

Are agent-based simulations robust? The wholesale electricity trading case*

Albert Banal-Estanol[†]

Augusto Rupérez Micola[‡]

March 30, 2010

Abstract

Agent-based computational economics is becoming widely used in practice. This paper explores the consistency of some of its standard techniques. We focus in particular on prevailing wholesale electricity trading simulation methods. We include different supply and demand representations and propose the Experience-Weighted Attractions method to include several behavioural algorithms. We compare the results across assumptions and to economic theory predictions. The match is good under best-response and reinforcement learning but not under fictitious play. The simulations perform well under flat and upward-sloping supply bidding, and also for plausible demand elasticity assumptions. Learning is influenced by the number of bids per plant and the initial conditions. The overall conclusion is that agent-based simulation assumptions are far from innocuous. We link their performance to underlying features, and identify those that are better suited to model wholesale electricity markets.

Keywords: Agent-based computational economics, electricity, market design, experience-weighted attraction (EWA), learning, supply functions, demand aggregation, initial beliefs.

1 Introduction

Companies and government agencies increasingly make use of simulation techniques, as part of an “engineering” approach in which game theory, experimental and computational economics complement each other (Roth, 2002). Examples of leading corporations using simulations include eBay, Google, IBM and Unilever. Simulations have also been crucial for designing the US National Resident Matching Program (Roth and Peranson, 1997), radio spectrum auctions (Rothkopf et al., 1998), kidney exchange mechanisms (Roth et al. 2004), school choice programmes (Abdulkadiroglu et al., forth.) and the electricity trading arrangements in many countries.

Agent-based computational economics (ACE) is becoming an important part of those market simulation efforts. In the ACE paradigm, markets are modelled as dynamic systems of interacting,

*Versions of this paper have been presented at the EURO (2009, Bonn) and INFORMS (2009, San Diego) conferences, as well as in seminars at the Darden School of Business, Universidad Carlos III and Universidade Católica Portuguesa. We thank seminar participants, Ido Erev for his earlier encouragement, as well as Derek Bunn, Pär Holmberg, Andreas Krause, Leigh Tesfatsion and Ann van Ackere for their comments.

[†]Universitat Pompeu Fabra and City University. c/ Ramon Trias Fargas, 25-27. 08005 Barcelona. albert.banalestanol@upf.edu

[‡]Universitat Pompeu Fabra and Barcelona GSE. c/ Ramon Trias Fargas, 25-27. 08005 Barcelona. augusto.ruperezmicola@upf.edu. The author thanks the Spanish Ministry of Education for financial help under grant ECO2009-09834.

boundedly-rational, learning agents. ACE has been used to model situations as diverse as competitive strategy (e.g. Denrell, 2004), innovation (e.g. Adner and Levinthal, 2001), financial markets (e.g. Noe et al. 2003; Pouget, 2007), business organisation (e.g. Rivkin and Siggelkow, 2001), or electricity trading.¹

One of the advantages of ACE is that the models are tailored to closely fit each situation. However, this is a disadvantage when it comes to understanding the factors driving the results. ACE models are often not comparable to each other (Fagiolo et al., 2007) and, probably as a consequence, ACE is struggling to reach its full potential. Leombruni et al. (2006) reveal that only eight papers among over 40,000 published in a list of well-regarded economics journals use it, and argue that this is due to the difficulty in interpreting and generalising the simulation outcomes and the lack of consensus on the techniques suitable for each situation.

This paper explores the consistency and the reliability of agent-based simulation techniques. For our case in point, we focus on wholesale electricity markets. ACE is prominent in electricity because wholesale electricity markets are huge and especially awkward for analytical methods. They feature imperfect competition, very low demand elasticity, discontinuously convex supply functions, high-frequency repeated trading, heterogeneous agents and high potential for collusion (Wilson, 2002). The ACE electricity literature documents several models commissioned by large energy players (e.g. Gaz de France, E.ON, Shell) and the UK's Competition Commission as well as some calibrations of the US market, like the "Electricity Market Complex Adaptive System" (EMCAS) (Macal and North, 2005) and the "Agent-based Modeling of Electricity Systems" (AMES) (Sun and Tesfatsion, 2007).

Unfortunately, and despite the large electricity simulations literature (see Marks, 2006; Weidlich and Veit, 2008 for surveys), there is still little consensus on which techniques are most appropriate. First, some papers assume that agents learn according to the reinforcement model while other papers use more complex forms of behaviour like fictitious play or best response. Second, few papers specify the initial conditions of the models. Third, demand is assumed to be elastic in some cases and inelastic in others. Finally, several papers use stepwise schedules to model the supply part of the market, while in others sellers bid linearly increasing functions.

We investigate the effects of these assumptions on simulation outcomes and how these outcomes compare to simple, empirically-supported, theoretical predictions. Specifically, we cast light on whether the simulations are consistent with the standard claim that pivotal dynamics determines the relationship between competition and prices. A firm is pivotal if the quantity demanded exceeds the sum of production capacities of all other firms and, as a result, it is necessary to fulfill demand. There is wide consensus on the importance of pivotal dynamics in spot electricity markets (see e.g. Rothkopf, 2002, for a discussion).² In our setting, all firms are pivotal when there are few of them but none of them is pivotal if there are many of them. As a consequence, our theoretical results predict that prices will be high under monopoly, will decrease with competition, drastically change at a pivotal dynamics "switching point", and will approach marginal costs beyond that point.

¹See Tesfatsion and Judd (2006) and Leigh Tesfatsion's website for extensive information sources about ACE (<http://www.econ.iastate.edu/tesfatsi/ace.htm>).

²Many authors have noted, for example, that no single supplier in California had a 20% market share during the California crises, although one (or few) firms could have exercised monopoly power in half of the hours (Blumsack et al. 2002). Bushnell et al. (1999) developed the Pivotal Supplier Index (PSI) for a supplier at a point in time, defined as a binary indicator which is set equal to one if the supplier is pivotal, and zero if it is not. The Federal Energy Regulatory Commission (FERC) has adopted a similar indicator as a market power screening device.

We adopt a stylised setting that allows us to include alternative implementations of demand, supply and agent behaviour, which yield many of the literature’s models as particular cases. Demand can be inelastic or price-sensitive, with a wide range of levels and elasticity specifications. Firms are allowed to submit either flat bids or increasing supply schedules, with single or multiple bids per plant. Agents’ behaviour is governed by Camerer and Ho’s (1999) Experience-Weighted Attraction algorithm (EWA) which includes reinforcement learning, fictitious play and best-response as particular cases and allows for the specification of different initial conditions.

Simulation outcomes are consistent with our theoretical predictions under flat and supply function bidding, and under several plausible elasticities. However, we show that not all simulations exhibit the predicted breaking points. The performance of fictitious play is poor, and it is clearly outperformed by best-response and reinforcement learning. The performance of the simulations are influenced by the number of bids per plant and the initial conditions. The results call into question a large part of the extant ACE electricity research and can potentially enhance the practical implementation of these techniques.

This paper is part of a new literature examining the consistency of ACE in various market settings (e.g. Fagiolo et al., 2007; Leombruni et al., 2007; Marks, 2007; and Midgley et al., 2007). In the electricity industry, we are only aware of two related working papers. Li et al. (2009) check the robustness of several reinforcement learning parameters, elasticity, and price caps in the AMES model, and Kimbrough and Murphy (2009) compare step and supply function bidding in a stylised setting. The question of validation, that is which models best fit real market data, is complementary to ours. Our approach focuses on theoretical reliability, and includes comparisons of demand, supply, and behavioural specifications.

The remainder of the paper is organised as follows. In part 2 we discuss the literature. In part 3, we present our framework and the alternative implementations of demand, supply and agent behaviour. In part 4 we derive the theoretical prediction. Part 5 includes the simulation results and we conclude in part 6. All proofs are in the Appendix.

2 Agent-based electricity modelling alternatives

The three main sets of assumptions in ACE electricity models are the representation of supply, demand, behavioural rules. This section is a survey of the choices made in previous work. In Table 1 we classify some of the most relevant papers.

<<Table 1: Published papers>>

2.1 Supply bidding

Bertrand and Bertrand with capacity constraints are generally not considered suitable in the ACE electricity literature because they do not fit the uniform pricing prevalent in power pools. Cournot quantity bidding is sometimes used as an alternative (e.g. Bunn and Oliveira, 2007 and 2008; Veit et al., 2006). However, a recurrent argument is that Cournot is also unsuitable because in real pools generators are allowed to submit multiple flat bids for sections of their capacity. Hence, most papers use either von der Fehr and Harbord’s (1993) stepwise auctions, or Green and Newbery’s (1992) adaptation of the “supply function” (SF) equilibrium due to Klemperer and Meyer (1989).

In the stepwise approach, the market is a sealed-bid, multiple-unit auction. Generators simultaneously submit single prices at which they are willing to supply sections of their capacity. An independent auctioneer ranks the bids according to their offer prices, intersects the demand and supply and determines the system marginal price.

The stepwise literature includes both per plant and overall firm bidding models. Stepwise auction papers with one bid per generator are more parsimonious and comparable to the theoretical literature. Nicolaisen et. al. (2001) and Richter and Sheblé (1998) create models similar to those of auction theory to study the structure and efficiency of electricity markets. Closer to industrial organisation, Rupérez Micola and Bunn (2008) and Rupérez Micola et al. (2008) examine how horizontal and vertical integration influence the firms' ability to exert market power. Nanduri and Das (2007) add a simple electricity network. Bagnall and Smith (2005) study how their model replicates human behaviour in the England and Wales market.

Other papers allow one bid per plant rather than per firm. Bower and Bunn (2000, 2001) simulate the transition between the England and Wales pool and the New Electricity Trading Arrangements (NETA) discriminatory auction and how this could affect market prices. Bunn and Martoccia (2005) also replicate the UK market and Bower et al. (2001) focus on Germany. Only García et al. (2005) and Banal-Estanol and Rupérez Micola (2009) include abstract models with multiple stepwise bids per plant.

Inherent inflexibilities in the operation of nuclear assets (e.g. safety concerns, very low marginal costs and high start-up and loss of volume costs) prompt generators to submit flat schedules at very low prices for each plant. However, the assumption is quite restrictive in comparison to most bid-based electricity markets where they can submit many bid steps per plant. Multi-bidding leads to the well-known "hockey stick" shape of the supply curve, with base-load plants submitting flat schedules and peak-load generators offering steeper step functions. Accordingly, a number of simulations include several bids per plant. For example, Day and Bunn (2001) and Bunn and Day (2009) developed detailed models of the England and Wales pool between 1990 and 2001. Bunn and Oliveira (2001, 2003) look into the related effect of NETA's introduction and test whether the incumbents could influence prices. However, these models are often computationally cumbersome for two reasons. First, the algorithm's operations grow with the number of bids. Second, agents' coordination is more difficult, which complicates learning and the convergence to a steady state.

The SF approach approximates actual bids with increasing supply functions relating quantities and prices. This is a compromise between providing realism and simplifying the simulation mechanics. Banal-Estanol and Rupérez Micola (2009) include an SF model with two stepwise bids per firm. Cincotti et al (2005) study the effect of market microstructure and costs on prices, and Visudhiphan and Ilic (1999) focus on dynamic learning. Day and Bunn (2001, 2009) propose an even more flexible approach in which firms submit several SF sections per plant.

However, the presence of multiple Nash equilibria complicates the comparison of ACE and equilibrium results. For example, the SF model has little predictive value if the range of variation in demand is small because almost anything between the Cournot and the competitive solution can be supported in equilibrium (see Bolle 1992). Further, the solution is undefined if there is no short-run demand elasticity (von der Fehr and Harbord, 1998). Similarly, there are often many non-Pareto ranked equilibria in stepwise auction settings (for discussions, see e.g. von der Fehr and

Harbord, 1993; Crawford et al., 2006).³

2.2 Demand representation

The current literature mostly represents demands with double-sided call auctions or fixed aggregated curves.

First, double-sided call auctions consist of supply and demand bidding and common valuations bounded between marginal cost and redemption values. Bids represent the price at which firms are willing to sell and buy all their capacity in a double version of the stepwise auction setting. Examples include the papers by Richter and Sheblé (1998), Nicolaisen et al. (2001), Bunn and Oliveira (2001, 2003) and Rupérez Micola and Bunn (2008).

However, many ACE simulations use aggregate demands. The literature often models short-run electricity demand as inelastic, in part due to the lack of real-time metering systems (e.g. Stoft, 2002). Examples include Bagnall and Smith (2005), Bunn and Martoccia (2005), Cincotti et al (2005), García et al. (2005), Nanduri and Das (2007), Rupérez Micola et al. (2008) and Sun and Tesfatsion (2007).

At first glance, this may seem like an appropriate representation of reality, but there are several reasons why relaxing that assumption can add value. First, markets have some level of bid-in demand, or implicit elasticity provided through the actions of system operators who may take out-of-market actions to effectively reduce demand when prices rise. Second, most volume is traded outside of balancing markets, either in exchanges or bilaterally. Third, financial derivatives increase demand elasticity. Finally, inelastic demand models tend to present a large number of non-Pareto ranked pure strategy equilibria. Papers with elastic demands include several by Derek Bunn and coauthors (e.g. Bower and Bunn, 2000, 2001; Bower et al., 2001; Bunn and Oliveira, 2007, 2008; Day and Bunn, 2001), and also Veit et al. (2006).

To our knowledge, only Banal-Estanol and Rupérez Micola (2009), Li et al. (2009) and Visudhiphan and Ilic (1999) use both elastic and inelastic demands. Still, they do not seek to explicitly explore the implications of the elasticity assumption.

2.3 Behavioural algorithm

ACE models require rules to govern agent behaviour. One of their main intentions is to realistically represent human decision-making, and its proponents frequently argue that existing deductive mechanisms often do poorly in experiments (Camerer and Ho, 1999, Roth and Erev, 1995; and van Huyck et al., 1990, are regularly used to support this claim). The electricity ACE literature is based on adaptive learning algorithms mainly derived from psychology.

Some previous work uses reinforcement learning (RL). In RL, agents tend to repeat actions that led to positive outcomes and avoid those that were detrimental. Several papers have used modified versions of the Roth and Erev (1995) algorithm, e.g. Banal-Estanol and Rupérez Micola

³Economic theory offers refinements to single out unique supply function equilibria. Klemperer and Meyer (1989) show that if outcomes with infinite demand occur with positive probability and there are no capacity constraints, then there is a unique SF equilibrium. Green and Newbery (1992) focus on the highest profit equilibrium. Baldick and Hogan (2002) rule out unstable equilibria and add a price cap and capacity constraints. Newbery (1991) considers entry and assumes bid-coordination. More recently, Holmberg (2007, 2008) shows that a unique SF equilibrium exists with binding capacity constraints.

(2009), Li et al. (2009), Nanduri and Das (2007), Nicolaisen et al. (2001), Rupérez Micola and Bunn (2008), Rupérez Micola et al. (2008), Sun and Tesfatsion (2007) and Veit et al. (2006). It is 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 flatten out. These are robust properties observed in the literature on human learning. One of RL's main strengths is that one does not need to make assumptions on the information that players have about each other's strategies, history of play and the payoff structure. This is consistent with the fact that, in many cases, electricity traders cannot observe one another's current strategies, and only imperfectly infer them from volatile prices. However, RL might be too simplistic to fully capture the strategic opportunities available to humans (Erev et al., 2007; Ert and Erev, 2007).

It is likely that players in real electricity markets attempt to engage in more sophisticated behaviour like best response to their competitors' actions. There are two main types of best response algorithms: fictitious play (FP) and "Cournot" best response (BR).⁴ In FP (Brown, 1951), each player assumes that her opponents play stationary, possibly mixed, strategies. In each round, the player best responds to her opponent's empirical frequency of play. Electricity studies using FP include those by Bunn and Oliveira (2001, 2003) and García et al. (2005). BR implies that the player only responds to her opponents' move in the directly precedent period. BR papers include those by Bunn and Oliveira (2007, 2008) and Day and Bunn (2001) and Bunn and Day (2009).⁵ To our knowledge, there is no research on whether the results obtained with RL, FP and BR differ substantially in the electricity context.⁶

Finally, most papers do not provide details about initial conditions. In those that do, the standard approach is to use a uniform initial probability distribution for all elements of the action space. Examples include Rupérez Micola and Bunn (2008), Rupérez Micola et al. (2008) and Banal-Estanol and Rupérez Micola (2009). We are not aware of any papers explicitly exploring the impact of alternative starting conditions.

3 Modelling specifications

Our model incorporates key features of electricity markets in the short-run. Although it could be easily extended to become more complex, it is stylised to facilitate the exposition as well as the comparison between theoretical predictions and simulation results. We first present the market structure and trading rules that form our framework. Then, we describe the alternative parameter implementations of demand, supply and agent behaviour, which yield many of the literature's

⁴The term "Cournot" in this context does not refer to quantity bidding but to the classic tâtonnement best-response process leading to market equilibrium. To avoid confusion, given that we do not use quantity bidding in our simulations, we refer to this behavioural algorithm as "best response" (BR).

⁵In addition, there are a number of papers who depart in varying degrees from those models. Bower and Bunn (2000, 2001), Bower et al. (2001) and Bunn and Martoccia (2005) all use an adjustment based on locally improving/not deteriorating the profit and of a desired level of market share, and Cincotti et al (2005) use a different version of reinforcement learning. Bagnall and Smith (2005) use hierarchical classifier systems and Richter and Sheblé (1998) use genetic algorithms.

⁶Pouget's (2007) analysis is similar to ours, but focuses on financial markets. A different aspect is the parameter choice within a learning paradigm. Banal-Estanol and Rupérez Micola (2009) test the robustness of their results to nine reinforcement learning parameter combinations and, more systematically, Li et al. (2009) use intensive parameter sweeps to determine suitable settings for two potentially critical learning parameters.

models as particular cases.

3.1 Market structure and trading rules

Let there be n symmetric generators, $i = 1, \dots, n$, with constant marginal production costs, c , up to capacity. Denoting the market capacity as K , the individual capacity of each firm is $k_n = K/n$. For a given K , n parametrises the degree of competition in the market, as the individual capacities decrease with the number of generators.

Prices are bounded between marginal costs and Ψ , with Ψ being the maximum “reasonable” price cap (e.g. Lin et al., 2009). This 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. It also reflects high cost back-up power generation facilities owned by many industrial users. There are no grid constraints. Although relevant in the long term, we do not deal with capacity expansion, long-term contracts, ancillary and capacity payments.

Trading takes place through a compulsory, uniform-price auction. Suppliers simultaneously submit individual schedules. An independent auctioneer adds them horizontally and creates an *ad hoc* market supply function. Then she intersects it with the market demand and determines the uniform price \hat{p} . Finally, she assigns individual quantities, q_i , to each of the bidders. Profits for each firm are

$$\pi_i = (\hat{p} - c) \cdot q_i \quad \text{for } i = 1, \dots, n. \quad (1)$$

3.2 Demand representations

We use alternative parameters to accommodate different demand levels and elasticities. Demand, $Q(p)$, can be inelastic or price-sensitive. In the inelastic case, demand is equal to a constant quantity \bar{Q} for any price between zero and Ψ , i.e. a vertical line at \bar{Q} , $Q(p) \equiv \bar{Q}$. We rotate this curve to obtain linear functions with different elasticities at the same point. We denote the vertical coordinate of the rotation point as v ($0 \leq v \leq \Psi$) and the deviation to the left of \bar{Q} at the price cap level as u ($0 \leq u \leq \bar{Q}$). Thus, all demand curves are linear, pass through (\bar{Q}, v) and $(\bar{Q} - u, \Psi)$ and can be written as

$$Q(p) \equiv \bar{Q} - \frac{u}{(\Psi - v)}(p - v).$$

In all cases, market demand is assumed to lead to system overcapacity, i.e. $Q(p) < K$ for all p .

Figure 1 shows examples of demand with $\bar{Q} = 80$, $\Psi = 200$, $v = 100$ and $u = 0, 5, 10$. The three demands are $Q(p) = 80$ ($u = 0$, in purple), $Q(p) = 80 - \frac{5}{100}(p - 100)$ ($u = 5$, in magenta), and $Q(p) = 80 - \frac{10}{100}(p - 100)$ ($u = 10$, in red).

<<Figure 1: Examples of alternative demands and supply schedules>>

3.3 Supply representations

Supply schedules vary along two dimensions. First, firms are allowed to submit either flat bids (“stepwise bidding” case) or increasing supply schedules (“SF bidding” case). Second, in line with other papers (e.g. Day and Bunn, 2001, and Hobbs and Pang, 2007), we consider the implications of allowing firms to submit multi-step schedules. To that purpose, we divide each firm’s capacity

into m equally-sized capacity bins, k_n/m . We use alternative values of m in the stepwise and supply bidding cases. In the former, firms submit flat bids for each segment, while in the latter, they use supply functions for each segment.⁷

We now provide details for each case and explain the market clearing process.

Stepwise bidding The feasible price offer domain is approximated by a discrete grid. Generators choose from S possible bids, equally spaced between c and Ψ , at which they are willing to supply each bin's capacity. That is, the set of possible bids is

$$S_m(q) \equiv \{c + s(\Psi - c)/S \mid s = 1, \dots, S\}.$$

Each possible bid corresponds to an “action” or choice variable s . Bids generated from lower actions are closer to c , i.e. more competitive. The individual schedules for each bin are flat but firms can submit multi-step schedules if $m > 1$.

Supply function bidding Supply schedules are non-decreasing. For each bin, the generators choose among S possible angles, $s = 1, \dots, S$, equally spaced between zero and ninety degrees (or $\pi/2$ radians). The schedules consist of the set of linear curves from the coordinates $(0, c)$ until $(k_n/m, b(s))$, capped at Ψ ,

$$S_m(q) \equiv \left\{ \min\left(c + \frac{b(s) - c}{k_n/m} q, \Psi\right) \mid s = 1, \dots, S \right\},$$

where

$$b(s) = c + \frac{\sin(s(\pi/2)/S)}{\cos(s(\pi/2)/S)} k_n/m.$$

The angle of the plant's supply schedule, s , is the “action” and therefore the choice variable. Schedules generated from lower actions are more competitive because they are flatter. The supply function for the lowest action ($s = 1$) is almost flat at c . The supply function when $s = S$ is the result of capping a vertical linear function at the origin.⁸ Schedules generated from high actions become flat at Ψ . Note that the amounts sold by each firm are always strictly positive.

Market clearing The auctioneer sets \hat{p} by intersecting the market demand and supply functions. Under stepwise bidding, she gives full capacity to bids below \hat{p} ; the remaining capacity to those equal to \hat{p} (in case of a tie, the selling bin is selected randomly); and zero sales to the bids above \hat{p} . Under SF bidding, she assigns full capacity to the parts of each schedule below \hat{p} . Parts above \hat{p} receive nothing. Appendix A includes a formal derivation of the market clearing process for each treatment.

The two panels in Figure 1 show hypothetical bidding examples with $n = 2$, $m = 1$, $K = 100$, $\Psi = 200$, and $c = 0$. The market supply function (black line) is the horizontal addition of the

⁷The models by Day and Bunn (2001 and 2009) are more elaborate because agents submit fixed slope and changing intercept, fixed intercept and changing slope, and changing both slope and intercept functions in a non-continuous supply function setting.

⁸In the simulations we add a marginal amount to the angle of the denominator to avoid the indeterminacy when $s = S$.

individual functions (blue and green).

3.4 Agent behaviour representations

We adopt the Experience-Weighted Attraction (EWA) adaptive learning mechanism put forward by Camerer and Ho (1999). This is a general learning model that nests RL, FP and BR as special cases. EWA assumes that each feasible action of each bin has a numerical attraction. The numerical attractions of all the actions of the same bin generate a bin-specific probability distribution. In each round, generators submit supply schedules according to these bin-specific probability distributions. Once the auctioneer determines the price and individual quantities, the attractions are adjusted with the behavioural rule and mapped into new probability distributions. This process is repeated until the simulation converges.

We now describe how agents use experience to update the attractions, and how these lead to choice probabilities. Then we specify the initial attractions and the convergence definition.

Updating rules and choice probabilities Each action s for bin j in generator i has an “attraction” $A_{i,s}^j(t) > 0$ after period t (≥ 1). Attractions are updated according to

$$A_{i,s}^j(t) = \frac{\phi N(t-1)A_{i,s}^j(t-1) + \left[\delta + (1-\delta)I(s, r_i^j) \right] \pi_i(s, r_i^{-j})}{N(t)}, \quad (2)$$

where r_i^j and r_i^{-j} denote the action taken in period t by bin j of firm i and by the rest of the bins, respectively; $I(x, y)$ is an indicator function with value 1 if $x = y$ and 0 if $x \neq y$; $N(t) = \rho N(t-1) + 1$ with $N(0) = 0$, representing the number of “observation-equivalents” of past experience; and the EWA parameters δ , ϕ , and ρ denote the weight placed on foregone payoffs, a discount factor to depreciate previous attractions, and a discount factor that weights the impact of previous against future experience, respectively.

When $\delta = 0$ and $\rho = 0$, EWA behaves like the widely used class of RL models (e.g. Roth-Erev, 1995). RL models 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 flatten out. When $\delta = 1$, and $\rho = \phi$, EWA is equivalent to the standard weighted belief-based models. In particular, it produces BR when $\rho = \phi = 0$ and FP when $\rho = \phi = 1$. In BR dynamics, players actions are determined by the best response to what her opponents did in the immediately preceding period, so that only the most recent observation counts. In FP, each player best responds to the empirical frequency of play of her opponents since the beginning of the game, and all observations count equally.

The algorithm linearly maps attractions into action choice probabilities. The probability of selecting an action in the next period is its attraction divided by the sum of attractions for all actions. That is, for $t \geq 0$,

$$P_{i,s}^j(t+1) = \frac{A_{i,s}^j(t)}{\sum_{k=1}^S A_{i,k}^j(t)}. \quad (3)$$

Prior beliefs and initial conditions The probabilities in the first period, $P_{i,s}^j(1)$, are generated from prior values of the attractions, $A_{i,s}^j(0)$. As explained by Camerer and Ho (1999), the prior values of the attractions may reflect pre-game experience. We construct $A_{i,s}^j(0)$ from four representative assumptions on prior beliefs. We assume that agents believe that the others will initially (i) choose the highest possible bid ($s = S$), (ii) choose the mid-point of the range ($s = S/2$), (iii) choose the lowest possible action ($s = 1$), or (iv) use a uniform distribution over all their actions. In each of the four treatments, $A_{i,s}^j(0)$ for all i and j is defined as the hypothetical profit that each action s would render if prior beliefs about the opponents were correct.

Figure 2 reports the impact of prior beliefs on the initial probabilities for any agent in markets with one and twelve agents (assuming stepwise bidding, $m = 1$, $K = 100$, $\Psi = 200$, $S = 50$ and $c = 0$). Actions are identified on the horizontal axis, with numbers ranging from 1 for the more competitive to 50 for the highest possible bid. Agents will use mixed strategies unless the probability of playing all but one action is zero. Probabilities concentrated on higher actions result in less competitive bids and vice-versa.

<<Figure 2: Initial probability distributions>>

By definition, prior beliefs on others have no impact when there is only one agent. Take the case of twelve agents. In treatment (i), that is when agents believe that their opponents will initially bid Ψ , they think that any bid below Ψ would allow them to sell full capacity at Ψ . Hence all the actions below S have initial attractions equal to the maximum profits, i.e. $A_{i,s}^j(0) = \Psi k_n$ for all $s < S$, and the same initial probability, $P_{i,s}^j(1) \approx \frac{1}{S-1}$.⁹ In treatment (ii), when agents believe that others will initially bid the minimum price ($s = 1$), any bid above the minimum would be out-of-the-money and earn them zero profits. Therefore, the initial attractions are concentrated on $s = 1$ and therefore $P_{i,1}^j(1) = 1$ and $P_{i,s}^j(1) = 0$ for $s > 1$. In treatment (iii), when agents assume the others will initially bid in the middle of the distribution, they randomise over their lower half and assign zero probability in the upper half. Finally, in treatment (iv), when players assume that the others will initially follow a uniform distribution, they bid more competitively than the diagonal to sell with higher probability their full capacity at the \hat{p} set by an opponent.

To our knowledge, the literature only includes uniform initial probability distributions over all the elements of the action space, which implicitly implies that agents believe that their opponents will bid Ψ . In this paper, we use that assumption as a reference and study whether the results are sensitive to the alternatives.

Convergence Following the theoretical literature on learning (see Fudenberg and Levine, 1998, for an overview), we define convergence in terms of strategy profiles. We say that a simulation run has converged if the maximum per-period change in the probability of playing any strategy is below a (small) threshold.

⁹To be precise,

$$P_{i,s}^j(1) = \frac{\Psi k_n}{\Psi k_n/n + (S-1)\Psi k_n} = \frac{n}{1+n(S-1)} \text{ for } s < S \text{ and } P_{i,\Psi}^j(1) = \frac{\Psi k_n/n}{\Psi k_n/n + (S-1)\Psi k_n} = \frac{1}{1+n(S-1)}.$$

If the number of firms is high, then $P_{i,s}^j(1) \approx \frac{1}{S-1}$ and $P_{i,\Psi}^j(1) \approx 0$.

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

$$\left| P_{i,s}^j(t+1) - P_{i,s}^j(t) \right| < \tau. \quad (4)$$

The simulation price is computed from the agents' mixed strategy profile z .

In practice, we select the action with the lowest probability in period t . Then, we compute the hypothetical probability that would result from assigning maximum profits to this action and minimum profits to all the other actions. The simulation has not converged as long as the difference between present and future probabilities is higher than τ . It has converged when it is lower. The smaller τ , the more stringent the threshold and the higher the necessary t .

Once there is convergence, we calculate expected end-of-simulation prices from the individual probability distributions. Note that convergence is compatible with the survival of several feasible trading actions, as in mixed strategies. Price volatility may not be equal to zero even if there is a steady state. Moreover, EWA bidding depends on the stochastic process and, as a result, simulation runs for the same parameters might lead to different end prices, i.e. the standard deviation of mean prices across simulations is not necessarily zero.

4 Theoretical predictions

We now derive predictions on the effect of competition on market prices. The predicted market prices shall depend on the number of “pivotal” plants (see e.g. Genc and Reynolds, 2005; Entriken and Wan, 2005; Perekhodtsev et al., 2002, and Banal-Estanol and Ruperez Micola, 2009). A firm is pivotal if it is necessary to satisfy the quantity in demand. In the inelastic case, the definition is straightforward as the demand is constant, but in the elastic case, the quantity demanded depends on the supply bids. In general, one has to define pivotality for an exogenous demand level.

Definition 2 A firm i is **pivotal** for a given level of demand Q' if this level exceeds the sum of production capacities of all other firms, i.e. if $\sum_{j \neq i} k_n = (n-1)k_n < Q'$.

Pivotal dynamics are simple in symmetric settings. In markets with few firms, they are all pivotal. In those featuring many firms, none is pivotal. We next define the level of competition at which the number of pivotal firms changes.

Definition 3 A level of competition \hat{n}^l is a **lower switching point** if (i) all firms are pivotal at the minimum demand ($Q' = Q(\Psi)$) for $n < \hat{n}^l$ and (ii) none of them is for $n \geq \hat{n}^l$.

Definition 4 A level of competition \hat{n}^u is an **upper switching point** if (i) all firms are pivotal at the maximum demand ($Q' = Q(c)$) for $n < \hat{n}^u$ and (ii) none of them is for $n \geq \hat{n}^u$.

If demand is inelastic, the upper and lower switching points coincide and, for simplicity, we call them “the switching point”. For example, if $Q(p) = 80$ for any p and $K = 100$, the switching point is $\hat{n}^u = \hat{n}^l = 5$ because if $n < 5$, $(n-1)k_n < 80$, and all firms are necessary to fulfill demand but no firm is pivotal when $n \geq 5$, $(n-1)k_n \geq 80$.

In the elastic case, the upper and lower switching points are different. Take for instance the case in which $\Psi = 200$, $v = 100$, $\bar{Q} = 80$, $u = 10$, $c = 0$ and $K = 100$. The minimum and maximum demands are $Q(200) = 70$ and $Q(0) = 90$. There is a lower switching point at $\hat{n}^l = 4$ since $(n - 1)k_n < 70 = Q(200)$ for $n < 4$ and $(n - 1)k_n \geq 70 = Q(200)$ for $n \geq 4$ and an upper switching point at $\hat{n}^u = 10$ because $(n - 1)k_n < 90 = Q(0)$ for $n < 10$ and $(n - 1)k_n \geq 90 = Q(0)$ for $n \geq 10$. Now, the equilibria.

Proposition 5 (a) For each K , $Q(p)$, there exists a unique upper switching point, $\hat{n}^u = K / (K - Q(c))$.
(b) If the number of firms is lower than this threshold ($n < \hat{n}^u$), then any given firm i bidding $p_m^*(n)$, where $p_m^*(n)$ is the monopoly price of the residual demand curve, $p_m^*(n) \equiv \arg \max_{0 \leq p \leq \Psi} \{(p - c)[Q(p) - (n - 1)k_n]\}$, is part of any equilibrium. The equilibrium price is $p_m^*(n)$.
(c) If the number of firms is higher than this threshold ($n \geq \hat{n}^u$), then all firms bidding c is an equilibrium. The equilibrium price is c .

The equilibrium price drops from the monopoly price of the residual demand to competitive levels at \hat{n}^u . The equilibrium is unique if $n = 1 < \hat{n}^u$ (monopoly pricing) and if $n > \hat{n}^u$ (marginal cost pricing). If $1 < n < \hat{n}^u$, however, there are multiple pure strategy equilibria with many payoff-equivalent actions as part of each of them.¹⁰ Consider, for example, the case of two firms, inelastic demand, one bin and stepwise bidding. A generator bidding Ψ and the other bidding close to c , or vice versa, are both equilibria in which one obtains a low profit and the other gets the maximum. The situation is similar to the standard ‘‘battle of the sexes’’ game. Experimental evidence shows that coordination in this game can be low (Cooper et al., 1990), and especially difficult due to the payoff asymmetry in each of the equilibrium (Crawford et al., 2008). In sum, if $1 < n < \hat{n}^u$, market prices are expected to be at most $p_m^*(n)$ but can be substantially lower because of multiple equilibria and coordination problems.

The next corollary shows that a higher n decreases the equilibrium price $p_m^*(n)$ and increases equilibrium profit asymmetries. This will further reduce the agents’ incentive to be the price-setter, and induce them to submit lower bids.

Corollary 6 If the number of firms below the threshold ($n < \hat{n}^u$),

- (a) the equilibrium price ($p_m^*(n)$) is non-increasing in n
(b) the relative profits of a price-setting firm with respect to a non-price setting firm in the equilibria $((p_m^* - c)[Q(p_m^*) - (n - 1)k_n] / [(p_m^* - c)k_n])$ decrease in n .

The proposition and the corollary allow us to sketch predictions about the effect of competition on prices. Monopoly prices should be equal to $p_m^*(1)$ as this leads to maximum profits. Prices should decrease in n as long as the firms are pivotal at the minimum demand, $n < \hat{n}^u$, due to lower equilibrium prices and growing profit asymmetries. Prices should drastically decrease at the upper switching point (\hat{n}^u) because the unique prediction is that prices are equal to marginal costs for $n > \hat{n}^u$.

Prediction 7 The dynamics pre- and post-switching point result in nonlinearities in the influence of n on prices. Pre \hat{n}^u , prices decrease with n from the monopoly price, $p_m^*(1)$. Post \hat{n}^u , prices are drastically reduced to c .

¹⁰There are also many mixed strategy Nash equilibria.

Note that this prediction is not specific to any of the supply, demand, and behavioural assumptions that we explore in this paper. Further, there is wide consensus on the importance of pivotal dynamics in real electricity markets (for a discussion on the role of pivotal dynamics, see e.g. Rothkopf, 2002). The prediction is so standard that any combination of agent-based modelling assumptions aiming to reproduce spot electricity markets should be able to fulfill it.

5 Simulations

This section presents the simulation results. We first introduce the simulation parameters. Then we graphically compare the simulation outcomes with the theoretical predictions. Following that, we formally test whether the data for each demand, supply and behavioural specification features breaking points at the predicted locations and, whether the predicted switching point is the best-fitting breaking point. Finally, we present some results on the impact of the behavioural assumptions on the stationary bidding strategies.

5.1 Parameters

We allow the number of firms, which parametrises the degree of competition, to vary from one to twelve, $1 \leq n \leq 12$. Marginal costs are set to zero, $c = 0$. Total capacity is $K = 100$, so individual capacities decrease from $k_1 = 100$ to $k_{12} = 8.33$. The price ceiling is $\Psi = 200$, with a grid of $S = 50$ possible actions and a demand rotation point $v = 100$.

To clarify the exposition, we focus on the demand, supply and behaviour assumptions in turn. We use as a reference specification an inelastic demand with 20% excess capacity ($u = 0$ and $\bar{Q} = 80$), firms using stepwise bidding with one bin ($m = 1$), and agents learning following RL with a uniform initial distribution (which, as shown in Section 3.4 (Figure 2), is equivalent to assuming prior beliefs equal to the maximum bid, Ψ).

To study the effect of the demand assumptions, we fix the reference supply representation (stepwise bidding and $m = 1$) and the reference behavioural algorithm (RL and uniform initial distribution). We run simulations for the combinations of expected demand levels $\bar{Q} = \{80, 85, 90\}$ and elasticity parameters $u = \{0, 5, 10\}$. We perform 50 simulations for each n and for each combination of demand level and elasticity parameters. The data set includes $3 \cdot 3 \cdot 12 \cdot 50 = 5,400$ observations.

When we focus on the supply side, we fix the reference demand specification ($\bar{Q} = 80$ and $u = 0$) and the reference behavioural algorithm (RL and uniform initial distribution). We perform simulations for each, the stepwise and supply function bidding, and for one, two and three bins. The resulting data includes $2 \cdot 3 \cdot 12 \cdot 50 = 3,600$ observations.

For the behavioural analysis, we fix the reference supply and demand specifications (stepwise, $m = 1$, $\bar{Q} = 80$ and $u = 0$). We perform simulations for each, BR, RL and FP, and for prior beliefs equal to Ψ , 0, $\Psi/2$ and random. In total, there are $3 \cdot 4 \cdot 12 \cdot 50 = 7,200$ observations in the data set.

Substituting into Proposition 5, the monopoly prices are equal to the maximum price, $p_m^*(1) = \Psi = 200$, in all the specifications. The upper switching points are $\hat{n}^u = \{5, 6, 10\}$, respectively for $\bar{Q} = \{80, 85, 90\}$, if $u = 0$ (inelastic demand), and $\hat{n}^u = \{7, 10, 12\}$ if $u = 5$ and $\hat{n}^u = \{10, 12, 12\}$ if $u = 10$.

5.2 Simulation results

Figures 3, 4 and 5 compare the impact of competition on prices for the different demand, supply and behavioural specifications. Each panel plots the long-run average price (\bar{p}) and two standard deviations across the 50 simulation runs for each n . Since the theoretical price predictions are corner solutions, the intervals sometimes exceed the simulation boundaries. The upper-left panel in each figure corresponds to the same reference specification ($\bar{Q} = 80$, $u = 0$, stepwise bidding, $m = 1$, RL and a uniform initial distribution).

Demand specifications Figure 3 reports the demand assumption results. Simulations use as demand levels $\bar{Q} = 80, 85$ or 90 (in columns one, two and three, respectively) and, as elasticity assumptions, $u = 0, 5$, or 10 (in rows one, two and three). In all the panels, the supply and behavioural assumptions are one-bin stepwise bidding and RL with uniform initial distribution.

<<Figure 3: The influence of demand specifications on prices>>

As predicted, simulated monopoly prices are close to $\Psi (= 200)$ in all instances and the relationship between n and \bar{p} is decreasing. In the inelastic cases ($u = 0$, first row), \bar{p} drops rapidly in n but flattens out around $n = 5$, $n = 6$ and $n = 10$, respectively for $\bar{Q} = \{80, 85, 90\}$, as predicted by the theory. More competition has a small effect beyond those values, as we predicted, but \bar{p} remains clearly above c , particularly in the case of tight capacity ($\bar{Q} = 90$, third column).

In the elastic cases (second and third rows), our theoretical section predicted a break at the *upper* switching points: $\hat{n}^u = \{7, 10, 12\}$ for $u = 5$ and $\hat{n}^u = \{10, 12, 12\}$ for $u = 10$ for $\bar{Q} = \{80, 85, 90\}$, respectively. That is, as elasticity increases, we predict a break for higher levels of n . Instead, the results seem to show a lower breaking point as the elasticity increases, particularly in the case of tight capacity. As we will see in the following subsection, this is consistent with a break at the *lower* switching point.

Comparing simulation results across specifications, the impact of the elasticity seems to be, at best, modest. Indeed, simulations are similar across rows, both in terms of price levels and in terms of the shape of the relationship between n and p . In the case of tight capacity, however, a higher elasticity seems to make the price more sensitive to n . Across demand levels, \bar{p} tends to be less sensitive to n as the levels of demand increase. Higher demand levels also lead to higher overall prices. This is also consistent with Proposition 5, which predicts that equilibrium prices are increasing in \bar{Q} (higher \hat{n}^u and higher p_m^*).

Overall, the demand results are quite, but not perfectly, consistent with the theory. First, although monopoly prices are close to Ψ and the relationship between n and \bar{p} is decreasing, post-threshold prices are far from c . Second, inelastic simulations fit the break predictions better than those with elasticity. Third, smaller excess capacity ($K - \bar{Q}$) results in higher prices. Fourth, results do not vary too much within our elasticity ranges, which are comparable to those in the literature.

Supply specifications Figure 4 reports the bidding assumption results. Simulations use stepwise (first row) or SF bidding (second row) with one, two or three bins (columns one, two and three). We use $\bar{Q} = 80$ and $u = 0$ and RL with uniform initial distribution in all cases.

<<Figure 4: The influence of supply specifications on prices>>

As predicted, the relationship between n and \bar{p} monotonously decreases both for the stepwise and the SF assumptions. Its shape changes around $n = 5$, consistent with the prediction of a switching point at $\hat{n}^u = 5$. In the stepwise case, however, prices remain above the competitive levels after the threshold. \bar{p} 's sensitivity to n decreases with the number of bins and therefore the deviation from theory grows. When $m = 1$, post-switching point prices are around 40, when $m = 2$ around 75 and when $m = 3$ around 100. If there are multiple bins, agents are able to use some of their bins to keep prices high while recouping some of the benefits of high prices with the other bins.

Although the price variability increases, SF yields a better fit in terms of average prices. This is probably because its higher “expressiveness” overcomes the difficulties to coordinate in the theoretical prediction. The standard deviation grows in m , especially around the switching point. Under SF, bids above the equilibrium price may be reinforced because they also obtain substantial profits. It is therefore more difficult for agents to tell good from bad bids and dispersion grows. In comparison, stepwise bids are either on- or out-of-the-money. Hence it is easier to identify the good bids and the simulations become crisper.

Overall, prices are less sensitive to n under stepwise bidding and SF is better at capturing the extreme monopoly and competitive predictions. Still, SF's price dispersions increase substantially around the pivotal breaks.

Behavioural specifications Figure 5 reports simulation results for the different behavioural specifications (assuming stepwise bidding with $m = 1$, and $\bar{Q} = 80$ and $u = 0$). RL is on the top, BR in the middle and FP in the bottom panel. The four columns correspond to the four prior belief assumptions in the following order: (i) the maximum price, (ii) randomly, (iii) the minimum price, and (iv) the medium point. As shown in Section 3.4 (Figure 2), this implies that the initial probability distributions are uniform, lower than the diagonal, concentrated in the lowest action, and uniform over the bottom half, respectively.

<<Figure 5: The influence of behavioural assumptions on prices>>

The panels consistently show that the relationships between n and \bar{p} are decreasing. All confidence intervals are narrow. In reinforcement learning (first row), monopoly prices are close to Ψ , decrease until the theoretical switching point ($\hat{n}^u = 5$) and approach c after it. The break at the switching point, however, is striking only if the prior is competitive (third column). In best response (second row), monopoly prices are instead far from Ψ (at about 140), but there is a clear breaking point for $n = 5$, after which they converge exactly to c . Initial conditions do not have an impact under BR due to the algorithm's lack of memory.

The differences between RL and BR results can be traced back to the algorithms' features. Under BR, prices have been shown to be competitive when no agent is pivotal. This is because under BR, the attractions of all the actions above those used by the other agents in the previous period are reduced to zero. Therefore agents choose lower actions and \bar{p} can only stay constant or decrease. The resulting unravelling yields $\bar{p} = c$. When all agents are pivotal, equilibrium forces tend to increase \bar{p} so that agents choose high actions with some probability. Simultaneously, unravelling prevents \bar{p} from staying very high. On balance, \bar{p} never reaches Ψ , not even in monopoly. In RL, instead, there is no in-built unravelling. Actions above those of the opponents are still played

because they might have generated profits in earlier periods. Thus, \bar{p} is higher than under BR for all n .

The fictitious play specification (third row) departs substantially from the predictions. Monopoly prices are around 140, decrease slowly in n and stay patently above competitive levels. FP keeps most of the initial noise as it weights heavily initial periods where outcomes are quasi-random. Hence, agents assign propensities to inadequate actions, adaptation slows down and there is strong path dependence. There are two countervailing forces at work. On the one hand, the higher n the more likely it is that there will be unravelling as in the case of best response. On the other hand, initial random prices are more likely to be high if there are more agents, so that reductions start from a higher base. Monopoly prices are far from Ψ because of unravelling. Prices when no firm is pivotal are far from zero owing to high initial prices. On balance, the n to p relationship decreases slowly and stays far from the theory extremes.

Overall, RL best matches the theory but is not as reactive to market conditions as BR. FP performs worst. Competitive prior beliefs render by far the most homogeneous post-switching prices across behavioural assumptions. This is because beliefs are self-fulfilling. Everyone's best response is to bid c when they believe that the others will bid c . BR and FP lock themselves up in that value while experimentation in RL is not powerful enough to depart from it. The theoretical soundness of RL with competitive initial beliefs is strikingly good and the best of the twelve. Monopoly prices are only slightly below Ψ , and decrease clearly in n for $n < 5$, so that prices are above 100. For $n > 5$, prices converge to c . These results are to our knowledge the first on the robustness of ACE techniques to behavioural choices in electricity markets.

5.3 Threshold regressions

In this section, we formalise the previous visual inspection with threshold regressions. Specifically, we carry out tests of whether the data features switching points in the predicted locations. To that purpose, we estimate a piecewise linear model between n and \bar{p} for each demand, supply and behavioural combination. The models are uniquely specified by a dummy variable associated with the upper switching point \hat{n}^u ,

$$p_i = \beta_0 + \beta_1 D_i + \beta_2 \alpha_i + \beta_3 D_i n_i + u_i, \text{ where } D_i = 0 \text{ if } n_i < \hat{n}^u, D_i = 1 \text{ when } n_i \geq \hat{n}^u. \quad (5)$$

The pre- and post-breaking points regression estimates are specified by

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

We test the null hypothesis of linearity against the alternative of structural breaks at the pivotal switching points \hat{n}^u . Evidence supporting the existence of a breaking point can come either from significant intercept or slope change coefficients, i.e. β_1 and β_3 different from zero.

Table 2 reports the results on the demand, supply and behavioural assumptions. On the left-hand side of each block, we specify the parameters used in each specification, together with the implied upper switching point. The parameters changing in each block are in boxes. The right-hand side reports regression estimates. The coefficients correspond to equation (5).

<<Table 2: Hypotheses' tests of breaking point regression estimates>>

In all cases, either β_1 or β_3 (or both) is significant at standard levels. There are three non-significant coefficients (one β_1 , two β_3) but the other coefficient in the same specification is always significant. The tests provide preliminary support for the hypotheses. One could therefore conclude that all simulations are consistent with the pivotal breaking point’s theory.

However, simple hypothesis tests might not be enough as there might be several points which can be identified as significant breaking points. This could explain why the tests identify breaks also in cases in which they are not visually identifiable (e.g. under fictitious play). One way to discriminate across breaking points is to single out those producing best-fitting models.

5.4 Identification of best-fitting breaking points

We create a procedure to select the “optimal” data breaking point and establish whether it is statistically equivalent to the theoretical one. We start by providing a definition for the optimal breaking point.

Definition 8 *An optimal breaking point \bar{n} satisfies $F(\bar{n}) \geq F(n)$ for any threshold n , where $F(n)$ denotes the F -statistic obtained from a piecewise linear regression with threshold n .*

We compute the optimal breaking point for each specification. We approximate the distribution of the optimal breaking points through a subsampling procedure.¹¹ For each specification, we extract 99 random subsamples of 240 observations stratified for n (i.e. twenty observations out of the fifty available for each n). Then, we use Definition 8 to obtain the subsamples’ optimal breaking points.

Figure 6 reports the comparison between the upper and lower theoretical switching points and the optimal data breaking points (subsamples’ mean, together with two standard deviations) for each specification. The upper panel (6a) provides results for the effects of the demand assumptions. In the inelastic case ($u = 0$), simulation results match the theoretical switching point if demand is low but are lower than the predicted switching points if demands are higher. Interestingly, in the elastic cases ($u > 0$), the simulated breaking points perfectly match the lower switching point (except in one case), instead of the upper switching point, as predicted by the theory.

<<Figure 6: Comparison of simulated and theoretical switching points>>

The middle panel (6b) reports breaking points for each of the supply assumptions. All combinations match the theoretical prediction well. Still, the confidence intervals for the supply function bidding cases are larger. The number of bins do not seem to have any effect on the simulated breaking points.

The lower panel (6c) includes the results for each of the behavioural assumptions. BR and RL match the theoretical prediction, independently of the initial condition used. FP, instead, generates much higher break point locations, except for one assumption on prior beliefs.

To summarise, the pivotal breaking point prediction is one of the simplest and most widely accepted in electricity markets. RL and BR specifications identify it well, but not FP. Notice that

¹¹One 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 (one in our case).

the optimal breaking point results do not coincide with the simple hypotheses tests of the previous subsection. Subsampling appears to be a good way to discriminate between several possible breaking points.

5.5 Latent intensity of competition

In a simulation environment, we can inspect the probability distributions from which agents choose their bids. Thus, it is possible to study how the parameters influence the agents’ “competitive attitude” and not only market outcomes.

Figure 7 depicts end-of-simulation individual latent probability distributions. It reports cumulative probabilities for one agent when $n = 1$ and $n = 12$ for the three behavioural assumptions, averaged across simulation runs (assuming $\bar{Q} = 80$, $u = 0$, stepwise bidding with one bin). The convergence definition implies that the probability distribution is largely invariant beyond the convergence point, so that these distributions are a good approximation of the agents’ long-term strategies.

<<Figure 7: Latent intensity of competition>>

The figure offers a number of general insights linking market structure, trading behaviour and prices. Distributions for $n = 1$ stochastically dominate those for $n = 12$, so that monopolist bids are less competitive than those of competitive agents. For a given n , there are substantial differences across behavioural algorithms. When $n = 12$, trading probabilities are extremely competitive under BR and RL. FP induces less competition. In contrast, when $n = 1$, BR and FP lead to the same probability distribution, which is more competitive than for RL.

We can also compare this end-of-simulation probability distributions with the initial probabilities, represented in Figure 2 (the case in which agents prior beliefs are the maximum). When $n = 1$, the final probability of distribution for BR and FP is the same as the initial one. In RL, instead, the monopoly agent learns to increase bids substantially. When $n = 12$, agents start from a uniform distribution but learn to submit more competitive bids in all cases, but particularly in BR and FP.

Overall, the figure provides solid support for the mechanisms underpinning the influence of each behavioural model on \bar{p} . It also shows that RL is more sensitive to n than the others, but also that it leads to less competitive outcomes than BR. As expected, n does not have a strong influence on FP distributions.

6 Discussion and conclusions

Firms and regulators alike have started to use agent-based computational economics (ACE) to study the properties of many markets. However, the literature has advanced little in creating a set of standards. This paper is an attempt to advance in that direction. We study the properties of different ACE techniques and how they compare to each other and to economic theory. As a case in point, we focus on a well-established claim in the wholesale electricity trading literature. Any model should be able to replicate it.

The demand, supply, and, particularly, behavioural comparisons call into question an important part of the extant ACE literature on electricity markets. Best-response and reinforcement learning

perform quite well, but fictitious play does not. Prior beliefs and initial conditions have an influence on the performance of the simulations. Competitive pre-game beliefs render the best match to theory. Flat and upward slopping supply functions yield similar results, and also several plausible elasticity assumptions. The simulations are influenced by the number of bids per plant. Simulations perform best when they combine inelastic demand, reinforcement learning, competitive initial beliefs and single-bin bids.

The paper has several implications. First, some preceding work makes choices that are not consistent with economic theory in our simple setting. There might be an important venue for research in checking the robustness of previous findings to alternative assumptions. Second, future modelling efforts should also incorporate more systematic robustness tests. Just as much as Cournot, Bertrand, Stackelberg or Supply Function Equilibria are widely used despite yielding different results, it might be appropriate for ACE researchers to use different algorithms. This paper underlines the need to understand, justify, and even sometimes exploit what each modelling choice entails.

Still, we have not included all possible assumptions. For example, we have left out variable marginal costs and behavioural rules like genetic algorithms and Q-learning that are also prominent, and we have not compared the implications of different convergence definitions. Furthermore, our strongest result relates to fictitious play. Its main difference with best response is mainly one of memory, so that how much memory to retain and how it should decay, are intriguing, and still unresolved, algorithmic questions.

Furthermore, this paper focuses on electricity markets. We should continue doing similar exercises in other settings, as in Fagiolo et al. (2007), Leombruni et al. (2007), Marks (2007) and Midgley et al. (2007). Finally, the question of which models fit best to real data is complementary to our research and deserves future attention both in the behavioural laboratory (e.g. Duffy, 2006) and empirically if ACE is to be more commonly used.

Appendix A: Supply functions

We show here how one can construct the market supply functions. For simplicity we assume that marginal costs are zero ($c = 0$) and firms can only use one bin ($m = 1$). First take the case of stepwise bidding. We proceed as follows. First, we order the “bids” from lowest to highest. Abusing of the notation, we denote them here as b_1, b_2, \dots, b_n . The horizontal sum of the individual supply functions $S(q)$ can be defined as

$$S(q) \equiv b_i \text{ for } \kappa_{i-1} < q < \kappa_i \text{ (} i = 1, \dots, n \text{) where } \kappa_i \equiv i k_n \text{ and } \kappa_0 \equiv 0.$$

Market clearing occurs at where this equation and the demand intersect. The market price \hat{p} is given by $\hat{p} = b_j$ where j is the unique index which satisfies $\kappa_{j-1} < Q(b_j) < \kappa_j$.

Second, take the case of supply function bidding. Again, order the “bids” ($b(s)$) from lowest to highest and abusing of the notation, denote them as b'_1, b'_2, \dots, b'_n . As shown in the text, the linear supply schedule for each firm i is given by $\min\{\frac{b'_i}{k_n}q, \Psi\}$, which is an increasing linear function from $(0, 0)$ until (k_n, b'_i) capped at Ψ . Denoting $b'_0 = 0$, the horizontal sum of the individual supply

functions $S'(q)$ can be defined as

$$S'(q) \equiv \min\left\{\frac{(b'_i - b'_{i-1})q + b'_{i-1}\kappa'_i - b'_i\kappa'_{i-1}}{\kappa'_i - \kappa'_{i-1}}, \Psi\right\} \text{ for } \kappa'_{i-1} < q < \kappa'_i \text{ (} i = 1, \dots, n),$$

where

$$\kappa'_i \equiv \left(i - 1 + \sum_{j=i}^n \frac{b'_j}{b'_j}\right) k_n \text{ and } \kappa'_0 \equiv 0.$$

Market clearing occurs at where this equation and the demand intersect. The market price \hat{p} can be found by solving the following implicit equation

$$\hat{p} = \frac{(b'_i - b'_{i-1})Q(\hat{p}) + b'_{i-1}\kappa'_i - b'_i\kappa'_{i-1}}{\kappa'_i - \kappa'_{i-1}}.$$

Appendix B: Proofs

Proof of Proposition 5

For part (a) note that $\hat{n}^u = K/(K - Q(0))$ is the unique solution to $(n - 1)k_n = Q(0)$.

For part (b), suppose first that firms can only use one bin and they bid using stepwise bidding functions. Suppose that all but firm i are bidding equal to the marginal costs and firm i is bidding at p_m^* . Firm i has no incentive to deviate since it is setting the monopoly price given the actions of the others. The other firms do not have an incentive to deviate either (1) at a price equal to the monopoly price because they would sell less at the same price; or (2) at a price below the monopoly because they would earn exactly the same. If $p_m^* < \Psi$, the remaining firms could set a price above the monopoly price. The optimal deviation should then be at a price marginally above p_m^* , but the deviator would sell less at the corresponding price.

For part (c), it is not profitable to unilaterally deviate because the remaining firms would serve the demand. Standard Bertrand arguments imply that this equilibrium is unique.

The equilibria is the same if firms can use multiple step bids per firm. Given the uniform pricing, the price-setting firm(s) cannot do better if the other firms submit the lowest bid for all their bins. The equilibria in the case of supply function bidding are essentially the same. Firms can again submit (almost) flat bids at the marginal cost level and almost flat bids at the price cap level.

Proof of Corollary 6

For part (a), substituting,

$$p_m^* = \arg \max_{0 \leq p \leq \Psi} (p - c) \left[\bar{Q} - \frac{u}{(\Psi - v)}(p - v) - (n - 1)\frac{K}{n} \right].$$

Clearly, p_m^* is never binding from below. The Kuhn-Tucker conditions,

$$\bar{Q} - \frac{u}{(\Psi - v)}(p - v) - (n - 1)\frac{K}{n} - (p - c)\frac{u}{(\Psi - v)} - \lambda = 0 \text{ and } \lambda(\Psi - p) = 0$$

and the second derivative is negative (and therefore a maximum). If $\lambda = 0$, then

$$p^* = \frac{(\Psi - v)}{2u} \left(\bar{Q} + \frac{u(v + c)}{(\Psi - v)} - (n - 1) \frac{K}{n} \right),$$

whereas if $\lambda > 0$ then $p^* = \Psi$. Deriving, in the interior case,

$$\frac{\partial p^*}{\partial n} = -\frac{K}{n^2} \frac{1}{2 \frac{u}{(\Psi - v)}} < 0.$$

For part (b), evaluating the demand at p^* , we have

$$Q(p^*) = \frac{1}{2} \left(\bar{Q} - \frac{u(c - v)}{(\Psi - v)} + \frac{(n - 1)K}{n} \right) = \frac{Q(c) + (n - 1)k_n}{2},$$

and deriving with respect to n

$$\frac{\partial Q(p_m^*)}{\partial n} = \frac{K}{2n^2} = \frac{k_n}{2n}.$$

Now, we have that the derivative of the relative quantities sold (and therefore profits) of a price-setting with respect to a non-price setting firm are

$$\frac{\partial [Q(p_m^*)/k_n - (n - 1)]}{\partial n} = \frac{\frac{\partial Q(p_m^*)}{\partial n} k_n - \frac{\partial k_n}{\partial n} Q(p_m^*)}{k_n^2} - 1$$

and given that $\frac{\partial k_n}{\partial n} = -k_n/n$, we have

$$\frac{\partial [Q(p_m^*)/k_n - (n - 1)]}{\partial n} = \frac{Q(c) + nk_n}{2nk_n} - 1 = \frac{Q(c) - nk_n}{2nk_n} < 0$$

since $Q(p) \leq K = nk_n$ for all p , and in particular that $Q(c) \leq nk_n$.

References

- [1] Abdulkadiroglu, Atila , Parag A. Pathak, and Alvin E. Roth (forthcoming): "Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match," *American Economic Review*.
- [2] Adner R. and Daniel Levinthal (2001): "Demand heterogeneity and technology evolution: implications for product and process innovation," *Management Science*, vol. 47, pp. 611 - 628.
- [3] Bagnall, A., Smith, G., 2005. A multi-agent model of the UK market in electricity generation. *IEEE Transactions on Evolutionary Computation* 9 (5), 522–536.
- [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] Banal-Estañol, A. and Augusto Rupérez Micola (2009): "Composition of electricity generation portfolios, pivotal dynamics and market prices". *Management Science* vol. 55, issue11, pp. 1813-1831.
- [6] Blumsack, S., Dmitri Perekhodtsev, Lester Lave (2002): "Market power in deregulated wholesale electricity markets: Issues in measurement and the cost mitigation", *The Electricity Journal*, vol. 15, issue 9, pp. 1-24.
- [7] 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.
- [8] Bower, J. and Derek W. Bunn (2000): "Model-based comparisons of pool and bilateral markets for electricity". *Energy Journal*, vol. 21, no. 3.
- [9] Bower, J. and Derek W. Bunn (2001): "Experimental analysis of the efficiency of uniform-price versus discriminatory auctions in the England and Wales electricity market," *Journal of Economic Dynamics and Control*, vol. 25, issues 3-4, pp. 561—592.
- [10] Bower, J., Derek W. Bunn and Claus Wattendrup (2001): "A model-based analysis of strategic consolidation in the German electricity industry." *Energy Policy*, volume 29, no. 12, pp. 987—1005.
- [11] Brenner, T. (2006): "*Agent Learning Representation: Advice on Modelling Economic Learning.*" in: L. Tesfatsion, K.L. Judd (eds.). *Handbook of Computational Economics*, Vol. 2. Amsterdam: Elsevier Science, pp. 895-947.
- [12] Brown, G.W. (1951): "Iterative solutions of games by fictitious play" In *Activity Analysis of Production and Allocation*, T.C. Koopmans (Ed.), New York: Wiley.
- [13] Bunn, D. W. and Christopher J. Day (2009): "Computational modelling of price formation in the electricity pool of England and Wales". *Journal of Economic Dynamics and Control*, vol. 33, pp. 363-376.
- [14] Bunn, D. W. and Maria Martoccia (2005): "Unilateral and collusive market power in the electricity pool of England and Wales", *Energy Economics*, Volume 27, Issue 2, Pages 305-315.
- [15] Bunn, D. and Fernando Oliveira (2001): "Agent-Based simulation: an application to the New Electricity Trading Arrangements of England and Wales", *IEEE Transactions on Evolutionary Computation*, Volume 5, Number 5, pp. 493-503.
- [16] Bunn, D. and Fernando Oliveira (2003): "Evaluating individual market power in electricity markets via agent-based simulation", *Annals of Operations Research*, vol. 121, Numbers 1-4, pp. 57-77.
- [17] Bunn, D. and Fernando Oliveira (2007): "Agent-based analysis of technological diversification and specialisation in electricity markets", *European Journal of Operational Research*, issue 181 (3), p. 1265-1278.

- [18] Bunn, D. and Fernando Oliveira (2008): "Modelling the impact of market interventions on the strategic evolution of electricity markets". *Operations Research*.
- [19] Bushnell, J., Christopher Knittel, and Frank Wolak (1999): "Estimating the opportunities for market power in a deregulated Wisconsin electricity market", *The Journal of Industrial Economics*, vol.47.
- [20] Camerer, C. and Teck, H. Ho (1999): "Experience-weighted attraction learning in normal form games". *Econometrica*, vol. 67, pp. 837-874.
- [21] Cincotti, S., Guerci, E., Raberto, M., (2005): "Price dynamics and market power in an agent-based power exchange." In: Abbott, D., Bouchaud, J.-P., Gabaix, X., McCauley, J.-L. (Eds.): Noise and Fluctuations in Econophysics and Finance. *Proceedings of the SPIE*, vol. 5848, pp. 233-240.
- [22] 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.
- [23] 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.
- [24] Day, C. and Derek Bunn (2001): "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, no. 2, pp. 123-141.
- [25] Denrell, J. (2004). "Random walks and sustained competitive advantage." *Management Science*, vol. 50, issue 7, pp. 922-934.
- [26] Duffy, J. (2006): "Agent-based models and human subject experiments," in: L. Tesfatsion and K.L. Judd, eds.: *Handbook of Computational Economics Vol. 2*. Handbooks in Economics Series, (Amsterdam: Elsevier), pp. 949-1011.
- [27] 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*.
- [28] Erev, I. and Alvin E. Roth. (1998): "Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria". *The American Economic Review*, vol. 88, no. 4, pp. 848-881.
- [29] Erev, I., Alvin E. Roth, Robert L. Slonim and Greg Barron (2007): "Learning and equilibrium as useful approximations: accuracy of prediction on randomly selected constant sum games". *Economic Theory*, vol. 33, pp. 29-51.
- [30] Ert, E. and Ido Erev (2007): "Replicated alternatives and the role of confusion, chasing and regret in decisions from experience", *Journal of Behavioural Decision Making*, vol. 20, pp. 305-322.

- [31] Fabra, N., Nils-Henrik von der Fehr and David Harbord (2006): "Designing electricity auctions". *RAND Journal of Economics*, Vol.37 (1), pp. 23-46.
- [32] Fagiolo, G., Christopher Birchenhall, and Paul Windrum (Eds.) (2007): "Empirical Validation in Agent-Based Models: Introduction to the Special Issue", *Computational Economics*, Volume 30, Number 3.
- [33] 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.
- [34] Fezzi, C. and Derek W. Bunn (2006): "Structural analysis of high frequency electricity demand and supply interactions". Available at SSRN: <http://ssrn.com/abstract=1241703>
- [35] Fudenberg, D. and D. Levine (1998): *The Theory of Learning in Games*, MIT Press.
- [36] Garcia, A., Campos-Nañez, E. and Reitzes, J. (2005): "Dynamic pricing and learning in electricity markets". *Operations Research*, vol. 53, no. 2, pp. 231—241.
- [37] Genc, T. and S. S. Reynolds (2005): "Supply function equilibria with pivotal electricity suppliers". Working paper, University of Guelph.
- [38] Green, R. and David M. Newbery (1992): "Competition in the British electricity spot market". *The Journal of Political Economy*, vol. 100, no. 5, pp. 929-53.
- [39] Hobbs, B. and J. S. Pang (2007): "Nash-Cournot equilibria in electric power markets with piecewise linear demand functions and joint constraints". *Operations Research*, vol. 55, no. 1, pp. 113-127.
- [40] Holmberg, P. (2007): "Supply function equilibrium with asymmetric capacities and constant marginal costs", *Energy Journal*, vol. 28, no. 2, pp. 55-82,.
- [41] Holmberg, P. (2008): "Unique supply function equilibrium with capacity constraints", *Energy Economics*, vol. 30, no 1, pp. 148-172.
- [42] Kimbrough, S. and Frederic H. Murphy (2009): "Strategic bidding of offer curves: an agent-based approach to exploring supply curve equilibria." Working paper.
- [43] Klemperer, P. and Margaret A. Meyer (1989): "Supply function equilibria in oligopoly under uncertainty", *Econometrica*, vol. 57, no. 6, pp. 1243-1277.
- [44] Li, H., Junjie Sun and Leigh S. Tesfatsion (2009): "Separation and volatility of locational marginal prices in restructured wholesale power markets," ISU Economics working paper #09009, June 9.
- [45] Leombruni, R., Richiardi, M., Saam, N. J., & Sonnessa, M. (2006): "A common protocol for agent-based social simulation." *Journal of artificial societies and social simulation*, no. 9, issue 1.
- [46] Macal and Michael North (2005): "Tutorial on agent-based modeling and simulation", *Proceedings of the 37th Winter simulation conference*, Pages: 2 - 15, ISBN:0-7803-9519-0.

- [47] Marks, R. (2006): "Market design using agent-based models," in: Leigh Tesfatsion & Kenneth L. Judd (ed.), *Handbook of computational economics*, vol. 2, Elsevier, pp. 1339-1380.
- [48] Marks, R.E. (2007): "Validating simulation models: A general framework and four applied examples". *Computational Economics*, 30(3), pp. 265-290.
- [49] 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.
- [50] Midgley D.F., Marks R.E., and Kunchamwar D. (2007): "The building and assurance of agent-based models: An example and challenge to the field", *Journal of Business Research, Special Issue: Complexities in Markets*, 60(8), pp. 884-893.
- [51] Nanduri, Vishnuteja and Tapas K. Das (2007): "A reinforcement learning model to assess the market power under auction-based energy bidding," *IEEE Transactions on Power Systems* 22(1):85-95.
- [52] Newbery, D. (1991): "Capacity-constrained supply function equilibria: competition and entry in the electricity spot market", mimeo University of Cambridge.
- [53] Nicolaisen, J., Valentin Petrov and Leigh 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.
- [54] Noe, T. H., Michael J. Rebello and Jun Wang (2003): "Corporate financing: An artificial agent-based analysis," *The Journal of Finance*, Vol. 58, No. 3, pp. 943-973.
- [55] Oliveira, F. and Derek W. Bunn (2007): "Agent-based analysis of technological diversification and specialisation in electricity markets." *European Journal of Operational Research*, 181 (3): 1265-1278, 2007.
- [56] 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.
- [57] Pouget, S. (2007): "Adaptive traders and the design of financial markets," *Journal of Finance*, vol. 62, issue 6, pp. 2835 - 2863.
- [58] Richter, C.W., Sheblé, G.B. (1998): "Genetic algorithm evolution of utility bidding strategies for the competitive marketplace." *IEEE Transactions on Power Systems* 13(1) (1), 256-261.
- [59] Rivkin, J. W., and Nicolaj Siggelkow (2003): "Balancing search and stability: interdependencies among elements of organizational design," *Management Science* vol. 49, no. 3 pp. 290-311.
- [60] Roth, A. E. and Elliott Peranson (1997): "The effects of the change in the NRMP matching algorithm," *Journal of the American Medical Association*, pp. 729-732.
- [61] Roth, A. E. (2002): "The economist as engineer: game theory, experimentation, and computation as tools for design economics," *Econometrica*, vol. 70(4), pages 1341-1378.

- [62] Roth, A.E. and Ido Erev (1995): "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term." *Games and Economic Behavior*, vol. 8, issue 1, 1995, Pages 164-212.
- [63] Roth, A. E., Tayfun Sonmez, and M. Utku Unver (2004): "Kidney exchange," *Quarterly Journal of Economics*, vol. 119, issue 2, pp. 457-488.
- [64] Rothkopf, M. H., Alexander Pekeč, Ronald M. Harstad (1998): "Computationally manageable combinatorial auctions". *Management Science*, vol. 44, issue 8, pp. 1131–1147.
- [65] Rothkopf, M. H. (2002): "Control of Market Power in Electricity Auctions," *The Electricity Journal*, vol. 15, issue 8, pp. 15-24.
- [66] Rupérez Micola, A. and Derek 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.
- [67] Rupérez Micola, A., Albert Banal-Estañol and Derek W. Bunn (2008): "Incentives and coordination in vertically related energy markets", *Journal of Economic Behavior and Organization*, vol. 67, issue 2, pp. 381—393.
- [68] Stoft, S. (2002): *Power System Economics: Designing Markets for Electricity*. Wiley, New York. ISBN: 978-0471150404.
- [69] Sun, J., Leigh Tesfatsion (2007): "Dynamic testing of wholesale power market designs: an open-source agent- based framework". *Computational Economics* 30 (3), 291–327.
- [70] Tesfatsion, L. and Kenneth Judd (2006): "*Handbook of computational economics, vol. 2: agent-based computational economics*". ISBN-13: 978-0-444-51253-6.
- [71] 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.
- [72] Veit, D. , Anke Weidlich, J. Yao, Shmuel Oren (2006): "Simulating the dynamics in two-settlement electricity markets via an agent-based approach", *International Journal of Management Science and Engineering Management*, volume 1(2), 83-97.
- [73] Visudhiphan, P., Ilic, M.D., (1999): "Dynamic games-based modeling of electricity markets." *Power Engineering Society 1999 Winter Meeting. IEEE*, vol. 1, pp. 274–281.
- [74] von der Fehr, N. and David Harbord (1993): "Spot market competition in the UK electricity industry". *Economic Journal*.
- [75] Weidlich, A. and Daniel Veit (2008): "A critical survey of agent-based wholesale electricity market models." *Energy Economics*, vol. 30, pp .1728–1759.
- [76] Wilson, R. (2002): "Architecture of power markets," *Econometrica*, vol. 70, issue 4, pp. 1299-1340.

No.	Paper	Journal	Research question	Wholesale supply representation	Wholesale demand representation	Behavioural algorithm
1	Bagnall and Smith (2005)	IEEE-Trans	Replication of human behaviour in the UK market	Single price bid per firm	Inelastic demand	Hierarchical classifier systems
2	Banal and Rupérez Micola (2009)	ManSci	Technological diversification and pivotal dynamics	Single price bid per plant and supply functions	Inelastic and elastic demands	Reinforcement learning (Roth-Erev)
3	Bower and Bunn (2000)	EnJor	Uniform price and discriminatory auctions	Single price bid per plant	Elastic demand	Reinforcement learning
4	Bower and Bunn (2001)	JEDC	Uniform price and discriminatory auctions	Single price bid per plant	Elastic demand	Reinforcement learning
5	Bower et al. (2001)	EnPol	Industrial consolidation in Germany	Single price bid per plant	Elastic demand	Reinforcement learning
6	Bunn and Martocchia (2005)	EnEcon	Tacit collusion	Single price bid per plant	Inelastic demand	Reinforcement learning
7	Bunn and Oliveira (2001)	IEEE - Trans	Provide pricing and strategic insights, ahead of NETA's introduction in the	Price / quantity bid per plant	Double-sided call auction	Best-response dynamics (Fictitious-play) with
8	Bunn and Oliveira (2003)	Annals of OR	Test whether two dominant generators could profitably influence wholesale price in the UK	Price / quantity bid per plant	Double-sided call auction	Combination of best-response dynamics (Fictitious-play) and reinforcement learning
9	Bunn and Oliveira (2007)	EJOR	Co-evolution of plant portfolios and spot prices	Quantities bidding	Elastic demand	Best-response dynamics (Cournot)
10	Bunn and Oliveira (2008)	OR	Co-evolution of plant portfolios and spot prices	Quantities bidding	Elastic demand	Best-response dynamics (Cournot)
11	Cincotti et al (2005)	Proc. SPIE	Effect of market microstructure and costs on prices	Supply functions	Inelastic demand	Reinforcement learning
12	Day and Bunn (2001)	JRE	Analyse divestures and their impact of market power in the UK pool	Piecewise supply functions	Elastic demand	Best-response dynamics (Cournot)
13	Day and Bunn (2009)	JEDC	Analyse market power in the UK pool	Piecewise supply functions	Elastic demand	Best-response dynamics (Cournot)
14	García et al. (2005)	OR	Dynamic price formation and hydropower behaviour	Single price bid per plant	Inelastic demand	Best-response dynamics (Fictitious-play)
15	Nanduri and Das (2007)	IEEE-Trans	Test of model on a simple electricity network	Single price bid per firm	Inelastic demand	Reinforcement learning (Roth-Erev)
16	Nicolaisen et. al. (2001)	IEEE-Trans	Market structure, market power and efficiency	Single price bid per firm	Double-sided call auction	Reinforcement learning (Roth-Erev)
17	Richter and Sheblé (1998)	IEEE-Trans	Wholesale market simulation	Single price bid per firm	Double-sided call auction	Genetic algorithm
18	Rupérez Micola and Bunn	JEBO	Horizontal cross-holdings	Single price bid per firm	Double-sided call auction	Reinforcement learning
19	Rupérez Micola et al. (2008)	JEBO	Vertical integration	Single price bid per firm	Inelastic demand	Reinforcement learning (Roth-Erev)
20	Sun and Tesfatsion (2007)	CompEcon	Interplay among market structure, protocols in relation to performance		Inelastic demand	Reinforcement learning (Roth-Erev)
21	Veit et al. (2006)	IJMEM	Dynamics in forward and spot electricity markets	Quantities bidding	Elastic demand	Reinforcement learning (Roth-Erev)
22	Visudhiphan and Ilic (1999)	IEEE meetings	Dynamic learning in power markets	Single price bid per firm and supply functions	Inelastic and elastic demands	Best-response dynamics

Table 1. Published papers

The table includes an alphabetical list of electricity agent-based modelling papers published as journal articles, with the year of publication and abbreviated journal title. In addition, the Table briefly summarises the research issue in each paper together with their supply bidding, demand representation and behavioural algorithm assumptions. Full citations appear in the references.

Demand assumptions

Parameters						Estimates								
Behaviour	Bid type	No. bins	u	Qbar	Upper sw. point	Beta0	t-stat. 0	Beta1	t-stat. 1	Beta2	t-stat. 2	Beta3	t-stat. 3	F - stat
RL	Stepwise	1	0	80	5	226,87***	136,73	-154,95***	-76,13	-38,70***	-50,38	35,13***	45,00	8212,06***
RL	Stepwise	1	0	85	7	214,96***	215,35	-135,23***	-69,86	-26,59***	-88,34	24,37***	69,50	8761,81***
RL	Stepwise	1	0	90	10	197,31***	254,29	-73,75***	-15,47	-13,30***	-86,59	9,69***	20,57	4928,32***
RL	Stepwise	1	5	80	7	206,12***	134,84	-155,16***	-52,35	-33,73***	-73,20	32,61***	60,74	3672,02***
RL	Stepwise	1	5	85	10	178,75***	107,08	-112,55***	-10,97	-17,54***	-53,06	16,40***	16,18	1340,62***
RL	Stepwise	1	5	90	12	173,95***	128,05	-66,17**	-2,04	-10,99***	-50,22	8,33**	2,95	1027,05***
RL	Stepwise	1	10	80	10	155,58***	70,59	-116,80***	-8,62	-17,06***	-39,08	17,32***	12,94	663,89***
RL	Stepwise	1	10	85	12	150,84***	78,53	-107,04**	-2,33	-11,64***	-37,60	12,63***	3,16	541,47***
RL	Stepwise	1	10	90	12	160,69***	102,99	-53,49	-1,43	-10,16***	-40,43	7,44**	2,29	634,33***

Supply assumptions

Parameters						Estimates								
Behaviour	Bid type	No. bins	u	Qbar	Upper sw. point	Beta0	t-stat. 0	Beta1	t-stat. 1	Beta2	t-stat. 2	Beta3	t-stat. 3	F - stat
RL	Stepwise	1	0	80	5	228,00***	145,26	-156,18***	-81,11	-39,09***	-53,80	35,54***	48,12	9246,67***
RL	Stepwise	2	0	80	5	221,17***	162,13	-113,89***	-68,06	-34,58***	-54,76	31,22***	48,64	6105,71***
RL	Stepwise	3	0	80	5	205,16***	161,18	-83,90***	-53,73	-24,35***	-41,32	21,61***	36,09	4301,83***
RL	SF	1	0	80	5	223,84***	61,97	-154,39***	-34,84	-22,90***	-13,70	16,32***	9,60	3705,21***
RL	SF	2	0	80	5	219,10***	35,21	-75,89***	-9,94	-20,85***	-7,24	6,94**	2,37	1208,01***
RL	SF	3	0	80	5	209,95***	31,60	-77,35***	-9,49	-14,75***	-4,79	1,70	0,54	1106,47***

Behavioural assumptions

Parameters						Estimates								
Behaviour	Bid type	No. bins	u	Qbar	Upper sw. point	Beta0	t-stat. 0	Beta1	t-stat. 1	Beta2	t-stat. 2	Beta3	t-stat. 3	F - stat
BR	Stepwise	1	0	80	5	149,14***	34,16	-88,99***	-16,61	-17,66***	-8,74	11,85***	5,77	1028,69***
RL	Stepwise	1	0	80	5	227,27***	141,33	-156,57***	-79,37	-38,55***	-51,79	35,14***	46,44	8817,67***
FP	Stepwise	1	0	80	5	136,34***	112,09	-7,92***	-5,31	-2,33***	-4,14	0,3	0,52	586,02***

Table 2: Hypotheses' tests of breaking point regression estimates

The parameters considered in each simulations batch are marked with boxes. The parameters outside those boxes stay constant. The estimates correspond to the threshold equation specified in eq. 5. Beta1 and Beta3 estimate the post-breaking point changes in intercept and slope. They support the existence of a breaking point when they are statistically significant: **significant at the 0.05 level; ***significant at the 0.01 level.

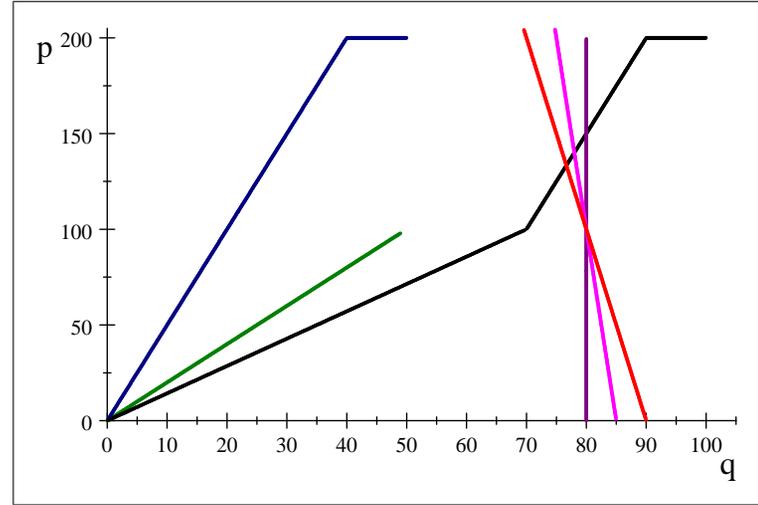
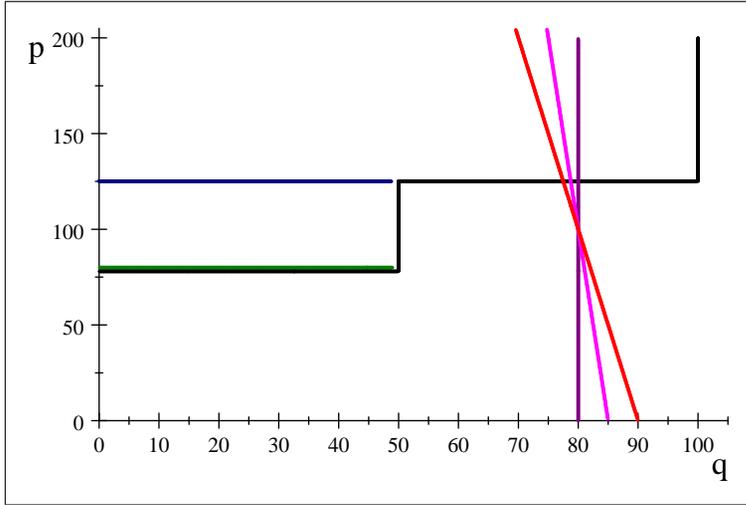


Figure 1: Examples of alternative demand and supply schedules

Three potential demands with no, low and high elasticity (purple, magenta, red, respectively) and two individual bids (green and blue) and their horizontal aggregation (black)

are included in each panel. The left one is for stepwise bidding and the right one supply function bidding

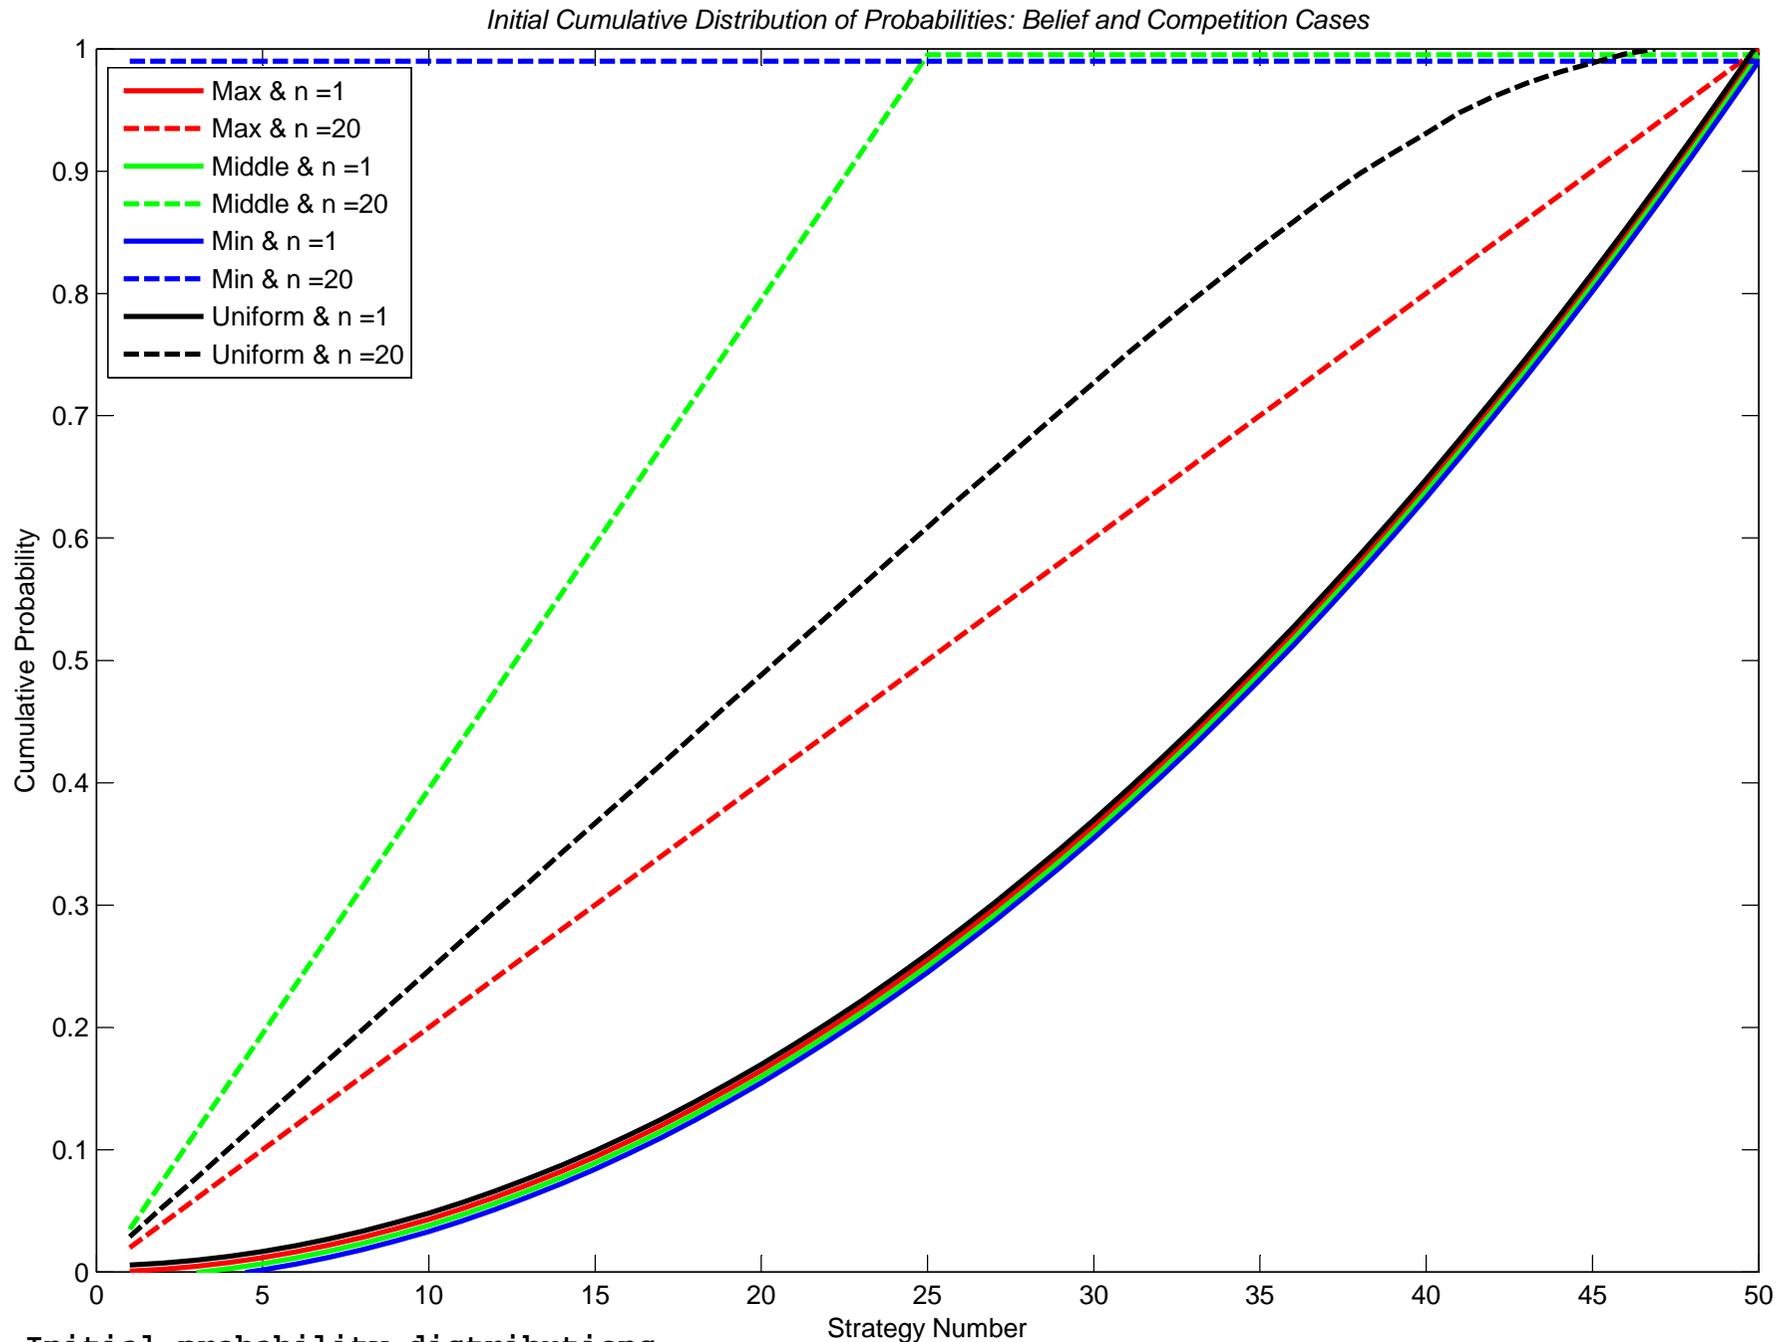


Figure 2: Initial probability distributions

Initial cumulative distributions if prior beliefs are that opponents will initially bid (i) maximum price, (ii) minimum price, (iii) middle of the distribution, and (iv) randomly. Competition cases: $n = 1$ and $n = 12$

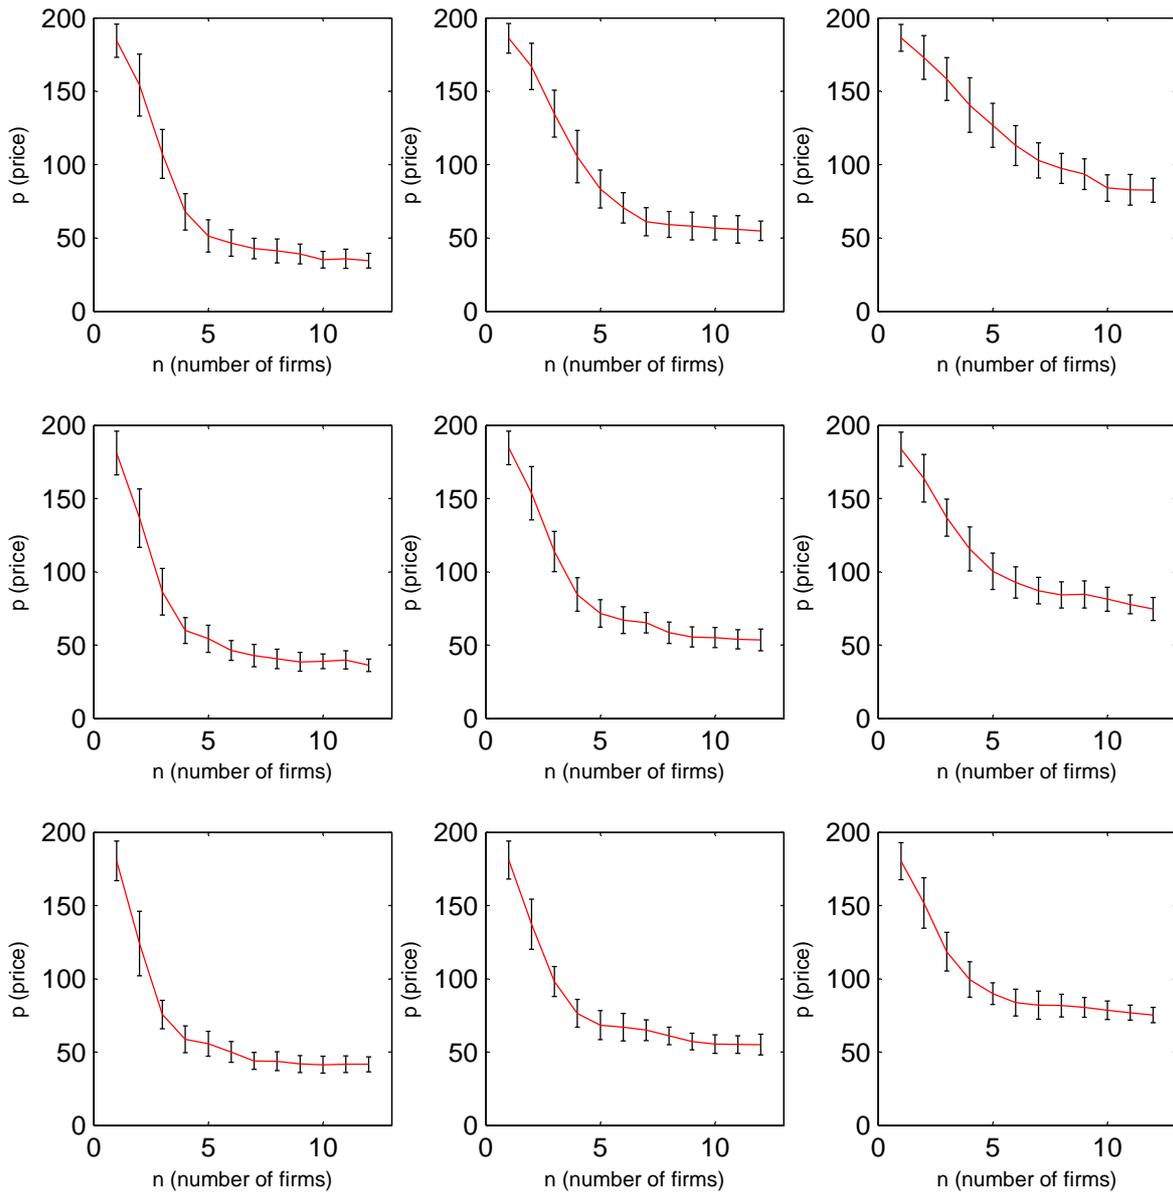


Figure 3: The influence of demand specifications on prices

Mean (+/- two standard deviations) of prices when demand levels are $Q = \{80, 85, 90\}$ (columns 1,2,3) and elasticities $u = \{0, 5, 10\}$ (rows 1,2,3)

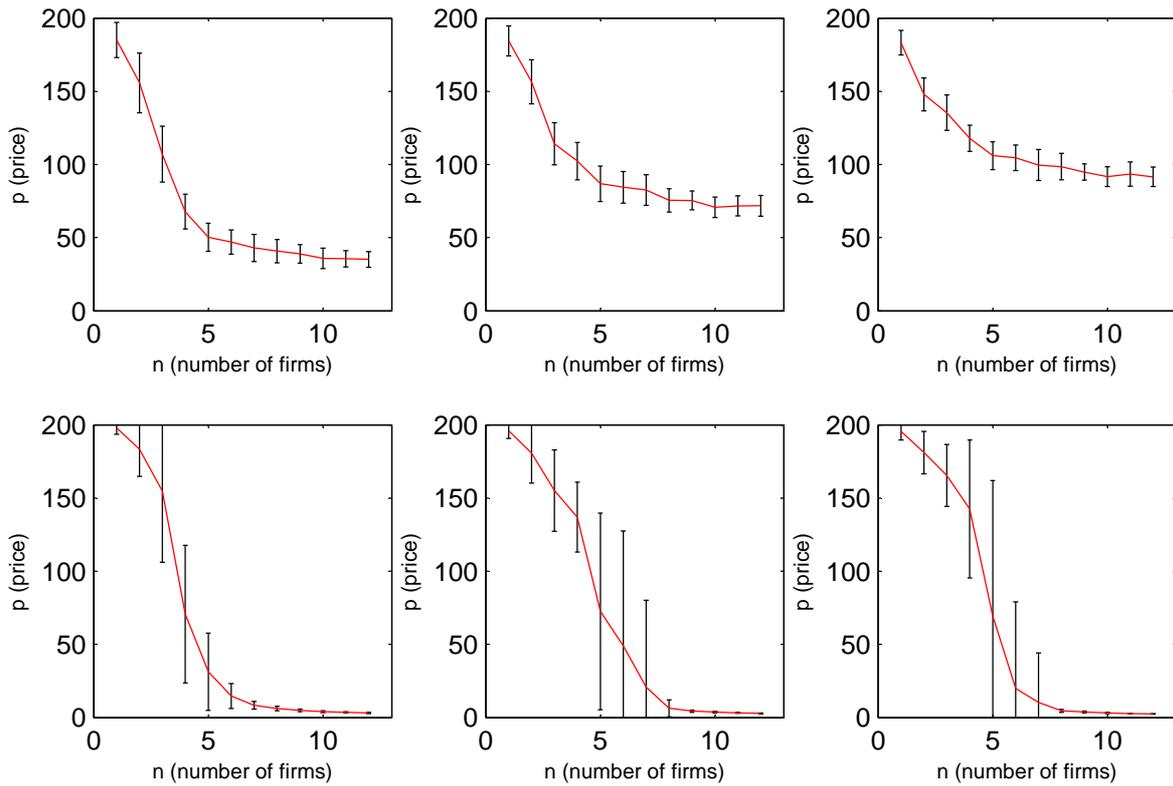


Figure 4: The influence of supply specifications on prices

Mean (+/- two standard deviations) of prices in the stepwise (upper) and supply function (lower row) cases with number of bins $m = \{1, 2, 3\}$ (columns 1, 2, 3).

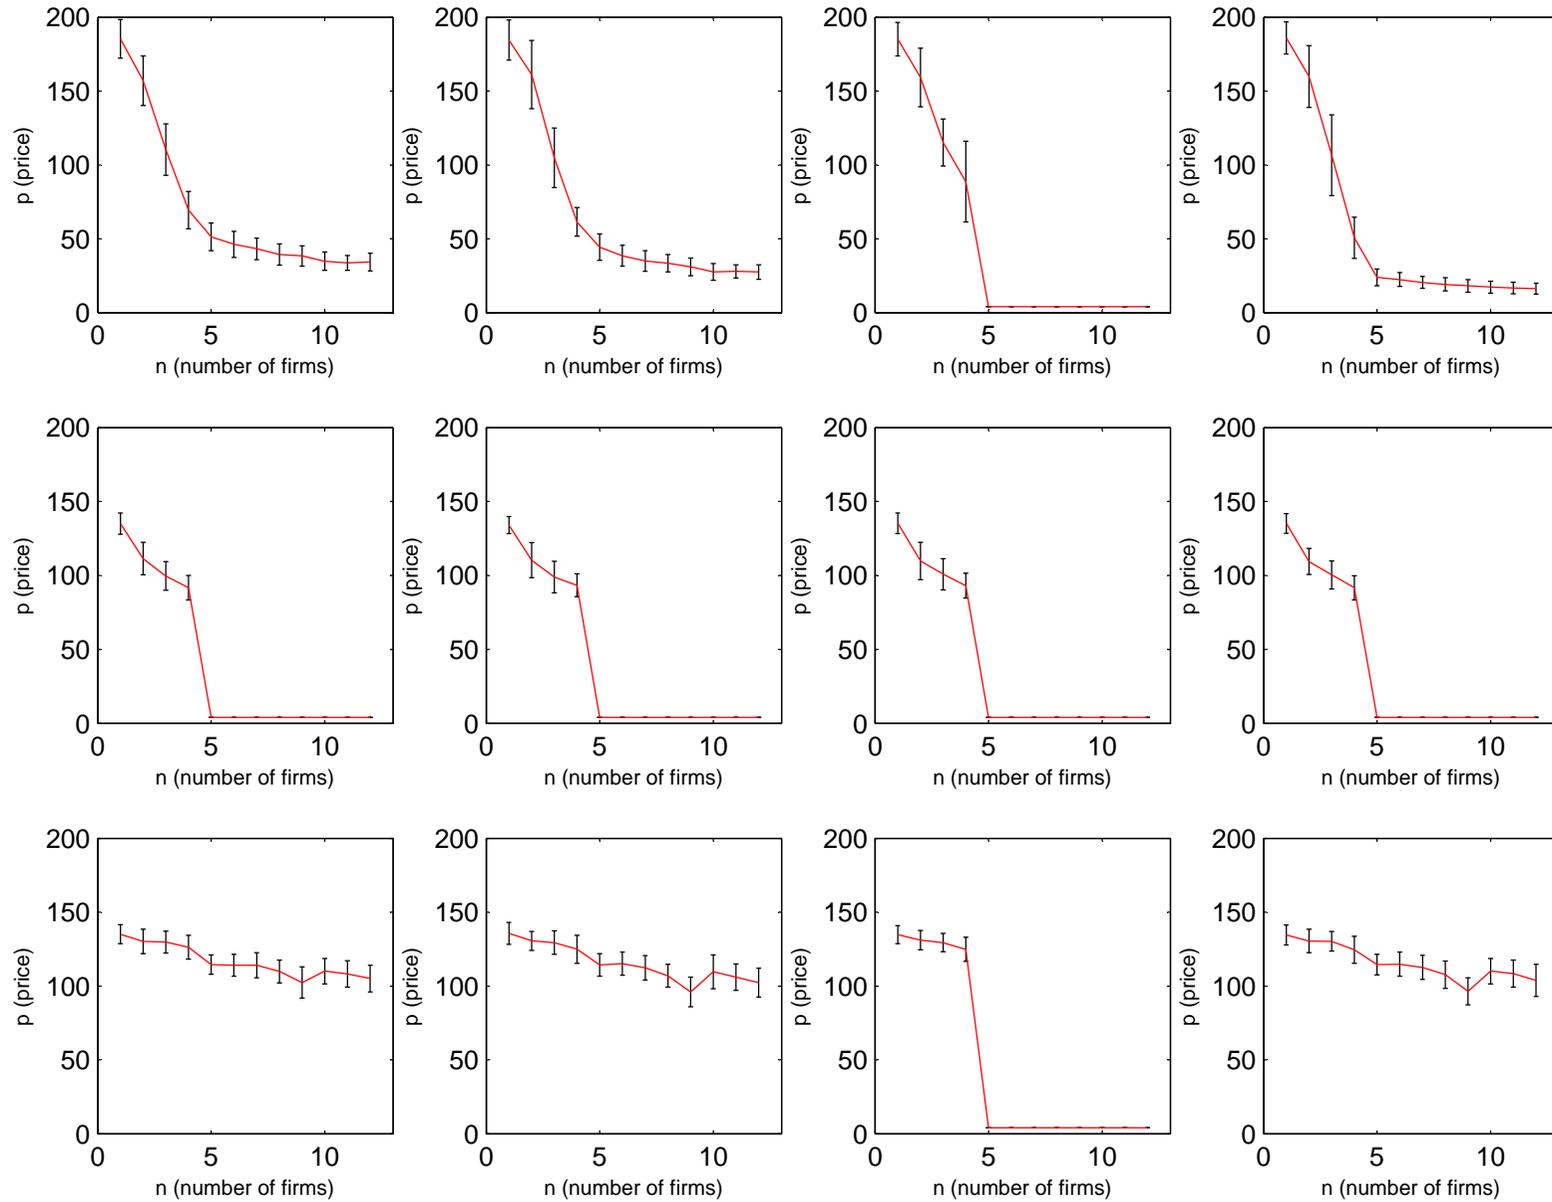


Figure 5: The influence of behavioral assumptions on prices

Mean (+/- two standard deviations) of prices in reinforcement learning (row 1), best-response (row 2), and fictitious play (row 3) assuming that prior beliefs are maximum price (column 1), random (column 2), minimum bid (column 3) and middle of the distribution (column 4)

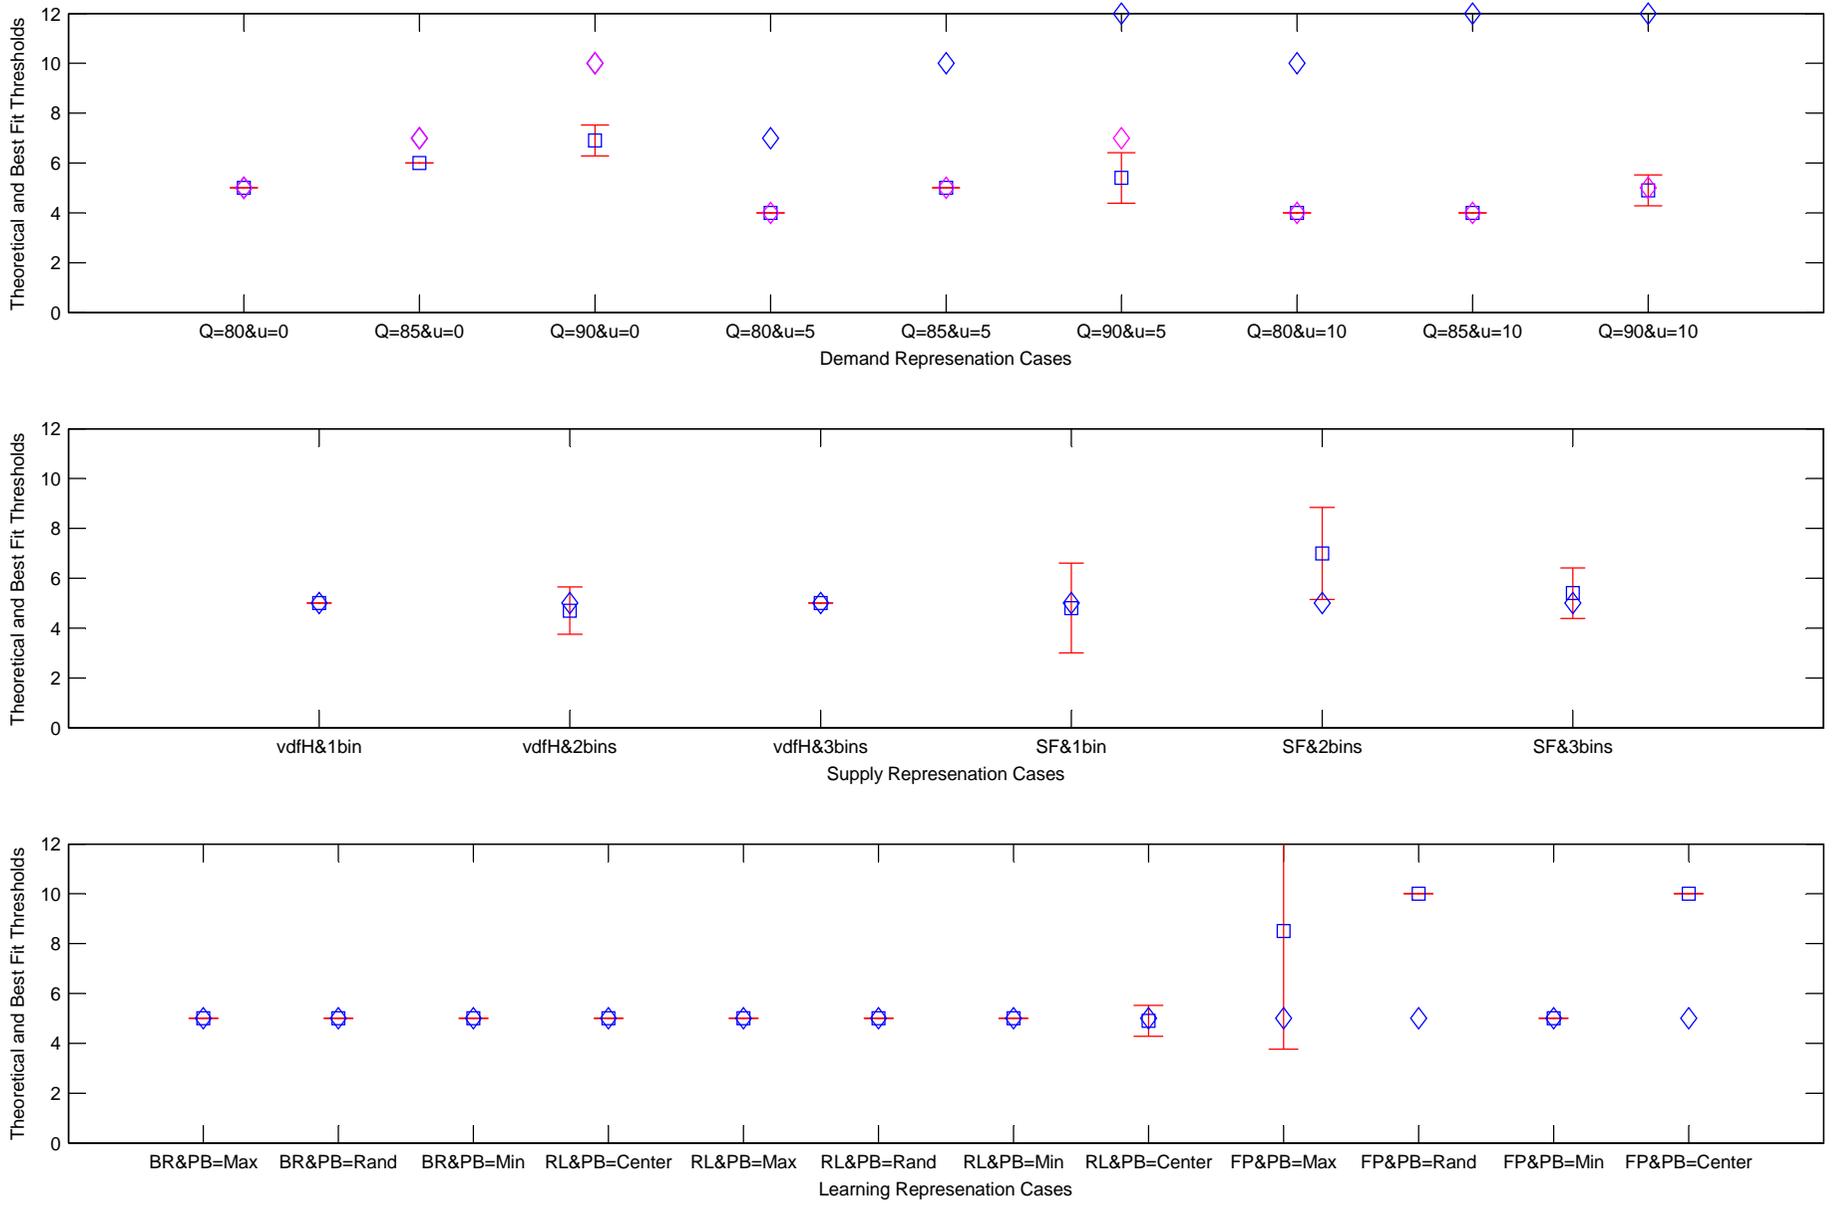


Figure 6: Comparison of theoretical and simulated switching points
 Optimal breaking point (+/- two standard deviations) for the demand (upper panel), supply (middle panel) and behavioural alternatives (lower panel)

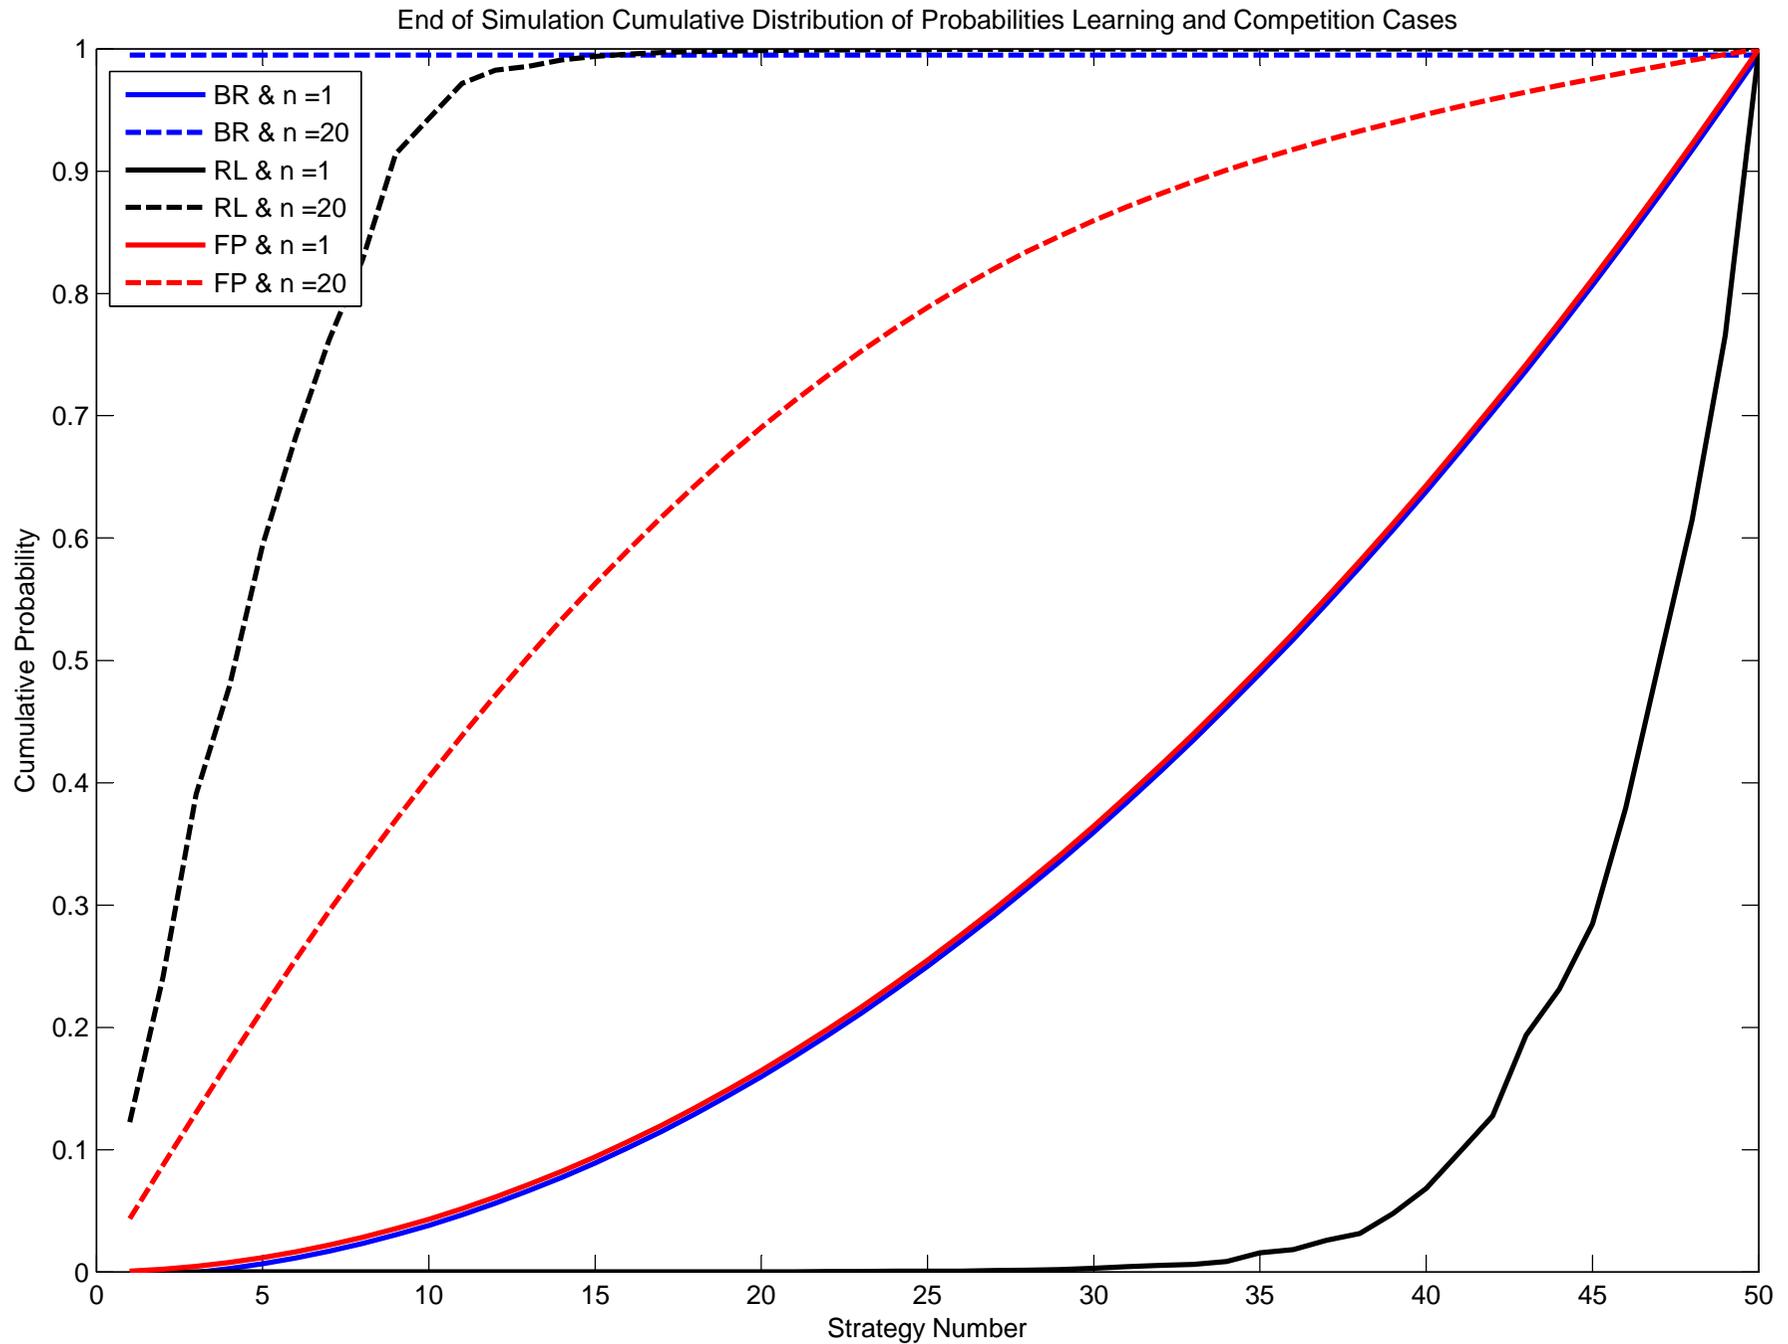


Figure 7: Latent intensity of competition

End of simulation cumulative distribution of probabilities if agents learn according to reinforcement learning, best response and fictitious play. Competition cases: $n = 1$ and $n = 12$