

Agent-Based Simulations for Electricity Market Regulation Advice: Procedures and an Example

Anke Weidlich and Daniel Veit, Mannheim

JEL C63, D43, L49

Agent-based computational economics, simulation, electricity markets, CO2 emissions trading, market interrelations.

Summary

This paper discusses the use of agent-based simulation models for regulatory advice in electricity market regulation. It briefly introduces the necessary procedures and the state-of-the-art of the methodology, and outlines its possible range of application. In a second part, the paper presents an agent-based simulation model developed by the authors. The model can be applied for analyzing different market designs and market structures in order to derive evidence for regulatory advice. This is exemplified through the analysis of two settlement rules in the balancing power market and of several divestiture scenarios of the German electricity sector.

1 Introduction

The electricity sector is characterized by technical constraints, multiple interlinked markets, and an oligopolistic structure with vertical integration. Asymmetric information, imperfect competition, strategic interaction, collective learning, and the possibility of multiple equilibria are further difficult aspects that characterize the electricity sector, like economies in general (Tesfatsion 2006). The ensemble of these aspects make electricity markets rank among the most complex of all commodity markets operated at present, and push most classical economic modeling methods to their limits.

The complexity of the electricity sector and its high importance for a competitive economy calls for new modeling methods that help gaining insights into the dynamics of power markets. The agent-based (AB) modeling methodology, or more specifically the field of Agent-Based Computational Economics (ACE), provides more flexibility for properly representing relevant complex and may help to overcome some of the limits of traditional modeling methods, like e.g. too strict assumptions. With ACE research, the focus in economic analysis is shifted from rational behavior and equilibrium towards heterogeneity and adaptivity. The tremendous availability of computational resources made it possible to set up large-scale and detailed computational models that allow a high degree of design flexibility. Populations of heterogeneous agents, feedback from interaction, and dynamic processes make up the core of computational economic models.

Although it constitutes a change in economic modeling paradigm, the agent-based approach does not need to be seen in sharp contrast with analytical modeling. Like Gulyás (2002) argues, an agent-based implementation is “rather a matter of degree than a binary choice”. When moving from top-down approaches in which firms or other actors are represented on an aggregate level, via an entity-level representation of actors, down

to the introduction of autonomous agents as actor representatives, frontiers between the modeling paradigms are fluent. However, with each step towards an agent-based implementation, the modeler gains flexibility, as more aspects, like e.g. heterogeneity of agents, can be accounted for.

During the last decade, many agent-based simulation models of electricity markets have been developed. A majority of these approaches use reinforcement learning algorithms for representing the adaptive behavior of players in the market, e.g. Sun/Tesfatsion (2007), Nicolaisen et al. (2001), Bower/Bunn (2001); the use of *genetic algorithms* (e.g. Cau/Anderson (2003) or *learning classifier systems* (e.g. Bagnall/Smith 2005) is less frequent in AB electricity market modeling. While earlier models represent the electricity sector as one single, strongly simplified market, some recent models comprise several markets (Sun/Tesfatsion 2007), and explicitly model interrelations between them (Rupérez Micola et al. 2006).

In most models, the agent's learning task is to set profit-maximizing offer prices or mark-ups. Capacity withholding strategies are mostly not modeled explicitly; however, setting a high offer price can also be interpreted as (economic) withholding. Another observation from surveying the agent-based electricity simulation literature is that the majority of models still neglect transmission grid constraints (although recent developments go into the direction of taking transmission into account, e.g. Sun/Tesfatsion (2007), and most models represent the demand side as a fixed, price-insensitive load.

ACE models serve for explaining, understanding, and analyzing socio-economic phenomena. Among the ACE research strands that Tesfatsion (2006) identifies, most electricity related research can be categorized into the normative strand. Agent-based simulation models are used as fully controllable virtual laboratories for testing economic design alternatives in order to determine the policies, institutions, or processes that perform best in an environment of selfish agents. This approach follows the postulation formulated by Roth (2002) that markets should be designed by using engineering tools, such as experimentation and computation. More specifically, most AB electricity market models center around questions of market power and market mechanisms. The comparison between pay-as-bid and uniform pricing is a very popular research question analyzed with AB approaches. Another important research issue for AB electricity modelers is the assessment of market power potential under different market structures or market mechanisms.

In this paper, the potential of agent-based modeling for analyzing market structures and market designs with the aim of advising decision makers in the electricity sector is discussed. Procedures of model building and model validation are presented, and current weaknesses of the methodology are pointed out in Section 2. In Section 3, an example of an agent-based electricity sector simulation model developed by the authors is described, and simulation results from this model are presented. Finally, Section 4 concludes and gives an outlook on further research directions to be conducted in the agent-based electricity market simulation field.

2 Procedure of agent-based modeling for regulation advice

In the following, the state of the art of the methodology is briefly discussed, with a special focus on its applications to electricity market modeling. Some basic concepts that characterize and motivate the use of agent-based approaches are presented in Section 2.1. The

model building process is described in Section 2.2. Section 2.3 discusses the important issue of (empirical) model validation.

2.1 Concepts and motivation for ACE models

According to Axtell (2000), one main motivation for agent-based models is the dissatisfaction with rational agents. Thus, he argues, all agent-based models involve some form of boundedly rational agents. In fact, many economists argue that people, although they might try to be rational, have natural limits to perception and information collection, memory, and computational capacity, and can consequently rarely meet the requirement of information or foresight that rational models impose. A realistic assumption in electricity market modeling is that agents do not have all information necessary for making an *optimal* choice of an action to take in order to maximize profits. Instead, they can *learn* to behave profitably through repeated interaction with the environment they are placed in.

Consequently, the most prominent approach of modeling bounded rationality that can be found in the AB electricity market modeling literature is the gradual amelioration of activities through *learning*, *dynamic adaptation*, or *evolution*. Learning and adaptation are obvious representations of human behavior in complex economic situations. As for their realism and usefulness for modeling agents in many economic systems, Batten (2000) postulates that [l]earning and adaptation should not be addenda to the central theory of economics. They should be at its core, especially in problems of high complexity.

Another essential element of agent-based modeling 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. Heterogeneity is an important characteristic of real-world electricity markets. Generator agents differ in size and spatial position, they own and operate different generating technologies (e.g. fossil, nuclear or renewable power plants) with different marginal costs and technical attributes, or they have different strategic characteristics (e.g. vertically integrated or not). During the process of trading, they further differentiate from each other through the individual experience they gain from trading. While no agent has complete information about the global state, each of them accumulates some knowledge of which strategies might be more or less successful when competing with the rest of the population. On this basis, each agent decides how best to act subsequently in order to maximize profits. By this means, aggregate system behavior, such as collusion, might be observed. Agent-based simulation is a natural and intuitive way of representing this agent heterogeneity, as it offers much more flexibility in defining each individual market participant than do analytical modeling methods.

2.2 Model building

There are only few guidelines proposed for the design and implementation process of agent-based simulation models. As a start, basic principles of simulation procedures used in other disciplines also apply to agent-based simulation. Law (2007) for example provides helpful guidance for building simulation models and analyzing output data;

Gilbert/Troitzsch (2005) describe simulation methodologies in the social sciences, including some hints for agent-based modeling.

According to Tesfatsion (2002), the ACE modeling procedure can be described as follows: 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 initial attributes of the agents (e.g. type characteristics, learning behavior, knowledge about itself and other agents); 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.

In AB electricity simulations, the most common agents that make up the population are generators, load serving entities, and a market or 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 e.g. utilities.

At the subsequent step of model building, that is model description and publication, AB electricity market modelers proceed in heterogeneous ways. It would be helpful if some standard way of model description, as it is conventional for other economic methodologies, became accepted for AB modeling in the medium-term. This description should include information about the number of runs conducted, the applied parameter values, and all other model details that allow for reproducing the described simulations. In order to make agent-based simulations better understandable and publishable, Richiardi et al. (2006) give some practical suggestions towards a standardized methodological protocol of AB simulations; they propose to:

- include references to the theoretical background of the economic phenomenon that is investigated, including simulation and non-simulation literature;
- state the main features of the simulation model (treatment of time: discrete or continuous, treatment of fate: stochastic or deterministic, coordination structure: centralized or decentralized, and others) clearly and immediately in order to facilitate understanding and model comparison;
- follow well-defined processes for data analysis, including accepted calibration, validation and sensitivity analysis techniques;
- use standard modeling languages such as the Unified Modeling Language (UML) for describing the static and dynamic properties of the model and use AB modeling toolkits, in order to make models more easily replicable.

Moreover, making model source codes publicly available would greatly benefit the research field, because researchers could revise and check the implementations of others and could also reuse parts of them¹.

¹ For a positive step into this direction see Leigh Tesfatsion's site at <http://www.econ.iastate.edu/tesfatsi/ElectricOSS.htm>.

2.3 Validation of agent-based simulation models

The difficulty of validating ACE model outcomes against empirical data remains one of the challenges for the ACE methodology.² Only few guidelines for calibrating and validating agent-based simulation models have yet been defined. While standard verification and validation techniques for simulation models (e.g. described by Sargent 2005, Gilbert/Troitzsch 2005, or Law 2007) can and should also be applied for AB simulations, it is difficult to establish credibility in the implemented agent behavior.

Many models presented in the agent-based electricity modeling literature are not empirically validated. In their survey of the relevant literature, Weidlich/Veit 2008a illustrate this finding, which is consistent with other surveys of agent-based electricity market models that we are aware of, such as Shun-Kun/Jia-Hai (2005) who present some selective papers and nicely show how the ACE methodology is ranged within other methods of economics, or Sensfuß et al. (2007) who also take into account models of long-term decision making in the power industry. The researchers who report about the (empirical) validation of their AB electricity market model Macal/North 2005 are one example) proceeded in heterogeneous ways. 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 (2008), 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 from the perspective of their usefulness for electricity modeling purposes. The development of guidelines for assuring the reliability and validity of AB electricity models would greatly benefit the research quality and diminish the heterogeneity of approaches in this field. It is a necessary step for reducing skepticism and ensuring the quality of the methodology.

Windrum et al. (2007) review some empirical validation techniques that could be used to ensure validity of AB models. They first examine how the output of an AB simulation can be analyzed. The model outcome is characterized by a set of statistics that are computed from data generated by the model data generation process. As most processes in AB simulations are stochastic in nature, several simulation runs with varying random number seeds are necessary. By exploring a sufficiently large number of points in the space of initial conditions and parameter values, and by computing a set of statistics at each point, one can gain an understanding of the behavior of the model data generation process.

In order to ameliorate the problem of empirical model validation, LeBaron (2006) makes three suggestions (his focus is on financial markets): The first is to attempt to construct an AB model such that it replicates empirical features which are not well replicated by standard models. The second one is to put as many parameters as possible under evolutionary control in order to find optimal values for crucial parameters. The third suggestion is to use insights gained from experimental economics in order to build more realistic learning dynamics.

² It should be noticed that some researchers argue that AB models are only suitable for qualitative analysis. This would entail that AB models can solely test theories in the form of causal relationships, and calibration would be less an issue. Validation could not be grounded on a comparison with empirical data in this case (see Pyka/Fagiolo 2005, or Windrum et al. 2007 for this discussion). In contrast, most of the latest papers on AB model validation consider empirical validation.

Moss/Edmonds (2005) propose to *cross-validate* AB models. They argue that the micro level of an AB model is best investigated qualitatively while the macro level should be investigated using quantitative methods. When the agents' behavior and interaction on the micro level is able to generate the macro level phenomena sharing the statistical characteristics of data from the real-world system, then the model is cross-validated: the micro level behavior is validated qualitatively by domain experts, and the macro level data is validated by comparing statistical properties of numerical outputs from the model with statistics of the real-world system.

Werker/Brenner (2004) propose an advanced methodology for calibrating and validating AB simulation models based on *Critical Realism*. It proceeds in three steps which involve (i) setting possible model specifications (parameters, interactions), where the assumptions on which the model is built should be induced from empirical data whenever this is possible; a set of several plausible models with certain parameter ranges results from this first step; (ii) running each model specification many times and comparing the tested model specifications to empirical observations, rejecting model specifications that are not confirmed by the empirical data; finally, (iii) identify the underlying mechanisms driving the part of the world that should be described and explained using abduction.

In this paper, we want to advertise for a two-level validation which considers both the micro and the macro level of an agent-based model. *Micro-validation* entails ensuring that the applied learning algorithm (or other behavioral representation) adequately reflects the behavior of the agents. Some papers describing AB electricity market models do not report on a careful micro-validation. In the described models, learning or adaptation is implemented in very heterogeneous ways. The choice of the learning algorithm itself is often not argued and justified; most authors do not answer the question why a specific learning model has been chosen and how good it performs in comparison to alternative behavioral representations. In contrast to this modeling practice, the implications of using specific learning models should be analyzed carefully before simulation outcomes are interpreted on a higher level of abstraction. It should be carefully tested – if necessary with a simplified scenario – whether bidding behavior patterns of the agents in the electricity market correspond to a desirable behavior, or to a behavior that can be expected in real-world markets.

During the *macro-validation* procedure, macro variables that result from the interaction of the agents participating in the modeled markets are analyzed. In electricity market models, these macro-variables are usually market prices that result from the agents' interactions, given the data input that characterizes the supply and demand side, and the market structure and rules. The simulated macro variable values are then compared to empirically observed values in order to verify if simulation outcomes resemble those observed at the real-world electricity markets. Law (2007) proposes several procedures for comparing real-world observations and simulation output data, such as the *basic* or the *correlated inspection approach*, the *confidence-interval approach* or the *time-series approach*. These are also applicable to AB electricity market simulations.

3 Example: An agent-based simulation model of electricity and emissions trading

As an example of an agent-based simulation model that can be applied for economic policy advice, we present here a model that represents the German electricity sector. It comprises three markets, which play an important role in short-term wholesale power trading:

a day-ahead electricity market at which hourly contracts for physical power delivery on the following day are traded, a market for positive minute reserve (balancing power market), where capacity that is held in reserve for regulating imbalances is procured by the transmission system operators, and an exchange for CO₂ emission allowances. The model has been developed within the research project *PowerACE* and represents the part of the implementation that is concerned with short-term wholesale power trading; some concepts of the whole implementation are described in Weidlich et al. (2008).

Following the agent-based paradigm, all relevant parts of the electricity sector simulation model are modeled as *agents*. The set of agents comprises market operators, electricity generators, load serving entities, and CO₂ market participants. Market operators collect supply and demand bids from registered agents and carry out the market clearing. Generator agents operate power plants and sell their generation output either on the day-ahead electricity market or – if their power plants meet the technical requirements for delivering minute reserve – on the balancing power market, and buy or sell CO₂ emission allowances if they operate fossil fuel fired plants. They are characterized by their power plant portfolios, where each plant is defined through several parameters such as its (constant) marginal generation cost, the net installed capacity, no-load costs, or its emission factor, denoting how much CO₂ emissions are associated with every MWh of output. Load serving entities demand electricity on the day-ahead market; they represent a fixed, price-insensitive hourly load. CO₂ market participants other than the generator agents constitute the external demand and supply of allowances. They are characterized by their demand or supply quantities and a valuation at which they like to sell or buy CO₂ emission allowances.

As depicted in Figure 1, all agents inherit basic methods from one abstract super class (*PowerACEAgent*). The three markets are subclasses of the *MarketOperator* class, and

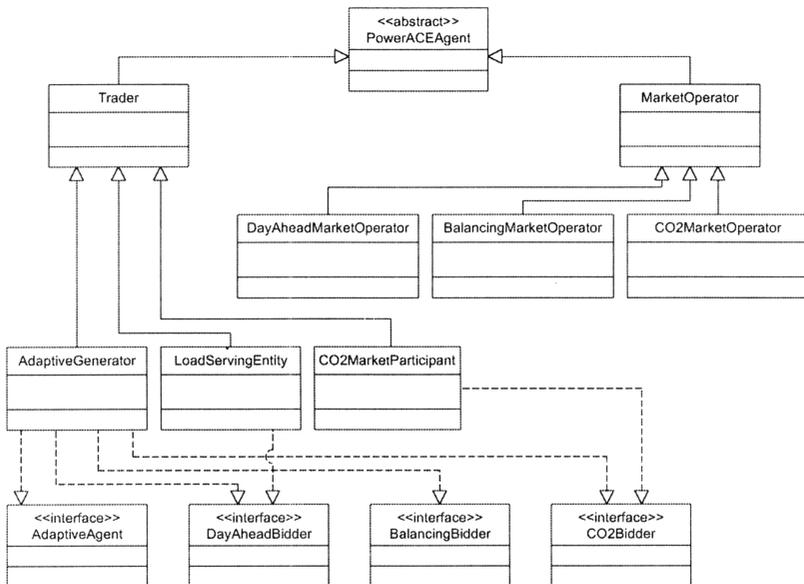


Figure 1 UML class diagram of agents in the simulation model

market participants inherit parts of the Trader agent class. In order to ensure that agents can submit correctly formed bids, they implement the corresponding interface of the market (DayAheadBidder, BalancingBidder, CO2Bidder). Agents of the AdaptiveGenerator class are able to trade on all three markets, and they can act strategically on the two electricity markets. Agents of the LoadServingEntity type do not act strategically.

The model implementation uses the *Recursive Porus Agent Simulation Toolkit* (Repast), which is a JAVA-based class library that facilitates agent-based simulations.³ A presentation of the simulation model is also given in Weidlich/Veit (2008b), in which results from other simulation scenarios are presented. The model description provided in the following reproduces parts of this previous paper and presents results from new simulation runs.

3.1 Simulated markets

The day-ahead market (DAM) is modeled as a sequence of 24 simple call markets for every delivery hour of the following day. Each agent i submits offers for each of its G_i generating units g . As in the simulations presented here the balancing power market is cleared first, the available capacity that an agent can govern in the day-ahead market depends on the trading results on this first market. The capacity an agent has committed on the balancing power market is subtracted from the net installed capacity q^{net} of this plant, resulting in the available capacity q^{avail} .

One supply offer consists of a quantity and the price at which the quantity is offered. The agent formulates the offers according to the output of the reinforcement learning algorithm (see Section 3.2 for a description of agent learning), thereby finding the best actions over a two-dimensional action domain. The price dimension allows offer prices from 0 to 100 EUR/MWh, in 21 discrete steps; the offer quantity that agents can choose ranges from fractions of $\beta = 0$ to 100 % of the available capacity, in six discrete steps. The set of offers submitted to the day-ahead market by agent i for delivery hour h , which contains separate offers for each generating unit g ($b_{h,i,g}^{DAM}$) is denoted $B_{h,i}^{DAM}$ and is defined as follows:

$$B_{h,i}^{DAM} = \left\{ \left\langle p_{h,i,g}^{DAM}, q_{h,i,g}^{DAM} \right\rangle : g = 1, \dots, G_i \right\} \quad (1)$$

$$\text{with } q_{h,i,g}^{DAM} = \beta_{h,i,g} \cdot q_{h,i,g}^{avail}, 0 \leq q_{h,i,g}^{DAM} \leq q_{h,i,g}^{DAM,avail} \quad (2)$$

Market clearing is effected by sorting all supply offers in ascending, and demand bids in descending price order. In case of equal bid prices, bids are sorted in descending volume order; priority between two identical bids is determined randomly.⁴ The intersection of the so-formed supply and demand curves sets the resulting price P_b^{DAM} at which all successful bids are remunerated.

The implemented balancing power market (BPM) represents procurement auctions for positive minute reserve, and is designed in a similar fashion as the minute reserve auctions operated by the four German transmission system operators. Agents with power

³ For a description and download of the toolkit, see <http://repast.sourceforge.net>.

⁴ Before the clearing process, all bids are randomized in order to avoid unintended recurrent priorities in cases of identical bids.

plants that are technically able to deliver minute reserve can sell capacity on this market. The capability of minute reserve delivery is expressed through the indicator function $I(g, MR)$ which is equal to zero if plant g is slowly controllable (not suitable for minute reserve delivery) and one if plant g is fast controllable and can, thus, be offered on the BPM.

An offer on the balancing power market contains a quantity in MW offered over the whole bidding period, and of two prices: the *capacity price*, i.e. the price for holding capacity in reserve during the bidding period, and the *energy price*; i.e. the price a generator is paid for produced minute reserve in case his plant is actually deployed for regulating purposes. Similarly to the day-ahead market, the domain of possible actions on the balancing power market is two-dimensional. While all agents offer their whole available capacity on the balancing power market, the two dimensions contain possible capacity prices and energy prices. Admissible prices range from 0 to 200 EUR/MW in 21 discrete steps for the capacity (*cap*) price and from 0 to 100 EUR/MWh in five steps for the energy price.

In the scenarios presented here, the BPM is carried out first, so the available capacity for all power plants is equal to their respective net capacities. After choosing an action from the learning algorithm, the agent formulates the offer that he wants to submit to the balancing market operator; the specification of a balancing power offer is as follows:

$$B_{k,i}^{BPM} = \left\{ \left\langle p_{k,i,g}^{BPM,cap}, p_{k,i,g}^{BPM,energy}, q_{k,i,g}^{BPM} \right\rangle : g = 1, \dots, G_i | I(g, MR) = 1 \right\} \quad (3)$$

$$\text{with } q_{k,i,g}^{BPM} = q_{k,i,g}^{BPM,avail} \quad (4)$$

The demand side of the balancing power market is represented as a predefined quantity of positive minute reserve that has to be procured. Six equally long bidding blocks of four hours length are differentiated for every trading day: from 0:00 to 4:00 am, from 4:00 to 8:00 am, and so forth. The tendered balancing capacity quantity Q_k^{BPM} is equal for every bidding block (for the simulations presented in this chapter, $Q_k^{BPM} = 3,500 \text{ MW } \forall k$).

During market clearing, offers are sorted by their capacity price; offers with the lowest capacity price are considered first for holding capacity in reserve, until demand is met. If minute reserve is actually needed for frequency control on the day of delivery, deployment is decided among the offers of those plants that are held in reserve on the trading day. A merit order is then constructed based on the energy price only, and offers with lowest energy prices are deployed first, until demand is met. Payments both for holding capacity in reserve and for delivering energy for regulating purposes are the respective offer prices.

All generator agents that own fossil fuel fired power plants are initially endowed with a certain amount of CO₂ allowances. In accordance with current emissions trading regulation, the initial allocation of allowances is calculated according to the grandfathering rule, i.e. based on past emissions for each single power plant. Agents can trade their need or surplus of allowances on the allowance market.

The CO₂ emission allowance market (CO₂M) is executed at the end of each trading day. It is assumed that all agents seek to even up their open positions every day. This entails that agents who sell electricity also make sure to have enough allowances for the carbon dioxide emissions associated to their generation output. Agents who have surplus allow-

ances try to sell them at the market price. The option of speculating for rising or falling allowance prices is neglected in this model.

The agent's daily trading quantities are calculated on the basis of initial endowments $q_y^{CO_2,init}$ issued for year y , and of trading success on the current trading day. The amount of carbon dioxide emitted during electricity generation, q^{emit} , is determined by volumes sold at the day-ahead electricity market $q^{DAM,sold}$ and by deployed minute reserve $q^{BPM,depl}$.⁵ The quantities are multiplied with the emission factor ω_g of plant g , quantifying the CO₂ emissions associated with every MWh of power output.

$$q_{t,i}^{CO_2,emit} = \sum_{G_i} \left(\sum_{b=1}^{24} \omega_g \cdot q_{b,i,g}^{DAM,sold} + \sum_{k=1}^6 \omega_g \cdot q_{k,i,g}^{BPM,depl} \right) \quad (5)$$

The remaining allowance budget that an agent has at its disposal at time t (day of the year y) is divided by the remaining days for which the allowances were issued, in order to calculate a daily budget. This budget is subtracted from the allowance quantity needed for power generation, thus resulting in the bid/offer quantity that agent i submits to the CO₂ market operator. In consequence, if an agent's budget for the current day is larger than the need for allowances, his bid quantity becomes negative, which corresponds to a selling offer.

$$q_{t,i}^{CO_2M} = q^{CO_2,emit_{t,i}} - \frac{q_{t,i}^{CO_2,bud}}{365 - t - 1} \quad (6)$$

Bids on the CO₂ allowance market contain a volume of allowances that is offered or asked, a bid price, and the compliance period (cp) for which the allowance should be valid (for simplicity, and because no speculation is considered, the compliance period is always equal to the current period). Buying bids have positive volumes, and selling offers have negative volumes. All agents submit one single daily bid on the allowance market, representing their requirement or surplus calculated over all power plants they own.

$$b_{t,i}^{CO_2M} = \langle p_{t,i}^{CO_2M}, q_{t,i}^{CO_2M}, cp \rangle \quad (7)$$

The CO₂ emission allowance market is modeled as a sealed bid double auction. Demand and supply bids are summed up to form supply and demand functions in the same way as on the day-ahead electricity market, and the uniform market clearing price is determined by the intersection of both curves.

The remaining allowance budget is updated at the end of each trading day by subtracting the amount of allowances used for emitted CO₂ quantities from the current budget and by adding the resulting trading volumes (positive for bought allowances, negative for sold volumes) to it. The budget at time $t = 0$ is the initial allowance quantity issued for year y through the grandfathering procedure.

⁵ If the deployment of minute reserve becomes necessary, this would occur one day after trading on the balancing power market. For simplicity, it is assumed that minute reserve quantities actually deployed are already known on the trading day. In real-world practice, only around 2 % of procured minute reserve capacity is actually deployed for frequency regulation (Bundesnetzagentur 2006) so these quantities are negligible.

Agents do not act strategically on the market for CO₂ emission allowances – they do not set bidding strategies through reinforcement learning. However, the costs incurred from allowance prices influence trading strategies on the electricity markets, as specified in the following section.

3.2 Agent learning and market interrelations

For each market in which they act strategically, agents choose new actions from instances of a reinforcement learning algorithm, and then formulate their bids according to this chosen action. Agents learn strategies separately for the day-ahead and for the balancing power market. Moreover, strategies for each bidding block on the balancing power market and for each hour on the day-ahead market are learned separately.

While optimizing their supply offers, agents consider opportunity costs that they could have achieved on the other market if they had sold their capacity there. Prices for CO₂ emission allowances are also included into the reinforcement as opportunity costs, because a generator would always have the opportunity to solely sell certificates, thereby realizing a profit.

The three markets that form the electricity sector simulation model are interrelated through the agents' bidding strategies. A power generator has the choice to offer his generating capacity on the DAM or on the BPM for the following day (for those plants that fulfill the technical requirements to deliver minute reserve), and has to trade off between these two options. After market clearing on the first market, an agent can offer his remaining unsold capacity on the second market. Through varying the offer quantity on the day-ahead market, agents can influence and optimize their joint strategy on both power markets.

The behavioral representation of the agent's search for profit maximizing strategies is modeled with Q-learning (Watkins 1989) and variants of the Erev/Roth (1998) algorithm. For this learning algorithm, different *states* have to be defined. In the simulations applying Q-learning, the *states* are based on bid prices and trading success. Bid prices are categorized as low (lower than or equal to one third of the maximum admissible bid price), high (higher than or equal to two thirds of the maximum admissible price) or medium (all remaining prices). A bid is further categorized as marginal or intra-marginal, in which case it is a successful bid, or as extra-marginal for a bid that was not successful. All combinations of bid prices and success form the six states that are differentiated in this model.

Reinforcements R that are fed back to the update function of the learning algorithm are calculated for both power markets after all markets have cleared on the current trading day. They are based on the profit π earned on the respective market. Reinforcements are set relative to the maximum possible profit π^{\max} that an agent can earn on the market. Consequently, the range of possible reinforcements $0 \leq R \leq 1$ is the same for all learning tasks; initial Q-values or propensities are set to $Q_0 = 1$. Maximum profits depend on variable costs c^{var} of the power plant deployed, and are different across plants.

Agents that own a portfolio of generating units set and learn offers separately for each plant. The reinforcement for learning the bidding strategy of one plant, however, contains some information about the performance of the whole portfolio. The influence of the whole portfolio profit on the reinforcement is set through the *portfolio integration* parameter ψ . Throughout the simulations presented here, it is set to $\psi = 0.5$. Reinforcements on the day-ahead market are defined as follows:

$$R_{b,i,g}^{DAM} = (1 - \psi) \cdot \pi_{b,i,g}^{DAM} / \pi_{i,g}^{DAM,max} + \psi \cdot \frac{\sum_{G_i} \pi_{b,i,g}^{DAM} / \pi_{i,g}^{DAM,max}}{G_i} \quad (8)$$

$$\text{with } \pi_{b,i,g}^{DAM} = q_{b,i,g}^{DAM,sold} \cdot P_b^{DAM} - c^{var} - c_{b,i,g}^{BPM,opp} - c_{b,i,g}^{CO2M,opp} \quad (9)$$

$$\pi_{i,g}^{DAM,max} = q_{i,g}^{net} \cdot P^{DAM,max} - c^{var} \quad (10)$$

$$c_{b,i,g}^{BPM,opp} = I(g, MR) \cdot P_{k(b)}^{BPM,cap} \cdot q_{b,i,g}^{DAM} \quad (11)$$

$$c_{b,i,g}^{CO2M,opp} = q_{b,i,g}^{CO2,emit} \cdot P_{t(b)}^{CO2M} \quad (12)$$

The reinforcement for trading on the balancing power market is defined in a similar manner as on the day-ahead market. If no minute reserve energy is actually delivered, profits are only defined by capacity prices and by *no load costs* (*nolc*), i.e. the cost for keeping the power plant in a stand-by state. For the calculation of $\pi^{BPM,max}$, only capacity prices are considered.

$$R_{k,i,g}^{BPM} = (1 - \psi) \cdot \pi_{k,i,g}^{BPM} / \pi_{i,g}^{BPM,max} + \psi \cdot \frac{\sum_{G_i} I(g, MR) \cdot \pi_{k,i,g}^{BPM} / \pi_{i,g}^{BPM,max}}{\sum_{G_i} I(g, MR)} \quad (13)$$

$$\text{with } \pi_{k,i,g}^{BPM} = q_{k,i,g}^{BPM,cap,sold} \cdot P_k^{BPM,cap} - c^{nolc} \cdot |h \in k| - c_{k,i,g}^{DAM,opp} \quad (14)$$

$$\pi_{i,g}^{BPM,max} = q_{i,g}^{net} \cdot P^{BPM,cap,max} - c^{nolc} \cdot |h \in k| \quad (15)$$

$$c_{k,i,g}^{DAM,opp} = q_{k,i,g}^{BPM,cap} \cdot \sum_{b \in k} P_b^{DAM} \quad (16)$$

3.3 Validation

As specified in Section 2.3, the micro-validation procedure requires a detailed analysis of which behavioral representation is best suited for modeling agents engaged in daily repeated electricity trading. In this context, tests of different learning algorithms have been carried out by the authors. This micro-validation procedure will not be reported here. However, we want to advertise one measure to enhance validity of results: if several learning models are deemed appropriate for representing agent behavior, comparative runs applying different learning models should be conducted, so that simulation results from one learning model are confirmed by runs applying alternative learning models.

From the micro-validation procedure we found that Q-learning and several variants of the learning model formulated by Erev/Roth (1998) are suitable to model learning in the developed agent-based simulation model of the German electricity sector. The Erev and Roth algorithm has proved to adequately represent empirically observed human learning in several kinds of games; Q-learning is used in several ACE papers and offers desirable properties from a machine learning perspective (convergence to optimality under certain conditions). In the following, results from Q-learning simulations are presented primarily. A more detailed comparison of the different learning algorithms applied in this model is given in Weidlich (2008).

On the basis of a properly micro-validated learning model, the next validation step can proceed. The macro-validation procedure entails the analysis of macro variables that result from the interaction of the agents participating in the modeled markets. Here, market prices that come out of simulations with the electricity sector model, run with data input that realistically represents the German electricity sector, are compared to empirically observed prices in the year 2006. The overall power plant structure is represented in an aggregate way, based on published data about installed capacities. The four dominant players in the German power market are represented in more detail, and further players are introduced so that the overall installed capacity and the proportions of different power plant technologies (coal-fired, gas-fired, hydro etc.) are properly represented. The merit order of the total installed capacity by all agents is depicted in the appendix (Figure 9).

The system's total load data provided by the *Union for the Co-Ordination of Transmission of Electricity* (UCTE) constitutes the data input for the demand side. Simulation results are compared to empirically observed prices at the European Energy Exchange (EEX) spot market and at the German balancing power markets for those days for which the system's total load is known from UCTE data (i.e. every third Wednesday of a month). In addition, average daily courses of prices over all workdays of the same month are calculated and compared to the simulation outcomes. The load profiles applied in the simulations presented in this paper are graphically displayed in the appendix (Figure 10).

Simulations that apply Q-learning as the agents' behavioral representation are run with a learning rate of $\alpha = 0.5$, a discount rate of $\gamma = 0.9$, and an ϵ -greedy action selection rule with $\epsilon = 0.2$. Other simulations apply the original Erev and Roth reinforcement learning algorithm or its modification as formulated by Nicolaisen et al. (2001), both with proportional action selection, a recency value of $\phi = 0.1$, and an experimentation value of $\epsilon = 0.2$. Simulations ran over 7,300 iterations. The outcome of one simulation run is defined as the average price over the last 365 iterations. Due to the stochastic nature of reinforcement learning, simulations are repeated ten times with different random number seeds at each run. The outcome of one simulation scenario is represented by the average over the ten repetitions for this scenario.

Figures 2 and 3 show examples for simulation results. Continuous lines plot the simulation outcome for the third Wednesdays of the simulated month; dashed lines plot the empirically observed prices of the same days, and chain dotted lines represent average

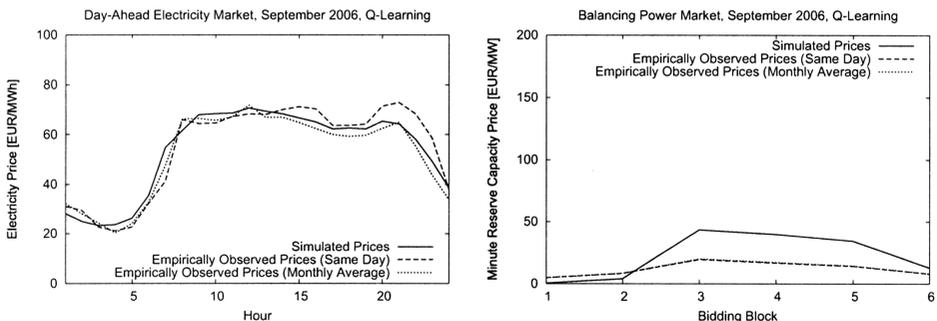


Figure 2 Example results for September 2006: simulated and real-world prices on the day-ahead (left) and balancing power market (right)

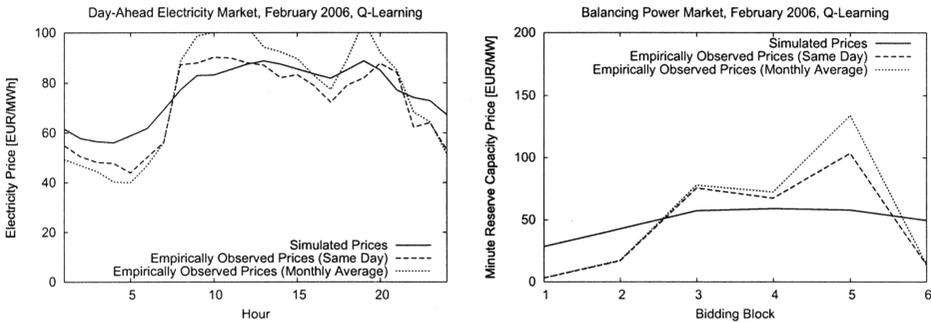


Figure 3 Example results for February 2006: simulated and real-world prices on the day-ahead (left) and balancing power market (right)

prices over all workdays of the specific month. Left figures display hourly results on the day-ahead market, where empirically observed prices correspond to prices for hourly contracts fixed in the daily EEX spot auction. Right figures show results from the simulated balancing power market for the six bidding blocks, and the empirically observed prices are the averages over the prices published by the four balancing market operators. Prices in summer, spring and fall months (see example in Figure 2) are quite closely reproduced by the simulation model. Prices in winter months (see example in Figure 3), however, show larger deviation between simulation outcomes and real-world prices. In some hours, prices even exceed 100 EUR/MWh; as the agents' action domain in the simulation model only allows for bid prices up to 100 EUR/MWh, prices above this level cannot be replicated by construction of the model. Enlarging the action domain might be a way to ensure that prices higher than 100 EUR/MWh can also be simulated. However, a sensitivity analysis reveals that simulation outcomes are quite sensitive towards the definition of the agents' action domain. If the range of possible bid prices on the day-ahead market is augmented to 110, 120 or 150 EUR/MWh, average yearly prices rise by 6.2, 12.5, and 31.3 %, respectively.

Variability between different runs (i.e. runs with different random number seeds) is very low for simulations with Q-learning. The standard deviation for the resulting prices of the ten repetitions ranges between 0.2 and 2.3 EUR/MWh for different hours on the day-ahead electricity market and between 0.05 and 3.9 EUR/MW for bidding blocks on the balancing power market. With these low variances, one single simulation run already delivers meaningful and reliable results. Variances for simulations applying Erev and Roth reinforcement learning are considerably higher. The standard deviations of the single observations are between 3.3 and 13.0 EUR/MWh on the day-ahead and 0.0 and 29.0 EUR/MW on the balancing power market when applying the original algorithm, and 3.4–17.9 EUR/MWh (DAM) and 0.0–24.5 EUR/MW (BPM) for simulations with the modified Erev and Roth algorithm. Thus, simulations applying Erev and Roth reinforcement learning should be repeated more often.

The simulated prices observed on the day-ahead market and on the balancing power market in one month stem from the same simulation run, and are a consequence of agents bidding subsequently on both markets (and in addition on the market for CO₂ emission allowances) and optimizing their strategies in face of market interrelations. The demand on the balancing power market, i.e. the tendered minute reserve quantity, is equal for all

bidding blocks. This market is cleared first, and the day-ahead market is operated subsequently. As the available supply capacity and the demand quantity in the balancing power market is the same in every hour, differences in prices between the bidding blocks can only result from the inclusion of opportunity costs in the agent's reasoning. The simulation outcome on the balancing power market shows characteristic daily courses of prices, with higher capacity prices in the third and fourth bidding block – and fifth in winter months – and lower prices in nocturnal bidding blocks. Similar characteristics can be observed in the real-world balancing power markets in Germany.

The confidence-interval approach as described in Law (2007) has been applied to the outcomes of this basic simulation scenario; a description of the methodology and the calculations for the reference scenario are described in Weidlich (2008). Results confirm that simulated prices are not significantly different from real-world prices of the same day. This is observed for Q-learning with ϵ -greedy action selection, and both the original and modified variant of the Erev and Roth reinforcement learning algorithm with proportional action selection. Thus, the developed model is deemed an appropriate and valid representation of the real-world system under study.

3.4 Measuring the impact of emissions trading

The data presented in the preceding section corresponds to simulations in which emission allowance trading was integrated – just like in the real-world market of the corresponding time frame. In further simulation runs, it is tested how emissions trading affects prices on the electricity markets. For this purpose, scenarios without CO₂ emissions trading are run and compared to the basic scenario results. The outcome of this comparison is depicted in Figures 4 and 5 for the day-ahead and balancing power market, respectively. As prices on both electricity markets are strongly influenced by the system's total load, simulated prices are sorted by demand quantities in the corresponding hours (DAM), or

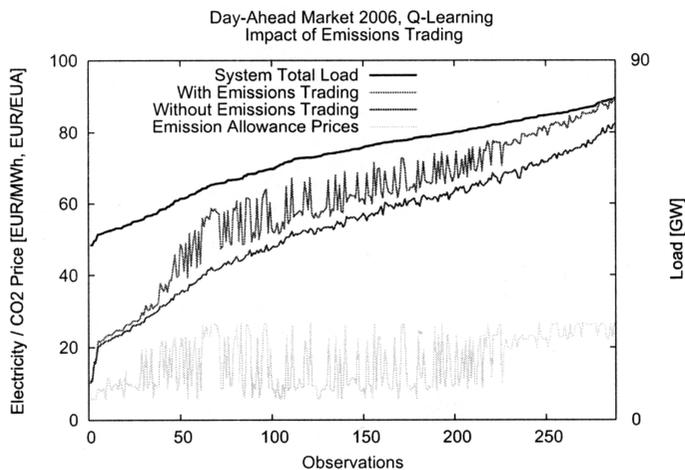


Figure 4 Impact of CO₂ emissions trading on day-ahead electricity prices

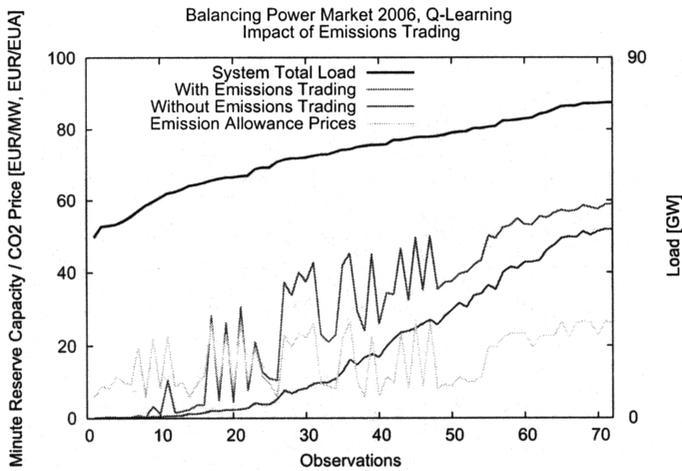


Figure 5 Impact of CO₂ emissions trading on minute reserve capacity prices

by average demand over the four hours constituting a bidding bloc (BPM). Demand is plotted at the second ordinate.

It can be seen that a large fraction of opportunity costs resulting from the possibility of selling CO₂ emission allowances is successfully passed over to electricity market offers, which ultimately raises prices at the day-ahead market and also at the balancing power market. Because of different emission situations in the single hours, the absolute increase in electricity prices is not constant across the observations. In hours of low demand, the introduction of emissions trading has hardly any effect on day-ahead electricity prices, whereas the difference in prices is considerable in high demand hours. Over a large range of intermediate demand situations, deviations between the scenarios with and without emissions trading fluctuate to some extent.

The intuition behind this result is that these hours with similar demand situations belong to different months, and CO₂ prices differ across months. Hours with very high demand all belong to the winter months in which demand is high and consequently many fossil fuel power plants are operated, resulting in (evenly) higher CO₂ allowance prices. As a consequence, it can be concluded that emissions trading considerably influences electricity prices and that it is the main cause for differences in prices resulting for hours with similar demand situations. This characteristic is observed on both the day-ahead and the balancing power market. Yearly average prices are 13.3 % higher for scenarios with emissions trading on the day-ahead market, and 56.8 % higher on the balancing power market (for simulations with Q-learning).

3.5 Varying the settlement rule at the balancing power market

In procurement auctions with two-part bids, the market clearing procedure distinguishes two aspects: the *scoring rule* and the *settlement rule* (Chao/Wilson 2002). The scoring rule defines how to compare bids, and the settlement rule determines payments. Basically, two options for payments of successful bids are discussed in the literature and em-

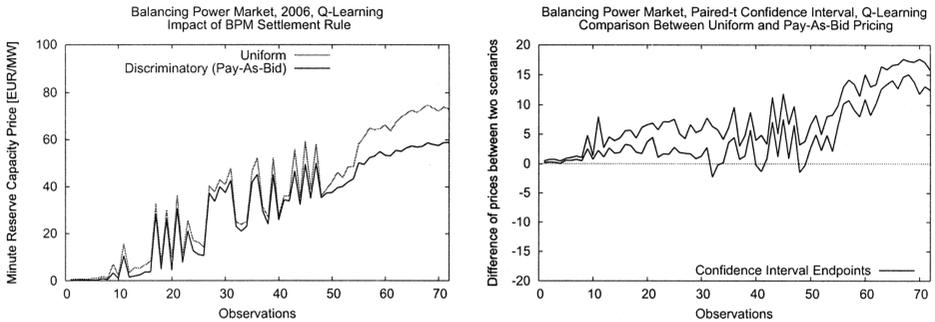


Figure 6 Impact of BPM settlement rule on minute reserve capacity prices; average prices (left) and confidence intervals (right)

ployed in practice: *uniform price* and discriminatory *pay-as-bid* settlement. In the former case, one market clearing price is determined – usually the price of the highest successful bid – which is paid to all successful bidders, whereas in the latter case, which is also applied in the model presented here, every successful bid is paid its bid price. In European balancing power markets, both settlement rules are employed. In Germany, Austria and the United Kingdom for example, pay-as-bid pricing is applied, whereas Sweden, Norway, Spain and the Netherlands opted for uniform pricing (Swider 2006).

The question which settlement rule is less vulnerable to strategic bidding is discussed in many research papers and leads to mixed results, e.g. Ausubel/Cramton (2002), Binmore/Swierzbinski (2000), Kahn et al. (2001). Also, different publications about agent-based electricity market simulations pursue comparisons between pay-as-bid and uniform pricing, e.g. Bakirtzis/Tellidou (2006), Bin et al. (2004), Bower/Bunn (2001), Cincotti et al. (2006), Weidlich/Veit (2006), Xiong et al. (2004). Most agent-based simulations come to the conclusion that the pay-as-bid mechanism leads agents to bid at higher prices, but overall average procurement costs are still lower under this mechanism than under uniform pricing.

Figure 6 shows a comparison of minute reserve prices for both settlement rules applied in our model. Results reveal that overall market prices are higher in the uniform pricing case than under pay-as-bid pricing. From Figure 6 (left) it becomes observable that price differences between minute reserve procurement auctions with uniform pricing and with pay-as-bid pricing are notable especially in the bidding blocs with high system load, with absolute differences around 14–16 EUR/MW (note that prices in this plot rank from 0 to 100 EUR/MW).

The application of the paired-*t* confidence interval method⁶ Law (2007) at a 90 % level confirms that resulting prices under uniform pricing are statistically significantly higher than under discriminatory pricing, with only some exceptions; the corresponding lower and upper confidence interval endpoints are depicted in Figure 6, for Q-learning simulations).

⁶ This method constructs a confidence interval around the differences between resulting prices of two scenarios in order to quantify how strong scenario outcomes deviate from each other. For all scenarios, the same set of random number seeds is used and the outcome from a run of one scenario is paired with the run of the other scenario that uses the same random number seed.

The yearly average of minute reserve prices in Q-learning simulations is 30.73 EUR/MW with pay-as-bid pricing, and 36.72 EUR/MW (+ 19.5 %) in the uniform pricing case. If Erev and Roth reinforcement learning is applied, results are slightly less definite, but generally confirm the findings from Q-learning simulations nonetheless. While with Q-learning, uniform prices are higher than average pay-as-bid prices in every single bidding bloc, simulations with Erev and Roth RL show some bidding blocs in which the respective average pay-as-bid price is higher than the uniform price. However, yearly average prices over all bidding blocs are also higher with uniform pricing (+ 15.15 % if the original Erev and Roth algorithm is applied, and + 28.9 % in simulations with the modified algorithm).

When comparing prices on the day-ahead market for the tested two settings with different settlement rules applied on the balancing power market, only very small differences can be observed. Yearly average prices are less than 1 % higher in the uniform pricing case than with pay-as-bid (Q-learning simulations). However, this tendency is only confirmed by simulations with the original Erev and Roth RL, while those applying the modified algorithm result in negligibly higher prices for pay-as-bid settlement. So, no significant difference of prices on the day-ahead market is found for the two tested scenarios. This result allows another conclusion, i.e. that opportunity costs resulting from the balancing power market hardly affect day-ahead prices, while, vice versa, balancing power prices are strongly influenced by resulting day-ahead prices. This observation is consistent with experts' assessments about (wholesale) market interrelations in the German electricity sector.

3.6 Assessing the implications of plant divestiture measures

In her 16th Main Report, the German Monopolies Commission (2006) “regards supervision of competition in the wholesale electricity and regulated energy markets as inadequate. The need for special supervision of competition on these markets is due to their particular vulnerability to supply strategies by generating companies that have sufficient market power to influence prices.” Generators with market power are usually understood as firms who can set prices above marginal generation costs while still making positive sales (e.g. Rassenti et al. 2003).

In the German electricity sector, the four largest generating companies taken together own 70–80 % of the total German generating capacity. In public debate, it is often presumed that these companies have market power and that they contribute to driving prices away from competitive levels. Under the assumption that the oligopolistic structure of the German electricity sector leads to excessive wholesale prices, it would be an adequate policy measure to increase supply side competition through forcing the dominant companies to divest parts of their generation assets. Similar measures have been carried out in the electricity industry of England and Wales.

The question to what extent the four large generating companies actually dispose of market power in the German electricity markets should not be discussed here. For the following analysis, it is assumed that they have some potential to exert market power, and it is tested whether divestiture measures can mitigate this potential to some extent. If results show that plant divestiture has no influence on market prices, this may either invalidate the assumption, or it may indicate that the tested divestiture options are not appropriate to mitigate the players' market power exertion.

In order to assess the impact of plant divestiture in the German electricity industry, six divestiture scenarios have been defined and run with the present model. Results are compared to the reference scenario without divestiture. In all of the tested divestiture scenarios, it is assumed that the largest four agents defined in the model, representing the four dominant players of the real-world industry, have to divest a given fraction of installed capacity of all their power plant types. New agents are then introduced into the market who each take over the whole or part of the divested capacity of one of the four large firms.

Three different fractions of divested capacity – 25 %, 33 % and 50 % – and two numbers of new agents – four and eight – are considered; the combination of these parameter values make up the six simulation scenarios. In the case of four new agents, they are allocated the whole divested capacity of one of the former large firms; in the case of eight new agents, two agents share the divested capacity of one large firm. Consequently, each new agent has a portfolio of power plants with the same proportion of technologies (gas-fired, coal-fired, hydro etc.) as the firm that had to divest part of its capacity. The new agents are also endowed with the share of emission allowances associated to the divested power plants.

Figures 7 and 8 depict resulting prices for the divestiture scenarios. In Figure 7, yearly average prices and average prices for two selected months are plotted; Figure 8 plots

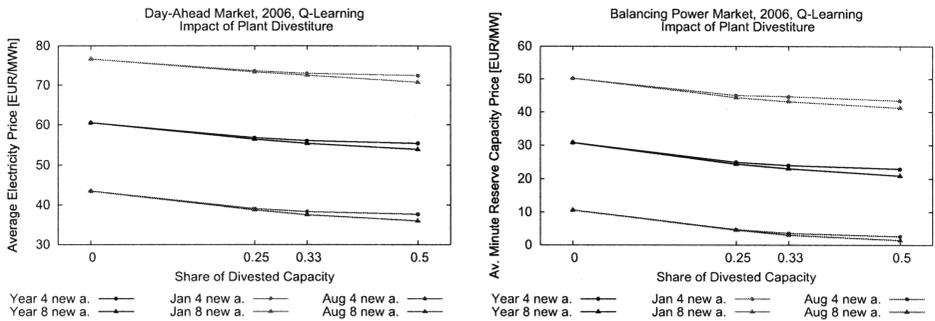


Figure 7 Monthly/yearly average prices on the day-ahead (left) and balancing power market (right) for different divestiture scenarios

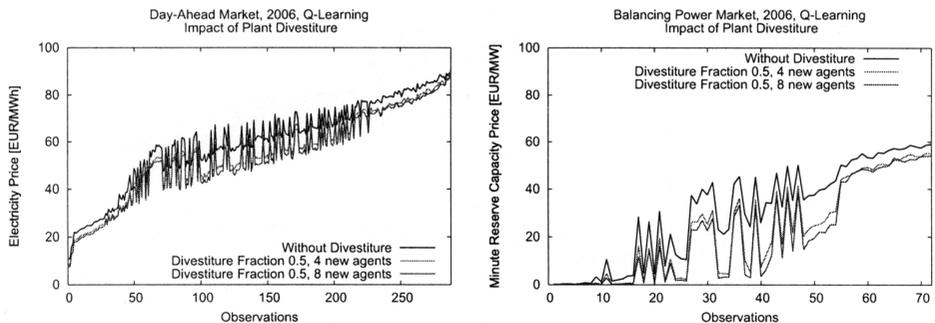


Figure 8 Impact of plant divestiture on prices on the day-ahead (left) and balancing power market (right)

prices for all observations from the two scenarios of 50 % capacity divestiture with four and eight newly introduced agents.

The simulation outputs suggest that divestiture leads, in fact, to falling market prices. Yearly average prices on the day-ahead market are reduced by 6.1 %, 7.4 % and 8.4 % for the scenarios with four new agents, and by 6.7 %, 8.5 % and 10.9 % for the scenarios with eight new agents, all compared to corresponding prices in the scenario without divestiture. On the balancing power market, the decline in prices lies between 19.1 % and 32.3 % for the six scenarios (for Q-learning simulations).

Market prices in all divestiture scenarios are lower than prices in the case without divestiture in a statistically significant manner, according to results from the paired-*t* confidence interval approach (at a 90 % confidence level, for Q-learning simulations). The very few (< 10) exceptions among the 288 observations for each setting on the day-ahead market, in which differences in prices from the scenario with and without divestiture are not statistically significant, are the hours of highest system load. On the balancing power market, prices in all observations are statistically significantly lower in simulations with divestiture than in those without divestiture.

It can be seen that average prices are lower in scenarios of higher amounts of divested capacity than they are in the scenarios with lower divestiture amounts. However, these differences in prices are only statistically significant for some observations. Moreover, the difference in resulting prices for divestiture scenarios with four and with eight new agents that enter into the market is only small, especially if a low share of capacity from the large generators is divested. The difference in resulting prices in the case of four and eight new players becomes more important when the share of divested capacity rises to 50 %. For all scenarios, differences in prices of scenarios with four new agents are only statistically significant for some observations.

Simulations applying Erev and Roth reinforcement learning generally confirm the Q-learning results presented here. Decreases in prices in the divestiture scenarios are between 5.0 % and 9.4 % on the day-ahead market and between 21.0 % and 34.0 % on the balancing power market, both for simulations applying the original Erev and Roth RL simulations; for simulations with the modified variant of Erev and Roth RL, price decreases lie between 5.7 % and 9.7 % (DAM) and between 11.0 % and 20.1 % (BPM), all compared to the reference scenario without divestiture.

4 Conclusion and outlook

In this contribution, the methodology of agent-based electricity market modeling for regulatory advice has been discussed and an example model developed by the authors has been presented. The model provides insights into the aggregate system characteristics resulting from dynamic strategic behavior in oligopolies under consideration of learning from daily interaction. It might help a regulator to prevent market power exertion through introducing suitable market designs or enforcing changes in the market structure.

Results from a realistic simulation scenario representing the German electricity sector have been assessed and the impact that emissions trading, different settlement rules on the balancing power market, and plant divestiture measures have on electricity prices has been analyzed through the application of well-defined simulation scenarios. It has been shown that all analyzed measures have a significant impact on electricity prices, and

that these effects can be quantified with the help of agent-based simulations. On the balancing power market, lower average prices are observed under pay-as-bid pricing than under uniform pricing. Changing the ownership structure of the generating capacities through divestiture leads to decreasing market prices on both the day-ahead and the balancing power market. However, the tested divestiture scenarios only entail rather small decreases in market prices. Given that a plant divestiture measure of the dimension implied by the six tested scenarios is presumably very difficult to enforce from a political perspective, the resulting effects on power market prices may not justify such drastic changes in the market structure.

The presented model may be enhanced in several ways so as to account for more aspects that are relevant in real-world power industries. For example, as the German electricity system is highly intermeshed with those of the neighboring countries, the national perspective followed here is a simplification that may be relaxed in favor of a European view in future work. Also, the current version of the model neglects transmission constraints; a more realistic transmission system representation should be embraced into future versions of the model.

Given the great flexibility that AB modeling allows, and given its potential to represent complex economic systems, ACE can surely become one important pillar of electricity market research. While the AB electricity simulation research field is constantly maturing, we still observe a large heterogeneity in simulation approaches and model descriptions. Choosing valid learning models that represent boundedly rational actors in electricity markets is one issue that needs further research and consolidation; the establishment of reliable and accepted validation procedures is a second important field of further research.

Another observation of the research field is that the agents' profits and the success of distinct trading strategies in power markets are rarely discussed. The focus of investigation is often placed on aggregate macro-level parameters. Moreover, aspects like bilateral trading or vertical integration could realistically be modeled in agent-based simulations, but have not yet been extensively analyzed by agent-based modelers. Future models might put a stronger focus on these aspects.

Although some issues of the agent-based methodology are still "under development", AB modelers have already been able to successfully simulate the core characteristics of today's electricity markets. Adaptive, self-seeking agents are an appropriate representation of strategic generating companies interacting in power trading activities. Some suggestions for questions that might be interesting to investigate with agent-based simulation approaches have been pointed out for future research. As for the encouraging potential of agent-based models for economic policy advice in the electricity sector, the proposed new applications and a further enhancement of the methodology are strongly encouraged.

A. Appendix

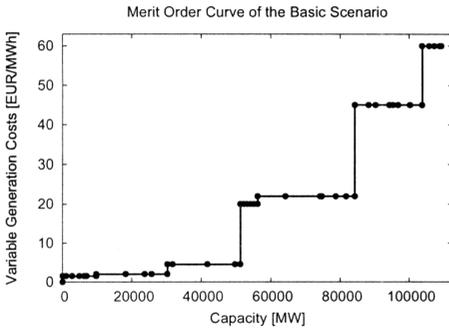


Figure A1 Merit order of installed generating capacity

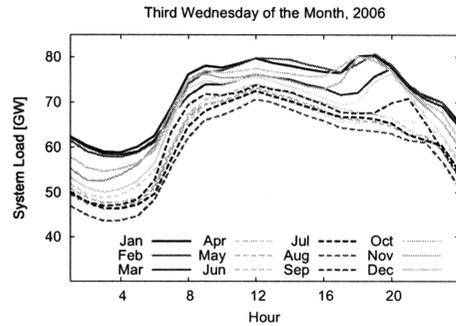


Figure A2 UCTE load for 2006

References

- Ausubel, L.M., P. Cramton (2002), Demand Reduction and Inefficiency in Multi-Unit Auctions. Papers of Peter Cramton from University of Maryland, Department of Economics.
- Axtell, R. (2000), Why Agents? On the Varied Motivations for Agent Computing in the Social Sciences. The Brookings Institution, Center on Social and Economic Dynamics, Working Paper No. 17.
- Bagnall, A., G. Smith (2005), A Multi-Agent Model of the UK Market in Electricity Generation. *IEEE Transactions on Evolutionary Computation* 9: 522–536.
- Bakirtzis, A.G., A.C. Tellidou (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, 1168–1173.
- Batten, D. (2000), *Discovering Artificial Economics – How Agents Learn and Economies Evolve*. Westview Press.
- Bin, Z., Y. Maosong, X. Xianya (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.
- Binmore, K., J. Swierzbinski (2000), Treasury auctions: Uniform or discriminatory? *Review of Economic Design* 5: 387–410.
- Bower, J., D. Bunn (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.
- Bundesnetzagentur (2006), Adjudication about the procedure of minute reserve procurement in Germany. 29. August 2006, Bonn.
- Cau, T.D.H., E.J. Anderson (2003), A co-evolutionary approach to the tacit collusion of generators in oligopolistic electricity markets: piecewise linear bidding structure case. Proceedings of Congress on Evolutionary Computation (CEC2003), Volume 4: 2306–2313.
- Chao, H.-P., R. Wilson (2002), Multi-Dimensional Procurement Auctions for Power Reserves: Robust Incentive-Compatible Scoring and Settlement Rules. *Journal of Regulatory Economics* 22: 161–183.
- Cincotti, S., E. Guerci, S. Ivaldi, M. Raberto (2006), Discriminatory versus uniform electricity auctions in a duopolistic competition scenario with learning agents. IEEE Congress on Evolutionary Computation, Vancouver, Canada.

- Erev, I., A.E. Roth (1998), Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed-Strategy Equilibria. *American Economic Review* 88: 848–881.
- Fagiolo, G., C. Birchenhall, P. Windrum (2007), Empirical Validation in Agent-based Models: Introduction to the Special Issue. *Computational Economics* 30: 189–194.
- Gilbert, N., K.G. Troitzsch (2005), *Simulation for the Social Scientist* (2nd ed.). Open University Press, McGraw-Hill.
- Gulyás, L. (2002), On the Transition to Agent-Based Modeling: Implementation Strategies From Variables to Agents. *Social Science Computer Review* 20: 389–399.
- Kahn, A.E., P.C. Cramton, R.H. Porter, R.D. Tabors (2001). Uniform Pricing or Pay-as-Bid Pricing: A Dilemma for California and Beyond. *The Electricity Journal* 14: 70–79.
- Law, A.M. (2007), *Simulation Modeling & Analysis* (4th ed.). McGraw-Hill Higher Education.
- LeBaron, B. (2006), Agent-Based Computational Finance. Pp. 1187–1233 in: L. Tesfatsion, K.L. Judd (eds.), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*. North-Holland.
- Macal, C.M., M.J. North (2005), Validation of an Agent-Based Model of Deregulated Electric Power Markets. Forschungsbericht, Los Alamos National Laboratory, presented at the North American Association for Computational and Social Organization (NAACSOS) Conference.
- Marks, R.E. (2007), Validating Simulation Models: A General Framework and Four Applied Examples. *Computational Economics* 30: 265–290.
- Midgley, D.F., R.E. Marks, D. Kunchamwar (2007), The Building and Assurance of Agent-Based Models: An Example and Challenge to the Field. *Journal of Business Research* 60: 884–893.
- Monopolies Commission (2006), More Competition in the Services Sector As Well. The Sixteenth Biennial Report 2004/2005, abbreviated version, Bonn.
- Moss, S., B. Edmonds (2005), Sociology and Simulation: Statistical and Qualitative Cross-Validation. *American Journal of Sociology* 110: 1095–1131.
- Nicolaisen, J., V. Petrov, L. Tesfatsion (2001), Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation* 5: 504–523.
- Pyka, A., G. Fagiolo (2005), Agent-Based Modelling: A Methodology for Neo-Schumpeterian Economics. Discussion Paper Series 272, Universität Augsburg, Institute for Economics. available at <http://ideas.repec.org/p/aug/augsbe/0272.html>.
- Rassenti, S.J., L.V. Smith, B.J. Wilson (2003), Controlling market power and price spikes in electricity networks: Demand-side bidding. *Proceedings of the National Academy of Sciences of the United States of America* 100: 2998–3003.
- Richiardi, M., R. Leombruni, N. Saam, M. Sonnessa (2006), A Common Protocol for Agent-Based Social Simulation. *Journal of Artificial Societies and Social Simulation* 9 (1).
- Roth, A.E. (2002), The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics. *Econometrica* 70: 1341–1378.
- Rupérez Micola, A., A. Banal Estañol, D.W. Bunn (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*.
- Sensfuß, F., M. Ragwitz, M. Genoese, D. Möst (2007), Agent-Based Simulation of Electricity Markets – A Literature Review. *Energy Studies Review* 15-2: 19–47.
- Shun-Kun, Y., Y. Jia-Hai (2005), Agent-based Computational Economics: Methodology and Its Application in Electricity Market Research. *Proceedings of the 7th International Power Engineering Conference*.
- Sun, J., L. Tesfatsion (2007), Dynamic Testing of Wholesale Power Market Designs: An Open-Source Agent-Based Framework. *Computational Economics* 30: 291–327.
- Swider, J.D. (2006), *Handel an Regelenergie- und Spotmärkten: Methoden zur Entscheidungsunterstützung für Netz- und Kraftwerksbetreiber*. Wiesbaden: Deutscher Universitäts-Verlag.
- Tesfatsion, L. (2002), Agent-Based Computational Economics: Growing Economies From the Bottom Up. *Artificial Life* 8: 55–82.

- Tesfatsion, L. (2006), Agent-Based Computational Economics: A Constructive Approach to Economic Theory. Pp. 831–880 in: L. Tesfatsion, K.L. Judd (eds.), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, North-Holland.
- Watkins, C.J.C.H. (1989), *Learning from Delayed Rewards*. Dissertation, Cambridge.
- Weidlich, A. (2008), *Engineering Interrelated Electricity Markets – An Agent-Based Computational Approach*. Contributions to Management Science. Springer Physica.
- Weidlich, A., F. Sensfuß, M. Genoese, D. Veit (2008), Studying the effects of CO₂ emissions trading on the electricity market – A multi-agent-based approach. Pp. 91–101 in: R. Antes, B. Hansjürgens, P. Letmathe (eds.), *Emissions Trading – Institutional Design, Decision Making and Corporate Strategies*. Springer Business and Economics.
- Weidlich, A., D. Veit (2006), Bidding in Interrelated Day-Ahead Electricity Markets: Insights from an Agent-Based Simulation Model. Proceedings of the 29th IAEE International Conference, Potsdam.
- Weidlich, A., D. Veit (2008a), A Critical Survey of Agent-Based Wholesale Electricity Market Models. *Energy Economics* 30: 1728–1759.
- Weidlich, A., D. Veit (2008b), Analyzing Interrelated Markets in the Electricity Sector – The Case of Wholesale Power Trading in Germany. IEEE Power and Energy Society General Meeting 2008.
- Werker, C., T. Brenner (2004), An Advanced Methodology for Heterodox Simulation Models Based on Critical Realism. Working Paper 0401, Papers on Economics & Evolution, Max Planck Institute of Economics, Evolutionary Economics Group, Jena, Germany.
- Windrum, P., G. Fagiolo, A. Moneta (2007), Empirical Validation of Agent-Based Models: Alternatives and Prospects. *Journal of Artificial Societies and Social Simulation* 10 (2, 8).
- Xiong, G., S. Okuma, H. Fujita (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.

Dr. Anke Weidlich, Universität Mannheim, Dieter Schwarz Stiftungslehrstuhl für ABWL und Wirtschaftsinformatik – E-Business und E-Government – Schloss, 68131 Mannheim, Germany.
Phone: +49(0)621 181 3336.
E-Mail: weidlich@bwl.uni-mannheim.de

Prof. Dr. Daniel Veit, Universität Mannheim, Dieter Schwarz Stiftungslehrstuhl für ABWL und Wirtschaftsinformatik – E-Business und E-Government – Schloss, 68131 Mannheim, Germany.
Phone: +49(0)621 181 3336.
E-Mail: veit@bwl.uni-mannheim.de