

A critical survey of agent-based wholesale electricity market models

Anke Weidlich*, Daniel Veit

*University of Mannheim, Chair of Business Administration and Information Systems,
Schloss, 68131 Mannheim, Germany*

1. Introduction

Because of the technical nature of the traded good, electricity markets rank among the most complex of all markets operated at present. Supply and demand have to be balanced in real time, considering transmission limits and unit commitment constraints. The electricity sector is characterized by multiple interlinked markets: fuel markets, markets for day-ahead scheduling and those for real-time dispatch or balancing energy, bilateral trading and auxiliary markets e.g. for emission allowances. Many energy firms are vertically integrated and act on several markets simultaneously, thus further complicating their trading strategies. Besides, and given the

* Corresponding author.

E-mail addresses: weidlich@uni-mannheim.de (A. Weidlich), veit@uni-mannheim.de (D. Veit).

oligopolistic structure of almost all electricity markets, participants have the potential to exert market power in many of these markets.

These complexities drive most classical modeling methods to their limits. Equilibrium models either do not consider strategic bidding behavior or assume that players have all relevant information about the other players' characteristics and behavior; they also disregard the consequences of learning effects from daily repeated interaction (Rothkopf, 1999). Game theoretical analysis is usually limited to stylized trading situations among few actors, and places rigid – oftentimes unrealistic – assumptions on the players' behavior. Human-subject experiments can be applied to electricity market research only with difficulties, because some expertise is necessary to realistically imitate the bidding behavior of a power generator. Thus, for many questions relevant in electricity market research, human-subject experiments are not an appropriate method.

Given the complexity of the electricity sector and also its high importance for a competitive economy, researchers and practitioners are increasingly willing to try new modeling methods in order to gain insights into various aspects of power markets. Agent-based (AB) modeling is one appealing new methodology that has the potential to overcome some shortcomings of traditional methods. Within the last ten years, more and more researchers have been developing electricity market models with adaptive software agents. This field of research is still growing and maturing. Some first attempts have already been revoked, others have gained popularity.

As the number of AB electricity models that have been published in journals starts to increase, and given the high attractiveness of agent-based approaches among researchers, this survey might help to get a sense of the state-of-the-art of the whole research field.¹ This literature review is supposed to guide newcomers or interested researchers through the intricate research field and points out the weaknesses and open issues that current approaches face. It is structured as follows: Section 2 gives a brief introduction to the methodology of Agent-Based Computational Economics (ACE); Section 3 presents the approaches and findings of relevant scientific papers in ACE electricity market research, and Section 4 summarizes the contributions made by the papers, points out some shortcomings of the current state-of-the-art, and suggests some lines of future work in the research field. Finally, Section 5 concludes.

2. Methodology of Agent-Based Computational Economics

2.1. Motivation for AB methods in economics

The electricity sector, and economies in general are characterized by difficult real-world aspects, such as asymmetric information, imperfect competition, strategic interaction, collective learning, and the possibility of multiple equilibria (Tesfatsion, 2006). Many of these factors can not – or only with difficulties – be accounted for with traditional economic modeling techniques. Analytical approaches usually have to put strong and constraining assumptions on the agents that make up the economic system under study, in order to set up elegant formal models.

When the concept of complexity came up, the focus in economic analysis shifted from rational behavior and equilibrium towards heterogeneity and adaptivity (a famous early example being the simulations of Axelrod, 1997). At the same time, the tremendous availability of computational resources made it possible to set up large-scale and detailed computational models that allow a

¹ We attempt to give a very broad overview of AB electricity market models and present the most relevant work in detail. The nature of a fast-growing and young research field entails the difficulty to account for all existing research; although great care has been taken to consider as many papers as possible, it cannot be guaranteed that the survey at hand is exhaustive.

high degree of design flexibility. AB models offered the possibility of not only describing relationships in complex systems, but growing them in an artificial environment (Epstein and Axtell, 1996). AB simulation is, thus, a third way between fully flexible linguistic models and more transparent and precise but highly simplified analytical modeling (Richiardi, 2004); the resulting models are dynamic and executable, so that their evolving behavior can be observed step by step (Holland and Miller 1991).

2.2. Procedure and main concepts

ACE researches the two-way feedback between regularities on the macro level and interaction of economic actors on the micro level. The actors are modeled as computational agents. The concept of (computational or soft-ware) “agents” stems from the fields of Distributed Artificial Intelligence (DAI) and Multi-Agent Systems (MAS). Common definitions of the term characterize them as autonomous, reactive, goal-oriented, or socially able, just to cite a few (for a discussion of the term, see e.g. Franklin and Graesser, 1997). However, as Drogoul et al. (2003) correctly annotate, these features do not all translate into computational properties in agent-based simulations. Most AB models do not require agents to exhibit all the characteristics of the software agents from the DAI or MAS world; instead, the most important features of agents in AB models is that they are goal-oriented and adaptive. All agents are assigned a value, like e.g. payoff, fitness, or utility, the amount of which is dependent on their actions in the environment they are placed in. Under goal-oriented, we understand that agents seek to maximize this value; adaptivity refers to the ability to learn which actions to take in order to increase this value over time, and so reach the goal.

In AB electricity simulations, the most common agents that make up the population are generators, load serving entities, and a market/system operator. Depending on the research questions, the simulation can also contain regulator agents, a transmission system representation, retail customers, or others. Agents can also be composed of other agents, thus permitting hierarchical constructions like utilities.

Another important aspect of agents in AB models is heterogeneity. AB modelers are not restricted to equally sized or symmetric firms, or to other constraints that arise from the limits of analytical modeling. Instead, every agent making up the modeled economy can be designed independently. The economy then evolves as a result of the interplay of these heterogeneous agents, i.e. from the bottom-up. The modeling procedure can be described as follows (Tsfatsion, 2002): After having (i) defined the research questions to resolve, the ACE modeler (ii) constructs an economy comprising an initial population of agents and subsequently (iii) specifies the initial state of the economy by defining the agents’ attributes (e.g. type characteristics, learning behavior, knowledge about itself and other agents) and the structural and institutional framework of the electricity market within which the agents operate; the modeler then (iv) lets the economy evolve over time without further intervention — all events that subsequently occur must arise from the historical time-line of agent-agent interactions, without extraneous coordination; this procedure is followed by (v) a careful analysis of simulation results and an evaluation of the regularities observed in the data.

2.3. Applications and open issues

Following Tsfatsion (2006), current ACE research can be divided into four strands: The empirical, or descriptive strand seeks to understand why and how global regularities result from the interplay of agents on the micro scale. Normative ACE research uses AB models as laboratories for economic design alternatives in order to test which policies, institutions, or

processes perform best in an environment of self-seeking agents. A third strand is theory generation, i.e. the structured analysis of dynamical behaviors of economic systems under alternative initial conditions, in order to find necessary conditions for global regularities to evolve. Finally, ACE researchers continually seek to improve the methodology itself and develop tools that facilitate setting up computational AB models. Most electricity related research is centered around the second strand, i.e. normative research with the aim of determining good market designs that leave little opportunity to exercise market power.

Given the great flexibility that AB modeling allows, and given its potential to represent complex economic systems, ACE can possibly become one important pillar of electricity sector research. Besides its advantages, however, some problems exist that researchers must be aware of and tackle in the future. Among the current weaknesses of the methodology, [Tesfatsion \(2006\)](#) states that it is not yet clear how ACE models will be able to scale up to provide empirically and practically useful models of large-scale systems with many thousands of agents. Besides, she annotates the difficulty of validating ACE model outcomes against empirical data. Especially the last point can clearly be observed when surveying the agent-based electricity modeling literature (see also discussion in Section 4, where some suggestions for further research into this problem are given).

3. Paper summaries

The modeling approaches and findings described in the most relevant studies are reviewed in the following. This review concentrates on AB simulation models for the analysis of market structures and market design in wholesale electricity trading. Aside from this, some researchers have applied agent-based methods for examining electricity consumer behavior at the retail level, e.g. [Hämäläinen et al. \(2000\)](#), [Roop and Fathelrahman \(2003\)](#), [Yu et al. \(2004\)](#), or [Müller et al. \(2007\)](#); others provide agent-based decision support tools for power market participants, e.g. [Praça et al. \(2004\)](#), [Bernal-Agustín et al. \(2007\)](#), or [Harp et al. \(2000\)](#). These will not be discussed here.

3.1. Simulations applying model-based adaptation algorithms

Some of the first AB simulation models of electricity systems define their own representation of how agents adapt to the system they are placed in. These learning representations are usually tailored for the specific design of the simulated market(s). They do not explicitly rely on findings from psychological research about learning or on developments from the DAI or MAS fields of agent learning. These – usually naïve or intuitive formulations – are termed model-based adaptation algorithms here. The most prominent work in this field has been conducted at the London Business School; other approaches, such as those by [Visudhiphan and Ilić](#), have also attracted interest by researchers.

3.1.1 Analyzing trading arrangements in England and Wales — [Bower, Bunn et al.](#)

[Bower and Bunn \(2000\)](#) present an AB simulation model of the England and Wales electricity market. The simulation is designed to compare different market mechanisms, i.e. daily versus hourly bidding and uniform versus discriminatory pricing.² Generator agents apply a simple

² The study is motivated by the proposition that the mandatory pool-based spot market, which existed since 1990, should be replaced by direct bilateral trading between generators and suppliers. This drastic change in the structure of wholesale electricity trading in England and Wales had been proposed by the Office of Electricity Regulation as part of the Revised Electricity Trading Arrangements in 1998, in an initiative to prevent the generating firms from exercising market power.

reinforcement learning algorithm which is driven by the goal to simultaneously maximize profits and reach a target utilization rate of the own power plant portfolio. The agents adjust their bidding strategies according to their last round's success; they either lower, raise, or repeat their last bid price, depending on whether their utilization and profit targets have been met in the last round, or not. The demand side of the market is modeled as a static aggregate load curve with limited price sensitivity. Transmission constraints or costs are neglected.

The results from different scenarios show that simulated market clearing prices are lowest in the case of daily bidding with uniform pricing, and highest in the case of hourly bidding with discriminatory *pay-as-bid* settlement. The authors explain this result by two phenomena: (i) in the pay-as-bid case, base load generators are forced to bid closer to the market clearing price in order to maximize profits (in contrast, they bid at prices close to zero in the uniform pricing case); this reduces competitive pressure on the mid-merit plants; (ii) hourly bidding allows generators to effectively segment demand into peak load and base load hours and, thus, to extract a greater portion of the consumer surplus than under daily bidding.

Another finding that the authors report from the simulation results is that for all simulated scenarios, bid prices fall as the target rate of utilization – which is a simple representation of the agents' hedging activity in the forward market – rises. At a target utilization rate of 100%, which can be assumed for a plant for which the output has already been contracted on the forward market, prices fall to marginal cost. This observation is consistent with theoretical considerations, i.e. that the optimal bidding strategy for a hedged power plant is to bid at short-run avoidable costs.

As generator agents in the model also learn across their portfolios and can transfer successful bidding strategies from one plant to all others, small agents with few power plants have an informational disadvantage over large firms who can submit more bids and, consequently, gather more market price information. In their simulation results, the authors find that agents with few plants perform better in a uniform price setting than they do in the pay-as-bid case. This can be explained by the fact that, in the uniform case, all agents receive the information of the system marginal price, i.e. the result of the industry's collective learning; thus the informational advantage that large generators have in the case of discriminatory pricing is mitigated.

In [Bower and Bunn \(2001\)](#) the computational analysis described above is complemented by a validation of the simulation model against classical models of monopoly, duopoly, and perfect competition. The mean simulated market clearing prices for pay-as-bid and uniform pricing are very close to the corresponding theoretical results in the monopoly and perfect competition models. In the case of duopoly, however, the difference between pay-as-bid average prices and uniform system marginal prices is much smaller in the simulation model than in the theoretical benchmark (£ 279.39 versus £ 262.10 in the simulation, £ 340.00 versus £ 257.50 in the theoretical model). The similar results for simulated and theoretical prices in the two extreme cases of monopoly and perfect competition gives confidence in the simulation model implementation. Nonetheless, it is not clear which conclusions the authors draw from the comparison of simulation and theory in duopoly and what their results imply for an oligopoly model (the reader would also have been interested in an evaluation of a simulated oligopoly case and whether prices in this case are closer to the competitive equilibrium or to the duopoly case).

[Bower et al. \(2001\)](#) apply the same basic model to the case of the German electricity sector. It is simulated as a day-ahead market in which plants are dispatched centrally and remunerated on a pay-as-bid basis. The target utilization rate is 60% for the major players and 100% for 14 small generator agents. The authors analyze the impact of four mergers of large German utilities that were probable at the time of the study (and have actually taken place shortly after). They find that electricity prices rise considerably as an effect of the mergers. Simulated prices are 16% higher in

on-peak and 5% in off-peak times when the two larger mergers take place. When all four mergers are realized, prices rise by 54% on-peak and 45% off-peak. When the merged firms also seek to rationalize their portfolios and shut down about 6% of the total system capacity, the effect on prices is even more severe: a 100% price rise in off-peak times and a 300% rise during winter peaks.

Bunn and Oliveira (2001) present a more detailed model of the *New Electricity Trading Arrangements* of England and Wales (NETA). In contrast to the approach described above, the authors explicitly model an active demand side and the interactions between two different markets, i.e. the bilateral market and the balancing mechanism. Trading in both markets is modeled as a call market with pay-as-bid settlement. Both generators and suppliers seek to maximize individual daily profits and simultaneously minimize the difference between their exposure to the balancing mechanism (BM) and their (fixed) objective for BM exposure. They learn to set mark-ups on their bid prices in both markets through reinforcement learning; the mark-up for the bilateral market is set relative to the price bid in the bilateral market price in the previous day, and the mark-up for the BM is set in relation to the bilateral market bid price of the same day. Consequently, all generator (supplier) agents learn three different policies for setting three different prices: the offer (bid) price for the bilateral market, and the bid prices of, respectively, *increments* and *decrements* for the balancing mechanism. In order to avoid inconsistent behavior during the learning process, the authors impose some “lower bounds of rationality” on the agents’ bidding strategies through the introduction of *operational rules*. Suppliers, for example, make sure that a more flexible power plant never undercuts the offer of a less flexible plant of their same portfolio.

The learning algorithm applied in the model is tailored for this specific trading arrangement. The range of actions that an agent can choose from, i.e. the possible mark-ups, differ for suppliers and generators (as an example suppliers can set mark-ups between 0.95 and 1.2 in the bilateral market, whereas generators choose from a range between -0.15 and 1.15 in the same market); the intervals are each partitioned into ten discrete mark-up values. At each trading day, agents calculate the *expected daily profit* and the *expected acceptance rate* of each possible mark-up using exponential smoothing of the previous days’ trading results. The *expected reward* for each mark-up is the product of the expected profit and acceptance rate. These expected rewards are ranked in descending order; the *perceived utility* of each mark-up j is calculated on the basis of its $rank(j)$ as follows:

$$Util_j = U \cdot \left(\frac{\text{Search Propensity} - n}{\text{Search Propensity}} \right)^{\text{Rank}(j)-1}.$$

The *search propensity* parameter expresses whether an agent has a rather conservative utility function (low value) or whether it is more adventurous in trying different mark-ups (high value); it is set to four for all agents in all simulations. U is equal to 1000 and n has a value of three. The probability of choosing mark-up j is defined on the basis of its perceived utility in the following way:

$$Pol_j = \frac{Util_j}{\sum_k Util_k}.$$

Unfortunately, the authors do not describe how the agents make sure not to violate the *operational rules* that have been defined, and to what extent the agents’ behavior results from the learned policy or from these operational rules. Moreover, one shortcoming of the reinforcement learning algorithm applied by the authors seems to be that the absolute values of $Util_j$

are the same in every round, but distributed on different mark-ups, as only the ranks are taken into account for its calculation. As a consequence, an agent has no information on how much better one mark-up is than the next best one. Also, as the rank of a mark-up enters exponentially into the determination of $Util_j$, the best mark-up has a very high probability of being chosen again (75% in the simulations presented in the paper), whereas for the last four mark-ups, this probability is close to zero.

The simulation model is run with power plant data that represents the UK wholesale electricity market. The authors observe very high prices in those 2 h of the day in which demand is highest.³ Another observation is that a wide spread between the *System Buy Price* and the *System Sell Price* emerges in the balancing mechanism and that the average bilateral price is centrally located between them. This corresponds to the intuitive system behavior.

An extension of the model and further analysis is presented by [Bunn and Oliveira \(2003\)](#). Here, the research question to be analyzed is whether two specific generation companies in the England and Wales electricity market are capable of manipulating market prices in order to increase their profits. The authors first analyze a simplified version of the bilateral market in the form of a one-shot pay-as-bid Bertrand game with capacity constraints. Based on the calculated results, they argue that this model does not allow evaluation individual market power abuse and that it neglects the impact of learning in repeated games, which might also be an important factor in market power analysis. Given these arguments, they apply the simulation model described above in order to find out whether the two generators can influence market prices to their own advantage. They compare six different withholding strategies, including the case of no withholding (benchmark), cases in which only one of the two generators withholds capacity, and one where both simultaneously withhold parts of their capacity. The reported results indicate that only one of the two generators, whose ability to influence market prices was studied, is capable of increasing electricity prices unilaterally. If both companies act together, they can also significantly increase power exchange prices. The second generator had no ability to manipulate prices alone. Moreover, prices in the balancing mechanism are found to be robust against manipulation from the two players.

One interesting remark that the authors state in their conclusion is that in these types of AB models, i.e. where agents learn to adapt their behavior to a stable environment, the potential for agents to collude on higher than marginal costs can be overestimated. In real markets where varying demand, fuel costs, and transmission constraints change the state of the world continuously, this coordination behavior might be harder to achieve.

An agent-based analysis of technological diversification and specialization is presented in [Bunn and Oliveira \(in press\)](#). The question that the authors want to answer in this paper is whether strategic generator agents in the electricity markets evolve into diversified players with a mix of base load, shoulder and peak load plants, or into specialized players that seek to dominate the market in their segment. They develop a model in which generators trade generating capacity among themselves and then, in a second stage, trade electricity from their plants, applying a Cournot strategy. Two mechanisms are compared: single-clearing and multi-clearing. The first mechanism corresponds to a power pool where one uniform price is set for every hour of the day; the latter is supposed to replicate trading in bilateral markets. In the multi-clearing setting, base load, shoulder, and peak load are traded separately in three different markets.

³ Note that electricity contracts are settled separately for every hour of the day in the model, but agents only learn one mark-up for the whole day and, in consequence, submit the same bids for all 24 h of the day.

Generators aim at maximizing the value of their whole portfolios. The plant trading game consists of three stages. In the *Initialization* stage, the Cournot game is solved, which gives the initial valuations of each plant. Then, in the *Identification* stage, the agents estimate which plants are most likely to be traded, in order to simplify the coordination problem. In the *Adaptation* stage, players decide which plants they will actually attempt to buy or sell; the resulting bids and offers are cleared in the *Trading* stage. Finally, the algorithm recalculates the capacities owned by each player and the respective cost structure in the *Updating* stage. The identification of the most likely trades is based on past observations. The adaptation procedure is based on stochastic search and learning. An *inertia* element of the adaptation algorithm defines whether the agent searches for new strategies or stays with the old ones; it decreases over the course of a simulation. New strategies are defined as the Best-Response to the given situation; the Best-Response action maximizes the sum of the utility (profit/reward) of an action and the discounted portfolio value.

The simulation model is run with plant data from the England and Wales electricity market. In the first set of simulations, base, shoulder and peak load plants are separated among three different players. The single-clearing mechanism leads to high concentrations in this setting: a single player (the base load player) becomes a monopolist, while the others sell all their capacity and are become extinct after less than 2000 iterations. In the multi-clearing case, several players coexist; the base load and shoulder players own most of the capacity at equal share, whereas the peak load player retains some capacity at the end of 2000 iterations. Prices are also higher in the single-clearing than in the multi-clearing case. In a second scenario, all three players have similar initial portfolios. In this case, the difference between the two clearing mechanisms is rather small. The authors conclude from this observation that if the industry is at a state of great diversification, it will tend to remain so, independently of the market-clearing mechanism.

3.1.2. Comparing different adaptation algorithms — Visudhiphan and Ilić

Another research initiative using AB models of wholesale electricity trading has been started by two researchers at the *Massachusetts Institute of Technology*. A first model implementation has been described in Visudhiphan and Ilić (1999). Here, three strategically interacting generator agents apply some form of *Derivative Follower* strategy (Greenwald et al., 1999) for learning to set profit-maximizing bid prices. In their extremely simplified model, the authors can show that in a market with price-inelastic load, generators extensively exercise market power, while a price-responsive demand side leads generators to bid more competitively, resulting in lower market prices.

In a later paper, Visudhiphan and Ilić (2001) report on simulation results from a model in which agents can strategically withhold capacity when their expected profit is higher than without withholding. Each agent records data about the market outcome in previous market rounds. The outcomes are each mapped to predefined discrete load ranges, so that each agent's memory can be represented as a matrix with rows corresponding to the different load ranges and columns corresponding to the market rounds. Agents also distinguish whether the resulting market price in one round has been a result of strategic or competitive behavior, and store this information likewise. Bid quantities and prices are defined separately in a two-step decision process. Each agent is assigned one out of six proposed strategies for setting the bid prices of anticipated marginal units: it can be set equal to (i) the maximum, (ii) the mean, or (iii) the minimum of historic prices, to (iv) the sum of weighted historic prices, (v) the last bid price plus the difference between the last market price and the last bid price, weighted by a constant β , or to (vi) a *target price* plus the absolute value of the difference between the last market price and this target price, weighted by a constant β ; the value of β depends on the success in the previous round. Simulation results are presented for two scenarios of available capacity. For each of the capacity scenarios, strategic and

competitive (i.e. marginal-cost) prices are compared. The authors come to the conclusion that generators are able to raise market prices if they bid strategically. This is observed not only for hours of high electricity demand, but also for low-demand hours. However, a distinction between the success of different price or quantity bidding strategies is not provided.

A more in-depth discussion of agent bidding representations is provided in [Visudhiphan \(2003\)](#). The thesis explores three different learning algorithms or bid selection strategies: (i) a modification of an algorithm formulated by [Auer et al. \(2002\)](#), (ii) a simple reinforcement learning algorithm using a Boltzmann distribution for defining the probabilities of choosing each action, and (iii) a model-based algorithm similar to the one presented in [Visudhiphan and Ilić \(2001\)](#). The Auer et al.'s algorithm assigns a probability $p_t(i)$ to each of the K possible actions. This probability is a mixture of a uniform distribution (γ/K) and a function of the weight factor $w_t(i)$ associated to each action i :

$$P_t(i) = (1 - \gamma) \frac{w_t(i)}{\sum_{j=1}^K w_t(j)} + \frac{\gamma}{K}.$$

The weight $w_t(i)$ is adjusted in every round, based on the rewards received from the chosen actions:

$$w_{t+1}(j) = w_t(t) \cdot \exp \left(\frac{\gamma}{3K} \left(\hat{x}_t(j) + \frac{\alpha}{p_t(j)\sqrt{KT}} \right) \right).$$

Here, $\hat{x}_t(j)$ is set to $x_t(j)/p_t(j)$ if $j=i_t$, and 0 otherwise, where $x_t(j)$ is the reward of the chosen action j in round t . Agents learn bid prices and bid quantities separately through applying this algorithm.

Similarly, agents learn prices and quantities separately in simulations applying a simple reinforcement learning algorithm where the probability of choosing an action j is defined as:

$$p_t(j) = \frac{e^{R_t(j)/\tau}}{\sum_{h=1}^K e^{R_t(h)/\tau}}.$$

The estimate value $R_t(j)$ of action j is updated in every round in the following way:

$$R_{t+1}(j) = \begin{cases} (1 - \alpha)R_t(j) + \alpha \cdot \Pi_t(j) & \text{if } j = i_t \\ R_t(j) & \text{otherwise} \end{cases}.$$

Various simulation runs with differing parameter combinations have been tested with these three algorithms, and results are compared. What is striking about the simulation results is that the daily load cycle seems to have a much stronger influence on resulting market prices than the learning representation. In all simulations, daily price cycles can clearly be distinguished, while for most of the learning algorithms prices do not exhibit any longer-term trends. The author does not provide any discussion about which learning algorithm is most appropriate for realistically modeling real-world behavior. She also concludes that her results cannot be validated against market results observed in any real-world market, because information on marginal-cost functions, bilateral contract obligations, operating constraints or other power system characteristics is not sufficiently available. Moreover, the thesis reports on comparative results from uniform versus pay-as-bid pricing; the conclusion of the result evaluation is that outcomes significantly depend on the learning algorithms that the agents employ.

With no quality measure for assessing the learning algorithms and with-out validation against empirical data, it becomes difficult to judge whether the model is appropriate for realistic electricity market modeling. It would have been desirable also to have a discussion of why the authors have obviously abandoned some of the concepts presented in their earlier papers. Among these is another publication by [Visudhiphan and Ilić \(2002\)](#), which describes an interesting approach to simulate distinct time scales of electricity trading (short-term bidding strategies, medium-term maintenance scheduling, and long-term new entry and shut-down or merger strategies). It has obviously not been implemented. As these authors have started to build an agent-based electricity market model very early, their experience – also with unsuccessful approaches – might be helpful for other researchers who are new in this field.

Two further publications by one coauthor describe results from another simulation model for analyzing market dynamics that arise from individual agent decision making. The strategic agents are two/three generators in a transmission system with and without congestion, respectively ([Ernst et al., 2004b](#)); other runs also include a profit-maximizing transmission line owner ([Ernst et al., 2004a](#)). Here, agents strategically set the bid price that will maximize their payoff under the assumption that the other agents repeat the same actions as in the precedent trading round. This approach seems to be a renunciation from the earlier approaches presented. Unfortunately, no discussion of the reasons for this turning away is provided.

3.2. Simulations applying genetic algorithms

Genetic algorithms (GA) are a class of heuristic search methods which are inspired by the biological process of evolution. In (electricity) market simulations, strategies that market participants can choose to apply are encoded into bitstrings which can be thought of as *chromosomes*. Most successful (or “fittest”) strategies are passed from one generation to the next by a mating process in which parent chromosomes produce *offsprings*. By mimicking *crossover* and *mutation*, GAs exploit the genetic dynamics underlying natural evolution to succeed in their environment.

Some AB models apply genetic algorithms for agents to search for optimal bidding strategies in electricity markets. [Curzon Price \(1997\)](#) has demonstrated the usefulness of GAs for simple standard games such as Bertrand and Cournot competition, price choice of a monopolist and a chain of monopolists, and also for very simplistic electricity market settings. In the following, some GA simulation models designed specifically for electricity market research are summarized.

3.2.1 Early GA approaches — [Richter, Petrov, Sheblé et al.](#)

[Richter and Sheblé \(1998\)](#), [Petrov and Sheblé \(2000\)](#), [Lane et al. \(2000\)](#), and [Nicolaisen et al. \(2000\)](#) present simple electricity market models that use genetic algorithms for representing the agents’ bidding behavior. The simulated market in the cited papers takes the form of a double auction where executable supply and demand bids are matched pairwise. Prices are determined as the midpoint between two matched bids or as a *competitive equilibrium* price.

The learning task for the agents differs in the cited papers. In [Richter and Sheblé \(1998\)](#) only generation companies are part of the GA population. They have three distinct evolving parts, or genes: one for determining the offer quantity, another for selecting the offer price and a third one for choosing a price forecast method. The latter includes strategies like e.g. moving average or linear regression. On the basis of the respective technique coded in their genes, agents determine their forecast price, then choose a bid price between their generating cost and the forecast equilibrium price, and determine an offer quantity between 0 and their maximum generation capacity. Standard GA methods are used for the evolution of the population.

Petrov and Sheblé (2000) propose a model where only one agent applies a genetic algorithm for evolving its trading strategy. All other agents display a very simple trading behavior: in a ten round auction, electricity sellers (buyers) increase (decrease) their bid price by one increment if the last offer was accepted; in the opposite case, they decrease (increase) their bid price for the next round. The GA agent develops more sophisticated trading strategies based on her last round's bid price and the equilibrium price. The strategy has the form of a decision tree whose functions comprise algebraic and logical operators (e.g. summation, division, greater, "if-then-else"). The authors find that the GA agent consistently surpasses the fixed rule agents during all separate runs of the simulation.

In Nicolaisen et al. (2000) and Lane et al. (2000) both buyer and seller strategies evolve with the help of a genetic algorithm. The GA is used to determine the agents' bid and ask prices. Possible bid (ask) prices are in the range of *marginal cost* and *marginal cost*+40 \$ (*marginal revenue*−40 \$ and *marginal revenue*); the bitstring representing the agent's gene codes a floating point number in the interval [0, 1] which is then multiplied by a dollar constant, resulting in the admissible bid (ask) price. The authors are interested in measuring the individual market power of buyers and sellers. When discussing their simulation results (most of which do not confirm their formulated hypotheses and even contradict theoretical economic considerations), the authors admit that their very simple GA implementation might not realistically represent human behavior in real markets. They propose to try other learning representations which allow agents to learn on the basis of their individual experience in the trading process.

The learning representations chosen in the aforementioned models are indeed very simplistic and do not deliver satisfactory results. For this reason, reinforcement learning is used for representing the agent behavior in AB electricity market models in later studies by the authors (Petrov and Sheblé, 2001; Nicolaisen et al., 2001; see Section 3.3.1).

3.2.2 Models of the Australian National Electricity Market — Cau et al.

Cau and Anderson (2002) develop a wholesale electricity market model similar to the Australian National Electricity Market. In Cau (2003), the model is developed further. It covers two bidding structures, i.e. stepwise and piece-wise linear bidding. In both cases, agents assign bid quantities to a number M of given price segments. In the piecewise linear case, bidding schedules are formed by linear interpolation between two bid points, whereas quantities are kept constant between two points in the stepwise case. The agents' task is to find the strategies that maximize their individual payoffs. A strategy is a set of bidding schedules for the possible environmental states. A state is defined by the previous spot market price, the past market demand and the forecast market demand. Demand is price-inelastic and can be classified as either *high* or *low*, while there is some uncertainty about the exact level. With two possible demand levels, the total number of states is $2 \times 2 \times$ the number of possible prices or price bands. Starting from randomly created bidding strategies, the agents evaluate their fitness (i.e. average payoff when the strategy has been played) and select the best strategies for further rounds. Standard evolutionary operations, such as crossover and mutation, are also effected on the population.

Simulations are run for a duopoly with two equally sized generators, one having lower marginal generation costs than the other. The authors observe that tacitly collusive strategies can be learned by the agents in this co-evolutionary environment, both in the stepwise and piecewise linear bidding structure. Cau (2003) further explores the effect of some market demand and market structure measures on the agents' ability to achieve tacit collusion. The author finds that on the demand side, high overall demand, high uncertainty and low price elasticity facilitates tacit collusion (measured as the joint profit ratio). On the supply side, situations in which tacit

collusion is easier to achieve are characterized by symmetry in cost and capacity, and small hedging contract quantity. Also, the influence of the number of competing generator agents on the success of collusive strategies is examined. As expected, the author finds that an increasing number of agents makes it more difficult for them to collude in a sustainable way. However, even in cases with many competing generators, tacit collusion can still occur.

3.3. Simulations applying Erev–Roth reinforcement learning

Based on psychological findings on human learning, Erev and Roth (1998) have developed a three parameter reinforcement learning algorithm. This learning model has gained much attention by AB modelers. Also, a considerable number of papers describing agent-based electricity models applying this learning algorithm can be found and are summarized in the following.

The learning formulation integrates the aspects of *experimentation* and *forgetting*:⁴

$$q_{nj}(t+1) = \begin{cases} (1 - \phi)q_{nj}(t) + R(x)(1 - \varepsilon) & \text{if } j = k \\ (1 - \phi)q_{nj}(t) + R(x)\frac{\varepsilon}{M-1} & \text{if } j \neq k \end{cases}$$

The choice of an action is probabilistic, and choice probabilities are derived from the propensities p_{nj} for action j in the next round in the following way (referred to as proportional action selection here):

$$p_{nj}(t+1) = \frac{q_{nj}(t+1)}{\sum_{k=1}^M q_{nk}(t+1)}.$$

Some researchers propose to determine the choice probabilities based on a Boltzmann distribution with a *Temperature* parameter T :⁵

$$p_{nj}(t+1) = \frac{e^{q_{nj}(t+1)/T}}{\sum_{k=1}^M e^{q_{nk}(t+1)/T}}.$$

Following Sutton and Barto (1998), this formulation will be referred to as *Softmax action selection* in this paper.

3.3.1 Developing the Erev and Roth algorithm — Nicolaisen, Petrov, Sheblé et al.

Petrov and Sheblé (2001) and Nicolaisen et al. (2001) present simple electricity market models applying Erev–Roth reinforcement learning. Both papers describe one problematic feature of the original algorithm formulation, i.e. that no propensity update occurs when profits are zero (or close to zero). Another flaw of the original algorithm formulation is accentuated by Koesrindartoto (2002): for some parameter combinations of ε and M , no learning occurs. These

⁴ Here, q_{nj} corresponds to the propensity of agent n to choose action j , $R(x)$ is the reinforcement from the last chosen action (k), M is the number of possible actions; f and ε denote the recency (or forgetting) experimentation parameters.

⁵ The *Temperature* parameter – also referred to as cooling parameter – determines the degree to which generator i focuses on actions with high propensity values. Usually, the temperature is decreased over the course of a simulation in order to allow more exploration at the beginning, while focusing on exploitation later on.

considerations lead to the formulation of the Modified Roth-Erev algorithm (MRE) (Nicolaisen et al., 2001):

$$q_{nj}(t+1) = \begin{cases} (1 - \phi)q_{nj}(t) + R(x)(1 - \varepsilon) & \text{if } j = k \\ (1 - \phi)q_{nj}(t) + q_{nj}(t)\frac{\varepsilon}{M-1} & \text{if } j \neq k \end{cases}$$

The MRE is used to simulate a double auction electricity market with discriminatory pricing. The aim of the study is to analyze market power and efficiency as a function of relative concentration and capacity of the market.

The simulation results do not support the hypotheses that the market power of sellers increases when relative capacity increases or when relative concentration decreases. However, the hypothesis that market efficiency (total auction profits in relation to total profits in competitive equilibrium) is high receives support by the simulation results. The authors compare their results to the market efficiency observed in simulations using genetic algorithms (Nicolaisen et al., 2000, see also Section 3.2.1) and come to the conclusion that individual MRE learning leads to higher market efficiency than GA social mimicry learning, because each agent learns to make his own bidding strategies on the basis of his own profits instead of mimicking strategies of structurally distinct traders.

3.3.2 Wholesale market reliability testing with AMES — Tesfatsion et al.

Based on the work at Iowa State University, Koesrindartoto and Tesfatsion (2004), Koesrindartoto et al. (2005), and Sun and Tesfatsion (2007) describe an electricity market model that encompasses the core features of the *Wholesale Power Market Platform*, a market design that has been proposed by the U.S. *Federal Energy Regulatory Commission* (FERC). The current implementation of their *AMES* model (*Agent-based Modeling of Electricity Systems*), as described in Sun and Tesfatsion (2007), comprises a two-settlement system consisting of a day-ahead market and a real-time market which are both cleared by means of locational marginal pricing. The market clearing is managed by an independent system operator (ISO) agent who operates an AC transmission grid.⁶ The authors report on initial simulation results that are run with a 5-node transmission grid test case and with only 1 day-ahead market (the real-time market is inactive). The demand side is simplified to a fixed and price insensitive daily load profile submitted to the ISO.

The generator agents learn to optimize a supply function. While reporting their true production limits (minimum and maximum capacity) to the ISO, they strategically set the prices (reported marginal costs) at these capacity levels;⁷ thus they have the ability to submit supply functions that are above their true marginal generating costs. The reinforcement learning algorithm applied in the described simulations is the MRE algorithm with Softmax action selection.

The simulation results show that all five generator agents learn to successfully submit bids above their true marginal cost. This leads to total variable costs of operation that are about three times higher than they are in the case in which generators report their true marginal costs. The authors conclude that the *Wholesale Power Market Platform* design features do not prevent the considerable exercise of market power by generators. Further extensions of the AMES model are

⁶ The representation of the transmission grid in this model is approximated by a bid-based DC optimal power flow (OPF) problem which is described in more detail in Sun and Tesfatsion (2006).

⁷ The resulting supply function is determined through linear interpolation between the prices at the minimum and maximum capacity.

envisaged. To encourage these extensions, the developers of the AMES test bed have released it as free open-source (Java) software.⁸

3.3.3 Vertical integration in the energy sector — Rupérez Micola et al.

Rupérez Micola et al. (2006) present a model that consists of three sequential oligopolistic energy markets representing a wholesale gas market, a wholesale electricity market and a retail electricity market. They analyze the effect of reward interdependence in vertically integrated energy firms. Trading in all three markets is modeled as a uniform-price auction with fixed inelastic demand to which seller agents submit their bids. Firms in each of the three tiers are identical and have constant marginal costs, normalized to 0. Agents always bid their full capacity and only choose bid prices; possible actions (prices) range from 0 to an upper price ψ in the retail market, which is cleared first. The possible bid price range in the wholesale electricity market is set from 0 to the resulting retail market price; in the gas market, which is cleared last, it goes from 0 to the resulting wholesale electricity market.

Vertically integrated firms are modeled as two agents that each trade in one market, and whose rewards are interdependent. The two agents can be thought of as two strategic business units within the same firm. In the case of vertical integration, the reinforcement that an agent perceives from trading results in one market is only partly based on the profit earned in this market; the other part consists of a fraction of the profits earned in the other market. The size of this fraction is expressed through the *reward interdependence parameter* $\alpha = \{0, 0.01, 0.02, \dots, 0.5\}$. Agents learn to set their bids with the help of a slightly modified Erev–Roth learning algorithm:

$$q_{nj}^i(t+1) = \begin{cases} (1 - \phi)q_{nj}^i(t) + R^i(x) & \text{if } j = k \\ (1 - \phi)q_{nj}^i(t) + (1 - \delta)R^i(x) & \text{if } j = k \neq 1 \\ (1 - \phi)q_{nj}^i(t) & \text{if } j \neq k \text{ and } j \neq k \pm 1 \end{cases} .$$

The index i , $i \in \{r, e, g\}$ indicates that propensities are built separately for each of the three markets. The parameter δ , $0 < \delta < 1$ determines the degree of local experimentation (through reinforcement of similar strategies). Proportional action selection is applied. The authors also introduce the aspect of “extinction in finite time”, which means that actions are removed from the action domain when their probability of being chosen falls below a fixed value μ .

Simulation runs are conducted with two gas shippers, three wholesale electricity traders, and four retail electricity traders. The reward interdependence parameter is varied from $\alpha = 0$ (no reward interdependence) to $\alpha = 0.5$ (strong reward interdependence) in 51 discrete steps. Results show that the presence of vertically integrated firms generally raises prices in at least two of the three markets considered. In the case of reward interdependence between a gas shipper and a wholesale generator, prices in the gas market (wholesale electricity market) rise from 59 (83) to 63 (93) units when α is increased from 0 to 0.5. The authors show that the vertically integrated firm can increase its overall profits. What the authors do not mention is that the other, not integrated firms in the gas and wholesale electricity market taken together can also increase their profits, to an extent even slightly higher than the integrated firm. So, all firms profit from the higher prices in the scenarios with high reward interdependence.

The question then is why the firms can achieve higher prices in the case in which one firm is vertically integrated. Here, the authors presume that agents coordinate overall profits in both markets and conclude that “vertically integrated firms give up profits downstream [in the

⁸ Available at <http://www.econ.iastate.edu/tesfatsi/AMESMarketHome.htm>.

wholesale electricity market] in order to increase the scope for upstream profits [in the gas market]”. However, as the agents learning algorithm in this model is not designed in a way so as to allow agents to develop strategies across several markets, it is not convincing that agents really coordinate overall profits. A closer look on the results reveals that profits are generally higher in the gas market than in the electricity market. As a consequence, vertically integrated gas shipper agents, whose reinforcements only contain a part of the profits earned in the gas market (the other part coming from the electricity market), are inevitably less satisfied with their trading results than their (otherwise identical) rivals. Thus, they search for other ways of attaining higher profits, which can lead to higher prices in the gas market. Inversely, the vertically integrated agent “feels” better off than his rivals in the electricity market. However, when prices in the gas market rise, margins in the wholesale electricity market shrink, so agents may seek to compensate this loss through higher bid prices. It seems more appropriate to explain observed market prices on a level that can be deduced directly from the agents’ learning tasks than interpreting them as higher level strategies.⁹

3.3.4. Further approaches applying Erev–Roth reinforcement learning

Bin et al. (2004) report on simulation results comparing three different pricing methods in electricity auctions. Agents in their model submit bids for their whole installed capacity; they learn to bid mark-ups on top of their marginal costs. The learning algorithm applied is similar to the Erev–Roth reinforcement learning algorithm with proportional action selection. The authors compare uniform pricing, pay-as-bid pricing and a mechanism called *Electricity Value Equivalent* (EVE) pricing. Simulations are carried out for two cases, i.e. one in which each generator agent owns only one power plant and the second where the same capacity belongs to less generators, each owning more than one power plant.¹⁰ Both the model and the result presentation are rather brief and in parts confusing; the conclusion of the simulation results is that EVE pricing leaves less room for generators to exercise market power than the other two considered pricing methods.

Cincotti et al. (2005) model a day-ahead electricity market which takes the form of a clearing-house double auction with uniform pricing. While the authors have used the original formulation of the Erev–Roth reinforcement learning algorithm in earlier work (Cincotti and Guerci, 2005), they now propose a new algorithm, which they say is inspired by Erev’s and Roth’s original work. The learning formulation they propose not only evaluates the profits gained from chosen actions, but also evaluates potential profits that would have been realized had other actions been chosen.¹¹ Propensities $f_s(t)$ for those strategies s that would have yielded a higher than the actual profit at time t are updated according to the simple rule

$$f_i(t + 1) = (1 - r) \cdot f_s(t) + G_s(t)$$

where r is a recency parameter and $G_s(t)$ is the potential profit. Propensities of strategies whose potential profits are lower than the realized profit are set to 0 and will, thus, not be considered in further rounds. Agents in this model learn to bid price quantity pairs that maximize their profits; bid prices can range from the generator’s marginal cost to the maximum admissible price, and bid

⁹ Following the same argument, it is not surprising that the authors do not find evidence for “raising rivals’ cost” behavior. Agents do not perceive the other agents’ costs and have no reasoning capability that allows them to devise such strategies.

¹⁰ It is not specified in the paper how agents strategically bid a portfolio of power plants in the second case.

¹¹ Potential profits are calculated under the (questionable) assumption that an agent would have sold the total bid volume if his bid price had been less than or equal to last round’s market clearing price.

quantities range from 0 to the generator's maximum installed capacity, both in steps of 1 EUR, or 1 MW respectively.¹² The authors conduct experiments for different cases of supply side competition; the number of generator agents is varied from 10 to 100, and the supply/demand ratio from 1.25 to 2. All generators have the same installed capacity and marginal costs. Among the findings is that prices quickly converge to the competitive equilibrium value (i.e. marginal cost) in most cases. The authors find that the equilibrium outcome is reached faster when the number of competing generators is small. This finding is rather unintuitive; it may be an indicator for the model not being a realistic representation of the market under study.

In [Cincotti et al. \(2006\)](#) a similar model is run with a learning algorithm as formulated by [Marimon and McGrattan \(1995\)](#). This algorithm assigns a strength $S_{i,t}$ to every action that an agent i can take, which is updated in every round as follows:

$$S_{i,t}(a_i) = \begin{cases} S_{i,t-1}(a_i) - \frac{1}{\eta_{i,t-1}(a_i)} \cdot [S_{i,t-1}(a_i) - \Pi_{i,t-1}(a_i)] & \text{if } i \text{ plays } a_i \\ S_{i,t-1}(a_i) & \text{otherwise} \end{cases}$$

$\eta_{i,t}(\alpha_i)$ is the number of times that strategy α_i has been played. The algorithm comprises some element of inertia, which means that agents can keep their mixed strategies constant over some time:

$$\bar{\sigma}_{i,t}(a_i) = \begin{cases} \sigma_{i,t-1}(a_i) \cdot \frac{\exp(S_{i,t-1}(a_i))}{\sum \sigma_{i,t-1}(a_i) \exp(S_{i,t-1}(a_i))} & \text{with probability } \rho_{i,t} \\ \sigma_{i,t-1}(a_i) & \text{with probability } 1 - \rho_{i,t} \end{cases}$$

The probability to choose a certain action in the next round is finally given by the following formula, which ensures that every action has at least a minimum probability $\varepsilon_{i,t} \in (0, 1)$ to be chosen:

$$\sigma_{i,t}(a_i) = \begin{cases} \varepsilon_{i,t} & \text{if } \bar{\sigma}_{i,t}(a_i) \leq \varepsilon_{i,t} \\ \frac{\bar{\sigma}_{i,t}(a_i)}{\sum \bar{\sigma}_{i,t}(a_i)} \left(1 - \varepsilon_{i,t} \cdot |\{\bar{\sigma}_{i,t}(a_i) \leq \varepsilon_{i,t}\}|\right) & \text{otherwise.} \end{cases}$$

The authors conduct experiments for a duopoly case in which both generators have the same installed capacity but different (constant) marginal costs. The domain of possible action contains only bid prices in one case (agents offer their maximum capacity), and both bid quantities and prices in second case; for both options, the two levels of *low demand* (one generator can satisfy demand alone) and *high demand* (both generators are called into operation) are distinguished. Uniform price and pay-as-bid mechanisms are compared on the basis of several simulation runs. The authors find that if agents compete only on the basis of prices, market efficiency, i.e. long-run profits gained by the two sellers, does not depend on the auction mechanism. However in the

¹² A few critical notes are in order about this learning representation: The proposed formulation has some similarities with the concept of experience-weighted attraction (EWA) ([Camerer and Ho, 1999](#)), in that it also partly bases its strategy evaluation on hypothetical payoffs. This raises the question why the authors have not considered to employ the EWA algorithm, which is better established and has been evaluated for different kinds of learning tasks, e.g. [Arifovic and Ledyard \(2004\)](#). Moreover, the early distinction of strategies once they do not perform better than the currently played strategy leads to a rapid and unjustified constriction of the action space and may hinder agents from finding the best strategies. As the bid price is set to 36 EUR for both agents in the first round, prices never rise above this level, as all higher bid price actions are eliminated after the first round. Finally, it should be noted that the proposed learning algorithm formulation does not have much left in common with the reinforcement learning algorithm formulated by Erev and Roth.

price-quantity bidding case, results for uniform and pay-as-bid pricing differ: for low demand, the uniform price auction offers higher profits for both generators. In the pay-as-bid auction, more competitive strategies are played more often, and profits are especially lower for the less efficient generator. In the high demand case, the uniform price auction also yields higher profits for the more efficient generator, but for the less efficient generator, profits are almost equal in pay-as-bid and uniform price auctions.

Weidlich and Veit (2006) simulate two markets that are cleared sequentially – a day-ahead electricity market and a market for balancing power. Agents place bids in both markets and evaluate their individual success in one market by integrating the opportunity cost of profits that could have been obtained in the other market. They learn from trading results using modified Erev and Roth reinforcement learning algorithm with proportional action selection. The day-ahead market is modeled as a call market with a fixed and price-insensitive demand side. The balancing power market replicates the rules in place in Germany. Both pay-as-bid and uniform price settlement have been simulated. The authors show that the order of market execution plays a significant role in resulting market prices: if the day-ahead market is cleared first, agents still have more available capacity to offer in the auction. Hence, competition is increased and prices are lower as compared to the case in which some generators have already sold (parts of) their capacity in the balancing power market. As potential profits that could have been earned on the day-ahead market play a more important role for setting bid prices in the balancing power market than vice versa, low prices in the day-ahead market also lead to lower prices in the balancing power market. When the balancing power market is cleared first, prices are higher in both markets. The authors also observe that agents tend to bid lower prices in the case of uniform pricing; however, resulting (average) prices are lower in the pay-as-bid case.

Veit, Weidlich, Yao, and Oren (2006) study the dynamics in two-settlement electricity markets. In these markets, energy producers sign strategic contracts in the forward market, and engage in oligopolistic competition in the spot market. While transmission constraints are ignored in the forward market (however, electricity is traded in different forward trading zones of the network), the spot market considers an underlying transmission network in form of a lossless DC power flow optimization problem. In the spot market, electricity is paid at nodal prices; forward contracts are settled at spot zonal settlement prices.

Agents in this model learn to set profit maximizing bids on the spot and forward market separately. They apply a modified Erev and Roth reinforcement learning algorithm with proportional action selection. Propensity update for spot bids is effected on the basis of achieved spot market profits. Bids are evaluated for each power plant and for each possible spot market state individually. The propensities for possible bids on the forward market are updated on the basis of the generators' total profit; these bids are evaluated globally, i.e. learning is not differentiated for single plants. The agents' learning task is to set profit maximizing bid quantities on both markets (Cournot game). The load is modeled as a linear demand function, so some price-responsiveness is assumed. Simulation results from this model demonstrate that the introduction of a forward market influences the supply agents' bidding strategies in the spot market. Forward trading leads to a more competitive behavior of the suppliers in the spot market, and thus to lower spot electricity prices.

3.4. Simulations applying *Q*-Learning

Krause et al. (2005) compare Nash equilibrium analysis and AB modeling with *Q*-learning for the case of a power pool with transmission constraints (the network representation is depicted in Fig. 1).

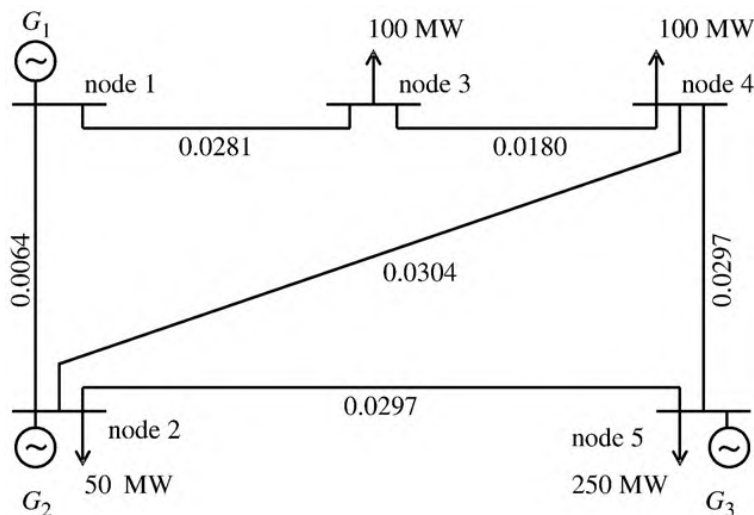


Fig. 1. Transmission system in the model described by Krause et al. (2005).

The model contains three agents who strategically set mark-ups to their electricity selling bids; the set of possible mark-ups is $\{0, 10, 20, 30\}$ \$/MW for the two agents whose cost function intercept is 10 \$/MW and $\{0,10,20\}$ \$/MW for the third agent, whose cost intercept is 20 \$/MW. The agents apply an update rule similar to the one used in Q -learning for learning the value of each possible action:

$$Q(a_t) \leftarrow Q(a_t) + \alpha(r_{t+1} - Q(a_t)).$$

However, unlike in the original Q -learning formulation by Watkins (1989), there is no differentiation between the states that agents are in at each iteration. So, one important element of Q -learning, e.g. taking into account future rewards of possible actions, is not included in this formulation. Despite this modification, the authors achieve satisfactory simulation results. They apply an ϵ -greedy strategy for action selection¹³ and compute Nash equilibria for simplified model settings. In those cases in which there exists one Nash equilibrium, the agents' actions quickly converge to this equilibrium, whereas in the case of two Nash equilibria, cyclic behavior is observed.

In a later study, Krause and Andersson (2006) apply the same model for analyzing the social welfare implications of the following three different congestion management methods:¹⁴

- Locational Marginal Pricing (LMP)
- Market Splitting
- Flow-based Market Coupling (FBMC).

The case in which all generators bid their true cost functions (referred to as *perfect competition*) is compared to a case in which all three generators apply the above described Q -learning-like algorithm (the *oligopolistic competition* case). Simulation results reveal that LMP results in the highest overall welfare in both the competitive and the oligopolistic case. Market splitting performs second best in both cases and FBMC shows lowest overall welfare. However,

¹³ For a description of the ϵ -greedy strategy, the reader is referred to Sutton and Barto (1998).

¹⁴ The congestion management methods are not described here. For references to these concepts, the interested reader is referred to the cited paper.

in the case with learning agents, generators' surplus is higher while consumer surplus is lower than in the competitive case. This is due to the exercise of market power by the generators.

Similar results for a different model specification are reported by [Naghbi-Sistani et al. \(2006\)](#). In their model, two different states are defined: (i) low production cost and (ii) high production cost. Agents strategically set the slope of their linear bid supply function; the possible bidding strategies are classified as high, middle and low (price/slope). In their simulations, only two bidders compete to satisfy demand in a power pool without transmission constraints. The authors use a Softmax action selection rule. The temperature parameter T is adjusted during the simulation in the following manner:

$$T = 20(1 - p(a_i|s))|\Delta Q(s, a_i)$$

where $p(\alpha_i|s)$ is generator I 's probability of choosing action α in state s , and $|\Delta Q(s, \alpha_i)|$ is the absolute change in Q -value from one iteration to the next. According to this definition, T is high when the Q -values have not converged to a final value, and low when changes in Q become small. The motivation behind this formulation is to reduce convergence time and sensitivity to learning parameters. However, as the choice of this particular formulation is not argued in detail, and is not compared to other possible formulations, it seems rather arbitrary.¹⁵ Reported results from an illustrative simulation runs show that agents quickly learn to play the strategy corresponding to a Nash equilibrium; the probability of this strategy converges to 1 after ~ 100 iterations.

[Xiong et al. \(2004\)](#) compare market prices and price volatility in uniform price and pay-as-bid electricity auctions with the help of a model with Q -learning generator agents. The environmental states are defined as last round's market prices; the set of possible states is $\{0,1,2,\dots,20\}$. Agents learn to set bid prices that maximize their payoffs; possible actions are also in the range of $\{0,1,2,\dots,20\}$ and bid volumes are always equal to the net capacity of a generator. The reinforcement for each hour h of the day comprises the profit r of that hour, and also the ratio of actual and target utilization rate (p is a parameter defining how strictly an agent tries to satisfy the target utilization rate):

$$r_h(s, a) = r(h) \times \left(\frac{\text{actual utilization}(h)}{\text{target utilization}(h)} \right)^p.$$

Action selection is effected with the ϵ -greedy strategy. Simulation runs are conducted with ten generator agents and with both price-inelastic demand and a responsive demand-side; both demand scenarios are run with uniform price and pay-as-bid clearing. The authors compare market prices¹⁶ and bid prices of one agent for the four simulated scenarios. The conclusion from their simulation results is that agents bid at higher prices in the pay-as-bid case, but overall prices are higher under uniform pricing. The introduction of interruptible loads causes prices to drop for both clearing mechanisms; however, the price decrease is stronger with uniform pricing.

¹⁵ The authors' proposition is even counterproductive: while aiming at reducing the influence of one parameter of the learning algorithm, they introduce a new one which is simply set to 20. This parameter does not have an interpretation related to psychological findings about learning, nor does it have a sensible machine learning interpretation. Instead, it abandons the distinction between action-value estimates on the one hand and choice probabilities on the other hand. In original Q -learning with Softmax action selection, the temperature parameter defines the preference of exploration over exploitation of successful strategies. With T values set individually for each action in each state, this principle is abandoned.

¹⁶ It is not specified how the "market price" is calculated in the pay-as-bid case, in which each generator faces different prices. It can be assumed that the authors take the weighted average over all successful bids as the market price.

Bakirtzis and Tellidou (2006) also compare uniform and pay-as-bid pricing with agents applying Q -learning. As in the previous paper, environmental states are defined by the last round's market price (uniform price or quantity weighted average price of winning bids, respectively). Agents also bid their full capacity at the price learned through the Q -learning algorithm; possible bid prices are bounded by the price cap and the generator's marginal generating cost. The action selection rule applied in this paper is adapted from the *Simulated Annealing Q-learning* algorithm developed by Guo et al. (2004). All simulated scenarios comprise 5-7 generators; load is constant and price-inelastic. Four cases are examined: (i) uniform pricing with one dominant generator whose installed capacity is enough to satisfy total demand alone and whose marginal generating costs are lowest; (ii) same as (i) but pay-as-bid pricing; (iii) uniform pricing with the dominant agent from the first two cases partitioned into three equally sized generators, whose marginal costs are lowest; (iv) same as (iii) but pay-as-bid pricing. Results show that the dominant firm in (i) and (ii) learns to set the bid price that fully maximizes its profits. The two next cheapest suppliers learn to stay below the bid price of the dominant firm, in order to be dispatched. In the pay-as-bid case, the bid prices of these two generators are closer to the dominant generator's bid than under uniform pricing, but resulting market prices are slightly higher in the latter case. In the more competitive cases (iii) and (iv), the three generators with the lowest marginal generating costs compete to serve demand, and prices mostly stay below the marginal cost of the next cheapest generator. Here, however, average prices under pay-as-bid are higher than market clearing prices under uniform pricing.

3.5. Simulations applying Learning Classifier Systems

Learning Classifier Systems (LCS) combine reinforcement learning (for increasing the probability of choosing successful actions) and genetic algorithms (for producing new rules from successful ones). A classifier is a rule of the form "if <condition> then <action>", where the condition describes the state of the environment. An LCS is thus suitable for agents to solve cognitive tasks.

Bagnall and Smith (2005) present a simplified AB simulation model of the pre-NETA UK electricity market that applies LCS for developing successful bidding strategies. The model and simulation results have also been described in several other papers, e.g. Bagnall and Smith (1999), Bagnall and Smith (2000), Bagnall (2000a), Bagnall (2000b), and Bagnall (2004). The agent population in the presented model consists of 21 generator agents, each owning one generating unit. The units are classified as one of the four types: nuclear, coal, gas, or oil/gas turbine. Every generation type has different characteristics of fixed cost, start-up cost, and generation cost. The market environment is characterized by the forecast demand for each half-hour time slot (representing typical days of summer/winter and weekday/weekend demand), by capacity premia for each half-hour time slot of the following day, and by transmission constraints. Agents are grouped into three different constraint groups: unconstrained, constrained on, or constrained off. During one daily trading cycle, the system marginal price and pool purchase price¹⁷ are calculated from all bids submitted by the generators (forming the unconstrained schedule). The constraint levels and the unconstrained schedule then produce the constrained schedule. Prices and payments are communicated to the agents at the end of each round.

¹⁷ The pool purchase price is the sum of the system marginal price and the capacity premium.

The agents follow two related objectives, quantified by two different reward functions: avoid losses and maximize profits. They apply two learning classifier systems (LCS) in order to learn rules for achieving these objectives. The learning task is to set the bid price, all other bid parameters being constant for each generation type. Each LCS has three components:

- the performance component produces a prediction array consisting of estimates of the expected reward for following specific actions, for any given input and rule set;
- the reinforcement component alters the parameters of the current rule set based on the environmental feedback;
- the rule discovery component generates new, potentially superior, rules from old ones using a genetic algorithm.

The three main research questions that the authors seek to analyze are whether the agents learn to behave in ways observable in the real world, how changes to the market mechanism alter agent behavior, and whether the agents can learn to cooperate. As regards the first question of interest, the authors conclude that the agents' behavior is broadly consistent with real-world strategies. This conclusion is grounded on the observation that nuclear units generally bid at low prices, regardless of the demand level, whereas oil/gas turbine units bid high in order to capture peak generation; coal and gas units bid close to the level required for profitability and increase their bids in times of high demand. An additional series of experiments has been run in order to answer the second question: the payment calculation has been changed from uniform pricing to a pay-as-bid scheme. The results show that agents tend to bid higher under pay-as-bid, but the overall cost of meeting demand is still higher under uniform pricing. The third research question was concerned with the evolution of cooperation. Cooperation has been defined here as situations in which two or more players make the same high bid. Although the agents are able to produce high market prices in some environments, the authors do not find many rounds in which the cooperation criterion is met. They conclude that the number of available actions, the exploitation/experimentation policy of the LCS and the potential incorrect generalization over environments makes it difficult for agents to maintain cooperative strategies.

Bagnall and Smith (2005) compare their approach to the models of Bunn and Oliveira (2001) and Bower and Bunn (2001). They emphasize that the complexity of the agent architecture and the information used for learning is the main difference of their model in comparison to related approaches. Unfortunately, they do not provide any arguments why this high complexity of agent behavior is actually needed and in what way their results might be more valuable than those generated from other simulation models. Neither do they examine any research questions that cannot be answered with the other referred approaches. As long as a simple model of agent behavior can realistically represent the real-world features of interest, there is no apparent reason to apply more complicated approaches. The authors of the presented papers have not yet convincingly made clear what failures of simpler learning representations they avoid with the very complex LCS applied in their model.

3.6. Simulations with supply function optimizing agents

While analytical evaluation of supply function equilibria in power markets either assumes continuous supply functions or restricts the analysis of various industry ownership structures to symmetric equally sized firms (or both), Day and Bunn (2001) present a method for

determining imperfectly competitive outcomes in electricity markets based on computational modeling. Their proposed model contains generation companies who bid individual piecewise linear supply functions into a market with uniform price clearing. The agents seek to optimize the value of an objective function, i.e. their daily profits from both spot sales and long-term financial contracts. The optimization routine that the agents apply works under the conjecture that the other agents submit the same bid as in the previous trading round. Agents in this model have a limited optimizing behavior in that they only change the bid price of one power plant per iteration (they choose the plant which increases the value of the objective function most when bid at another price). Electricity demand is represented by an aggregate demand function with a defined demand elasticity; simulations are run with different demand elasticities in order to determine the influence of demand response on the generators' ability to exercise market power.

The authors evaluate their computational model by comparing it with the equilibrium in continuous supply functions that can be obtained through the approach formulated by [Klemperer and Meyer \(1989\)](#). They find that results from the two approaches are reassuringly close for a simplified market scenario which models competition between three symmetric generating companies who have linear marginal costs. Based on this finding, the authors are confident that the computational approach can also deliver realistic results for more complex scenarios that cannot be represented by an analytical supply function equilibrium model. Following this argument, they then use the computational model for analyzing different options for the second round of plant divestiture in the England and Wales electricity market in 1999. Results from various runs with different demand elasticities and different volumes of financial forward contracts show that the analyzed divestiture options result in lower average percentage bids above marginal cost. For a 25% and 50% plant divestiture, the average bid above marginal cost in different periods (summer, spring/autumn, and winter) significantly decreased as compared to the initial conditions without divestiture. One result that the authors point out is that the main reduction in mark-up is caused by the creation of five generators (from initially three); the difference in divestiture percentage has only a small effect on observed mark-ups. However, the authors also find that the proposed divestiture still leaves considerable market power with generators in the short term, and could result in prices more than 20% above short-run marginal costs.

In a later paper, [Bunn and Day \(2002\)](#) present this model as a competitive benchmark against which to assess generator conduct and to diagnose the separate causes of market structure and market conduct in situations in which prices appear to be above marginal costs. As the England and Wales electricity market cannot be characterized by perfect competition, but rather as a daily profit-maximizing oligopoly, the authors argue that fully competitive (marginal cost) baselines are not appropriate for assessing market power abuse. Instead, they argue that the result from their computational model can serve as a realistic baseline for imperfect (oligopoly) competition, where agents learn to compete, but not to collude. For the tested scenarios, the simulated system supply functions are shown to lie above the marginal cost function and significantly below the system supply curve observed in the England and Wales pool on an exemplary day, except at low demand levels. This leads the authors to the conclusion that the extent to which the simulated supply functions are above the marginal cost function is caused by the market structure. The extent to which observed system supply functions in the real-world market are still above the simulated system supply functions is then interpreted as the degree of collusion within the market, and identifies a problem of market conduct.

3.7. Large-scale national agent-based electricity simulations

To our knowledge, four national U.S. and one Australian laboratories are currently developing large-scale agent-based electricity system models which are intended to serve as tools for reliability and market design analysis for power markets:

- EMCAS, the Electricity Market Complex Adaptive System developed by Argonne National Laboratory (Conzelmann et al., 2005);
- Marketecture from Los Alamos National Laboratory¹⁸ (Atkins et al., 2004);
- N-ABLE™, the Agent-Based Laboratory for Economics developed at Sandia National Laboratory (Ehlen and Scholand, 2005);
- Simulations for Coupled Systems within the GridWise™ program at Pacific Northwest National Laboratory (Widergren et al., 2004); and
- NEMSIM, the Australian National Electricity Market Simulator which is under development at CSIRO (Batten and Grozev, 2006).

These models rely on very detailed databases of the regions under study, which includes the topology of the transmission grid and other physical constraints, differentiated load data and detailed cost data of the power plants. The issues that most of the models address are questions of best market designs to prevent the exercise of market power, transmission system reliability, or also environmental regulation measures. The models seem to be designed as a decision support for concrete policy making and not primarily for academic research. Consequently, the model descriptions are rather vague and much of the exact implementation of agent behavior or simulated scenarios remains unclear.

Some of the cited papers contain pointers to the literature about learning algorithms though no actual model implementation using any kind of agent learning is presented. The agents' behavior has been described in more detail for two models: In one scenario simulated with the *N-ABLE™* model, agents apply a heuristic planning process in the form of a *greedy* scheduling algorithm. The research question of this scenario was the influence of real-time pricing contracts on consumption and profitability in the retail electricity market (Ehlen et al., 2007). In the *Markecture* model agents follow one out of three possible fixed strategies: they set bid prices and quantities according to the *competitor*, *oligopolist*, or the *competitive-oligopolist* strategy. The *competitor* strategy refers to bidding at marginal cost, whereas the *oligopolist* bids at the point where the marginal revenue and marginal cost functions intersect; the *competitive-oligopolist* strategy lies at a random point in the range between the two other strategies. Buyers' and sellers' surplus, efficiency and market clearing prices/quantities are compared for three different market clearing algorithms (Atkins et al., 2004). However, as agents do not adapt to the different market clearing rules, no well-grounded conclusions can be drawn on the efficiency of these rules.

In summary, the scientific usefulness and academic contribution of large-scale AB models that integrate an enormous amount of details (for example, the demand representation in the *Markecture* model goes down to the level of every individual and its activity and mobility profile)

¹⁸ Leading scientists who developed the Marketecture model at the Los Alamos National Laboratory have now gone to the *Network Dynamics and Simulation Science Laboratory* (NDSSL) at Virginia Tech, so it is not clear whether they continue the work on Marketecture there. AB electricity market simulation seems to be one out of many other topics investigated at the NDSSL lab.

Table 1
Summarized overview of agent-based electricity market modeling approaches

References	Agents' actions	Market	Demand side	Transmission system	Research questions
<i>Model-based learning algorithms</i> (Section 3.1)					
Bower and Bunn (2000); Bower and Bunn (2001)	Set bid prices for portfolio of plants	Day-ahead market with uniform or pay-as-bid clearing	Static price responsive load	No transmission constraints	Comparison of pay-as-bid vs. uniform price clearing, and daily vs. hourly bidding (impact on prices)
Bower et al. (2001)	Set bid prices for portfolio of plants	Day-ahead market with uniform or pay-as-bid clearing	Static price responsive load	No transmission constraints	Impact of merger options in the German electricity market on wholesale prices
Bunn and Oliveira (2001); Bunn and Oliveira (2003)	Set mark-ups on bid prices in both markets separately	Forward market with pay-as-bid clearing, balancing mechanism	Active demand side bidding	No transmission constraints	Evaluation of generators' conduct in the England and Wales market: can either one of the generators, or two of them together exercise market power?
Bunn and Oliveira (in press)	Trade power plants with other agents; play Cournot strategy in the electricity market	Market for power plant capacity; day-ahead electricity market	Static hourly linear demand functions	No transmission constraints	Do electricity markets tend towards technological diversification or specialization? Which influence does the market clearing have on this question?
Visudhiphan and Ilić (2001); Visudhiphan (2003)	Set bid prices and quantities (step-wise bid functions)	Uniform price call market	Fixed inelastic demand	No transmission constraints	Realistic representation of market price dynamics and participants' bidding behavior in electricity markets; role of generator learning for strategic bidding
<i>Genetic algorithms</i> (Section 3.2)					
Cau and Anderson (2002); Cau (2003)	Assign bid quantities to price segments	Cost minimizing ISO	Inelastic, uncertain demand (high /low with equal probability)	No transmission constraints	Analysis of collusive strategies by the agents
Nicolaisen et al. (2000); Lane et al. (2000)	Set bid prices	Double auction with uniform pricing	Active demand side bidding	No transmission constraints	Measure market power exerted in a double auction
Richter and Sheblé (1998); Petrov and Sheblé (2000)	Set bid prices	Double auction with uniform pricing	Price inelastic, no demand-side bidding	No transmission constraints	Examination of bidding strategies
<i>Erev and Roth reinforcement learning</i> (Section 3.3)					
Bin et al. (2004)	Set bid prices	Call market with uniform, pay-as-bid and electricity value equivalent pricing	Fixed inelastic demand	No transmission constraints	Comparison of resulting prices for the three pricing mechanisms
Cincotti et al. (2005); Cincotti et al. (2006)	Set bid prices, or price quantity pairs	Call market	Fixed inelastic demand	No transmission constraints	Comparison of bidding strategies and resulting prices in pay-as-bid and uniform price auctions
Nicolaisen et al. (2001)	Set bid prices	Double-auction with discriminatory midpoint pricing	Adaptive demand side (both buyers and sellers of electricity bid in the auction)	No transmission constraints	Analysis of buyers and sellers market power under different concentration conditions; distinction between structural market power and market power due to agent learning

(continued on next page)

Table 1 (continued)

References	Agents' actions	Market	Demand side	Transmission system	Research questions
Rupérez Micola et al. (2006)	Set bid prices on three subsequent markets	Wholesale and retail electricity market, natural gas market	Fixed inelastic demand	No transmission constraints	Can agents benefit from vertical integration? How can agents exert vertical market power?
Sun and Tesfatsion (2007)	Set supply functions	Real-time market, day-ahead market with locational marginal pricing	Fixed inelastic demand	Bid-based DC optimal power flow problem	Market design reliability
Veit et al. (2006)	Set bid quantities	ISO for spot and forward trading	Demand function with intercept and slope	Nodal prices (DC power flow problem)	Dynamics between forward trading and a spot market
Weidlich and Veit (2006)	Set price quantity bid pairs	Call market, procurement auction for balancing power	Fixed inelastic demand	No transmission constraints	Dynamics between two interrelated markets, comparison prices in pay-as-bid and uniform price settlement case
<i>Q-learning (Section 3.4)</i>					
Bakirtzis and Tellidou (2006)	Set bid prices	Uniform price and pay-as-bid call market	Fixed inelastic load	No transmission constraints	Comparison of market prices in uniform price and pay-as-bid auctions
Krause et al. (2005); Krause and Andersson (2006)	Set bid prices	Generation cost minimizing ISO (considers network constraints)	Fixed demand; linear function of demand side marginal benefit	DC network representation with transmission capacity constraints	Comparison between <i>Q</i> -learning and Nash equilibrium strategies; evaluation of different congestion management mechanisms
Naghbi-Sistani et al. (2006)	Set slope of linear bid function	Uniform price call market	Fixed price-elastic demand	No transmission constraints	Comparison between Nash equilibria and the proposed <i>Q</i> -learning algorithm with temperature variation
Xiong et al. (2004)	Set bid prices	Uniform price and pay-as-bid call market	Both fixed inelastic demand and interruptible loads (demand response)	No transmission constraints	Comparison of market prices in uniform price and pay-as-bid auctions
<i>Learning classifier systems (Section 3.5)</i>					
Bagnall (2000b); Bagnall and Smith (2005)	Set bid prices	ISO who integrates unit commitment constraints into the allocation calculation	Fixed half-hourly forecast demand	No transmission constraints	Can the agents evolve behaviors observable in the real world? How do market mechanisms (uniform price vs. pay-as-bid) effect bidding behavior?
<i>Supply function optimization heuristic (Section 3.6)</i>					
Day and Bunn (2001); Bunn and Day (2002)	Allocate plant capacity to price bins	ISO calculating the system marginal price	Fixed, price-elastic demand (locally linear functions)	No transmission constraints	Differentiation between market structure and market conduct as a reason for high electricity prices; analysis of the impact of divestiture proposals on the possibility of players to exert market power

has not yet been proven. The practical usefulness cannot be judged here, as hardly any – or only illustrative – simulation results from the cited models are publicly available.

4. Summary and discussion

In order to allow a better comparison and a brief overview of the different modeling approaches, the presented papers are shortly summarized in [Table 1](#) (entries in this table are categorized according to the applied learning model applied; within one category, the cited papers are ordered alphabetically). In those cases in which the same authors have described similar or enhanced models in several papers, only the most relevant reference has been specified. As for their rather practical application and not primarily academic focus, the large-scale simulation models presented in Section 3.7 are not included in this overview.

The comparison of the different models shows the similarities and differences between current AB electricity models:

- The large majority of models neglect transmission grid constraints.
- Most of the models represent the demand side as a fixed, price-insensitive load.
- In most models, the agents' learning task is to set profit-maximizing bid prices or mark-ups. Capacity withholding strategies are mostly not modeled explicitly; however, setting a high bid price can also be interpreted as (economic) withholding.
- No preferred learning representation or trend towards specific models of behavior can be observed. A number of models rely on the reinforcement learning algorithm formulated by [Erev and Roth \(1998\)](#) (and modifications of this algorithm); however, they do not form a considerable majority. Genetic algorithms seem to be left apart, though not completely abandoned.
- Most research questions of AB modelers center around market power and market mechanisms. The comparison between pay-as-bid and uniform pricing is a very popular question. Here, results from different models seem to be consistent; most authors find that agents bid higher under pay-as-bid, but overall prices are higher under uniform pricing. Another important research issue in for AB electricity modelers is the assessment of potential for market power under different market structures or market mechanisms.

The amount of papers reviewed in this survey shows that electricity market research applying AB simulation is a very active field of research. One might describe it as adolescent — it has departed from its infancy which began in the last years of the past century. This is documented by the appearance of the first notable papers that have successfully been published in energy-related or other journals. However, we still observe a large heterogeneity in representing boundedly rational actors in electricity markets, and also in validation techniques, result evaluation and quality assessment, or simply in labeling. Agent-based modeling allows for great flexibility in specifying how agents behave; the reverse of this medal is that models are rarely comparable, and can sometimes not be described in all necessary detail. On its way to adulthood, hence, several methodological questions will have to be discussed by researchers who are active in this field, in order to increase the comparability of different models. Some of these issues are enumerated in the following.

4.1. Agent learning behavior

The common element of all models presented here is adaptivity. Agents are able to learn to achieve their goals (high profits, high plant utilization etc.) given the environment they are placed

in. The way in which learning takes place is implemented differently in almost all models. Even if two researchers use the same basic learning algorithm, they may define and describe it in a different manner. Parameter values are often not justified properly and are not fully revealed. A precise description of an agent's action domain or the space of possible environmental or internal states is not defined clearly in every paper. Especially in the models applying Q -learning, the definition of environmental states has implications on the model results; however, none of the cited papers using Q -learning argues why the states have been set as they are, and whether they have the Markov property.

Moreover, the choice of the learning algorithm itself is hardly argued and justified in any paper. Some algorithms like e.g. the Erev and Roth reinforcement learning formulation, are popular and used by many researchers. Others are hardly used at all, without apparent reason.¹⁹ Most papers do not answer the question why they chose a specific learning model and how good it performs in comparison to others. Also, it might be interesting to discuss if there is any meaningful minimum level of rationality that agents participating in electricity markets should be endowed with.

4.2. Market dynamics and complexity

As stated in the introduction, the electricity sector is characterized by the interlinking of multiple markets, and by additional complexities through limited transmission capacities and unit commitment constraints. Most papers considered here simplify real-world markets significantly; they consider only one market or neglect technical constraints. Some AB models are so highly stylized that they cannot claim to be more realistic than traditional equilibrium models. Also, the focus of most researchers is placed on convergence towards stable market outcomes. The out-of-equilibrium dynamics or the way towards an equilibrium are not considered. It might also be interesting to examine under which circumstances agents reach an equilibrium outcome and when they fail to do so; or, in case of multiple equilibria, which outcome occurs more frequently and how robust these outcomes are against some changes in parameter values.

These considerations also lead to the question of what can at all be considered the outcome of an agent-based simulation. Most researchers let the simulation run a specified number of iterations and then calculate some aggregate values from the late iterations (considering early iterations as a settling phase for the learning algorithm). Usually the results are boiled down to one convergence price per simulation. We are not aware of any paper in which the characteristics of the time series of prices are examined in more detail (in order to answer questions like "How volatile are prices?", "Can price spikes be observed?", "Are prices mean reverting?", etc.). Also, the agents' profits and success of different trading strategies are rarely discussed.

One important aspect of the electricity sector that can perhaps best be represented in AB models is not considered in any of the presented papers: bilateral trading. It might be interesting to compare the efficiency of market outcomes in a bilateral setting with that of a centralized auction. Moreover, vertical integration could realistically be modeled in agent-based simulations. With few exceptions, however, this aspect is neglected in current agent-based electricity sector models. We would suggest that these still neglected factors should be stressed more in future modeling approaches.

¹⁹ To give an example for a learning representation that has hardly gained any attention by AB electricity researchers, one might mention Experience-Weighted Attraction, which is a learning model that combines aspects of reinforcement learning and belief-based learning. It has been formulated by [Camerer and Ho \(1999\)](#) and has shown a good fit when compared to three classes of games.

4.3. Calibration and validation

Calibrating and validating agent-based electricity models is a challenging task, and only few guidelines for this process have yet been defined. To our knowledge, only [Macal and North \(2005\)](#) have reported the process of validating their model using different techniques. Many of the reviewed papers in this survey lack information about empirical model validation; those researchers who have undertaken some empirical validation proceeded in heterogeneous ways.

While standard verification and validation techniques for simulation models (e.g. described by [Sargent 2005](#) or [Law 2007](#)) can and should also be applied for AB simulations, it is difficult to establish credibility in the implemented agent behavior. Very recently, the need for reliable validation techniques has obviously been recognized. AB researchers have analyzed and suggested procedures and guidelines for calibrating and validating agent-based simulation models, e.g. [Windrum et al. \(2007\)](#), [Marks \(2007\)](#), [Richiardi et al. \(2006\)](#), [Midgley et al. \(2007\)](#), and a whole journal special issue is devoted to this topic ([Fagiolo et al., 2007](#)). These general suggestions should now be assessed with regard to their usefulness for electricity modeling purposes. The development of guidelines for assuring the validity of AB electricity models would greatly benefit the research quality and diminish the heterogeneity of approaches in this field. It should thus be one of the main tasks for future work.

4.4. Model description and publication

Just as model verification and validation is done very differently in the cited papers, so is the description of the model, its parameters, and the results. Some papers do not deliver information about the number of runs they have conducted, some do not even publish all model parameters so that it is not possible to replicate the reported results. It would be helpful if some standard way of model description, as it is conventional for other economic methodologies, became accepted in the medium-term.

Presumably, many models are still used by their developers, so that these are reluctant to make their source code available. However, this would greatly benefit the research field, because researchers could revise and check the implementations of others and could also reuse parts of them. Leigh Tesfatsion has taken the initiative in this direction by setting up a website with links to published sets of AB electricity model source code²⁰ (and by publishing her source code as well).

5. Conclusion

This paper has critically reviewed a considerable amount of relevant papers in agent-based electricity market research. [Table 1](#) summarizes the core characteristics of the cited work and displays the similarities and differences between the approaches. In Section 4 we have identified some of the current problems facing this research methodology that require further effort and a consolidation of the approaches pursued by different research groups. Especially sound argumentations for the choice of specific learning algorithms, more careful and well documented validation and verification procedures as well as the appropriate publication of details of concrete simulation models are crucial for the further development of agent-based electricity market modeling.

²⁰ <http://www.econ.iastate.edu/tesfatsi/ElectricOSS.htm>.

Despite the open issues and problems, AB electricity research has been successful in recent time. Many AB researchers have successfully replicated core characteristics of today's electricity markets using models with adaptive, self-seeking agents. With a decrease of heterogeneity between competing models, and with increasing consensus on important methodological questions, the field of AB modeling can soon become one major strand of research for the analysis of complex electricity systems.

Acknowledgements

This research has been conducted within the PowerACE project (www.powerace.de). We gratefully thank the VolkswagenStiftung for funding this work.

References

- Arifovic, J., Ledyard, J., 2004. Scaling up learning models in public good games. *Journal of Public Economic Theory* 6 (2), 203–238.
- Atkins, K., Barrett, C.L., Homan, C.M., Marathe, A., Marathe, M.V., Thite, S., 2004. Agent based economic analysis of deregulated electricity markets. 6th IAEE European Conference. *Modelling in Energy Economics and Policy*, Zürich, Switzerland.
- Auer, P., Cesa-Bianchi, N., Freund, Y., Schapire, R., 2002. The nonstochastic multiarmed bandit problem. *Siam Journal on Computing* 32 (1), 48–77.
- Axelrod, R., 1997. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton University Press, Princeton, New Jersey.
- Bagnall, A., 2004. Applications of learning classifier systems, Chapter A multi-agent model of the UK market in electricity generation. *Studies in Fuzziness and Soft Computing*. Springer Verlag, pp. 167–181.
- Bagnall, A., Smith, G., 1999. Proceedings of artificial evolution. *Lecture Notes in Computer Science*, Chapter An Adaptive Agent Model for Generator Company Bidding in the UK Power Pool, vol. 1829. Springer Verlag, pp. 191–203.
- Bagnall, A., Smith, G., 2000. Game playing with autonomous adaptive agents in a simplified economic model of the UK market in electricity generation. *IEEE International Conference on Power System Technology*, pp. 891–896.
- 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.
- Bagnall, A.J., 2000a. A multi-adaptive agent model of generator bidding in the UK market in electricity. In: Whitely, D., Goldberg, D., Cantú-Paz, E., Spector, L., Parmee, I., Beyer, H.-G. (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2000)*. Morgan Kaufmann, San Francisco, USA.
- Bagnall, A.J., 2000b. *Modelling the UK Market in Electricity Generation with Autonomous Adaptive Agents*. Ph. D. thesis, University of East Anglia, School of Information Systems.
- Bakirtzis, A.G., Tellidou, A.C., 2006. Agent-based simulation of power markets under uniform and pay-as-bid pricing rules using reinforcement learning. *Proceedings of the IEEE Power Systems Conference and Exposition PSCE '06*, Atlanta, USA, pp. 1168–1173.
- Batten, D., Grozev, G., 2006. Complex science for a complex world: exploring human ecosystems with agents. Chapter NEMSIM: Finding Ways to Reduce Greenhouse Gas Emissions Using Multi-Agent Electricity Modelling. ANU E Press, pp. 227–252.
- Bernal-Aguistin, J.L., Contreras, J., Martín-Flores, R., Conejo, A.J., 2007. Realistic electricity market simulator for energy and economic studies. *Electric Power Systems Research* 77, 46–54.
- Bin, Z., Maosong, Y., Xianya, X., 2004. The comparisons between pricing methods on pool-based electricity market using agent-based simulation. *Proceedings of the IEEE International Conference on Electric Utility Deregulation, Restructuring and Power Technologies (DRPT2004)*, Hong Kong.
- Bower, J., Bunn, D.W., 2000. Model-based comparisons of pool and bilateral markets for electricity. *Energy Journal* 21 (3), 1–29.
- Bower, J., Bunn, D., 2001. Experimental analysis of the efficiency of uniform-price versus discriminatory auctions in the England and Wales electricity market. *Journal of Economic Dynamics & Control* 25, 561–592.
- Bower, J., Bunn, D.W., Wattendrup, C., 2001. A model-based analysis of strategic consolidation in the German electricity industry. *Energy Policy* 29 (12), 987–1005.

- Bunn, D.W., Oliveira, F.S., 2001. Agent-based simulation: an application to the new electricity trading arrangements of England and Wales. *IEEE Transactions on Evolutionary Computation*, Special issue: Agent Based Computational Economics 5 (5), 493–503.
- Bunn, D.W., Day, C.J., 2002. Computational modelling of price-formation in the electricity pool of England and Wales. London Business School Decision Sciences Working Paper.
- Bunn, D.W., Oliveira, F.S., 2003. Evaluating individual market power in electricity markets via agent-based simulation. *Annals of Operations Research* 121 (1-4), 57–77.
- Bunn, D.W., Oliveira, F.S., in press. Agent-based analysis of technological diversification and specialization in electricity markets. *European Journal of Operational Research*, 1–13.
- Camerer, C., Ho, T.H., 1999. Experienced-weighted attraction learning in normal form games. *Econometrica* 67 (4), 827–874.
- Cau, T.D.H., 2003. Analysing Tacit Collusion in Oligopolistic Electricity Markets Using a Co-evolutionary Approach. Ph. D. thesis, Australian Graduate School of Management.
- Cau, T.D.H., Anderson, E.J., 2002. A co-evolutionary approach to modelling the behaviour of participants in competitive electricity markets. *IEEE Power Engineering Society Summer Meeting* 3, 1534–1540.
- Cincotti, S., Guerci, E., 2005. Agent-based simulation of power exchange with heterogeneous production companies. *Computing in Economics and Finance 2005*, Society for Computational Economics. Number 334. available at <http://ideas.repec.org/p/sce/scecf5/334.html>.
- 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.
- Cincotti, S., Guerci, E., Ivaldi, S., Raberto, M., 2006. Discriminatory versus uniform electricity auctions in a duopolistic competition scenario with learning agents. *IEEE Congress on Evolutionary Computation*, Vancouver, Canada.
- Conzelmann, G., Boyd, G., Koritarov, V., Veselka, T., 2005. Multi-agent power market simulation using EM-CAS. *IEEE Power Engineering Society General Meeting*, vol. 3, pp. 2829–2834.
- Curzon Price, T., 1997. Using co-evolutionary programming to simulate strategic behaviour in markets. *Journal of Evolutionary Economics* 7, 219–254.
- Day, C.J., Bunn, D.W., 2001. Divestiture of generation assets in the electricity pool of England and Wales: a computational approach to analyzing market power. *Journal of Regulatory Economics* 19 (2), 123–141.
- Drogoul, A., Vanbergue, D., Meurisse, T., 2003. Multi-agent based simulation: where are the agents? In: Sichman, J.S., Bousquet, F., Davidsson, P. (Eds.), *Multi-Agent-Based Simulation*, Proceedings of the Third International Workshop MABS 2002, Revised Papers. Lecture Notes in Computer Science, vol. 2581. Springer, pp. 1–15.
- Ehlen, M., Scholand, A., 2005. Modeling interdependencies between power and economic sectors using the N-ABLE agent-based model. *Proceedings of the IEEE Power Engineering Society General Meeting*, pp. 2842–2846.
- Ehlen, M.A., Scholand, A.J., Stamber, K.L., 2007. The effects of residential real-time pricing contracts on transco loads, pricing, and profitability: Simulations using the N-ABLE™ agent-based model. *Energy Economics* 29, 211–227.
- Epstein, J.M., Axtell, R.L., 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. The MIT Press.
- Erev, I., Roth, A.E., 1998. Predicting how people play games: reinforcement learning in experimental games with unique, mixed-strategy equilibria. *American Economic Review* 88 (4), 848–881.
- Ernst, D., Minoia, A., Ilić, M., 2004a. Market dynamics driven by the decision-making of both power producers and transmission owners. *IEEE Power Engineering Society General Meeting*, vol. 1, pp. 255–260.
- Ernst, D., Minoia, A., Ilić, M., 2004b. Market dynamics driven by the decision-making of power producers. *Proceedings of Bulk Power System Dynamics and Control - IV*, Cortina d'Ampezzo, Italy.
- Fagiolo, G., Birchenhall, C., Windrum, P., 2007. Empirical validation in agent-based models: introduction to the special issue. *Computational Economics* 30 (3), 189–194.
- Franklin, S., Graesser, A., 1997. Is it an agent or just a program? A taxonomy for autonomous agents. In: Müller, J., Wooldridge, M., Jennings, N. (Eds.), *Intelligent Agents III: Agent Theories, Architectures, and Languages*, ECAI96 Workshop Proceedings. Springer-Verlag, pp. 21–35.
- Greenwald, A.R., Kephart, J.O., Tesauro, G.J., 1999. Strategic pricebot dynamics. *EC '99: Proceedings of the 1st ACM conference on Electronic commerce*, New York, NY, USA. ACM Press, pp. 58–67.
- Guo, M., Liu, Y., Malec, J., 2004. A new Q-learning algorithm based on the metropolis criterion. *IEEE Transactions on Systems, Man, and Cybernetics* 34 (5), 2140–2143.
- Hämäläinen, R.P., Mäntysaari, J., Ruusunen, J., Pineau, P.O., 2000. Cooperative consumers in a deregulated electricity market — dynamic consumption strategies and price coordination. *Energy* 25, 857–875.
- Harp, S.A., Brignone, S., Wollenberg, B.F., Samad, T., 2000. SEPIA — a simulator for electric power industry agents. *IEEE Control Systems Magazine* 20 (4), 53–69.

- Holland, J.H., Miller, J.H., 1991. Artificial adaptive agents in economic theory. *American Economic Review* 81 (2), 365–370.
- Klemperer, P.D., Meyer, M.A., 1989. Supply function equilibria in oligopoly under uncertainty. *Econometrica* 57 (6), 1243–1277.
- Koesrindartoto, D., 2002. Discrete double auctions with artificial adaptive agents: a case study of an electricity market using a double auction simulator. Department of Economics Working Papers Series. Working Paper, vol. 02005.
- Koesrindartoto, D., Tesfatsion, L., 2004. Testing the reliability of FERC's wholesale power market platform: an agent-based computational economics approach. Proceedings of the 24th USAEE/IAEE North American Conference, Washington, D.C., USA.
- Koesrindartoto, D., Sun, J., Tesfatsion, L., 2005. An Agent-Based Computational Laboratory for Testing the Economic Reliability of Wholesale Power Market Designs. Proceedings of the IEEE Power Engineering Society General Meeting, vol. 3, pp. 2818–2823.
- Krause, T., Andersson, G., 2006. Evaluating congestion management schemes in liberalized electricity markets using an agent-based simulator. Proceedings of the Power Engineering Society General Meeting, Montreal.
- Krause, T., Andersson, G., Ernst, D., Beck, E.V., Cherkaoui, R., Germond, A., 2005. A comparison of Nash equilibria analysis and agent-based modelling for power markets. Proceedings of the Power Systems Computation Conference (PSCC), Liege, Belgium.
- Lane, D., Kroujiline, A., Petrov, V., Sheblé, G., 2000. Electricity market power: marginal cost and relative capacity effects. Proceedings of the 2000 Congress on Evolutionary Computation, vol. 2. La Jolla, USA, pp. 1048–1055.
- Law, A.M., 2007. *Simulation Modeling & Analysis*, fourth ed. McGraw-Hill Higher Education.
- Macal, C.M., North, M.J., 2005. Validation of an agent-based model of deregulated electric power markets. Technical report, Los Alamos National Laboratory, presented at the North American Association for Computational and Social Organization (NAACSOS) Conference.
- Marimon, R., McGrattan, E., 1995. Learning and Rationality in Economics, Chapter On adaptive learning in strategic games. Blackwell, pp. 63–101.
- Marks, R.E., 2007. Validating simulation models: a general framework and four applied examples. *Computational Economics* 30 (3), 265–290.
- Midgley, D.F., Marks, R.E., Kunchamwar, D., 2007. The building and assurance of agent-based models: an example and challenge to the field. *Journal of Business Research* 60, 884–893.
- Müller, M., Sensfuß, F., Wietschel, M., 2007. Simulation of current pricing-tendencies in the German electricity market for private consumption. *Energy Policy* 35 (8), 4283–4294.
- Naghibi-Sistani, M., Akbarzadeh - Tootoonchi, M., Javidi-Dashte Bayaz, M., Rajabi-Mashhadi, H., 2006. Application of Q-learning with temperature variation for bidding strategies in market-based power systems. *Energy Conversion and Management* 47, 1529–1538.
- Nicolaisen, J., Smith, M., Petrov, V., Tesfatsion, L., 2000. Concentration and capacity effects on electricity market power. Proceedings of the 2000 Congress on Evolutionary Computation, vol. 2. La Jolla, USA, pp. 1041–1047.
- Nicolaisen, J., Petrov, V., Tesfatsion, L., 2001. Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation* 5 (5), 504–523.
- Petrov, V., Sheblé, G.B., 2000. Power auctions bid generation with adaptive agents using genetic programming. Proceedings of the 2000 North American Power Symposium, Institute of Electrical and Electronic Engineers, Waterloo-Ontario, Canada.
- Petrov, V., Sheblé, G., 2001. Building electric power auctions with improved roth-erev reinforced learning. Proceedings of the North American Power Symposium, Texas, USA.
- Praça, I., Ramos, C., Vale, Z., Cordeiro, M., 2004. Intelligent agents for the simulation of competitive electricity markets. *International Journal of Modelling & Simulation* 2, 73–79.
- Richiardi, M., 2004. The Promises and Perils of Agent- Based Computational Economics. EconWPA Working Paper, Report- no mgr040107.
- Richiardi, M., Leombruni, R., Saam, N., Sonnessa, M., 2006. A common protocol for agent-based social simulation. *Journal of Artificial Societies and Social Simulation* 9 (1).
- 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.
- Roop, J.M., Fathelrahman, E., 2003. Modeling electricity contract choice: an agent-based approach. Proceedings of the ACEEE Summer Study Meeting, Rye Brook, New York. Rye Brook, New York.
- Rothkopf, M.H., 1999. Daily repetition: a neglected factor in the analysis of electricity auctions. *The Electricity Journal* 12 (3), 60–70.
- Rupérez Micola, A., Banal Estañol, A., and Bunn, D.W., 2006. Incentives and Coordination in Vertically Related Energy Markets. Discussion Paper SP II 2006 – 02, Wissenschaftszentrum Berlin.

- Sargent, R.G., 2005. Verification and validation of simulation models. Proceedings of the Winter Simulation Conference.
- Sun, J., Tesfatsion, L., 2006. DC optimal power flow formulation and solution using QuadProgJ. Technical report, Economics Working Paper, Department of Economics. Iowa State University.
- Sun, J., Tesfatsion, L., 2007. Dynamic testing of wholesale power market designs: an open-source agent-based framework. *Computational Economics* 30 (3), 291–327.
- Sutton, R.S., Barto, A.G., 1998. Reinforcement Learning: An Introduction. MIT Press.
- Tesfatsion, L., 2002. Agent-based computational economics: growing economies from the bottom up. *Artificial Life* 8, 55–82.
- Tesfatsion, L., 2006. Handbook of computational economics. Chapter Agent-Based Computational Economics: A Constructive Approach to Economic Theory. *Agent-Based Computational Economics*, vol. 2. North-Holland, pp. 831–880.
- Veit, D.J., Weidlich, A., Yao, J., Oren, S.S., 2006. Simulating the dynamics in two-settlement electricity markets via an agent-based approach. *International Journal of Management Science and Engineering Management* 1 (2), 83–97.
- Visudhiphan, P., 2003. An Agent-Based Approach to Modeling Electricity Spot Markets. Ph.D. thesis, Massachusetts Institute of Technology.
- Visudhiphan, P., Ilić, M.D., 1999. Dynamic games-based modeling of electricity markets. *Power Engineering Society 1999 Winter Meeting. IEEE*, vol. 1, pp. 274–281.
- Visudhiphan, P., Ilić, M.D., 2001. An agent-based approach to modeling electricity spot markets. *Proceedings of IFAC Modeling and Control of Economic Systems*, Klagenfurt, Austria, pp. 407–412.
- Visudhiphan, P., Ilić, M.D., 2002. On the Necessity of an Agent-based Approach to Assessing Market Power in the Electricity Markets. *Proceedings of the Tenth International Symposium on Dynamic Games and Applications*, vol. 2. Sankt Petersburg, Russia.
- Watkins, C.J.C.H., 1989. Learning from Delayed Rewards. Ph. D. thesis, Cambridge.
- Weidlich, A., Veit, D., 2006. Bidding in interrelated day-ahead electricity markets: insights from an agent-based simulation model. *Proceedings of the 29th IAEE International Conference*, Potsdam.
- Widergren, S.E., Roop, J.M., Guttromson, R.T., Huang, Z., 2004. Simulating the dynamic coupling of market and physical system operations. *Proceedings of the IEEE Power Engineering Society General Meeting*, vol. 1, pp. 748–753.
- Windrum, P., Fagiolo, G., Moneta, A., 2007. Empirical validation of agent-based models: alternatives and prospects. *Journal of Artificial Societies and Social Simulation* 10 (2, 8).
- Xiong, G., Okuma, S., Fujita, H., 2004. Multi-agent based experiments on uniform price and pay-as-bid electricity auction markets. *Proceedings of the IEEE International Conference on Electric Utility Deregulation, Restructuring and Power Technologies (DRPT2004)*, Hong Kong.
- Yu, J., Zhou, J.Z., Yang, J., Wu, W., Fu, B., Liao, R.T., 2004. Agent-based retail electricity market: modeling and analysis. *Proceedings of the Third International Conference on Machine Learning and Cybernetics*, Shanghai.