \right]. \tag{2}$$

The corresponding system of $|\mathcal{J}_i|$ first-order conditions is

$$\mathcal{E}_{i,t} \left[q_k(b_t, x_t, \xi_t; \theta) + \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_{j,t}) \frac{\partial q_j(b_t, x_t, \xi_t; \theta)}{\partial b_{k,t}} \middle| \Omega_{i,t} \right] = 0, \quad k \in \mathcal{J}_i, \tag{3}$$

and the solution is the best course of action for firm i given its knowledge $\Omega_{i,t}$ and its expectations $\mathcal{E}_{i,t}$. These expectations need not be correct. For example, in an environment in which bids are generally increasing from period-to-period, a firm that expects its rivals to hold its bids at last period's levels (i.e. adaptive best response) will consistently underestimate the demand for its product. An important question then is what model firms use to guide their expectations. We will return to this question below. But first, we outline how we estimate the demand parameters θ.

Table 2: Summary Statistics (top 10 firms only)

	Mean	Std. Dev.	Min	Max
Unit share	0.0124	0.0164	0	0.131
Indicator for positive share	0.769	0.422	0	1
Bid	7.797	22.20	1.515	238.6
Average position	0.714	0.318	0	1
Number of observations	4752			

Summary statistics on the frequency response market. An observation is a bmunit-month, and the sample is restricted to units owned by the top 10 biggest firms (ranked by revenue over the sample period). Average position is the average over half-hour periods (in a month) of the unit's declared operating position (as a fraction of its maximum generating capacity).

4.1 Demand

We estimate a logit demand model for monthly frequency response at the unit level using data from the top ten firms. Together the 65 units that are owned by the top 10 firms account for just over just over 80% of the market (see Table 1). We choose to estimate a logit mainly to reduce the dimensionality of the demand system: a linear demand model would require estimating $(66 \times 65)/2 = 195$ own and cross-price elasticity terms. Another advantage is that we can focus on modeling a unit's share of the market rather than the quantity it is allocated. The former is much more stable because the total market size is relatively erratic (see figure 4). One downside is that we have to explicitly deal with units that aren't active in a given month and thus have a zero share in our data, as zero shares are impossible to rationalize in a standard logit model.

The units owned by smaller firms are aggregated into an "other" unit which is treated as the "outside" good. Since the characteristics of this "outside good" are potentially time-varying, we capture them in estimation with month specific fixed effects (the μ_t below). This allows the characteristics of the outside good to vary from period to period in an arbitrary fashion.

We summarize the data used in estimation in table 2. We have 4752 unit-month observations, although around 23% of the time a unit gets a zero share. Shares of individual units are reasonably small (1.25% on average), although this is highly variable, with one unit getting 13% of the FR market in one month. The average bid is 7.8 £/MWh, but there are a number of outlying bids (as high as 238.6 £/MWh) that we will need to be thoughtful in handling.

Since we presumably have less information at our disposal than the firms do, we choose to

err on the side of flexibility in our modeling. We allow for unit fixed effects (the γ_j below) to capture the preferences of the grid operator for particular units as well as transmission constraints. Each unit has two time-varying characteristics that capture the main forces at work in the market: their bid $b_{j,t}$ and their average position in the main market $m_{j,t}$. The latter is defined as the (average) ratio of the unit's operating position to its maximum production level, where the average is across half-hourly periods in month t. From our conversations with market participants, we know that these are the most important time-varying determinants of demand.

Let $s_{j,t}$ be the share of unit j in month t. The logit model implies that we have:

$$\log s_{j,t} - \log s_{0,t} \equiv \delta_{j,t} = \alpha \log b_{j,t} + \beta m_{j,t} + \gamma_j + \mu_t + \xi_{j,t} \tag{4}$$

To be flexible, we allow these error terms to follow an $AR(1)$ process:

$$\xi_{j,t} = \rho \xi_{j,t-1} + \nu_{j,t}$$

We assume that the innovation $\nu_{j,t}$ is iid and satisfies the exclusion restrictions:

$$E[\nu_{j,t+\tau} | b_{j,t}, m_{j,t}] = 0 \text{ for } \tau = 0, 1, 2, \ldots$$

This setup allows for firms to know $\xi_{j,t-1}$ before bidding (i.e. for bids to be correlated with $\xi_{j,t}$), but requires that firms do not know the innovation $\nu_{j,t}$ when they bid.

In the case where $\rho = 0$ (i.e. no autocorrelation), we can estimate equation (4) by OLS. But for $\rho > 0$, there is a potential endogeneity problem, since $\xi_{j,t-1}$ may be correlated with $b_{j,t}$ (units of high unobserved quality may bid more). To deal with this, we quasi-first-difference:

$$\delta_{j,t} - \rho \delta_{j,t-1} = \alpha(\log b_{j,t} - \rho \log b_{j,t-1}) + \beta(m_{j,t} - \rho m_{j,t-1}) + \widetilde{\gamma}_j + \widetilde{\mu}_t + \nu_{j,t} \tag{5}$$

where $\widetilde{\gamma}_j = (1-\rho)\gamma_j$ and $\widetilde{\mu}_t = \mu_t - \rho \mu_{t-1}$. We estimate this equation by non-linear least squares.[10]

As noted earlier, around 23% of our units get zero share in each period, so that $\delta_{j,t}$ is not well-defined in every period (and nor is $\delta_{j,t-1}$ for the quasi-first-differencing). We choose to

[10] We ignore an incidental parameter problem in the estimation of the model with $\rho > 0$. Following Nickell (1981), the bias in a linear (balanced) panel model with fixed effects is of the order ρ/T. The median number of time periods in which we observe a unit is $T = 58$, so the bias is probably relatively small.

Table 3: Demand System Estimates

	Log share ratio	
	OLS	QFD
Log bid	-1.690***	-1.684***
	(0.136)	(0.135)
Average position	2.538***	2.416***
	(0.144)	(0.122)
Unit and Month FE	yes	yes
ρ	–	0.41
s.e. ρ	–	0.03
R^2 (in shares)	0.52	0.66
N	3600	3319

In all regressions the dependent variable is the log ratio of the unit share to the outside good share (an observation is a bmunit-month), coded as missing where the share is zero and omitted in estimation. In the first column, the regression is by OLS; in the second column, the 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 (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

simply drop these observations in the remainder of the analysis, resulting in a smaller sample size. The main concern with doing so is that this potentially introduces a selection problem: in particular, it may be that units get zero share when they bid too high, and so our demand system understates the responsiveness of residual demand to price. In the appendix we show that this does not appear to be a significant problem, as the probability of getting a positive share is not significantly correlated with a unit's bid, over a large range of bids.[11]

Table 3 shows the estimated demand system parameters. The estimated price coefficient $\hat{\alpha}$ is negative and significantly less than -1. The coefficient on main market position $\hat{\beta}$ is significant and positive, consistent with the fact that only generators who are currently operating can supply frequency response. In the QFD specification we find some evidence of persistence in the unobservables, as $\rho > 0$. However, our estimate of α is not significantly different across the specifications.

To assess the fit of the demand system, we simulate out the shares $\{s_{j,t}\}_{j=1}^{65}$ in each period

[11] We do however find that extremely high bids ($b_{j,t} > 23$) are an almost perfect predictor of not getting called. Accordingly, in the rest of the paper we assume that demand is zero for bids in this range (implying, for example, that such bids are never optimal); and we drop these observations when inverting the first order conditions to obtain costs.

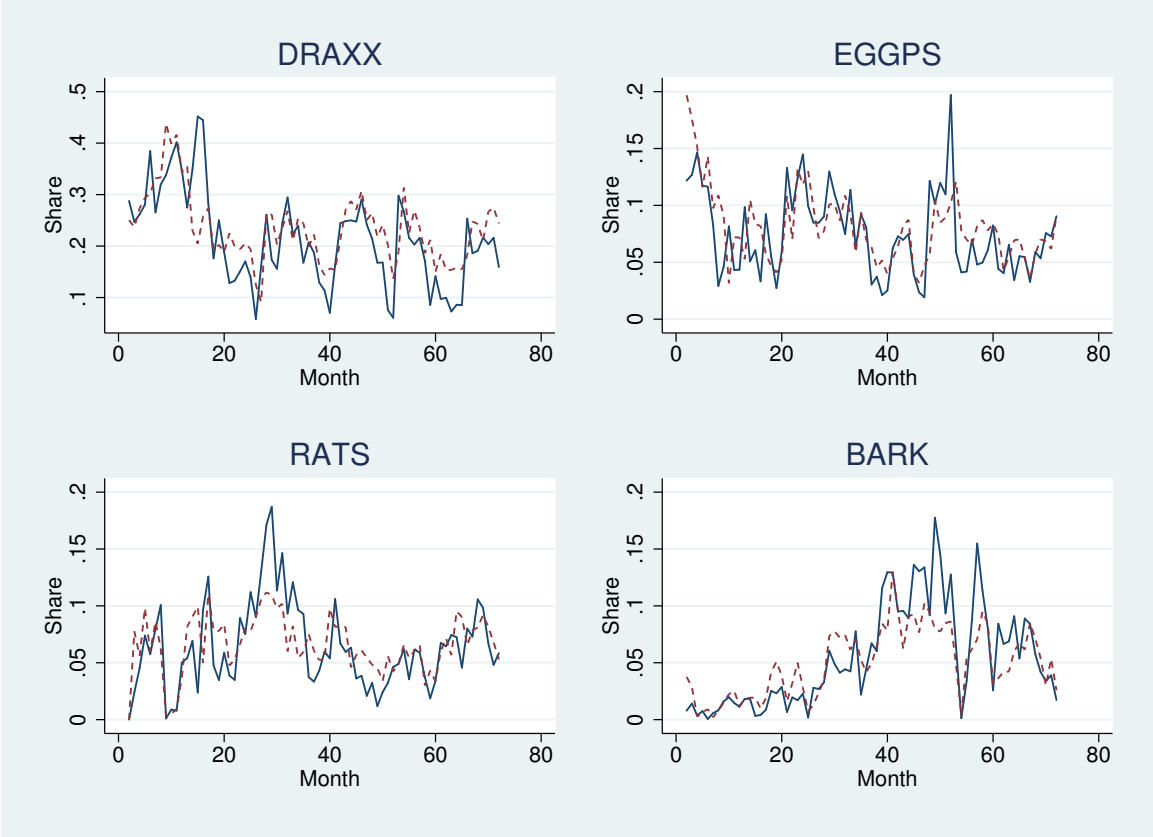


Figure 12: **Fit of the demand system.** Shown are the actual (blue, solid) and predicted (red, dashed) shares for the 4 biggest stations (Drax, Eggborough, Ratcliffe and Barking).

from $t = 2 \ldots 72$. Since each unit's share depends on the period-specific utility from all available options, this requires integrating out the vector of iid errors $\{\nu_{j,t}\}_{j=1}^{65}$, which we do by sampling independently and uniformly from the empirical distribution of residuals $\hat{\nu}_{j,t}$ in equation (5). The fit of the demand system is reasonably good. Comparing observed and predicted shares we get an R^2 of 0.66 in our preferred QFD specification.[12] Figure 12 shows that the fit is good even for the biggest stations, whose shares vary quite dramatically over time. This indicates that the good fit is not solely a consequence of having unit fixed effects.

[12]The R^2 is computed on the sample of units that have positive share, as we have no model for predicting zero share outcomes.

4.2 Costs

Now that we have an estimate of the demand system, we would like to infer the marginal costs of each generating unit. In many other electricity markets, a generator has to remove capacity from the main market in order to provide FR. As a result the main cost of participation in the FR market is the opportunity cost of lost production in the main market, and so the holding payments are a magnitude of order bigger. In the UK, however, it is NG's responsibility to create sufficient FR capacity by re-positioning units in the main market, and so firms can — and do — contract out all of their generating capacity in the forward market while still actively participating in the FR market. So the main source of marginal costs is wear and tear, which should be relatively stable. Thus we let each unit have a mean marginal cost of $\bar{c}_j$, and allow for time-varying iid errors $\varepsilon_{j,t}$:

$$c_{j,t} = \bar{c}_j + \varepsilon_{j,t} \tag{6}$$

To recover the mean costs $\bar{c}_j$ we take an approach that is standard in empirical IO: we invert the first order conditions. There are a couple of complications here. First, we believe the market is out of equilibrium for a large part of the sample period, and so we only perform this inversion during the "equilibrium period" (from month 43 onwards). Second, we don't want to commit to any one fully specified model of firm behavior, since testing these models against each other is the subject of the rest of the paper. Instead, we will simply assume that in equilibrium firms have rational expectations about demand (i.e. they make no predictable mistakes), an approach often taken in the macroeconomics literature.

To spell this out formally, we return to the equilibrium first order conditions (3). The quantity allocated to any given firm is equal to total market size Q_t times share: $q_j(b_t, x_t, \xi_t; \theta) = Q_t s_j(b_t, x_t, \xi_t; \theta)$. Because total demand for frequency response is determined exogenously by regulatory requirements, it is independent of the bids, so that $\partial Q_t / \partial b_{k,t} = 0$.[13] This means Q_t drops out of the first order condition in $b_{k,t}$ and we have:

$$\mathcal{E}_{i,t}\left[s_k(b_t, x_t, \xi_t; \theta) + \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_{j,t}) \frac{\partial s_j(b_t, x_t, \xi_t; \theta)}{\partial b_{k,t}} \middle| \Omega_{i,t} \right] = 0 \tag{7}$$

Let $s_j(b_t, x_t, \xi_t; \theta)$ be our estimated logit share function, which we will take to be the truth.

[13] Moreover it is presumably known to be independent of bids, so that $\mathcal{E}_{i,t}[\partial Q_t / \partial b_{k,t} | \Omega_{i,t}] = 0$ as well.

Recall that the characteristics x_t include the vector of availability for each unit m_t, as well as unit fixed effects. We now introduce assumptions on the information sets and the expectation operators.

Assumption 1 (Information Sets) $\Omega_{i,t} = \{c_{i,t}, \bar{c}_{-i}, x_t, D_t\}$ *where* D_t *is the publicly available data:*

$$D_t \equiv \{s_\tau, b_\tau, x_\tau, \xi_\tau\}_{\tau<t}$$

This assumption does a decent job of capturing the firms' information sets. Firms know their own period-specific costs when they bid, as well as the average costs of rival units, which seems reasonable. They also know everything that has happened in the past. It is a little misleading in two respects: first it assumes they know x_t when they bid, and thus implicitly their own future operating position and that of their rivals; and second it assumes they know the volumes for the current month when bidding on the next month, which is impossible since they bid during the middle of the month. We plan to fix this up in future versions of the paper. Next we make the rational expectations assumption:

Assumption 2 (Rational Expectations) *Firms have rational expectations in equilibrium:*

$$\mathcal{E}_{i,t}\left[s_j(b_t, x_t, \xi_t; \theta)|\Omega_{i,t}\right] = E\left[s_j(b_t, x_t, \xi_t; \theta)|\Omega_{i,t}\right] \ \forall i \ \forall t \geq 43$$

We show in the appendix that this implies the first order conditions should on average hold when evaluated at the demand realizations, where the averaging is over time:

$$E\left[s_k(b_t, x_t, \xi_t; \theta) + \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_{j,t}) \left.\frac{\partial s_j(b, x_t, \xi_t; \theta)}{\partial b_{k,t}}\right|_{b=b_t}\right] = 0 \tag{8}$$

The basic idea is that although the demand realizations may be different from what any one firm expected (the vector of rival bids depends on rival cost shocks; demand depends on realizations of ξ), the difference between the realization and the expectation must be independent of the bid, and thus average out over time.

This motivates two different estimation approaches. The simpler one is to solve the system of equations (8) separately by firm-month to get a vector of implied costs $c_{i,t} = \{c_{j,t}\}_{j \in \mathcal{J}_i}$. These cost estimates will be somewhat noisy because of both the cost shocks ε_t and differences

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

Station	# Units	# Obs	Fuel	Vintage	Mean	Std. Dev.	GMM Estimate
Barking	2	59	CCGT	1994	1.39	.17	1.39
Connah's Quay	4	109	CCGT	1996	1.17	.18	1.18
Cottam	4	87	Coal	1970	1.55	.19	1.55
Drax	6	173	Coal	1974	1.36	.4	1.4
Eggborough	4	108	Coal	1968	1.73	.21	1.74
Peterhead	1	29	CCGT	2000	1.72	.11	1.72
Ratcliffe	4	97	Coal	1969	1.54	.23	1.55
Seabank	2	58	CCGT	2000	1.77	.12	1.77

Summary statistics on the cost estimates derived from solving the firm first order condition arising from the demand system, separately by time period and firm. An observation is a cost estimate at the station-month level, obtained by averaging the unit-month specific cost estimates. Means and standard deviations of these cost estimates are reported separately by station. In the final column, we report the within-station average of the unit-specific cost estimate obtained from the GMM cost estimation approach, to show that it is similar.

between expected demand and its realization. We average those costs over time for each unit to get an estimate of the mean cost $\bar{c}_j$ for each unit.

The problem with this approach is that (8) is only guaranteed to hold when we average across many periods, and if we use a single firm-month, we're essentially not averaging at all. So as an alternative second approach, we pool all the data from a single firm during the equilibrium period, and estimate the costs $\bar{c}_i$ by minimizing the square of the empirical analogue of (8) (i.e. the average first order condition for firm i over the equilibrium period). This is equivalent to GMM with the constant as an instrument.

Table 4 shows our cost estimate for the top 8 stations (ranked by revenue).[14] Both approaches produce similar estimates of the mean costs. Moreover, the estimates are pretty reasonable: they vary between 1.39 and 1.77 for the top 8 stations, and are on average 1.54 across the whole sample. By comparison, the "cost reflective" regulated prices before CAP047 were around 1.7.[15] Since one would have expected some markup to be built into the regulated

[14]As mentioned earlier, we only include units when the firm received a positive share, as the FOC is only required to hold on the interior of the support of bids that are accepted. As a result the number of observations for each firm is not exactly equal to number of units times number of periods (30), although for many stations the total number of observations is pretty close.

[15]We have two sources for this: figure 2, which shows holding payments in the high market alone, and a policy document prepared by NG for Ofgem prior to the implementation of CAP047 (`https://www.ofgem.gov.uk/ofgem-publications/62273/8407-21104ngc.pdf`). The latter notes in paragraph 5.3 that the current response prices are "of the order of £5/MWh", for the bundle of primary, secondary and high, implying an average price of 1.67 for a single unit of response.

Table 5: Projecting costs onto unit characteristics

	Cost estimate		
Unit vintage	-0.013*	-0.013*	-0.013*
	(0.007)	(0.007)	(0.007)
Dual Fuel	-0.617***	-0.617***	-0.543**
	(0.207)	(0.208)	(0.208)
Large Coal	-0.150	-0.150	-0.169
	(0.183)	(0.184)	(0.177)
Medium Coal	-0.507**	-0.506**	-0.431**
	(0.205)	(0.207)	(0.203)
Oil	-0.683***	-0.680***	-0.618***
	(0.136)	(0.138)	(0.131)
Pumped Storage	0.163	0.162	0.164
	(0.226)	(0.230)	(0.223)
Fixed effects	none	month-of-year	month-of-year × fuel type
R^2	0.15	0.16	0.17
N	1531	1531	1531

The dependent variable is the cost estimate $c_{j,t}$. Standard errors are clustered by bmunit. The omitted fuel type is combined cycle gas turbines (CCGT). Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

prices, our costs are in the right ballpark. This is evidence that the market is indeed in a static Nash equilibrium during the final phase of the data, as we have assumed.

We also project our estimated costs $c_{j,t}$ onto unit characteristics, showing the results in table 5. We show that the (typically larger) CCGT or Large Coal plants have higher costs than plants with other fuel types, while units of later vintage have lower costs (although the latter result is only statistically significant at 10%). In the second column we add month-of-year fixed effects, and in the third column a full set of interactions between fuel type and month-of year. While jointly statistically significant, the estimated fixed effects are small and add little to the fit of the model. This is further evidence against significant interaction between the bids in this market and the economics of the main electricity market, since the main market is certainly seasonal — see figure 5.

5 Learning

We are now in a position to start evaluating different learning models. Most sensible models will make similar predictions during the equilibrium period. Indeed, they should do very well at predicting behavior since the mean costs were estimated to fit bidding patterns during the equilibrium period, and most models will generate first order conditions similar to (8). And without introducing further heterogeneity across firms (e.g. in their prior beliefs on the price elasticity), we will struggle to explain the divergent strategies during the early period of the data. For that reason we will compare their predictive abilities during the middle period, where firms appear to start to compete more aggressively on price.

5.1 Learning models

As we said earlier, firms face two kinds of uncertainty: they are uncertain as to how their competitors will play, and also as to how NG will trade off price against other factors, such as main market position. This leads to uncertainty about the demand for their services. In line with this, we allow for a set of models in which firms know the functional form of the demand system $s(b, x, \xi; \theta)$ as well as the stochastic process for the demand unobservables ξ, but are uncertain about competitor behavior and potentially also the demand parameters θ.

We model strategic uncertainty through the well-known model of *fictitious play* (see e.g. Fudenberg and Levine (1998)). This is one of the leading strategic learning models (though there are important alternatives such as reinforcement learning and experience-weighted attraction learning). In this model, firms believe that rival bids are sampled from the empirical distribution of their play in the past, with sampling weights that typically decrease as one goes back in time. For example, the dynamic best response model is one in which each player believes that each opponent will play their last action with certainty (i.e. they believe that $b_{m,t} = b_{m,t-1}$ for all $m \notin \mathcal{J}_i$). This can lead to interesting dynamics, which may or may not converge to equilibrium.

A more general formulation allows sampling weights that geometrically decay (i.e. $w_{t-k} = \delta^k / \sum_{l=1}^{t-1} \delta^l$). Dynamic best response is then a limit case as $\delta \to 0$. Another interesting limit case is the model with $\delta = 1$, so that all past observations are equally weighted. We consider the two limit cases as well as the case $\delta = 1/2$ below.

While the above notion of fictitious play is conceptually well-defined for 2-player or popu-

lation games, it requires additional assumptions on the sampling procedure with multiple distinct opponents. We assume that the fictitious joint distribution of rival bids is a product of the marginal distributions (i.e. independent sampling). This seems natural, as in a stable and competitive environment, firms should not expect rival actions to be coordinated.

All of these fictitious play models assume that firms know the demand system parameters θ. But as we have argued above, it is quite possible that the firms are learning the price sensitivity α (and possibly other demand parameters) over time. So we consider an additional set of models with both strategic and demand uncertainty, where firms engage in *statistical learning*: they estimate α_t according to the best information at their disposal as of time t, and then use this estimate in choosing their bids.

We operationalize this idea as follows. Recall that the firms are assumed to commonly observe the public data $D_t = \{s_\tau, b_\tau, x_\tau, \xi_\tau\}_{\tau<t}$. We estimate the demand system using the dataset D_{T+1} (i.e. all the data), but the firms must make do with less. The predictable part of demand is a function of bids, main market position, lagged unobservables and unit fixed effects (as well as a monthly fixed effect that controls the size of the inside-good market). We assume that the firms know the fixed effects from the beginning of the market, and are only learning about the price parameter α (statistical learning model 1), or about both the price parameter α and the main market position parameter β (statistical learning model 2).[16] It is natural that firms are learning α, since there wasn't any price variation before. It is less obvious why they should be learning β, since the firms should know NG's responsiveness to availability from the pre-CAP047 period. But we find that the model fits better when we allow for this; more investigation of this point is needed.

In concrete terms, when firms are learning both α and β we assume that their estimates $(\hat{\alpha}_t, \hat{\beta}_t)$ are the OLS estimates of α and β in the regression equation:

$$y_{j,t} = \alpha \tilde{b}_{j,t} + \beta \tilde{m}_{j,t} + \nu_{j,t}$$

estimated on the dataset D_t, where $y_{j,t} = \delta_{j,t} - \rho\delta_{j,t-1} - \tilde{\gamma}_j - \tilde{\mu}_t$, $\tilde{b}_{j,t} = b_{j,t} - \rho b_{j,t-1}$ and $\tilde{m}_{j,t} = m_{j,t} - \rho m_{j,t-1}$. When firms learn α only, we run the same regressions imposing the

[16] It may seem odd to assume that firms know the time fixed effects for the previous periods. What we have in mind is that the time fixed effects stand in place of a more detailed model of the latent utility of the outside good. Part of this is — like the unit fixed effects — known to the firms; while part is uncertain (e.g. the effect of price changes in the outside good). So without further modeling, we can either assume they know more than they should (they know the time fixed effects), or less (they estimate the time fixed effects). We have started with the simpler case, but we plan to return to this point.

constraint that β is equal to our estimated value. We will refer to these demand estimates as the "sequential" demand estimates in what follows.

Notice that knowing the parameter estimates at each point in time would not be enough to guide the behavior of a sophisticated Bayesian firm: they would want to know the entire posterior distribution on the parameters. We assume that our firms are more naïve, and have beliefs that place a point mass on the estimated parameters (i.e. ignore residual parameter uncertainty). This is a simplification that reduces the computational complexity of the simulation exercises that follow.

Taking the three fictitious play models ($\delta = 0, 0.5, 1$) with the three statistical learning models (no learning, learning about α and learning about α and β), we obtain 9 models to take to the data (although we only estimate a subset of them).

5.2 A general approach to predicting bids

Now that we have obtained estimates of the primitives, we can start to compare the fit of different learning models. Each learning model M implies a particular form for the operator $\mathcal{E}_{i,t}$. This pins down a bidding function $\phi^M_{i,t}(c_{i,t}|\Omega_{i,t})$ via the first order conditions (3). For example, in the case of best response with known demand parameters we have $\phi^{BR}_{i,t}(c_{i,t}|\Omega_{i,t})$ as the solution $b_{i,t}$ of:

$$E\left[s_k(b_{i,t}, b_{-i,t-1}, x_t, \xi_t; \theta) + \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_{j,t}) \frac{\partial s_j(b_{i,t}, b_{-i,t-1}, x_t, \xi_t; \theta)}{\partial b_{k,t}} \middle| \xi_{t-1}\right] = 0$$

This simply replaces the unknown rival bids with last period's bids, implying that the only remaining source of uncertainty is the demand unobservables ξ, which are integrated out. We can analogously get bidding functions for all the different models by replacing rival bids with an expectation over the empirical distribution of their past bids (fictitious play) or replacing the estimated demand parameter θ with estimates $\hat{\theta}_t$ (statistical learning).

Since we don't know the time-varying costs $c_{j,t}$ for each unit — and since the costs we recovered in the earlier step conflate cost shocks and differences between expected and realized demand — we will make our predictions using estimated mean costs $\bar{c}_j$. As a result, we don't expect any of the models to fit perfectly — even during the equilibrium period — and our focus will mainly be on their relative fit.

5.3 Evaluating fit

We follow the experimental literature in testing the fit of the model with "one-step-ahead" predictions, comparing the observed and predicted bids for each time period t, where the predicted bids use (for example) data on the observed bids at $t-1$ (see e.g. Erev and Roth (1998)). This is a natural test, as it corresponds to the thought experiment of predicting the next move of a player in a game. This should be contrasted with "k-step-ahead" predictions, in which we would compare observed and predicted bids at t, where the predictions are based on forward simulation of a model from period $t-k$. We hope to return to this latter exercise in a later version of this paper, as this would be a good test of whether these models are useful in making counterfactual predictions of entire price paths.

We evaluate the fit of the various models in three ways. The most direct is graphical. Examining the fit for each unit separately is impractical, and so we compare the volume-weighted average bids under each model with the observed price path (the weight for each unit in each time period is just their realized share $s_{j,t}$). As mentioned earlier, we do this only for the middle period of the data, where we believe we have the best chance of explaining the patterns in the data with an "off-the-shelf" model.

A second measure of fit is an R^2, defined as usual as $R^2 = 1 - RSS/TSS$, where the RSS is $\sum_{t=17}^{42}(b_{j,t} - \hat{b}_{j,t}^M)^2$ and the TSS is $\sum_{t=17}^{42}(b_{j,t} - \bar{b})^2$ where $\bar{b}$ is the mean bid over this period. This tests whether the model does a good job of fitting the bids by a randomly chosen unit, but it may be more interesting to look at a measure that weights bigger units more. For this reason we compute a second R^2 measure along the same lines, replacing the RSS with $\sum_{t=17}^{42}(b_{j,t} - \hat{b}_{j,t}^M)^2 s_{j,t}$ and the TSS with $\sum_{t=17}^{42}(b_{j,t} - \tilde{b})^2$ where $\tilde{b}$ is the share-weighted average bid over this period.

A final measure is the likelihood of observing the data given the model. This requires that make some additional parametric assumptions. We assume that the cost shocks $\varepsilon_{j,t}$ are iid normally distributed, with parameters $\theta_j^c = (\bar{c}_j, \sigma_j^2)$, where the standard deviation is assumed equal to sample standard deviation of the $\hat{c}_{j,t}$. In this case we can compute the log-likelihood of observing a bid vector $b_{i,t}$ under model M as:

$$\ell^M(b_{i,t}) = \log \ell^* \left(\beta^{-1,M}(b_{i,t}|\Omega_{i,t}, \theta_j^c)\right) - \log|\det[J^M]|$$

where ℓ^* is the multivariate normal likelihood function and J^M is the Jacobian of β^M.

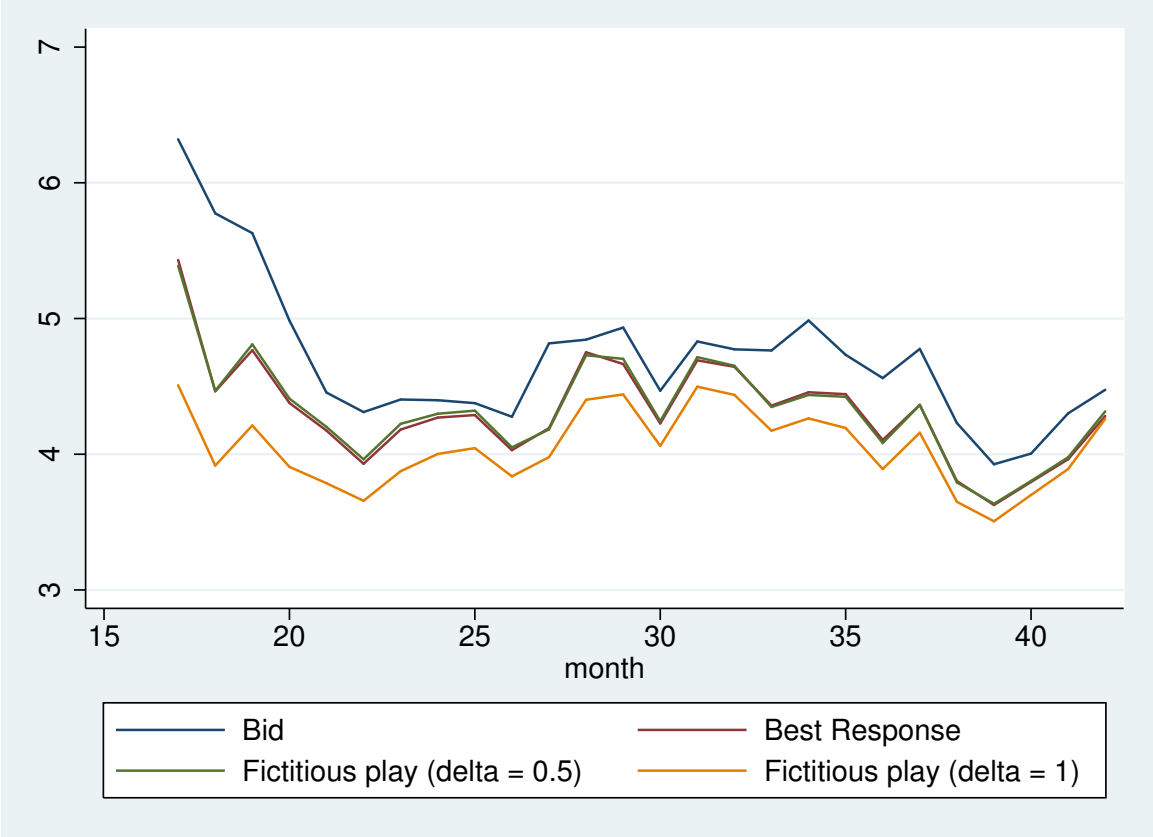


Figure 13: **Predicted bids: different strategic models.** Shown are the actual bids, and predicted bids from three different strategic models: fictitious play with $\delta = 1$ (i.e. equal weighting of past bids); fictitious play with $\delta = 0.5$; and best response (i.e. $\delta = 0$). The bids are weighted by contemporaneous shares to produce a single price path for each model.

An observation is a firm-month pair. Under our assumptions on the DGP, the observations are independent since differences between the observed and predicted bids by a firm are solely attributable to the realizations of the iid cost shocks $\{\varepsilon_{j,t}\}_{j \in \mathcal{J}_i}$. Thus the log likelihood of the data is the sum of the individual log-likelihood contributions across firm-month pairs.

An advantage of the likelihood-based fit measure is that it opens up possibilities for testing whether model A fits better than model B in a statistically significant sense. We apply the testing framework of Vuong (1989) to test pairs of models by comparing the likelihood ratio of the two models with the standard deviation in the likelihood ratios across observations (much like a paired 2-sample t-test).

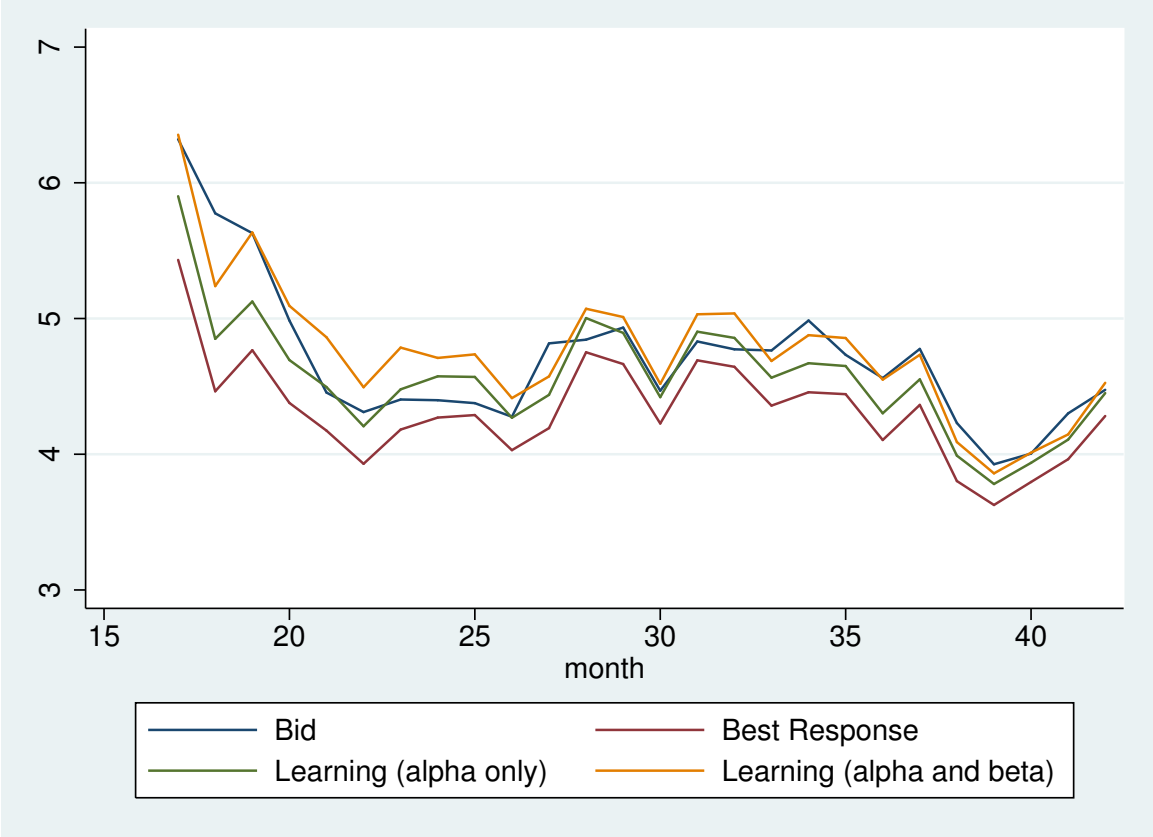


Figure 14: **Predicted bids: different models of demand uncertainty.** Shown are the actual bids, and predicted bids from three different models of demand uncertainty: best response with known demand parameters, best response with statistical learning of α and best response with statistical learning of α and β. The bids are weighted by contemporaneous shares to produce a single price path for each model.

5.4 Results

Figure 13 shows the comparative fit of the fictitious play models with no statistical learning. All of the models under-predict bids during the middle period. The best response and fictitious play model with $\delta = 0.5$ appear basically indistinguishable, whereas the fictitious play model with $\delta = 1$ performs significantly worse. This is evidence that firms pay relatively more attention to recent bids by their opponents.

The graphical analysis is backed up by the other measures of fit shown in Table 6. We cannot statistically distinguish between the first two models, but the model with $\delta = 1$ is worse (p-value ≈ 0). Notice also that all three models have *negative* R^2. This occurs because their

Table 6: Fit of the Learning Models

Strategic Model	Statistical Learning	R^2 (unweighted)	R^2 (share-weighted)	Log likelihoood
Best Response	α and β	0.10	0.37	-8023.92***
Best Response	α	0.07	0.36	-9077.53***
Fictitious play ($\delta = 0.5$)	none	-0.05	0.25	-11263.89
Best Response	none	-0.05	0.25	-11289.96***
Fictitious play ($\delta = 1$)	none	-0.27	-0.06	-13186.57

Fit of the different learning models, arranged from best-fitting to worst. For the fit measures based on R^2, an observation is a unit-month during the second period of the data (between months 17 and 42 inclusive). The unweighted R^2 is equal to one minus the ratio of the residual to the total sum of squares, while the weighted measure weights each observation by its share $s_{j,t}$ in computing the RSS and TSS. For the log likelihood measures, an observation is a firm-month during the second period of the data, and the latent cost shocks are assumed to be normally distributed, independently and identically across units and months. We test whether the fit of adjacent models is statistically different using the likelihood ratio test for non-nested models of Vuong (1989): asterisks next to the likelihood indicate whether the model fits significantly better than the model immediately below it (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

prediction of the mean bid is too low, whereas the total sum of squares is at least computed around the correct mean. In this sense the fit of the models is not that bad; while the mean bid level is too low, many of the pattens in the data are replicated by the models. Evidence in favor of this is given by the share-weighted measures, which are positive for the first two models.

The statistical learning models perform much better (see figure 14). In particular, the model in which firms learn both α and β does a very good job of reproducing the observed price path, and has a share-weighted R^2 of 0.37 — extremely respectable for a model with no parameters! Another way of interpreting this is that the out-of-sample fit is pretty good, which is surprising given that the middle period is so different to the equilibrium period. This is statistically the preferred model.

We have to be a little careful in how we interpret these results. It seems reasonable to conclude that firms pay more attention to recent behavior by their competitors, and that we are unable to distinguish the different fictitious play models much further. Here we have ten big players, and since the individual bidding idiosyncrasies of rivals may "average out", this may not be an environment in which it is profitable for firms to do more than watch for trends in the overall behavior of their opponents.

We also see that the statistical learning model fits significantly better. This could be taken

as evidence that firms learn in the manner we suggest, which would be natural but also quite interesting. On the other hand, our logit demand system is an approximation to the true demand system, and so it may that while the true demand system is stationary, our best approximation to it is not. In particular, our sequential demand estimates may be better approximations to the demand system at that point in time than our final set of estimates, and so it is possible that the agents aren't learning about demand at all and we fit better with the statistical learning models because we get better approximations to the demand system. In this case, the strategic learning models alone are sufficient to fit the patterns in the data. Unraveling these different stories is something we hope to do in future versions of this paper.

6 Conclusion

[to be completed]

References

Ackerberg, D. A. (2003): "Advertising, learning, and consumer choice in experience good markets: an empirical examination," *International Economic Review*, 44(3), 1007—1040.

Bolton, P., and C. Harris (1999): "Strategic Experimentation," *Econometrica*, 67(2), 349–374.

Börgers, T., and R. Sarin (1997): "Learning through reinforcement and replicator dynamics," *Journal of Economic Theory*, 77(1–14).

Boylan, R. T., and M. A. El-Gamal (1993): "Fictitious Play: A Statistical Study of Multiple Economic Experiments," *Games and Economic Behavior*, 5, 205–222.

Brown, G. (1951): "Iterative Solutions of Games by Fictitious Play," in *Activity Analysis of Production and Allocation*, ed. by T. Koopmans. Wiley.

Camerer, C., and T.-H. Ho (1999): "Experienced-Weighted Attraction Learning in Normal Form Games," *Econometrica*, 67(4), 827–874.

Camerer, C., T.-H. Ho, and J.-K. Chong (2002): "Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games," *Journal of Economic Theory*, 104, 837–874.

Cheung, Y.-W., and D. Friedman (1997): "Individual Learning in Normal Form Games: Some Laboratory Results," *Games and Economic Behavior*, 19, 46–76.

Conley, T., and C. Udry (2010): "Learning about a new technology: Pineapple in Ghana," *American Economic Review*, 100(1), 35–69.

Covert, T. (2013): "Learning about hydraulic fracturing in the North Dakota Bakken Shale," working paper, Harvard University.

Dickstein, M. (2013): "Efficient Provision of Experience Goods: Evidence from Antidepressant Choice," working paper, Stanford University.

Erdem, T., and M. P. Keane (1996): "Decision-Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 15(1), 1–20.

Erev, I., and A. E. Roth (1998): "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *The American Economic Review*, 88(4), 848–881.

Fudenberg, D., and D. Kreps (1993): "Learning mixed equilibria," *Games and Economic Behavior*, 5, 320–67.

Fudenberg, D., and D. Levine (1998): *The Theory of Learning in Games*. MIT Press, Cambridge, MA.

Goldfarb, A., and M. Xiao (2011): "Who Thinks about the Competition: Managerial Ability and Strategic Entry in US Local Telephone Markets," *American Economic Review*, 101(5), 3130–3161.

Griliches, Z. (1957): "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica*, 25(4), 501–522.

Hart, S., and A. Mas-Colell (2000): "A simple adaptive procedure leading to correlated equilibrium," *Econometrica*, 68(5), 1127–1150.

Hortaçsu, A., and S. 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.

Milgrom, P., and J. Roberts (1991): "Adaptive and Sophisticated Learning in Normal-Form Games," *Games and Economic Behavior*, 3, 82–100.

Mirman, L. J., L. Samuelson, and A. Urbano (1993): "Duopoly Signal Jamming," *Economic Theory*, 3, 129–149.

Newberry, P. (2013): "An Empirical Study of Observational Learning," Working paper, Penn State economics department.

Nickell, S. (1981): "Biases in Dynamic Models with Fixed Effects," *Econometrica*, 49(6), 1417–1426.

Townsend, R. M. (1983): "Forecasting the Forecasts of Others," *Journal of Political Economy*, 91(4).

Vuong, Q. (1989): "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica*, 57(2), 307–333.

Zhang, J. (2010): "The Sound of Silence: Observational Learning in the U.S. Kidney Market," *Marketing Science*, 29(2), 315–335.

Appendix

A. Data Appendix

Data sources: The majority of the data on frequency response (FR) and the electricity main market comes from the operators of the transmission grid and the balancing mechanism. NG plc (NG) owns and operates the transmission grid between generators and distributors in the United Kingdom. Since a redesign on November 1, 2013, the data website of NG is available at `http://www2.nationalgrid.com/UK/Industry-information/Electricity-transmission-operational-data/Data-explorer/Outcome-Energy-Services/`.

All FR data is available under the tab "Frequency Response – FFR and Mandatory". We downloaded our FR data from a previous version of the NG data website. In those cases detailed below where the original data is no longer available on the NG data website, it is available from the authors on request. NG used to publish Seven Year Statements detailing their projections of energy supply and demand and upcoming challenges. These used to be available at `http://www.nationalgrid.com/uk/Electricity/SYS/archive/`.

- **Bids**: We obtained FR bid data directly from the NG data website.The relevant file is labeled "Prices". Currently, a version is available that starts in January 2007 and is updated every month. From the old version of the data website, we downloaded one file for the period from November 2005 to January 2010, and another file for January 2007 to July 2013. These files contain monthly bids (in £/MWh) by every BM unit with mandatory FR provision requirements separately for the market segments primary, secondary, and high. The combined data period from the two files is November 2005 to July 2013.

- **Capabilities**: We obtained FR capabilities data directly from the NG data website. The relevant file is labeled "Capabilities". Currently, a version is available that starts in January 2006 and is updated every month. From the old version of the data website, we downloaded one file for the period from November 2005 to January 2010, and another file for January 2006 to August 2013. The former file reports that November and December 2005 are not available, so only the latter file is relevant, since it contains all the data that is available. The file contains monthly response capabilities by every BM unit with mandatory FR provision requirements separately for the market segments

primary, secondary, and high. For the market segment primary, response capabilities in MWh are given at 0.2 Hz, 0.5 Hz, and 0.8 Hz, while for the market segments secondary and high, only response capabilities at 0.2 Hz and 0.5 Hz are listed. In each case, the column on the right represents the maximum over the operating range. These values are constant over the sample period for more than 80% of the generators. The data period is January 2006 to August 2013.

- **Volumes**: We obtained FR volume data directly from the NG data website. Unfortunately, the new data website no longer provides historic volumes, and only a file that holds volumes from August 2013 is available. We downloaded monthly volume files for November 2005 thru June 2013. Each of these files contains one month of daily holding volumes in MWh by every BM unit with mandatory FR provision requirements separately for the market segments primary, secondary, and high. The combined data period of these monthly files is November 2005 to June 2013.

- **Main market position**: Elexon publishes all messages submitted to the Balancing Mechanism Reporting System on a given day at `http://www.bmreports.com/`. An example for a daily file is `http://www.bmreports.com/tibcodata/tib_messages.2003-01-01.gz`. Each file collects the messages submitted as part of the balancing mechanism on a given day. These messages contain information on Final Physical Notification (FPN), Maximum Export Limit (MEL), Bid-Offer Data (BOD), or Bid-Offer Acceptance Level (BOAL) for typically a half-hour interval.

- **Fuel type**: We take fuel type information from appendix F1 of the Seven-Year Statement prepared by NG in 2011: `http://nationalgrid.com/NR/rdonlyres/3B1B4AE4-2368-4B6E-8D47211/NETSSYS2011AppendixF1.xls` The sheet "F-2", corresponding to table F.2, provides fuel type for every BM unit listed under the column "Plant type". For an additional eleven stations, we take information on fuel type from the Variable Pitch project: `http://www.variablepitch.co.uk/grid/`

- **Vintage**: We take fuel type information from appendix F1 of the Seven-Year Statement prepared by NG in 2011: `http://nationalgrid.com/NR/rdonlyres/3B1B4AE4-2368-4B6E-8D47211/NETSSYS2011AppendixF1.xls` The sheet "F-2", corresponding to table F.2, provides vintages for most BM units under the column "Commissioning Year". The cell is empty for almost all hydro plants, so we take this information from the website of

the British Hydropower Association: `http://www.british-hydro.org/` For an additional eleven stations, we take this information from Wikipedia (5), from press releases prepared by the respective operator (5), and the website `www.scottish-places.info` (1). We are missing vintage for FAWN-1, which is connected with the Esso refinery in Fawley.

- **Ownership**: After registration on https://www.elexonportal.co.uk/, information on the registered party is contained in the file "reg_bm_units.csv" available under "Operational Data" → "Registration Information" → "Registered BM units" or under `https://www.elexonportal.co.uk/REGISTEREDBMUNITS`. It is based on registration data at the Central Registration Agency and under "Party Name", it lists the registered party. We downloaded a version of this file on December 29, 2009, and July 15, 2013, but there were no conflicts.

Sample selection and data construction: The paper restricts attention to data for the time period November 2005 to November 2011. We include BM units if they provided positive FR volume in at least one of these months. We collapse volumes for the three market segments primary, secondary, and high into one by summing (daily) volume across segments and days to get a unit-month level observation. Thus:

$$q_{j,t} = \sum_{k=P,S,H} \sum_{d \in M_t} q_{k,d,j,t}$$

where M_t is set of days in month t, k indexes market segment and d indexes days. The monthly bids are constructed as quantity-weighted averages of segment bids, where the weights are constant and given by the overall volumes of the three segments over the sample period:

$$b_{j,t} = \left(\sum_{k=P,S,H} Q_k b_{k,j,t} \right) / Q$$

where $Q_k = \sum_j \sum_t \sum_{d \in M_t} q_{k,d,j,t}$ and $Q = Q_p + Q_s + Q_h$. For the demand system, we classify BM units into "inside goods" if they were owned by one of the ten largest parties, where size is measured by revenue, and aggregate them into one "outside good" otherwise. Shares for inside goods and the outside good are computed at the monthly level.

Variables used: The unit of observation is BM unit by month, and we use the following variables:

Table 7: Sources and Definitions of Variables used in the Analysis

Variable	Unit	Sample	Source	Definition
Bid	£/MWh	All	Bids	Monthly bid
Volume	MWh	All	Volumes	Sum of segment volumes
Average MEL ratio	Fraction in $[0, 1]$	$>$ 75% in merged_data	Main market position	Average MEL in given month divided by maximum MEL over sample period
Fraction on	Fraction in $[0, 1]$	$>$ 90% in merged_data	Main market position	Fraction of time unit position is strictly greater than zero according to BOAL information
Fuel type	Categorical	All but one BM unit	Fuel type	Type of fuel (e.g., Oil, Large Coal, OCGT)
Vintage	Year	All	Vintage	Commissioning year
Owner	Categorical	All	Ownership	Registered party

B. Selection on positive share

In the main text, we estimate the demand system on the subsample of units that have positive share in a period (i.e. $s_{j,t} > 0$). Here we offer evidence that this does not significantly bias our estimate of the price elasticity of demand. Table 8 shows the outcome of a probit where the binary dependent variable is an indicator for a unit getting positive share (equivalently positive volume) in period t, and the regressors are an indicator for the bid being above 23, the log bid, and the unit position. The coefficient on log bid is small both economically and statistically, so that a high bid does not appear to affect the probability of getting a positive share. The exception is a bid above 23; of the 78 cases where bids in this range were submitted, only 2 received positive share (and both were very small).

Table 8: Determinants of positive volume

	Indicator for positive share
Bid ≥ 23	-1.798**
	(0.775)
Log bid	-0.106
	(0.223)
Average position	3.016***
	(0.149)
Unit and Month FE	yes
N	4752

An observation is a unit-month, and the dependent variable is an indicator for a unit having positive volume. The regressors are an indicator for the unit's bid being higher than 23, the log bid, and the average position of the unit, where average position is the average over half-hour periods (in a month) of the unit's declared operating position (as a fraction of its maximum generating capacity). Standard errors are clustered by bmunit. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

C. Equilibrium first order conditions

By the rational expectations assumption 2, we may substitute out the firm operator $\mathcal{E}_{i,t}$ in (3) to get:

$$E\left[s_k(b_t, x_t, \xi_t; \theta) + \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_{j,t}) \frac{\partial s_j(b_t, x_t, \xi_t; \theta)}{\partial b_{k,t}} \middle| \Omega_{i,t}\right] = 0$$

Let the term in brackets be denoted $G_{k,t}(b_t, c_{i,t}, x_t, \xi_t; \theta)$, and let $\tilde{G}_{k,t}(b_t, c_{i,t}, x_t, \xi_t; \theta)$ be its realization (as in (8)). Let $u_{k,t} = \tilde{G}_k(b_t, c_{i,t}, x_t, \xi_t; \theta) - G_k(b_t, c_{i,t}, x_t, \xi_t; \theta)$. Then by the properties of conditional expectations, it must be that $u_{k,t}$ is mean zero, independent of anything in the information set $\Omega_{i,t}$ and (thus) independent over time. So we have $E[\tilde{G}_{k,t}(b_t, c_{i,t}, x_t, \xi_t; \theta)] = E[G_{k,t}(b_t, c_{i,t}, x_t, \xi_t; \theta)] + E[u_{k,t}] = 0$ for each k, which we can use to form moment conditions.