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

October 22, 2013
(Preliminary and incomplete, please do not cite)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 Nash equilibrium in prices. We examine how well theoretical models of strategic learning do in explaining 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: firms use the public data on demand for their and their rivals products to update their beliefs about overall demand. A model that combines simple best response dynamics with statistical learning fits the data surprisingly well.

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 demand 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 learning and competitive dynamics would be helpful in providing more nuanced conclusions.

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 repeated static games i.e. on stable environments in which players are trying to maximize static payoffs. An organizing question in that literature is which learning models are guaranteed to converge to equilibrium, and under what conditions this occurs. But there has been little empirical work on strategic learning, and consequently little evidence to discriminate among the many models that have been offered by the theory community. 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 adheres relatively closely to the modeling assumptions of the theory literature. 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 implicit collusion is difficult and repeated game considerations may be less relevant.

Our paper attacks two main questions: 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 find convincing evidence that the market reaches equilibrium after three and a half years. Generators set wildly different prices in the early stages of the market, and updated 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 speed of convergence seems moderate: 42 periods is a reasonable amount of time to reach equilibrium.

The rest of the paper brings structural econometrics to bear in trying to assess 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, providing further evidence that the later period reflects equilibrium play.

Next, we consider a series of learning models. In some models firms have the correct model of overall demand throughout the sample period (i.e. their estimate coincides with ours), but still face strategic uncertainty about how their competitors will bid. We allow firms to place differential weight on the past play of their competitors (in the theory literature, this is "fictitious play"). A second set of models allows 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 (more limited) set of data available to them at each point in time. We assess each model by comparing the simulated bids to the observed bids in the period prior to equilibrium.

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). But 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, 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 just under 0.4, 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.

Another implication is that the information structure plays an important role in how players update their strategies. Here, the system operator is required by regulation to make all prices and volume data available as soon as practically possible, allowing firms to make inferences from rivals' outcomes. This may have increased the speed of convergence.

Literature Review. [to be written]

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.

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 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. 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.

A BM 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. 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.[1] The *quantity* of FR capacity offered by a unit depends on the type of BM unit (fuel type, vintage, etc.) and its current operating position.[2]

CAP047. 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

[1] 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.

[2] For example, the FR capacity of a BM unit operating near the top of its range is is limited by its ability to temporarily increase energy production.

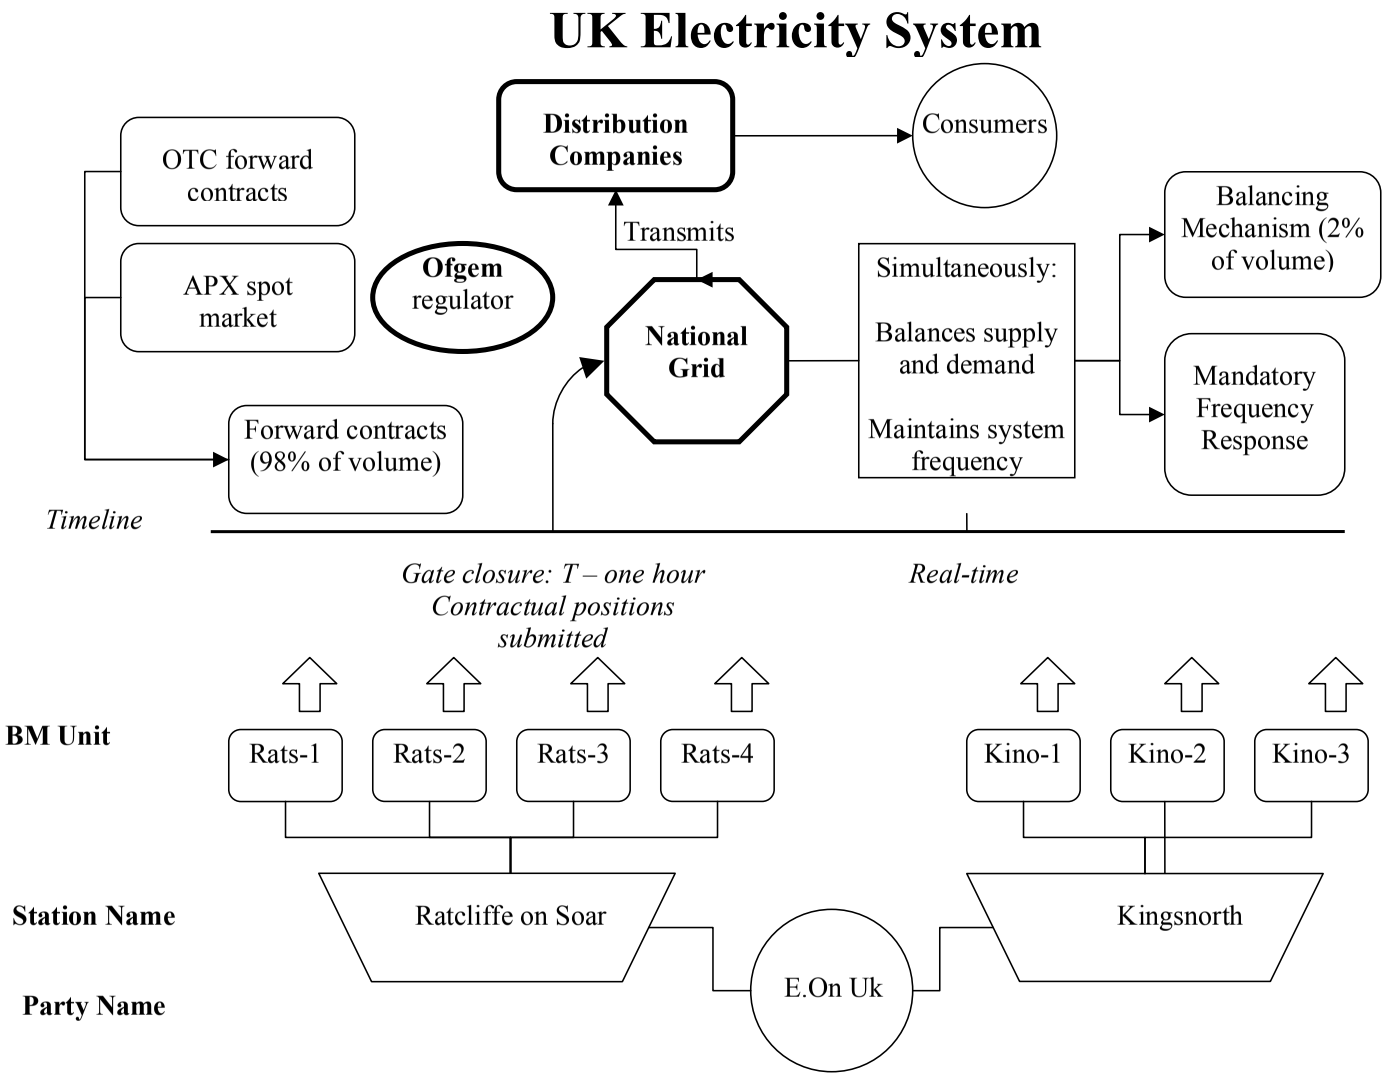


Figure 1: Overview of the British electricity market.

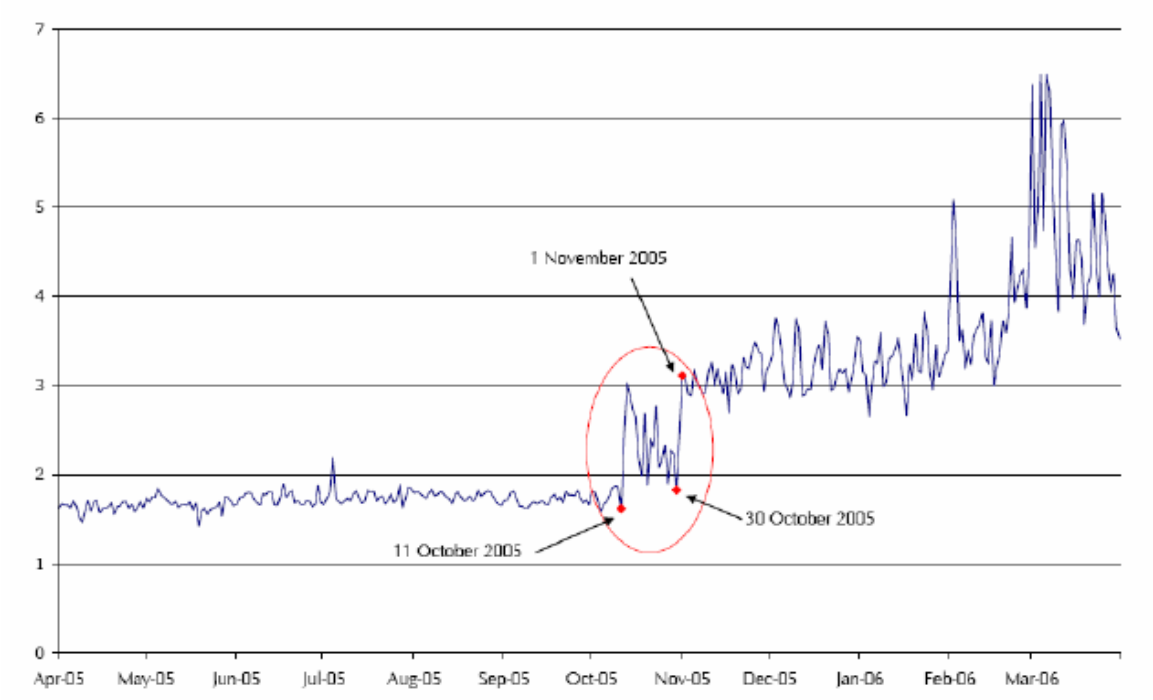


Figure 2: **Holding payments for high frequency response pre and post CAP047.** Source: National Grid.

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 the "prices" 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.[3] 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.[4] A bid commits the BM unit to offer FR at a fixed price over the next month. The quantity that it delivers if instructed into FR mode varies with its technological capabilities (fixed over the

[3]BM units are also allowed to specify maximum capacities for each type of FR for deviations of .2, .5, and .8 from 50Hz. These maximum FR capacities were initially negotiated by NG and the firms owning the BM units, and never change in our data.

[4]CAP047 left the energy response payment unchanged. Since this is a relatively small source of FR profits, we focus on the holding payments in the paper.

sample period) and operating position (time-varying).

National Grid 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 high FR almost immediately doubled with 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 data from two main 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.[5]

The second data source is Elexon Ltd. They are contracted by the government regulator to

[5]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.

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.

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.[6]

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, 2012).[7] 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-plant firm, while the next two largest participants, E.ON and RWE, are multi-plant firms. Anecdotally, Drax's disproportionate revenue share is attributable to having a relatively new plant, with accurate governor controls, making them attractive for providing

[6]The BM data comes from bmreports.com and the ownership data is under static market data on the website of Elexon Ltd.

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

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.[8] Within a year, this increased Drax's annual revenue from the FR market more than threefold.

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.

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 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. This is consistent with the fact that NG's response requirement derives from government regulation and risk-response curves that did not change over the sample period. 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.

Turning from the demand for FR to the supply of FR, 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

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

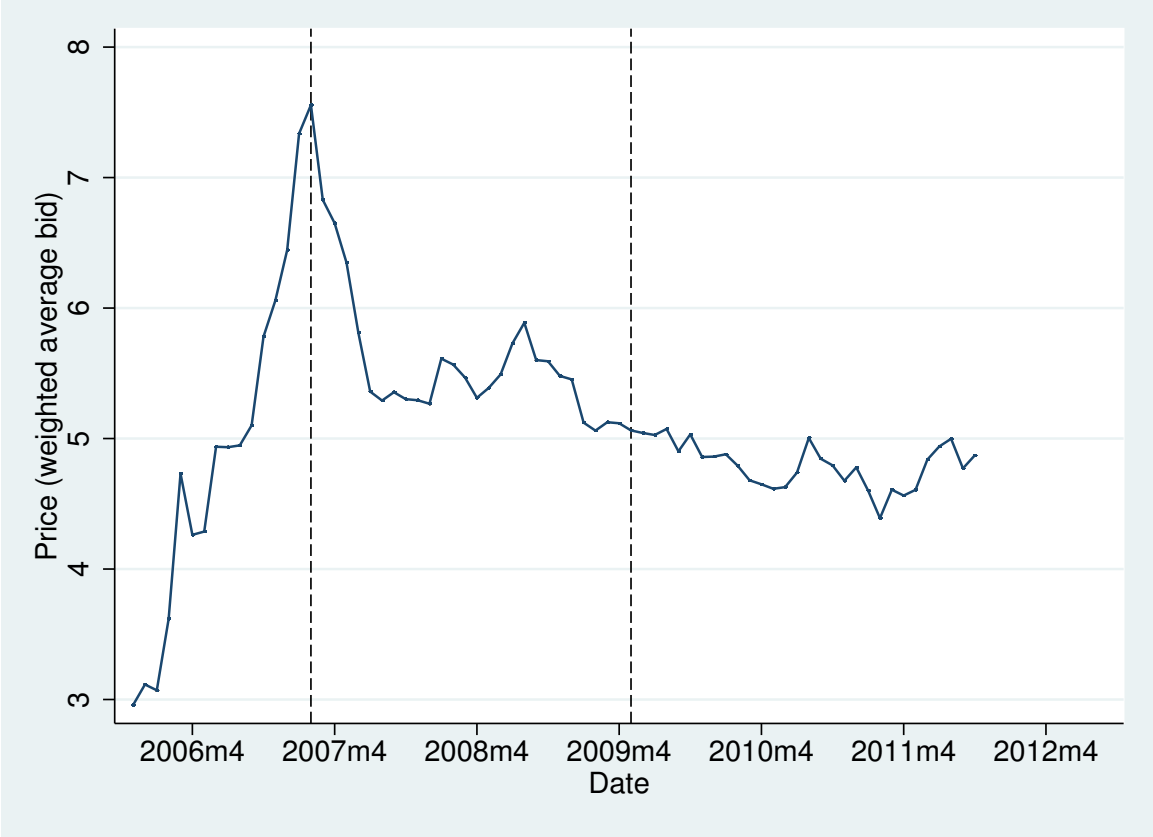


Figure 3: **Monthly prices over time.** Prices are weighted averages of bids across units, where the weights for each period are the volumes received by each unit.

the sample period, the movements are relatively small and none of the stations who become active or inactive are particularly large stations.

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

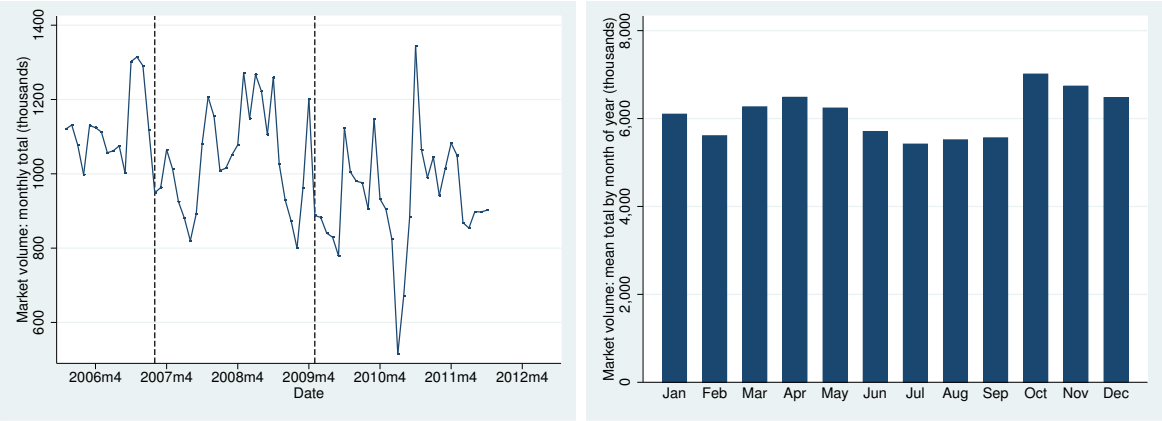


Figure 4: **Volumes over time and by season.** The left panel shows the total volume of frequency response for each month in the sample period, while the right panel shows the total volume by month-of-year (4 observations per month-of-year).

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.[9] The levels and 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 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

[9]We define "biggest" by revenue rank during the pre-equilibrium period, as we want to include all the big players during this period.

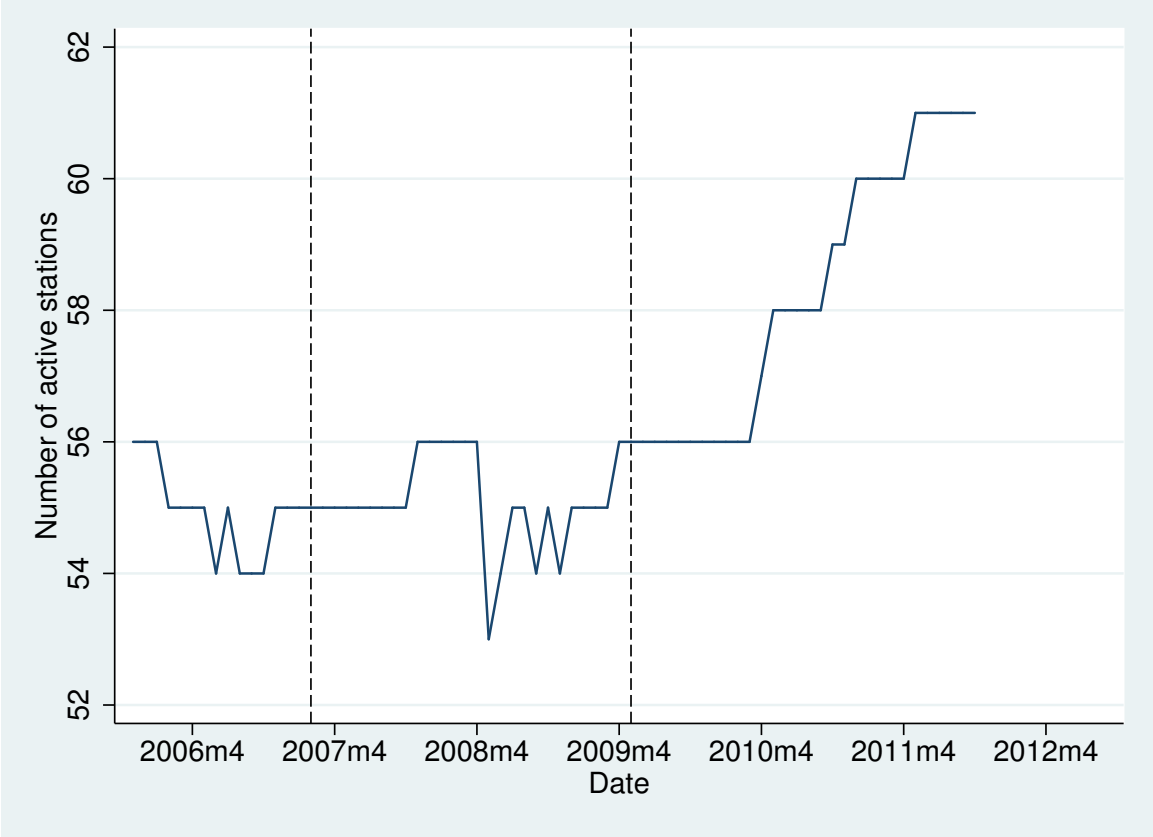


Figure 5: **Number of active stations over time.** This plots the number of active stations over time, where a station is "active" if one of its units submits a competitive bid (bids ≤ 23).

how others would bid and therefore the volume of FR they would attract. In each 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

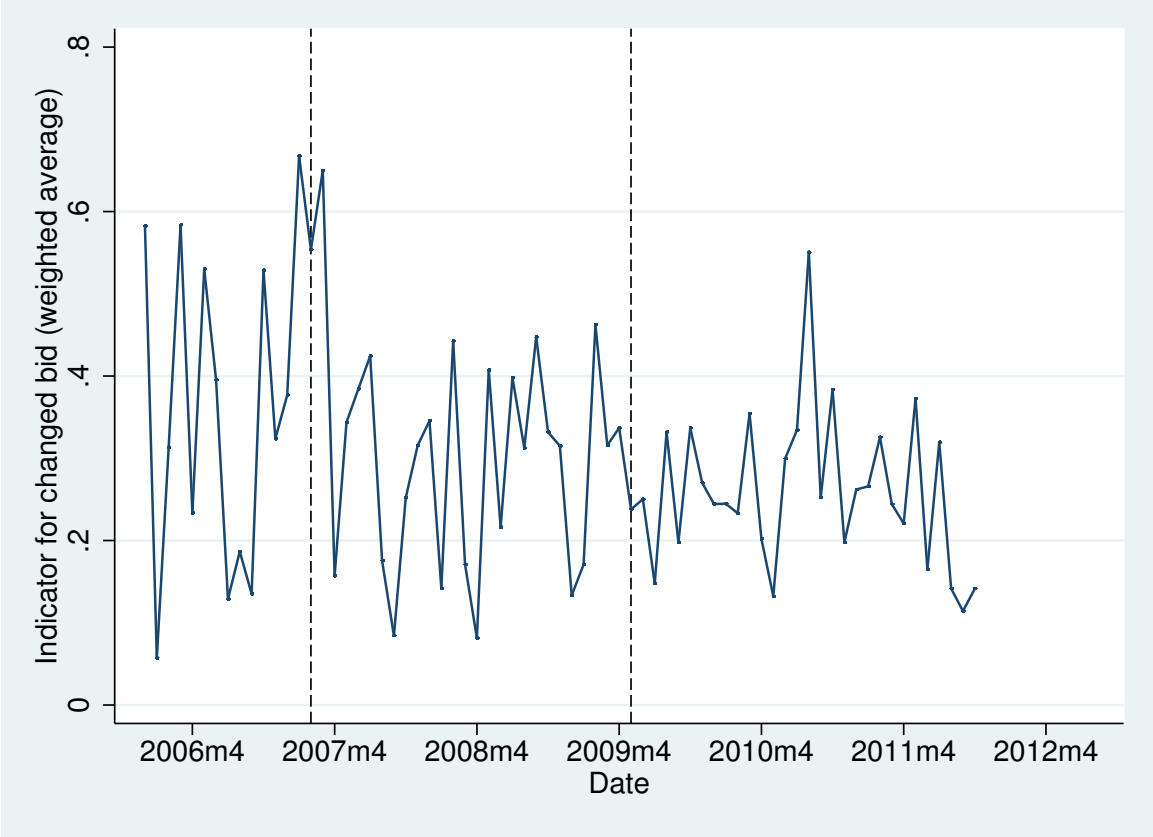


Figure 6: **Probability of a bid change over time.** This is a weighted average across units of an indicator for a change in their bid in period t, using quantity-weights from period $t-1$.

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.[10]

[10]Drax increases its bid at the end of 2007 for *exactly* two periods, giving everyone an opportunity to see its bid increase and respond, presumably with the aim of some sort of tacit collusion. When no one follows suit, they drop it back down.

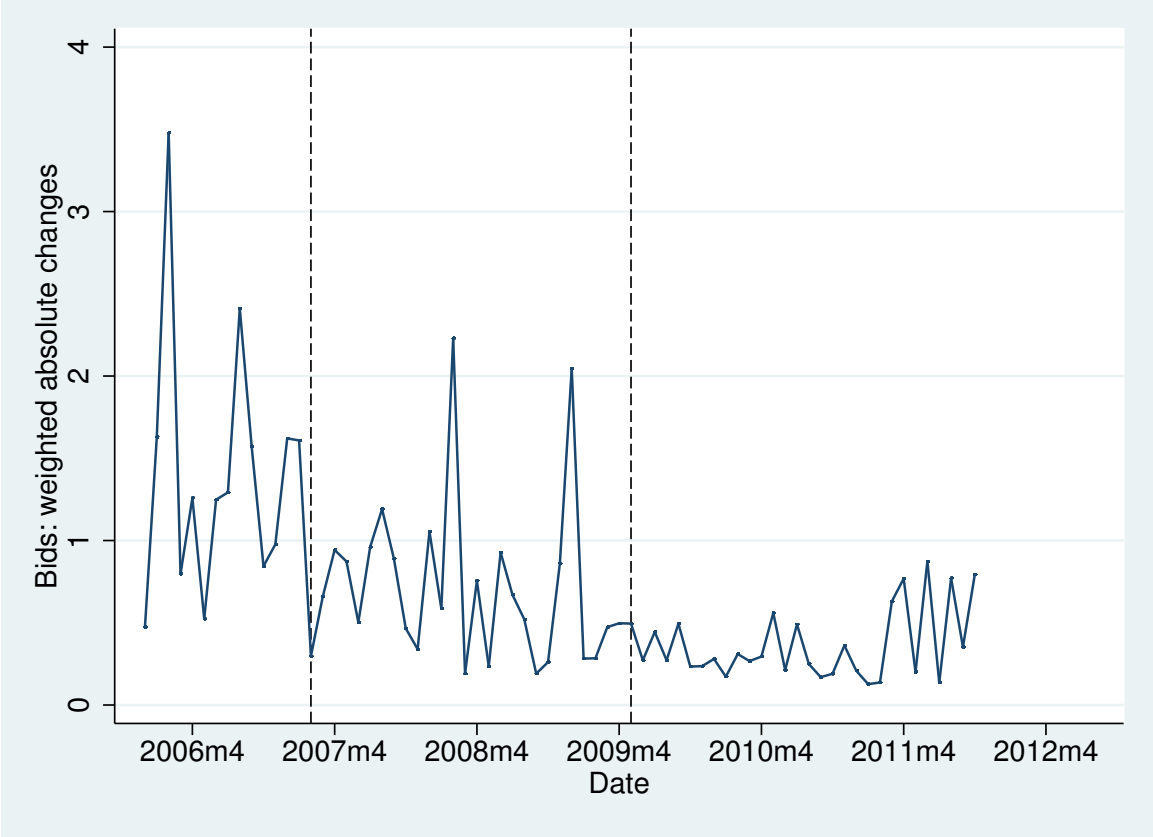


Figure 7: **Absolute bid changes over time conditional on changing.** This is a weighted average across units of the absolute change in their bid in period t, using quantity-weights from period $t-1$ and assigning a weight of zero if their bid did not change.

Final period (June 2009 – October 2012). 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, for 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. Prices are periodically "shocked", perhaps by unforeseen changes in market conditions, but the bid changes that occur thereafter mitigate the shocks and bring the market price back to a relatively stable £4.8/MWh.

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

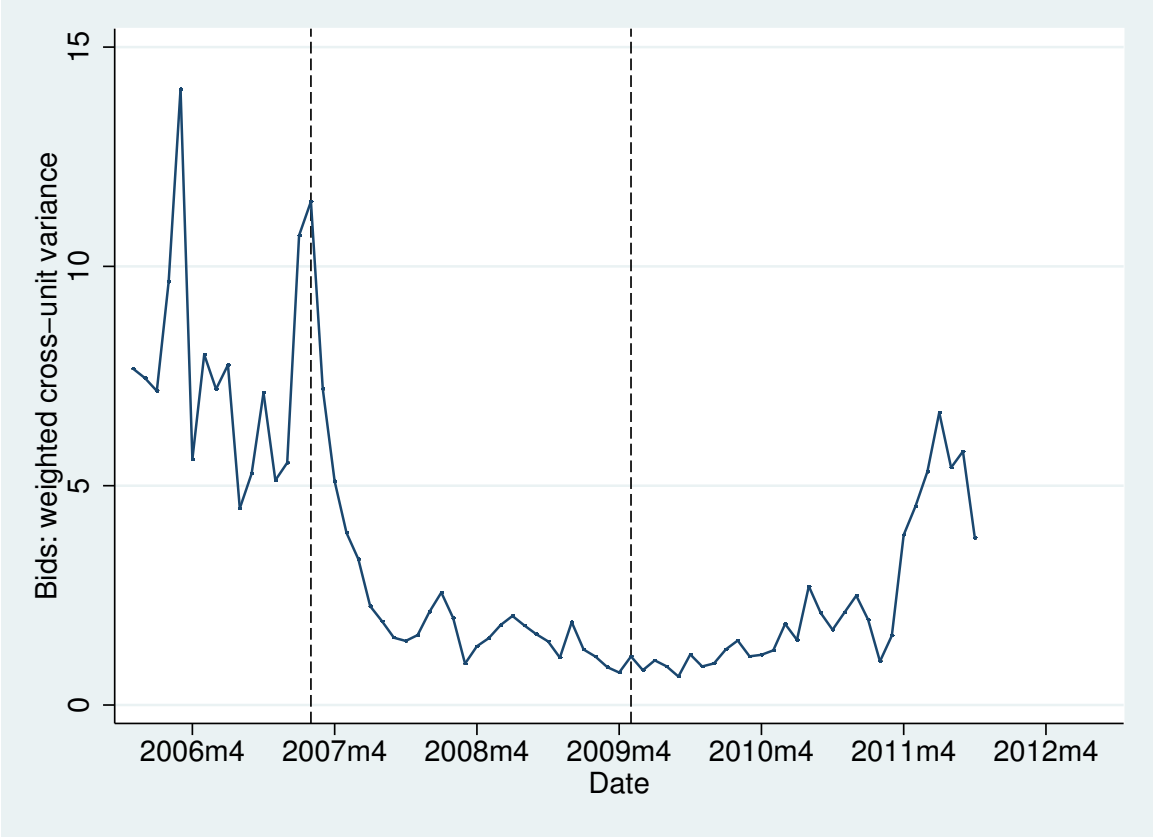


Figure 8: **Variance in bids across units.** Variance in the cross-section of bids at each point in time, where bids are quantity-weighted using contemporaneous quantities.

more erratic or rise (Drax, 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, 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 profit given their model of the world and their current information

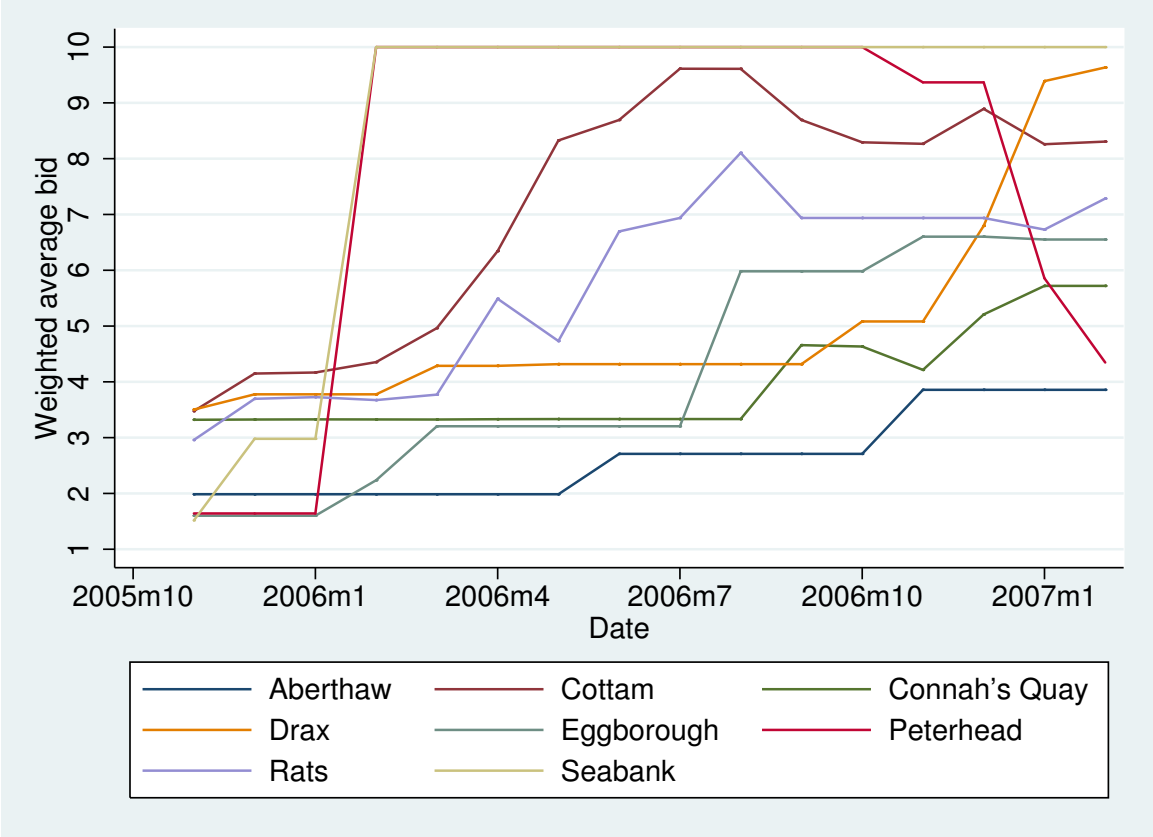


Figure 9: **Bids of the top 8 biggest stations (November 2005 – February 2007)** Stations are ranked by total revenues in the pre-equilibrium period. Bids are censored above at 10 to compress the axes and improve the visual presentation.

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.

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

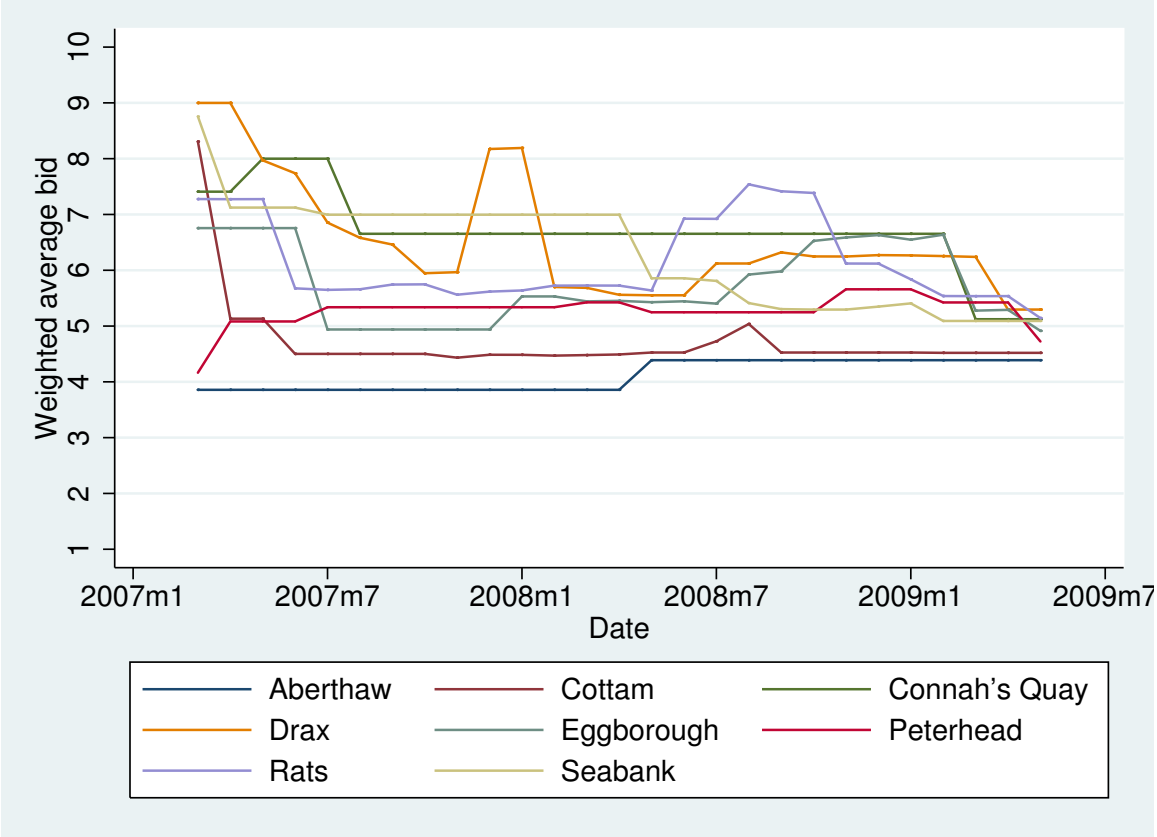


Figure 10: **Bids of the top 8 biggest stations (March 2007 – May 2009)** Stations are ranked by total revenues in the pre-equilibrium period

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

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

While firms know their bids and costs they are uncertain at the time of bidding exactly what the future demand for their services is. To capture this, we assume that National Grid allocates frequency response volumes to unit j according to a function $q_j \equiv q_j(b_t, x_t, \xi_t)$. This function may be unknown to the firms. The vector x_t stands in for a set of unit characteristics that are mutually known to all firms at the time of bidding (e.g. each other's costs), while ξ_t is a set of mutually unknown characteristics (e.g. the future positions of each unit during the month to come). Firms have expectations about demand summarized in an operator $\mathcal{E}_{i,t}$

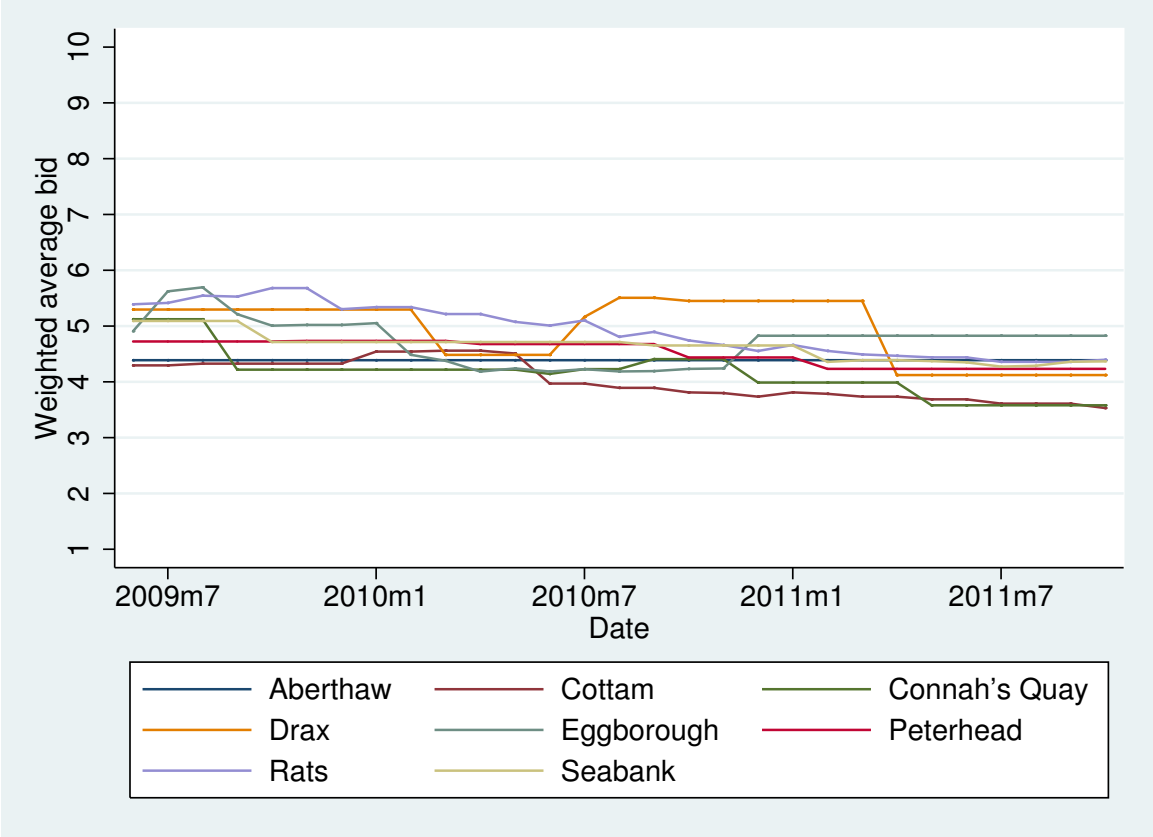


Figure 11: **Bids of the top 8 biggest stations (June 2009 – October 2012)** Stations are ranked by total revenues in the pre-equilibrium period

and maximize expected profits given their information set $\Omega_{i,t}$.[11] Thus the bids $b_{i,t}$ solve the maximization problem

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

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

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

[11]NG publishes bids and quantities as soon as they becomes available, so at the time they bid for the next month, firms know their rivals' bids in the current month, and quantities for the last month.Current quantity data is not yet available as firms bid halfway through the month. Firm expectations must be measurable with respect to their information set.

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).

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: we allow for the possibility that firms under or overestimate the level and slope of the demand curve in arbitrary ways. The central question then is what model firms use to guide their expectations. We will return to this question soon. But first, we need to estimate the true demand function $q_j(b_t, x_t, \xi_t)$.

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:

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

$\xi_{j,t}$ captures demand uncertainty that arises from, for example, transmission constraints in the main market. 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:

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

where L denotes the lag operator. Fixing ρ, we can estimate the remaining parameters by OLS. So we use a nested loop estimation procedure, estimating ρ in the outer loop via a grid search for the ρ that minimizes the sum of squared residuals.[12]

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 that get zero share are inactive because they bid too high, or because they are badly positioned in the main market. Dropping them may cause us to underestimate the relevant coefficients in the demand system.

We have examined this with a probit in which the dependent variable is an indicator for getting a positive share, and the controls are the same as in the logit specification. We find that extremely high bids ($b_{j,t} > 23$) are an almost perfect predictor of not getting called; but after throwing out these outliers, bids do not significantly predict that a unit gets a positive share, which is reassuring.[13]

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 as one would expect. In the QFD specification we find some evidence of persistence in the unobservables, as $\rho > 0$. However, since our estimate of α is not significantly different across the specifications, it appears that bid endogeneity — at least of the form permitted here — is not particularly important.

To assess the fit of the demand system, we simulate out the shares $\{s_{j,t}\}_{j=1}^{65}$ in each period from $t = 2 \ldots 72$. This requires integrating out the iid errors $\{\nu_{j,t}\}_{j=1}^{65}$, which we do by sampling uniformly from the empirical distribution of residuals $\hat{\nu}_{j,t}$ in (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.[14] Figure 12 shows that the fit is good even for the biggest stations, whose shares vary quite dramatically over time. This indicates that the

[12] 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_0/T. The median number of time periods in which we observe a unit is $T = 58$, so the bias is probably relatively small.

[13] Omitting the zero share observations also eliminates most of the outlier bids, but we choose to drop the few remaining cases with $b_{j,t} > 23$, as they may have undue influence on our estimates of the price coefficient (particularly since we would otherwise include them *only* when they get positive share).

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

Table 3: Demand System Estimates

	Log share ratio	
	OLS	QFD
Log bid	-1.690***	-1.684***
	(0.136)	(0.135)
Position / Capacity (average over month)	2.538***	2.416***
	(0.144)	(0.122)
Unit and Month FE	yes	yes
ρ	–	0.41
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 estimation is by quasi-first-differencing (an estimate of the autocorrelation ρ is given in this case). 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 (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

good fit is not solely a consequence of having unit fixed effects.

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. As we said earlier, we believe these should be reasonably stable. So we let each unit have a mean marginal cost of $\bar{c}_j$, and allow for time-varying iid errors $\varepsilon_{j,t}$:

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

These errors may arise as a consequence of variation in opportunity costs in the main market: when a firm bids low in the frequency response market, National Grid may decide to hold the unit's position below its maximum generating capacity in the balancing mechanism so that it is available for frequency response. This may hurt the firm's profit in the balancing mechanism, which is the "opportunity cost" of providing frequency response. But since the balancing mechanism is a relatively small part of the main market, with most electricity traded in advance, such opportunity costs may be small.

To recover the mean costs $\bar{c}_j$ we take an approach that is standard in empirical IO: we invert

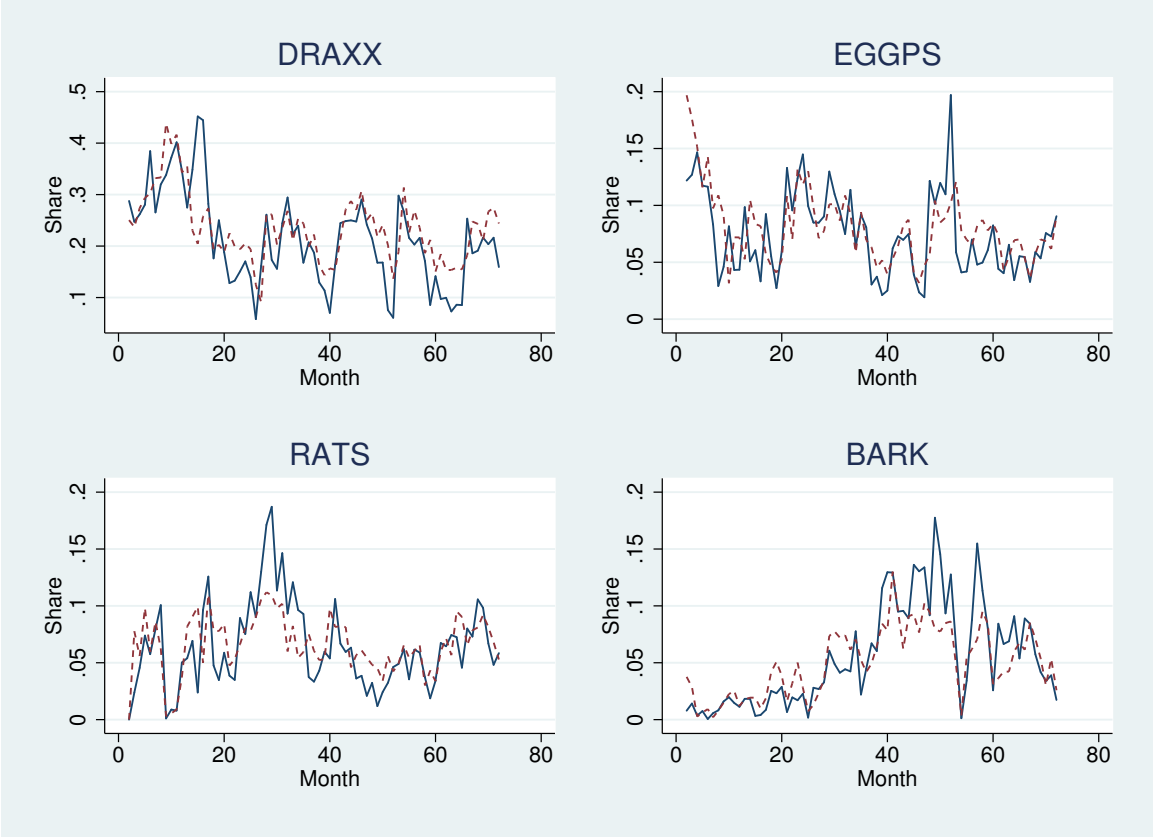


Figure 12: **Fit of the demand system.** Shown are the actual (blue, solid) and predicted (red, dashed) shares for the 4 biggest stations (Drax, Eggborough, Ratcliffe and Barking).

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: $q_j(b_t, x_t, \xi_t) = Q_t s_j(b_t, x_t, \xi_t)$. Because total demand for frequency response is a known function of main market generation, it is independent of the bids, so that $\partial Q_t / \partial b_{k,t} = 0$.[15] This means Q_t

[15]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.

drops out of the first order condition in $b_{k,t}$ and we have:

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

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

Assumption 1 (Information Set) $\Omega_{i,t} = \{D_t, x_t\}$ *where* D_t *is the publicly available data:*

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

This assumption does a decent job of capturing the firms' information sets. 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:*

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

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:

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

Given our demand estimates, we 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 $\varepsilon_{j,t}$ and differences between expected demand and its realization. Thus we average those costs over time for each unit to get an estimate of the

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

Station	# Units	# Observations	Mean	Std. Dev.	Std. Dev. (within time period)
Barking	2	59	1.39	.17	.05
Connah's Quay	4	109	1.17	.18	0
Cottam	4	87	1.55	.19	0
Drax	6	173	1.36	.4	0
Eggborough	4	108	1.73	.21	.08
Peterhead	1	29	1.72	.11	0
Ratcliffe	4	97	1.54	.23	.03
Seabank	2	58	1.77	.12	.01

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 unit-month. Means and standard deviations are reported separately by station. The within time period standard deviation is the average standard deviation in costs within a time period for each station.

mean cost $\bar{c}_j$ for each unit.

Table 4 shows our cost estimate for the top 8 stations (ranked by revenue).[16] The mean costs we back out are actually pretty reasonable: they vary between 1.17 and 1.77. By comparison the "cost reflective" regulated prices before CAP047 look to be around 1.6-1.8 (see figure 2). Since one would have expected a reasonable markup to be built into the regulated prices (perhaps 20% as in the energy response payments), our costs are in the right ballpark. This suggests that the market may indeed be in a non-cooperative Nash equilibrium during the final phase of the data, as we have assumed.

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).

[16]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.

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 national grid 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)). 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 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 firms should not expect rival actions to be coordinated.

All of these fictitious play models assume that firms know the demand system parameters θ. An alternative hypothesis would be one of *statistical learning*, in which firms estimate $\hat{\theta}_t$ 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 $D_t = \{s_\tau, b_\tau, x_\tau\}_{\tau<t}$. 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). Thus in the latter case 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 $y_{j,t} = \delta_{j,t} - \rho L.\delta_{j,t} - (1-\rho)\gamma_j - \tilde{\mu}_t$, $\tilde{b}_{j,t} = b_{j,t} - \rho L.b_{j,t-1}$ and $\tilde{m}_{j,t} = m_{j,t} - \rho L.m_{j,t-1}$. The estimated $\hat{\alpha}_t$ in the former case is derived analogously, except that β is known rather than estimated. We will refer to these demand estimates as the "sequential" demand estimates in what follows.

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 uniquely pins down a bidding function $\beta^M_{i,t}(c_{i,t}|\Omega_{i,t})$ via the first order conditions (3). For example, in the case of best response we have $\beta^{BR}_{i,t}(c_{i,t}|\Omega_{i,t})$ as the solution $b_{i,t}$ of:

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

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, and our focus will mainly be on their relative fit.

5.3 Evaluating fit

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 $s_{j,t}$).[17] As mentioned earlier, we do this only for the middle period of the data.

A second measure of fit is an R-squared, 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-squared 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. For this measure we need to 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 estimated from the 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.

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

[17] An alternative would be to directly predict the price path under each model — i.e. using weights $\hat{s}_{j,t}^M$ predicted by each model. We plan to implement this in future iterations of the paper.

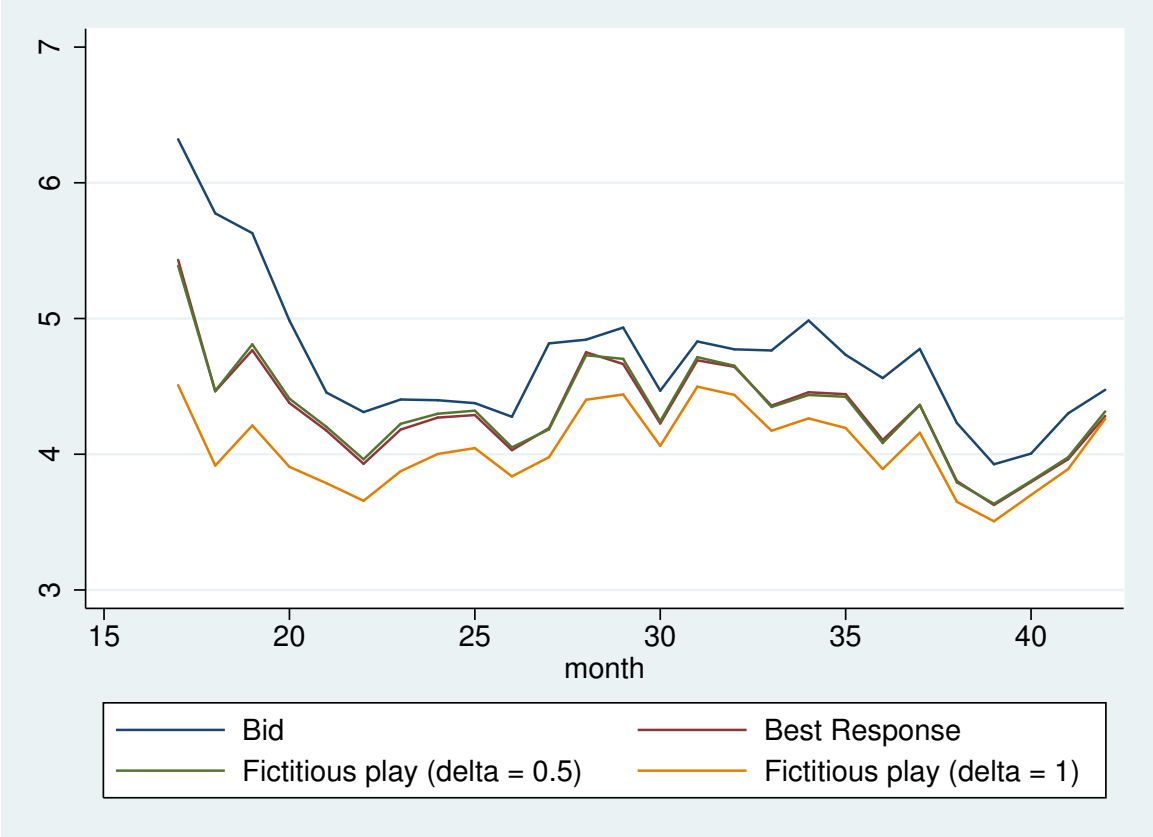


Figure 13: **Predicted bids: different strategic models.** Shown are the actual bids, and predicted bids from three different strategic models: fictitious play with $\delta = 1$ (i.e. equal weighting of past bids); fictitious play with $\delta = 0.5$; and best response (i.e. $\delta = 0$). The bids are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

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-likelihoods 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 of the individual log likelihoods (much like a paired 2-sample t-test).

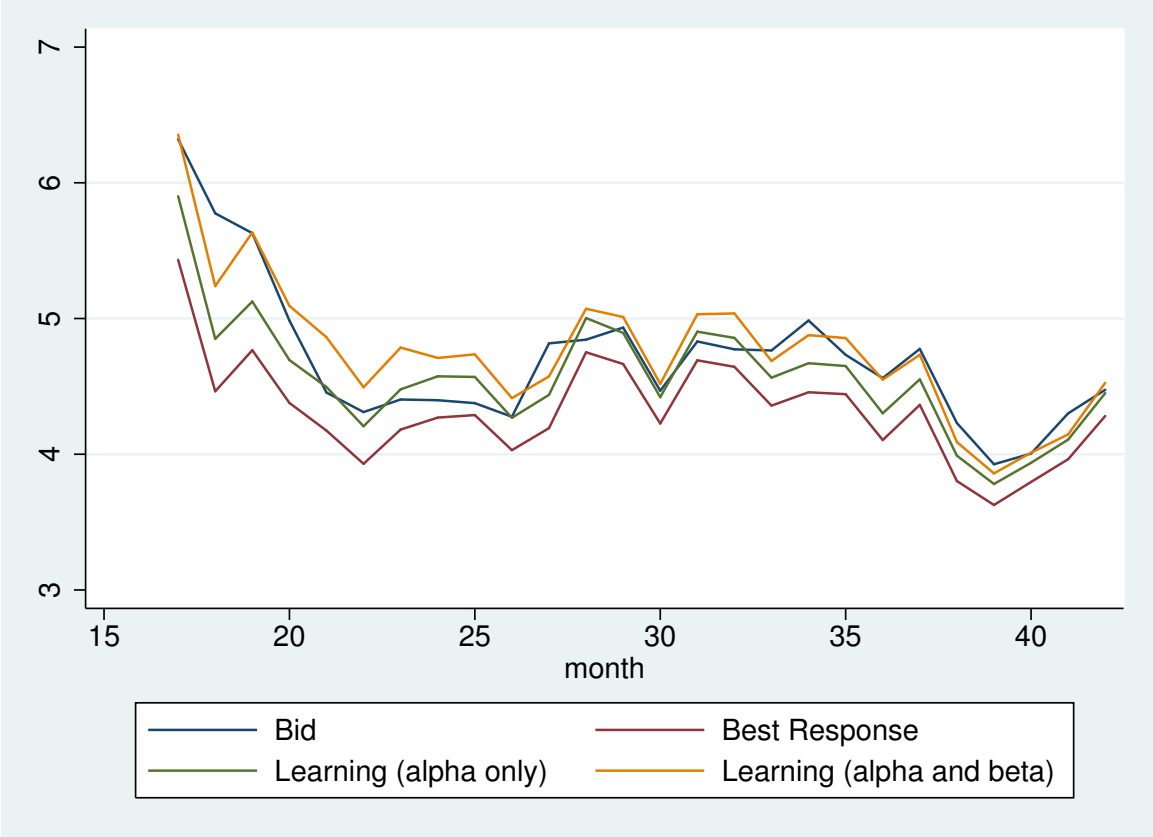


Figure 14: **Predicted bids: different models of demand uncertainty.** Shown are the actual bids, and predicted bids from three different models of demand uncertainty: best response with known demand parameters, best response with statistical learning of α and best response with statistical learning 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 models which only admit strategic uncertainty (i.e. 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 5. 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 is possible

Table 5: 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 (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

because they might get the mean bid completely wrong (and indeed we can see that they do), 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 0.4 — 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 that our final set of estimates, and so it is possible that the agents aren't learning at all and we fit better with the statistical learning models because we get better approximations to the demand system. Unraveling these different stories is something we hope to do in future versions of this paper.

6 Conclusion

[to be written]

References

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

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

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

Appendix

A. Data Appendix

1. Explain construction of the following variables:
   (a) bid, measured in £/MWh, observed monthly
   (b) quantity, measured in MWh, observed daily
   (c) operating position, measured in MW, observed every 5 minutes?
   (d) capacity, measured in MW, constructed how?
2. Explain how we aggregate over P, S, and H.
3. Explain how we aggregate BM units into power stations and into firms.

B. More efficient demand estimation

Estimators that only use observations that are present in consecutive years are consistent but inefficient. We would like to be able to use observations at lags greater than one . But this can be be problematic since $\nu_{j,t}$ can be correlated with $b_{j,t+\tau}$ for $\tau \geq 1$. Let $t^*(t,j)$ equal the last period prior to t in which j was active. We have

$$\delta_{j,t} - \rho^{t-t^*(t,j)}\delta_{j,t-t^*(t,j)} = \alpha\Big(\log b_{j,t} - \rho^{t-t^*(t,j)} \log b_{j,t-t^*(t,j)}\Big) + \beta\Big(m_{j,t} - \rho^{t-t^*(t,j)} A_{j,t-t^*(t,j)}\Big) + \gamma_j\Big(1 - \rho^{t-t^*(t,j)}\Big) + \mu_t - \rho^{t-t^*(t,j)}\mu_{t-t^*(t,j)} + \sum_{\tau=0}^{t-t^*(t,j)+1} \rho^\tau \nu_{j,t-\tau}.$$

If $t^*(t,j) = t-1$ the disturbance is just $\nu_{j,t}$. Since $\nu_{j,t}$ is not realized until after the bid is in, it is uncorrelated with the RHS variables in the above equation. However if $t^*(t,j) = t-2$ then $\nu_{j,t-1}$ is in the disturbance and it can be correlated with the bid. So to use observations with $t^*(t,j) < t-1$ we need instruments. Lagged bids would suffice.

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 $E[u_{k,t}] = 0 \Rightarrow E[\tilde{G}_{k,t}(b_t, c_{i,t}, x_t, \xi_t; \theta)] = 0$ for each k, which we can use to form moment conditions.