

Composition of Electricity Generation Portfolios, Pivotal Dynamics and Market Prices*

Albert Banal-Estanol[†]

Augusto Rupérez Micola[‡]

November 3, 2008

Abstract

We use simulations to study how the diversification of electricity generation portfolios influences wholesale prices. We find that the relationship between technological diversification and market prices is mediated by the supply to demand ratio. In each demand case there is a threshold where pivotal dynamics change. Pivotal dynamics pre- and post-threshold are the cause of non-linearities in the influence of diversification on market prices. The findings are robust to changes in the main market assumptions.

Keywords: Electricity, market power, simulations, technology diversification.

1 Introduction

Electricity is a non-storable, undifferentiated commodity, delivered into a market with low demand elasticity, high security of supply requirements and wide seasonal variations. As a result, the industry accommodates a wide range of generating technologies. Some generators are technologically diversified and own nuclear plants on the base-load as well as higher cost thermal units. For example PG&E, a large US utility, owns hydro, nuclear, thermal and renewable plants (PG&E Corporation, 2006). Others are specialists, focusing on only one technology. Until recently, British Energy's generation portfolio was formed exclusively by eight nuclear generating units (British Energy, 2006).¹

*Previous versions of this paper have been presented at INFORMS Annual Meetings (Pittsburgh and Seattle), ICCMS (Geneva), Swiss IO Day (Bern), the CCRP Regulation and Competition Workshop (London), the Trento Summer School in Adaptive Dynamics, EURO (Prague), EARIE (Valencia), and seminars at City University (London) and EPFL (Lausanne). We acknowledge comments from the participants, Ann van Ackere, Antonio Conejo, Joanne Evans, John Garcia, Mireia Jofre-Bonet, Andreas Krause, Tarjei Kristiansen, Giulia Iori, Ángel López, Fernando Oliveira, José Olmo, Omiros Paspiliopoulos, William Pouliot, Willem Smit, Anke Weidlich, Klaus Zauner, an associate editor and three referees.

[†]Department of Economics, City University and Department of Business and Economics, Universitat Pompeu Fabra and Barcelona GSE (c/ Ramon Trias Fargas, 25-27. 08005 Barcelona, Spain. Tel. (+34) 93 542 1766. E-mail: albert.banalestanol@upf.edu. Web-site: <http://www.econ.upf.es/ca/personal/professor.php?id=p4134>

[‡]Department of Economics and Business, Universitat Pompeu Fabra and Barcelona GSE (c/ Ramon Trias Fargas, 25-27. 08005 Barcelona, Spain. Tel. (+34) 93 542 1755. E-mail: augusto.ruperezmicola@upf.edu). Web: <http://www.ruperezmicola.net>

¹British Energy has recently acquired one coal-fired plant.

A market in which generators are specialised could exhibit more market power because the price-setting part of the merit order is more concentrated. However, in electricity pools, specialised high-cost generators have less incentives to exert market power because they lack base-load plants to reap the benefits (Ausubel and Cramton, 2002). In contrast, diversified firms have incentives to use their high-cost plants to increase market prices and thereby increase the profit on the base-load, but may not have enough price-setting capacity to do so.

Our paper addresses the general questions of “what is the shape of the diversification to prices relationship?” and “what are its determinants?” Specifically, we consider different markets where a generation duopoly own varying amounts of base- and peak-load capacity that is bundled into low- and high- cost plants. In order to isolate the portfolio effects, we keep market concentration as well as market base-load and high-cost capacities constant.

The trading environment is a multi-unit, compulsory, uniform-price auction. Despite being close to real spot markets, this model is characterised by multiple non-Pareto ranked Nash equilibria (von der Fehr and Harbord, 1993), as in the "battle of the sexes" game. By inspecting equilibrium payoffs, however, we are able to identify crucial differences across equilibria and draw predictions about the effects of diversification on prices for each demand level.

The main prediction is that changes in the number of “pivotal” plants will cause nonlinearities.² In our setup, there is one, and only one, pivotal plant under little or no diversification. After a threshold, the number of pivotal plants changes. In low-demand situations, we predict that the change will result in further competitive pressures. In high-demand cases, there should be increased implicit coordination and higher prices. To test the hypotheses, we use a simulation method based on the adaptive theory of reinforcement learning put forward by Roth and Erev (1995).³

We find support for the non-monotonic diversification to market prices relationship. For each demand to supply ratio, we identify a diversification breaking point where dynamics change. In low-demand situations, prices drop after the breaking point, whereas in high-demand cases further diversification leads to higher prices. We show that the non-monotonicity is caused by regime changes in the firms’ incentives and ability to exert market power. The estimated breaking points are shown to statistically match the theoretical thresholds at which the number of pivotal plants changes. The findings are robust to a series of alternative modelling specifications, including single price and supply function peak-load bidding as well as inelastic and elastic demands.

Despite its importance, the literature on generation portfolios as a source of market power is relatively sparse. Borenstein et al. (1999) carry out Cournot numerical simulations to highlight the weaknesses of market power concentration measures and also analyse the impact of several divestitures on market power in California. Bushnell (2003) analyses competition among several firms owning a mixture of hydroelectric and thermal generation resources and concludes that firms may find it profitable to allocate more hydro production to off-peak periods than they

²A high-cost plant is pivotal if the quantity demanded exceeds the sum of production capacities of all other plants and, as a result, the plant is necessary to fulfill demand.

³Papers which also use the Roth and Erev method in the electricity context include Nicolaisen et al. (2001), Rupérez Micola and Bunn (2008) and Rupérez Micola et al. (2008).

would under perfect competition, or if they did not act strategically.⁴ Crawford et al. (2006) characterise the generators' asymmetric behaviour in equilibrium and analyse the consequences of some divestitures in England and Wales. Perhaps the closest paper to ours is by Bunn and Oliveira (2007), who use a simulation to model the interaction between an electricity market and a plant swapping game. They identify a symbiotic interaction between the two markets: initial situations where firms are perfectly diversified evolve, via plant trading, into lower electricity prices than those in which firms were originally specialised. To the best of our knowledge, though, ours is the first paper that studies explicitly the diversification to prices relationship.⁵

More attention has been devoted to quantifying the impact of marginal costs on prices and estimating the generators' ability to exert market power.⁶ Rudkevich et al. (1998), for example, characterise the pricing behaviour of identical firms and show that the equilibrium price increases with the system's production cost curve. Wolfram (1995) quantifies the markups in the England and Wales pool by computing the short-run marginal cost functions for each generator using fuel cost estimates. In a follow-up paper, Wolfram (1998) finds that generators with more inframarginal capacity submitted higher bids for units with comparable costs. Borenstein et al. (2002) and Joskow and Kahn (2002) decompose the dramatic California price and expenditure increases in the late Nineties and early 2000's into rising production costs, scarcity rents and ability to exercise of market power. Borenstein et al. (2002) find that 21% of the increased electricity expenditure is due to increased production costs, 20% is due to increased competitive rents, and the remaining 59% is attributable to increased market power.⁷

The remainder of the paper proceeds as follows. In Section 2, we discuss the main approaches that have been used to model electricity markets. Section 3 includes the model and an outline of our theoretical predictions. In Section 4 we introduce the simulation procedure. The main econometric results and robustness checks are presented in Sections 5 and 6. We conclude with a short discussion in Section 7. All theoretical proofs are in the Appendix.

2 Modelling competition in spot markets

There is some controversy about which model best fits competition in spot electricity markets (see Newbery, 1997, von der Fehr and Harbord, 1998, and Borenstein et al., 1999). Borenstein et al. (1999) argue that the Cournot model is an appropriate starting point since generating plants can be rendered "unavailable" due to maintenance and other reliability considerations. Other strategic models, like standard Bertrand or Bertrand with capacity

⁴In a related paper, Garcia et al. (2005) analyse the price-formation process in an infinite-horizon model where hydroelectric generators engage in dynamic price-based competition and show how simulations with a basic learning algorithm converge to the Markov Perfect Equilibrium.

⁵Arellano and Serra (2007) show how, in cases where a regulator uses peak-load marginal costs to determine wholesale prices, generators can exercise market power by increasing the share of peak technology in their portfolio.

⁶There have also been some attempts to analyse other possible ways of exercising market power, besides the unilateral effects. Fabra and Toro (2005) and Puller (2007) analyse collusive attempts to exercise market power in a dynamic context. Mansur (2007) shows that vertical integration mitigates market power.

⁷Market power not only increases prices but also creates market inefficiencies as more expensive competitive production substitutes less expensive production by firms with market power (Wolak and Patrick, 1997).

constraints, are unsuitable due to the use of uniform pricing in power pools. As a result, the electricity literature has suggested two main alternative approaches (von der Fehr and Harbord, 1998). One of them is the “auction approach”, suggested by von der Fehr and Harbord (1993) – vdF-H – and the other is Green and Newbery’s (1992) adaptation of the “supply function” equilibrium (SFE) due to Klemperer and Meyer (1989).

In the auction approach, the market is modelled as a sealed-bid, multiple-unit auction. Generators simultaneously submit single prices at which they are willing to supply each generating set. Bids are ranked according to their offer prices and the system marginal price is determined by the intersection of demand and supply. VdF-H show that for some demand levels there are no pure strategy equilibria and for others there are multiple equilibria (see also Crawford et al. 2006).

One key restriction of most vdF-H implementations (e.g. Nicolaisen et al., 2001; Rupérez Micola and Bunn, 2008, and Rupérez Micola et al., 2008) is that firms are allowed only a single bid price per plant. Inherent inflexibilities in the operation of nuclear plants (e.g. safety concerns, very low marginal costs and high start-up and loss of volume costs) prompt them to submit flat schedules at very low prices. This restriction, however, is quite unrealistic to model peak-load thermal plants, which typically bid several steps per unit. At the market level this leads to the well-known “hockey stick” shape of the supply curve, with base-load plants submitting flat bid schedules and more expensive peak-load generators offering steeper supply functions.

The SFE approach addresses this issue with a model in which each firm chooses a supply function relating quantity to price. Klemperer and Meyer (1989) show that the steeper the equilibrium supply functions, the more closely competition resembles the Cournot model; with flatter equilibrium supply functions, competition is closer to marginal cost pricing. Green and Newbery (1992) show that the Nash equilibrium in supply functions implies a high mark-up on marginal cost and substantial deadweight losses in the early Nineties England and Wales market.

The SFE model is not without its limitations, either. If the range of variation in demand is finite then it appears to have little predictive value, since almost anything between the Cournot and the competitive solution can be supported in equilibrium (see Bolle 1992).⁸ Further, the solution is undefined if there is no short-run demand elasticity (von der Fehr and Harbord, 1998). Another difficulty with this approach is the assumption that generators submit continuously differentiable supply functions.⁹ Restricting strategy sets to discrete step functions has the merit of being more realistic.

In summary, although the auction and the supply function approaches seem to correspond better with real spot markets, both of them exhibit multiple equilibria. As an alternative, some papers have used the Cournot model

⁸The literature offers refinements as a way of singling out a unique supply function equilibria. Klemperer and Meyer (1989) show that if outcomes with infinite demand occur with positive probability, and if demand can be met with non-binding capacity constraints, then there is a unique SFE. Green and Newbery (1992) focus on the highest profit equilibrium, rather than including other equilibrium possibilities. Alternatively, Baldick and Hogan (2002) choose a unique equilibrium by ruling out unstable equilibria and adding a price cap and capacity constraints. Newbery (1991) obtains a unique SFE by considering entry and assuming bid-coordination. More recently, Holmberg (2007, 2008) shows that a unique SFE exists if there are binding capacity constraints with a positive probability.

⁹Holmberg (2007) allows for horizontal segments in a supply function context but shows that these segments are not consistent with equilibrium behaviour.

(e.g. Bushnell, 2003, and Borenstein et al., 1999). In this paper, we explore a different route. We draw from both the auction and the supply function models and try to obtain results in the multiple equilibrium environment. We use the auction approach in the main part of the paper and leave the supply function approach for the extensions.

3 Basic model

3.1 Market structure

Our model incorporates key features of electricity markets in the short run. Two companies that compete for the supply of the market own a mix of low (e.g. nuclear) and high (e.g. thermal) marginal cost capacity.¹⁰ Denoting the generating companies as 1 and 2 and the overall market capacity as K , the capacities of their respective low (l) and high (h) cost plants are

$$k_1^l = k_2^h = \frac{(1-\alpha)K}{2} \quad \text{and} \quad k_1^h = k_2^l = \frac{\alpha K}{2},$$

where $\alpha \in]0, 0.5[$ represents the degree of portfolio diversification. In the case of specialisation ($\alpha = 0$), company 1 is a low-cost specialist and company 2 is specialised in high-cost technology. Portfolio diversification increases with α , a growing proportion of the base-load generator's capacity is exogenously replaced with high-cost units. Symmetrically, the generator's high-cost capacity is replaced with base-load. In the case of full diversification ($\alpha = 0.5$), each company holds the same amount of low- and high-cost generating capacity. This formulation isolates the effects of portfolio diversification because, while allowing for different degrees of diversification, the total capacity of each company is kept constant ($k_i^l + k_i^h = K/2$ for $i = 1, 2$) as are the market aggregates of low- and high-cost capacities ($k_1^j + k_2^j = K/2$ for $j = l, h$).

Marginal costs are assumed to be constant; normalised to 0 for the low-cost plants ($c_l = 0$) and equal to c for the high-cost plants ($c_h = c$),¹¹ and there are no grid constraints.¹² Although relevant in the long term, we do not deal with entry and exit of firms, capacity expansion, the use of long-term contracts (as in e.g. Baldick et al., 2006), ancillary and capacity payments.

3.2 Market rules

Trading takes place through a multi-unit, compulsory, uniform-price auction. In the basic model, firms submit simultaneous single-price bids at which they are willing to sell up to the capacity of each plant.¹³ Possible bids are

¹⁰The model could easily be extended to other more realistic market configurations, including all sorts of oligopolies and the existence of a competitive fringe. However, our analysis in a stylised market is more transparent and comparable to previous literature.

¹¹In the robustness section, we allow for strictly positive marginal costs for the base-load unit. The main variable cost component is fuel costs (e.g. enriched uranium, natural gas), which we assume to be constant.

¹²The addition of network constraints would undoubtedly make the analysis richer but it would also make it more complicated to disentangle effects due exclusively to technology diversification from those arising from local market power exerted by relatively small players.

¹³In the robustness section we relax this assumption and allow for supply function bidding. Piecewise supply functions have been introduced, among others, by Hobbs and Pang (2007).

bounded between marginal costs and Ψ , the maximum “reasonable” price.¹⁴

We model the market demand Q as fully inelastic,¹⁵ drawn from a uniform distribution in the interval $[\bar{Q} - \varepsilon, \bar{Q} + \varepsilon]$, where \bar{Q} is its expected value and ε accounts for the small uncertainty typical in day-ahead forecasting.¹⁶ We assume that there is always some system overcapacity, $\bar{Q} + \varepsilon < K$, but that demand exceeds the market aggregate of low-cost capacity, $\bar{Q} - \varepsilon > K/2$, consistent with the normal operations of many deregulated energy markets.¹⁷

An independent auctioneer determines the uniform market price P by intersecting the *ad hoc* supply function with the realised demand. She assigns full capacity, $q_i^j = k_i^j$, to the M plants with bids below the market price; the remaining capacity, $q_i^j = Q - \sum_{\{i,j\} \in M} k_i^j$, to the plant with a bid equal to the market price;¹⁸ and zero sales, $q_i^j = 0$, to those bidding above the market price. Profits for each company are

$$\pi_i = P q_i^l + [P - c] q_i^h \quad \text{for } i = 1, 2. \quad (1)$$

3.3 Nash equilibria and predictions

In this subsection we compute the equilibria of the model, which depend crucially on the number of “pivotal” plants (see e.g. Genc and Reynolds, 2005; Entriken and Wan, 2005; Perekhodtsev et al., 2002).

Definition 1 *A plant is pivotal if (i) it is high-cost and (ii) if the quantity demanded exceeds the sum of production capacities of all other plants.*

Definition 2 *A level of diversification α_v is a switching point if the number of pivotal plants for $\alpha < \alpha_v$ is different than that for $\alpha \geq \alpha_v$.*

If $\bar{Q} = 240$ and $K = 300$, for example, the switching point is $\alpha_v = 0.40$ because the number of pivotal plants changes from one to two at this level. If $\alpha < 0.40$, the peak-load plant of Firm 2 is the only pivotal since $k_1^l + k_2^l + k_1^h < 240$ and $k_1^l + k_2^l + k_2^h > 240$, whereas if $\alpha > 0.40$ both peak-load plants are pivotal because $k_1^l + k_2^l + k_1^h < 240$ and $k_1^l + k_2^l + k_2^h < 240$. If $\bar{Q} = 180$ (and $K = 300$), the number of pivotal plants is reduced from one to none at a switching point $\alpha_v = 0.13$.

The level of excess capacity generates two pivotal switching regimes:

Proposition 3 *For any \bar{Q}*

(a) if $K/2 < \bar{Q} \leq 3K/4$, the switching point is $\alpha'_v(\bar{Q}) \equiv (2\bar{Q} - K)/K$, at which the number of pivotal plants is reduced from one to none.

¹⁴This upper price ceiling can be understood as a limit triggering regulatory intervention or the cost of alternative, expensive, load fuels to which the system administrator could switch at short notice if prices exceed Ψ . It also reflects high-cost back-up power generation facilities owned by many industrial users.

¹⁵Several contributions show that price elasticity of short-term electricity demand is typically quite low (e.g. Stoft, 2002). In the robustness section we allow for short-term demand elasticity.

¹⁶The small uncertainty is introduced for the sake of realism, but its absence would not alter our findings.

¹⁷For example, the UK energy system includes a reserve margin of about 20% of expected peak demand.

¹⁸In case of a tie, the selling plant is selected randomly.

(b) if $3K/4 < \bar{Q} \leq K$, the switching point is $\alpha_v''(\bar{Q}) \equiv 2(K - \bar{Q})/K$, at which the number of pivotal plants is increased from one to two.

For each demand level, there is a switching point where the number of pivotal plants changes. There is always one pivotal plant when diversification is low. In contrast, there are no pivotal plants in relative spare capacity cases (i.e. $\alpha > \alpha_v'$ when $\bar{Q} \leq 3K/4$), and two, if capacity is tight (i.e. for $\alpha > \alpha_v''$ when $\bar{Q} > 3K/4$). Table 1 summarises the pivotal switching points for ten-unit step demand levels when $K = 300$, together with the capacities of each plant and the number of pivotal plants before and after the thresholds.

<<TABLE 1: PIVOTAL SWITCHING POINTS FOR EACH DEMAND LEVEL>>

Our trading setting often presents multiple non-Pareto-ranked equilibria (von der Fehr and Harbord, 1993; Crawford et al., 2006). Denoting the equilibrium bids of each firm as $(b_1^l, b_2^l, b_1^h, b_2^h)$, the following proposition characterises the pure strategy Nash equilibria and shows how they depend on the pivotal dynamics.

Proposition 4 *There exist $\alpha_1(\bar{Q})$ and $\alpha_2(\bar{Q})$ such that*

(a) *if no firm is pivotal ($\alpha_v' < \alpha \leq 0.5$), then*

$$\begin{aligned} b_1^l &= \Psi - \varepsilon, b_1^h = \Psi \text{ is part of an equilibrium if } \alpha \geq \alpha_1, \\ b_2^l &= \Psi - \varepsilon, b_2^h = \Psi \text{ is part of an equilibrium if } \alpha \leq \alpha_2 \text{ and} \\ b_1^l &= b_2^l = 0 \text{ and } b_1^h = b_2^h = c \text{ is an equilibrium if } \alpha > \alpha_2, \end{aligned}$$

(b) *if only one high-cost plant is pivotal ($\alpha \leq \min\{\alpha_v', \alpha_v''\}$), then*

$$\begin{aligned} b_1^l &= \Psi - \varepsilon, b_1^h = \Psi \text{ is part of an equilibrium if } \alpha \geq \alpha_1 \text{ and} \\ b_2^h &= \Psi \text{ is part of an equilibrium for any } \alpha, \end{aligned}$$

(c) *if both high-cost plants are pivotal ($\alpha_v'' < \alpha \leq 0.5$), then*

$$\begin{aligned} b_1^h &= \Psi \text{ is part of an equilibrium for any } \alpha \text{ and} \\ b_2^h &= \Psi \text{ is part of an equilibrium for any } \alpha. \end{aligned}$$

For any class of equilibrium, there are many payoff-equivalent, lower bids of the other plants which are part of a pure strategy equilibria.

Figure 1 depicts the number of pivotal plants and the bidding equilibria for different demand and diversification level when $K = 300$. The number of pivotal plants is determined by a centered inverted V-shape curve. Below the curve (areas b.1 and b.2), there is one pivotal plant. In the northwest area (areas a.1, a.2 and a.3), no plant is pivotal. In the northeast (area c), the two high-cost plants are pivotal. Within the area in which no plant is pivotal, the first equilibrium is possible in area (a.1), the second in areas (a.1) and (a.2) and the third in area (a.3). Within the area in which one firm is pivotal, the first equilibrium exists in area (b.1) and the second one exists in

areas (b.1) and (b.2). Finally, both equilibria in the case in which both high-cost plants are pivotal exist in all of area (c).

<<FIGURE 1: NASH EQUILIBRIA FOR EACH α AND \bar{Q} WHEN $K = 300$ >>

Proposition 4 shows that non-Pareto-ranked multiple equilibria are widespread.¹⁹ For most parameter values, one of the firms keeps prices high while the other submits lower bids and sells its full capacity at the maximum price. As a result, one of the firms obtains a relatively low profit whereas the other gets its maximum profit. Our model is similar to the “battle of the sexes”, a standard coordination game with two pure strategy equilibria with asymmetric payoffs. Experimental evidence shows that coordination, even in the simple cases, can be lower than 50% for this game (Cooper et al., 1990). Coordination is especially difficult when there are equilibrium payoff asymmetries because the point of coordination is not clear (Crawford et al., 2008).²⁰ In our setting, firms will have problems to coordinate over who is price-setting, and this will be more difficult the higher the equilibrium profit asymmetry. More asymmetry will reduce the firms’ incentives to set the price and induce them to submit lower bids.

This allows us to draw some predictions about the effects of diversification on prices for each demand level.²¹ Prices should decrease when the number of pivotal plants moves from one to zero. First, with no pivotal plants, there is an equilibrium in which prices are equal to marginal costs, while the equilibrium price is always equal to the maximum in the one-pivotal cases. Second, the equilibria are very asymmetric when there are no pivotal plants because one firm uses both plants to set the price and obtains a relatively low profit while the other gets its maximum profit. In one of the one-pivotal plant equilibria, though, the price-setting firm sells at least its full low-cost capacity and profits are more symmetric. Hence, prices should decrease when the number of pivotal plants moves from one to zero, i.e. in low-demand cases.

Hypothesis 1 *The dynamics pre- and post-switching point result in nonlinearities in the influence of diversification on market prices. A breaking point occurs at the switching point, where the number of pivotal plants changes.*

Hypothesis 2 *In spare capacity cases, prices drop at the breaking point.*

Similarly, we would expect higher prices as the number of plants increases from one to two (high demand cases). Coordination problems should be less severe in this case since the price-setting firm uses the peak load plant to set the price and that results in more symmetric profits in the two equilibria. In contrast, when there is only one pivotal plant, there is an equilibrium in which the price-setting firm uses both plants.

¹⁹Furthermore the game also has many mixed strategy Nash equilibria.

²⁰As argued by Iriberry (2006), the only symmetric equilibrium, in which identical players choose the same action and obtain an equal expected payoff, is the mixed strategy equilibrium. But each player prefers even his worst pure-strategy equilibrium to the symmetric mixed-strategy equilibrium

²¹Absolute prices are not comparable as demand should be a fundamental factor influencing prices, albeit for most of the parameter regions of the proposition the equilibrium price is the maximum price.

Hypothesis 3 *When capacity is tight, prices increase at the breaking point.*

Although the three regions in Figure 1 can be divided further into subregions, we do not expect the diversification to prices relationship to present other breaking points. The presence of equilibria in which the price-setting firm obtains very low profits should be less important. However, as a robustness check, we will test the existence of a second breaking point.

4 Simulation Procedure

4.1 Behavioural learning

In equilibrium multiplicity cases, a selection method is necessary to choose amongst them. In broad terms, there are two schools of thought in the area of equilibrium selection (Haruvy and Stahl, 2004). On the one hand, we have deductive selection — based on reasoning and coordination in focal points — and, on the other hand, we have inductive selection — based on adaptive dynamics. Until recently, deductive principles have dominated the equilibrium selection literature. Existing deductive mechanisms, however, have been shown to do poorly in experiments (see e.g. van Huyck et al., 1990). Simple adaptive learning dynamics, instead, often yield successful equilibrium predictions (see e.g. Camerer and Ho, 1999, and Roth and Erev, 1995).

Reinforcement models are widely used adaptive learning mechanisms (see e.g. Erev and Haruvy, 2001, Sandholm, 2002, and Hopkins, 2008), especially in energy markets (Nicolaisen et al., 2001, Veit et al., 2006, Sun and Tesfatsion, 2007, Rupérez Micola and Bunn, 2008, Rupérez Micola et al., 2008). They are based on the law of effect, whereby actions that result in more positive consequences are more likely to be repeated in the future, and on the law of practice, whereby learning curves tend to be steep initially and then flatter.

We adopt the well-known reinforcement learning method put forward by Roth and Erev (1995) – R-E. This method has some advantages. Since it is widely used, and is simple and more parsimonious than other algorithms, our results are more easily comparable to previous models. Moreover, its principles fit some features of energy market trading well.²² It is based on the law of effect and the law of practice, which are robust properties observed in the literature on human learning. One of its main strengths is that one does not need to make assumptions on the information that players have about strategies, history of play and the payoff structure of the other players. In many cases energy market players cannot observe one another’s current strategies, and only imperfectly infer them from volatile prices.²³

Models of reinforcement learning rule out dominated strategies and, in particular situations like constant-sum games with a unique equilibrium, strategies converge to the Nash equilibrium (Beggs, 2005). Van Huyck et al. (1997) show that, in a generic coordination game with multiple Pareto-ranked-equilibria, the selected equilibrium

²²Other simulation algorithms are either completely naive (e.g. zero intelligence) or difficult to interpret in an energy market context (genetic algorithms, Q-learning).

²³Other algorithms like best response, fictitious play or experience-weighted attraction require an amount of information in the agents’ hands that we find difficult to justify.

depends on historical accidents, rather than on deductive concepts of efficiency. In our more complex game with multiple non-Pareto ranked equilibria, we also expect R-E to depend on the stochastic process. The convergence to a unique steady state, pure or mixed strategy Nash equilibria, cannot be guaranteed in general (van Huyck et al., 1997). However, and due to the law of practice, R-E stabilises at one point and a pattern of actions emerges. Thus our hypotheses’ tests proceed in two steps. First, we perform many simulation for each combination of parameters. Then, we use econometrics to determine whether the stochastic regularities are systematic and, if so, how they depend on the parameters.

4.2 Implementation

In our implementation, learning takes place by repeating the following three steps in each period.

Step 1 *Generators submit price offers for each plant according to a plant-specific probability distribution over the set of possible bids.*

In the main simulation, the feasible price offer domain for each plant is approximated by a discrete grid. For each plant, generators choose among S possible prices, equally spaced between the minimum and the maximum price offer. That is, the sets of possible bids for the low- and high-cost plants, B^l and B^h , range from 0 and c ,²⁴ respectively, up to Ψ ,

$$B^l = \{s(\Psi/S) \mid s = 1, \dots, S\}, \quad (2)$$

$$B^h = \{c + s(\Psi - c)/S \mid s = 1, \dots, S\}. \quad (3)$$

Each bid is generated by an “action s ”. Bids generated from lower actions are more competitive, i.e. closer to marginal costs. Notice though that the same action s implies a higher bid for a peak-load generator than for a base-load generator.

In each round t , each generator i selects an action s for plant j with a likelihood or “propensity” $r_{i,s}^j(t) > 0$. The probability of an action being played is given by its propensity divided by the sum of the propensities of all possible actions,

$$p_{i,s}^j(t) = \frac{r_{i,s}^j(t)}{\sum_{u=1}^S r_{i,u}^j(t)}. \quad (4)$$

Propensities for all actions are initialised to the plants’ maximum per-period profit, i.e. $r_{i,s}^j(1) = \Psi k_i^j$, so that all actions have the same initial probability, $p_{i,s}^j(1) = \frac{1}{S}$ for all s , i and j .

Step 2 *The auctioneer determines the market price by intersecting the ad hoc supply function with the realised demand.*

²⁴An alternative is to allow expensive plants to bid below c so that they have to find out for themselves that this is not profitable. This slows the learning process down but does not alter our results. Thus, as most of the electricity simulation literature, we do not allow firms to bid below marginal costs.

The realised demand is assumed to be independently distributed across periods. The price and the individual quantities are communicated independently to each generator.

Step 3 *Each plant-specific probability distribution is adjusted based on the performance of the bid used.*

At the end of each round, plants reinforce the selected action, \bar{s} , through an increase in its propensity equivalent to the performance of the *company as a whole*, $\pi_i(t)$. Actions that are similar, i.e. $\bar{s} - 1$ and $\bar{s} + 1$, are reinforced to a lesser extent, by $(1 - \delta)\pi_i(t)$ where $0 < \delta < 1$ (“persistent local experimentation” in the terminology of R-E). All propensities are discounted by γ (“gradual forgetting”) and actions whose probability falls below a certain threshold are removed from the space of choice (“extinction in finite time”). The pre-extinction propensities for the following period $r_{i,s}^{j'}(t+1)$ are

$$r_{i,s}^{j'}(t+1) = \begin{cases} (1 - \gamma) r_{i,s}^j(t) + \pi_i(t) & \text{if } s = \bar{s} \\ (1 - \gamma) r_{i,s}^j(t) + (1 - \delta) \pi_i(t) & \text{if } s = \bar{s} - 1 \text{ or } s = \bar{s} + 1 \\ (1 - \gamma) r_{i,s}^j(t) & \text{if } s \neq \bar{s} - 1, s \neq \bar{s} \text{ and } s \neq \bar{s} + 1, \end{cases}$$

and the final propensities, corrected by the extinction feature, are

$$r_{i,s}^j(t+1) = r_{i,s}^{j'}(t+1) I_{\left\{ \frac{r_{i,s}^{j'}(t+1)}{\sum_{u=1}^S r_{i,u}^{j'}(t+1)} > \mu \right\}}, \quad (5)$$

where I is an indicator function that takes value 1 if the condition between brackets is satisfied and is otherwise zero.

Algorithm 1 *Steps 1 to 3 are repeated until convergence has been achieved.*

Following the theoretical literature on learning (see Fudenberg and Levine, 1998, for an overview), we define convergence in terms of mixed strategy profiles. In our case, this is done using the empirical frequencies of each player’s actions.²⁵ We consider that a simulation run has converged if the maximum attainable per-period change in the probability of playing any strategy is below a (small) threshold τ .

Definition 5 *For a given τ (small), a simulation run has converged to a mixed strategy profile z at time t if for any potential action profile a in time $t+1$, the probability distribution adjustment of any action s of any plant j of any generator i is such that*

$$\left| p_{i,s}^j(t+1) - p_{i,s}^j(t) \right| < \tau.$$

The resulting price is computed from the agents’ mixed strategy profile z .

²⁵In this sense, R-E is a particularly attractive algorithm because it is based on the "power law of practice" psychological principle: learning curves tend to be steep initially, and then flatter. By construction agents operating under Roth-Erev learn fast in the beginning of the simulation but their learning slows down with increasing practice.

We set an exogenous (small) τ .²⁶ In each period, we identify the action with the lowest probability in a given period and calculate the hypothetical probability (at $t+1$) that would result from that action leading to the highest attainable profit. The simulation has not converged as long as the difference between the hypothetical probability (at $t+1$) and the true (at t) is higher than τ . It has converged when it is lower. When τ is smaller, the threshold is more stringent and the simulation will run for more periods.

The resulting mean price is computed from the agents' choice distributions.²⁷ Note that the convergence definition is compatible with the survival of several trading actions, as in mixed strategies. Price volatility is not equal to zero even if there is convergence to a steady state. Further, reinforcement learning depends on the stochastic process and, as a result, simulations for the same parameters might lead to different mean prices, i.e. the standard deviation of mean prices across simulations is not necessarily equal to zero.

4.3 Parameters and data set

The end-user reasonable price ceiling is set at $\Psi = 200$, with a discrete grid of $S = 100$ possible prices. Total capacity is set to $K = 300$, so that each generator's capacity is $K/2 = 150$. Marginal costs for the high-cost plants are $c = 100$ and equal to zero for the low-cost plants. With these parameters, the firms' maximum attainable profits range between those for a specialised base-load firm (30,000 monetary units) and those for a specialised peak-load firm (15,000 monetary units), depending on the capacity configuration. Firm 1's maximum profit decrease from 30,000 when $\alpha = 0$ to 22,500 when $\alpha = 0.5$ whereas Firm 2's maximum profit increase from 15,000 when $\alpha = 0$ to 22,500 when $\alpha = 0.5$. The maximum industry profits remain constant at 45,000.

We perform simulations for a discrete grid of fourteen expected demand cases, $\bar{Q} = \{160, 170, \dots, 290\}$, corresponding to market expected excess capacity of 46.66% through 3.33%, with a small uncertainty ($\varepsilon = 5$). If, for example, $\bar{Q} = 240$ then $Q \sim U[235, 245]$. For each instance, we consider fifty-one diversification levels, $\alpha = \{0, .01, .02, \dots, .50\}$. Further, we check the robustness of the analysis to changes in R-E parameters by taking nine combinations of the learning parameters, $\gamma = \{0.0025, 0.005, 0.0075\}$ and $\delta = \{0.25, 0.50, 0.75\}$, with $\mu = 0.0005$ and $\tau = 0.00025$ throughout. For each specification, we perform fifty simulation runs. To create our data set we take, for each simulation run, α , \bar{Q} , γ and δ , the expected price resulting from the agents' choice distributions once convergence has been achieved. Hence, our data set includes $50 \times 51 \times 14 \times 3 \times 3 = 321,300$ observations.

As representative cases, we focus on the demand cases of $\bar{Q} = \{240, 180 \text{ and } 280\}$, with expected excess capacities of 20%, 40% and 6.66%, respectively. In those examples, we approximate power systems under normal operations, spare and tight capacity conditions.

²⁶As explained below, we set $\tau = 0.00025$ in our parametrisation, which implies that simulations run for between 500 and 2000 periods. We have performed several robustness checks to the specification of this parameter but they do not alter our results.

²⁷Similar results are obtained if we take the average price obtained by letting the simulations run for an additional number of periods, beyond the convergence point. In all cases, the one-lag Augmented Dickey-Fuller test-statistic for the price series was typically lower than -10. Given that the 95% critical value is -3.43, the null hypothesis of non-stationarity is always rejected.

5 Econometric results

We present our estimation results in several steps. We start by exploring the impact of diversification on prices and show, with some examples, that the relationship is not linear. Then we use the theoretical predictions to explore the nonlinearities. First, we test whether the data set features breaking points in the locations predicted by the theory. Then, we use an intuitive procedure to test whether the theoretically derived piecewise model is best fitting. Subsequently, we test via a formal statistical model the existence of a breaking point at the theoretical location and discard the existence of a second breaking point. Finally, we present some results on the impact of diversification on the stationary bidding strategies.

5.1 Linear diversification/price relationship

Table 2 reports the results of a linear regression between α and stationary market prices, with \bar{Q} as a covariate, and fixed effects for δ and γ . The results show a positive relationship between demand and prices as expected. The diversification estimate α is negative and strongly significant. A naïve observer might be tempted to see this as evidence of diversification leading to lower prices. One of the main purposes of this paper, though, is to show that nonlinearities are central.

<<TABLE 2: FIXED EFFECTS REGRESSION RESULTS>>

Figures 2-4 report the mean price and its 95% confidence interval as a function of portfolio diversification (α) for $\bar{Q} = \{240, 180 \text{ and } 280\}$, with $\gamma = 0.005$ and $\delta = 0.50$. In Figure 2 ($\bar{Q} = 240$), prices remain around 168 in the interval $\alpha \in (0, 0.40)$. At $\alpha = 0.40$, prices increase to around 176 and stay there until $\alpha = 0.50$. Figures 3 ($\bar{Q} = 180$) and 4 ($\bar{Q} = 280$) reinforce the view of demand, or its analogue “excess capacity”, mediating on the influence of portfolio diversification. In the spare capacity situation (Figure 3) prices are lower than in the baseline case. The specialization price is 154.1. Further, the relationship between prices and diversification is flat until about $\alpha = 0.19$, where there is a discontinuity. Prices drop to 142 and stay low in spite of a slight upward trend. In the tight capacity situation (Figure 4), prices are slightly higher but they seem to follow a similar pattern to the baseline case. They start at 173.7 and stay flat until around $\alpha = 0.13$, where there is an increase to about 179.1. Beyond that point, they are flat once again, albeit at the higher level.

<FIGURES 2, 3, 4: PORTFOLIO DIVERSIFICATION AND MEAN PRICE FOR $\bar{Q} = \{240, 180 \text{ and } 280\}$ >

Taken together, Table 2 and Figures 2 to 4 provide interesting insights. Table 2 suggests that the shape of the portfolio diversification to market prices relationship could be thought of as decreasing if one had only considered a linear model. More careful analysis, though, reveals two further elements: First, the relationship is not monotonic and second, there is a significant variation in its shape, depending on \bar{Q} . Structural breaks seem to occur for $\alpha = 0.40$ in the $\bar{Q} = 240$, case, for $\alpha = 0.19$ for $\bar{Q} = 180$ and $\alpha = 0.13$ for $\bar{Q} = 280$. Note that these points are visually close to the respective pivotal switching points identified in the theory section. Finally, notice that prices

move downwards when we move from one to zero pivotal plants ($\bar{Q} = 180$) and upwards when the movement is from one to two pivotal plants ($\bar{Q} = 240$ and $\bar{Q} = 280$).

5.2 Non-linear relationship and pivotal regime switching points

In this section, we carry out a test of whether the data set features pivotal regime switching points in the locations predicted by the theory. To that purpose, we use the pivotal switching points from Table 1 and estimate a piecewise linear model between diversification and prices for each demand and R-E combination. The models are uniquely specified by a dummy variable associated with the threshold value α_v ,

$$P_i = \beta_0 + \beta_1 D_i + \beta_2 \alpha_i + \beta_3 D_i \alpha_i + u_i, \quad (6)$$

where $D_i = 0$ if $\alpha_i < \alpha_v$, $D_i = 1$ when $\alpha_i \geq \alpha_v$. The pre- and post- breaking points regression estimates are specified by

$$E(P_i|D_i = 0, \alpha_i) = \beta_0 + \beta_2 \alpha_i \text{ and } E(P_i|D_i = 1, \alpha_i) = (\beta_0 + \beta_1) + (\beta_2 + \beta_3) \alpha_i.$$

We test the null hypothesis of linearity against the alternative of structural breaks at the pivotal switching points α_v . Evidence for the existence of a breaking point can come either from significant intercept change or slope change coefficients, i.e. β_1 and β_3 different from zero. Table 3 summarises the results obtained through Fama-MacBeth-type (1973) regressions averaging across R-E treatments for each demand.²⁸ In all fourteen \bar{Q} cases at least one of the two coefficients is significant at standard levels. There are five non-significant coefficients (one intercept, four slopes) but in all cases the other coefficient is significant at the .01 level. These results indicate that the data presents the regime switching points suggested by the theory. The tests provide preliminary support for our hypotheses.

<<TABLE 3 PARAMETER ESTIMATES FOR REGIME SWITCHING POINTS>>

5.3 Testing the location of the breaking points

This section uses an intuitive procedure to test whether the theoretical piecewise model is best fitting. The procedure selects the ‘‘optimal’’ breaking point from the data and seeks to establish whether it is statistically equivalent to the one obtained through theory. We start by providing a definition of an optimal breaking point.

Definition 6 *The optimal breaking point $\hat{\alpha}$ satisfies $F(\hat{\alpha}) \geq F(\alpha)$ for any threshold α , where $F(\alpha)$ denotes the F-statistic obtained from a piecewise linear regression with threshold α .*

Figure 5 reports the optimal piecewise linear models for each demand. Each line corresponds to an R-E parameter combination and each panel corresponds to a different demand specification. If $160 \leq \bar{Q} \leq 220$ (*low* demand) the thresholds result in price reductions of about 10 monetary units. When $230 \leq \bar{Q} \leq 290$ (*high* demand),

²⁸Fama and MacBeth (1973) popularised the following two-step procedure. In the first step, for each treatment a regression is performed. In the second step, the final coefficient estimates are obtained as the average of the first step coefficient estimates.

the breaking points coincide with clear price jumps followed by flat diversification/prices relationships. The size of the high demand jumps is approximately 6 monetary units. This piece of evidence supports Hypotheses 2 and 3. Note that the R-E assumptions have little influence on price levels.

<<FIGURE 5: AVERAGE BEST-FIT REGRESSIONS FOR R-E COMBINATIONS AND \bar{Q} >>

As an additional step, we compare the theoretical and estimated values with confidence intervals. We approximate the distribution of the structural breaks through a subsampling procedure.²⁹ For each demand and R-E assumptions, we extract 99 random subsamples of 1,020 observations stratified for α (we use twenty observations for each α). Then, we use Definition 6 to obtain the subsamples' optimal breaking points.

In Figure 6, we combine mean (squares) and the 95% confidence intervals of the optimal breaking points with their corresponding theoretical values from Table 1 (diamonds). Each panel reports one R-E combination and, for each of them, we represent the simulated and theoretical pivotal breaking points (vertical axis) in each demand level (horizontal axis). They co-move following an inverted V-shape. Intervals are very narrow, even indistinguishable, and visually overlap with the theoretical predictions. The tests reject the hypothesis of equal theoretical and simulation-derived pivotal switching points in only one instance out of 126 (table available upon request). Moreover, since the theoretical prediction does not depend on the R-E parameters (it is the same across Figure 6 panels), these results also confirm that the breaking point for each demand does not depend on the R-E specification.

<<FIGURE 6: THEORETICAL AND ESTIMATED BREAKING POINTS FOR R-E COMBINATIONS AND \bar{Q} >>

5.4 Testing the number of breaking points

In this section, use the framework developed by Gombay et al. (1996) and extended by Orasch (1999) to formally test the existence of the breaking point at the theoretical prediction and discard the existence of a second one.³⁰ Following their notation, we redefine the theoretical and estimated breaking points, k_v and \tilde{k} , as the integers immediately below the change. The predicted pivotal points for $\bar{Q} = 180$, $\bar{Q} = 240$, and $\bar{Q} = 280$ are 19, 39 and 12, respectively.

Gombay et al. (1996) show that the statistic \tilde{k} , defined as

$$\tilde{k} \equiv \arg \max_k M(k) \text{ where } M(k) \equiv n^{-3/2} k(n-k) \left\{ \frac{1}{k} \sum_{1 \leq i \leq k} \hat{u}_i^2 - \frac{1}{n-k} \sum_{k+1 < i \leq n} \hat{u}_i^2 \right\},$$

²⁹One of the main complications in this type of models arises in the case when there are an unknown number of change points. However, the problem is easier when (like in our case) there is prior theory suggesting the number of breaking points. In our case, one or at most two change points is a sensible possibility. To test that our method does not lead to biases in the location of the breaking points, we carry out tests to discard the existence of a second breaking point.

³⁰Gombay et al. (1996) have developed a test to detect a possible change in the variance of independent observations, a framework which has been extended by Orasch (2004) to detect multiple changes. The tabulation of the asymptotic distributions can be found in Orasch and Pouliot (2004). Pouliot (2008) brings it to a linear regression context.

where \hat{u}_i are the estimated residuals from a linear regression and n is the number of observations, can be used to detect the location of an unknown breaking point. The statistic $M(\tilde{k})$ can be used to test the null hypothesis of no breaking point. Orasch (1999) shows that this can be extended to test for the existence and detect the location of at most two breaking points, located in

$$(\tilde{k}_1, \tilde{k}_2) \equiv \arg \max_{k_1 \leq k_2} N(k_1, k_2) \text{ where } N(k_1, k_2) \equiv [(n - k_1)M(k_2) + k_2M(k_1)] / n.$$

The statistic $N(\tilde{k}_1, \tilde{k}_2)$ can be used to test the null hypothesis of no breaking point.

We run the test for the average prices when $\bar{Q} = 180$, $\bar{Q} = 240$, and $\bar{Q} = 280$, with $\gamma = 0.005$ and $\delta = 0.50$. Orasch's (1999) estimator rejects the null hypothesis of the non-existence of any breaking point versus at-most-two breaking points. The statistic $N(\tilde{k}_1, \tilde{k}_2)$ is equal to 1.72, 2.013 and 1.94, respectively, and the critical value is 1.54. The estimated breaking points $(\tilde{k}_1, \tilde{k}_2)$ are (19, 20), (41, 42) and (1, 14). As we can see, they are all very close to our predicted pivotal values, 19, 39 and 12. The only exception is \tilde{k}_1 in the last case. As argued by Mason and Schuenemeyer (1983), however, change point statistics exhibit poor sensitivity to deviations that may occur in the tails, so that \tilde{k}_2 is the meaningful change point in practice.

We are now going to show (1) that \tilde{k}_1 is statistically equal to \tilde{k}_2 in the first two cases and therefore there is only one breaking point, and (2) that that they are all statistically equal to the predicted theoretical value. Suppose that this is true. Then, as shown by Orasch (1999), \tilde{k}_1 is equal to the estimator of the at-most-one breaking point case, \hat{k} . Gombay et al. (1995) shows that the asymptotic distribution of

$$\frac{\delta^2}{\gamma^2}(\hat{k} - k_v),$$

where γ is the fourth moment of u_i and δ is an arbitrarily chosen function such that $\delta(n) \rightarrow 0$ as $n \rightarrow \infty$, has a density function equal to

$$f(t) = \begin{cases} h(-t, 1 - k_v/n, k_v/n) & \text{if } t \leq 0 \\ h(-t, k_v/n, 1 - k_v/n) & \text{if } t > 0 \end{cases}$$

where

$$h(t, x, y) = 2x(x + 2y)\{1 - \Phi((x + 2y)t^{1/2}) \exp(2y(x + y)t) - 2x^2(1 - \Phi(xt^{1/2}))\}.$$

We can now derive the confidence intervals for the three \tilde{k}_1 . If we take $\delta(n) = c/\sqrt{n}$ for a given constant $c > 0$, then the 95% confidence intervals are (14.82, 20.19), (38.96, 51.18) and (11.45, 14.19) for $c = 200$, $c = 50$ and $c = 100$, respectively.³¹ As we can see even for a large constant c (and therefore for small intervals), the theoretical and estimated values for the second breaking point fall within the intervals.

To summarise, our econometrical examination supports for the theoretical hypotheses: (1) the dynamics pre- and post-breaking point result in nonlinearities in the influence of diversification on market prices, (2) a breaking point occurs at the switching point, when the number of pivotal plants changes, (3) in spare capacity cases prices drop at the breaking point and (4) when capacity is tight, prices increase at the breaking point.

³¹As usual, the intervals are such that $[\tilde{k}_1 - z_i^+ \gamma^2 / (c/\sqrt{n})^2, \tilde{k}_1 - z_i^- \gamma^2 / (c/\sqrt{n})^2]$ where, in the first case, $z_i^- = -6$, $z_i^+ = 21$, $\gamma^2 = 155.96$, $n = 51$, $\tilde{k}_1 = 19$ and $c = 200$.

5.5 Diversification and latent intensity of competition

In a simulation environment, it is possible to inspect the probability priors from which bids are chosen. It is therefore possible to study how market structures (excess demand, generation diversification, etc.) influence the firms’ “competitive attitude” and not only market outcomes.³² Through their trading interaction and the R-E algorithm, firms learn to prioritise those bidding strategies that achieve higher payoffs and choose them more often. Price regularities follow once marginal supply patterns are established.

The panels in Figure 7 depict the end-of-simulation individual latent probability distributions from which firms choose bids. Figure 7 summarises the probabilities under $\bar{Q} = 180$, and $\bar{Q} = 280$ for specialisation ($\alpha = 0$), diversification ($\alpha = 0.5$) and at the breaking points ($\alpha_v = 0.20$ and $\alpha_v = 0.13$, respectively), averaged across the 50 simulation runs for $\delta = 0.5$ and $\gamma = 0.005$. On the horizontal axes, actions are identified with numbers ranging from 1, for the more competitive, to 100, for the highest possible bid. Cumulative probabilities are calculated on the vertical axes for each element of the action space.

From its definition we know that the probabilities’ concentration is largely invariant beyond the convergence point, so that its distributions are a good approximation to the plants’ long-term mixed strategies. Probabilities concentrated on higher actions result in the plants bidding less competitively, and viceversa. Curve movements to the upper-left and lower-right corners suggest that the market becomes more and less competitive, respectively. More linear distributions do not mean that there has been no learning, only that agents have learnt that they are better off with a mixed strategy in which all actions have similar probabilities.

<<FIGURE 7: LATENT INTENSITIES OF COMPETITION>>

Figure 7 offers a number of general insights linking individual probability distributions to market outcomes. Peak-load plants’ distributions stochastically dominate the base-load distributions, which means that peak-load bids are less competitive than those from base-load plants.³³ Under $\alpha = 0$, bids are lower for the base-load specialist. However, when $\alpha = 0.50$, bidding priors are very similar for each plant type but different across types. Thus, there is a clear identification between generation technology and the competitiveness of the plant’s trading prior.

Moreover, we find evidence of a link between portfolio diversification, learning, trading behaviour and market outcomes. For a given α , a demand increase from $\bar{Q} = 180$ to $\bar{Q} = 280$ makes the bids less competitive (i.e. lower-right movements). Trading priors shift in the competitive direction (upper left) when diversification grows from $\alpha = 0$ to α_v with resulting lower bids and, hence, lower market prices. When the movement is from α_v to $\alpha = 0.5$, though, the curves move to the lower-right corner, which suggests a less competitive attitude and higher prices. The latent intensity of competition figures provide support for the theoretical mechanisms underpinning the hypotheses.

³²We have also analysed latent competition intensities in Rupérez Micola and Bunn (2008) and Rupérez Micola et al. (2008).

³³In particular, the expected action of a peak-load generator is higher than for a base-load generator. This will translate into even higher expected bids since the same action s implies a higher bid for a peak-load generator than for a base-load generator. Notice that unless the probabilities of playing a high action with the low-capacity plant and a low action with the high-capacity plant are reduced to zero, there will be some instances in which an agent might bid higher with the low-cost plant, although this will occur rarely.

6 Robustness

In this section we carry out robustness checks against an alternative supply function bidding procedure, as well as cases in which the marginal costs of the low-cost plants are positive and in which there is some short-term demand elasticity.

6.1 Supply function bidding

One restriction of the vdF-H implementation is that firms are allowed to submit a single-price bid per plant. This might be a reasonable approximation in the case of nuclear power. Inherent inflexibilities in the operation of such plants (e.g. safety concerns, very low marginal costs and high start-up and loss of volume costs) prompt them to submit flat schedules. However, the restriction is quite unrealistic when modelling high-cost units such as thermal plants, which typically have many increasing steps. In this section, we check the robustness of our results to an alternative model in which firms submit single-price bids for their base-load capacity and bid with increasing supply functions for their peak-load. This model takes into account how technological diversification affects the “expressivity” of bid functions.³⁴

The only changes with respect to the vdF-H implementation relate to the definition of supply function schedules and market clearing. Now, generators submit supply schedules for each peak-load plant according to a probability distribution over the set of possible angles, equally spaced between the minimum (zero degrees) and the maximum (ninety degrees or $\pi/2$ radians). Each schedule corresponds to a line between the marginal cost c and

$$B^h = \left\{ c + \frac{\sin(s(\pi/2)/S)}{\cos(s(\pi/2)/S)} k_i^h \mid s = 1, \dots, S \right\}, \quad (7)$$

bounded between zero and Ψ . The “action s ” is the variable to reinforce and represents the angle of the plant’s linear supply schedule. Higher actions represent schedules with increasing slopes, rotating around the marginal cost point. Schedules generated from lower actions are more competitive, i.e. flatter. Base-load plants bid flat schedules as in the basic model ($B^l = \{s(\Psi/S) \mid s = 1, \dots, S\}$).

We run the simulation for the same α , \bar{Q} , γ and δ parameters, and obtain a new data set with 321,300 observations. The left-hand block in Table 4 summarises the Fama-MacBeth regression estimates testing the existence of a breaking point at α_v for each \bar{Q} in this setting. In all fourteen \bar{Q} cases, at least one of the two coefficients is significant at the .01 level, and there is only one non-significant coefficient at standard confidence levels (slope coefficient for $\bar{Q} = 290$). These results suggest that the relationship between diversification and prices not only appears when firms bid as in vdF-H but also when peak-load firms submit more expressive supply functions.

<TABLE 4: ROBUSTNESS OF REGIME SWITCHING POINT ESTIMATES TO SUPPLY BIDDING AND $c_l=50$ >

³⁴By expressivity we mean the firm’s technological flexibility to submit supply schedules with varying slopes. We thank the Associate Editor for providing us with this term.

6.2 Positive base-load marginal costs

Another restriction of our implementation is that base-load costs are equal to zero. This is clearly a simplification: direct costs – such as fuel costs – are always different from zero and, in this section, we check the robustness of our results to the assumption. To be conservative, we have chosen a relatively high base-load cost ($c_l = 50$), corresponding to one-half of the peak-load cost ($c_h = c = 100$).

We run simulations for the α , \bar{Q} , γ and δ parameters and obtain a new data set consisting of 321,300 observations. The right hand side block in Table 4 summarises the Fama-MacBeth estimates testing the existence of a breaking point at α_v for each \bar{Q} and the alternative cost value. In all fourteen \bar{Q} cases at least one of the two coefficients is significant at the .01 level. The relationship between diversification and prices also seems to appear when there are positive base-load costs.

6.3 Demand elasticity

We have restricted our main implementation to have fully inelastic demand. The literature has established the extremely low price-elasticity of short-term demand, originating among others from the lack of real-time metering (e.g. Stoft, 2002). However, several electricity market features might effectively reduce demand when prices rise. Large industrial users might be able to turn on back-up generation assets if prices move beyond a certain threshold, and even the residential demand elasticity might be larger than zero in some circumstances. Moreover, the existence of forward contracts is economically similar to an increase in demand elasticity (Bushnell, 2007).

We examine the robustness of our results to a setting with some demand elasticity. We run simulations in a linear market demand case ($P_t = 2400 - 8Q_t$) implying these two bounds: $P = 0$ if and only if $Q = 300$ and $P = 200$ if and only if $Q = 275$. We run simulations for the usual α , γ and δ parameters, and obtain a data set with 36,720 observations. Figure 8 shows the average best-fit regressions for all nine R-E parameter combinations, which clearly present a breaking point around $\hat{\alpha} = 0.17$. Interpolating, $\hat{\alpha}$ is equivalent to the breaking point one would expect in the case of inelastic demand $\bar{Q} = 275$. Prices are also similar to the ones obtained for inelastic demands of $\bar{Q} = 270$ and $\bar{Q} = 280$. The Fama-MacBeth estimates are statistically significant.

<<FIGURE 8: AVERAGE BEST FIT REGRESSIONS WITH ELASTIC DEMAND>>

7 Conclusion

The main research question in this paper concerns the shape of the diversification versus market price relationship. Theoretical and computational results suggest that the relationship is not monotonic but that it includes structural stability breaks. The correspondence between close-form and simulation switching points was very strong.

The findings might help explain some features of the electricity time-series, e.g. high volatility and seasonality. Low- and high-demand periods might not only lead to more or less supply competition but also change its nature. When demand is low there are no pivotal plants but in high-demand cases there are two of them. Hence, exogenous

weather patterns might explain changes in the firms' trading behaviour, which in turn might contribute to cause the dramatic regime switching observed in econometric studies of electricity prices (for example, Karakatsani and Bunn, 2004).

Further, our paper contributes to clarifying the influence of generation technologies with a characterisation of the role played by pivotal players. The composition of a firm's generation portfolio is a market power instrument because it modifies pivotal dynamics and changes the intensity of competition. In low-demand cases, there is a regime-switching point of the diversification level where the market moves from one to no pivotal plants, there is a sudden loss of market power and prices drop. In high-demand cases, there is a regime-switching point where the market setting changes from one to two pivotal players, there is an increase in the new pivotal plant market power which facilitates some implicit cooperation and prices increase.

The robustness tests suggest that more expressive bidding functions, positive costs and some demand elasticity do not qualitatively change these findings. The analysis, however, relies on a number of additional assumptions. First, they stem from the R-E algorithm, which is only one of the models one could use. R-E reinforcement learning is shown to be a fruitful alternative where standard theoretical methods are impractical. Moreover, where there are unique theoretical predictions (e.g. switching points), R-E simulations match them well. However, alternative behavioural models (e.g. Day and Bunn, 1999; García et al., 2005) and experiments can still further our understanding of real-life electricity markets. Second, the various simulation parameters – including number of firms, technology stocks, etc. – were defined as exogenous and independent of one another. It is possible that in reality they would be endogenously determined, and complex simulations might contribute to study their reciprocal dynamics (e.g. Bunn and Oliveira, 2007). Here, we have focused on a stylised, relatively standard market model, whose outcomes are more directly comparable to those of close-form approaches.

References

- [1] Arellano, S. and P. Serra (2007): “A Model of Market Power in Electricity Industries subject to Peak Load Pricing”, *Energy Policy*, vol. 25, pp. 5130-5135.
- [2] Ausubel, L. and P. Cramton (2002): “Demand Reduction and Inefficiency in Multi-Unit Auctions.” Working paper, Department of Economics, University of Maryland.
- [3] Baldick, R., S. Kolos and S. Tompaidis (2006): “Interruptible Electricity Contracts from an Electricity Retailer's Point of View: Valuation and Optimal Interruption”. *Operations Research*, vol. 54 (4).
- [4] Baldick, R., Ross and W. Hogan (2002) “Capacity Constrained Supply Function Equilibrium Models of Electricity Markets: Stability, Non-decreasing constraints, and Function Space Iterations”. POWER Working Paper PWP-089, University of California Energy Institute.

- [5] Beggs, A. W. (2005): "On the convergence of reinforcement learning", *Journal of Economic Theory*, vol. 122, no. 1, pp.1-36
- [6] Bolle, F. (1992): "Supply function equilibria and the danger of tacit collusion The case of spot markets for electricity". *Energy Economics*, vol. 14, no. 2, pp. 94-102.
- [7] Borenstein, S., J. Bushnell, E. Kahn and S. Stoft (1995): "Market Power in California Electricity Markets". *Utilities Policy* vol 5 (3-4), pp. 219-236.
- [8] Borenstein, S., J. Bushnell and F. Wolak (2002): "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market". *American Economic Review*, vol. 92 (5), p.p. 1376-1405.
- [9] British Energy (2006): Annual Report. Downloaded from: <http://www.british-energy.com/pagetemplate.php?pid=324>
- [10] Bunn, D.W. and F. Oliveira (2007): "Agent-based Analysis of Technological Diversification and Specialization in Electricity Markets". *European Journal of Operations Research*.
- [11] Bushnell, J. (2003): "A Mixed Complementarity Model of Hydrothermal Electricity Competition in the Western United States". *Operations Research*, vol. 51, No. 1, January-February 2003, pp. 80-93.
- [12] Bushnell, J. (2007): "Oligopoly equilibria in electricity contract markets", *Journal of Regulatory Economics*, vol. 32, no.3, pp 225-245
- [13] Camerer, C. and T. H. Ho (1999): "Experience-weighted Attraction Learning in Normal Form Games". *Econometrica*, vol.67, issue 4, pp. 827 - 874.
- [14] Cooper, D. J., D. DeJong, B. Forsythe, T. Ross (1990): "Alternative Institutions for Resolving Coordination Problems: Experimental Evidence on Forward Induction and Preplay Communication", in J. Friedman (Ed.), *Problems of Coordination in Economic Activity*. Dordrecht: Kluwer.
- [15] Crawford, G. S., J. Crespo and H. V. Tauchen (2006): "Bidding Asymmetries in Multi-Unit Auctions: Implications of Bid Function Equilibria in the British Spot Market for Electricity". Working paper.
- [16] Crawford, V. P., U. Gneezy and Y. Rottenstreich, (2008): "The Power of Focal Points is Limited: Even Minor Payoff Asymmetry yields Massive Coordination Failures", *American Economic Review*, Vol. 98, pp. 1443-1458.
- [17] Day, C. and D. Bunn (1999): "Divestiture of Generation Assets in the Electricity Pool of England and Wales: A Computational Approach to Analyzing Market Power". *Journal of Regulatory Economics*, Vol. 19, Number 2 / March, 2001, pp. 123-141.
- [18] Entriken, R. and S. Wan (2005): "Agent-Based Simulation of an Automatic Mitigation Procedure" Proceedings of the 38th Annual Hawaii International Conference on System Sciences.

- [19] Erev, I., Haruvy, E. (2001): "Variable Pricing: a Customer Learning Perspective", working paper Columbia Business School.
- [20] Fabra, N. and J. Toro (2005): "Price Wars and Collusion in the Spanish Electricity Market", *International Journal of Industrial Organization*, vol. 23, pp. 155-181.
- [21] Fama, E.F. and J. D. MacBeth (1973): "Risk, Return, and Equilibrium: Empirical Tests". *Journal of Political Economy*, vol. 81, no. 3, pp. 607-636.
- [22] Fudenberg, D. and D. Levine (1998): *The Theory of Learning in Games*, MIT Press.
- [23] Garcia, A., E. Campos-Nañez and Reitzes, J. (2005): "Dynamic Pricing and Learning in Electricity Markets". *Operations Research*, Vol. 53, No. 2, March–April 2005, pp. 231–241.
- [24] Genc, T. and S. S. Reynolds (2005): "Supply Function Equilibria with Pivotal Electricity Suppliers". Working paper, University of Guelph.
- [25] Gombay, E., Horvath, L. and M. Huskova (1996): "Estimators and Tests for Change in Variances", *Statistics & Decisions*, vol. 14, pp. 145-159.
- [26] Green, R. and D. Newbery (1992): "Competition in the British Electricity Spot Market". *The Journal of Political Economy*, vol. 100, pp. 929-953.
- [27] Haruvy, E. and D. O. Stahl (2004): "Deductive versus inductive equilibrium selection: experimental results". *Journal of Economic Behaviour and Organization*, vol. 53, pp. 319-331.
- [28] Hobbs, B. and J. S. Pang (2007): "Nash-Cournot Equilibria in Electric Power Markets with Piewise Linear Demand Functions and Joint Constraints". *Operations Research*, vol. 55, no. 1, pp. 113-127.
- [29] Holmberg, P. (2007): "Supply Function Equilibrium with Asymmetric Capacities and Constant Marginal Costs", *Energy Journal*, vol. 28, no. 2, pp. 55-82,.
- [30] Holmberg, P. (2008): "Unique Supply Function Equilibrium with Capacity Constraints", *Energy Economics*, vol. 30, no 1, pp. 148-172.
- [31] Hopkins, E. (2008): "Adaptive Learning Models of Consumer Behavior", *Journal of Economic Behaviour and Organization*, forthcoming.
- [32] Iriberry, N. (2006): "The Dynamics of Coordination Rates in Experimental Battle of the Sexes Game", working paper University of California San Diego.
- [33] Joskow, P. and Edward Kahn (2002): "A Quantitative Analysis of Pricing Behavior in California's Wholesale Electricity Market During Summer 2000". *The Energy Journal*, vol 23, no. 4.

- [34] Karakatsani, N. and D. Bunn (2004): “Modelling Stochastic Volatility in High-frequency Spot Electricity Prices”. Energy Markets Group Working Paper, London Business School.
- [35] Klemperer, P. and Margaret Meyer (1989): “Supply Function Equilibria in Oligopoly under Uncertainty”. *Econometrica*, vol. 57, pp. 1243-1277.
- [36] Mansur, E. (2007): “Upstream Competition and Vertical Integration in Electricity Markets”. *Journal of Law and Economics*, vol. 50, issue 1, pp. 125-156.
- [37] Mason, D. M. and J. H. Schuenemeyer (1983): “A Modified Kolmogorov-Smirnov Test Sensitive to Tail Alternatives”. *The Annals of Statistics*, Vol. 11, No. 3.
- [38] Newbery, D. (1991): “Capacity-constrained Supply Function Equilibria: Competition and Entry in the Electricity Spot Market”, mimeo University of Cambridge.
- [39] Newbery, D. (1997): “Privatisation and Liberalisation of Network Utilities”, *European Economic Review*, vol. 41, pp. 357-383.
- [40] Nicolaisen, J., V. Petrov and L. Tesfatsion (2001): “Market Power and Efficiency in a Computational Electricity Market with Discriminatory Double-Auction”. *IEEE Transactions on Evolutionary Computation*, vol. 5, no. 5, pp. 504–523.
- [41] Orasch, M.R. (1999): “On Detecting Multiple Changes”. Laboratory for Research in Statistics and Probability, Carleton University and University of Ottawa.
- [42] Orasch, M.R. and W. Pouliot (2004): “Tabulating Weighted Sup-norm Functionals Used in Change-point Analysis”, *Journal of Statistical Computation and Simulation*, Vol 74, pp. 249-276.
- [43] Perekhodtsev, D., L. B. Lave, S. Blumsack (2002): “The Model of Pivotal Oligopoly Applied to Electricity Markets”. Carnegie Mellon Electricity Industry Center, CEIC Working Paper 02-06
- [44] PG&E Corporation (2006): Annual Report.
Downloaded from: <http://www.pgecorp.com/investors/pdfs/2006AnnualReport.pdf>
- [45] Pouliot, W. (2008): “Multiple Change-Point Detection in Linear Regression Models via U-Statistic Based Process”. mimeo.
- [46] Puller, s. (2007): “Pricing and Firm Conduct in California’s Deregulated Electricity Market”, *Review of Economics and Statistics*, Vol. 89 (1), pp. 75-87.
- [47] Roth, A. and I. Erev (1995): Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. *Games and Economic Behavior*, Vol 8, pp. 164–212.

- [48] Rudkevich, A., M. Duckworth, and R. Rosen (1998): "Modeling Electricity Pricing in a Deregulated Generation Industry: The Potential for Oligopoly Pricing in a Poolco", *Energy Journal*, vol. 19(3), pp. 19-48.
- [49] Rupérez Micola, A. and D. W. Bunn (2008): "Crossholdings, information and prices in capacity constrained sealed bid-offer auctions", *Journal of Economic Behavior and Organization*, vol. 66, issue 3-4, pp. 748—766.
- [50] Rupérez Micola, A. A. Banal-Estanol and D. W. Bunn (2008): "Incentives and coordination in vertically related energy markets", *Journal of Economic Behavior and Organization*, vol. 67, issue 2, pp. 381—393.
- [51] Sandholm, W. (2002): "Evolutionary Implementation and Congestion Pricing", *Review of Economic Studies*, vol. 69, pp. 667-689.
- [52] Stoft, S. (2002): *Power System Economics: Designing Markets for Electricity*. Wiley. ISBN: 978-0471150404.
- [53] Sun, J. and L. Tesfatsion (2007): "Dynamic Testing of Wholesale Power Market Designs: An Open-Source Agent-Based Framework". *Computational Economics* 30 (3), 291–327.
- [54] Van Huyck, J.B., Raymond C. Battalio, Richard O. Beil (1990): "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure", *The American Economic Review*, Vol. 80, no. 1, pp. 234-248
- [55] Von der Fehr, N. and D. Harbord (1993): "Spot Market Competition in the UK Electricity Industry". *Economic Journal*.
- [56] Von der Fehr, N. and D. Harbord (1998): "Competition in Electricity Spot Markets Economic Theory and International Experience", University of Oslo Memorandum in Economics, 5.
- [57] Veit, D., A. Weidlich, J. Yao and S. Oren (2006): "Simulating the Dynamics in Two-Settlement Electricity Markets via an Agent-Based Approach". *International Journal of Management Science and Engineering Management*, vol. 1, pp. 83-97.
- [58] Wolak, F. A. and R. H. Patrick (1997): "The Impact of Market Rules and Market Structure on the Price Determination Process in the England and Wales Electricity Market.". POWER Working Paper PWP-047, University of California Energy Institute.
- [59] Wolfram, C. (1998): "Strategic Bidding in a Multiunit Auction: An Empirical Analysis of Bids to Supply Electricity in England and Wales". *RAND Journal of Economics*, Winter 1998, 29(4), pp. 703–25.
- [60] Wolfram, C. (1999): "Measuring Duopoly Power in the British Electricity Spot Market." *American Economic Review*, September 1999, 89(4), pp. 805–26.

TABLE 1: Expected pivotal switching points for each demand level

Demand	Switching point	Capacity Low, 1	Capacity Highcost 1	Capacity Lowcost 2	Capacity Highcost 2	No. Pivotal plants post-threshold
160	0.07	140	10	10	140	0
170	0.13	130	20	20	130	0
180	0.20	120	30	30	120	0
190	0.27	110	40	40	110	0
200	0.33	100	50	50	100	0
210	0.40	90	60	60	90	0
220	0.47	80	70	70	80	0
230	0.47	80	70	70	80	2
240	0.40	90	60	60	90	2
250	0.33	100	50	50	100	2
260	0.27	110	40	40	110	2
270	0.20	120	30	30	120	2
280	0.13	130	20	20	130	2
290	0.07	140	10	10	140	2

TABLE 2: Fixed effects regression of prices on demand and diversification

Variable	Parameter	95% conf. Interval	
Intercept	103.25	103.15	103.36
Diversification	-0.0542	-0.0554	-0.053
Demand	0.2707	0.2702	0.2711
R ²	0.8219		

TABLE 3: Hypothesis tests for the breaking point estimates

Change parameters under supply function bidding							
Demand	Switching point	Beta1	Std(Beta1)	t-stat	Beta3	Std(Beta3)	t-stat
160	0.07	-8.28***	0.5	-49.31	-0.2***	0.12	-5.04
170	0.13	-9.43***	0.34	-83.33	-0.1***	0.03	-8.86
180	0.2	-13.96***	0.63	-66.18	0.03***	0.02	4.66
190	0.27	-15.83***	0.72	-65.79	0.07***	0.02	13.24
200	0.33	-1.81**	0.82	-6.6	-0.25***	0.02	-40.63
210	0.4	-13.26***	2.26	-17.63	-0.01	0.05	-0.55
220	0.47	-13.08***	4.04	-9.7	-0.02	0.08	-0.61
230	0.47	4.62**	7.60	1.81	0.06	0.15	1.22
240	0.4	-10.3***	1.62	-19.05	0.38***	0.04	29.56
250	0.33	0.01	0.87	0.05	0.17***	0.01	26.01
260	0.27	5.94***	0.97	18.33	0.03***	0.02	4.15
270	0.2	4.01***	0.78	15.32	0.07***	0.02	10.54
280	0.13	4.59***	0.85	16.05	0.06***	0.04	4.31
290	0.07	5.64***	0.65	26.03	0.00	0.07	-0.08

*** Significant at the .01 level; ** significant at the .05 level; * significant at the .10 level

TABLE 4: Robustness of regime switching point estimates to supply bidding and c = 50

Demand	Switching point	Change parameters under supply function bidding						Change parameters under c = 50					
		Beta1	Std(Beta1)	t-stat	Beta3	Std(Beta3)	t-stat	Beta1	Std(Beta1)	t-stat	Beta3	Std(Beta3)	t-stat
160	0.07	-5.44***	0.43	-37.57	0.77***	0.1	22.08	-6.17***	0.56	-32.77	-0.1***	0.1	-3.04
170	0.13	-12.17***	0.28	-127.46	0.85***	0.03	65.8	-7.43***	0.32	-70.53	-0.05***	0.03	-5.19
180	0.2	-18.74***	1.15	-48.48	0.88***	0.05	53.79	-11.78***	0.48	-73.34	0.07***	0.02	11.36
190	0.27	-22.63***	1.29	-52.31	0.81***	0.04	51.65	-13.99***	0.28	-150.79	0.12***	0.01	26.24
200	0.33	-23.91***	1.48	-48.18	0.70***	0.04	51.81	1.23***	0.68	5.48	-0.24***	0.02	-37.77
210	0.4	-21.79***	3.34	-19.56	0.53***	0.08	20.87	-7.02***	2.29	-9.22	-0.08***	0.05	-4.93
220	0.47	-20.17***	5.87	10.29	0.43***	0.12	10.8	-10.46***	5.86	-5.36	-0.03	0.12	-0.64
230	0.47	-15.41***	9.57	-4.83	0.31***	0.19	4.71	7.54***	6.44	3.51	-0.01	0.13	-0.25
240	0.4	-3.94***	1.88	-6.3	0.08***	0.04	5.29	-7.31***	1.53	-14.37	0.3***	0.03	27.95
250	0.33	-1.92***	0.97	-5.91	0.03***	0.02	3.26	2.38***	0.6	11.84	0.09***	0.01	19.39
260	0.27	-0.75***	0.36	-6.22	0.01*	0.01	1.71	7.54***	0.65	34.92	-0.04***	0.01	-8.84
270	0.2	-0.82***	0.19	-13.1	0.01***	0.01	3.34	5.5***	0.69	23.9	-0.01	0.03	-0.33
280	0.13	-0.71***	0.13	-15.74	0.01***	0.01	4.29	5.94***	0.73	24.58	-0.04***	0.04	-3.03
290	0.07	-0.54***	0.16	-10.15	0.02	0.05	1.15	6.5***	0.53	36.6	-0.08***	0.06	-3.9

*** Significant at the .01 level; ** significant at the .05 level; * significant at the .10 level

Figure 1: Nash equilibria for each α and Q when $K=300$

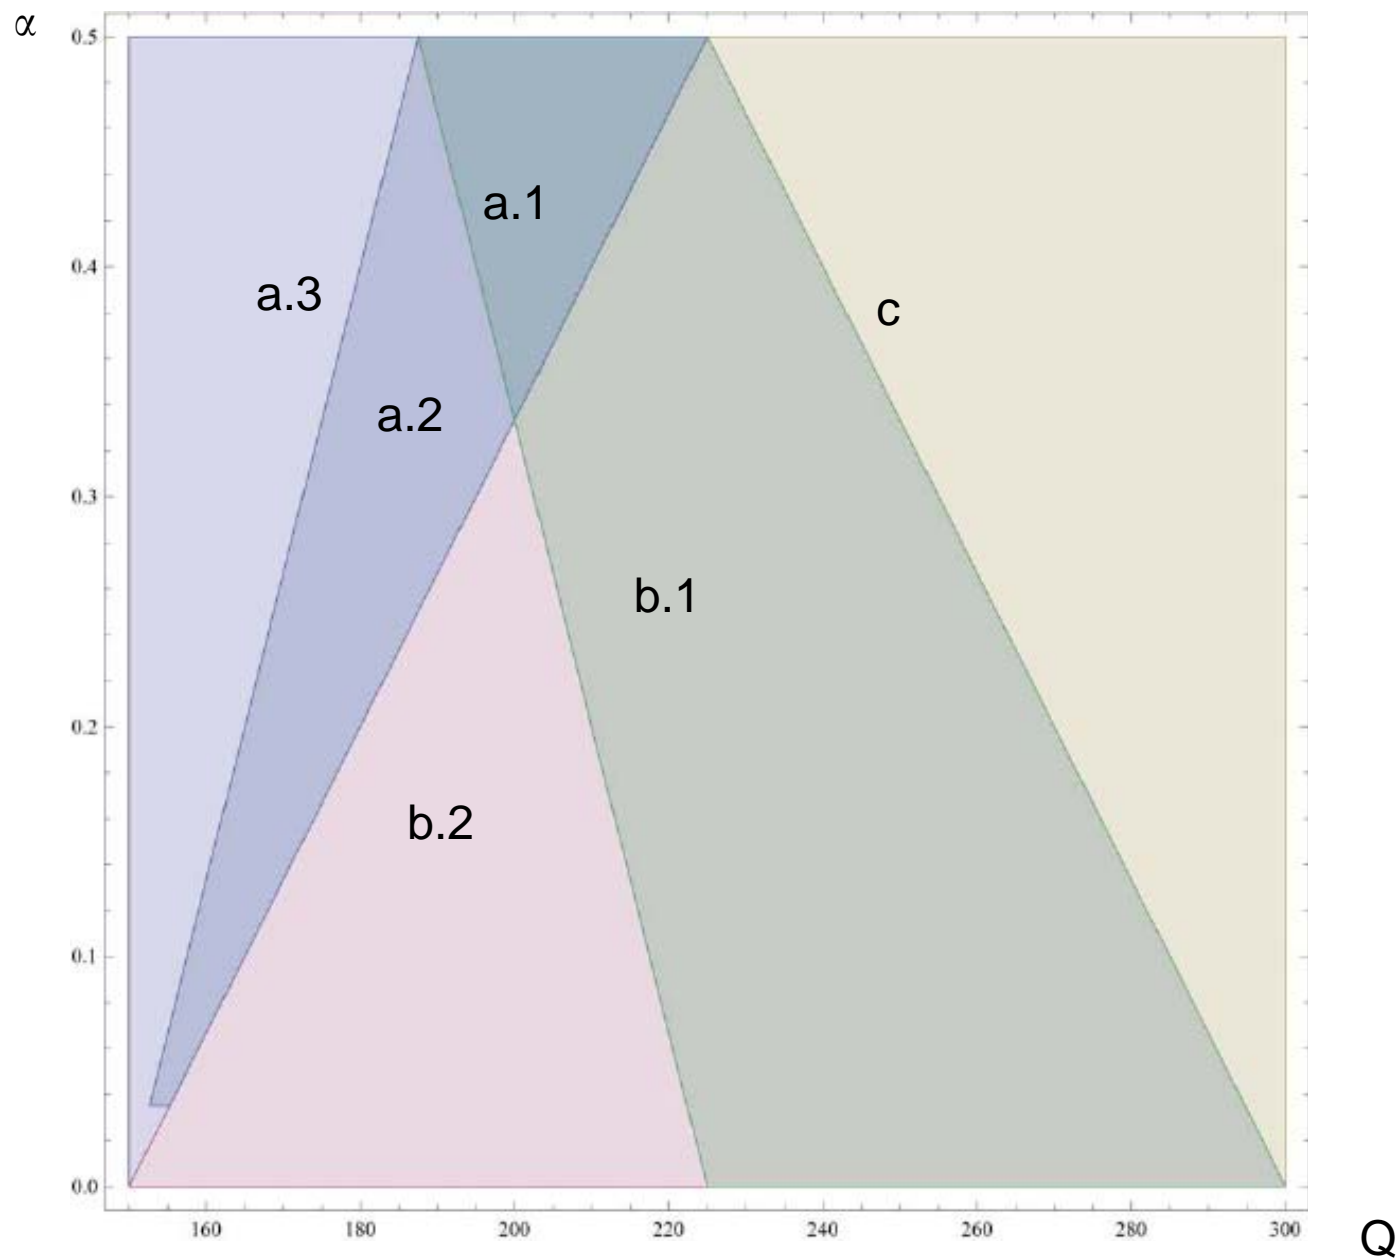


Figure 2: Mean plus and minus two standard deviation of prices (Q=240)




Figure 3: Mean plus and minus two standard deviation of prices (Q=180)





Figure 4: Mean plus and minus two standard deviation of prices (Q=280)



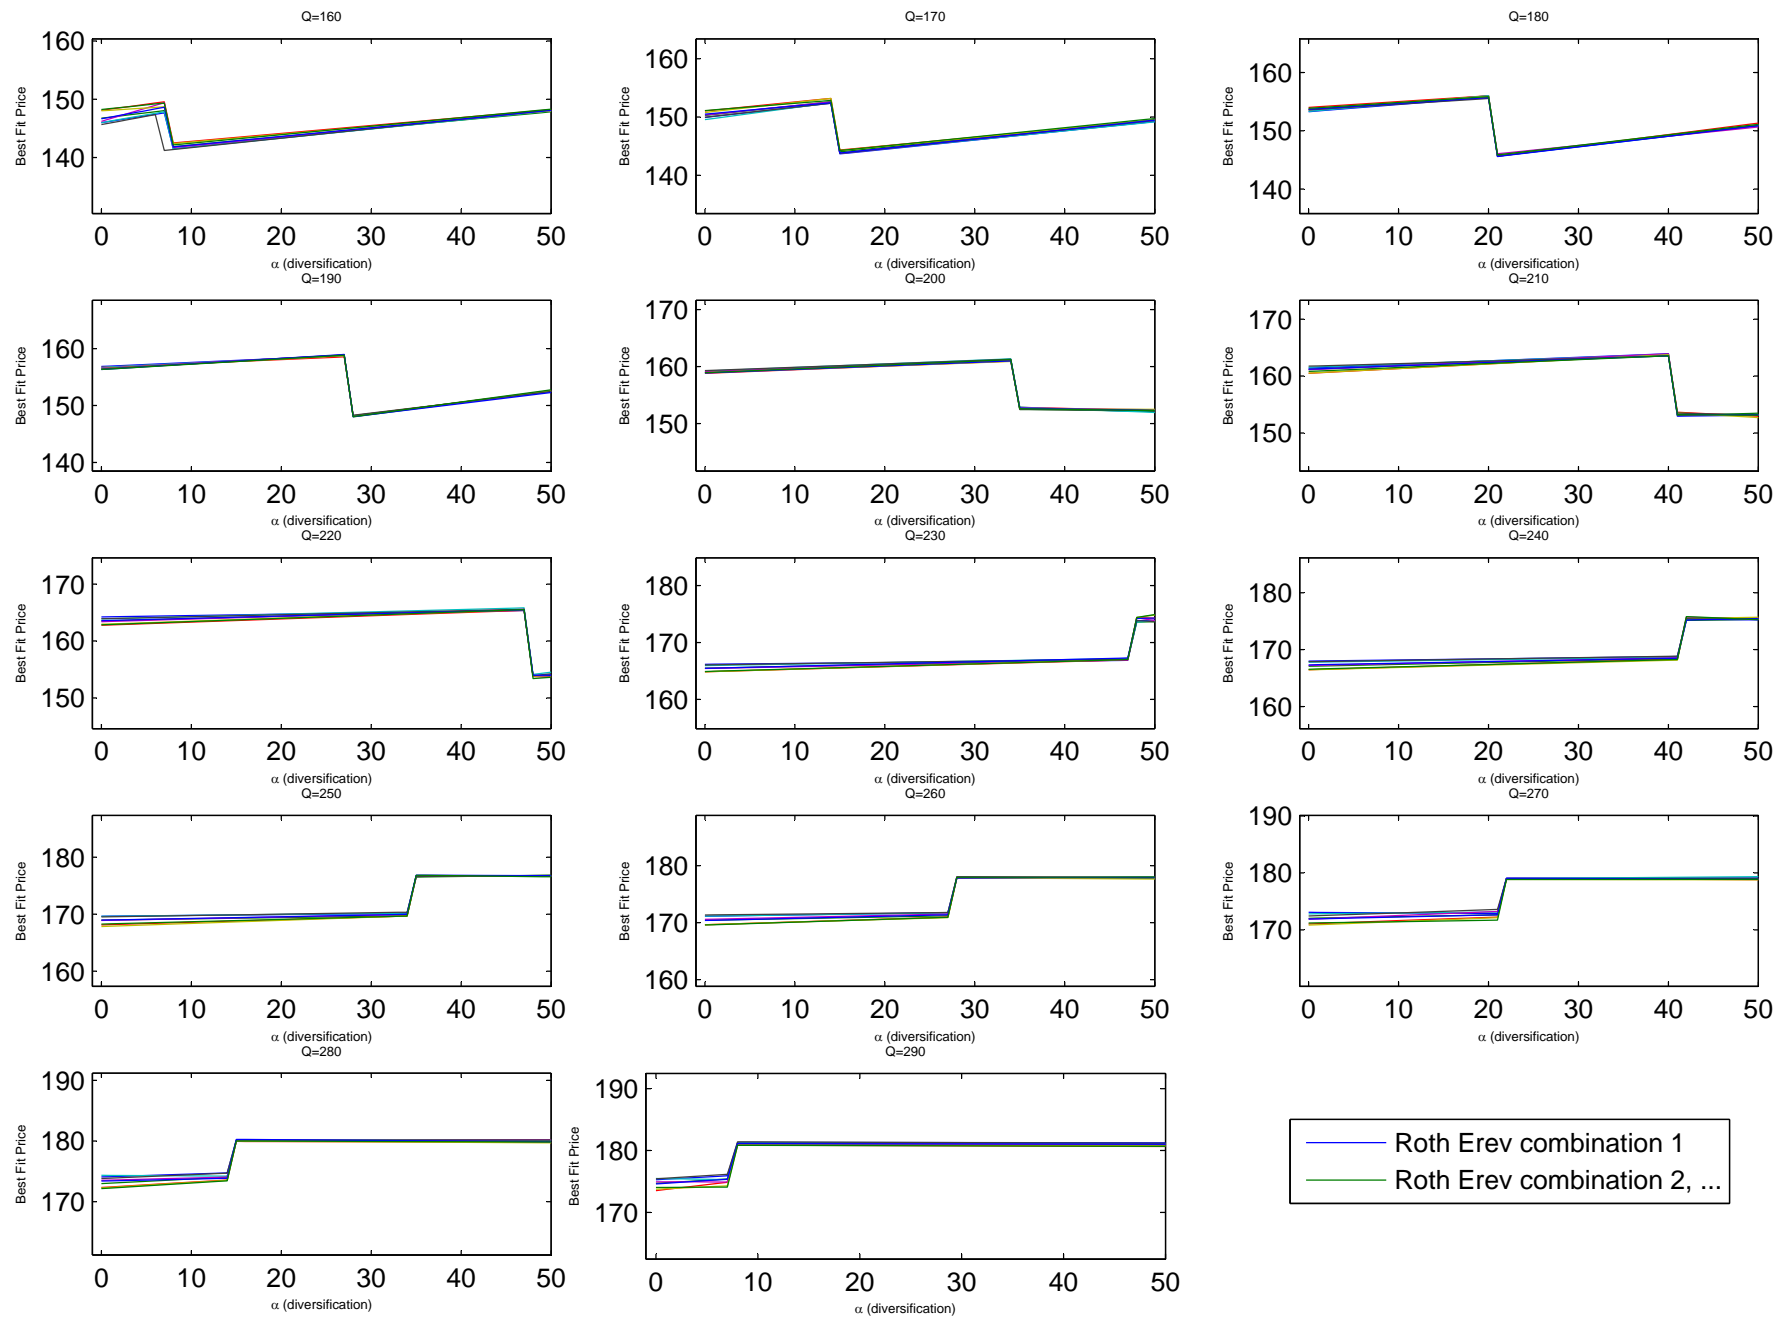


Figure 5: Average best fit regressions (over all samples)

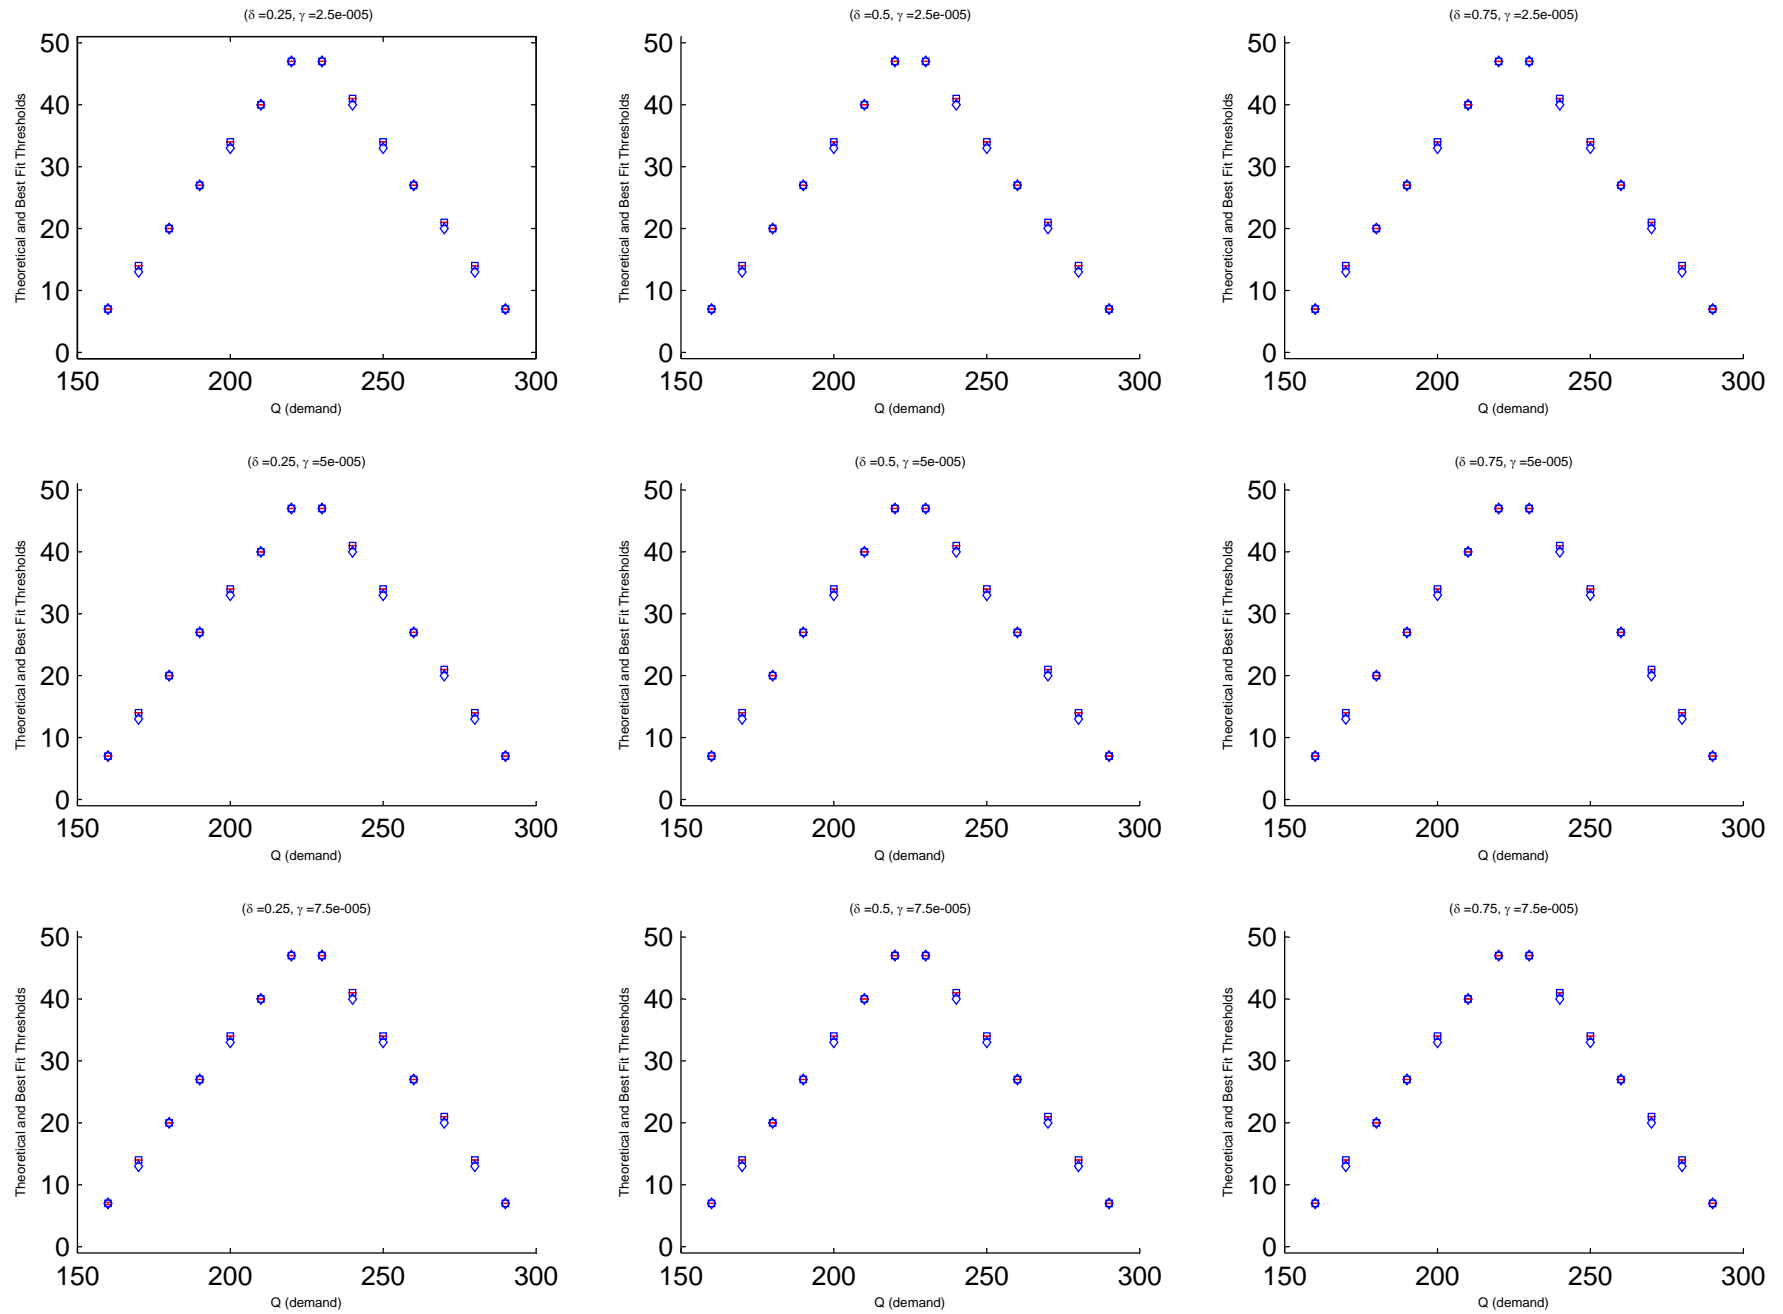
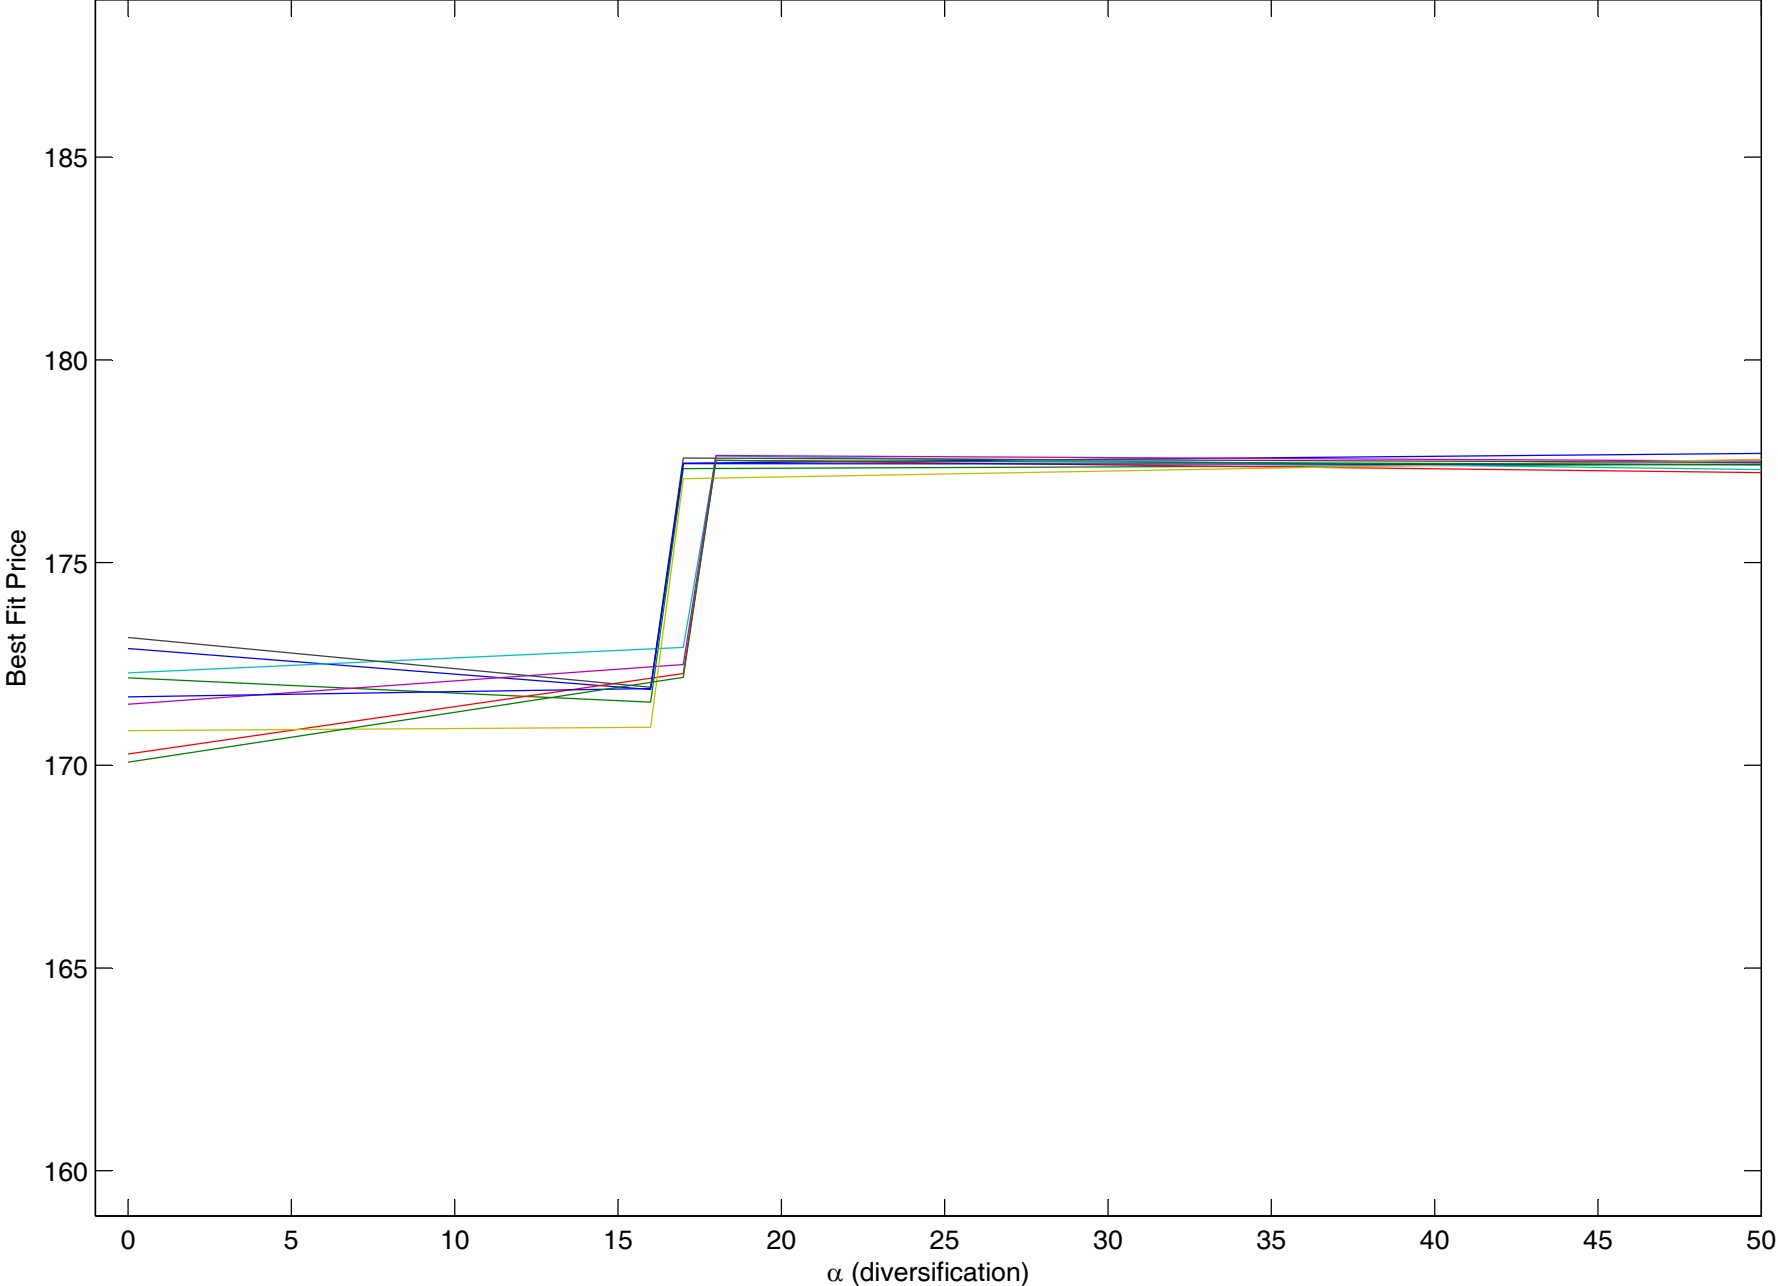


Figure 6: Mean and two std of simulated best-fit thresholds

Figure 7 Average Best Fit Regressions (over all samples) Elastic Demand



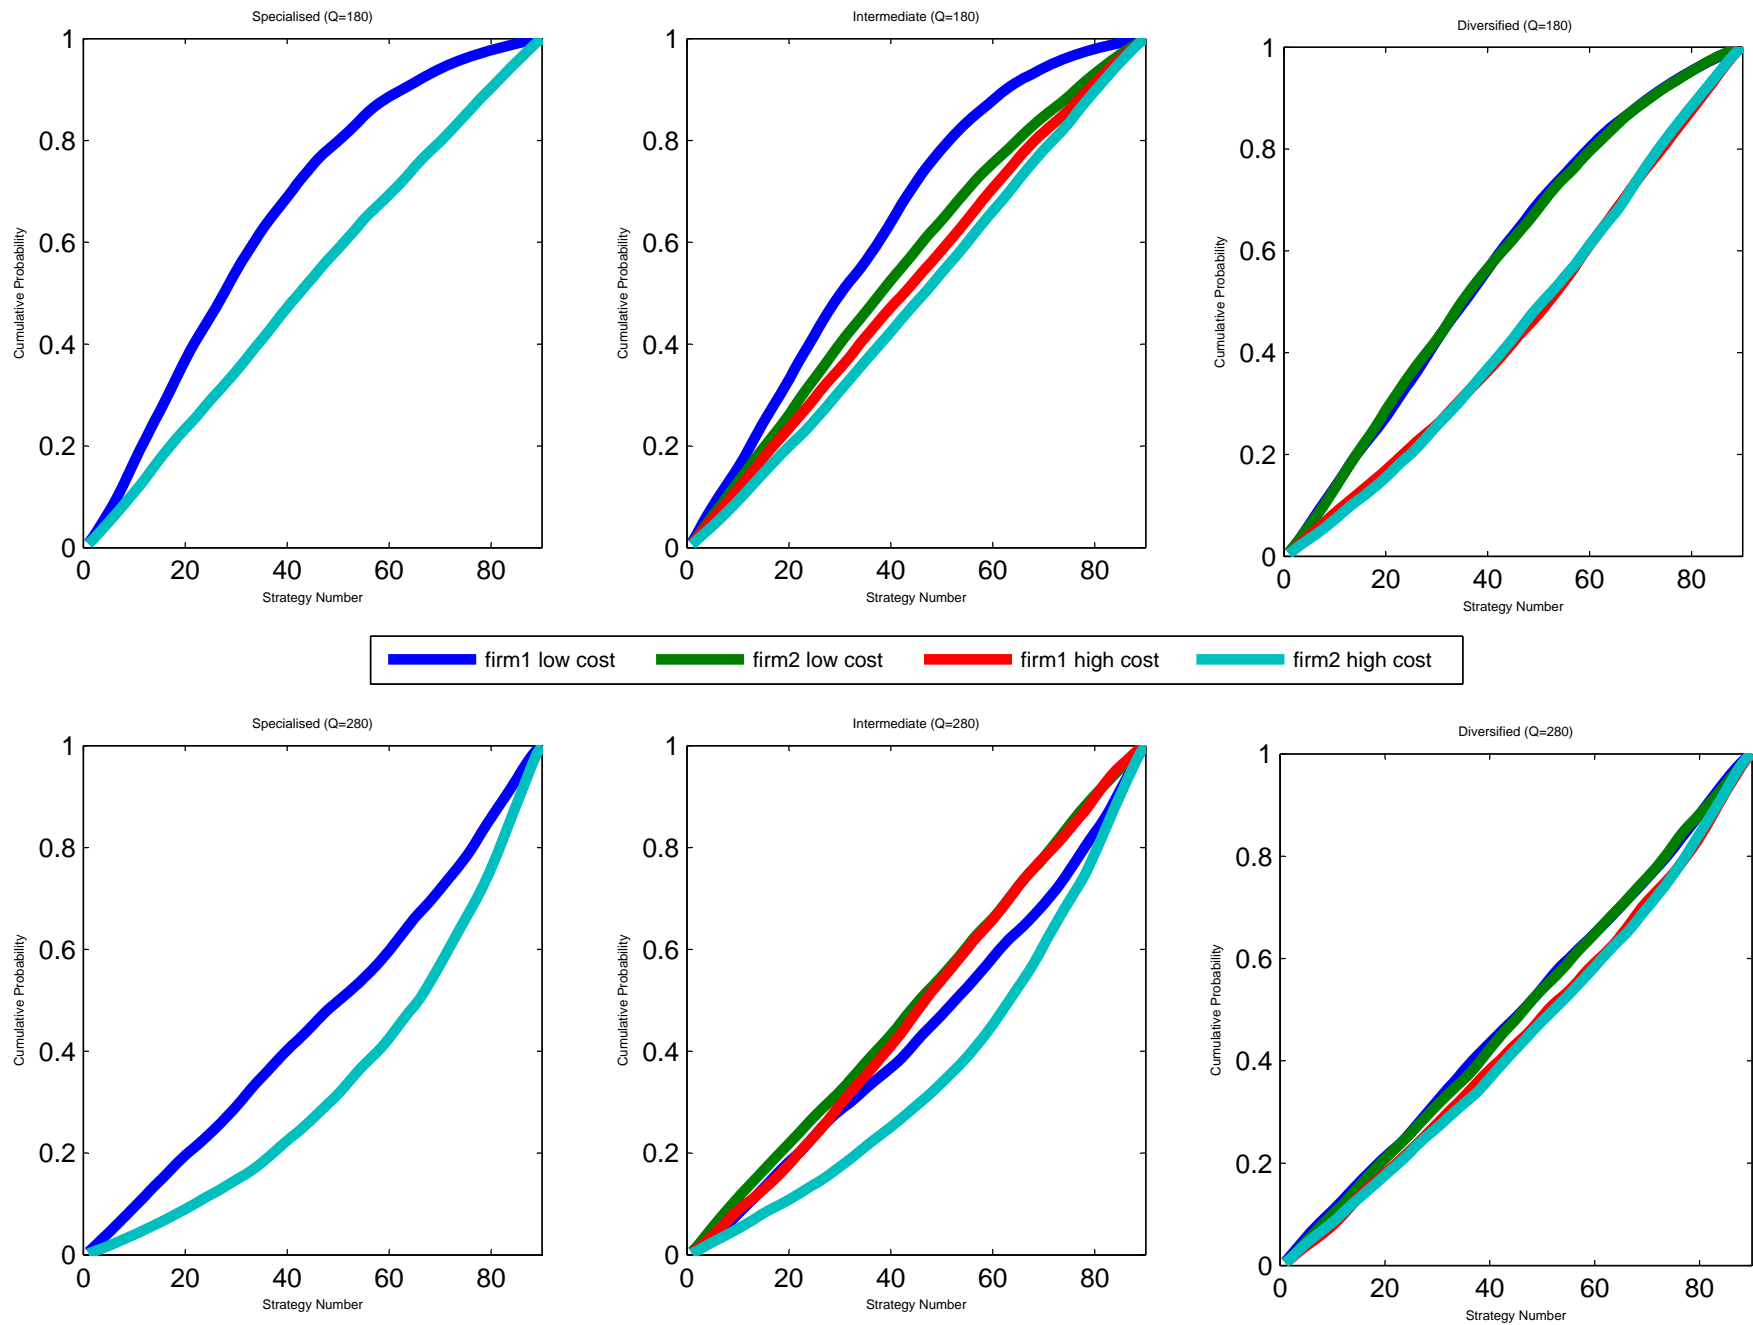


Figure 8: End-of-simulation cumulative distribution of probabilities