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 21, 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(Preliminary and incomplete, please do not cite)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. We provide an analysis of the evolution of the deregulated market from the date it started. Initial activity was volatile, with some firms exploring different prices, while others made few price changes. This was followed by a period in which prices fell and the variance in the cross sectional distribution of bids declined markedly. By the end of our study price changes had became relatively rare and small, consistent with convergence to a static Nash equilibrium. We examine how well models of learning do in predicting play during the period prior to convergence but after the initial volatility. Models where perceptions of competitors' play depend on past play suggest that firms' weight recent play disproportionately. We also find evidence of statistical learning about the underlying demand parameters conditional on competitors' play. A model that combines these two features fits quite well: it is able to explain 37% of the share-weighted variation in prices, even though none of the model parameters are chosen to fit the pricing behavior.

*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, Duke, 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. Firms entering these markets face considerable uncertainty. This is partly because overall demand is typically uncertain, and partly because the demand for their product depends on the decisions of competitors facing similar uncertainties. This paper explores how this uncertainty manifests itself in the bidding behavior of firms from "day one" of a new market.

An understanding of how firms form strategies in new markets — or indeed in existing markets following a sudden change in conditions — is necessary in any investigation of the implications of changes in market institutions. The empirical industrial organization literature has focused on simulation of new equilibria following counterfactual policy changes (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 might help narrow down the set of possibilities in cases where there are multiple counterfactual equilibria (see, e.g. Lee and Pakes (2009)).

There has been little empirical work on how new markets (or existing markets with sharp changes in the institutions which govern them) evolve, and consequently little field evidence to discriminate among models of how firms adapt to both the strategic and demand uncertainty presented by changes in the environment. 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 to bid into an auction market, setting the stage for price competition. We study what happened next.

Though the introduction of the auction essentially started a new market, it is a new market with features which should make it relatively easy to analyze. 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, so the rules of the game are reasonably clear. Moreover the same firms had been supplying frequency response in the regulatory period preceding the change, so the bidders had quite a bit of prior knowledge of the vagaries of demand. Finally 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.

The paper is primarily an attempt to bring descriptive evidence to the question of how markets adapt to changes in institutions. We show that the FR 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, price changes became infrequent and small, and the cross-sectional variance in prices had shrunk markedly. 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 in favor of the rest point being a static Nash equilibrium.

When the market starts up we see different firms following different strategies which do not seem to be coordinated in any way. The behavior of some looks to be exploratory, and when the exploration leads to large gains in returns, others seem to follow. There are also firms who hardly change their prices at all. After just over a year competition between firms drives the highest prices down, leading to dramatically lower variance in the cross-sectional bid distribution. We consider a series of learning models that might explain the convergence in this period. Recall that the firms may be uncertain about both the behavior of their rivals and 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 weight the past play of their competitors differentially for different lagged periods (in the theory literature, this is "fictitious play", as in Brown (1951)). In our second set of models we allow firms to be uncertain about the demand parameters as well as rival strategies. We model how firms form perceptions about demand parameters by a statistical learning process, assuming firms run regressions given the set of data available to them at each point in time. This approach originates in the macro literature on rational expectations (e.g. Townsend (1983), Marcet and Sargent (1989), Evans and Honkapohja (2001)).

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. By contrast, when we allow for statistical learning on the properties of the demand system given the bids of others, as well as use the fictitious play assumptions for the perceptions of competitors' bids, we obtain a much better fit. Perhaps not surprisingly, firms seem to optimize given the currently available information and their perceptions evolve over time.

In the context of the fictitious play model we reject a model in which firms think 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; in the extreme this leads to a best response model. A model with firms both re-estimating demand each period and best responding to their rival's actions fits the data quite well: the (volume-weighted) R^2 is 0.37, which is reasonably impressive considering that the demand and cost parameters were estimated in a previous step, and no parameters were chosen to maximize the fit of the learning process per se.

We view the current version of the paper as making two distinct contributions. The first is to carefully document what happens to prices and bidding over time when a new market is started in a fairly stable environment. Since it is common in empirical industrial organization to assume that an equilibrium will be re-asserted following a change in conditions, and yet not clear from the theory on strategic learning that this should be the case, it seems useful to bring empirical evidence to bear on this issue. We find that there is convergence to some sort of rest point, albeit only after three and a half years of monthly strategic interaction. This rest point is consistent with a static Nash equilibrium in bids. To the best of our knowledge, ours is the first paper to demonstrate convergence and describe how it occurs.

A second contribution is to test different learning models against each other by comparing their relative fit during the competitive period. Our results here are preliminary, but suggest that firms more heavily weight recent competitor actions and engage in statistical learning.

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)); usually in stable and known environments. 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. 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 are now 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. But recently Salmon (2001) has argued using Monte Carlo simulations that in fact there is an "identification failure": different models are hard to tell apart statistically because they make similar predictions.

A second, distinct, theoretical literature considers behavior when there is 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)).

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, big firms made bids that were best responses to rival play, but small firms failed to optimize fully (although their behavior improved over time). Goldfarb and Xiao (2011) show that managers with different experience and education levels make different entry decisions in local US telephone markets after deregulation. They rationalize this with a cognitive hierarchy model, in which more experienced managers think more steps ahead.

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).

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

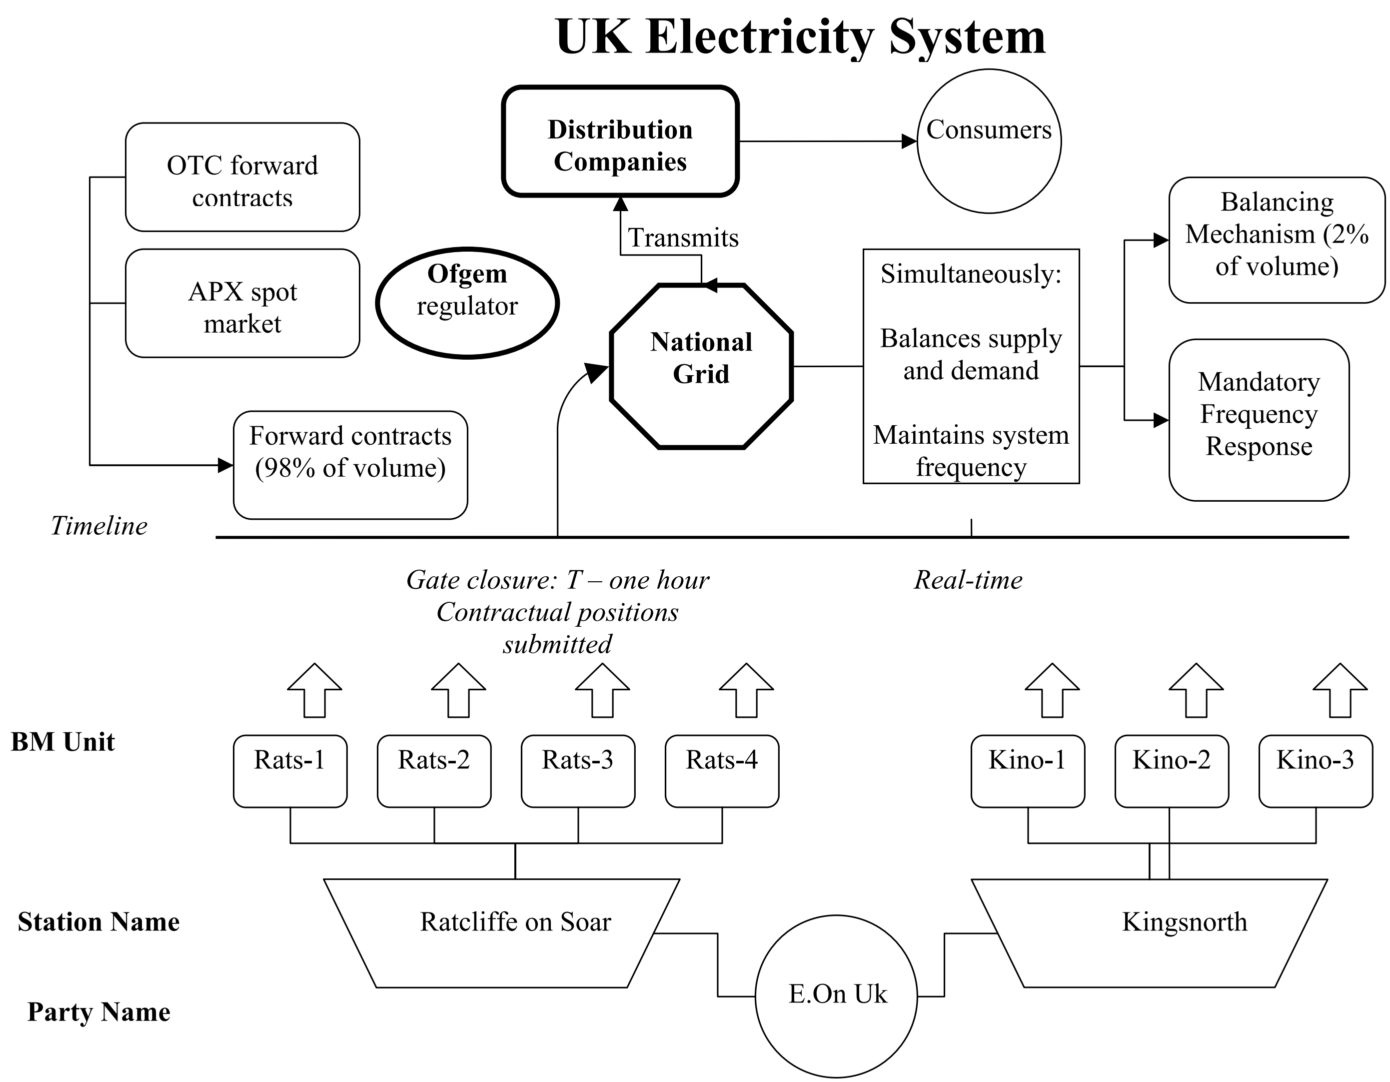

Figure 1: Overview of the British electricity market.

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 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 Con-

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

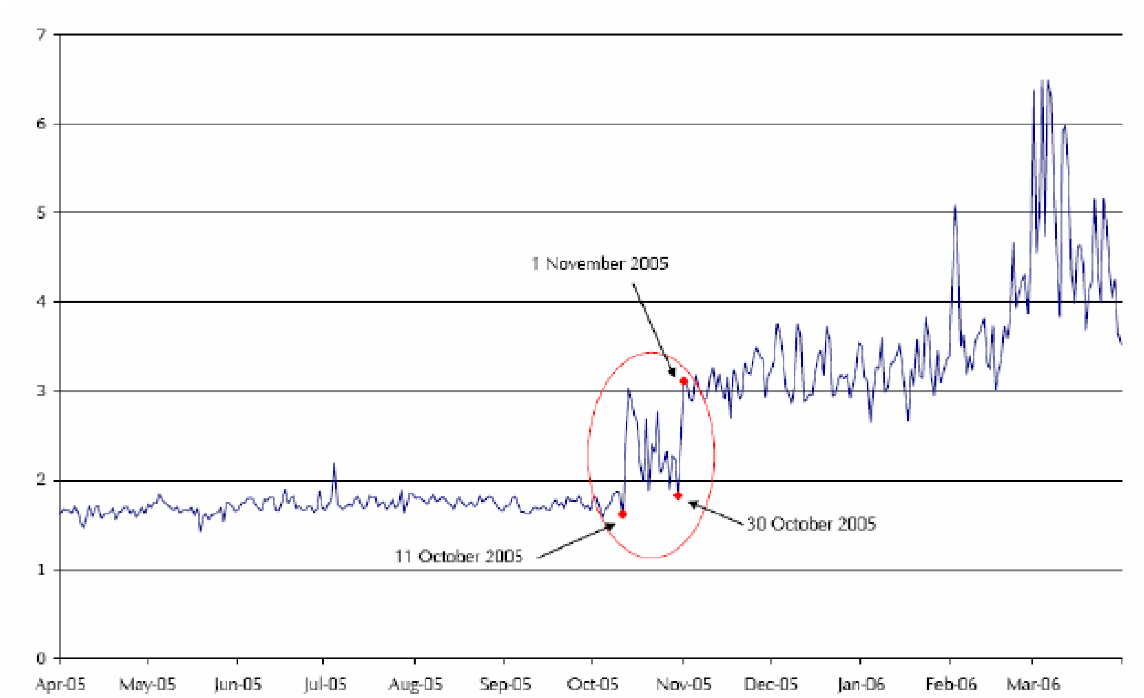

Figure 2: **Holding payments for high frequency response pre and post CAP047.** Source: National Grid.

nection 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.

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]

[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,

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 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.

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.

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.

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 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.

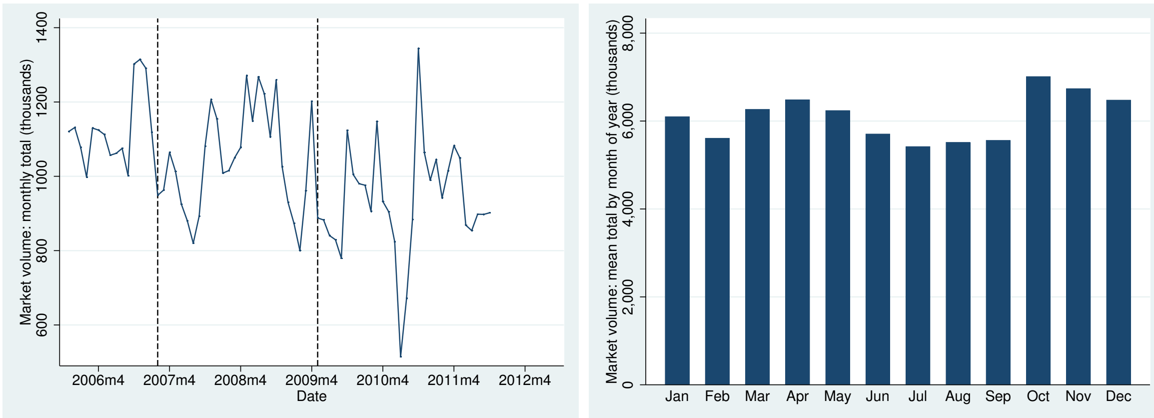

Figure 4: **Volumes over time and by season.** The left panel shows the total volume of frequency response for each month in the sample period, while the right panel shows the total volume by month-of-year (4 observations per month-of-year).

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.

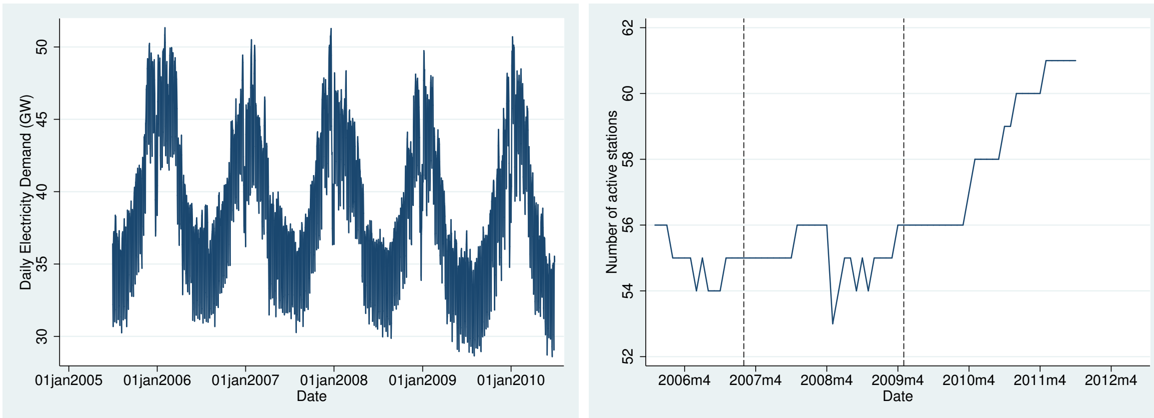

Figure 5: **Main market volume and number of active stations over time.** The left panel plots the total daily realized electricity demand in the UK over the sample period. This plots the number of active stations over time, where a station is "active" if one of its units submits a competitive bid (bids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

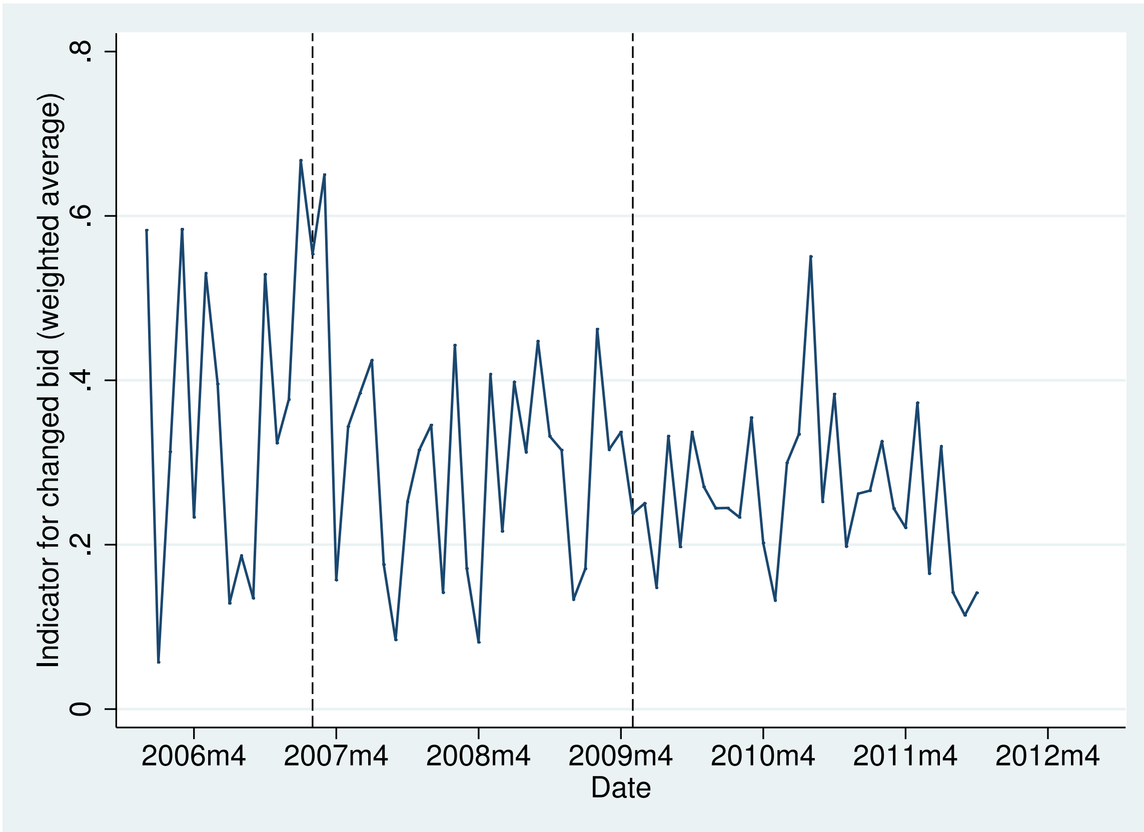

Figure 6: **Probability of a bid change over time.** This is a weighted average across units of an indicator for a change in their bid in period t, using quantity-weights from period $t-1$.

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.

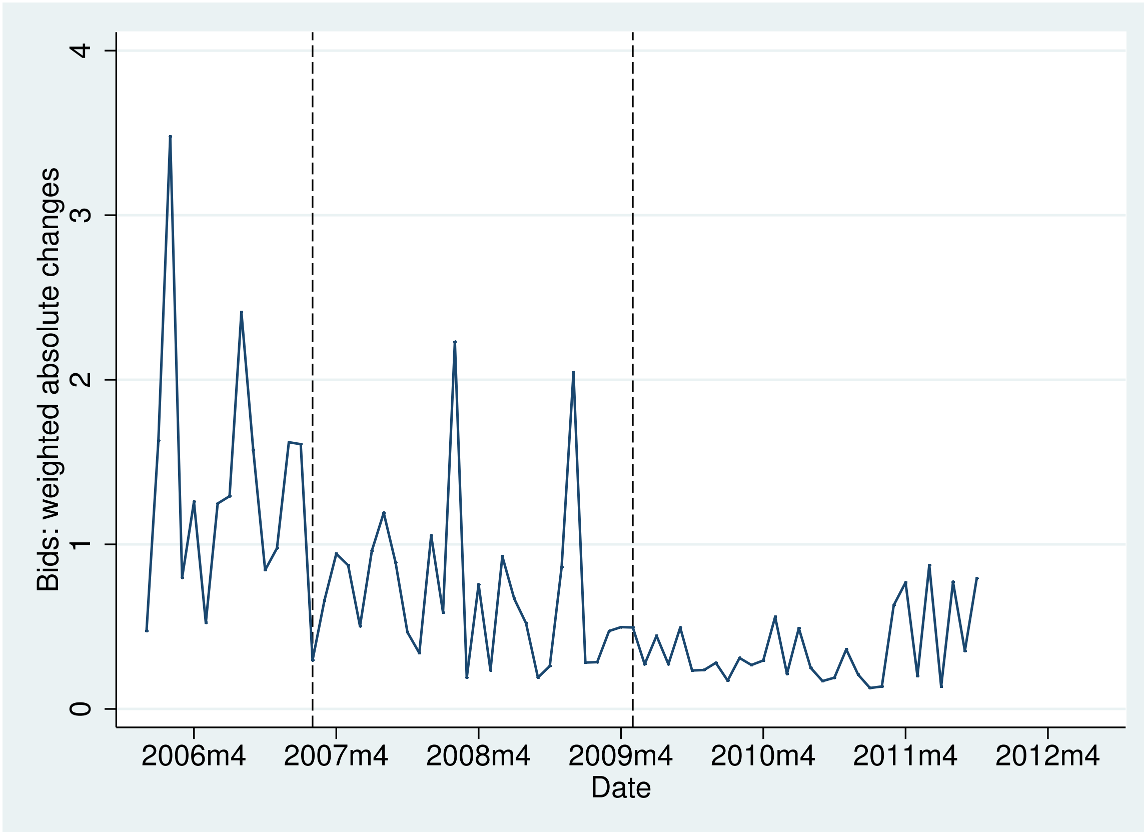

Figure 7: **Absolute bid changes over time conditional on changing.** This is a weighted average across units of the absolute change in their bid in period t, using quantity-weights from period $t-1$ and assigning a weight of zero if their bid did not change.

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

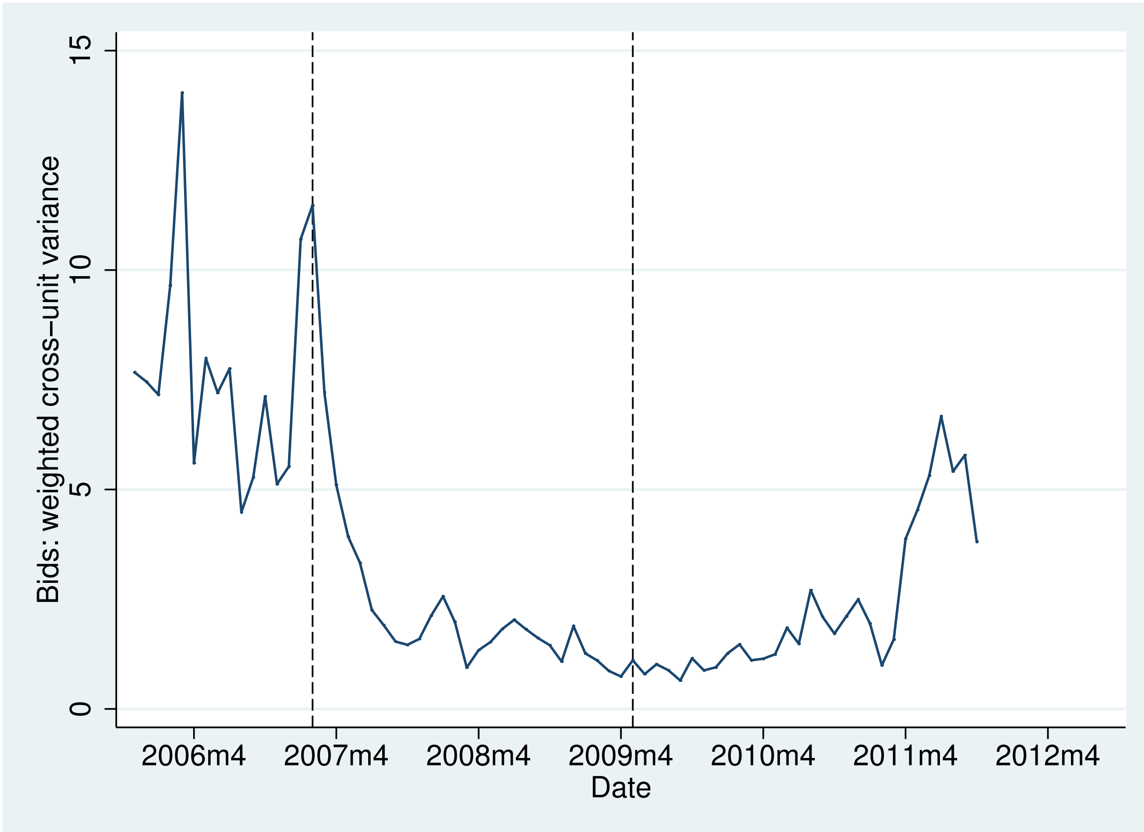

Figure 8: **Variance in bids across units.** Variance in the cross-section of bids at each point in time, where bids are quantity-weighted using contemporaneous quantities.

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.

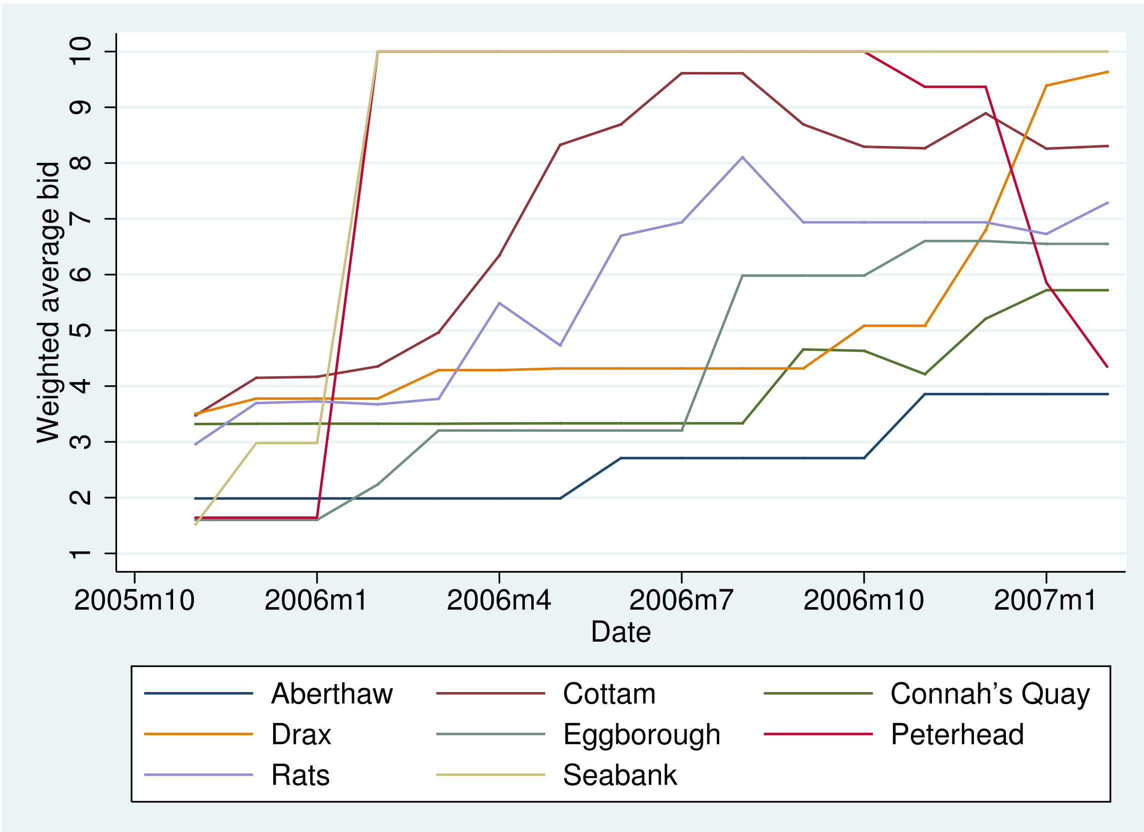

Figure 9: **Bids of the top 8 biggest stations (November 2005 – February 2007)**
Stations are ranked by total revenues in the pre-equilibrium period. Bids are censored above at 10 to compress the axes and improve the visual presentation.

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,

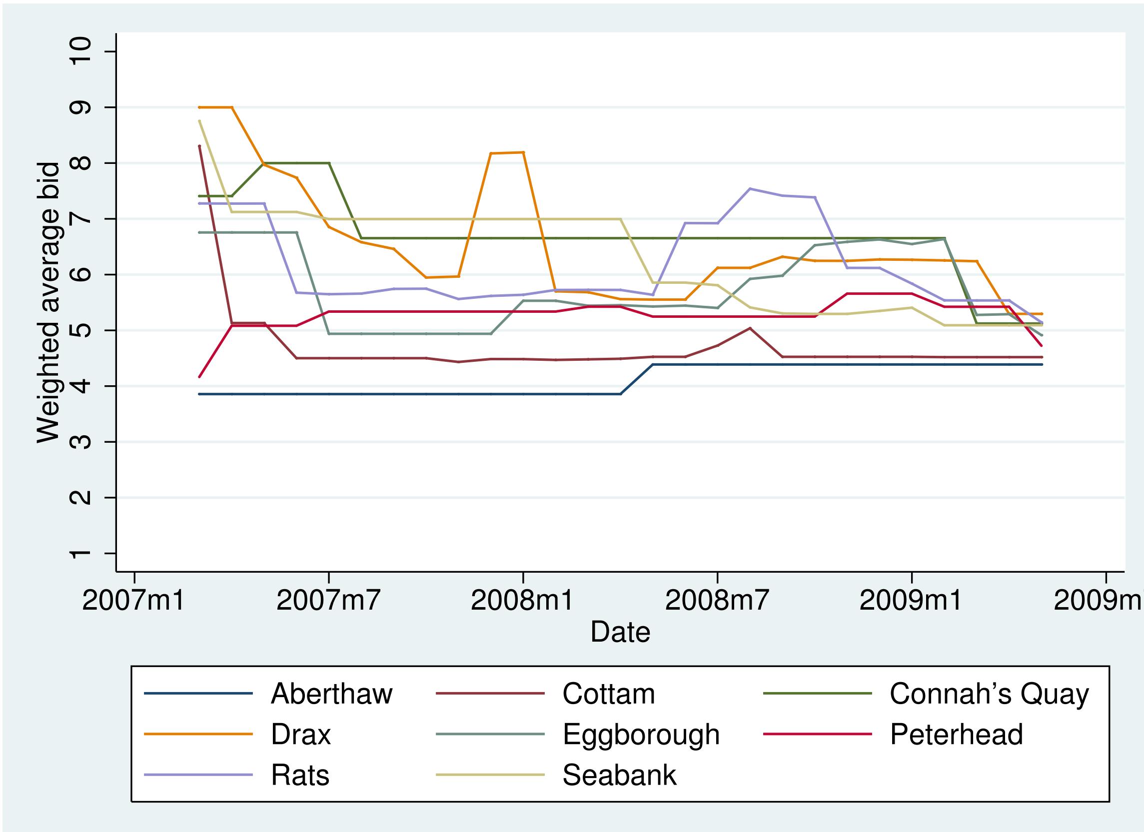

Figure 10: **Bids of the top 8 biggest stations (March 2007 – May 2009)** Stations are ranked by total revenues in the pre-equilibrium period

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

Summary. We have shown that the environment is stable. Despite this, the opening phase of the market was characterized by heterogeneity in bidding strategies. Over time, the bids grew closer, the cross-sectional bid variance declined, and the frequency of bid changes decreased. In the final period, the bids stabilized.

One interpretation of this final period is that the data has reached a static Nash equilibrium. Notice that in a stable game with a unique static Nash equilibrium, in every period players will take the same actions.[10] So the data from the final period is consistent with a static Nash. But it may also be consistent with other notions of equilibrium (e.g. collusion),

[10]With multiple equilibria, this statement remains true provided there is a deterministic and time-invariant equilibrium selection mechanism.

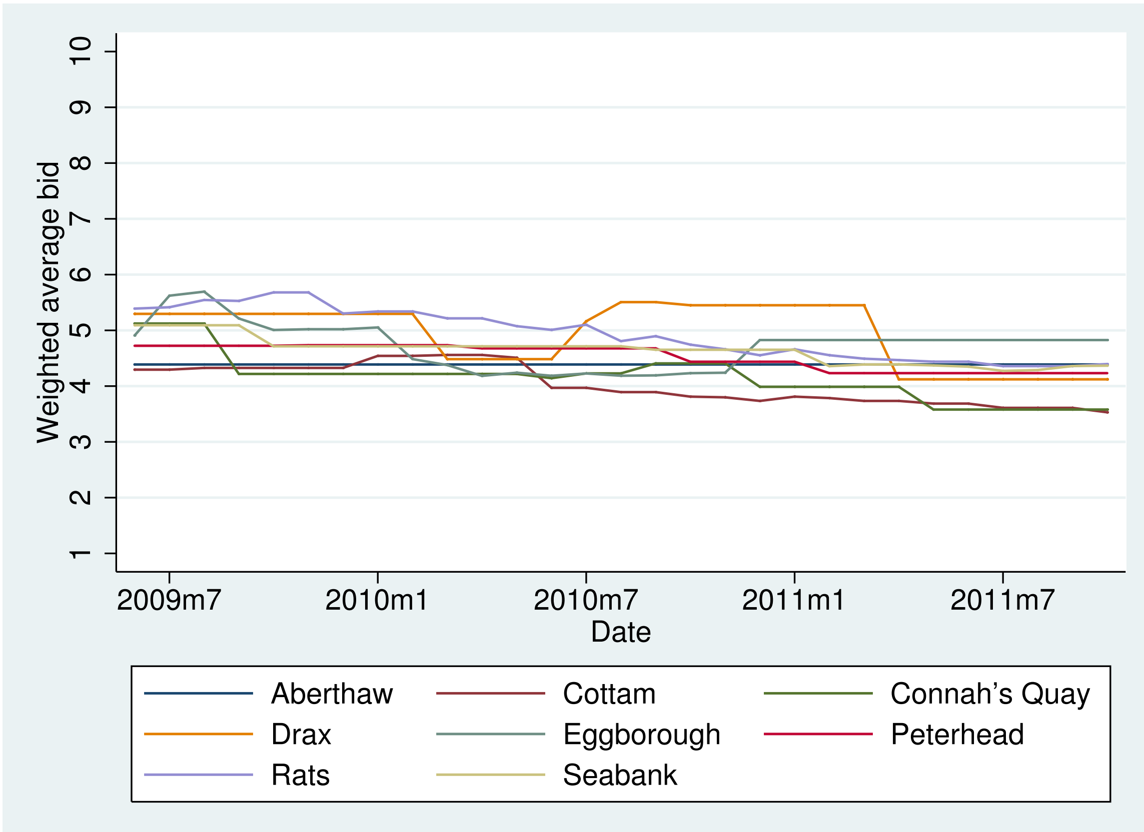

Figure 11: **Bids of the top 8 biggest stations (June 2009 – October 2011)** Stations are ranked by total revenues in the pre-equilibrium period

or even certain disequilibrium models. To shed further light on this, we need to directly examine whether firms are mutually best responding in a static sense, which requires us to start modeling and estimating the demand and cost primitives.

4 Estimating demand and cost primitives

Our goal is to understand how the market evolved. The prior data analysis described above indicates that it would be useful to divide the period of analysis into three sub-periods and allow firm behavior to be different in each. We treat the first period as a period of experimentation. We have already illustrated why we think this is the case, but a comment by Ian Foy reinforces our interpretation. His 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". The experimentation appears to be quite complex; apparently different firms had different strategies and some firms reacted to the experience of other firms as well as their own past experience. As a result further analysis of this period is beyond the scope of this paper.

We treat the second period as a period of learning about how best to maximize *current* profits; i.e. we treat it as a period of exploitation rather than experimentation. It is this period which we consider in light of the available learning models. The learning models require a demand system and the bidders' perceived costs of supplying FR. The demand generated by any vector of prices is generated by NG's complex computer program, and since, to our knowledge, this has not changed over the period, we form an approximation to it from the observed quantity responses over the entire period of our data. The perceived costs are not observable and we also have to estimate them before turning to an exploration of learning.

To obtain cost estimates we treat the third period as a stable period in which, at least on average, firms' bids do satisfy the first order condition obtained from maximizing expected profits. As a result we use the third period's best response condition to estimate costs. Those costs are then used in conjunction with the estimated demand system in our investigation of the second period's learning about how best to maximize profits.

4.1 The Demand System.

We approximate the demand system with a logit model for market shares. Were we not to use the logit structure, and instead use, say, a log-linear approximation to demand, it would require estimation of at least $(66 \times 65)/2 = 195$ own and cross-price elasticity terms, which is more than our data can handle. Also by using a model for market shares we eliminate the need for a model for the total quantity allocated. Market shares are more stable than total market size because the total market size varies by month of the year as illustrated in figure 4. Notice that by using a logit model for shares we have to explicitly deal with units that aren't active in a given month (i.e. have a zero share), as zero shares are impossible to rationalize in a standard logit model.

We estimate the logit demand model for monthly frequency response at the unit level focusing on 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). The units owned by smaller

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).

firms are aggregated into an "other" unit which is treated as the "outside" good. We allow the characteristics of this "outside good" to be time-varying capturing there effects with with month specific fixed effects (the μ_t below).

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.

Notationally, let i index firms, j index BM units, and t index months. In month $t-1$ firm i submits a bid $b_{j,t}$ for BM unit j in month t. 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 $b_t = (b_{i,t}, b_{-i,t}) = (b_{j,t}, b_{-j,t})$. Since we presumably have less information at our disposal than the firms do, we choose to err on the side of flexibility in modeling demand. We allow for unit fixed effects (the γ_j below) to capture both the preferences of the grid operator for particular units and 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.

If $s_{j,t}$ be the share of unit j in month t. Our logit model has

$$\log s_{j,t} - \log s_{0,t} \equiv \delta_{j,t} = \alpha \log b_{j,t} + \beta m_{j,t} + \gamma_j + \mu_t + \xi_{j,t}. \tag{1}$$

We allow these error terms, the $\{\xi_{j,t}\}$, to follow an $AR(1)$ process, or

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

where the innovations, the $\{\nu_{j,t}\}$, are mean independent of current and past bids and positions in the main market, and past disturbances. This setup allows the firms to know $\xi_{j,t-1}$ before bidding (the bids are made prior to the operating period), but requires that firms do not know the innovation $\nu_{j,t}$ when they bid.

If $\rho = 0$ (i.e. no autocorrelation), we can estimate equation (1) by OLS. However if $\rho > 0$, OLS will be biased since the $\xi_{j,t}$ may be correlated with bids. That is $b_{j,t}$ can be a function of $\xi_{j,t-1}$ which, in turn, is a determinant of $\xi_{j,t}$ (units of high unobserved quality may bid more). To deal with this, we quasi-first-difference to obtain

$$\delta_{j,t} - \rho\delta_{j,t-1} = \alpha(\log b_{j,t} - \rho \log b_{j,t-1}) + \beta(m_{j,t} - \rho m_{j,t-1}) + \widetilde{\gamma}_j + \widetilde{\mu}_t + \nu_{j,t} \tag{2}$$

where $\widetilde{\gamma}_j = (1-\rho)\gamma_j$ and $\widetilde{\mu}_t = \mu_t - \rho\mu_{t-1}$ and estimate by non-linear least squares.[11]

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 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. We address this problem in the appendix where we show that it does not appear to be significant. [12]

[11] We ignore an "incidental parameter" problem in the estimation of the model with $\rho > 0$. That is the "within" estimator generated by the BM unit specific fixed effect generates an error which is a function of mean of the quasi first difference in ξ' which in turn can be correlated with the quasi first difference in bids. 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.

[12] The probability of getting a positive share is not significantly correlated with a unit's bid, over a large range of bids. However 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

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 (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

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 from $t = 2 \dots 72$. 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 (2). 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.[13] Figure 12 shows that the fit is quite good even for the biggest stations, whose shares vary quite dramatically

conditions to obtain costs.

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

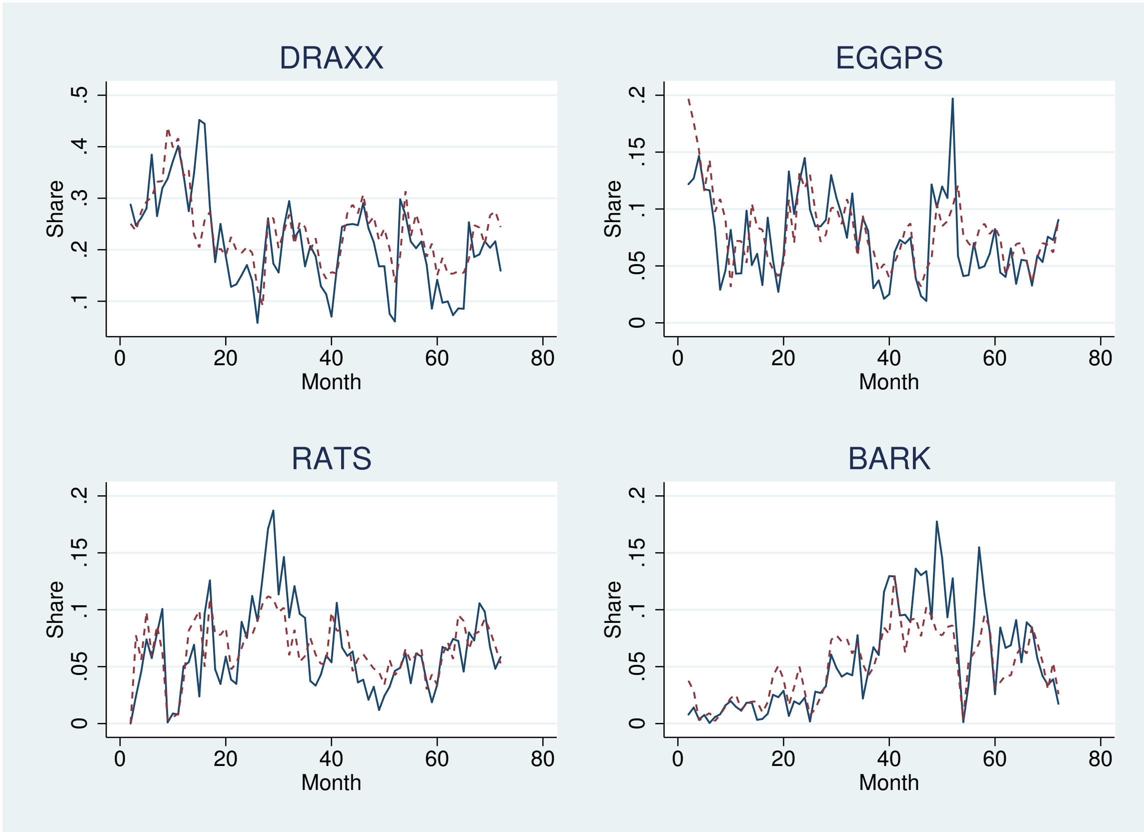

Figure 12: **Fit of the demand system.** Shown are the actual (blue, solid) and predicted (red, dashed) shares for the 4 biggest stations (Drax, Eggborough, Ratcliffe and Barking).

over time. This indicates that the good fit is not solely a consequence of having unit fixed effects.

4.2 Costs.

The next step is to use the data from the third and final period to estimate costs. We assume that each BM unit has a marginal cost of providing frequency response c_j. 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 over time for each unit.

The realized profit for firm i in period t is then equal to

$$\pi_{i,t} = \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_j) M_t s_{j,t}(b_t, x_t, \xi_t; \theta) \tag{3}$$

where M_t is the total market size in period t. In the next section, where we use the data from the second period to investigate how learning occurs, we will entertain the possibility that the firm knows the form of the demand function but does not know either the bids of its competitors, $b_{-i,t}$, or θ (the response of demand to its bid given $b_{-i,t}$). In this section we can be less specific, as all we will need is that the weak form of rationality introduced below holds in the final, or equilibrium, period.

To recover the costs c_j we take an approach that is standard in empirical IO: we invert the first order conditions, but only for our "equilibrium period" (from month 43 onwards). This requires bidding assumptions for that period, and we turn to those assumptions now. We let the information the agent has available to form its expectations of the profits that would accrue from different bids be $\Omega_{i,t}$, and assume that firms form their bids to maximize their expectation of static profits conditional on this information set in each period. We let $\mathcal{E}[\cdot|\Omega_{i,t}]$ be the operator which provides the firm's perceptions of expected profits conditional on $\Omega_{i,t}$.

Assumption 1 (static maximizing behavior) *The firm choses its bid to*

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

This assumption is stronger than it may at first appear. While it seems completely reasonable to ask that agents choose their bids to maximize profits given their perceptions (i.e. that they "do their best"), we are restricting them to static profit maximization. This rules out collusion, for example. It also rules out bid experimentation, although by the final stage of the data the incentives to do this may be small.

Since total demand for frequency response is determined exogenously by regulatory requirements, it is independent of the bids, so that $\partial M_t/\partial b_{k,t} = 0$.[14] Thus M_t drops out of the first order condition in $b_{k,t}$ and we have the system of $|\mathcal{J}_i|$ first-order conditions for firm i

$$\mathcal{E}_{\theta,b_{-i,t},\xi_t}\left[s_k(b_t, x_t, \xi_t; \theta) + \sum_{j\in\mathcal{J}_i} (b_{j,t} - c_j)\, \frac{\partial s_j(b_t, x_t, \xi_t; \theta)}{\partial b_{k,t}} \middle| \Omega_{i,t} \right] = 0, \quad k \in \mathcal{J}_i, \tag{4}$$

and the solution is the course of action taken by firm i.

In order to back out costs, we next place some restrictions on agent perceptions during the final period. Our assumption on the relationship between the firm's expectations and the data generating process in the last period is only that agents expectations are on average correct. More formally if we let $t = 1, \dots, T^e$ index the equilibrium time periods than our assumption is that

Assumption 2 (behavior in the equilibrium period)

$$(T^e)^{-1} \sum_{t=1}^{T^e} \left[s_k(b_t, x_t, \xi_t; \theta) + \sum_{j\in\mathcal{J}_i} (b_{j,t} - c_j)\, \frac{\partial s_j(b_t, x_t, \xi_t; \theta)}{\partial b_{k,t}} \right] = 0.$$

Given the estimated demand system, Assumption 2 lets us back out c_j for each BM unit simply by solving for it (i.e. method-of-moments estimation). This is a weak assumption in the sense that it is consistent with many different models of play during the last period. This makes using the assumption attractive for inference. For example, suppose the firms knew each other's costs and characteristics, and had perfect foresight over the demand innovations $\{\nu_{j,t}\}$. Then if they were playing the standard Bertrand-Nash equilibrium of the resulting game of complete information, we would expect the first order conditions to hold exactly in every period: stronger than requiring that they hold on average, as in assumption 2.

With this in mind, as an alternative estimation approach we also solve the system of equations (4) separately by firm-month, replacing the firm's expectations with the realizations, to get a vector of implied costs $\hat{c}_{i,t} = \{c_{j,t}\}_{j\in\mathcal{J}_i}$. These cost estimates should be identical over time if the more restrictive assumptions above hold, but otherwise they will be a bit noisy.

Table 4 shows our cost estimate for the top 8 stations (ranked by revenue).[15] The mean and standard deviation of the time-varying estimates are shown, as are the estimates from

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

[15] 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

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

Station	# Units	# Obs	Fuel	Vintage	Mean	Std. Dev.	MM 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	1969	1.55	.19	1.55
Drax	6	173	Coal	1974	1.36	.4	1.4
Eggborough	4	108	Coal	1969	1.73	.21	1.74
Peterhead	1	29	CCGT	2000	1.72	.11	1.72
Ratcliffe	4	97	Coal	1968	1.54	.23	1.55
Seabank	2	58	CCGT	1998	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 method-of-moments cost estimation approach.

the method-of-moments approach based on assumption 2. The mean of the time-varying estimates are approximately equal to the method-of-moments estimates.

They are quite 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.[16] Since one would have expected some markup to be built into the regulated prices, our costs are in the right ballpark. Note that this also suggests that, at least when we average over time, market outcomes are well approximated by a static Nash equilibrium during the equilibrium period of the data.

Table 5 shows the results from projecting our estimated costs $c_{j,t}$ onto unit characteristics. As one might have expected 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

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.

[16] 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 $\times$ 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 (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

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 (4). As noted above without introducing further heterogeneity across firms, we would struggle to explain the divergent strategies during the early period of the data. This leaves us with comparing the predictive abilities of different learning models during the middle period, where firms appear to start to compete more aggressively on price.

5.1 Learning models

Baseline: equilibrium play. It is useful to set up a baseline to compare the learning models to. The obvious baseline is the static Nash equilibrium, since typically when an IO economist wants to predict outcomes, they assume some form of equilibrium play in the counterfactual. In each period, we find a static Nash by iterating the best response function (with some damping) until convergence. In contrast to the simpler belief formation in the strategic learning models below, equilibrium play demands that agents mutually believe that their rivals will play the equilibrium actions, and best respond given those beliefs.

Strategic Learning. 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 $m \notin \mathcal{J}_i$). 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 special case ($\delta = 0$). Another interesting special case is the model with $\delta = 1$, so that all past observations are equally weighted.

We are just starting our analysis of learning models and we begin by considering these two special cases as well as the case $\delta = 1/2$. While the above notion of fictitious play is conceptually well-defined for 2-player or population 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. This is another assumption we will eventually relax.

Statistical Learning. The fictitious play models used in the literature assume that firms know the demand system parameters θ. However 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 now detail how we operationalize the idea of demand uncertainty. Recall that we estimated the demand system using all of our data. When the firm's formulated their bids they did not have all of this information available. The predictable part of demand is a function of bids, main market position, lagged unobservables and unit and month fixed effects. 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).[17] 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. However 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 that uses the data available prior to the bid and the regression equation

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

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 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.

[17] 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.

Predicting bids: Each learning model M implies a particular form for the operator $\mathcal{E}_{\theta,b_{-i,t},\xi_t}$. This pins down a bidding function $\phi^M_{i,t}(c_{i,t}|\Omega_{i,t})$ via the first order conditions (4). 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:

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

Best response with known demand parameters is equivalent to putting point mass at the estimated demand parameters and at last period's bids, implying that the only remaining source of uncertainty is the realization of the demand unobservables ξ. As noted earlier, we assume that firms know the stochastic process for ξ_t, and optimize given their knowledge of ξ_{t-1} and the distribution of the innovations ν_t. For our calculations, we sample the innovations $\{\nu_{j,t}\}$ from their estimated empirical distribution independently across units. Each draw gives us a draw of ξ_t according to $\xi_t = (\xi_{1,t-1} + \nu_{1,t} \ldots \xi_{J,t-1} + \nu_{J,t})$, and we solve the system of first order conditions so that they are all zero when averaged across draws. 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).

5.2 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 conditional on the information available to the player at the time of its decision.

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 $s_{j,t}$). For the reasons noted earlier we do this only for the middle period of the data.

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.

Many of the units in our data do not change their bids from period-to-period (see figure 6), even though the environment is continually changing (e.g. rival bids have changed). None of our models will do well in matching this "bid stickiness". So as a last measure, we calculate an R^2 that assigns zero weight to units that don't change their bid from the last period, and gives share-weight to the remaining units (i.e. RSS is $\sum_{t=17}^{42}(b_{j,t} - \hat{b}_{j,t}^M)^2 s_{j,t} 1(b_{j,t} \neq b_{j,t-1})$).[18]

5.3 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. Notice that even though the fit isn't great, the best response and fictitious play models perform better than the baseline of equilibrium play.

The graphical analysis is backed up by the other measures of fit shown in Table 6. All the models have *negative* R^2. This occurs because their prediction of the mean bid is too low, whereas the total sum of squares is at least computed around the correct mean. There is a sense in which 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 respectable for the first two models, and even better once they are restricted to active units.

The statistical learning models perform better still (see figure 14). In particular, the model in which firms learn both α and β does quite a good job of reproducing the observed price

[18] We thank Pat Bayer for suggesting that we restrict attention to active players in measuring fit.

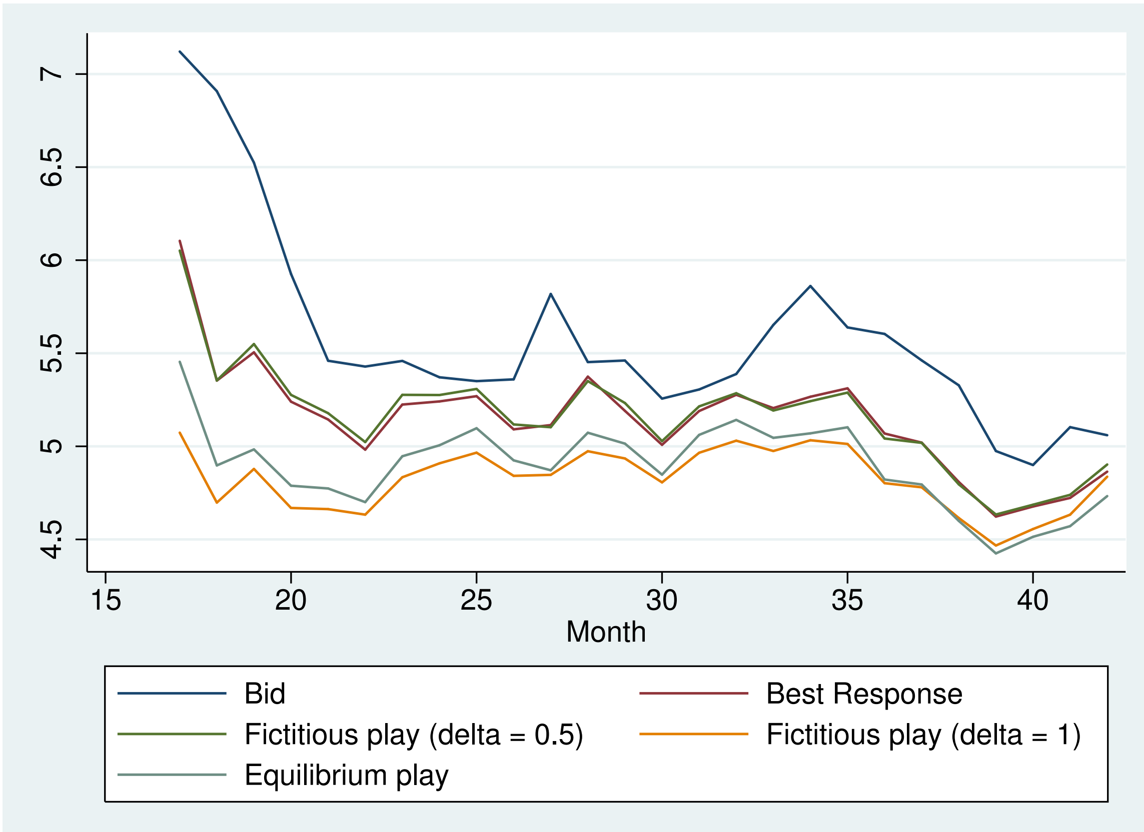

Figure 13: **Predicted bids: different strategic models.** Shown are the actual bids, and predicted bids from three different strategic models: fictitious play with $\delta = 1$ (i.e. equal weighting of past bids); fictitious play with $\delta = 0.5$; and best response (i.e. $\delta = 0$). The bids are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

path, and has a share-weighted R^2 of 0.37, and an activity-and-share-weighted R^2 of 0.57 — respectable for a model where no parameters were directly chosen to maximize fit. 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.

What have we learned about learning in this market from our research to date? The data favor models in which firms pay more attention to recent behavior by their competitors. 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 be that while the true demand system is stationary, our best approximation to it is not. In particular, our sequential demand estimates may be better

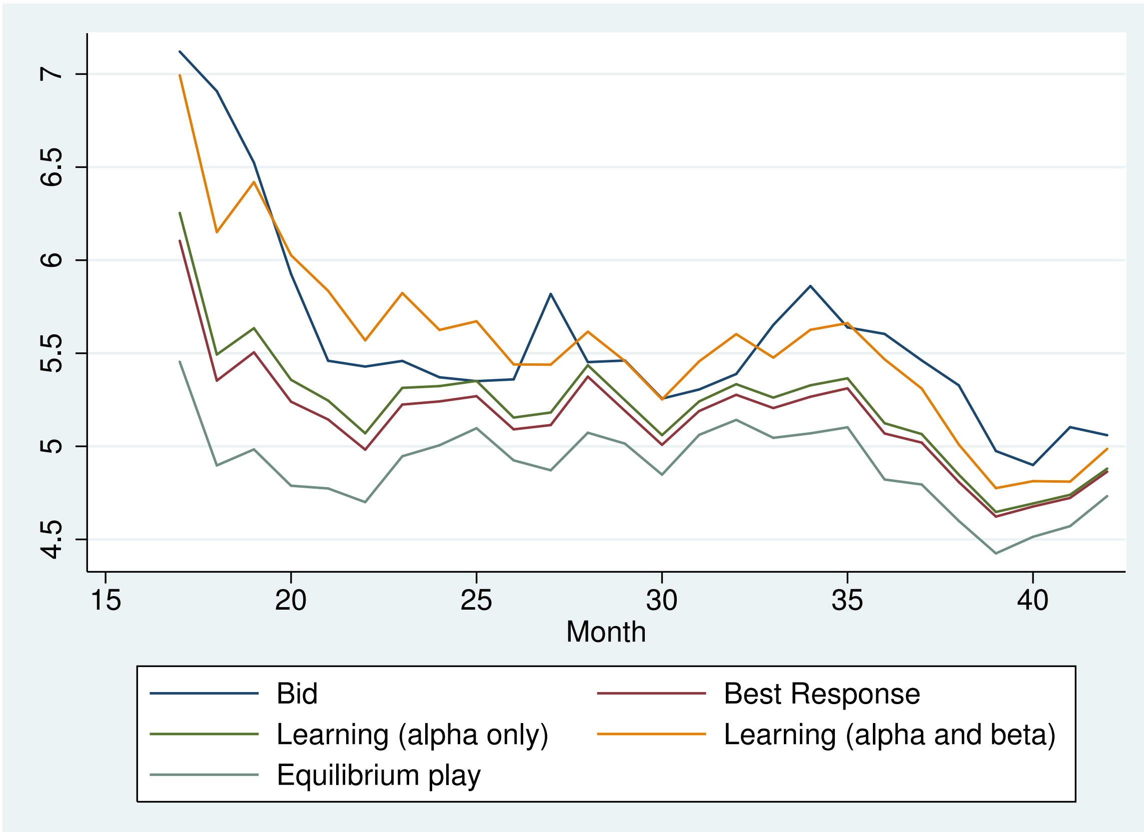

Figure 14: **Predicted bids: different models of demand uncertainty.** Shown are the actual bids, and predicted bids from three different models of demand uncertainty: best response with known demand parameters, best response with statistical learning of α and best response with statistical learning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

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.

Table 6: Fit of the Learning Models

Strategic Model	Statistical Learning	R^2 (unweighted)	R^2 (share-weighted)	R^2 (active only, share-weighted)
Best Response	α and β	0.08	0.37	0.57
Best Response	α	-0.01	0.28	0.40
Fictitious play ($\delta = 0.5$)	none	-0.04	0.25	0.34
Best Response	none	-0.04	0.24	0.34
Fictitious play ($\delta = 1$)	none	-0.18	0.06	0.04
Equilibrium play	none	-0.26	-0.06	-0.12

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 share-weighted measure weights each observation by its share $s_{j,t}$ in computing the RSS and TSS. In the final column we share-weight, but in addition set the weight of units whose bid is unchanged (i.e. they are "inactive' this period) to zero.

6 Conclusion

We have documented what happened to the British frequency response market following its deregulation. All our analysis suggests that the market converged to a rest point that appears to be a static Nash equilibrium, although this process took over 3 years. The opening period of the market was characterized by substantial uncertainty, and firms took different approaches, with some exploring different bids, while others did little at all. This period seems hard to formally model.

By contrast, the middle period is more amenable to analysis. We find that simple fictitious play models fit the data better than assuming firms play according to the static Nash equilibrium. Assuming that firms are uncertain about the price elasticity and update over time as data comes improves fit still further.

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.

Evans, G. W., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press.

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.

Lee, R. S., and A. Pakes (2009): "Multiple Equilibria and Selection by Learning in an Applied Setting," *Economics Letters*, 104(1), 13–16.

Marcet, A., and T. J. Sargent (1989): "Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information," *Journal of Political Economy*, 97(6), pp. 1306–1322.

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.

Salmon, T. C. (2001): "An Evaluation of Econometric Models of Adaptive Learning," *Econometrica*, 69(6), 1597–1628.

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

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-8D` `47211/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-8D` `47211/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 $\geq$ 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 (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).