

Just starting out: Learning and equilibrium in a new market*

Ulrich Doraszelski

University of Pennsylvania

Gregory Lewis

Microsoft Research and NBER

Ariel Pakes

Harvard University and NBER

December 18, 2015

Abstract

We document the evolution of the newly created market for frequency response within the UK electricity system over a six-year period from 2005 to 2011. Firms competed in price while facing considerable initial uncertainty about market demand and rival behavior. We show that over time prices stabilized, converging to a rest point that is consistent with equilibrium play. We draw on models of fictitious play and adaptive learning to further analyze how this convergence occurs, showing that these models fit the data well and indeed outperform equilibrium analysis during the middle two years of our data.

*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 benefited from discussions with Joseph Cullen, Drew Fudenberg, Mar Reguant and Frank Wolak, and from the comments of seminar participants at Boston College, Cornell, Duke, Kellogg, and the NBER Productivity Lunch and IO Program Meeting. Rebecca Diamond, Duncan Gilchrist, Matt Hlavacek, Daniel Pollmann, Sean Smith, Amanda Starc, and Wei Sun have all provided excellent research assistance.

1 Introduction

What do competing agents or firms do when their environment changes? Answering this question is necessary for predicting what will happen following adjustments to market institutions, such as government regulation or market structure. Existing approaches are based on the computation of counterfactual equilibria, but convergence to equilibrium after a perturbation may not be swift or indeed certain. Understanding how firms learn and when and how convergence occurs is therefore important for such predictions.

This paper offers a case study of a newly deregulated market, the frequency response (FR) market in the UK. Initially, firms faced tremendous uncertainty both about demand for their services and about what their rivals would do. We explore how this demand and strategic uncertainty manifest themselves in the behavior of firms from “day one”, tracing their behavior over the next six years.

Broadly speaking, FR is a product required by the system operator to keep the electricity system running smoothly. Historically, electricity generating firms had been obligated to provide FR to the system operator at a fixed price. Deregulation created a market in which firms are allowed to bid for providing FR, thus setting the stage for price competition. An attractive feature of this market is that the demand for FR and the set of market participants are relatively stable, so that bid changes can be plausibly attributed to learning rather than changes in the environment.

The first part of the paper carefully documents bidding behavior over time. We distinguish three phases in the evolution of the FR market. The early phase of the FR market is characterized by heterogeneous bidding behavior and frequent and sizeable adjustments of bids. Some firms appear to experiment with their bids, and other firms appear to follow suit. Yet other firms do not change their bids at all for many months. The price of FR exhibits a noticeable upward trend during the early phase that culminates in a “price bubble.” During the middle phase of the FR market, this trend reverses itself. Competition between firms drives the highest bids down, leading to a dramatic reduction in the range of bids. Adjustments of bids are less frequent and smaller than in the early phase. By the time the FR market enters its late phase, it appears to have reached a “rest point.” This rest point is consistent with a complete information Nash equilibrium, and we show that firms adjust quickly to periodically occurring small changes in the market environment. Since the industrial organization literature routinely assumes that equilibrium reasserts itself after

changes in market institutions, it seems useful to be able to point to a concrete example even if the FR market converged to a rest point only after three and a half years of monthly strategic interaction. To the best of our knowledge, ours is among the first papers to demonstrate convergence.

The second part of the paper analyzes in more detail how this convergence occurs and provides evidence on how agents learn. To do so, we first estimate the demand and cost primitives under a relatively weak rationality assumption that we view as appropriate for the late phase of the FR market. Our assumption encompasses complete information Nash equilibrium along with much else.

We then consider a series of learning models. To capture strategic uncertainty about rival bids, we consider fictitious play models in which firms form their beliefs based on past observed rival behavior (Brown 1951). To model the firms' perceptions about demand, we add adaptive learning in which these perceptions are grounded in statistical analysis of the available data (Sargent 1993, Evans and Honkapohja 2001, Evans and Honkapohja 2013). We assess the ability of our learning models to make “one-step-ahead” predictions by comparing actual bids with those predicted by each model.

The experimentation by some firms in the early phase of our data is hard to rationalize with these models and so we do not analyze their fit. During the middle phase, the best-fitting models are those in which firms more heavily weight recent rival behavior in forming beliefs; and adaptively learn about the price elasticity. These predictions are substantially better than those from a complete information Nash equilibrium where all agents know the demand parameters: our preferred model explains 42% of the variance in bids versus only 34% from the equilibrium model. Our work is thus broadly supportive of these learning models, which have previously only been tested in lab experiments. In contrast, during the late phase, all models — including the equilibrium model — fit the data well.

By demonstrating convergence, our results confirm the value of equilibrium analysis in predicting long-run play in a stable environment. On the other hand, for short-run counterfactual predictions our results suggest that learning models such as those used here may sometimes be more appropriate. Assuming adaptive best response, for example, may more accurately predict short-run dynamics while also being easier to compute and selecting which equilibrium will be played when there are multiple equilibria (Lee and Pakes 2009). Our paper suggests that researchers should be thoughtful in how they generate their counterfactual

predictions, paying particular attention to the relevant time horizon for their predictions.

Related literature. Our paper is closely related to a large body of work in micro, macro and experimental economics. Going back to Cournot (1838), there has been work on the theory of learning in normal-form and, more recently, extensive-form games. This literature mainly aims to derive conditions on the underlying game under which the canonical models of belief-based learning (including fictitious play (Brown 1951)), and reinforcement learning imply convergence to equilibrium (Milgrom and Roberts 1991, Fudenberg and Kreps 1993, Börgers and Sarin 1997, Hart and Mas-Colell 2000). Belief-based learning starts with the premise that players keep track of the history of play and form beliefs about what their rivals will do in the future based on their past play. Reinforcement learning assumes that strategies are “reinforced” by their past payoffs and that the propensity to choose a strategy depends in some way on its stock of reinforcement.

Experimental economists have pushed this theoretical literature further by using lab experiments to determine which learning models best describe how people actually learn (Erev and Roth 1998). On the one hand, this has resulted in the development of more general models such as experience-weighted attraction learning (Camerer and Ho 1999) and models with sophisticated learners who try to influence how other players learn (Camerer, Ho and Chong 2002). On the other hand, there is a growing consensus that telling apart belief-based learning from reinforcement learning is difficult in practice (Salmon 2001).

A second, distinct, theoretical literature considers behavior when agents have only partial knowledge of the environment in which they operate. There is a long literature in applied mathematics and statistics analyzing bandit problems (Robbins 1952), 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 a partially known environment, noting informational free-riding incentives (Bolton and Harris 1999, Keller, Rady and Cripps 2005) and incentives to “signal jam” (Riordan 1985, Mirman, Samuelson and Urbano 1993).

Macroeconomists largely think about learning in terms of expectation formation. The influential idea of adaptive learning (Sargent 1993, Evans and Honkapohja 2001, Evans and Honkapohja 2013) posits that agents proceed like an econometrician does and use the available data to estimate a model of the economy and a rule for forming expectations. The central question then is if under these learning rules the economy reaches a rational-expectations

equilibrium. There is a corresponding experimental literature on expectation formation (Fehr and Tyran 2008, Anufriev and Hommes 2012).

Our paper contributes to these literatures by bringing them together and providing empirical evidence by assessing the fit of different models. We combine fictitious play with adaptive learning to model both strategic and demand uncertainty. There is existing theoretical work on how firms learn about demand (Rothschild 1974, Bergeman and Välimäki 1996, Bergeman and Välimäki 2006, Bernhardt and Taub 2015), but little empirical work. What empirical work there is in the industrial organization and marketing literatures has largely been about how consumers experiment to learn their demand for experience goods (Erdem and Keane 1996, Akerberg 2003, Dickstein 2013), possibly through observational or social learning (Griliches 1957, Porter 1995, Conley and Udry 2010, Zhang 2010, Covert 2013, Newberry 2013)

There has also been little empirical work assessing whether behavior in new markets converges to some notion of equilibrium. Joskow, Schmalensee and Bailey (1998) study the emissions rights market that was created by the 1990 Clean Air Act Amendments, concluding that the market “had become reasonably efficient” (p. 669) within four years. Sweeting (2007) examines the electricity spot market in England and Wales between 1995 and 2000, and finds evidence of tacit collusion between the two largest generators. Hortaçsu and Puller (2008) look at the electricity spot market in Texas from 2001 to 2003, following a restructuring that introduced a uniform-price auction. They find that firms with large stakes made bids that were close to optimal, while small players deviated significantly. 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.

Structure of paper. The paper proceeds as follows. In Sections 2 and 3 we describe the FR market, our data on it, and some descriptive evidence on how this market evolved over time. Section 4 outlines our strategy for estimating the demand and cost primitives. In Section 5 we consider how well different learning models fit the data, before concluding in Section 6. Additional information on the construction of the data are contained in the data appendix. The online appendix presents several robustness checks and extensions.




Figure 1: Overview of the UK electricity market.

2 The FR market

We begin by with an overview of the UK electricity market. It is a network of generators and distributors, connected by a transmission grid. This grid is owned and operated by a company called National Grid plc (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. Figure 1 summarizes the UK electricity market.

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 bilateral forward contracts between generators and distributors. These contracts 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

generators and distributors must submit their contracted positions to NG, as well as bids and offers indicating the terms under which they are willing to be repositioned. NG then acts to equate supply and demand over the settlement period by accepting bids and offers, in something akin to a multi-unit discriminatory auction. This process is called the balancing mechanism (BM), and it accounts for the remaining 2% of electricity sales. The generators bidding in the BM are called BM units. A power station typically consists of multiple BM units, and multiple stations may be owned by the same firm. The BM units belonging to the same station tend to be identical.

Frequency response. NG is obligated by government regulation to maintain a system frequency within a one-percent band of 50 Hertz (Hz, the number of cycles per second). 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 supply and demand in real time, NG instructs one or more BM units into 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 a 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 to satisfy its response requirement. To the best of our knowledge, the response requirement remained unchanged over the sample period.²

FR services are thus a second product, distinct from electricity, that BM units can sell to NG, and the FR market is distinct from the main market (comprised of the BM and bilateral

¹There are in fact three types of FR. 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 the combinations primary-high and primary-secondary-high. To simplify the presentation and analysis, we aggregate the three types of FR; see the data appendix for details.

²We have checked the publicly available minutes of all meetings of the Balancing Services Standing Group (comprising representatives of the generators and NG) and found no discussion of a change in the response requirement.




Figure 2: Holding payment for high response by day pre and post CAP047. Source: National Grid.

forward contracts). Providing FR is costly: 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 also runs less efficiently, with a degraded heat rate. The BM unit 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 supply and demand shocks. The energy response payment compensates the BM unit for actual adjustments to its energy production.³ The energy response payment is considered by industry insiders to be a relatively small source of profit.

Deregulation. Our interest in FR stems from a change in the way the holding payment is determined. This change occurred with the enactment of an amendment to the Connection and Use of System Code called CAP047 and “went live” on November 1, 2005. Pre CAP047, providing FR was mandatory, and the holding payment was at an administered price which had been fairly constant over time (see Figure 2). CAP047 replaced the mandatory provision of FR with a market.

In this market, a BM unit tenders a bid each month for providing FR. The bid for the

³If 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.

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 (measured in £/MWh). 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., it gets “paid-as-it-bids”). Its bid commits the BM unit to offer FR at a fixed price over the next month.⁴

NG can call upon any BM unit at any time, and often does not choose the lowest bidders to provide FR. Instead, it simultaneously accepts bids in the BM and instructs BM units into FR mode 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 that jointly governs the FR market and the BM is solved in real time by a proprietary linear program running on a supercomputer. NG may not choose the lowest bids for at least two reasons. First, BM units differ in the precision of their governor controls, and NG may prefer to call upon more expensive but more precise BM units. The precision of a BM unit is thus a source of product differentiation and market power. Second, because the FR capacity of a BM unit depends on its operating position, NG may prefer to call upon a BM unit operating in the middle of its range, and thus with plenty of FR capacity, rather than a BM unit operating at the extremes of its range. Indeed, NG may first alter the operating position of the BM unit by taking over part of its obligations in the BM before instructing the BM unit into FR mode. As a result, a BM unit does not have to withhold generating capacity from the main market in order to participate in the FR market.⁵

The market for FR was proposed by RWE Npower Renewables Ltd., one of the largest firms in the UK electricity market. This proposal was opposed by NG, who argued that since its demand for FR is regulated and thus inelastic, firms would be able to exploit their market power and the price of FR would rise. The government regulator dismissed these concerns,

⁴The quantity that the BM unit delivers if instructed into FR mode varies with its current operating position and system deviation according to a specific contract between the BM unit and NG that is generally fixed over the sample period. This contract takes the form of a 5×3 matrix for each type of FR (see footnote 1) that specifies the quantity delivered at five deload points (operating positions) and three system deviations (0.2Hz, 0.5Hz, and 0.8Hz away from 50Hz). At other deload points and deviations, the quantity is determined by linear interpolation. The matrices are proprietary information, but selected entries are published by NG in the capability data (see the data appendix). For over 80% of the BM units, the observed entries do not change over the sample period.

⁵Our data shows that BM units can — and do — contract out all of their capacity in the forward market while still actively participating in the FR market. We thank Frank Wolak for pointing out to us that in many other countries the FR market is run separately from the BM. As a result, a BM unit has to withhold generating capacity to participate in the FR market. Because of the resulting opportunity cost, the holding payment is an order of magnitude larger than in the UK.

and on November 1, 2005 introduced CAP047. Figure 2 shows that NG had every reason to worry about CAP047, as the holding payment doubled within the year.

From the pre-CAP047 period, firms had an understanding of the response requirement NG is obligated to satisfy and the relative desirability of their BM units, as well as the cost of providing FR. However, firms faced uncertainty as to the demand for their FR services because they did not know how their rivals would bid in the auction. In addition to this strategic uncertainty, the firms faced demand uncertainty in that they did not know how price sensitive NG was. Our goal is to understand how firms learned to bid in the presence of this uncertainty, and how this contributed to the evolution of the holding payment in Figure 2.

Data. Our analysis focuses on the first six years of the operation of the FR market from November 2005 to October 2011. 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, see footnote 1) by each BM unit at a daily level. The combination of bid and quantity data allows us to calculate the holding payment received by each BM unit.

Our data on the BM comes from Elexon Ltd. Elexon is contracted by the government regulator to manage measurement and financial settlement in the BM. For every BM unit we have data on the bids and acceptances in the BM every half-hour. In combination with data on the contracted position that the BM unit submits to NG one hour prior to the settlement period, this allows us to assess the operating position of the BM unit.

Finally, we collected data on ownership and characteristics of power stations and fuel prices from various sources. See the data appendix for further details on data sources as well as sample and variable construction.

Market participants. There are 130 BM units grouped into 61 power stations owned by 29 firms. The FR market is mildly concentrated with a ten-firm-concentration ratio of just over 80% and an HHI of 76.5. Table 1 summarizes revenue in the FR market for the ten largest firms over the first six years of the market’s existence.

The largest firm, Drax, had over 20% of the FR market and earned about £100,000,000 over the sample period, or about £1,400,000 per month. Drax is a single-station firm,

Table 1: Firms with the largest frequency response revenues

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

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

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

Supply and demand of FR. The demand for and supply of FR are relatively stable over most of the sample period, which makes studying learning and convergence to equilibrium much easier. We argue this using a sequence of figures. Starting with the demand for FR, the left panel of Figure 3 plots the monthly quantity of FR. Though this series is clearly volatile, it is no more volatile at the beginning than at the end of the period we study (and as we show in Section 3, the bids are). The right panel of Figure 3 shows some evidence of modest seasonality.

In addition to the mandatory frequency response (MFR) that is the focus of this paper,

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




Figure 3: MFR quantity by month (left panel) and on average by month-of-year.




Figure 4: MFR and FFR quantities by month (left panel) and fuel prices (right panel).

NG uses long-term contracts with BM units to procure FR services. This is known as firm frequency response (FFR). Figure 4 plots the monthly quantity of FFR and, for comparison purposes, that of MFR (see also the left panel of Figure 3). The quantity of FFR remains relatively stable over our sample period up until July 2010, when it almost doubles and thereafter remains stable at the new level.

Turning from the demand to the supply of FR, the right panel of Figure 4 plots quarterly fuel prices paid by power stations in the UK over time. Fuel prices may matter for the FR market in that they change the “merit order” in the main market. For example, when gas is relatively expensive, gas-powered BM units may be part-loaded and therefore available for FR, whereas coal-powered BM units may be operating at full capacity and thus require




Figure 5: Number of active power stations by month (left panel) and market share of always-active power stations (right panel).

repositioning in the BM in preparation for providing FR. Though there are some upward trends in oil and — to a lesser extent — gas prices, they are largely confined to the end of the sample period.

Finally, a BM unit can opt out of the FR market by submitting an unreasonably high bid. The left panel of Figure 5 plots the number of “active” power stations over time, where we define a station as active if one of its BM units submits a competitive bid of less than or equal to £23/MWh (see Appendix A.2 for details). The number of active stations fluctuates a bit, ranging from 53 to 61 over the sample period. In the first four years of the FR market, the fluctuations are relatively small and none of the stations who become active or inactive is particularly large. The right panel of Figure 5 shows that the share of stations that are always active is steady at around 95%. There are some larger fluctuations in last two years of the FR market.

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



Figure 6: Quantity-weighted and unweighted FR price by month. Weights are based in month t .

3 Evolution of the FR market

Our discussion divides the evolution of the FR market into three phases that differ noticeably in bidding behavior. Figure 6 shows the average monthly price of FR, computed as the quantity-weighted average bids, with vertical lines separating the three phases. For comparison purposes, Figure 6 also shows the unweighted average bids.

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

The movements in the price of FR in the early phases in Figure 6 occurred despite the relative stability of the demand and supply conditions (see Section 2), and are too persistent to be driven by seasonality in the demand for FR. Although there are some changes in FFR and an upward trend in the number of active power stations as well as in the oil and gas

prices, most of that action occurs towards the end of the sample period, when the price of FR has become quite stable. We therefore look for an alternative explanation for the changes in bidding behavior over time. In particular since none of the participants in this market had any experience bidding into it, it seems unlikely that they had strong priors about either how their competitors would bid, or about how their allocation of FR would vary with their bid conditional on how their competitors would bid. We begin with a summary of how bidding behavior changed from one phase to the next. After providing the overview, we look more closely at the role of individual power stations.

Early or rising-price phase (November 2005 – February 2007). In the early or rising-price phase, firms change the bids of their BM units more often and by larger amounts (in absolute value) than in the middle and late phases. On average, the bids of 4 out of 10 BM units change each month by between £1/MWh and £3/MWh (conditional on changing). This is illustrated in Figures 7 and 8.

In addition to changing their bids more often and by noticeably larger amounts, firms tender very different bids in the early phase. Figure 9 shows that the range of bids as measured by the variance of bids across BM units is an order of magnitude larger than in the middle and late phases.

Comparing the left and right panels of Figure 10 shows that most of the variance stems from differences in bids between firms (across-firm variance, right panel) rather than from differences between BM units within firm (within-firm variance, left panel). What within-firm variance there is, is highest in the early phase and then declines, suggesting that firms initially experimented by submitting different bids for their BM units, and that such experimentation became less prevalent over time.

Figure 11 shows the monthly bids of the eight largest power stations by revenue in the FR market. The top left panel provides a more detailed look at the early phase. In line with the wide range of bids documented in Figure 9 and the right panel of Figure 10, the levels and trends of the bids are quite different across stations. Firms seem to experiment with different bids during the early phase of the FR market. Barking, Peterhead and Seabank bid very high early on — pricing themselves out of the market — and then drift back down into contention. The remaining stations start low and then gradually ramp up. The big increase in bids by Drax during late 2006 and early 2007 leads to the “price bubble” in Figure 6.




Figure 7: Quantity-weighted and unweighted probability of a bid change between month t and $t - 1$. Weights are based in month $t - 1$.

Middle or falling-price phase (March 2007 – May 2009). In the middle or falling-price phase, firms change the bids of their BM units less often and by much smaller amounts (in absolute value) than in the early phase. As Figures 7 and 8 illustrate on average the bids of 3 out of 10 BM units change each month by around £1/MWh (conditional on changing). Figure 9 shows that the range of bids is much narrower than in the early phase.

The top right panel of Figure 11 provides more detail. The “price bubble” bursts when Seabank and Barking sharply decrease their bids and steal significant market share from Drax. Drax follows Seabank and Barking down, and this inaugurates intense competition and the noticeable downward trend in the price of FR in Figure 6. Experiments with increased bids are not successful. Drax, for example, increased its bid at the end of 2007 for exactly two months, giving its rivals an opportunity to see its increased bid and follow suit. When no one did, Drax decreased its bid.⁷

The dominant trend in the top right panel of Figure 11 is for the bids of the different power stations to move toward one another. The way this happens is that the stations that entered the middle phase with relatively high bids decreased their bids while the firms that entered

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




Figure 8: Quantity-weighted and unweighted absolute value of bid change conditional on changing between month t and $t - 1$. Weights are based on month $t - 1$ and are zero if the BM unit’s bid did not change.

the phase with relatively low bids maintained those bids. This intense competition generated the marked decrease in the range of bids in Figure 9.

Late or stable-price phase (June 2009 – October 2011). In the late or stable-price phase, firms change the bids for their BM units as often as in the middle or falling-price phase, but by much smaller amounts (in absolute value). As Figures 7 and 8 illustrate, on average, the bids of 3 out of 10 BM units change each month by around £0.5/MWh (conditional on changing). Figure 9 shows that the range of bids is again much narrower than in either of the earlier phases. The bottom panel of Figure 11 provides more detail. While bids at some power stations continue to fall (Rats and Cottam), others are more erratic or rise (Drax and Eggborough), and others are almost completely flat (Peterhead). Overall, however, the bids of the different stations are noticeably closer to one another in this phase. By the time the FR market has entered its late phase, the impression prevails that it has reached a “rest point” that is periodically perturbed by small changes in the physical environment.

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




Figure 9: Quantity-weighted and unweighted variance in bids across BM units by month. Weights are based in month t .

bids grow closer and the frequency and size of adjustments to bids falls.

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

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

Finally, we view the late or stable-price phase as the FR market having reached an understanding of the behavior of competitors, the resulting allocation of FR, and the likely impact




Figure 10: Quantity-weighted and unweighted variance in bids within a firm (left panel) and across firms (right panel). The right panel shows quantity-weighted variance across firms in the quantity-weighted mean firm bids and the unweighted variance across firms in the unweighted mean firm bids.

of changes in the physical environment. As a result, firms are able to adjust with quick small changes to the perturbations which occurred in the late phase.

4 Demand and cost estimation

In this section we model and estimate the demand and cost primitives under a relatively weak rationality assumption. These serve as an input to the learning models we use in Section 5 to better understand the data from the middle and later phases of the FR market.

4.1 Demand

We estimate a generously parameterized logit model at the BM unit-month level to approximate the market shares that are being generated by the proprietary linear program that NG solves in real time to satisfy its response requirement by instructing BM units into FR mode. We focus on the $J = 72$ BM units owned by the ten largest firms in Table 1.⁸ Together these “inside goods” account for just over 80% of revenue in the FR market. We treat the remaining BM units as parts of the “outside good.”

⁸Due to non-competitive or missing bids, we subsume 10 of the 82 BM units into the outside good.

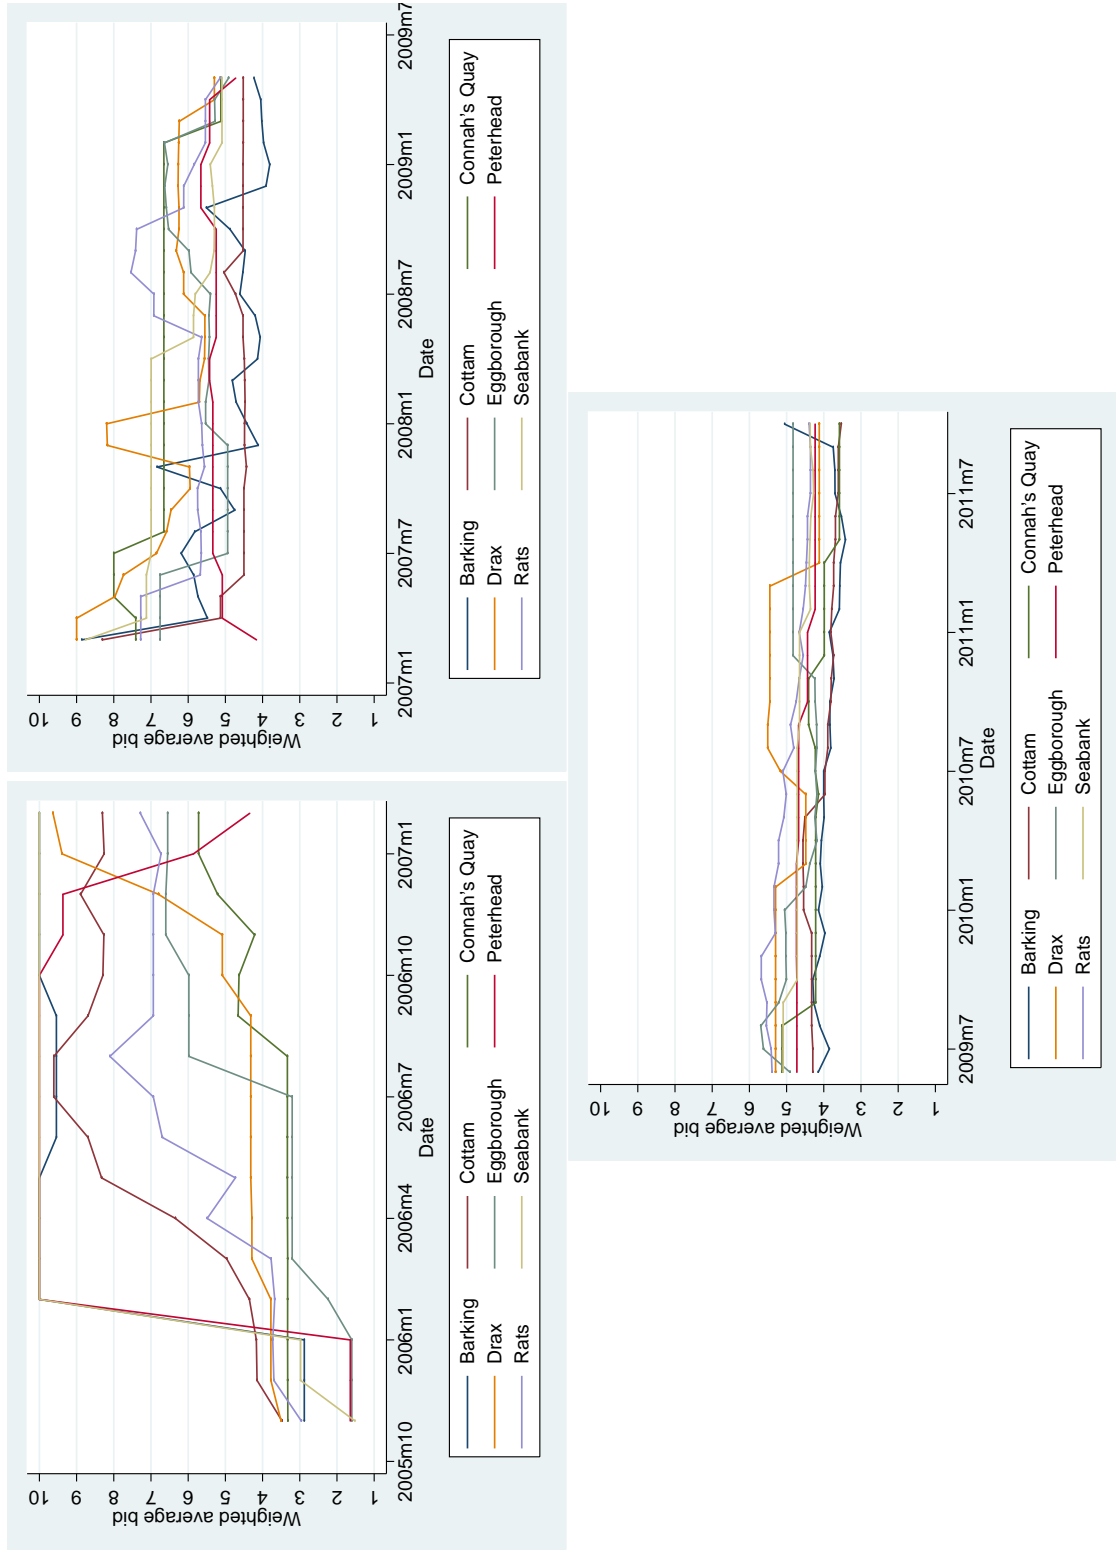


Figure 11: Quantity-weighted average bids of the largest power stations by month. November 2005 – February 2007 (top left panel), March 2007 – May 2009 (top right panel), and June 2009 – October 2011 (bottom panel). Stations ranked by revenue in the FR market during early and middle phases. Bids are censored above at £10/MWh to improve visual presentation.

In addition to parsimoniously parameterizing own- and cross-price elasticities when there are this many goods, an advantage of using a logit model for market shares is that it avoids having to model market size. As the right panel of Figure 3 shows, the monthly quantity of FR is seasonal. A disadvantage of using a logit model is that we have to explicitly deal with a BM unit that has a zero share in a month. To do so, we combine our logit model with a probit model that captures whether the BM unit is eligible for providing FR services.

Model. Let i index firms, j BM units, and t 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 $b_t = (b_{i,t}, b_{-i,t})$.

Let $s_{j,t}$ denote the market share of BM unit j in month t and $s_{0,t} = 1 - \sum_j s_{j,t}$ the market share of the outside good. Let $e_{j,t} = 1(s_{j,t} > 0)$ be the indicator for BM unit j being eligible for providing FR services — and thus having a positive market share — in month t . Accounting for eligibility, we specify a logit model for the market share of BM unit j in month t as

$$s_{j,t} = \frac{e_{j,t} \exp(\alpha \ln b_{j,t} + \beta x_{j,t} + \gamma_j + \mu_t + \xi_{j,t})}{1 + \sum_k e_{k,t} \exp(\alpha \ln b_{k,t} + \beta x_{k,t} + \gamma_k + \mu_t + \xi_{k,t})}, \quad (1)$$

where γ_j and μ_t are BM-unit and month fixed effects and $x_{j,t}$ and $\xi_{j,t}$ are observable and unobservable (to the econometrician) characteristics of BM unit j in month t .

The month fixed effect μ_t subsumes any time-varying characteristics of the outside good. The BM-unit fixed effect γ_j captures the time-invariant preferences of NG for a BM unit due to, e.g., the precision of its governor controls and possibly also transmission constraints. In addition to its bid $b_{j,t}$, BM unit j has time-varying observed characteristics, $x_{j,t}$, and a time varying unobserved characteristic, $\xi_{j,t}$, in month t which are meant to capture the main time varying forces that influence demand in the FR market. The observable characteristics $x_{j,t}$ include two controls for the operating position of the BM unit, namely the fraction of the month the BM unit is fully loaded and the fraction of the month it is part-loaded. As discussed in Section 2, NG uses long-term contracts to procure FFR services that may be a substitute for MFR services. To capture this, $x_{j,t}$ further includes a dummy for whether BM unit j is under contract with NG in month t and provides positive FFR volume. Finally, we

allow the unobservable characteristics $\xi_{j,t}$ to follow an $AR(1)$ process with

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

where the innovation $\nu_{j,t}$ is iid across BM units and months and mean independent of current and past bids $(b_{j,\tau})_{\tau \leq t}$ and observable characteristics $(x_{j,\tau})_{\tau \leq t}$. This setup allows a firm to condition its current bid on past unobservable (to the econometrician) characteristics but not on the current innovation, in line with the fact that the bid for the current month is submitted before the 20th of the previous month.

Our probit model for BM unit j being eligible for providing FR services in month t is

$$e_{j,t} = 1(\check{\beta}x_{j,t} + \check{\gamma}_j + \check{\mu}_t + \eta_{j,t} > 0),$$

where $\check{\gamma}_j$ and $\check{\mu}_t$ are BM-unit and month fixed effects, $x_{j,t}$ are the same observable characteristics of BM unit j in month t as in equation (1), and $\eta_{j,t} \sim N(0, 1)$ is a standard normally distributed disturbance that is iid across BM units and months and, similar to $\nu_{j,t}$, mean independent of current and past bids and observable characteristics.⁹ It follows that

$$\Pr(e_{j,t} = 1 | x_{j,t}) = 1 - \Phi(-\check{\beta}x_{j,t} - \check{\gamma}_j - \check{\mu}_t) = \Phi(\check{\beta}x_{j,t} + \check{\gamma}_j + \check{\mu}_t), \quad (2)$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function (CDF). We estimate equation (2) by maximum likelihood (ML).

Equation (1) implies

$$\ln s_{j,t} - \ln s_{0,t} \equiv \delta_{j,t} = \alpha \ln b_{j,t} + \beta x_{j,t} + \gamma_j + \mu_t + \xi_{j,t} \quad (3)$$

as long as $e_{j,t} = 1$. We can estimate equation (3) by ordinary least squares (OLS) if $\rho = 0$ and $\nu_{j,t}$ is independent of $\eta_{j,t}$.

However, if $\rho \neq 0$, then OLS is biased to the extent that $\xi_{j,t}$ is correlated with $b_{j,t}$ or $x_{j,t}$.

⁹While we allow the probability of having a positive market share to differ across BM units and months, we assume that it is not affected by the bid itself. In the online appendix we include the log bid $\ln b_{j,t}$ in a number of ways and show that although it is statistically significant, it is economically small: in our preferred specification, a £1/MWh increase in bid (corresponding to 18% of the mean and 36% of the standard deviation of bids) decreases the probability of being eligible by -0.021 on a baseline of 0.75, or by about 2.8%. Including the bid as a determinant of eligibility therefore has very little effect on the results of interest but would considerably complicate the analysis of leaning and equilibrium.

Because $\xi_{j,t-1}$ is at least partially known to the firm when it chooses $b_{j,t}$, we may expect $b_{j,t}$ to be a function of $\xi_{j,t-1}$, which is in turn correlated with $\xi_{j,t}$ if $\rho \neq 0$. To deal with this, we quasi-first-difference equation (3) to obtain

$$\delta_{j,t} - \rho\delta_{j,t-1} = \alpha(\ln b_{j,t} - \rho \ln b_{j,t-1}) + \beta(x_{j,t} - \rho x_{j,t-1}) + \tilde{\gamma}_j + \tilde{\mu}_t + \nu_{j,t}, \quad (4)$$

where $\tilde{\gamma}_j = (1 - \rho)\gamma_j$ and $\tilde{\mu}_t = \mu_t - \rho\mu_{t-1}$. As long as $e_{jt} = e_{jt-1} = 1$ and $\nu_{j,t}$ is independent of $\eta_{j,t}$, we can estimate equation (4) by non-linear least squares (NLLS).¹⁰ We maintain this independence assumption in the main text for ease of presentation (allowing for correlation has little effect on the estimated parameters, see the online appendix).

Data. Table 2 summarizes the data used in the estimation. Over the first six years of the operation of the FR market, we have 5175 observations at the BM unit-month level. Market shares are small with an average of 1%, although there is considerable heterogeneity and the maximum over months and BM units is 13%. In about 25% of observations, the market share is zero. Bids are £5.5/MWh on average. Some data on operating position is missing, and where it is, we include a dummy for missing operating position in $x_{j,t}$ and interact it with the controls for being fully loaded and part-loaded.

Results. The first column of Table 3 shows OLS estimates from equation (3) and the second column NLLS estimates from equation (4). The number of observations differs because we require $s_{j,t} > 0$ for OLS and $s_{j,t} > 0$ and $s_{j,t-1} > 0$ for NLLS.

The estimates are remarkably similar across specifications. Because market shares are small, the coefficient on log bid closely approximates the price elasticity of demand. It is negative and significantly less than -1 , as one would expect. The coefficients on fully loaded and part-loaded in $x_{j,t}$ are positive and significant. This makes sense because a BM unit can provide FR only if it is currently operating. The coefficient on part-loaded is larger than that on fully loaded in line with our expectation that NG prefers to call upon a BM unit in the middle of its operating range. The coefficient on positive FFR volume in $x_{j,t}$ is negative and significant, indicating that a BM unit has a smaller share of the MFR market

¹⁰Due to the BM-unit fixed effects equation (4) is estimated using a “within” estimator. The transformation used in estimation has the average of both $\nu_{j,t}$ and $\ln b_{j,t}$ on the “right hand side” of the estimation equation. Correlation between these terms is a possible source of bias in the parameter estimates. The econometrics literature shows that this bias in a linear (balanced) panel model is of the order ρ/T (Nickell 1981), and since we observe a BM unit for a median of $T = 72$ months we ignore it.

Table 2: Summary Statistics (top 10 firms only)

	Mean	Std. Dev.	Min	Max
Share	0.0115	0.0160	0	0.131
Eligibility	0.752	0.432	0	1
Bid	5.453	2.759	1.515	21.00
Fully loaded	0.133	0.236	0	0.997
Part loaded	0.551	0.373	0	1
Missing operating position	0.115	0.319	0	1
Positive FFR volume	0.00734	0.0854	0	1
Number of observations	5175			

Summary statistics on the frequency response market. An observation is a bmonth-month, and the sample is restricted to units owned by the top 10 biggest firms (ranked by revenue over the sample period). Fully loaded is the fraction of time the unit's final physical notification is that it is fully loaded (i.e. operating at or close to capacity). Part loaded is the corresponding fraction when it is operating below capacity. FFR volume is the quantity of FR provided through firm frequency response contracts (i.e. outside of this market).

if it is already under contract with NG, also as expected. Finally, the NLLS estimates from equation (4) in the second column of Table 3 provide evidence of persistence in the unobservable characteristics $\xi_{j,t}$ as the $AR(1)$ coefficient ρ is positive and significant.

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

To assess goodness of fit, we predict the market share of BM unit j in month t conditional on $s_{j,t} > 0$. To do so, we sample independently and uniformly from the empirical distribution of residuals $\hat{\xi}_{j,t}$ for the OLS specification in equation (3) and from the empirical distribution of residuals $\hat{\nu}_{j,t}$ for the NLLS specification in equation (4).¹¹ In both cases we repeatedly

¹¹In the latter case, we proceed as follows: We first obtain the residuals $\hat{\nu}_{j,t}$ along with the estimated parameters $\hat{\alpha}$ and $\hat{\beta}$ from equation (4). We then rewrite equation (3) as $\delta_{j,t} - \alpha \ln b_{j,t} - \beta x_{j,t} = \gamma_j + \mu_t + \xi_{j,t}$, substitute in $\hat{\alpha}$ and $\hat{\beta}$, and estimate by OLS. This yields the residuals $\hat{\xi}_{j,t}$ along with the estimated BM-unit and month fixed effects $\hat{\gamma}_j$ and $\hat{\mu}_t$. We simulate $\xi_{j,t}$ by substituting $\hat{\xi}_{j,t-1}$ and a draw from the empirical distribution of residuals $\hat{\nu}_{j,t}$ into the law of motion $\xi_{j,t} = \rho \xi_{j,t-1} + \nu_{j,t}$. If BM unit j has a zero share in month $t-1$ so $\xi_{j,t-1}$ is missing, then we go back to the first month $\tau_1 < t-1$ such that $s_{j,\tau_1} > 0$ and we go forward to the first month $\tau_2 > t-1$ such that $s_{j,\tau_2} > 0$. We assume that $\nu_{j,l} = \nu$ for all $l = \tau_1, \dots, \tau_2$ and solve the equations $\xi_{j,t-1} = \rho^{t-1-\tau_1} \xi_{j,\tau_1} + \nu \sum_{l=0}^{t-1-\tau_1-1} \rho^l$ and $\xi_{j,\tau_2} = \rho^{\tau_2-t+1} \xi_{j,t-1} + \nu \sum_{l=0}^{\tau_2-t} \rho^l$ for ν and $\xi_{j,t-1}$. If missing for a stretch at the beginning so that τ_1 is not defined, then we use the second equation alone with $\nu = 0$; if missing for a stretch at the end so that τ_2 is not defined, then we use the first equation

Table 3: Demand System Estimates

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

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

sample to integrate out over the empirical distribution of residuals. The logit model fits the data reasonably well. Comparing realized and predicted market shares from equation (3) and equation (4), we get an R^2 of 0.49 and 0.67. This reinforces the importance of persistence in the unobservable characteristics $\xi_{j,t}$ and prompts us to take the NLLS estimates from equation (4) in the second column of Table 3 as our leading estimates.

Figure 12 shows that the fit is good even for the largest power stations, whose market shares change quite dramatically from one month to the next. This indicates that the good fit is not solely a consequence of having BM-unit fixed effects.¹²

alone with $\nu = 0$.

¹²We have done a number of robustness checks. The most notable is that the estimate for α decreases from -1.614 in the middle column of Table 3 to -1.801 if we instrument for $\ln b_{j,t}$ in equation (4). To do so, we first regress $\ln b_{j,t}$ on $\delta_{j,t-1}$, $\ln b_{j,t-1}$, $x_{j,t}$, $x_{j,t-1}$, and BM-unit and month fixed effects $\tilde{\gamma}_j$ and $\tilde{\mu}_t$ and predict $\widehat{\ln b_{j,t}}$. We then replace $\ln b_{j,t}$ by $\widehat{\ln b_{j,t}}$ in equation (4) and estimate by NLLS. While the estimate for α decreases, the estimate for β remains virtually unchanged. We find the same if we additionally include $b_{j,t-2}$ in the first-stage regression. These changes are not large enough to affect our conclusions.




Figure 12: Goodness of fit. Realized (blue, solid) and predicted (red, dashed) market share by month for the four largest power stations Drax (top left panel), Eggborough (top right panel), Ratcliffe (bottom left panel), and Barking (bottom right panel).

4.2 Cost

Since the firms we are modeling have been providing FR for a long time, we assume that they know their cost. However, we as researchers do not. With demand estimated, we therefore turn to estimating cost as both are inputs into the learning models in Section 5.

The main source of cost is the additional wear and tear that a BM unit incurs while in FR mode, which we expect to be relatively stable over time. Let c_j denote the constant marginal cost of BM unit j for providing FR. The realized profit of firm i in month t is

$$\pi_{i,t} = \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_j) M_t s_j(b_t, x_t, \xi_t, e_t; \theta^0), \quad (5)$$

where M_t is market size in month t and our notation emphasizes that the market share of BM unit j in month t depends on the bids b_t , characteristics x_t and ξ_t , eligibilities e_t of all BM units, as well as on the true parameters θ^0 of the demand system. In contrast to market share, market size M_t is independent of bids b_t because the response requirement NG is obligated to satisfy is exogenously determined by government regulation as a function of the demand for electricity.

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

$$\max_{b_{i,t}} \mathcal{E}_{b_{-i,t}, \xi_t, e_t, \theta_t}^{i,t} \left[\sum_{j \in \mathcal{J}_i} (b_{j,t} - c_j) M_t s_j(b_t, x_t, \xi_t, e_t; \theta_t) \middle| \Omega_{t-1} \right]. \quad (6)$$

The notation in equation (6) is designed to stress the two main sources of uncertainty that a firm faces, namely (i) strategic uncertainty about its rivals' bids $b_{-i,t}$ and (ii) demand uncertainty generated by the realizations of ξ_t and e_t and the fact that the parameters θ_t of the logit model may not be known. From the perspective of the firm, the demand parameters θ_t are therefore a random variable. Using the information available to it, the firm forms perceptions about $b_{-i,t}$, ξ_t , e_t , and θ_t .¹³ These perceptions underlie the expectation operator $\mathcal{E}_{b_{-i,t}, \xi_t, e_t, \theta_t}^{i,t} [\cdot | \Omega_{t-1}]$ in equation (6). How perceptions are formed is the central question for the learning models that we turn to in Section 5, but for now we remain agnostic. Finally, we assume that the marginal cost $c = (c_j)_{j=1, \dots, J}$ of all BM units are common knowledge, so that at time t all firms have a common information set Ω_{t-1} . The firms may nevertheless have different perceptions about $b_{-i,t}$, ξ_t , e_t , and θ_t as indicated by the superscripts of the expectation operator $\mathcal{E}_{b_{-i,t}, \xi_t, e_t, \theta_t}^{i,t} [\cdot | \Omega_{t-1}]$.

Equation (6) implies that current bids do not impact future profit, and because of this rules out most models of experimentation. It is therefore not an appropriate characterization of the bidding behavior in the early phase of the FR market. It also rules out collusive equilibria, since in that case firms act to maximize joint profits rather than those of their

¹³We make the simplifying assumption that the firm has perfect foresight about market size M_t and the characteristics x_t to avoid modeling their perceptions about these less interesting objects.

own BM units alone. Below we come back to the possibility of collusion.

Equation (6) implies that the bids $b_{i,t}$ of firm i in month $t \geq 44$ solve the first-order conditions

$$\mathcal{E}_{b_{-i,t}, \xi_t, e_t, \theta_t}^{i,t} \left[M_t s_k(b_t, x_t, \xi_t, e_t; \theta_t) + \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_j) M_t \frac{\partial s_j(b_t, x_t, \xi_t, e_t; \theta_t)}{\partial b_{k,t}} \Big|_{\Omega_{t-1}} \right] = 0, \quad \forall k \in \mathcal{J}_i. \quad (7)$$

The system of first-order conditions (7) does not suffice to estimate marginal cost c_i , however, for the simple reason that we have not specified how the firm forms perceptions. So to derive an estimation equation for c_i , we use a relatively weak rationality assumption to restrict perceptions in a way that we view as appropriate for the late phase.

By the time the FR market enters the late phase, a firm has had ample opportunity to observe how its rivals bid as well as the resulting allocation of market shares. There are changes in the physical environment during the late period, and these do cause changes in bids, but we assume that the firm's bids are by now free of bias. In particular, we assume that if we substitute in realized market sizes $\{M_t\}$ and market shares $\{s_{k,t}\}$ as well as our estimate $\hat{\alpha}$ from Table 3, then the bids $b_{i,t}$ of firm i in month $t \geq 44$ satisfy the time-averaged first-order conditions

$$\frac{1}{29} \sum_{t=44}^{72} \left[M_t s_{k,t} + \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_j) M_t (1(k=j) - s_{k,t}) \frac{\hat{\alpha} s_{j,t}}{b_{k,t}} \right] = 0, \quad \forall k \in \mathcal{J}_i, \quad (8)$$

where we have substituted out for the derivatives in (7) using the properties of the logit ($1(\cdot)$ is the indicator function). We estimate c_i by solving the system of $|\mathcal{J}_i|$ equations (8) for the $|\mathcal{J}_i|$ unknowns. This is straightforward because the equations are linear in the unknowns.

A sufficient condition for our estimation procedure to yield consistent estimates of c_i as the time horizon T grows large beyond $T = 72$ in equation (8) is that the firm's perceptions about $b_{-i,t}$, ξ_t , e_t , and θ_t and our demand parameter estimates $\hat{\theta}$ converge to the truth. Formally, assume (i) that our estimates $\hat{\theta}$ are consistent for θ^0 ; (ii) that the Jacobian of the time-average of (7) in c_i is almost surely invertible and its inverse has a uniformly bounded spectral norm (uniformly in T) almost surely, a technical identification condition that ensures both that bids could be inverted to recover costs if perceptions were known and that cost vectors far from the truth cannot approximately rationalize the bids; (iii) that the subjective probability measure underlying firm perceptions converges weakly to the objective probability distribution (uniformly across information sets). Then as we show

in the appendix, these assumptions suffice for the consistency of our estimation procedure (under some additional weak regularity conditions).

This estimation approach is unusual and deserves some explanation. The first assumption is entirely standard: the econometrician must understand the incentives facing the players in the game (at least asymptotically) in order to back out their costs. The second is almost standard: it must be possible to invert from bids to costs given perceptions, otherwise identification is hopeless. What is unusual here is that the identification condition must hold for all sequences of perceptions the firms may form. Restricting perceptions in some way seems necessary: for example, if a firm is allowed to believe that their bids are irrelevant, then they may tender “crazy” bids that are completely uninformative as to their costs.

Requiring a strong identification condition is the price we pay for the surprising weakness of the final assumption. Instead of assuming that the market is in equilibrium, which entails rational expectations, we just assume that the firm’s perceptions converge to the truth over time. This doesn’t require that the environment necessarily reaches some sort of rest point (the objective distribution need not be stationary), nor does it place any restrictions on how firms form perceptions (e.g., that firms update their perceptions according to Bayes rule). Still, determining which learning models satisfy this convergence condition is a complicated problem that we leave to future research.

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

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

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

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

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

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

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

the same order as the standard error of the estimates. Most of the variation in marginal cost is therefore across stations.

Table 5 shows the result of regressing marginal cost on the characteristics of the BM units. As expected, a (typically smaller) BM unit using dual fuel or oil has lower cost than a BM unit using other fuel types. Moreover, although not statistically significant, the estimates suggest that a BM unit of later vintage has lower cost.

Results: residuals. Using our estimates, we evaluate month-for-month the realized value of the profit derivative $M_t s_k(b_t, x_t, \xi_t, e_t; \hat{\theta}) + \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_j) M_t \frac{\partial s_j(b_t, x_t, \xi_t, e_t; \hat{\theta})}{\partial b_{k,t}}$ in equation (8). For simplicity, we call this value a “residual.” By construction, the residual is zero on average across months for all BM units in the late phase of the FR market. Figure 13 shows the time series of the average residual across BM units. It contrasts the early and middle phases in the left panel with the late phase of the FR market in the right panel. The average residual starts well above zero in the early phase before falling below zero in the middle phase. The standard deviation falls throughout, consistent with our earlier discussion of convergence.

In the late phase, the average residual is above zero in some months and below zero in others and the standard deviation remains relatively constant. Interestingly, even after the substantial increase in FFR volume that occurs in July 2010 (indicated with a dotted line in the right panel, see Figure 4) and changes in participation during this phase (see Figure 5), the average residual continues to be zero and the standard deviations are still an order of

Table 5: Projecting costs onto unit characteristics

	Cost estimate
Unit vintage	-0.015 (0.017)
Dual Fuel	-0.819* (0.466)
Large Coal	-0.463 (0.429)
Medium Coal	-0.683 (0.544)
Oil	-0.967** (0.397)
R^2	0.13
N	71

The dependent variable is the cost estimate c_j . The omitted fuel type is combined cycle gas turbines (CCGT). One observation is dropped because of missing vintage data. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

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

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

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

At the same time, it is striking how the R^2 falls over the three phases of the FR market, indicating that the lagged value explains progressively less of the variation in the residual.




Figure 13: Average and standard deviation of residuals during the early and middle phases (left panel) and late phase (right panel). In the right panel, the dotted line indicates July 2010, when FFR volume nearly doubles.

In the third phase, we find that we find significant autocorrelation using all observations (column 3), but that it essentially disappears when we restrict attention to observations in which the BM unit’s bid changed between months (column 6). Our interpretation of this is that by the end firms have reasonably accurate perceptions, and when they choose to update their bids they do so in a way that accounts for the information contained in the lagged residual.

Collusion? We now briefly examine the possibility of collusion between the firms in our data. We try two different approaches. The first is to look for coordination in the timing and direction of bid changes across BM units, as this could be a sign of collusion being established or breaking down. To capture timing, we define a dummy for BM unit j changing its bid between months $t - 1$ and t and, to capture direction, another dummy for the BM unit increasing its bid. We compute all pairwise correlations between BM units in the dummy for a BM unit changing its bid and in the dummy for a BM unit increasing its bid (conditional on both BM units in the pair changing their bids). In Figure 14 we plot the distribution of correlation coefficients separately for BM units owned by the same firm (“within firm”, left panels) and for BM units owned by different firms (“across firms”, right panels). Note that we expect some across-firm correlation in both the timing and the direction of bid changes due to common shocks to demand.

The within-firm correlations for the timing and direction of bid changes in the left panels

Table 6: Autocorrelation in residuals

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

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

are positive and substantial. This reinforces our contention that decisions are centralized at the level of the firm rather than made at the level of the BM unit. The right panels show correlations evenly distributed around zero, consistent with independent decision making across firms.

Our second approach is more direct: we assume particular collusive arrangements and infer costs given the assumed conduct. Specifically, we re-solve equation (8) for the cost vector that is consistent with observed play during the late phase of the FR market under the assumption that the top 10 firms colluded and maximized the combined profits of all their BM units. This yields an estimated average cost of £-9.8/MWh for the BM units, which is negative and clearly implausible. The estimates are negative because demand is relatively inelastic, and so rationalizing the bids in the face of increased market power requires firms to have low costs. When we repeat the exercise assuming that only the top 3 firms collude, the implied average costs are £-0.25/MWh, still negative. This suggests that there is little or ineffective collusion during the late phase of the FR market. While we cannot rule out collusion in the earlier phases — and indeed we believe Drax attempted to establish tacit collusion in the middle phase — the lack of significant bid correlation is suggestive evidence against this.

Repositioning in the BM. One might worry that equation (5) does not reflect the full set of incentives a firm faces, as it does not account for the profit that accrues to a BM unit as it is repositioned in the BM in preparation for providing FR. In the online appendix, we incorporate these incentives. Using additional data on the BM, we first model and estimate




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

demand for repositioning. Extending equation (8), we then simultaneously estimate the marginal cost of providing FR and a markup on repositioning. The estimated markup is very small and not statistically different from zero, and the marginal cost of providing FR does not change materially from that reported earlier in this section.

The markup in part reflects the amount of attention paid to the profit from repositioning when deciding on the FR bid. Our estimate may thus be explained by the fact that FR bids and bids in the BM are made by different people within the firm, and those deciding on FR bids may not pay attention to the BM. This is consistent with our conversations with Ian Foy, who told us that people in the industry do not think of repositioning incentives when deciding on FR bids.

5 Learning and equilibrium

In this section we consider how well different learning models fit the data. We noted in Section 3 that realistically accounting for the bidding behavior in the early phase of the FR market requires an explanation for the heterogeneity in the way firms learn and a model that allows for experimentation. These are topics we do not tackle here. Instead, we consider models of the bidding behavior in the middle and late phase of the FR market. The middle phase is characterized by a convergence of bids in a relatively stable physical environment whereas there are several environmental changes in the late phase.

Our learning models capture the two main sources of uncertainty that a firm faces, namely strategic uncertainty about its rivals' bids $b_{-i,t}$ and demand uncertainty generated by the realizations of ξ_t and e_t and by the fact that the parameters θ of the logit model may not be known. As noted in Section 1, the literature traditionally uses different types of models for how a firm forms perceptions about rivals' bids and for how the firm forms perceptions about demand, and so do we. Our learning models combine fictitious play and adaptive learning about demand.

Recall that prior to CAP047 providing FR was mandatory, so at the start of our study firms already had quite a bit of experience with demand and cost. We therefore assume throughout that firms know the marginal cost c of all BM units, the $AR(1)$ process generating ξ_t , the objective probability distribution of e_t , and the BM-unit fixed effects $\gamma = (\gamma_j)_{j=1,\dots,J}$. The latter capture the time-invariant preferences of NG for the different BM units. However, there is reason to think firms had to learn about other aspects of demand. In particular, since pre CAP047 the holding payment was at an administered price which had been fairly constant over time, firms may not have been able to assess the sensitivity of NG to the bids they submit, a sensitivity captured by the parameter α in our model. They may also have been uncertain about the parameter β to the extent that the time-varying characteristics $x_{j,t}$ mattered differently for NG post CAP047. Finally, firms may have been uncertain about the month fixed effect μ_t that subsumes any time-varying characteristics of the outside good.

We measure fit by predicting bids under each of our candidate models, and comparing our predictions to the observed bids. We make two kinds of predictions. The first is “one-period predictions”: for each month t during the second and third phases, we take the data available to the firms at the time they bid (which includes bid vectors $\{b_r\}_{r \leq t-1}$ and market shares $\{s_r\}_{r \leq t-2}$) and predict their bids. This corresponds to the thought experiment of predicting

the next move of a player in a game and is analogous to the one-step-ahead predictions used to assess predictive accuracy in the experimental and computational literatures (Erev and Roth 1998, Fershtman and Pakes 2012).

But in industrial organization, analysts are often asked to predict how pricing will evolve for many periods following some market change. So we also consider a second “multi-period” prediction exercise, in which we sequentially predict bid vectors and market shares and then overwrite the observed data with our predictions as we go. To see how these differ, consider the following example: if the strategic model specifies that firms best respond to last period’s bids, and we predict in period t that unit 1 will bid 7 but it actually bids 11, then in the one-period prediction exercise we will predict that in period $t+1$ the other firms will best respond to the actual bid of 11, but in the multi-period prediction exercise they will best respond to the simulated bid of 7. This allows the possibility that our predictions will increasingly diverge from the observed data.

As a baseline we also compute a complete information Nash equilibrium. The literature often refers to this rational-expectations model, so how it compares to our learning models is of some interest.

5.1 Complete information Nash equilibrium

The bids b_t^* in a complete information Nash equilibrium solve the system of J equations

$$\mathbb{E}_{\xi_t, e_t} \left[M_t s_k(b_t^*, x_t, \xi_t, e_t; \theta^0) + \sum_{j \in \mathcal{J}_i} (b_{j,t}^* - c_j) M_t \frac{\partial s_j(b_t^*, x_t, \xi_t, e_t; \theta^0)}{\partial b_{k,t}} \middle| M_t, x_t, \xi_{t-1}, \theta^0, c \right] = 0, \quad \forall k = 1, \dots, J. \quad (9)$$

The notation $\mathbb{E}_{\xi_t, e_t} [\cdot]$ in equation (9), in contrast to $\mathcal{E}_{b_{-i,t}, \xi_t, e_t, \theta_t}^{i,t} [\cdot]$ in equation (6), is meant to indicate that the expectation is computed with respect to the objective probability distribution of ξ_t and e_t conditional on the information available to firm i in month $t-1$ when it submits its bids $b_{i,t}$ for month t . We assume that the firm has perfect foresight about market size M_t , the observable characteristics x_t , and the marginal cost c of all BM units. In addition, we condition the expectation operator on the unobservable characteristics ξ_{t-1} with the implicit understanding that $\xi_{j,t}$ follows the $AR(1)$ process $\xi_{j,t} = \rho \xi_{j,t-1} + \nu_{j,t}$ with ρ known. Finally, we abstract from uncertainty about the parameters of the demand system for now and condition the expectation operator on θ^0 . Later on we extend the complete

information Nash equilibrium to accommodate adaptive learning about demand.

Note that the expectation operator in equation (9) does not condition on $b_{-i,t}$. While the best response of firm i depends on its perceptions of its rivals' bids, in a complete information Nash equilibrium these perceptions are consistent with actual play. A complete information Nash equilibrium is thus obtained by solving the above J equations in the J unknowns b_t^* .

To make equation (9) practical we replace ξ_{t-1} , θ^0 , and c by our estimates of those objects and evaluate the expectation operator using Monte Carlo integration. To this end we generate a random sample $\left(\xi_t^{(s)}, e_t^{(s)}\right)_{s=1,\dots,S}$ with $S = 50,000$ and replace the expectation operator by the corresponding sample average. To obtain $\xi_{j,t}^{(s)}$, we sample independently and uniformly from the empirical distribution of residuals $\hat{v}_{j,t}$ from the NLLS specification in equation (4) (for further details see again footnote 11). To obtain $e_{j,t}^{(s)}$, we sample independently from a Bernoulli distribution with success probability $\Pr(e_{j,t} = 1|x_{j,t})$ as specified in equation (2).

5.2 Learning models

We now explain in detail how in our learning models a firm forms perceptions about rivals' bids and perceptions about demand.

Perceptions about rivals' bids: fictitious play. Belief-based learning starts with the premise that players keep track of the history of play and form beliefs about what their rivals will do in the future based on their past play. We consider fictitious play as a leading example of belief-based learning.¹⁶ In particular, we assume that in month $t - 1$ when firm i chooses its bids $b_{i,t}$ it believes that its rivals' bids $b_{-i,t}$ are sampled from the empirical distribution of their past play. To account for correlation in the bids of firm i 's rivals, we sample an entire vector of rivals' bids. Since firm i may believe that its rivals' bids $b_{-i,t}$ are more similar to more recent observations, we allow for geometrically declining sampling weights. In particular, we assign sampling weight $\delta^{t-\tau-1}$ to rivals' bids $b_{-i,\tau}$ in month $\tau \leq t - 1$, where $\delta \in [0, 1]$ is a decay parameter, and then normalize so the sampling weights sum to one. For month $\tau = t - 1, t - 2, t - 3, \dots$ the sampling weight is therefore proportional to $1, \delta, \delta^2, \dots$ ¹⁷

We let $F(\delta)$ denote our model of fictitious play with decay parameter δ . Note that $\delta = 1$

¹⁶For a discussion of fictitious play and its possible variants see Fudenberg and Levine (1998).

¹⁷In nine instances the bid $b_{j,t}$ of BM unit j in month t is missing. To facilitate sampling, we impute it by going back in time to the first month $\tau < t$ such that $b_{j,\tau}$ is not missing.

indicates no decay and $\delta = 0$ full decay. Adaptive best response thus arises as a special case of fictitious play for $\delta = 0$. Under $F(0)$ firm i believes that its rivals' bids $b_{-i,t}$ in month t are equal to its rivals' bids $b_{-i,t-1}$ in month $t - 1$ with certainty. In the other extreme, under $F(1)$ all past observations are weighted equally.

Perceptions about demand: adaptive learning. To account for uncertainty about demand in addition to uncertainty about rivals' bids, we allow a firm to adaptively learn about the parameters θ of the logit model. Developed in the macroeconomics literature, adaptive learning stipulates that agents learn about parameters in the same way as econometricians by using the available data to estimate them. As more data becomes available over time, agents update their estimates (for a detailed treatment of adaptive learning see Evans and Honkapohja (2001)).

We focus on a simple form of adaptive learning that assumes that firms use the regression procedure described in Section 4.1 and the data available to them to estimate the parameters they are uncertain about. In month $t - 1$ when preparing their bids b_t for month t firms only have data for month $\tau \leq t - 2$ at their disposal because NG does not publish quantities for a month until the very end of the month. We refer to the estimates obtained in this manner as the “sequential” estimates for month t and distinguish them from the “full-sample” estimates obtained in Section 4.1. Throughout we assume that firms ignore any uncertainty in the sequential estimates so that their perceptions are a point mass on those estimates.

We noted above that there is reason to think that firms may have been uncertain about the price sensitivity parameter α , about the coefficient β on the time-varying characteristics $x_{j,t}$, and about the month fixed effect μ_t . Accordingly, we distinguish four combinations of parameters that firms may have been uncertain about, namely α , α and β , α and μ_t , and α , β , and μ_t . Our notation for and implementation of adaptive learning for these four cases plus the baseline case that abstracts from demand uncertainty is as follows:

1. $A(\alpha)$: Using data for month $\tau \leq t - 2$ and fixing β , ρ , $\tilde{\gamma}$, and $(\tilde{\mu}_\tau)_{\tau \leq t-2}$ at the full-sample estimates in Table 3, estimate equation (4) by OLS to obtain the sequential estimate $\hat{\alpha}^{(t)}$ for month t .
2. $A(\alpha, \beta)$: Using data for month $\tau \leq t - 2$ and fixing ρ , $\tilde{\gamma}$, and $(\tilde{\mu}_\tau)_{\tau \leq t-2}$ at the full-sample estimates, estimate equation (4) by OLS to get the sequential estimates $\hat{\alpha}^{(t)}$

and $\hat{\beta}^{(t)}$ for month t .

3. $A(\alpha, \mu_t)$: Using data for month $\tau \leq t - 2$ and fixing β , ρ , and $\tilde{\gamma}$ at the full-sample estimates, estimate equation (4) by OLS to obtain the sequential estimates $\hat{\alpha}^{(t)}$ and $\left(\hat{\mu}_\tau^{(t)}\right)_{\tau \leq t-2}$ for month t . Using data for month $\tau \leq t - 2$ and fixing β and γ at the full-sample estimates and α at the sequential estimate for month t , estimate equation (3) by OLS to obtain the sequential estimates $\left(\hat{\mu}_\tau^{(t)}\right)_{\tau \leq t-2}$ (for more details see footnote 11). Extrapolate $\hat{\mu}_t^{(t)} = \hat{\mu}_{t-2}^{(t)}$ to obtain the sequential estimate $\hat{\mu}_t^{(t)}$ for month t .
4. $A(\alpha, \beta, \mu_t)$: Using data for month $\tau \leq t - 2$ and fixing ρ and $\tilde{\gamma}$ at the full-sample estimates, estimate equation (4) by OLS to obtain the sequential estimates $\hat{\alpha}^{(t)}$, $\hat{\beta}^{(t)}$, and $\left(\hat{\mu}_\tau^{(t)}\right)_{\tau \leq t-2}$ for month t . Using data for month $\tau \leq t - 2$ and fixing γ at the full-sample estimate and α and β at the sequential estimates for month t , estimate equation (3) by OLS to obtain the sequential estimates $\left(\hat{\mu}_\tau^{(t)}\right)_{\tau \leq t-2}$. Extrapolate $\hat{\mu}_t^{(t)} = \hat{\mu}_{t-2}^{(t)}$ to obtain the sequential estimate $\hat{\mu}_t^{(t)}$ for month t .
5. $A(\emptyset)$: Fix all parameters θ at the full-sample estimates in Table 3.

Predictions. We combine fictitious play with adaptive learning to make predictions. A fictitious play model $F(\delta)$ with $\delta \in [0, 1]$ and an adaptive learning model $A(y)$ with $y \in \{\alpha, (\alpha, \beta), (\alpha, \beta, \mu_t), \emptyset\}$ together with available data D_t determine a predicted bid $b_{i,t}^{(\delta,y)}$ for firm i in month t . That bid $b_{i,t}^{(\delta,y)}$ solves the system of \mathcal{J}_i equations

$$\begin{aligned} & \mathcal{E}_{b_{-i,t}, \xi_t, e_t, \theta_t}^{i,t} \left[M_t s_k(b_{i,t}^{(\delta,y)}, b_{-i,t}, x_t, \xi_t, e_t, \theta_t) \right. \\ & \left. + \sum_{j \in \mathcal{J}_i} \left(b_{j,t}^{(\delta,y)} - c_j \right) M_t \frac{\partial s_j(b_{i,t}^{(\delta,y)}, b_{-i,t}, x_t, \xi_t, e_t, \theta_t)}{\partial b_{k,t}} \Big| F(\delta), A(y), D_t \right] = 0, \quad \forall k \in \mathcal{J}_i. \end{aligned} \quad (10)$$

The notation $\mathcal{E}_{b_{-i,t}, \xi_t, e_t, \theta_t}^{i,t} [\cdot]$ indicates that the expectation operator is with respect to the subjective probability distributions of $b_{-i,t}$ and θ_t induced by $F(\delta)$ and $A(y)$. To make equation (10) practical we replace ξ_{t-1} and c in D_t by our estimates. In the single-period prediction exercises, we estimate θ_t for each model $A(y)$ as detailed above using the observed data for month $\tau \leq t - 2$. We evaluate the remaining random variables in the expectation operator by drawing samples $\left(b_{-i,t}^{(s)}, \xi_t^{(s)}, e_t^{(s)} \right)_{s=1, \dots, S}$ with $S = 50,000$, proceeding as in Section

Table 7: Different Learning Models for Middle phase: Single-period predictions.

	Average Diff: Predicted and Actual Bid					Mean Square Error				
	$A(\emptyset)$	$A(\alpha)$	$A(\alpha, \mu)$	$A(\alpha, \beta)$	$A(\alpha, \beta, \mu)$	$A(\emptyset)$	$A(\alpha)$	$A(\alpha, \mu)$	$A(\alpha, \beta)$	$A(\alpha, \beta, \mu)$
F(0)	-0.21	-0.11	1.42	0.61	1.02	1.23	1.21	4.67	2.20	3.22
F(0.5)	-0.20	-0.10	1.44	0.62	1.03	1.23	1.22	4.72	2.23	3.26
F(1)	-0.47	-0.38	1.10	0.31	0.70	1.55	1.43	3.01	1.53	2.09
Eq.	-0.42	-0.29	1.90	0.67	1.29	1.36	1.26	6.70	2.29	4.13

* $F(x)$ denotes fictitious play with $\delta = x$. $A(x)$ denotes adaptive learning where all parameters except x are known and x is sequentially estimated. All statistics are share-weighted.

5.1 for sampling $\xi_{j,t}^{(s)}$ and $e_{j,t}^{(s)}$. We sample $b_{-i,t}$ according to $F(\delta)$ on the data $\{b_s\}_{s \leq t-1}$.

The multi-period prediction exercises are identical except for the data D_t : for any model $(F(\delta), A(y))$ we replace the observed bids and market shares during the second and third phases with our predictions $\{b_t^{(\delta,y)}\}$, thus changing what firms observe.¹⁸ This has two effects. First, when sampling $b_{-i,t}$ according to $F(\delta)$, we will now sample our predicted bids. Second, when estimating θ_t , the dataset used for estimation will now contain our predicted bids and market shares. These effects will generate differences between the single and multi-period predictions. For the multi-period prediction exercises, we only consider two of the adaptive learning models, $A(\emptyset)$ and $A(\alpha)$.

5.3 Results

We measure the predictive accuracy of our various fictitious play and adaptive learning models in two ways: using (i) the share weighted average difference between the predicted and actual bids, where we average over BM units and months and (ii) the share weighted mean square error of prediction. The share weighted average error may be a useful measure if the goal is to predict the overall cost of the FR market to NG. The mean square error is a traditional way of assessing goodness of fit. We begin with discussing how well different learning models fit the data from the middle phase of the FR market and then move on to the late phase.

¹⁸When predicting market shares, we use the estimated parameter vector $\hat{\theta}$ regardless of the adaptive learning model $A(y)$, as this is our best guess as to the true DGP.

Table 8: Different Learning Models for Middle phase: Multi-period predictions.

	Average Diff: Predicted and Actual Bid		Mean Square Error	
	$A(\emptyset)$	$A(\alpha)$	$A(\emptyset)$	$A(\alpha)$
F(0)	-0.41	-0.21	1.31	1.22
F(0.5)	-0.39	-0.20	1.26	1.19
F(1)	-0.53	-0.38	1.46	1.33
Eq.	-0.42	-0.29	1.36	1.26

* $F(x)$ denotes fictitious play with $\delta = x$. $A(x)$ denotes adaptive learning where all parameters except x are known and x is sequentially estimated. All statistics are share-weighted.

Middle or falling-price phase (March 2007 – May 2009). The two panels of Tables 7 and 8 summarize the share weighted average and mean square error for the middle phase of the FR market in single-period and multi-period predictions respectively. Each panel provides a two-way classification: The columns pertain to the baseline case $A(\emptyset)$ that abstracts from demand uncertainty and the adaptive learning models $A(\alpha)$, $A(\alpha, \beta)$, $A(\alpha, \mu_t)$, and $A(\alpha, \beta, \mu_t)$ (last three only in the single-period case). The rows pertain to the fictitious play models $F(0)$ (full decay or adaptive best response), $F(0.5)$ (some decay), and $F(1)$ (no decay) and the complete information Nash equilibrium.

Starting with the single-period predictions, the top left corner of each panel of Table 7 has the smallest (in absolute value) entries. This pattern is starkest in the average error in the left panel but also clearly present in the mean square error in the right panel. Hence, under both accuracy measures, the preferred model for perceptions about rivals' bids is either $F(0)$ or $F(0.5)$. The preferred model for perceptions about demand is either the baseline case $A(\emptyset)$ that abstracts from demand uncertainty or the adaptive learning model $A(\alpha)$ that accounts for firms being uncertain about the price sensitivity parameter α , with the average error indicating a slight preference for $A(\alpha)$ over $A(\emptyset)$. Extending adaptive learning to additional parameters noticeably worsens goodness of fit under both accuracy measures.

We have attempted to more finely determine the decay parameter δ in the fictitious play model by conducting a grid search over $\delta = 0, 0.1, \dots, 1$. While we obtain the smallest mean square error at $\delta = 0.3$, there is very little difference in mean square error as δ ranges from 0 to 0.6. Hence, all we can say is that the data from the middle phase of the FR market appears to favor a fictitious play model in which firms rely disproportionately on more recent observations to form beliefs about rivals' bids. A striking feature of Table 7 is that under both accuracy measures the fictitious play models $F(0)$ and $F(0.5)$ fit the data better than

the complete information Nash equilibrium irrespective of the particular adaptive learning model.

Turning to the multi-period predictions in Table 8, we see a qualitatively similar pattern. The fictitious play models $F(0)$ and $F(0.5)$ continue to fit better than equilibrium, though the differences in fit are in most cases smaller. Once again, $F(0)$ and $F(0.5)$ are hard to distinguish. One can more clearly see the comparative performance of the single and multi-period predictions in Figure 15, where we compare the time series of the predicted price under the model $(F(0), A(\alpha))$ to the actual price and the complete information Nash Equilibrium (the plot with $F(0.5), A(\alpha)$ is similar). As we can see in the figure, adaptive best response along with adaptive learning about the price sensitivity parameter α tracks the actual price of FR fairly closely, especially after month 20.

The single-period model matches the path better, which makes sense since it uses realized rather than predicted bids in making its predictions. Indeed, it is remarkable that the multi-period model fits so well given that errors in predicting behavior have the potential to “blow up” as they form the basis for subsequent predictions. Instead, the multi-period model roughly matches the slow decline in prices that we see in the data. This is not true of the model without adaptive learning (not shown): the multi-period model $(F(0), A(\emptyset))$ predicts that bids reach equilibrium levels within a few periods.

The complete information Nash equilibrium limits the range of the predicted price of FR. This is perhaps unsurprising because ultimately only changes in the demand and cost primitives can effect changes in the Nash equilibrium. In contrast, because it is based on observed past behavior, adaptive best response is better able to capture trends in the actual price of FR. Adaptive learning adds another source of fluctuations in the predicted price of FR as perceptions about demand change with more data becoming available to firms over time.

Late or stable-price phase (June 2009 – October 2011). Tables 9 and 10 are the analogs for the late phase to Tables 7 and 8. Since the patterns in the single and multi-period predictions are also similar in this phase, we discuss them together. Perhaps the first thing to note from Table 9 and Table 10 is that the mean square errors are much smaller than those for the middle phase, sometimes even one third or one tenth of the value. This is especially true for the complete information Nash equilibrium. The second thing to note is that the average errors in Table 9 are similar (in absolute value) to some of the smaller average errors for the middle phase in Table 7.




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

We have again attempted to more finely determine the decay parameter δ in the fictitious play model. While there is a slight preference for $\delta = 0$ in single-period prediction, i.e., adaptive best response, the difference only becomes noticeable for $\delta > 0.6$. In the late phase of the FR market, any fictitious play model for perceptions about rivals' bids with $\delta \leq 0.6$, including adaptive best response, appears to provide an adequate description of the data, a finding which is perhaps not surprising once perceptions have “settled down.”

It is difficult to determine a preferred model for perceptions about rivals' bids from Table 9 and 10. While the data appear to favor the fictitious play models $F(0)$ and $F(0.5)$ over $F(1)$, both $F(0)$ and $F(0.5)$ are essentially indistinguishable from the complete information Nash equilibrium. In contrast to the middle phase of the FR market, perceptions about rivals' bids may well have been in equilibrium during the late phase.

It is similarly difficult to determine a preferred model for perceptions about demand. Looking at the single-period models in Table 9, under both accuracy measures, the baseline case $A(\emptyset)$ that abstracts from demand uncertainty and the adaptive learning models $A(\alpha)$, $A(\alpha, \beta)$, and $A(\alpha, \mu_t)$ are very similar to one another. The surprise is that the adaptive learning model

Table 9: Different Learning Models for Late phase: Single Period Predictions*.

	Average Diff: Predicted and Actual Bid					Mean Square Error				
	$A(\emptyset)$	$A(\alpha)$	$A(\alpha, \mu)$	$A(\alpha, \beta)$	$A(\alpha, \beta, \mu)$	$A(\emptyset)$	$A(\alpha)$	$A(\alpha, \mu)$	$A(\alpha, \beta)$	$A(\alpha, \beta, \mu)$
F(0)	0.21	0.21	0.23	0.25	0.15	0.40	0.40	0.44	0.48	0.43
F(0.5)	0.22	0.22	0.24	0.26	0.15	0.41	0.41	0.44	0.50	0.44
F(1)	0.32	0.33	0.35	0.37	0.26	0.52	0.53	0.58	0.64	0.55
Eq.	0.25	0.25	0.29	0.30	0.16	0.42	0.42	0.49	0.52	0.42

* $F(x)$ denotes fictitious play with $\delta = x$. $A(x)$ denotes adaptive learning where all parameters except x are known and x is sequentially estimated. All statistics are share-weighted.

Table 10: Different Learning Models for Late phase: Multi-period predictions*.

	Average Diff: Predicted and Actual Bid		Mean Square Error	
	$A(\emptyset)$	$A(\alpha)$	$A(\emptyset)$	$A(\alpha)$
F(0)	0.23	0.31	0.40	0.46
F(0.5)	0.23	0.30	0.40	0.45
F(1)	0.23	0.31	0.42	0.49
Eq.	0.25	0.25	0.42	0.42

* $F(x)$ denotes fictitious play with $\delta = x$. $A(x)$ denotes adaptive learning where all parameters except x are known and x is sequentially estimated. All statistics are share-weighted.

$A(\alpha, \beta, \mu_t)$ has the smallest average error irrespective of the particular fictitious play model. Recall that our logit model for market shares is an approximation of the true underlying demand. During the late phase, small changes in the physical environment occur periodically. Although the true underlying demand is stationary, our best approximation to it may not be in light of these changes. A possible explanation for the fact that $A(\alpha, \beta, \mu_t)$ has the smallest average error may therefore be that $A(\alpha, \beta, \mu_t)$ has more free parameters than the other adaptive learning models and is thus able to provide a better time-varying approximation of the true underlying demand.

Figure 16 is the analog for the late phase to Figure 15. It shows clearly that in this phase both the single and multi-period versions of the model ($F(0), A(\emptyset)$) make similar predictions to the complete information Nash equilibrium. Thus the multi-period model converges to the complete information Nash prediction, and at a rate that matches what we observe in the data. Convergence is not generally guaranteed under fictitious play (Fudenberg and Levine 1998), so this is an indication that our pricing game is in some sense “well-behaved”.




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

Summary. Perhaps the most important point to take away is that different models may be appropriate for different periods in an industry’s evolution. In periods a short time after a major change in the industry’s environment, it may be better to rely on a model that allows for learning than to rely on equilibrium for one’s understanding of, or predictions of, behavior. This was true of the second period of our data, in which fictitious play models (including adaptive best response) provided a substantially better fit. This was true both of the single-period and multi-period models, and indeed the multi-period model did well in predicting the overall path of prices during this time period. On the other hand, in the third period firms had experience in predicting the bids of their competitors and NG’s demand, and the equilibrium model performed well despite some changes in the market.

We also observe that learning models that incorporated both demand and strategic uncertainty performed well. Models in which firms learned about α uniformly performed better than those in which it was assumed that they faced no demand uncertainty, and in the third phase allowing additional free parameters generated better fit. Bringing fictitious play models from the microeconomic literature together with adaptive learning models from the

Table 11: Profit Statistics

Rank	Firm name	Months Changed	Matched Direction (%)	Total Profit	Ex-Post Lost Profit (%)	Ex-Post Bid Diff.	Ex-Post Bid Diff. (%)
1	Drax Power Ltd.	23	88	68.4	2.03	1.1	20
2	E.ON UK plc	52	59	44.4	2.01	.76	14
3	RWE plc	15	87	25.1	3.06	.93	23
4	Eggborough Power Ltd	18	58	18.2	1.43	.84	17
5	Keadby Generation Ltd	17	73	14.9	5.12	1.03	19
6	Barking Power Ltd	57	56	11.8	6.17	.65	15
7	Jade Power Generation Ltd	15	60	10	4.74	1.24	24
8	SSE Generation Ltd	17	71	9.4	.47	.42	8
9	Seabank Power Ltd	9	78	9.1	8.21	1.14	20
10	Centrica plc	42	74	8.7	2.59	.92	18

Any single change in bid against the previous month constitutes a bid change by the firm. Matched direction indicates the percent of bid changes that are in the same direction as our estimated ex-post-optimal bid. Ex-post lost profit indicates the percent increase in profit expected by taking the estimated ex-post-optimal bid instead of the actual bid. Ex-post bid difference indicates the share-weighted average distance between the ex-post-optimal and actual bids, in absolute and relative terms.

macroeconomic literature, as we have done here, may prove productive in modeling learning.

5.4 Timing of bid changes

One of the striking empirical regularities of the data is some firms take a far more active approach to bidding than others. This is documented in Table 11, where we count the number of months in which a firm updated a bid of *any* of their units. Firms like E.ON, Barking, and Centrica change their bids in more than half the periods, but the remainder of the top 10 make changes far less frequently (notably including Drax). Conditional on making a change, they generally adjust their bids in the direction predicted by our preferred model, i.e., $F(0)$ in combination with $A(\alpha)$, although the most active firms are less accurate, consistent with our account of some firms “exploring” during the early phase.

The frequency of bid changes is not significantly correlated with firm size (as measured by total profit): the three most active firms are ranked second, sixth and tenth on this size metric. We have also had little success in explaining the timing of bid changes. Define for each firm and time period the “profit opportunity” as the difference between the profit

they would make by bidding optimally and the profit they would make by keeping their bids unchanged (both calculated under our preferred model). One would expect that firms would be more likely to adjust their bids when the profit opportunity is large, yet we find no statistically significant support for this in any of a number of probit regressions where we explore different plausible specifications.

In the sixth column we show the “ex-post lost profit”, defined as the percentage of their observed total profit each firm could additionally have gained by making the ex-post optimal bid given the realizations of $(b_{-i,t}, \xi_t, e_t)$ and evaluated using our demand and cost estimates. The magnitudes are generally small, both as percentages and absolute numbers, which is striking given that the optimal bid is computed using data the firms did not know. Contrast this with the average absolute difference between a firm’s (share-weighted) average bid and the ex-post optimal bid, shown in the last two columns in absolute and relative terms. These magnitudes are much bigger, with most firms placing bids that are around 20% away from their ex-post optimum.

For example, consider RWE plc, a firm that owns many of the power stations who participate in this market, though none of the top 8 shown in Figure 11. They rarely updated their bids, and (with rare exceptions) each of the BM units they owned converged to their equilibrium bids from below, slowly tracking the market upward. As a result, we estimate that RWE plc made 3.06% less, or £768,000 over six years than they could have. This is enough money that one might have expected them to pay more attention and update more frequently (e.g., by tracking more closely), but perhaps not enough to motivate hiring a full-time employee to study the market and optimize their bidding.

This suggests another reason why we have been unable to explain bid levels and changes more precisely: most firms are paying infrequent attention to this market and making sensible but imperfect adjustments when they do. This combination of inattention and “satisficing” behavior (Simon 1955) may be optimal for them, given the human resource and institutional constraints they face (this also has the flavor of rational inattention as in Sims (2003)). Yet as Akerlof and Yellen (1985) have noted, these small departures from perfect rationality may lead to aggregate behavior that is quite different from equilibrium and thus hard to predict with equilibrium models (or indeed with learning models that do not account for these frictions).

6 Conclusion

We have carefully documented the evolution of the FR market in the UK following its deregulation, using a variety of methods. We find that the market converged to a rest point, after three and a half years or 42 periods of interaction. This rest point is consistent with a complete information Nash equilibrium, in that cost estimates derived using this assumption are plausible. There is substantial heterogeneity in how the main market participants approach this rest point. Early on, there is substantial uncertainty about demand and so some firms experiment, while others are more cautious and make infrequent adjustments. During the middle phase, firm behavior is more predictable: bid predictions from fictitious play models in which firms best respond to recent rival behavior are able to explain a substantial share of the bid variance. The fit is further improved by allowing for some adaptive learning about the price parameter. These models match the observed bids better than assuming equilibrium play. In the final phase, all models fit approximately equally well.

This is a case study, and so some caution is warranted. Still, we view the paper as reaching a number of meaningful conclusions. Our research supports the idea that in stable environments play will generally converge to an approximate Nash equilibrium. This is consistent with the theory of learning in games, and the corresponding idea that equilibrium play is a good prediction for the long-run outcome of a game with strategic players.

But for predicting short-run play following a change in the environment, there are a variety of potential approaches. Here we find that models of fictitious play and adaptive learning outperform full-information equilibrium predictions, at least in the second phase of the data. This lends empirical support to these learning models, which to the best of our knowledge have thus far only been tested in the lab. We would thus argue that some thought is warranted before imposing equilibrium play in counterfactual analysis, especially when a short-term prediction is needed. This may be a welcome conclusion for many researchers since equilibrium analysis presents computational challenges that are not present in simulating best response dynamics, for example; and imposing a particular learning model will allow for selection among multiple equilibria.

Finally, we note that Drax changed its bidding strategy only after hiring an employee to analyze the market, and firms were heterogeneous in their activity levels in a way that is hard to explain through profit motives alone. Firms did not fully optimize their bids, but this cost them little in forgone profits. All of this suggests that the standard analysis

of market incentives should ideally be combined with institutional analysis — or at least some robustness to institutional inertia and imperfect optimization — in predicting market evolution. Developing robust tools for counterfactual simulation based on these observations is an interesting topic for future work.

References

- Akerberg, Daniel A.**, “Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination,” *International Economic Review*, 2003, 44 (3), 1007–1040.
- Akerlof, George and Janet Yellen**, “Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria?,” *The American Economic Review*, 1985, 75 (4), 708–720.
- Anufriev, Mikhail and Cars Hommes**, “Evolutionary Selection of Individual Expectations and Aggregate Outcomes in Asset Pricing Experiments,” *American Economic Journal: Microeconomics*, 2012, 4 (4), 35–64.
- Bergeman, Dirk and Juuso Välimäki**, “Learning and Strategic Pricing,” *Econometrica*, 1996, 64 (5), 1125–1149.
- and —, “Dynamic Price Competition,” *Journal of Economic Theory*, 2006, 127, 232–263.
- Bernhardt, Dan and Bart Taub**, “Learning About Common and Private Values in Oligopoly,” *Rand Journal of Economics*, 2015, 46 (1), 66–85.
- Bolton, Patrick and Christopher Harris**, “Strategic Experimentation,” *Econometrica*, 1999, 67 (2), 349–374.
- Borenstein, Severin, James B. Bushnell, and Frank A. Wolak**, “Measuring Market Inefficiencies in California’s Restructured Wholesale Electricity Market,” *American Economic Review*, 2002, 92 (5), 1376–1405.
- Börgers, Tilman and R. Sarin**, “Learning Through Reinforcement and Replicator Dynamics,” *Journal of Economic Theory*, 1997, 77 (1–14).

- Brown, G.W.**, “Iterative Solutions of Games by Fictitious Play,” in T.C. Koopmans, ed., *Activity Analysis of Production and Allocation*, Wiley, 1951.
- Camerer, Colin and Teck-Hua Ho**, “Experienced-Weighted Attraction Learning in Normal Form Games,” *Econometrica*, jul 1999, 67 (4), 827–874.
- , — , and **Juin-Kuan Chong**, “Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games,” *Journal of Economic Theory*, 2002, 104, 837–874.
- Conley, Timothy and Christopher Udry**, “Learning About a New Technology: Pineapple in Ghana,” *American Economic Review*, 2010, 100 (1), 35–69.
- Cournot, Antoine Augustine**, *Recherches sur les principes mathématiques de la thorie des richesses*, Paris: Hachette, 1838.
- Covert, Thomas**, “Learning About Hydraulic Fracturing in the North Dakota Bakken Shale,” Working Paper, Harvard University 2013.
- Dickstein, Michael**, “Efficient Provision of Experience Goods: Evidence from Antidepressant Choice,” Working Paper, Stanford University 2013.
- Erdem, Tülin and Michael P. Keane**, “Decision-Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets,” *Marketing Science*, 1996, 15 (1), 1–20.
- Erev, Ido and Alvin E. Roth**, “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria,” *American Economic Review*, 1998, 88 (4), 848–881.
- Evans, George W. and Seppo Honkapohja**, *Learning and Expectations in Macroeconomics*, Princeton University Press, 2001.
- and — , “Learning as a Rational Foundation for Macroeconomics and Finance,” in Roman Frydman and Edmund S. Phelps, eds., *Rethinking Expectations: The Way Forward for Macroeconomics*, Princeton University Press, 2013.
- Fehr, Ernst and Jean-Robert Tyran**, “Limited Rationality and Strategic Interactions: The Impact of the Strategic Environment on Nominal Inertia,” *Econometrica*, 2008, 76 (2), 353–394.

- Fershtman, Chaim and Ariel Pakes**, “Dynamic Games with Asymmetric Information: A Framework for Empirical Work,” *Quarterly Journal of Economics*, 2012, 127 (4), 1611–1661.
- Fudenberg, Drew and David Kreps**, “Learning Mixed Equilibria,” *Games and Economic Behavior*, 1993, 5, 320–67.
- and **David Levine**, *The Theory of Learning in Games*, Cambridge, MA: MIT Press, 1998.
- Goldfarb, Avi and Mo Xiao**, “Who Thinks About the Competition: Managerial Ability and Strategic Entry in US Local Telephone Markets,” *American Economic Review*, 2011, 101 (5), 3130–3161.
- Griliches, Zvi**, “Hybrid Corn: An Exploration in the Economics of Technological Change,” *Econometrica*, 1957, 25 (4), 501–522.
- Hart, Sergiu and Andreu Mas-Colell**, “A Simple Adaptive Procedure Leading to Correlated Equilibrium,” *Econometrica*, 2000, 68 (5), 1127–1150.
- Heckman, James**, “Sample Selection Bias as a Specification Error,” *Econometrica*, 1979, 47 (1).
- Hortaçsu, Ali and Steven L. Puller**, “Understanding Strategic Bidding in Multi-unit Auctions: A Case Study of the Texas Electricity Spot Market,” *Rand Journal of Economics*, 2008, 39 (1), 86–114.
- Joskow, Paul L., Richard Schmalensee, and Elizabeth M. Bailey**, “The Market for Sulfur Dioxide Emissions,” *American Economic Review*, 1998, 88 (4), 669–685.
- Keller, Godfrey, Sven Rady, and Martin Cripps**, “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 2005, 73 (1), 39–68.
- Lee, Robin S. and Ariel Pakes**, “Multiple Equilibria and Selection by Learning in an Applied Setting,” *Economics Letters*, 2009, 104 (1), 13–16.
- Milgrom, Paul and John Roberts**, “Adaptive and Sophisticated Learning in Normal-Form Games,” *Games and Economic Behavior*, 1991, 3, 82–100.

- Mirman, Leonard J., Larry Samuelson, and Amparo Urbano**, “Duopoly Signal Jamming,” *Economic Theory*, 1993, *3*, 129–149.
- Newberry, Peter**, “An Empirical Study of Observational Learning,” Working Paper, Penn State 2013.
- Newey, Whitney K.**, “Uniform Convergence in Probability and Stochastic Equic,” *Econo*, July 1991, *59* (4), 1161–1167.
- Nickell, Stephen**, “Biases in Dynamic Models with Fixed Effects,” *Econometrica*, 1981, *49* (6), 1417–1426.
- Porter, Robert H.**, “The Role of Information in U.S. Offshore Oil and Gas Lease Auctions,” *Econometrica*, 1995, *63* (1), 1–27.
- Riordan, Michael H.**, “Imperfect Information and Dynamic Conjectural Variations,” *Rand Journal of Economics*, 1985, *16* (1), 41–50.
- Robbins, Herbert**, “Some Aspects of the Sequential Design of Experiments,” *Bulletin of the American Mathematical Society*, 1952, *58* (5), 527–535.
- Rothschild, Michael**, “A Two-Armed Bandit Theory of Market Pricing,” *Journal of Economic Theory*, 1974, *74*, 185–202.
- Salmon, Timothy C.**, “An Evaluation of Econometric Models of Adaptive Learning,” *Econometrica*, 2001, *69* (6), 1597–1628.
- Sargent, Thomas J.**, *Bounded Rationality in Macroeconomics*, Oxford University Press, 1993.
- Simon, Herbert A.**, “A Behavioral Model of Rational Choice,” *Quarterly Journal of Economics*, 1955, *69* (1).
- Sims, Christopher A.**, “Implications of rational inattention,” *Journal of Monetary Economics*, 2003, *50*, 665–690.
- Sweeting, Andrew**, “Market Power in the England and Wales Wholesale Electricity Market 1995–2000,” *Economic Journal*, 2007, *117* (520), 654–685.

- Wolak, Frank A.**, “Identification and Estimation of Cost Functions Using Observed Bid Data: An Application to Electricity,” in Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, eds., *Advances in Econometrics: Theory and Applications*, Cambridge University Press, 2003, pp. 133–169.
- , “Quantifying the Supply-Side Benefits from Forward Contracting in Wholesale Electricity Markets,” *Journal of Applied Econometrics*, 2007, *22* (7), 1179–1209.
- Zhang, Juanjuan**, “The Sound of Silence: Observational Learning in the U.S. Kidney Market,” *Marketing Science*, 2010, *29* (2), 315–335.

A Appendix

A.1 Consistency proof

The proof requires some notation. We begin by more succinctly describing our estimation equation (8). First consider BM unit k owned by firm i . Given demand parameters θ and marginal cost vector c_i define

$$g_{i,k}(M_t, s_{i,t}, b_{i,t}, c_i; \theta) = M_t s_{k,t} + \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_j) M_t (1(k = j) - s_{k,t}) \frac{\alpha s_{j,t}}{b_{k,t}},$$

Consider next all BM units $k \in \mathcal{J}_i$ owned by firm i and define g_i to be the corresponding vector-valued random function. Finally, define the vector-valued random function

$$\Psi_i^T(c_i) = \frac{1}{T - 43} \sum_{t=44}^T g_i(M_t, s_{i,t}, b_{i,t}, c_i; \hat{\theta}^T),$$

where we substitute in our estimate $\hat{\theta}^T$ and explicitly indicate the dependence of this estimate on the data through the superscript T . Note that $g_{i,k}(\cdot)$, $g_i(\cdot)$, and $\Psi_i^T(c_i)$ are linear in c_i . Our estimate \hat{c}_i^T solves the linear system of equations $\Psi_i^T(c_i) = 0$; the system has a unique solution provided all units \mathcal{J}_i get positive share during the sample period $44 \dots T$. Denote the true marginal cost as c_i^0 . Our goal is to show that $\text{plim}_{T \rightarrow \infty} \hat{c}_i^T = c_i^0$.

Turning to the data generating process in equation (7), given marginal cost c_i and the subjective probability measure $P^{i,t}(b_{-i,t}, \xi_t, e_t, \theta_t | \Omega_{t-1})$ underlying $\mathcal{E}_{b_{-i,t}, \xi_t, e_t, \theta_t}^{i,t}[\cdot | \Omega_{t-1}]$ define

$$h_{i,k}^{i,t}(b_{i,t}, c_i, \Omega_{t-1}) = \mathcal{E}_{b_{-i,t}, \xi_t, e_t, \theta_t}^{i,t}[g_{i,k}(b_t, x_t, \xi_t, e_t, M_t, c_i; \theta_t) | \Omega_{t-1}] = E_{P^{i,t}}[g_{i,k}(b_t, x_t, \xi_t, e_t, M_t, c_i; \theta_t) | \Omega_{t-1}].$$

Consider all BM units $k \in \mathcal{J}_i$ owned by firm i and define $h_i^{i,t}$ to be the corresponding vector-valued random function. Further define

$$\Phi_i^T(c_i) = \frac{1}{T - 43} \sum_{t=44}^T h_i^{i,t}(b_{i,t}, c_i, \Omega_{t-1}).$$

Note that $h_{i,k}^{i,t}(\cdot)$, $h_i^{i,t}(\cdot)$, and $\Phi_i^T(c_i)$ are linear in c_i . Because the true marginal cost c_i^0 is known to the firm and the firm optimizes by choosing its bids $b_{i,t}$ to satisfy its first-order conditions in equation (7), it must be that $h_i^{i,t}(b_{i,t}, c_i^0, \Omega_{t-1}) = 0$ and $\Phi_i^T(c_i^0) = 0$. Moreover,

$\Phi_i^T(c_i) = 0$ is a linear system of equations in c_i with unique solution whenever the (square) matrix of coefficients on c_i is non-singular. The first part of our identification assumption guarantees this: the Jacobian J^T is almost surely non-singular for every T , where the kj -th element of J^T is $\frac{\partial}{\partial c_j} \frac{1}{T-43} \sum_{t=44}^T h_{i,k}^{i,t}(b_{i,t}, c_i, \Omega_{t-1})$.

Now to prove $\text{plim}_{T \rightarrow \infty} c_i^T = c_i^0$, we first show that $\Psi_i^T(c_i)$ uniformly converges to $\Phi_i^T(c_i)$ in probability using a triangle inequality argument:

$$\begin{aligned} \|\Psi_i^T(c_i) - \Phi_i^T(c_i)\| &\leq \left\| \frac{1}{T-43} \sum_{t=44}^T g_i(b_t, x_t, \xi_t, e_t, M_t, c_i; \hat{\theta}^T) - g_i(b_t, x_t, \xi_t, e_t, M_t, c_i; \theta_0) \right\| \\ &\quad + \left\| \frac{1}{T-43} \sum_{t=44}^T g_i(b_t, x_t, \xi_t, e_t, M_t, c_i; \theta_0) - h_i^{0,t}(b_{i,t}, c_i, \Omega_{t-1}) \right\| \\ &\quad + \left\| \frac{1}{T-43} \sum_{t=44}^T h_i^{0,t}(b_{i,t}, c_i, \Omega_{t-1}) - h_i^{i,t}(b_{i,t}, c_i, \Omega_{t-1}) \right\|. \end{aligned}$$

where $h_i^{0,t}(b_{i,t}, c_i, \Omega_{t-1}) = E[g_i(b_t, x_t, \xi_t, e_t, M_t, c_i; \theta_t) | \Omega_{t-1}]$. We will assume that c_i belongs to a compact convex parameter set; that firm's perceptions about the price parameter α lie in some bounded interval $[\underline{\alpha}, 0]$, that α^0 is negative and finite; that M_t is bounded above by M ; and that all bids placed are bounded below by a strictly positive constant $\underline{b} > 0$. Then since $\text{plim}_{T \rightarrow \infty} \hat{\theta}^T = \theta^0$, the first term converges to zero by the continuous mapping theorem. The second term converges in probability to zero by a WLLN, since each term in the summation is a mean zero random variable, independent of the previous term because of the conditioning on Ω_{t-1} . The third term converges in probability to zero since we assume that $P^{i,t}(\cdot | \Omega_{t-1})$ converges in distribution to the truth, uniformly in Ω_{t-1} ; g_i is continuous and bounded; and so $h_i^{i,t} \equiv P^{i,t} g_i \rightarrow_p E[g_i] \equiv h_i^{0,t}$. Convergence of the sequence of individual terms implies convergence of the sequence of averages (Cesàro means). Taken together, we have that $\Psi_i^T(c_i)$ converges in probability pointwise to $\Phi_i^T(c_i)$.

Notice also that both Ψ_i^T and Φ_i^T are linear in c_i with Jacobians whose entries are globally bounded in absolute value by $\hat{\alpha}^T M / \underline{b}$ and $\underline{\alpha} M / \underline{b}$ respectively, and since $\hat{\alpha}^T \rightarrow_p \alpha^0$, both sequences of functions satisfy a global stochastic Lipschitz condition. This implies that the convergence is uniform in c_i (Newey 1991).

Lastly, we show that this convergence result is sufficient for $\text{plim}_{T \rightarrow \infty} c_i^T = c_i^0$. This is a little non-standard because the limiting system of equations $\Phi_i^T(c_i)$ is random. Our second identification assumption is that the set of equations $\Phi_i^T(c_i) = 0$, viewed as a random linear

system $A^T c_i = b^T$, has the property that the inverse of A^T denoted $A^{T,-1}$ has a uniformly bounded spectral norm almost surely. This implies that any approximate solution of the set of equations (i.e. $c_i = A^{T,-1} b$ for b in a neighborhood of b^T) must be close to c_i^0 (since $\|c_i - c_i^0\| \leq \|A^{T,-1}\| \|b - b^T\|$ for $\|A^{T,-1}\|$ the spectral norm of $A^{T,-1}$). This gives the separation condition we need: for all $\delta > 0$ $\inf_{\|c_i - c_i^0\| > \delta} \|\Phi_i^T(c_i)\| > \delta\epsilon$ with certainty, for some $\epsilon > 0$. And since c_i^0 is the unique solution of $\Phi_i^T(c_i) = 0$ and $\Psi_i^T(c_i)$ converges to $\Phi_i^T(c_i)$ uniformly in probability, for every $\delta > 0$ there exists $T(\delta)$ such that $T > T(\delta) \Rightarrow \|\Psi_i^T(c_i) - \Phi_i^T(c_i)\| < \delta\epsilon$ with probability one, so that the minimizer of $\|\Psi_i^T(c_i)\|$ must be within δ of c_i^0 with probability one, completing the consistency proof.

A.2 Data appendix

Data sources. Since a redesign on November 1, 2013, the data website of NG is available at <http://www2.nationalgrid.com/UK/Industry-information/Electricity-transmission-operations/Data-explorer/Outcome-Energy-Services/>. The data on the FR market is available under the tab “Frequency Response — FFR & Mandatory.” We downloaded our data from a previous version of the NG data website. In those cases detailed below where the 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 FFR 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 0.2Hz, 0.5Hz, and 0.8Hz, while for the market segments secondary and high, only response capabilities at 0.2Hz and 0.5Hz 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.

- **Quantities:** We obtained FR quantity data directly from the NG data website. Unfortunately, the new data website no longer provides historic quantities, and only a file that holds quantities from August 2013 is available. We downloaded monthly quantity files for November 2005 thru June 2013. Each of these files contains one month of daily holding quantities 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:** 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 BM 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.
- **Electricity demanded:** We take information on electricity demanded from NG at <http://www.nationalgrid.com/uk/Electricity/Data/Demand+Data/>. The data is stored in a sequence of excel spreadsheets, each of which has the quantity demanded on a given day.
- **Firm frequency response:** We obtain information on FFR from the reports published at <http://www2.nationalgrid.com/UK/Industry-information/Electricity-transmission>

Report-explorer/Services-Reports/. The data is stored in a sequence of excel spreadsheets published monthly, each of which has FFR volumes by day.

- **Fuel type:** We take fuel type information from appendix F1 of the Seven-Year Statement prepared by NG in 2011 at <http://nationalgrid.com/NR/rdonlyres/3B1B4AE4-2368-4B6E-8DA4-539A67EAD41F/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 Variable Pitch at <http://www.variablepitch.co.uk/grid/>.
- **Fuel prices:** The UK Department of Energy and Climate Change publishes quarterly and annual prices of fuels purchased by power generators and of gas at UK delivery points. A file titled “Average prices of fuels purchased by the major UK power producers and of gas at UK delivery points (QEP 3.2.1)” is available at <https://www.gov.uk/government/statistical-data-sets/prices-of-fuels-purchased-by-major-power-producers>. The sheet “Quarterly” contains the quarterly price of coal, oil, and gas, measured in pence per kilowatt hour (KWh), in columns D, F, and G.
- **Vintage:** We take fuel type information from appendix F1 of the Seven-Year Statement prepared by NG in 2011 at <http://nationalgrid.com/NR/rdonlyres/3B1B4AE4-2368-4B6E-8DA4-539A67EAD41F/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 power stations, so we take this information from the website of the British Hydropower Association at <http://www.british-hydro.org/>. For an additional eleven power 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 and variable construction. The unit of observation is BM unit by month. We consider the time period November 2005 to October 2011. We include BM units in the analysis if they provided positive FR quantity in at least one of these months.

We aggregate quantities for the three market segments primary, secondary, and high (see footnote 1) by summing daily quantities across segments and days. For BM unit j in month t we thus obtain FR quantity as

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

where k indexes market segments and τ days, and we abuse notation to denote as $\tau \in t$ the days in month t . The FR bids are constructed as quantity-weighted averages of segment-specific bids, where the weights are constant and given by the overall quantities 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 $Q_k = \sum_j \sum_t \sum_{\tau \in t} q_{k,\tau,j,t}$ and $Q = Q_P + Q_S + Q_H$.

Because a bid above £23/MWh is only accepted 12 times in our dataset of over 9000 observations, we label such a bid non-competitive; we otherwise label the bid competitive. One reason to opt out of the FR market by submitting a non-competitive bid is that the BM unit undergoes maintenance that month. Modelling maintenance and other reasons a BM unit opts out of the FR market is beyond the scope of this paper, and throughout we simply drop the corresponding observations. We also drop observations if the bid is missing.

Online Appendix for “Just starting out: Learning and equilibrium in a new market”

Ulrich Doraszelski

University of Pennsylvania

Gregory Lewis

Microsoft Research and NBER

Ariel Pakes

Harvard University and NBER

December 18, 2015

A.1 Selection

Selection on observables. To investigate selection on observables, we extend the probit model in equation (2) to include the log bid $\ln b_{j,t}$:

$$\Pr(e_{j,t} = 1 | b_{j,t}, x_{j,t}) = 1 - \Phi \left(-\check{\alpha} \ln b_{j,t} - \check{\beta} x_{j,t} - \check{\gamma}_j - \check{\mu}_t \right) = \Phi \left(\check{\alpha} \ln b_{j,t} + \check{\beta} x_{j,t} + \check{\gamma}_j + \check{\mu}_t \right). \quad (11)$$

Table 13 shows ML estimates. In the first column, we exclude the observable characteristics $x_{j,t}$; in the second column, we include them. In the third column, the bid enters more flexibly through a series of dummies for $b_{j,t}$ being in each decile of the distribution of bids. The coefficient on log bid $\ln b_{j,t}$ is statistically significant, as are half of the decile coefficients in the flexible specification. However, as noted in the main text, the impact of the log bid $\ln b_{j,t}$ is economically small.

Selection on unobservables. To examine selection on unobservables, we revert to the probit model in equation (2). We allow for correlation between $\nu_{j,t}$ and $\eta_{j,t}$ (and hence $\xi_{j,t}$ and $\eta_{j,t}$) and assume that they are iid across BM units and months and jointly normal distributed as

$$\begin{pmatrix} \nu_{j,t} \\ \eta_{j,t} \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2 & \lambda\sigma \\ \lambda\sigma & 1 \end{pmatrix} \right).$$

It follows that

$$\begin{aligned} & \mathbb{E}(\nu_{j,t} | e_{j,t} = e_{j,t-1} = 1, x_{j,t}) \\ &= \mathbb{E} \left(\nu_{j,t} | \eta_{j,t} > -\check{\beta} x_{j,t} - \check{\gamma}_j - \check{\mu}_t, \eta_{j,t-1} > -\check{\beta} x_{j,t-1} - \check{\gamma}_j - \check{\mu}_{t-1}, x_{j,t} \right) \\ &= \mathbb{E} \left(\nu_{j,t} | \eta_{j,t} > -\check{\beta} x_{j,t} - \check{\gamma}_j - \check{\mu}_t, x_{j,t} \right) \\ &= \lambda\sigma \frac{\phi \left(-\check{\beta} x_{j,t} - \check{\gamma}_j - \check{\mu}_t \right)}{1 - \Phi \left(-\check{\beta} x_{j,t} - \check{\gamma}_j - \check{\mu}_t \right)} = \lambda\sigma \frac{\phi \left(\check{\beta} x_{j,t} + \check{\gamma}_j + \check{\mu}_t \right)}{\Phi \left(\check{\beta} x_{j,t} + \check{\gamma}_j + \check{\mu}_t \right)}, \end{aligned}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal probability density function (PDF) and CDF. Hence, $\mathbb{E}(\nu_{j,t} | e_{j,t} = e_{j,t-1} = 1, x_{j,t}) \neq 0$ as long as $\lambda \neq 0$ and there is correlation between $\nu_{j,t}$

and $\eta_{j,t}$.

Estimating equation (4) requires adding an inverse Mills ratio selection correction (Heckman 1979). Table 14 shows the resulting NLLS estimates. The coefficient on the inverse Mills ratio is significant but the remaining coefficients are very similar to our leading estimates in Table 3.

Table 13: Determinants of positive volume

Indicator for positive share			
Log bid	-0.302*	-0.526***	
	(0.165)	(0.203)	
Fully loaded		2.604***	2.591***
		(0.365)	(0.349)
Part loaded		2.277***	2.436***
		(0.344)	(0.300)
Positive FFR volume		-0.581	-0.527
		(0.481)	(0.451)
Bid decile 2			-0.003
			(0.360)
Bid decile 3			0.442
			(0.314)
Bid decile 4			-0.430
			(0.360)
Bid decile 5			-0.699**
			(0.326)
Bid decile 6			-0.959***
			(0.335)
Bid decile 7			-0.729**
			(0.356)
Bid decile 8			-0.693**
			(0.341)
Bid decile 9			-0.443
			(0.317)
Bid decile 10			-0.866***
			(0.320)
Unit and Month FE	yes	yes	yes
Flexible bid controls	no	no	yes
N	5175	5175	5175

An observation is a unit-month, and the dependent variable is an indicator for a unit having positive volume. Inactive units are omitted. The regressors are the bid (either in logs, or with indicators for the bid being in decile bins), the fraction of time the unit's final physical notification is that it is fully loaded (i.e. operating at capacity) and part loaded (i.e. operating below capacity); and whether the unit is under a firm frequency response contract. Standard errors are clustered by bmunit. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

Table 14: Demand System Estimates

	Log share ratio QFD
Log bid	-1.649*** (0.117)
Fully loaded	1.580*** (0.226)
Part loaded	1.927*** (0.185)
Positive FFR volume	-0.573** (0.246)
Millsratio	-0.517*** (0.182)
Unit and Month FE	yes
ρ	0.40
s.e. ρ	0.03
N	3509

The dependent variable is the log ratio of the unit share to the outside good share (an observation is a unit-month), coded as missing where the share is zero and omitted in estimation. 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). Standard errors are clustered by bmunit. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

A.2 Repositioning in the BM

In this section, we account for the profit that accrues to a BM unit as it is repositioned in the BM in preparation for providing FR. The BM is a multi-unit discriminatory auction that is held every half-hour. Prior to this auction, a BM unit submits its contracted position to NG along with its bid. A bid in the BM is essentially a supply curve that is centered at the BM unit's contracted position. This supply curve is described by price-quantity pairs through which the BM unit can offer to increase its energy production in up to five increments above its contracted position. If NG accepts an offer, the BM unit is paid by NG accordingly. The supply curve is further described by up to five price-quantity pairs through which the BM unit can bid to decrease its energy production below its contracted position. If NG accepts a bid, the BM unit pays NG accordingly.

The BM in other countries has been studied in great detail by Borenstein, Bushnell and Wolak (2002), Wolak (2003, 2007), Sweeting (2007), and Hortaçsu and Puller (2008). In line with our focus on the FR market, we work with a much simpler model of the BM that is designed to merely give us a sense of the profit that accrues to a BM unit as it is repositioned in the BM and how that profit changes with its bid for providing FR. We proceed in two steps. First, we estimate a demand model for repositioning. To account for the interdependency between the BM and the FR market, we include the bid for providing FR in the demand model. Second, to obtain profit, we estimate the markup in the BM jointly with the cost of providing FR.

Data. For every BM unit we have data on bids and offers (up to ten price-quantity pairs), contracted position, and actual position every half-hour. The quantity of upward repositioning $q_{j,\tau}^+$ of BM unit j in half-hour τ effected through the BM is therefore the larger of zero and the difference between actual and contracted position; the quantity of downward repositioning $q_{j,\tau}^-$ is the larger of zero and the difference between contracted and actual position. Market size $M_\tau^+ = \sum_j q_{j,\tau}^+$ and $M_\tau^- = \sum_j q_{j,\tau}^-$ is the total amount of upward, respectively, downward repositioning in half-hour τ .

We face two problems with the data. First, if BM unit j is not repositioned up or down in the BM in half-hour τ , then $q_{j,\tau}^+ = 0$, respectively, $q_{j,\tau}^- = 0$. This happens quite frequently, and we account for it in our demand model. Second, the bids and offers can take on extreme values. This sometimes happens even though the BM unit is repositioned so that $q_{j,\tau}^+ > 0$

or $q_{j,\tau}^- > 0$. Hence, taken at face value, the bids and offers imply an implausibly huge profit. We deal with this by directly estimating the markup rather than marginal cost in the BM.

The only place in which the offers are used in what follows is to construct a grid of 24 prices for upward repositioning as follows: Pooling across all BM units and half-hours, we consider the distribution of offers and take the 4th through 96th percentiles. We proceed analogously to fix a grid of 24 prices for downward repositioning.

Demand. As with the FR market, the “inside goods” are the $J = 72$ BM units owned by the ten largest firms in Table 1 and the “outside good” encompasses the remaining BM units. To simplify the exposition, we focus on the demand for upward repositioning. The demand for downward repositioning is analogous.

Let $s_{j,\tau}^+$ denote the market share of upward repositioning of BM unit j in half-hour τ and $s_{0,\tau}^+ = 1 - \sum_j s_{j,\tau}^+$ the market share of the outside good. Let $e_{j,\tau}^+ = 1(s_{j,\tau}^+ > 0)$ be the indicator for BM unit j being eligible for repositioning in the BM — and thus having a positive market share — in half-hour τ . Accounting for eligibility, we use a logit model for the market share of BM unit j in half-hour τ with

$$s_{j,\tau}^+ = \frac{e_{j,\tau}^+ \exp(\alpha^+ \ln b_{j,t} + \beta^+ x_{j,\tau}^+ + \gamma_j^+ + \xi_{j,\tau}^+)}{1 + \sum_k e_{k,\tau}^+ \exp(\alpha^+ \ln b_{k,t} + \beta^+ x_{k,\tau}^+ + \gamma_k^+ + \xi_{k,\tau}^+)}. \quad (12)$$

γ_j^+ is a BM-unit fixed effect. $b_{j,t}$ is the bid for providing FR of BM unit j in the month t to which half-hour τ belongs. $x_{j,\tau}^+$ are controls that parsimoniously represent the supply curves that the BM units bid in the BM. We include in $x_{j,\tau}^+$ the hypothetical market share of BM unit j in half-hour τ at each of the 24 prices in the grid for upward repositing.¹ Finally, $\xi_{j,\tau}^+$ is a disturbance that, we assume, is mean independent of $b_{j,t}$ and $x_{j,\tau}^+$. This rules out that a firm conditions its bid in the BM on $\xi_{j,\tau}^+$.

We use a probit model for BM unit j being eligible for repositioning in the BM in half-hour τ with

$$e_{j,\tau}^+ = 1(\check{\alpha}^+ \ln b_{j,t} + \check{\beta}^+ \check{x}_{j,\tau}^+ + \check{\gamma}_j^+ + \eta_{j,\tau}^+ > 0).$$

$\check{\gamma}_j^+$ is a BM-unit fixed effect. $b_{j,t}$ is the bid for providing FR of BM unit j in the month t to

¹From its supply curve we can infer a hypothetical quantity of upward repositioning for BM unit j in half-hour τ at any given price. We compute the hypothetical market share of BM unit j in half-hour τ from the hypothetical quantities of all BM units, irrespective of whether they are part of the inside or outside goods.

which half-hour τ belongs. $\check{x}_{j,\tau}^+$ contains additional half-hour-of-day (same for each day), day-of-week (same for each week), week-of-year (same for each year), and year fixed effects and controls that parsimoniously represent the supply curves that the BM units bid in the BM. We include in $\check{x}_{j,\tau}^+$ the lowest offer of BM unit j in half-hour τ along with the corresponding quantity. Next we compute the distribution of lowest offers of all BM units (irrespective of whether they are part of the inside or outside goods) in half-hour τ . We include in $\check{x}_{j,\tau}^+$ ten dummies for the decile in which the lowest offer of BM unit j in half-hour τ falls. We proceed similarly for the quantity corresponding to the lowest offer and include in $\check{x}_{j,\tau}^+$ another ten dummies for the decile in which the quantity corresponding to the lowest offer of BM unit j in half-hour τ falls. Finally, $\eta_{j,\tau}^+ \sim N(0, 1)$ is a standard normally distributed disturbance that, we assume, is mean independent of $b_{j,t}$ and $\check{x}_{j,\tau}^+$ and independent across BM units and half-hours.

It follows that

$$\Pr(e_{j,\tau}^+ = 1 | b_{j,t}, \check{x}_{j,\tau}^+) = 1 - \Phi \left(-\check{\alpha}^+ \ln b_{j,t} - \check{\beta}^+ \check{x}_{j,\tau}^+ - \check{\gamma}_j^+ \right) = \Phi \left(\check{\alpha}^+ \ln b_{j,t} + \check{\beta}^+ \check{x}_{j,\tau}^+ + \check{\gamma}_j^+ \right), \quad (13)$$

where $\Phi(\cdot)$ is the standard normal CDF. We estimate equation (13) by ML. Moreover, equation (12) implies

$$\ln s_{j,\tau}^+ - \ln s_{0,\tau}^+ \equiv \delta_{j,\tau}^+ = \alpha^+ \ln b_{j,t} + \beta^+ x_{j,\tau}^+ + \gamma_j^+ + \xi_{j,\tau}^+$$

as long as $e_{j,\tau}^+ = 1$. We assume $\xi_{j,\tau}^+$ and $\eta_{j,\tau}^+$ are independent of each other and estimate by OLS.

Results. Tables 16 and 17 show our estimates for the logit model in equation (12) and the probit model in equation (13). In the first and third columns, we exclude the controls $\check{x}_{j,\tau}^+$ and $x_{j,\tau}^+$; in the second and fourth columns, we include them. The number of observations differs because we require $s_{j,t} > 0$ for OLS.

The coefficient on log FR bid $\ln b_{j,t}$ is significantly different from zero and negative in the logit model in equation (12) and the probit model in equation (13), both for upward and downward repositioning. This indicates that a BM unit that submits a low FR bid is more likely to be repositioned in the BM and also by larger amounts, presumably so that it can provide FR services. However, the impact is economically small. For example, in the logit

Table 16: Repositioning Share Analysis

	Log relative share of repositioning(if positive)			
	Upward repositions		Downward repositions	
logFRbid	-0.086***	-0.108***	-0.076***	-0.100***
N	260482	260482	885659	885659
R^2	0.57	0.58	0.53	0.54

In the first pair of regressions, the dependent variable is the log share ratio of upward repositioning volume; in the next two columns, it is the corresponding log share ratio of downward repositioning volume. Controls for the share of volume that a uniform auction would assign this unit based on its offers (upward) and bids (downward) at a set of 24 increasing prices are included in the second and fourth columns, although their coefficients are omitted. Unit fixed effects are included in all specifications. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$), and standard errors are clustered by half-hour periods (an observation is a unit-half-hour).

Table 17: Repositioning Probit Analysis

	Probability of Repositioning			
	Upward repositions		Downward repositions	
logFRbid	-0.05789***	-0.04886***	-0.20162***	-0.10188***
Closest bid/offer price		-0.00000*		0.00000***
Closest bid/offer quantity		-0.00011***		0.00162***
N	1511736	1511735	1511740	1508635

Estimates from a 20% random sample of observations. In the first pair of regressions, the dependent variable is the indicator variable of whether a unit get repositioned upward; in the next two columns, it is the corresponding indicator for downward repositions. Controls for the bid/offer closest to the current contracted position are shown; dummies for the percentile of the bid and offer (relative to contemporaneous offers) are included in the second and fourth columns but suppressed. Month-of-year, day-of-week and hour-of-day dummies fixed effects are included in all specifications. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

model in equation (12), the elasticity of market share with respect to FR bid is on the order of -0.1, compared to around -1.6 in the FR market.

Markup and profit. To simplify the exposition, we again focus on upward repositioning. Conditional on eligibility (or in realization), the market share of BM j in half-hour τ is $s_j^+(b_t, x_\tau^+, \xi_\tau^+, e_\tau^+; \theta^+)$, as defined on the right-hand side of equation (12). We use the shorthands $x_\tau^+ = (x_{j,\tau}^+)_{j=1,\dots,J}$, $\xi_\tau^+ = (\xi_{j,\tau}^+)_{j=1,\dots,J}$, and $e_\tau^+ = (e_{j,\tau}^+)_{j=1,\dots,J}$. θ^+ denotes the parameters of the logit model in equation (12). Unconditionally (or in expectation), the

market share of BM j in half-hour τ is

$$s_j^+(b_t, x_\tau^+, \xi_\tau^+, \check{x}_\tau^+; \theta^+, \check{\theta}^+) = \sum_{e_\tau^+ \in \{0,1\}^J} s_j^+(b_t, x_\tau^+, \xi_\tau^+, e_\tau^+; \theta^+) w^+(b_t, \check{x}_\tau^+, e_\tau^+; \check{\theta}^+),$$

where

$$w^+(b_t, \check{x}_\tau^+, e_\tau^+; \check{\theta}^+) \equiv \prod_{l=1, \dots, J} \Phi \left(\check{\alpha}^+ \ln b_{l,t} + \check{\beta}^+ \check{x}_{l,\tau}^+ + \check{\gamma}_l^+ \right)^{e_{l,\tau}^+} \left(1 - \Phi \left(\check{\alpha}^+ \ln b_{l,t} + \check{\beta}^+ \check{x}_{l,\tau}^+ + \check{\gamma}_l^+ \right) \right)^{1-e_{l,\tau}^+} \quad (14)$$

and the summation is over all 2^J possible values of e_τ^+ . $\check{\theta}^+$ denotes the parameters of the probit model in equation (13).

We assume that the profit that accrues to BM unit j as it is repositioned in the BM over the course of month t (again unconditionally or in expectation) can be written as

$$\mu_j \sum_{\tau \in t} \left(M_\tau^+ s_j^+(b_t, x_\tau^+, \xi_\tau^+, \check{x}_\tau^+; \theta^+, \check{\theta}^+) + M_\tau^- s_j^-(b_t, x_\tau^-, \xi_\tau^-, \check{x}_\tau^-; \theta^-, \check{\theta}^-) \right),$$

where we abuse notation to denote as $\tau \in t$ the half-hours in month t . μ_j is a common markup for upward and downward repositioning. If NG accepts an offer to increase energy production, then the BM unit is paid by NG according to its offer but bears the cost of the additional fuel. If NG accepts a bid to decrease energy production, then the BM unit pays NG according to its bid but saves on fuel cost. Because bids and offers are under the control of the firm owning the BM unit, we expect the markup to be nonnegative.

Recalling that \mathcal{J}_i denotes the indices of the BM units that are owned by firm i , the profit of firm i in the BM over the course of month t (again unconditionally or in expectation) is

$$\sum_{j \in \mathcal{J}_i} \mu_j \sum_{\tau \in t} \left(M_\tau^+ s_j^+(b_t, x_\tau^+, \xi_\tau^+, \check{x}_\tau^+; \theta^+, \check{\theta}^+) + M_\tau^- s_j^-(b_t, x_\tau^-, \xi_\tau^-, \check{x}_\tau^-; \theta^-, \check{\theta}^-) \right).$$

We are interested in how this profit changes with the bid for providing FR. Recall that the bid for the current month is submitted before the 20th of the previous month while bidding in the BM takes place during the current month. We simplify and assume that in preparing its bid for providing FR a firm ignores $\frac{\partial x_\tau^+}{\partial b_{j,t}}$, $\frac{\partial \check{x}_\tau^+}{\partial b_{j,t}}$, $\frac{\partial x_\tau^-}{\partial b_{j,t}}$, and $\frac{\partial \check{x}_\tau^-}{\partial b_{j,t}}$ for all $\tau \in t$. In essence, this says that the firm ignores that through its bid for providing FR it can influence the

competitive landscape for the subsequent bidding in the BM. Under some conditions the envelope theorem ensures that this assumption is satisfied with respect to the bids and offers for the BM units that are owned by the firm. We emphasize, however, that this assumption has bite with respect to the bids and offers for the BM units that are owned by the firm's rivals.

It remains to compute $\frac{\partial s_j^+(b_t, x_\tau^+, \xi_\tau^+, \check{x}_\tau^+; \theta^+, \check{\theta}^+)}{\partial b_{k,t}}$ and $\frac{\partial s_j^-(b_t, x_\tau^-, \xi_\tau^-, \check{x}_\tau^-; \theta^-, \check{\theta}^-)}{\partial b_{k,t}}$. We have

$$\begin{aligned} \frac{\partial s_j^+(b_t, x_\tau^+, \xi_\tau^+, \check{x}_\tau^+; \theta^+, \check{\theta}^+)}{\partial b_{j,t}} &= \sum_{e_\tau^+ \in \{0,1\}^J} \left(s_j^+(b_t, x_\tau^+, \xi_\tau^+, e_\tau^+; \theta^+) (1 - s_j^+(b_t, x_\tau^+, \xi_\tau^+, e_\tau^+; \theta^+)) \frac{\alpha^+}{b_{j,t}} \right. \\ &\quad \left. + s_j^+(b_t, x_\tau^+, \xi_\tau^+, e_\tau^+; \theta^+) \frac{\check{\alpha}^+ \phi \left(\check{\alpha}^+ \ln b_{j,t} + \check{\beta}^+ \check{x}_{j,\tau}^+ + \check{\gamma}_j^+ \right)}{b_{j,t} \left(\Phi \left(\check{\alpha}^+ \ln b_{j,t} + \check{\beta}^+ \check{x}_{j,\tau}^+ + \check{\gamma}_j^+ \right) + e_{j,\tau}^+ - 1 \right)} \right) w^+(b_t, \check{x}_\tau^+, e_\tau^+; \check{\theta}^+) \end{aligned}$$

for $k = j$ and

$$\begin{aligned} \frac{\partial s_j^+(b_t, x_\tau^+, \xi_\tau^+, \check{x}_\tau^+; \theta^+, \check{\theta}^+)}{\partial b_{k,t}} &= \sum_{e_\tau^+ \in \{0,1\}^J} \left(-s_j^+(b_t, x_\tau^+, \xi_\tau^+, e_\tau^+; \theta^+) s_k^+(b_t, x_\tau^+, \xi_\tau^+, e_\tau^+; \theta^+) \frac{\alpha^+}{b_{k,t}} \right. \\ &\quad \left. + s_j^+(b_t, x_\tau^+, \xi_\tau^+, e_\tau^+; \theta^+) \frac{\check{\alpha}^+ \phi \left(\check{\alpha}^+ \ln b_{k,t} + \check{\beta}^+ \check{x}_{k,\tau}^+ + \check{\gamma}_k^+ \right)}{b_{k,t} \left(\Phi \left(\check{\alpha}^+ \ln b_{k,t} + \check{\beta}^+ \check{x}_{k,\tau}^+ + \check{\gamma}_k^+ \right) + e_{k,\tau}^+ - 1 \right)} \right) w^+(b_t, \check{x}_\tau^+, e_\tau^+; \check{\theta}^+) \end{aligned}$$

for $k \neq j$. Note that these derivatives are themselves expectations over eligibility e_τ^+ using probability weights $w^+(b_t, \check{x}_\tau^+, e_\tau^+; \check{\theta}^+)$.

To jointly estimate the marginal cost of providing FR and the markup on repositioning operations, we adjust the estimation equation (8) as follows: When we substitute in realizations and parameter estimates, then the bids $b_{i,t}$ of firm i in month $t \geq 44$ during the late phase

satisfy the system of equations²

$$\begin{aligned}
& \frac{1}{29} \sum_{t=44}^{72} \left[\left(M_t s_k(b_t, x_t, \xi_t, e_t; \theta) + \sum_{j \in \mathcal{J}_i} (b_{j,t} - c_j) M_t s_j(b_t, x_t, \xi_t, e_t; \theta) (1(k=j) - s_k(b_t, x_t, \xi_t, e_t; \theta)) \frac{\alpha}{b_{k,t}} \right. \right. \\
& \quad + \sum_{j \in \mathcal{J}_i} \mu_j \sum_{\tau \in t} \left(M_\tau^+ \left(s_j^+(b_t, x_\tau^+, \xi_\tau^+, e_\tau^+; \theta^+) (1(k=j) - s_k^+(b_t, x_\tau^+, \xi_\tau^+, e_\tau^+; \theta^+)) \frac{\alpha^+}{b_{k,t}} \right. \right. \\
& \quad \quad \left. \left. + s_j^+(b_t, x_\tau^+, \xi_\tau^+, e_\tau^+; \theta^+) \frac{\check{\alpha}^+ \phi \left(\check{\alpha}^+ \ln b_{k,t} + \check{\beta}^+ \check{x}_{k,\tau}^+ + \check{\gamma}_k^+ \right)}{b_{k,t} \left(\Phi \left(\check{\alpha}^+ \ln b_{k,t} + \check{\beta}^+ \check{x}_{k,\tau}^+ + \check{\gamma}_k^+ \right) + e_{k,\tau}^+ - 1 \right)} \right) \right. \\
& \quad \quad \left. + M_\tau^- \left(s_j^-(b_t, x_\tau^-, \xi_\tau^-, e_\tau^-; \theta^-) (1(k=j) - s_k^-(b_t, x_\tau^-, \xi_\tau^-, e_\tau^-; \theta^-)) \frac{\alpha^-}{b_{k,t}} \right. \right. \\
& \quad \quad \left. \left. + s_j^-(b_t, x_\tau^-, \xi_\tau^-, e_\tau^-; \theta^-) \frac{\check{\alpha}^- \phi \left(\check{\alpha}^- \ln b_{k,t} + \check{\beta}^- \check{x}_{k,\tau}^- + \check{\gamma}_k^- \right)}{b_{k,t} \left(\Phi \left(\check{\alpha}^- \ln b_{k,t} + \check{\beta}^- \check{x}_{k,\tau}^- + \check{\gamma}_k^- \right) + e_{k,\tau}^- - 1 \right)} \right) \right) \right) \otimes (1, f_{k,t-1}) \Big] = 0, \quad \forall k \in \mathcal{J}_i,
\end{aligned}$$

where \otimes denotes the Kronecker product, 1 the constant, and $f_{k,t-1}$ the fuel price relevant for BM unit k in month $t-1$. We omit distinguishing between parameters and estimates to simplify the notation.

These $2|\mathcal{J}_i|$ equations not only require that the first-order conditions are on average correct in the late phase but also that they are uncorrelated with the lagged fuel price that is known to the firm at the time it prepares its current FR bid. To facilitate the estimation, we assume the markup is common across BM units and firms and solve the resulting overdetermined system of linear equations by OLS.

Results. Accounting for repositioning incentives has a relatively small impact on the estimated marginal costs of providing FR: as Table 18 shows, the average across BM units falls from £1.41/MWh to £1.36/MWh. The estimated markup is not significantly different from zero.

²We make the simplifying assumption that the firm has perfect foresight about $M_t^+ = (M_\tau^+)_{\tau \in t}$, $M_t^- = (M_\tau^-)_{\tau \in t}$, $x_t^+ = (x_\tau^+)_{\tau \in t}$, $\check{x}_t^+ = (\check{x}_\tau^+)_{\tau \in t}$, $x_t^- = (x_\tau^-)_{\tau \in t}$, and $\check{x}_t^- = (\check{x}_\tau^-)_{\tau \in t}$.

Table 18: Cost Estimates with Repositioning Incentives

	Without repositioning	With repositioning
Average marginal cost	1.40	1.36
Main market markup	–	-0.0014
s.e. markup	–	0.0041

Cost and markup estimates, with and without accounting for repositioning. Estimation is by generalized method of moments. The first column estimates are the same as those discussed in the main text.