Online appendix for "Just starting out: Learning and equilibrium in a new market"

Ulrich Doraszelski
University of Pennsylvania, CEPR, and NBER

Gregory Lewis
Microsoft Research and NBER

Ariel Pakes
Harvard University and NBER

September 26, 2017

A.3 Sufficient condition for belief convergence

We show that if our estimate of θ is consistent for that parameter and if the firm's subjective probability distribution converges weakly to the objective probability distribution (uniformly across information sets), then

$$\left\| T^{-1} \sum_{t=44}^{T} \left(h_{i,t}^{T}(c_i) - h_{i,t}^{e}(c_i) \right) \right\| = o_p(1),$$

where the notation follows the discussion of consistency in the appendix of the paper. This is one of the sufficient conditions for obtaining consistent estimates of c_i (the other one being an identification condition).

Let firm i have a subjective probability measure $P^{i,t}(b_{-i,t}, \xi_t, e_t, \theta_t \big| \Omega_{i,t-1})$ underlying $\mathcal{E}_{b_{-i,t},\xi_t,e_t,\theta_t}[\cdot|\Omega_{i,t-1}]$, so that $\mathcal{E}_{b_{-i,t},\xi_t,e_t,\theta_t}[f(u)|\Omega_{i,t-1}] = \int f(u) dP^{i,t}(u|\Omega_{i,t-1})$. Let α_0 denote the true price parameter. By the triangle inequality

$$\begin{aligned}
\left\| T^{-1} \sum_{t=44}^{T} \left(h_{i,t}^{T}(c_i) - h_{i,t}^{e}(c_i) \right) \right\| \leq & \left\| T^{-1} \sum_{t=44}^{T} \left(h(c_i, \hat{\alpha}_T, y_{i,t}) - h(c_i, \alpha_0, y_{i,t}) \right) \right\| \\
& + \left\| T^{-1} \sum_{t=44}^{T} \left(h(c_i, \alpha_0, y_{i,t}) - h_{i,t}^{0}(c_i) \right) \right\| \\
& + \left\| T^{-1} \sum_{t=44}^{T} \left(h_{i,t}^{0}(c_i) - h_{i,t}^{e}(c_i) \right) \right\|,
\end{aligned}$$

where $h_{i,t}^{0}(c_i) = E[h(c_i, \alpha, y_{i,t})|\Omega_{i,t-1}]$ and $E[\cdot|\Omega_{i,t-1}]$ denotes the expectation with respect to the objective probability measure (which puts point mass on $\alpha = \alpha_0$). By assumption, $\text{plim}_{T\to\infty} \hat{\alpha}_T = \alpha_0$, so 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 $\Omega_{i,t-1}$. The third term converges in probability to zero since $P^{i,t}(\cdot|\Omega_{i,t-1})$ weakly converges to the objective probability measure, uniformly in $\Omega_{i,t-1}$, h is continuous and bounded, and so $h_{i,t}^{e}(c_i) \equiv \int h(c_i, \alpha, y_{i,t}) dP^{i,t}(\alpha, y_{i,t}|\Omega_{i,t-1}) \to_p E[h(c_i, \alpha, y_{i,t})|\Omega_{i,t-1}] = h_{i,t}^{0}(c_i)$. Convergence of the sequence of individual terms implies convergence of the sequence of averages. Then, since the right-hand side converges in probability to zero, so does the left-hand side.

A.4 Selection

Selection on observables: persistence in eligibility. We extend the probit model in equation (2) to include lagged eligibility $e_{j,t-1}$:

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

The first column of Table 14 shows ML estimates for this model. There is statistically significant and economically meaningful evidence of persistence in eligibility. Footnote 15 in the main text explains why we decided not to model this in the main paper.

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

$$\Pr(e_{j,t}=1|b_{j,t},x_{j,t}) = 1-\Phi\left(-\breve{\alpha}\ln b_{j,t}-\breve{\beta}x_{j,t}-\breve{\gamma}_j-\breve{\mu}_t\right) = \Phi\left(\breve{\alpha}\ln b_{j,t}+\breve{\beta}x_{j,t}+\breve{\gamma}_j+\breve{\mu}_t\right). \tag{12}$$

The second and third columns of Table 14 show ML estimates. In the third column, we allow the bid to enter 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 footnote 15 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

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

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal probability density function (PDF) and CDF. Hence, $\mathrm{E}\left(\nu_{j,t}|e_{j,t}=e_{j,t-1}=1,x_{j,t}\right)\neq 0$ 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 15 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 14: Demand estimates with selection on observables.

	Eligibility		
Lagged eligibility	1.708		
	(0.074)		
Log bid		-0.526	
		(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	-0.009	-0.581	-0.527
	(0.233)	(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)
N obs	5099	5175	5175

ML estimates of probit model for eligibility with various controls for selection on observables. All models include BM-unit and month fixed effects. The unit of observation is a BM unit-month. Standard errors are clustered by BM unit.

Table 15: Demand estimates with selection on unobservables

	Market share
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)
Mills ratio	-0.517
	(0.182)
Autocorrelation coefficient	0.397
	(0.031)
N obs	3509

NLLS estimates of logit model for market share with Mills ratio to control for selection on unobservables. The model includes BM-unit and month fixed effects. NLLS estimates allow the unobservable characteristic $\psi_{j,t}$ to follow an $AR(1)$ process with autocorrelation coefficient ρ. The unit of observation is a BM unit-month. Standard errors are clustered by BM unit.

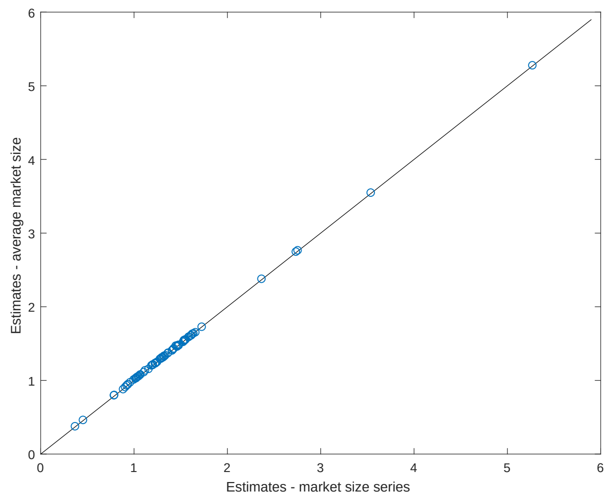

Figure 17: Comparison of cost estimates with market size series versus market size average.

A.5 Market size

As previously mentioned in Section 4.2, if a firm does not have perfect foresight about market size, M_t cannot be canceled out of equations (7) and (8). In our cost estimation, this implies up-weighting months with high values of M_t and down-weighting months with low values of M_t. As a robustness check, we repeat the cost estimation, canceling M_t from the first order condition. The results are nearly the same, as Figure 17 illustrates.

A.6 Fuel price

We model the marginal cost $c_{j,t}$ of BM unit j in month t as $c_{j,t} = c_j + \mu f_{j,t}$, where c_j is a BM-unit fixed, $f_{j,t}$ is the fuel price that the BM unit faces, and μ is a parameter.

Estimation. To estimate the $J+1$ parameters $c = (c_j)_{j=1,\ldots,J}$ and μ, we replace equation (8) by

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

and we add the equation

$$\frac{1}{29J}\sum_{k=1}^{K}\sum_{t=44}^{T=72}\left[\left[M_t s_{k,t} + \sum_{j\in\mathcal{J}_i}(b_{j,t}-c_{j,t})\,M_t\,(1(k=j)-s_{k,t})\,\frac{\hat{\alpha}s_{j,t}}{b_{k,t}}\right]f_{k,t-1}\right] = 0, \tag{13}$$

where $f_{k,t-1}$ is the fuel price relevant for BM unit k in month $t-1$. These $J+1$ equations are linear in the $J+1$ unknowns.

Results. We estimate μ to be -0.0137 with a standard error of 0.0040. This is economically small: on average across BM units, marginal cost decreases from £1.44/MWh to £1.35/MWh over the final phase of the FR market.

To probe this estimate, we re-specify $c_{j,t} = c_j + \lambda t$, where t is a time trend that is common across BM units. To estimate, we replace equation (13) by

$$\frac{1}{29J}\sum_{k=1}^{K}\sum_{t=44}^{T=72}\left[\left[M_t s_{k,t} + \sum_{j\in\mathcal{J}_i}(b_{j,t}-c_{j,t})\,M_t\,(1(k=j)-s_{k,t})\,\frac{\hat{\alpha}s_{j,t}}{b_{k,t}}\right]t\right] = 0. \tag{14}$$

We estimate λ to be -0.00295 with a standard error of 0.0012. We finally re-specify $c_{j,t} = c_j + \mu f_{j,t} + \lambda t$. To estimate, we use equations (13) and (14). We estimate μ to be -0.01928 wit a standard error of 0.0104 and λ to be 0.001550 with a standard error of 0.0030. Hence, neither coefficient is statistically significant. We conclude that the impact of fuel price is indistinguishable from a downward time trend in cost.

Table 16: Middle phase with time-varying marginal cost: Average absolute prediction error $CE^{(\delta,y)}$

	Single-period prediction					Multi-period prediction	
	$A(\emptyset)$	$A(\alpha)$	$A(\alpha,\mu)$	$A(\alpha,\beta)$	$A(\alpha,\beta,\mu)$	$A(\emptyset)$	$A(\alpha)$
F(0)	0.30	0.26	1.35	0.56	0.98	0.44	0.29
F(0.5)	0.30	0.27	1.37	0.57	0.99	0.42	0.28
F(1)	0.48	0.40	1.10	0.42	0.73	0.55	0.35
Eq.	0.53	0.42	1.87	0.63	1.29		

Computed from bids predicted by fictitious play model $F(\delta)$ or complete information Nash equilibrium (rows) in combination with adaptive learning model $A(y)$ (columns).

Table 17: Middle phase with time-varying marginal cost: Mean square prediction error $MSE^{(\delta,y)}$

	Single-period prediction					Multi-period prediction	
	$A(\emptyset)$	$A(\alpha)$	$A(\alpha,\mu)$	$A(\alpha,\beta)$	$A(\alpha,\beta,\mu)$	$A(\emptyset)$	$A(\alpha)$
F(0)	1.23	1.20	3.91	1.84	2.79	1.31	1.19
F(0.5)	1.23	1.20	4.00	1.88	2.84	1.28	1.17
F(1)	1.43	1.34	2.89	1.48	2.04	1.47	1.27
Eq.	1.45	1.32	6.35	1.94	3.86		

Computed from bids predicted by fictitious play model $F(\delta)$ or complete information Nash equilibrium (rows) in combination with adaptive learning model $A(y)$ (columns).

.

Learning models. As a robustness check, we re-ran our analysis of learning models in Section 5 under the assumption that the marginal cost $c_{j,t}$ of BM unit j in month t is $c_{j,t} = c_j + \mu f_{j,t}$, as specified and estimated using the procedure outlined above.

Tables 16, 17, 18, and 19 correspond to 8, 9, 11, and 12 in the main text. The broad conclusions are robust to allowing for time-varying marginal cost: the best fitting models remain those with $A(\emptyset)$ or $A(\alpha)$ and $F(0)$ or $F(0.5)$, and these fit substantially better in the middle phase of the FR market and only slightly better in the late phase.

A.7 Repositioning in the BM

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

Table 18: Late phase with time-varying marginal cost: Average absolute prediction error $CE^{(\delta,y)}$

	Single-period prediction					Multi-period prediction	
	$A(\emptyset)$	$A(\alpha)$	$A(\alpha,\mu)$	$A(\alpha,\beta)$	$A(\alpha,\beta,\mu)$	$A(\emptyset)$	$A(\alpha)$
F(0)	0.15	0.15	0.24	0.19	0.18	0.16	0.27
F(0.5)	0.16	0.16	0.25	0.19	0.18	0.16	0.27
F(1)	0.24	0.24	0.34	0.27	0.26	0.17	0.29
Eq.	0.17	0.17	0.28	0.21	0.19		

Computed from bids predicted by fictitious play model $F(\delta)$ or complete information Nash equilibrium (rows) in combination with adaptive learning model $A(y)$ (columns).

Table 19: Late phase with time-varying marginal cost: Mean square prediction error $MSE^{(\delta,y)}$

	Single-period Prediction					Multi-Period Prediction	
	$A(\emptyset)$	$A(\alpha)$	$A(\alpha,\mu)$	$A(\alpha,\beta)$	$A(\alpha,\beta,\mu)$	$A(\emptyset)$	$A(\alpha)$
F(0)	0.32	0.32	0.41	0.38	0.39	0.31	0.40
F(0.5)	0.32	0.32	0.42	0.38	0.40	0.31	0.39
F(1)	0.40	0.40	0.54	0.48	0.50	0.33	0.42
Eq.	0.31	0.31	0.44	0.38	0.38		

Computed from bids predicted by fictitious play model $F(\delta)$ or complete information Nash equilibrium (rows) in combination with adaptive learning model $A(y)$ (columns).

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

$$s^+_{j,\tau} = \frac{e^+_{j,\tau} \exp\left(\alpha^+ \ln b_{j,t} + \beta^+ x^+_{j,\tau} + \gamma^+_j + \mu^+_t + \xi^+_{j,\tau}\right)}{1 + \sum_k e^+_{k,\tau} \exp\left(\alpha^+ \ln b_{k,t} + \beta^+ x^+_{k,\tau} + \gamma^+_k + \mu^+_t + +\xi^+_{k,\tau}\right)}. \tag{15}$$

γ^+_j is a BM-unit fixed effect and μ^+_t is a half-hour fixed effect, to control for changes in the share of the outside good. $b_{j,t}$ 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 repositioning.[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

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

$\breve{\gamma}^+_j$ is a BM-unit fixed effect. $b_{j,t}$ is the bid for providing FR of BM unit j in the month t to which half-hour τ belongs. $\breve{x}^+_{j,\tau}$ contains additional hour-of-day (same for each day), day-of-week (same for each week), and month-of-year (same for each year) fixed effects and controls that parsimoniously represent the supply curves that the BM units bid in the BM. We include in $\breve{x}^+_{j,\tau}$ 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 $\breve{x}^+_{j,\tau}$ 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 $\breve{x}^+_{j,\tau}$ 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 $\breve{x}^+_{j,\tau}$ and independent across BM units and half-hours.

It follows that

$$\Pr(e^+_{j,\tau} = 1 | b_{j,t}, \breve{x}^+_{j,\tau}) = 1 - \Phi\left(-\breve{\alpha}^+ \ln b_{j,t} - \breve{\beta}^+ \breve{x}^+_{j,\tau} - \breve{\gamma}^+_j\right) = \Phi\left(\breve{\alpha}^+ \ln b_{j,t} + \breve{\beta}^+ \breve{x}^+_{j,\tau} + \breve{\gamma}^+_j\right), \tag{16}$$

where $\Phi(\cdot)$ is the standard normal CDF. We estimate equation (16) by ML. Moreover, equation (15) implies

$$\ln s^+_{j,\tau} - \ln s^+_{0,\tau} \equiv \delta^+_{j,\tau} = \alpha^+ \ln b_{j,t} + \beta^+ x^+_{j,\tau} + \gamma^+_j + \xi^+_{j,\tau}$$

[1] From 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.

Table 20: Demand estimates: Market share in upward and downward repositioning in BM

	Upward repositions		Downward repositions	
Log FR bid	-0.086	-0.108	-0.076	-0.100
	(0.027)	(0.027)	(0.009)	(0.009)
R^2	0.57	0.58	0.53	0.54
N obs	260482	260482	885659	885659

Separate OLS estimates of logit model for market share for upward and downward repositioning in BM. All models include BM-unit and half-hour fixed effects. The second and fourth models include the hypothetical market share of BM unit j in half-hour τ at each of the 24 prices in the grid for upward, respectively, downward repositioning (estimates not reported). The unit of observation is a BM unit-half hour. Standard errors are clustered by half-hour.

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 20 and 21 show our estimates for the logit model in equation (15) and the probit model in equation (16). In the first and third columns, we exclude the controls $\breve{x}^+_{j,\tau}$ 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 (15) and the probit model in equation (16), 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 model in equation (15), 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 (15). We use the shorthands $x^+_\tau = \left(x^+_{j,\tau}\right)_{j=1,\dots,J}$, $\xi^+_\tau = \left(\xi^+_{j,\tau}\right)_{j=1,\dots,J}$, and $e^+_\tau = \left(e^+_{j,\tau}\right)_{j=1,\dots,J}$. θ^+ denotes the parameters of the logit model in equation (15). Unconditionally (or in expectation), the

Table 21: Demand estimates: Eligibility for upward and downward repositioning in BM

	Upward repositions		Downward repositions	
Log FR bid	-0.04915	-0.03726	-0.19963	-0.09871
	(0.00816)	(0.00856)	(0.00670)	(0.00737)
Closest bid/offer price		-0.00000		0.00000
		(0.00000)		(0.00000)
Closest bid/offer quantity		-0.00009		0.00156
		(0.00004)		(0.00004)
N obs	1511766	1511766	1511765	1508785

Separate ML estimates of probit model for eligibility for upward and downward repositioning in BM. All models include BM-unit, month-of-year, day-of-week, and hour-of-day fixed effects. The second and fourth models include the lowest bid, respectively, offer of BM unit j in half-hour τ along with the corresponding quantity (estimates reported), ten dummies for the decile in which the lowest bid, respectively, offer of BM unit j in half-hour τ falls (estimates not reported). The unit of observation is a BM unit-half hour. Standard errors are clustered by half-hour. The sample is restricted to 20% of observations.

market share of BM j in half-hour τ is

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

where

$$w^+(b_t, \breve{x}_\tau^+, e_\tau^+; \breve{\theta}^+) \equiv \prod_{l=1,\ldots,J} \Phi\left(\breve{\alpha}^+ \ln b_{l,t} + \breve{\beta}^+ \breve{x}_{l,\tau}^+ + \breve{\gamma}_l^+\right)^{e_{l,\tau}^+} \left(1 - \Phi\left(\breve{\alpha}^+ \ln b_{l,t} + \breve{\beta}^+ \breve{x}_{l,\tau}^+ + \breve{\gamma}_l^+\right)\right)^{1-e_{l,\tau}^+} \tag{17}$$

and the summation is over all 2^J possible values of e_τ^+. $\breve{\theta}^+$ denotes the parameters of the probit model in equation (16).

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

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

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

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

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 $\frac{\partial x_\tau^+}{\partial b_{j,t}}$, $\frac{\partial \breve{x}_\tau^+}{\partial b_{j,t}}$, $\frac{\partial x_\tau^-}{\partial b_{j,t}}$, and $\frac{\partial \breve{x}_\tau^-}{\partial b_{j,t}}$ 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 $\frac{\partial s_j^+(b_t,x_\tau^+,\xi_\tau^+,\breve{x}_\tau^+;\theta^+,\breve{\theta}^+)}{\partial b_{k,t}}$ and $\frac{\partial s_j^-(b_t,x_\tau^-,\xi^-\tau,\breve{x}_\tau^-;\theta^-,\breve{\theta}^-)}{\partial b_{k,t}}$. We have

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

for $k=j$ and

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

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

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 equations[2]

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

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 cost of providing FR: as Table 22 shows, the average across BM units falls

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

Table 22: Marginal cost estimates and estimated markups

	Without repositioning	With repositioning
Average marginal cost	1.40	1.36
Main market markup		-0.0014
		(0.0041)

Average of marginal cost estimate c_j across BM units and estimated markup μ, with and without accounting for repositioning incentives. The estimates without accounting for repositioning incentives are the ones presented in Section 4.2.

from £1.41/MWh to £1.36/MWh. The estimated markup is not significantly different from zero.

A.8 Collusion

We try three different ways of examining the data for evidence of collusion. The first is to look for coordination in the timing and direction of bid changes across BM units. 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 18 we plot the distribution of correlation coefficients separately for BM units owned by the same firm ("within firm", left panels) and for BM units owned by different firms ("across firms", right panels). Note that we expect some across-firm correlation in both the timing and the direction of bid changes due to common shocks to demand.

The within-firm correlations for the timing and direction of bid changes in the left panels are positive and substantial. This reinforces our contention that decisions are centralized at the level of the firm rather than made at the level of the BM unit. The right panels show correlations pretty much evenly distributed around zero, consistent with independent decision making across firms. While we cannot rule out collusion in the earlier phases — and indeed we believe Drax attempted to establish a tacitly collusive arrangement in the middle phase — the lack of significant bid correlation is suggestive evidence against this.

Our second approach is more direct. We assume particular collusive arrangements and

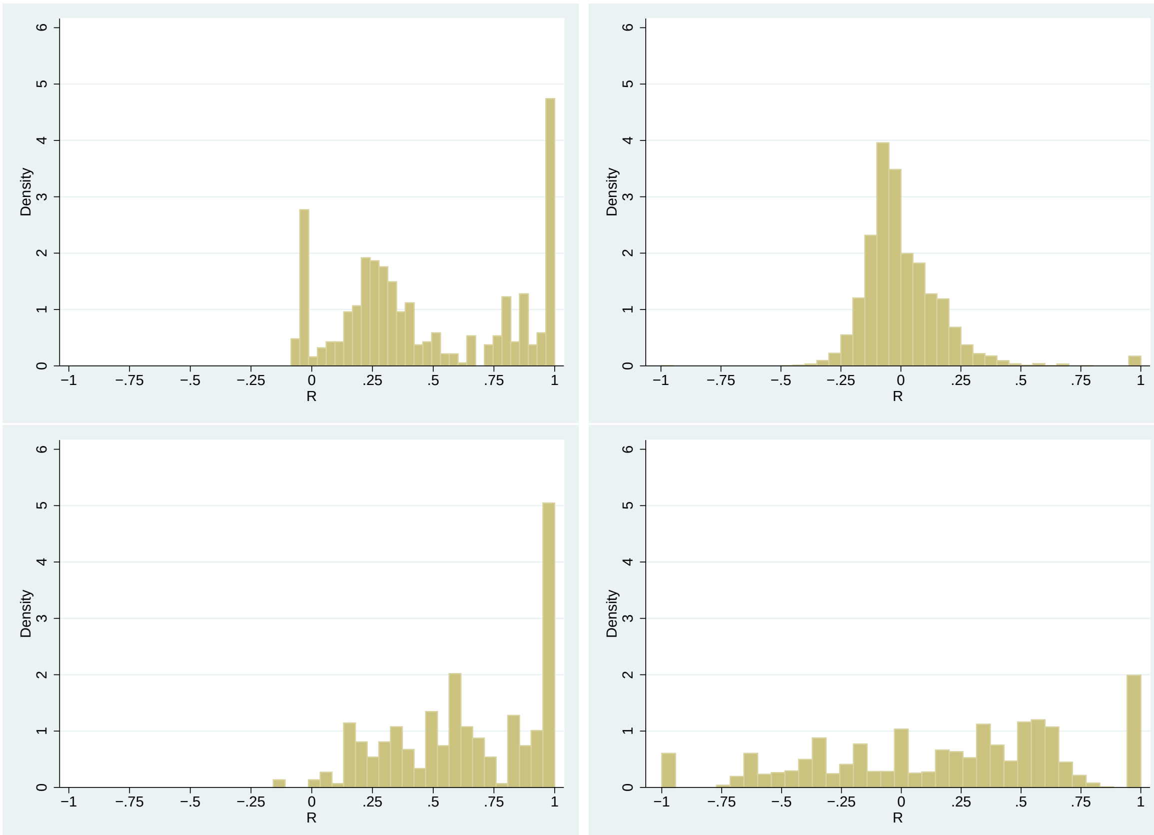

Figure 18: Top left is within-firm correlation in bid changes; top right is across-firm correlation in bid changes; bottom left is within-firm correlation in direction of change (conditional on both changing); bottom right is across-firm correlation in directions.

infer cost given the assumed conduct. Specifically, we re-solve equation (8) for the cost c_i 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 clearly implausible. The estimates are negative because demand is relatively inelastic, and so rationalizing the bids in the face of increased market power requires low cost. When we repeat the exercise assuming that only the top 3 firms collude, the implied average cost is £-0.25/MWh, still negative.

Finally, we ask at how much weight μ we can have a firm put on its rivals' profits while still ensuring that the observed bids b_t for months $t \geq 44$ are consistent with non-negative cost

$c = (c_j)_{j=1,\ldots,J}$. We formalize this question as the program

$$\max_{\mu,c} \mu$$

subject to

$$\frac{1}{29}\sum_{t=44}^{T=72}\Big[M_t s_{j,t} + \sum_{k\in\mathcal{J}_i}(b_{k,t}-c_k)\,M_t\,(1(j=k)-s_{j,t})\,\frac{\hat{\alpha}s_{k,t}}{b_{j,t}}$$
$$+\mu\sum_{k\in\mathcal{J}_{-i}}(b_{k,t}-c_k)\,M_t\,(1(j=k)-s_{j,t})\,\frac{\hat{\alpha}s_{k,t}}{b_{j,t}}\Big] = 0, \quad \forall j \in \mathcal{J},$$

and

$$\mu \geq 0,$$
$$c_j \geq 0, \quad \forall j = 1, \ldots, J.$$

Note that μ pertains to BM units $k \in \mathcal{J}_{-i}$ that are *not* owned by firm i.

Conditional on μ the program is linear in c. We thus start with $\mu = 0$ and solve for c. Then we successively increase μ and re-solve for c. The implied cost estimates for various values of μ are shown in Figure 19. The horizontal axis is μ and the vertical axis is the share of BM units. The blue line shows the share of BM units with negative cost and the orange line is the share for which the 95% confidence interval is entirely negative. By $\mu = 0.28$, over 5% of BM units have negative cost with 95% confidence. Consistent with the first two approaches, we see little evidence of collusion.

A.9 Switching costs and inattention

Suppose there is a fixed costs to adjusting bids or that a firm is occasionally inattentive to the FR market for exogenous reasons. Suppose further that firms are myopic. Then the first-order conditions in equation (7) hold in those months in which the firm actually adjusted its bids. This allows us to re-estimate cost by simply restricting the sample. These resulting estimates do not differ significantly from our baseline estimates, as Figure 20 illustrates.

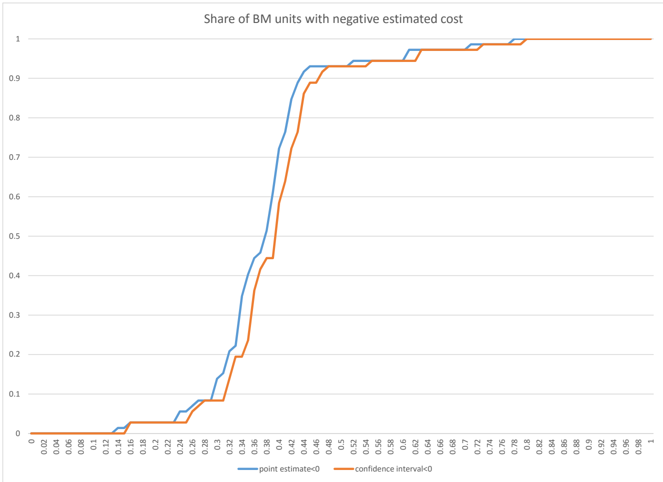

Figure 19: Share of BM units with negative estimated cost.

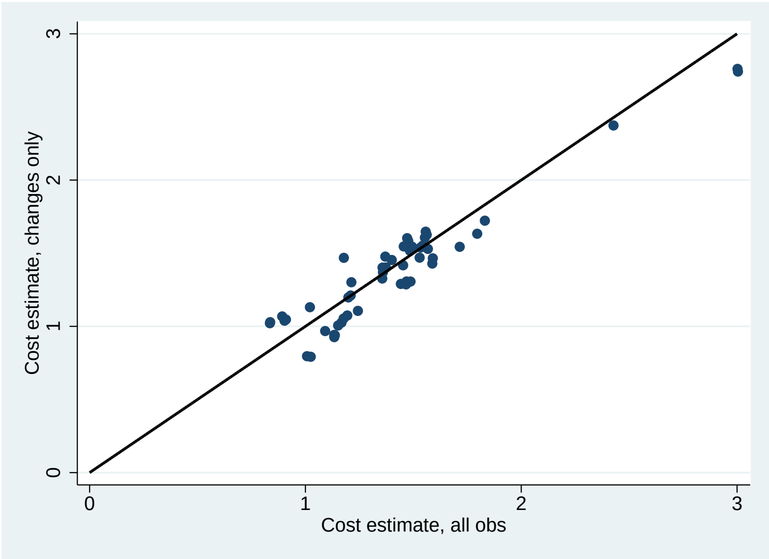

Figure 20: Comparison of cost estimates using all months versus only months in which the firm adjusted its bids.