



City Research Online

City St George's, University of London

Citation: Banal-Estanol, A. & Rupérez Micola, A. (2011). Behavioural simulations in spot electricity markets. *European Journal of Operational Research*, 214(1), pp. 147-159. doi: 10.1016/j.ejor.2011.03.041

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/7693/>

Link to published version: <https://doi.org/10.1016/j.ejor.2011.03.041>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Behavioural simulations in spot electricity markets*

Albert Banal-Estañol[†]

Augusto Rupérez Micola

February 7, 2011

Abstract

We study the consistency of behavioural simulation methods used to model the operations of wholesale electricity markets. We include different supply and demand representations and propose the Experience-Weighted Attractions method (Camerer and Ho, 1999) to encompass several behavioural paradigms. We compare the results across assumptions and to standard economic theory predictions. The match is good under flat and upward-slopping supply bidding, and also for plausible demand elasticity assumptions. Learning is influenced by the number of bids per plant and the initial conditions. The simulations perform best under reinforcement learning, less well under best-response and especially poorly under fictitious play. The overall conclusion is that simulation assumptions are far from innocuous. We link their performance to underlying features, and identify those that are better suited to model liberalised electricity markets.

1 Introduction

The electricity industry is undergoing fundamental changes leading to a more liberal regime and the alteration of its business logic. As part of the process, vertically integrated utilities and simple transfer pricing rules are substituted by sophisticated financial trading arrangements. De-regulated electricity markets feature imperfect competition, very low demand elasticity, discontinuously convex supply functions, high-frequency repeated trading, several production technologies and high potential for collusion (Wilson, 2002). As a result, prices in the new electricity hubs are volatile and often characterised by strategic behaviour and learning, which poses new challenges both for the operations and scholarly study of the industry.

Simulations have hence emerged as a natural way to study the operations of de-regulated electricity markets. An important part of the literature employs behavioural methods, with firms modelled as

*This paper was presented at the EURO (2009, Bonn), INFORMS (2009, San Diego), Computational Economics and Finance (2010, London) conferences, as well as in seminars at Cass, Darden, UCH, Universidade Católica Portuguesa and UPF. We thank seminar participants, two referees, Ido Erev, Derek Bunn, Pär Holmberg, Andreas Krause, Leigh Tesfatsion and Ann van Ackere for their comments.

[†]Both authors are with Universitat Pompeu Fabra and the Barcelona GSE. Albert Banal-Estañol is also with City University, albert.banalestanol@upf.edu, augusto.ruperezmicola@upf.edu

interacting, boundedly-rational agents (see Marks, 2006; Weidlich and Veit, 2008 for surveys).¹ The literature includes models commissioned by large firms (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).

One of the advantages of behavioural simulations is that they are tailored to fit closely the operations in each industry. However, this is a disadvantage when it comes to understanding the factors driving the results because there is no consensus on the techniques appropriate for each situation. As a consequence, simulation results are often not comparable (Fagiolo et al., 2007) and these methods are struggling to reach their full potential (e.g. Leombruni et al., 2006). First, some papers assume that firms behave 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. 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.

This paper explores the consistency of the behavioural simulation techniques used in the literature to model the operations of the new wholesale electricity auctions. We investigate the effects of the assumptions on simulation outcomes and how these outcomes compare to simple, empirically-supported, theoretical predictions. Specifically, we cast light on whether the results are consistent with the standard claim that pivotal dynamics determine 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. 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.

We adopt a stylised setting that allows us to include alternative implementations of demand, supply and firm 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. Firms’ 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.

¹This trend is part of a “economic engineering” approach (Roth, 2002). Behavioural simulations have also been used to model competitive strategy (e.g. Denrell, 2004), innovation (e.g. Adner and Levinthal, 2001), financial markets (e.g. Noe et al. 2003; Pouget, 2007), or business organisation (e.g. Rivkin and Siggelkow, 2001). See Tesfatsion and Judd (2006) and <http://www.econ.iastate.edu/tesfatsi/ace.htm> for information about agent-based methods.

Simulation outcomes are consistent with our theoretical predictions under flat and supply function bidding, and under several plausible elasticities. However, the performance of the simulations is influenced by the number of bids per plant and the initial conditions. The performance of fictitious play is poor, and it is clearly outperformed by best-response and especially reinforcement learning. The results call into question a large part of the extant behavioural electricity research and can potentially enhance the practical implementation of these techniques in the operation of the energy industry. We also find some evidence suggesting that experimental research can help us identify the most suitable assumptions in market simulations. Weighted fictitious play, and especially power choice rules and regret-feedback models improve over some of the standard models.

This paper is part of a new literature examining the consistency of behavioural simulations in various de-regulated 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 mainly 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 firm 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 supplementary material.

2 Behavioural electricity modelling alternatives

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

2.1 Supply bidding

Bertrand and Bertrand with capacity constraints are generally not considered suitable in the 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.

Others allow one bid per plant. Bower and Bunn (2000, 2001) simulate the transition between the British pool and the New Electricity Trading Arrangements (NETA) 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, firms' 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 equilibria complicates the comparison of simulation and equilibrium results. For example, the SF model has little predictive value if 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 (e.g. von der Fehr and Harbord, 1993; Crawford et al., 2006).

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 behavioural 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. 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

Behavioural simulation models require rules to govern firm 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 behavioural simulations literature is based on adaptive learning algorithms mainly derived from psychology.

Some previous work uses reinforcement learning (RL). In RL, firms 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 (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 reality 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).² 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).³ To our knowledge, there is no research on whether the results obtained with RL, FP and BR differ substantially in the electricity context. The use of weighted fictitious play, power choice rules and regret models may improve the models' quality.

Finally, most papers do not report initial conditions. In those that do, the standard approach is to use a uniform initial probability distribution for all elements of the action space (e.g. 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 alternative starting conditions.

²The term "Cournot" does not refer to quantity bidding but to the classic *tatônnement* process leading to equilibrium. To avoid confusion, we refer to this algorithm as "best response" (BR).

³Several papers depart from those models. Bower and Bunn (2000, 2001), Bower et al. (2001) and Bunn and Martoccia (2005) all use local adjustment. Bagnall and Smith (2005) use hierarchical classifier systems and Richter and Sheblé (1998) use genetic algorithms.

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 firm behaviour, which yield many of the literature’s 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. Although relevant in the long term, we do not deal with capacity expansion, long-term contracts, ancillary and capacity payments. For ease of comparison to existing research, we have also left out the network issues inherent in the operations of electricity utilities, i.e. we assume an un-congested network as is done in most of the existing simulations’ literature.

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 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 of the Supplementary Appendix shows examples of demand with $\bar{Q} = 8$, $\Psi = 20$, $v = 10$ and $u = 0, 0.5$, and 1 . The three demands are $Q(p) = 8$ ($u = 0$, in purple), $Q(p) = 8 - \frac{0.5}{10}(p - 10)$ ($u = 0.5$, in magenta), and $Q(p) = 8 - \frac{1}{10}(p - 10)$ ($u = 1$, in red).

3.3 Supply representations

Supply schedules vary along two dimensions. First, firms are allowed to submit either flat bids (“step-wise bidding” case) or increasing supply schedules (“SF bidding” case). Second, in line with Day and Bunn, 2001, and Hobbs and Pang, 2007, we consider multi-step schedules. 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. 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 0 and $\pi/2$ radians. The schedules consist of the linear curves from $(0, c)$ 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.⁴ 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 intersects the market demand and supply functions to set $\hat{p}$. Under stepwise bids, 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 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. The Supplementary Appendix includes a formal derivation of the market clearing process. The two panels in Figure 1 of the Supplementary Appendix show hypothetical bidding examples with $n = 2$, $m = 1$, $K = 10$, $\Psi = 20$, and $c = 0$. The market supply function (black line) is the horizontal addition of the individual functions (blue and green).

3.4 Firm behaviour representations

We use the Experience-Weighted Attraction (EWA) adaptive learning mechanism (Camerer and Ho, 1999). This behavioural model nests RL, FP and BR as special cases. It assumes that each feasible action of each bin has a numerical attraction. The attractions generate a bin-specific probability distribution. In each round, generators submit supply schedules according to these bin-specific probability distributions. Once the market clears, 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 firms 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

⁴We add a marginal amount to the denominator to avoid indeterminacy when $s = S$.

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.

Camerer and Ho (1999) estimated several parameters in stylised games, and preferred models that fall between the three extreme examples. We focus mainly on the extreme RL, BR and FP parametrizations for ease of comparison to the electricity literature. In the last section of the results, we also consider an alternative regret model similar to the one in Erev and Ert (2007) and Marchiori and Warglien (2008).

As is done in most of the electricity literature, we linearly map 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,

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

Note that this is a particular case of the power probability function used by Camerer and Ho (1999) in which the exponent is equal to one. In the same paper, Camerer and Ho (1999) also propose a power function, which we employ in the last section of our analysis.

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 firms 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 in the Supplementary Appendix reports the impact of prior beliefs on the initial probabilities in markets with one and twelve firms (assuming stepwise bidding, $m = 1$, $K = 10$, $\Psi = 20$, $S = 50$ and $c = 0$). Actions are identified on the horizontal axis. Firms 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. By definition, prior beliefs on others have no impact when there is only one firm.

To our knowledge, the literature only includes uniform initial probability distributions over all the elements of the action space, which implicitly implies that firms believe that their opponents will bid Ψ . Thus, take the case of twelve firms. In treatment (i), that is when firms 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, $P_{i,s}^j(1) \approx \frac{1}{S-1}$. In treatment (ii), when firms 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 firms 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 if they used a uniform distribution, assigning higher probability to lower prices.

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

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 firms' 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; 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 $Q(p) = 8$ for any p and $K = 10$, the switching point is $\hat{n}^u = \hat{n}^l = 5$ because if $n < 5$, $(n-1)k_n < 8$, and all firms are necessary to fulfill demand but no firm is pivotal when $n \geq 5$, $(n-1)k_n \geq 8$.

In the elastic case, the upper and lower switching points are different. Take for instance the case in which $\Psi = 20$, $v = 10$, $\bar{Q} = 8$, $u = 1$, $c = 0$ and $K = 10$. The minimum and maximum demands are $Q(20) = 7$ and $Q(0) = 9$. There is a lower switching point at $\hat{n}^l = 4$ since $(n-1)k_n < 7 = Q(20)$ for $n < 4$ and $(n-1)k_n \geq 7 = Q(20)$ for $n \geq 4$ and an upper switching point at $\hat{n}^u = 1$ because $(n-1)k_n < 9 = Q(0)$ for $n < 1$ and $(n-1)k_n \geq 9 = Q(0)$ for $n \geq 1$.

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 $\hat{n}^u$. It is unique if $n = 1 < \hat{n}^u$ (monopoly) 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.⁵ 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 “battle of the sexes” game. Coordination in this game can be low (Cooper et al., 1990), and especially difficult due to the payoff asymmetry in each of the equilibria (Crawford et al., 2008). 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 issues. 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 firms’ 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 .*

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 modelling assumptions aiming to reproduce spot electricity markets should be able to fulfill it.

5 Simulations

We first introduce the simulation parameters. and graphically compare the simulation outcomes with the theory. Then, we formally test whether the data for each specification features breaking points at

⁵There are also many mixed strategy Nash equilibria.

the predicted locations and, whether the predicted switching point is best-fitting.

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 $K = 10$, so individual capacities decrease from $k_1 = 10$ to $k_{12} = 0.833$. The price ceiling is $\Psi = 20$, with a grid of $S = 50$ possible actions and a demand rotation point $v = 10$. The convergence parameter is $\tau = 0.004$. 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 $\bar{Q} = 8$), firms using stepwise bidding with one bin ($m = 1$), and firms learning following RL with a uniform initial distribution.

To study the effect of the demand assumptions, we fix the reference supply (stepwise bidding and $m = 1$) and behavioural representations (RL and uniform initial distribution). We run simulations for the combinations of expected demand levels $\bar{Q} = \{8, 8.5, 9\}$ and elasticity parameters $u = \{0, 0.5, 1\}$. 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 demand ($\bar{Q} = 8$ and $u = 0$) and behavioural specifications (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 supply and demand specifications (stepwise, $m = 1$, $\bar{Q} = 8$ 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.

From Proposition 5, the monopoly prices are equal to the maximum price, $p_m^*(1) = \Psi = 20$. The upper switching points are $\hat{n}^u = \{5, 6, 10\}$, respectively for $\bar{Q} = \{8, 8.5, 9\}$, if $u = 0$ (inelastic demand), and $\hat{n}^u = \{7, 10, 12\}$ if $u = 0.5$ and $\hat{n}^u = \{10, 12, 12\}$ if $u = 1$.

5.2 Simulation results

Figures 1, 2 and 3 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 ($\bar{Q} = 8$, $u = 0$, stepwise bidding, $m = 1$, RL and a uniform initial distribution).

Demand specifications Figure 1 reports the demand results. Simulations use as demand levels $\bar{Q} = 8, 8.5$ or 9 (in columns one, two and three, respectively) and, as elasticity assumptions, $u = 0$,

0.5, or 1 (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.

Monopoly prices are close to $\Psi (= 20)$ 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 $\bar{Q} = \{8, 8.5, 9\}$. More competition has a small effect beyond those values, but $\bar{p}$ remains clearly above c , particularly in the case of tight capacity ($\bar{Q} = 9$, third column).

In the elastic cases (second and third rows), the theory predicted a break at the *upper* switching points: $\hat{n}^u = \{7, 10, 12\}$ for $u = 0.5$ and $\hat{n}^u = \{10, 12, 12\}$ for $u = 1$ for $\bar{Q} = \{8, 8.5, 9\}$, respectively. That is, as elasticity increases, we predict a break for higher levels of n . Instead, the simulations show a lower breaking point as the elasticity increases, particularly under tight capacity. As we will see, 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. Simulations are similar across rows, in terms of both price levels and the shape of the n to p relationship. However, under tight capacity, 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 2 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 $\bar{Q} = 8$ and $u = 0$ and RL with uniform initial distribution in all cases.

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 4, when $m = 2$ around 7 or 8 and when $m = 3$ around 10. If there are multiple bins, firms 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 firms 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 3 reports simulation results for the different behavioural specifications (assuming stepwise bidding with $m = 1$, and $\bar{Q} = 8$ 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. This implies that the initial probability distributions are uniform, lower than uniform, concentrated in the lowest action, and uniform over the bottom half, respectively.

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 RL break at the switching point is striking only if the prior is competitive (third column).

The pre-break relationships between n and $\bar{p}$ are not decreasing for BR and FP. In BR (second row), monopoly prices are far from Ψ (at about 14), 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 can be traced back to the algorithms' features. Under BR, prices have been shown to be competitive when no firm is pivotal. This is because under BR, the attractions of all the actions above those used by the other firms in the previous period are reduced to zero. Therefore firms choose lower actions and $\bar{p}$ can only stay constant or decrease. The resulting unravelling yields $\bar{p} = c$. When all firms are pivotal, equilibrium forces tend to increase $\bar{p}$ so that firms 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, and may even increase slightly. In RL, instead, there are no in-built unravelling and best-response mechanisms. 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 .

Fictitious play (third row) departs significantly from the predictions. Monopoly prices are around 14, are slightly increasing in n and stay patently above competitive levels. FP keeps most of the initial noise as it weights heavily initial periods with quasi-random outcomes. Firms 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 firms, 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. They are not decreasing due to the in-built best-response features of the algorithm. On balance, the n to p relationship is quite flat 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 10. For $n > 5$, prices converge to c . These results are to our knowledge the first on the robustness of simulation techniques to behavioural choices in electricity markets.

5.3 Threshold regressions

In this section, we carry out tests of whether the simulation data confirm our hypotheses. 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 $\hat{n}^u$,

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

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 first 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. Second, we test whether prices are decreasing in the number of firms before the breaking point and flatter thereafter, i.e. β_2 negative and β_3 positive.

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).

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 support for the first hypothesis. However, both best-response and fictitious play exhibit increasing firms-price relationships before the breaking point, i.e. $\beta_2 > 0$. Reinforcement learning, on the other hand, satisfies the second hypothesis. It displays a negative relationship before the breaking point and a flatter relationship thereafter, as predicted by the theory.

5.4 Weighted fictitious play, power choice rules and regret models

We now check whether the poor performance of the belief-based models (BR, FP) is due to the particular specifications used in the available electricity models. The preceeding work in this literature motivated three simplifications:

First, we have used extreme behavioral assumptions about how agents form beliefs about others' behavior. Players either disregard all periods except the last one (BR) or give equal weight to all of them (FP). We consider now an alternative weighted fictitious play rule (WFP). Under this assumption, agents consider all their previous experience, but their opponents' earlier behavior carries less weight than more recent behavior. Second, we have assumed a linear probability choice rule, a particular type of the power probability rule proposed by Camerer and Ho (1999). As an alternative, we also study their exponential probability rule. Third, we have not considered alternative behavioural rules based on recent additions to the literature. For example, Ert and Erev (2007) and Marchioria and Warglien (2008) include some regret-based feedback and argue that it might more accurately predict actual human behavior in games with mixed strategy equilibria.

In this section, we address the three alternative specifications. We use our baseline case (stepwise bidding with $m = 1$, and $\bar{Q} = 8$, $u = 0$ with uniform initial distribution) for the WFP model ($\delta = 1$ and $\phi = \rho \in (0, 1)$), both with linear and exponential probability rules, and with and without regret. We use a model with an exponential regret-based choice rule, similar to Ert and Erev's (2007) model, and use their estimated parameter combinations in the benchmark specification. Our implementation assumes that each action s for bin j in generator i is updated with

$$P_{i,s}^j(t+1) = \frac{e^{\lambda A_{i,s}^j(t)/\theta_i^j(t)}}{\sum_{k=1}^S e^{\lambda A_{i,k}^j(t)/\theta_i^j(t)}}, \quad (6)$$

where $\theta_i^j(t)$ is the regret parameter. $\theta_i^j(t)$ is updated according to

$$\theta_i^j(t) = 1 + \frac{\phi N(t-1) \theta_i^j(t-1) + \max_s \{\pi_i(s, r_i^{-j})\} - \pi_i(r_i^j, r_i^{-j})}{N(t)}, \quad (7)$$

where r_i^j and r_i^{-j} denote again the action taken in period t by bin j of firm i and by the rest of the bins, respectively, and $\theta_i^j(0) = 1$.

Figure 4 provides the simulation results. The first row reports simulations of the WFP model with a linear choice rule, as in (3). The WFP specifications are, from left to right, $\phi = 0.52$ (as

estimated in Ert and Erev, 2007), $\phi = 0.25$ and $\phi = 0.75$. The simulation outcomes are similar to each other and to those under BR ($\phi = 0$) (second row, first column in Figure 3). Prices are close to the competitive levels after a breaking point. However, they are not sensitive to the number of firms before the breaking point and far from the maximum (Ψ), even under monopoly. Prices are not as close to the competitive levels as in BR, especially when $\phi = 0.75$, but clearly closer than under FP ($\phi = 1$).

The second row provides the results of the WFP model with an exponential choice rule, as in (6), but without regret (that is, $\theta_i^j(t) = 1$ for all t and i and j). We assume $\phi = 0.52$ and, from left to right, $\lambda = 2.75$ (as estimated in Ert and Erev, 2007), $\lambda = 1$ and $\lambda = 5$. Prices are decreasing and close to Ψ under monopoly. Still, post-breaking point prices are high, particularly when $\lambda = 1$. Also, the confidence intervals for $\lambda = 2.75$ and $\lambda = 5$ are quite wide around the breaking point.

The third row provides the results of the WFP model with an exponential choice rule, as in (6), with regret ($\theta_i^j(t)$ updated with (7)). We use the same parameters as in the second row ($\phi = 0.52$ and, from left to right, $\lambda = 2.75$, $\lambda = 1$ and $\lambda = 5$). Confidence intervals are narrower than without regret but the breaking point is less clear, especially for $\lambda = 1$. Post-breaking point prices also remain well above competitive levels.

Overall, it seems that one can improve on the basic belief-based algorithms of the electricity modelling literature. As an illustration, some of alternatives in this section are clearly better than the previous best-response and fictitious play outcomes. Further tests and refinements are necessary but this evidence indicates that one way to improve the performance of simulations models is by using experiments to identify the appropriate simulation algorithms.

6 Discussion and conclusions

Firms and regulators alike have started to use behavioural simulations 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 simulation techniques and how they compare to each other and to a standard economic theory benchmark. As a case in point, we focus on a well-established claim in the wholesale electricity auctions 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 behavioural literature on electricity markets. Reinforcement learning performs quite well, but best-response and especially fictitious play do 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. Second, future models should incorporate more systematic robustness tests. Third, 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. Fourth, 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. We have also left out the network congestion issues inherent in the operations of electricity utilities. Finally, 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).

The question of which models fit best to real data is complementary to our research and deserves future attention. The alternative belief-based models proposed in the last section of the paper perform better than best-response and fictitious play. Empirical and especially experimental research can help in identifying the appropriate specifications to model electricity markets. Still, most researchers would agree in that there is no single universal set of assumptions that can be applied to all situations. This is for example one of the interpretations one can draw from Camerer and Ho's (1999) behavioural algorithm. We think that this is a valid option as long as we know the modelling choices' implications. The ACE community has made a lot of progress in recent years and we believe it is now time to take stock of what has been achieved, consolidate and move forward.

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] 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.
- [5] 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.
- [6] Bolle, F. (1992): "Supply function equilibria and the danger of tacit collusion The case of spot markets for electricity". *Energy Economics*, vol. 14, no. 2, pp. 94-102.

- [7] Bower, J. and Derek W. Bunn (2000): "Model-based comparisons of pool and bilateral markets for electricity". *Energy Journal*, vol. 21, no. 3.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] Bunn, D. W. and Maria Martocchia (2005): "Unilateral and collusive market power in the electricity pool of England and Wales", *Energy Economics*, Volume 27, Issue 2, Pages 305-315.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] Bunn, D. and Fernando Oliveira (2008): "Modelling the impact of market interventions on the strategic evolution of electricity markets". *Operations Research*.
- [18] Camerer, C. and Teck, H. Ho (1999): "Experience-weighted attraction learning in normal form games". *Econometrica*, vol. 67, pp. 837-874.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] 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.

- [23] Denrell, J. (2004). "Random walks and sustained competitive advantage." *Management Science*, vol. 50, issue 7, pp. 922-934.
- [24] 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*.
- [25] 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.
- [26] 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.
- [27] 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.
- [28] Fabra, N., Nils-Henrik von der Fehr and David Harbord (2006): "Designing electricity auctions". *RAND Journal of Economics*, Vol.37 (1), pp. 23-46.
- [29] 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.
- [30] 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.
- [31] 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>
- [32] Fudenberg, D. and D. Levine (1998): *The Theory of Learning in Games*, MIT Press.
- [33] 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.
- [34] Genc, T. and S. S. Reynolds (2005): "Supply function equilibria with pivotal electricity suppliers". Working paper, University of Guelph.
- [35] Hobbs, B. and J. S. Pang (2007): "Nash-Cournot equilibria in electric power markets with pieewise linear demand functions and joint constraints". *Operations Research*, vol. 55, no. 1, pp. 113-127.
- [36] Kimbrough, S. and Frederic H. Murphy (2009): "Strategic bidding of offer curves: an agent-based approach to exploring supply curve equilibria." Working paper.
- [37] Klemperer, P. and Margaret A. Meyer (1989): "Supply function equilibria in oligopoly under uncertainty", *Econometrica*, vol. 57, no. 6, pp. 1243-1277.
- [38] 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.
- [39] Leombruni, R., Richiardi, M., Saam, N. J., & Sonnessa, M. (2006): "A common protocol for agent-based social simulation." *JASS*, no. 9-issue 1.
- [40] 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.
- [41] Davide Marchiori and Massimo Warglien (2008): "Predicting human interactive learning by regret-driven neural networks," *Science* 319, 1111.
- [42] 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.

- [43] Marks, R.E. (2007): "Validating simulation models: A general framework and four applied examples". *Computational Economics*, 30(3), pp. 265-290.
- [44] 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.
- [45] 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.
- [46] 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.
- [47] 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.
- [48] 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.
- [49] 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.
- [50] Pouget, S. (2007): "Adaptive traders and the design of financial markets," *Journal of Finance*, vol. 62, issue 6, pp. 2835 – 2863.
- [51] 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.
- [52] 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.
- [53] 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.
- [54] 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.
- [55] 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.
- [56] Roth, A. E., Tayfun Sonmez, and M. Utku Unver (2004): "Kidney exchange," *Quarterly Journal of Economics*, vol. 119, issue 2, pp. 457-488.
- [57] Rothkopf, M. H. (2002): "Control of Market Power in Electricity Auctions," *The Electricity Journal*, vol. 15, issue 8, pp. 15-24.
- [58] 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.
- [59] 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.
- [60] Sun, J., Leigh Tesfatsion (2007): "Dynamic testing of wholesale power market designs: an open-source agent- based framework". *Computational Economics* 30 (3), 291–327.

- [61] Tesfatsion, L. and Kenneth Judd (2006): "*Handbook of computational economics, vol. 2: agent-based computational economics*". ISBN-13: 978-0-444-51253-6.
- [62] 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.
- [63] 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.
- [64] 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.
- [65] von der Fehr, N. and David Harbord (1993): "Spot market competition in the UK electricity industry". *Economic Journal*.
- [66] Weidlich, A. and Daniel Veit (2008): "A critical survey of agent-based wholesale electricity market models." *Energy Economics*, vol. 30, pp .1728–1759.
- [67] Wilson, R. (2002): "Architecture of power markets," *Econometrica*, vol. 70, issue 4, pp. 1299-1340.




Figure 1: 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 2: 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 3: The influence of behavioral assumptions on prices

Mean (+/- two standard deviations) of prices in RL (row 1), BR (2), and FP (3). Prior beliefs are max. price (column 1), random (2), min. price (3) and middle of the distribution (4)




Figure 4: Some alternative behavioral assumptions and their impact on prices

Mean (+/- two standard deviations) of prices in Weighted Fictitious Play (WFP - row 1, phi parameters 0.52; 0.25; 0.75), WFP with power choice rule (row 2, phi 0.52; lambda parameters 2.75; 1; 5), and WFP, power choice rule and regret (row 3, phi 0.52; lambda parameters 2.75; 1; 5).

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 Rup��rez 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 Garc��a 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 Shebl�� (1998)	IEEE-Trans	Wholesale market simulation	Single price bid per firm	Double-sided call auction	Genetic algorithm
18 Rup��rez Micola and Bunn	JEBO	Horizontal cross-holdings	Single price bid per firm	Double-sided call auction	Reinforcement learning
19 Rup��rez 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 Ilıc (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 modeling papers published as journal articles, with the year of publication and abbreviated journal title. In addition, the Table briefly summarizes the research issue in each paper together with their supply bidding, demand representation and behavioral algorithm assumptions. Full citations appear in the references list.

Demand				Estimates										
Behavior	Bid type	No. bins	u	Qbar	Theory br. 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	22.6915***	120.1925	-15.8361***	-68.3801	-3.7236***	-42.6075	3.3883***	38.14074	6770.1818
RL	Stepwise	1	0	85	7	20.8404***	149.9273	-8.9893***	-33.6750	-1.5780***	-38.0172	1.2370***	25.9813	3571.9040
RL	Stepwise	1	0	90	10	20.1647***	228.4622	-8.9473***	-16.5035	-1.4150***	-80.9688	1.1326***	21.1356	4325.2052
RL	Stepwise	1	5	80	7	22.1599***	175.1506	-17.0345***	-49.4377	-3.5936***	-94.2047	3.4454***	77.5304	6818.3850
RL	Stepwise	1	5	85	10	17.4470***	94.7108	-10.8443***	-9.5838	-1.6442***	-45.0715	1.5931***	14.2440	929.3660
RL	Stepwise	1	5	90	12	19.5195***	200.8121	-8.9316***	-3.8525	-1.2370***	-78.9680	1.0249***	5.0782	2651.2966
RL	Stepwise	1	10	80	10	15.9032***	74.4074	-11.8689***	-9.0407	-1.7786***	-42.0241	1.79439***	13.8278	750.4089
RL	Stepwise	1	10	85	12	16.1842***	80.9220	-10.4379***	-2.1882	-1.2457***	-38.6473	1.2756***	3.0716	566.0308
RL	Stepwise	1	10	90	12	16.4728***	129.1202	-7.0974***	-2.3325	-1.0077***	-48.6707	0.8540***	3.2238	961.1434

Supply				Estimates										
Behavior	Bid type	No. bins	u	Qbar	Theory br. 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	22.7623***	117.1292	-16.0214***	-67.2072	-3.7875***	-42.1032	3.46129***	37.8502	6387.31
RL	Stepwise	2	0	80	5	22.5093***	137.7616	-11.9539***	-59.6407	-3.5819***	-47.3575	3.2362***	42.0915	4698.30
RL	Stepwise	3	0	80	5	20.9970***	152.6804	-8.8664***	-52.5584	-2.5914***	-40.7079	2.3111***	35.7144	3995.64
RL	SF	1	0	80	5	23.8829***	29.5219	-11.8191***	-11.9099	-3.3377***	-8.9128	2.1521***	5.6534	698.92
RL	SF	2	0	80	5	21.6725***	21.9375	-7.5188***	-6.2043	-1.5723*	-3.4381	0.2282	0.4909	500.63
RL	SF	3	0	80	5	20.9316***	34.3007	-13.7422***	-18.3579	-1.1605	-4.1082	0.4488*	1.5632	1445.53

Behavioural choices				Estimates										
Behavior	Bid type	No. bins	u	Qbar	Theory br. 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	13.0663***	19.2451	-3.2282***	-3.8761	0.64344***	2.0473	-1.5568***	-4.8731	616.48
RL	Stepwise	1	0	80	5	23.2553***	127.6799	-16.4564***	-73.6547	-4.0425***	-47.9472	3.7094***	43.2801	7348.12
FP	Stepwise	1	0	80	5	13.0186***	55.7207	3.6323***	12.6736	0.5914***	5.4685	-1.5801***	-14.3719	1801.16

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 0.10 level; **significant at the 0.05 level; ***significant at the 0.01 level.