

Viehmann, Johannes; Lorenczik, Stefan; Malischek, Raimund

Working Paper

Multi-unit multiple bid auctions in balancing markets: An agent-based Q-learning approach

EWI Working Paper, No. 18/03

Provided in Cooperation with:

Institute of Energy Economics at the University of Cologne (EWI)

Suggested Citation: Viehmann, Johannes; Lorenczik, Stefan; Malischek, Raimund (2018) : Multi-unit multiple bid auctions in balancing markets: An agent-based Q-learning approach, EWI Working Paper, No. 18/03, Institute of Energy Economics at the University of Cologne (EWI), Köln

This Version is available at:

<https://hdl.handle.net/10419/203683>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Multi-unit multiple bid auctions in balancing markets: an agent-based Q-learning approach

AUTHORS

Johannes Viehmann

Stefan Lorenczik

Raimund Malischek

EWI Working Paper, No 18/03

December 2018

Institute of Energy Economics at the University of Cologne (EWI)
www.ewi.uni-koeln.de

**Institute of Energy Economics
at the University of Cologne (EWI)**

Alte Wagenfabrik
Vogelsanger Str. 321a
50827 Köln
Germany

Tel.: +49 (0)221 277 29-100
Fax: +49 (0)221 277 29-400
www.ewi.uni-koeln.de

CORRESPONDING AUTHOR

Johannes Viehmann
Institute of Energy Economics at the University of Cologne (EWI)
johannes.viehmann@gmail.com

ISSN: 1862-3808

The responsibility for working papers lies solely with the authors. Any views expressed are those of the authors and do not necessarily represent those of the EWI. Stefan Lorenczik and Raimund Malischek work as energy analysts at the International Energy Agency (IEA). The content of this article reflects the authors' personal views and does not necessarily represent the views of the IEA Secretariat or the IEA Member countries.

Multi-unit multiple bid auctions in balancing markets: an agent-based Q-learning approach

Johannes Viehmann^{a,*}, Stefan Lorenczik^b, Raimund Malischek^b

^a*Institute of Energy Economics at the University of Cologne; Vogelsanger Strasse 321, 50827 Cologne, Germany*

^b*International Energy Agency; 31-35 rue de la Fédération, 75739 Paris Cedex 15, France*

Abstract

There is an ongoing debate on the appropriate auction design for competitive electricity balancing markets. Uniform (UPA) and discriminatory price auctions (DPA), the prevalent designs in use today, are assumed to have different properties with regard to prices and efficiencies. These properties cannot be thoroughly described using analytical methods due to the complex strategy space in repeated multi-unit multiple bid auctions. Therefore, using an agent-based Q-learning model, we simulate the strategic bidding behaviour in these auctions under a variety of market conditions. We find that UPAs lead to higher prices in all analysed market settings. This is mainly due to the fact that players engage in bid shading more aggressively. Moreover, small players in UPAs learn to free ride on the price setting of large players and earn higher profits per unit of capacity owned, while they are disadvantaged in DPAs. UPAs also generally feature higher efficiencies, but there are exceptions to this observation. If demand is varying and players are provided with additional information about scarcity in the market, market prices increase only in case asymmetric players are present.

Keywords: Agent-based computational economics; Auction design; Electricity markets

JEL classification: C63; D43; D44; L94

1. Introduction and motivation

The relative performance of different auction designs both in terms of efficiency and prices has been a controversial issue for many years. This is particularly true when it comes to more complex settings such as repeated Multi-Unit Multiple Bid (MUMB) auctions. In multi-unit auctions the auctioneer buys several units of the same good and bidders are allowed to place several bids.¹ In contrast to single-bid auctions, classical closed-form solutions are not available for multiple bid auctions as bidders might engage

*Corresponding author, *phone:* +49 (0)221.277 29-100

Email address: johannes.viehmann@gmail.com (Johannes Viehmann)

¹This process is also referred to as competitive bidding, in which bidders compete for the right to sell. Even though not covered in our analysis, our results are valid also for the typical auction case in which bidders compete for the right to buy.

in bid shading, i.e., they might increase bid prices in order to maximise expected profits. Another layer of complexity is added if auctions are held repeatedly, information about auction results vary, secondary markets exist or players are asymmetric, either in terms of their size or costs. The most common auction designs are Uniform Price Auctions (UPAs) in which every successful bid receives the marginal clearing price and Discriminatory Price Auctions (DPAs)² in which every successful bid is paid its bidding price. Vickrey auctions (see Vickrey, 1961) on the other hand are less common.

The above mentioned controversy about performance of alternative auction designs is mirrored in the development of the electricity industry ever since deregulation started and markets began to evolve. Nowadays, most Day-Ahead (DA) electricity markets in Europe are operated by means of transparent, repeated UPAs. However, the picture looks more diverse when it comes to procurement of balancing capacity which is also referred to as Ancillary Services (AS).³ The European Network of Transmission System Operators for Electricity (ENTSO-E) provides a comprehensive survey on how AS are currently procured and how balancing markets are designed across European countries (ENTSO-E, 2017). The procurement schemes range from several mandatory designs to bilateral arrangements between *Transmission System Operators (TSOs)* and market players to organized markets and hybrid schemes. The United Kingdom and most central European countries such as Belgium, Germany, Austria, Switzerland, the Czech Republic or Slovakia currently apply repeated DPAs to procure most of their balancing capacities (ENTSO-E, 2017). Portugal, Spain, Norway, Greece and Romania on the other hand are preferably using UPAs. Some countries such as France and the Netherlands use both auction types for different types of balancing capacities. However, even if the same auction type is applied, significant differences exist with regard to the detailed auction rules and the publication of auction results, not only among different countries but also within countries for different types of balancing capacity.

Our motivation is threefold. First, we aim to contribute to the current discussion on European harmonisation of procurement rules by showing how different levels of market concentration affect prices in UPAs and DPAs. This is particularly interesting as policy makers in Europe seem to be undecided which scheme to prefer and as the costs of procurement of balancing capacities are eventually paid by the consumers of electricity. In an earlier version of the Network Code on Electricity Balancing (NCEB), ENTSO-E (2014) stated that the procurement of balancing energy shall be based on marginal pricing (UPAs). However, according to the latest version of the guideline (ENTSO-E, 2017b), each TSO is free to define its rules for procurement of balancing capacity. On the other hand, ENTSO-E (2017a) published a consultation report

²Also referred to as Pay-as-Bid Auctions (PABAs).

³In the draft network code on electricity balancing, ENTSO-E (2013) defines balancing as "*all actions and processes, on all time-scales, through which Transmission System Operators ensure, in a continuous way, to maintain the system frequency within a predefined stability range [...]*". Any deviations from the planned schedule are defined as *imbalances* and will be balanced by the Transmission System Operator (TSO). In order to do so, the TSO procures beforehand *balancing reserves*, also referred to as *balancing capacity*, from market players (*Balancing Service Providers (BSPs)*). In case of real-time imbalances, the TSO will then call BSPs for the activation of *balancing energy*.

on the “FCR Cooperation”⁴ in which the partner TSOs propose a change from the current DPA to a UPA settlement scheme in the future.

Second, we investigate how the auction types perform in term of efficiency. In the context of our auction game, total welfare is at its maximum if costs are minimized. Hence, we define the 100% efficiency benchmark as an auction result in which only those bids that are associated to the lowest cost capacities have been accepted. While regulators and TSOs in Europe focus mainly on consumer surplus which is equivalent to lower consumer prices in markets with a non-elastic demand, social welfare doesn’t seem to be the main objective. However, we believe that the efficiency of the competing auction schemes is an essential feature that should be analysed thoroughly as well.

The third part of our motivation stems from a decision taken by the Bundesnetzagentur for Electricity, Gas, Telecommunications, Post and Railway (BNetzA) in April 2011. With this decision, several auctioning rules for the procurement of balancing capacity, precisely, of Secondary Reserve (SR)⁵, were modified (BNetzA, 2011). Of particular interest was the decision to reduce information given to Balancing Service Providers (BSPs) about auction results. Whilst, prior to the decision, accepted (infra-marginal) and non-accepted (extra-marginal) bids have been published, the BNetzA decided that TSOs would only publish accepted bids in the future (Figure 1). Even though the agency admitted that this step would reduce transparency in the market, it stated that the expected benefits were likely to outweigh the negative effects. Especially, the risk of strategic behaviour by pivotal players ought to be reduced. As the market was dominated by a limited number of big players (see Heim and Götz, 2013), the agency believed that the knowledge of prices and volumes of extra-marginal bids might have led to an increase of bid prices by pivotal players (BNetzA, 2011). To the present day, extra-marginal bids are published for Minute Reserve (MR) auctions, but not for SR and PR auctions in Germany.

In the paper at hand we develop an agent-based Q-learning model that allows for comparing UPAs and DPAs in repeated MUMB auctions. The model enables us to investigate a wide range of market settings such as a varying number of players, or player characteristics, such as symmetric or asymmetric players both in terms of size or cost. In a dynamic setting, we vary the demand and control the amount of information that is available to market players about past auction results. Thus, we can test whether information about the supply-demand ratio increases strategic behaviour and consumer prices. Vickrey auctions are not considered in our analysis as they are not found in electricity markets to our knowledge.

⁴The FCR cooperation is a cooperation between several European TSOs including the Austrian TSO APG, the Belgian TSO ELIA, the Danish TSO Energinet, the Dutch TSO Tennet, the French TSO RTE, the German TSOs 50Hertz, Amprion, Tennet and TransnetBW as well as the Swiss TSO SwissGrid. Currently, these TSOs hold a common weekly FCR (Frequency Containment Reserve, also referred to as Primary Reserve (PR)-auction.

⁵Also referred to as automatic Frequency Restoration Reserve (aFRR).




Figure 1: Schematic illustration of balancing capacity bids

First, our results indicate that marginal prices of UPAs turn out to be higher than average prices of DPAs, a result which is valid for all settings. Second, and not surprisingly, we find that with a decreasing number of players and increasing asymmetry between players in terms of size and increasing demand to supply ratios, prices are increasing. With regard to efficiency we observe that UPAs are generally more efficient, but there are some exceptions to this observation. Our analysis with respect to different information regimes shows that prices tend to increase with additional information about the supply/demand ratio only if the number of players is limited and a large asymmetry in terms of size exists.

The remainder of this paper is structured as follows. Section 2 provides an overview of related literature, section 3 introduces the applied agent-based model and the learning algorithm. Section 4 presents the results of our analysis and section 5 concludes.

2. Literature

2.1. Discriminatory (DPAs) vs. uniform pricing auctions (UPAs)

There is a vast amount of literature comparing uniform and discriminatory price auctions. While most analytical papers favour discriminatory pricing in terms of lower consumer prices, experimental and empirical publications find a high degree of collusion in repeated DPAs resulting in possibly higher consumer prices and lower overall efficiency. Fabra et al. (2006) analytically look into how bidding behaviour is affected by the auction format. They start their analysis with a model of two suppliers owning one unit of capacity each of which has to be submitted to the market by one single offer for the entire capacity. Player's capacities and costs are known but asymmetric and demand is certain and inelastic. They find that for high demand, DPA consumer prices are lower in comparison to UPA, while the effect on welfare is ambiguous and subject to the model parameters.⁶ Expanding their model to a symmetric oligopoly, Fabra et al. (2006) show by a

⁶For low-demand, prices are competitive and dispatch is efficient for both DPA and UPA.

numerical example that for a given number of suppliers in the DPA, roughly twice as many suppliers are required for the UPA in order to reduce consumer prices to the same level. Frederico and Rahman (2003) analytically compare UPA and DPA with uncertain and elastic demand and perfect information about the cost structure of the industry. In case of perfect competition, they find that a switch from UPA to DPA leads to lower consumer prices, but also to a reduction in welfare. In case of perfect collusion (monopoly), consumer prices and output are lower as well but the effect on welfare is ambiguous. Additionally, they conclude that abuse of market power by the monopolist is harder under DPA. However, if demand is certain, UPA and DPA yield identical outcomes.

On the contrary to most analytical papers, Kahn et al. (2001) strongly argue against a switch from UPA to DPA. They believe that firms will change their bidding behaviour immediately after the introduction of DPA, trying to bid at or slightly below the marginal price. As a result, average DPA prices are equal to UPA prices. However, firms face higher costs in order to forecast the marginal price and overall efficiency is likely to decrease as some low-marginal cost bids might be rejected *“because their bidders have overestimated the market-clearing price”* (Kahn et al., 2001). In the case of imperfect or oligopolistic markets, Kahn et al. (2001) argue that smaller bidders are disadvantaged in DPAs as they have higher forecasting costs per unit of production and are likely to benefit less from exertion of market power by bigger players. Hence, market entry of new players and long-term disappearance of market power becomes less likely. Finally, Kahn et al. (2001) criticise the lower transparency in DPAs which makes it difficult to detect collusive behaviour.

To our knowledge, a publication by Rassenti et al. (2003) is the most detailed experimental study comparing UPA and DPA in complex electricity market environments. Strikingly, they find that *“DPA in a no market power environment is as anti-competitive as a UPA with structural induced market power”*. Rassenti et al. (2003) argue that in an environment with cyclic and revealed inelastic demand, also present in PABAs in German balancing markets, *“the DPA invites sellers to tacitly collude, coordinating their offers without explicit communication at the highest previously observed price in a similar period”*. Similar results were obtained in an empirical study by Heim and Götz (2013) who find collusive behaviour in the DPA for SR in Germany. Using data provided by BNetzA, Heim and Götz (2013) first show that a high degree of market concentration and pivotal players exist. They believe that observed price increases can be traced back to *“repeated pretended bad guessing”* of the clearing price. As a result, firms can profit from increased price levels in later periods of the repeated DPA. Finally, they stress that regulatory authorities are unable to take legal action against abusive behaviour in DPAs as firms can hide behind the *“guess the clearing price principle”* (Heim and Götz, 2013).

2.2. Information regimes

Next to the market design, the question which information is available to market participants is of great importance when evaluating the performance of different auction designs, especially in case of repeated

auctions. However, most publications choose one particular set of information without variation. To our knowledge, there is little literature that sheds light on the effect of different information regimes in repeated MUMB auctions. Müsgens and Ockenfels (2011) present a qualitative assessment on information feedback in repeated DPAs. The article is written in the context of European balancing power markets and gives an overview of different information regimes. While in some markets there was no feedback at all about past auction results, others publish only the marginal clearing price, the volume weighted average price of all accepted bids or both. The complete bid curve including non-accepted extra-marginal bids (see Figure 1) is published only for very few markets. Müsgens and Ockenfels (2011) argue that the publication of the marginal bid price is important for the efficiency of the market. However, they reject that the benefit of additionally publishing extra-marginal bids outweighs the risk of pivotal players increasing their bid prices.

Even though Bower and Bunn (2001) do not vary the information available to players, they explicitly set up a case in which no information about the market outcome or other players' bids is made public. Merely private information and success of own bids at previous auctions is known to players. In the England and Wales electricity market with asymmetric firms, they find significantly higher prices for DPA when compared to UPA as larger firms with more bids can gather more information about past market outcomes due to their sheer size. With UPAs however, all firms with successful bids have the same information about the marginal clearing price.

2.3. Agent-based models

Agent-based models have become increasingly popular in economic studies. The applications range from microeconomic topics like the exploration of the supply function equilibrium in Kimbrough and Murphy (2013) or the Cournot equilibrium in Waltman and Kaymak (2008) to more macro-based analysis as for instance presented in Geanakoplos et al. (2012) on systemic risk in the eye of the housing bubble. Applications to energy markets are diverse, naming just exemplarily Bunn and Oliveira (2007) for an analysis of technology diversification or Naghibi-Sistani et al. (2006) for an analysis of bidding behaviour in market based power systems.⁷

One reason for this increasing popularity is that agent-based models allow to analyse situations and problems in which classical closed-form solutions are not available. Further, they allow to increase models complexity and to take into consideration more realistic modelling assumptions and real life market features that are usually excluded from economic analyses. They are also particularly suited in situations in which there is the opportunity for learning due to repeated action. This is for instance the case in daily electricity market auctions.

Agent-based models have also been used to explore the relative performance of UPAs and DPAs. These auction formats are prevalent in electricity day-ahead and balancing market auctions and serve as motivation

⁷Weidlich and Veit (2008) provide an overview of the vast applications of agent-based modelling in energy markets.

for our analysis. In MUMB UPAs and DPAs, closed-form solutions are no longer available and hence a particular interesting field to study via agent-based models. A common mistake in analysing these auctions is to directly transfer results from the corresponding single-unit auction to the multi-unit case. In general, results do not transfer from single-unit to multi-unit auctions as shown for instance in an overview by Krishna (2002) or the literature cited therein. Therefore, in the absence of closed-form solutions, agent-based models provide a valuable option to analyse these auction types. Previous agent-based analysis of UPAs and DPAs include the works of Hailu and Thoyer (2007), Bower and Bunn (2001), Bakirtzis and Tellidou (2006) as well as Xiong et al. (2004). These analyses have painted an unclear picture of the relative performance of the different auctions. Whereas for instance Bower and Bunn (2001) find evidence that prices are higher in DPAs due to the non-availability of market prices, Hailu and Thoyer (2007) argue that there is no clear ordering of the auction formats and results are dependent on the population and the characteristics of supply relative to demand. All studies have assumed a fixed information regime.

Our analysis adds to the existing stream of literature on three key aspects: First and most importantly, we are to our knowledge the first to extend the Q-learning algorithm to a Multi-Unit Multiple Bid (MUMB) set-up, other than Bakirtzis and Tellidou (2006) and Xiong et al. (2004) who use a multi-unit single bid approach. Second, we incorporate demand uncertainty for each consecutive auction. And third, we analyse the effect of different information regimes, i.e., the information about market outcomes that is provided to the players after each round of play. To our knowledge we are the first to incorporate the different information that can be provided to the players into the learning algorithm and to explore its effect systematically.

3. The model

We explore bidding strategies of players under different auction and information regimes using an agent-based model. In the model, players bid price-quantity pairs. In doing so, players take into account their costs for providing the good as well as a player specific capacity constraint, i.e., the amount of the good a player can supply. Players have two separate blocks of a fixed size, for which they can make separate bids. We incorporate different bid configurations for each player, ranging from rather flat supply curves to more hockey stick shaped bids. The maximum bid for a block is restricted by an upper bound. The introduction of a price cap is necessary in order to prevent prices to approach infinity as demand is assumed to be inelastic. The minimum possible bid is the player's marginal cost. The inelastic demand is either deterministic or stochastic, based on the setting.

Each auction is repeated many times to enable learning. Also each set-up is repeated several times to check robustness of results. We use the Q-learning algorithm which is a variant of the reinforcement learning approach. Players get feedback on their actions and improve their behaviour in successive rounds of play

allowing them to learn from the past through memory.⁸ In general, the Q-learning framework consists of a memory-state space which is given by a set S , elements of which are represented by s . s_t then represents the state a certain player is in at time-point t . Agents choose actions $a \in A$ which lead to a transition from state s_t to s_{t+1} . Actions are chosen by drawing from the probability distribution over the action space given by

$$Prob(a) = \frac{\exp(Q_t(s_t, a)/\beta)}{\sum_{a' \in A} \exp(Q_t(s_t, a')/\beta)} \quad (1)$$

This representation is also known as the Boltzman exploration strategy and corresponds to a logit model. $Q_t(s_t, a)$ denotes the Q-value of the agent when he is in state s_t and action a . Intuitively, the Q-value represents how favourable the execution of action a in state s_t is. $\beta > 0$ represents the experimentation parameter: the higher the value of β , the more experimentation is performed as the probability $Prob(a)$ of choosing any action a is closer to being evenly distributed. In our implementation, following the literature (see, e.g., Waltman and Kaymak, 2008), we use a gradually decreasing parameter β of the form

$$\beta(t) = 1000 * 0.99995^t \quad (2)$$

The parameter β steers the exploration phase and the subsequent transition to the exploitation phase. Using the Q-learning algorithm – as for many other learning algorithms – it is necessary for players to have a sufficiently long exploration phase in which they randomly choose all actions many times to learn about possible pay-offs. After the transition to the exploitation phase is completed, parameter β ensures that players select those actions with the highest expected pay-offs only.

The Q-values of the players are updated after each round of play according to the following rule:

$$Q_{t+1}(s, a) = \begin{cases} (1 - \alpha)Q_t(s, a) + \alpha(\pi_t + \gamma \max_{a' \in A} Q_t(s_{t+1}, a')) & \text{if } s = s_t \text{ and } a = a_t \\ Q_t(s, a) & \text{otherwise.} \end{cases} \quad (3)$$

$0 < \alpha \leq 1$ and $0 \leq \gamma < 1$ represent the learning parameter and the discount factor. In our setting they are chosen to be 0.5, similar to other analysis in the literature (see, e.g., Waltman and Kaymak, 2008). A lower value of the learning parameter α implies that more weight is put on the old Q-value in the updating process, which can be interpreted as putting more weight on the player's history as compared to recent experience. The γ parameter represents the time preference of the players, with smaller values of γ indicating more myopic behaviour. π_t represents the players pay-off after round t . As a result of the Q-learning algorithm, each players individually learns the optimal behaviour that maximises its pay-off in the long run.

⁸(see, e.g., Bakirtzis and Tellidou, 2006; Xiong et al., 2004; Kutschinski et al., 2003) for previous applications of Q-learning in economic research.

Applied to our MUMB auction, the state s_t of a player at time t is defined by the auction results of the previous round such as the marginal price, the volume weighted average price of all accepted bids and for our uncertainty case (see section 4.4) also the level of total demand. Given the state s_t , each player now selects an action a_t according to equation 1 which means, he decides at which prices he is going to place his capacity via two bids of equal size into the market. Next, all bids of all players are collected, the market clears and moves to a new landing state s_{t+1} . Each player can now compute its profits and also the maximum Q-value of the new state. These information are now used to update the corresponding Q-value of the initial state (see equation 3). After one auction round is concluded, a new auction round begins, this time with the landing stage s_{t+1} as starting point. The less profitable an action is, the lower the Q-value gets and the less often a player selects the corresponding strategy. In the final exploitation phase, a player will only select the strategy with the highest Q-value.

4. Simulation and results

4.1. Overview

Unlike in energy only markets, there is only very little fundamental information in balancing markets. German TSOs for example merely publish a list of all players on a firm level, but not of the individual units prequalified to provide balancing capacity. Firms might be able to obtain some information regarding the general availability of generation units which are assumed to be able to provide balancing capacity via mandatory messages about planned and unplanned non-usabilities of generation units. However, even if units are available to produce power on the energy-only market, specific technical reasons that do not need to be published might prevent them from providing balancing capacity. Cost calculations are even more complex. First, players are likely to face opportunity costs for selling production units into competing markets like the energy-only market or other balancing markets. Second, it is unknown how competitors split costs between positive and negative balancing capacity in case a unit is sold for positive and negative reserve at the same time, how they estimate added profits from energy calls, how they can reduce capacity costs using the portfolio effect (see also Müsgens et al., 2012) and how they calculate back-up costs in case of unplanned outages. The Q-learning approach is particularly suitable to simulate strategic behaviour in balancing markets as players have very little fundamental information. Therefore, they mainly use historical auction prices as basis for their bidding. The same is true for our Q-learning approach in which prices (and demand levels) are the only information required to define the initial states s_t of each auction round.⁹

We distinguish between three main model settings. We start with the *Base Case* in section 4.2, in which demand is constant and all players have capacity costs of zero¹⁰. We analyse how the number of players,

⁹It is important to work with a limited number of criteria (such as volume weighted average price, marginal price and level of demand) that define states in Q-learning. Otherwise, the memory-state space grows too large and computing time increases.

¹⁰Cost of zero were chosen arbitrarily, we could have also selected costs of 1 or 2. In the base case it is merely important that costs are constant and equal for all players.

their symmetry in terms of size and different demand levels affect prices of UPAs and DPAs. In the *Cost Case* (section 4.3), we allocate costs to the individual players capacities in order to study how the competing auction designs perform in terms of efficiency. Finally, we present the *Uncertainty and Information Case* in section 4.4. In this case, we vary the demand from auction to auction within a model run and also modify the information players receive about total demand. In all model settings, players receive information about the state they are in, which is defined by the previous auctions marginal price and the volume weighted average price of all accepted bids. Additionally, information about the demand level is added in our *Uncertainty and Information Case*. Size and costs of a player are always constant within a model run.

On the more technical details: if two or more players bid capacities at the same price and this price is equal to the marginal price, the principle of pro-rata allocation is applied. The price cap is set to 10, the price floor to zero. All players have costs of zero except in the cost case. Each scenario is calculated with either four and/or eight players and different (average) demand levels. We run each case 50 times as a general rule with 400,000 to 700,000 successive auctions in each run. At the beginning of the subsequent auctions, players choose their actions randomly while learning which ones have the highest pay off (exploration phase). Later on, players base their choice on previous experiences: They have learned successful strategies and start to exploit this knowledge (exploitation phase, see also section 3). We run each model configuration 50 times with varying seeds for the random number generator. This way we are able to check for more multiple stable outcomes and the robustness of results. The number of auctions per run depends on the complexity of the model settings. In our basic scenario with symmetric players, merely 400,000 successive auctions per run are required before a stable equilibrium is found. With a higher level of (demand) uncertainty, agents require up to 700,000 runs to commit themselves to a limited set of strategies. For our static *Base Case* and *Cost Case* we denote the equilibrium as stable if the average price of the last 10,000 auctions deviates by less than 2% from the average price of the 10,000 auctions that were hold 100,000 rounds prior to the last 10,000 auctions¹¹. However, for the *Uncertainty and Information Case*, we change the definition for a stable equilibrium as this case is more dynamic due to the varying demand. Therefore, stability of the equilibrium is defined by the stable formation of Q-values of the individual players.¹² As the players are exploring strategies at the beginning, we truncate the results of the majority of the auctions. The average prices presented in the following subsections take into account merely the last 10,000 auctions of each model run. Also, the term *average prices* always refers to the marginal price in UPAs and the volume weighted average price of accepted bids in DPAs across all 50 runs.

¹¹As an example, if the average price of the successive auctions 390,001 to 400,000 deviate by less than 2% from the average price of auction rounds 290,001 to 300,000, the result is considered to be stable.

¹²Prices have been used as convergence criterion for the Base and Cost Case merely for the simplicity of the approach. Checking for convergence of the Q-values, as done in the Uncertainty and Information Case, would have yielded the same result.

4.2. Base Case

We start by analysing the base case without uncertainty, i.e., both demand and supply are constant for all subsequent auctions of each model run. For this case, we will solely focus on the influence of the market clearing scheme given different settings with regard to the level of demand, the number and the symmetry in terms of size of participants. Aggregated supply is always set to 100. Hence, in the symmetric base case, there are 4 (or 8) identical players competing with a capacity of 25 (12.5) units each. Costs are set to zero for all players. Figure 2(a) illustrates the results for different levels of demand for 4 (dashed lines) and 8 (dotted lines) symmetric players. The uniform price scheme is indicated in grey, discriminatory pricing in black. Focussing on the influence of the number of players, we observe that – as expected – with a higher number of market participants prices decrease¹³. The opposite is true for the market demand: the higher the demand levels, the higher the market prices observed. At a demand level of 100 (demand to supply ratio of 1), prices always converge to the price cap of 10, hence this data point is not shown in any of the figures.

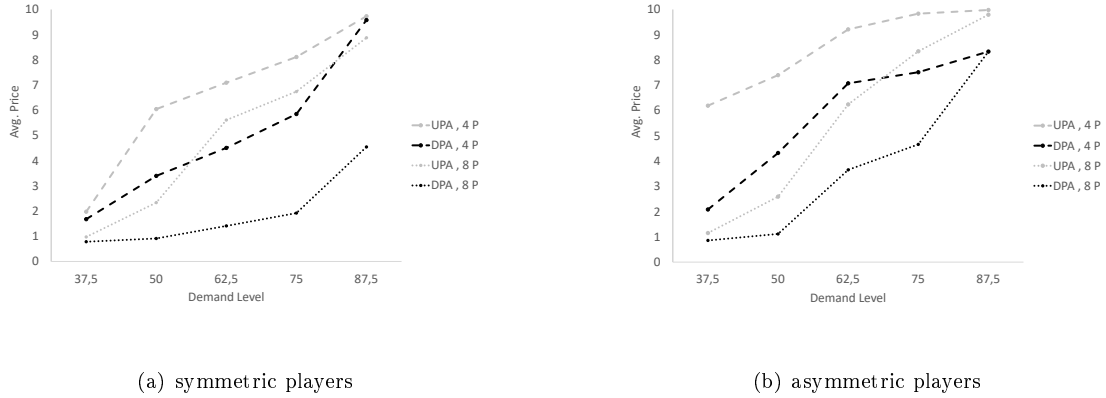


Figure 2: Avg. price in UPA and DPA with symmetric and asymmetric players

With regard to the market clearing scheme and focussing on the symmetric case (Figure 2(a)), we can observe generally higher prices in UPAs than in DPAs. This observation is independent from the number of market participants and the level of demand. Merely for low levels of demand at which prices are close to the price floor and high demand at which prices are close to the price cap, both auction types converge. The shape of the price increase from low to higher levels of demand indicates that strategic bidding is easier at an earlier stage in the case of UPAs. Especially in the market with 8 players, market prices with DPA remain closer to the price floor while UPAs prices have already increased. Regarding the market power of individual players, demand levels at which (in the symmetrical case all) players become pivotal¹⁴ are at 75

¹³Some sample runs with less than 4 and more than 8 players clearly confirm this observation.

¹⁴If a player is pivotal, demand cannot be covered without some capacity of this player. Hence, the barrier at which a player becomes pivotal serves as a measure of market concentration and indicates a potential for collusive behaviour.

for 4 and 87.5 for 8 players. In UPAs, the most rapid price increases can already be observed at demand levels without any pivotal player (between 37.5 and 50 for 4 and between 50 and 62.5 for 8 players). In case of DPA, the most rapid price increase occurs around the pivotal levels. Thus, players appear to be able to collaborate more easily with UPAs and can coordinate on higher prices even if capacity is hardly scarce.

A closer look into bidding strategies reveals some possible explanations for the price patterns observed. Comparing the individual average profits of the players, it can be noted that they hardly deviate from each other. The average standard deviation of observed profits is negligibly small. This is noteworthy as individual players have no information about other players' bidding strategies or realised profits, yet they find an equilibrium in which profits are nearly evenly distributed amongst them. Therefore, we consider average bid prices and acceptance rates of all players in the subsequent analysis on bidding strategies. In Figure 3(a), the average bid prices of the first and second bid are shown. The first bid is lower than or equal to the second bid by definition. It is very obvious that the price spread between the first and second bid is much lower in DPAs when compared to UPAs, independent of the demand level. As derived by Krishna (2002), players in UPAs learn to heavily shade their bids and place the second bid significantly higher than the first one on average. This effect is less pronounced in DPAs, but still existing, which is an interesting finding by itself. Our data also reveals that the rate at which players bid both bids at the same price is twice as high for DPAs in the 4 player auctions (55 vs. 27%) and almost three times as high in the 8 player auctions (58 vs. 21%) across all demand levels.




Figure 3: Avg. bid prices and bid acceptance rates in UPA and DPA with symmetric players

The corresponding average bid acceptance rates of the first and second bid are displayed in Figure 3(b).¹⁵ The bid acceptance rate is generally higher for the first and lower for the second bid in UPAs compared to DPAs. This seems plausible as players engage more aggressively in bid shading and the first bid potentially

¹⁵The bid acceptance rate shows which share of the capacity of the first and second bid is accepted on average.

profits more from a high second bid, even if the acceptance rate of the second bid is very low. In the 4 player UPA case at a demand level of 50, the highest differential between first bid acceptance rate (98.4%) and second bid acceptance rate (1.6%) can be observed. At the very same observation point, the average difference between bid prices (see Figure 3(a)) reaches its maximum (first bid at 2.0, second bid at 7.9) and the average second bid price even its absolute maximum. As shown earlier in Figure 2(a), UPA prices rise steeply towards the demand level of 50 and the difference between UPA and DPA is at its maximum. For the 8 player case, the picture looks very similar at the demand level of 62.5. Hence, the higher UPA prices are a result of complex interactions between second bid acceptance rate and more aggressive bid shading.

As a first modification from our symmetric base case, we also consider the case of asymmetric players as shown in Figure 2(b). In our 4 player setting, one large player is endowed with half of the capacity, while the remaining 3 players are identical and own a capacity of one third of 50 each. In our 8 player setting, we have two large players are endowed with a capacity of 25 each, while the remaining 6 identical players merely own one sixth of 50 each. All other parameters remain unchanged. Generally, prices exceed those of the symmetric case¹⁶. This is to be expected as large players can exercise more market power than small players. Also, prices tend to rise faster towards the price cap at lower demand levels, at high demand levels the gradient of the price increase slows down. The general influence of the auction type remains unchanged, meaning that UPAs result in higher price levels than DPAs independent from the level of demand. Demand levels at which the large player(s) become pivotal are at 50 for 4 and at 75 for 8 players.

In the asymmetric case, the analysis of the different bidding strategies of the large and small players is of particular interest. The average bid acceptance rates of the single large player (L) and the three small players (S) in the 4 player case¹⁷ are displayed in Figure 4(a). At low and medium demand levels, the acceptance rates of the small players' first bids are much higher than those of the large player, the same is true for the second bids. In fact, the first bid acceptance rate of the small players at a demand level of 37.5 in the UPA case is already at 97.8% (large player at 46.1%). At a demand level of 75, the second bid acceptance rate of the small player is at 100%, while the second bid of the large player is not accepted at all. These differences are due to the fact that the large player places on average higher bids than the small players and takes the role of setting the marginal price. Our data shows that the large player is placing the marginal bid in more than 99.5% of all auctions for demand levels of 62.5 or higher for both auction types. The difference in bid acceptance rates is generally smaller in DPAs.

In UPAs, the small players learn to free ride on the high price setting of the large player. Hence, their profits per unit of capacity owned¹⁸ are significantly higher than those of the large player for all demand levels (see Figure 4(b)). This is due to the fact that they earn the same marginal price for a higher share of their

¹⁶One exception is the 4 player DPA at a demand level of 87.5. An explanation is given at the end of this subsection.

¹⁷In this section, we focus on the 4 player case exclusively. While the results of the 8 player case are similar, they are less pronounced as there are two large players that have less market power than the single large player in the 4 player case.

¹⁸If a share of a player's capacity is not sold, it is valued at a profit of zero.

capacities. The distribution of profits looks quite different when it comes to DPAs, in which the differences in profits are much less pronounced. The small players still manage to secure higher bid acceptance rates, but this comes at the cost of bidding at lower prices. In fact, they maintain a considerable safety distance to the large players' bids in order for the large player not to be tempted to underbid the small players. As this safety distance is particularly large at high demand levels (75 and 87.5), the large players' profit per unit of capacity owned even exceeds the average profits of the small players¹⁹. These findings confirm a statement by Kahn et al. (2001) who argue that smaller bidders are disadvantaged in DPAs as they are likely to benefit less from the exertion of market power by bigger players.




Figure 4: Avg. bid acceptance rates and profits in UPA and DPA with 4 asymmetric players

4.3. Cost Case

In the base case above, capacity costs were set to zero. Now, we allocate different costs to the individual capacities of the individual players. Hence, this scenario additionally allows for analysing the efficiency of the competing auction types. In the context of our auction game, total welfare is at its maximum if costs are minimised. Hence, we define the 100% efficiency benchmark as an auction result in which only those bids that are associated to the lowest cost capacities have been awarded. As soon as one low capacity cost bid is not awarded and replaced by a higher capacity cost bid, the cost base increases to above 100% and the efficiency decreases accordingly.

We conducted a wide range of trial runs with different cost allocations and found two main cases whose results show different characteristics. In the first set, costs are allocated in ascending order (cost scenario CA) as shown in Table 1. In the 4 player case, this translates into the first player having cost of zero allocated to one half of its capacity and costs of 1 to its other half, the second player having costs of 2 and

¹⁹This is also the reason for the single exception observed (4 player DPA, demand level of 87.5), at which prices of the symmetric case are higher than in the asymmetric case.

3 and so on. In the second case, costs are mixed among players (cost scenario CM). In the 4 player case, the first player is now endowed with cost of zero and 7, the second player with 1 and 6 and so on²⁰.

		P1	P2	P3	P4	P5	P6	P7	P8
4 Player	Cost ascending (CA)	0/1	2/3	4/5	6/7				
	Cost mixed (CM)	0/7	1/6	2/5	3/4				
8 player	Cost ascending (CA)	0/0	1/1	2/2	3/3	4/4	5/5	6/6	7/7
	Cost mixed (CM)	0/7	0/7	1/6	1/6	2/5	2/5	3/4	3/4

Table 1: Cost allocation schemes among players (1st half of capacity/2nd half of capacity)

The results of the 4 player auction are displayed in Figure 5(a). In the ascending cost scenario (CA), prices increase mainly in parallel with the minimum cost. Only at a demand level of 75 for UPAs (87.5 for DPAs), the price increase accelerates. The price differences between UPA and DPA are small when compared to our zero cost case as displayed in Figure 2(a), but UPA prices again exceed those of DPAs for all demand levels. The corresponding cost base is generally decreasing (efficiency is increasing) with higher demand levels. There are two systematic reasons for this trend. First, the relative cost difference is higher at low cost (and low demand) levels even though the absolute cost difference between our costs steps is constant and amounts to 1²¹. Second, the higher the demand level, the less high-cost extra-marginal capacities are available that might potentially increase the cost base. At a demand level of 100 (not shown in the graphs), the cost base and efficiency are per definition at 100% as all capacities available are required to cover the demand. With regard to the auction types, the cost base of UPA is 4 to 10 percentage points lower when compared to DPA for demand levels of 50 or higher. Merely at low demand levels, UPAs feature a slightly higher cost base. Comparing the mixed cost scenario (CM) to CA reveals some remarkable differences, even though the overall costs and hence the minimum cost are the same. The modified allocation of costs leads to a steep increase of UPA prices at demand levels of 62.5 and 75, while DPA prices are very similar to CA scenarios and rise in parallel with the minimum cost for the most part. The price increase coincides with a sudden rise of the cost base. At a demand level of 62.5, the UPA cost base exceeds DPA by almost 5 percentage points, whereas UPA cost base is equal to DPA at a demand level of 75 and 5 to 23 percentage points lower than DPA for all other demand levels.

The results of the 8 player auction are shown in Figure 5(b). Again, prices in the CA scenario rise in parallel with the minimum cost, the UPA price increase at demand levels of 62.5 and 75 in the CM scenario is much less pronounced than in the 4 player case but still visible. The corresponding increase in the cost

²⁰We also conducted extensive runs for a mixed cost case in which the first player is endowed with capacity costs of zero and 4, the second player with 1 and 5 and so on. As the results are very similar to the CM scenario as shown in Table 1, we choose not to present the data here.

²¹As an example, if a low capacity cost bid (cost of 1) is replaced by the next higher bid (cost of 2), the cost base increases by 100%. If a bid with associated costs of 4 is replaced by the next higher bid (cost of 5), the cost base merely increases by 25%.

base is less steep as well. Therefore, UPA cost base remains below DPA for all demand levels by up to 34 percentage points. The same holds true for the CA scenarios, in which DPA cost base exceeds UPA by up to 16 percentage points.

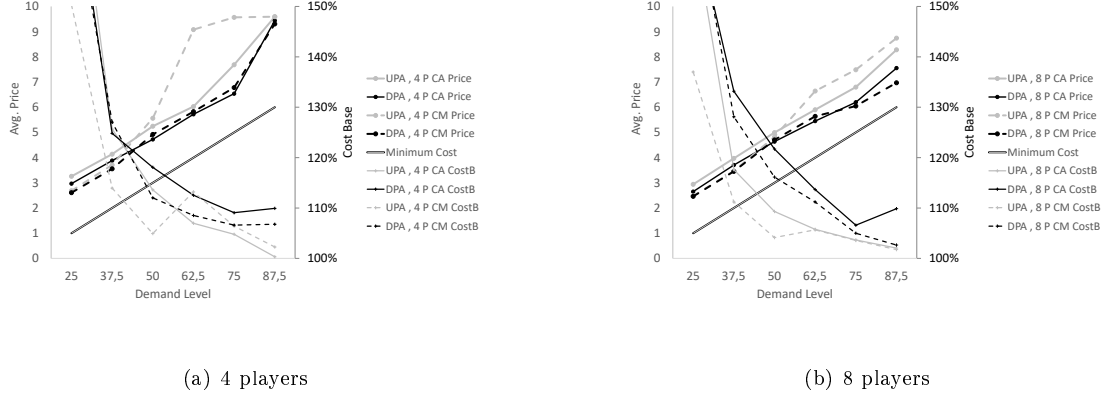



Figure 5: Avg. price and cost base in UPA and DPA with 4 and 8 players

Taking a closer look, we first explore why the UPA cost base is lower than DPA in most cases and second, why prices in the UPA CM scenario are significantly higher than in CA at demand levels of 62.5 and 75. The average bid acceptance rates of the 4 player auctions are shown in Figure 6. The subfigures are now arranged by the underlying capacity costs in order to put the competing cost allocation scenarios on a comparable basis²². The average bid acceptance rate in the UPA ascending cost CA scenario (see Figure 6(a)) displays an interesting pattern. Players learn to push in their first (low capacity cost) bid very aggressively, while they place their second bid at a price in between the next (higher capacity cost) players' first and second bid. Also, they rise the bid acceptance rate of their second bid close to 100% before the second next player is able to increase its first bid well above 20%. At a demand level of 50, the average first bid acceptance rate of the third player with underlying capacity costs of 4 (C4, P3, B1) is significantly higher (59%) than the second players' second bid (C3, P2, P2, 31%), while the rate of the first players' second bid (C1, P1, B2) is already above 99%. This pattern is true for all players and all demand levels. Low capacity cost players largely exclude higher cost players from the market (except of the first bid of the next player), which comes at the expense of relatively low UPA prices as seen in Figure 5(a) and of a certain degree of inefficiency, as the higher cost first bids of the next player have higher acceptance rates than the previous players lower cost second bids.


As shown in Figure 6(b), the pattern of the acceptance second bid rates looks quite different in the DPA CA scenarios. At demand levels of 50 or higher, the second bid acceptance rates of lower cost capacities

²²For explanation: the first line in Figure 6(a) shows the first bid (B1) acceptance rate with underlying capacity cost of zero (C0) of player 1 (P1).


are significantly lower, giving room to more costly capacities, while the first bid acceptance rates are only slightly lower than in UPA. At a demand level of 50, the average second bid acceptance rate of the first player (C1, P1, B2) is down from 99 to 82% when compared to UPA, while the rate of the second player (C3, P2, B2) is up from 35 to 56%. At this point, the first player sets the marginal bid in 76% and plays a same bid price strategy in 71% of all auctions. In UPAs on the other hand, the same player sets the marginal bid in only 6% of all auctions as the average second bid price is considerably lower (2.7 versus 4.8). In DPAs, low cost players cannot hide behind higher cost players by placing low bid prices, hence they are forced to play the same bid price strategy and set the marginal price more often. This ultimately results in a higher cost base (lower efficiency) for DPAs at demand levels of 50 or higher and for all demand levels in the 8 player case as shown in Figure 5(b).




(a) UPA and ascending costs CA



(b) DPA and ascending costs CA



(c) UPA and mixed costs CM



(d) DPA and mixed costs CM

Figure 6: Avg. bid acceptance rates in UPA and DPA with 4 symmetric players for cost scenarios CA and CM

Due to the different allocation of costs, players modify their bidding strategies in the UPA CM scenarios as displayed in Figure 6(c). The acceptance rates of the 4 lowest capacity costs bids (C0 to C3), which are equivalent to the players first bids in CM, are decreasing as costs increase. The lowest cost bid features the

highest acceptance rates, the second lowest the second highest acceptance rates and so on. As the difference of the first and second capacity is larger now, the strategy of pushing more expensive players entirely out of the market does not apply any more. At a demand level of 50, the 4 cheapest bids, which are sufficient to cover demand, possess a combined market share of 97% (as compared to 84% in UPA CA), which also explains the lower cost base of 105% (114% in UPA CA). However, at a demand level of 62.5, the pattern suddenly changes. The average acceptance rate of the next expensive capacity bid (C04, P4, B2) drops to 31% (as compared to 93% in UPA CA for bid (C04, P3, B1)) and the rates of the three most costly capacity bids rise to on average 23% (9% in UPA CA). As a consequence, the cost base rises to 113%, exceeding the cost of UPA CA (107%) and even the cost of DPA CM (109%). As the bidding strategy from UPA CA is not applicable any more, players learn to jointly raise their second bids. Even though efficiency decreases, this strategy pays off in terms of higher UPA prices and higher profits per player.

Finally, we compare the results above with those of the DPA CM scenarios as shown in Figure 6(d). At demand levels of 50 or lower, low capacity cost bids have on average lower acceptance rates than in UPA CM, resulting in a higher cost base. However, compared to DPA CA, the capacity cost allocation leads to a more distinct segregation, with the order of acceptance rates of individual capacities mostly in line with their underlying costs. Hence, the cost base is lower at demand levels of 50 or higher by 2 to 6 percentage points.

To sum up, the argument of Kahn et al. (2001) that efficiency in DPAs is worse compared to UPAs as some low cost bids might be rejected as bidders overestimate the market clearing price can only partly be confirmed by our results. We have shown several examples, in which the efficiency of UPAs is worse than in DPAs. This is particularly true for low supply to demand ratios and for certain cost allocations among the players. Plus, there seems to be a general misconception about efficiency in MUMB UPAs. As bidders engage in bid shading and increase prices of their second bids, UPA results are generally inefficient as well.

4.4. *Uncertainty and Information Case*

In this section, we start with our initial setting from section 4.2 in which all players have zero costs. We vary the market set-up in two ways: first, we introduce demand uncertainty, i.e., in each auction round, demand is randomly adjusted within the interval $[-10;10]$ from the indicated average level of demand and demand changes from auction to auction are randomly chosen from $\{-1,0,1\}$.²³ Second, the market participants are provided with two different levels of information about the previous auction: either the demand of the previous auction is known (Info 2) or unknown (Info 1). If it is unknown (Info 1), players have to develop a bidding strategy not knowing the current level of demand, i.e., a strategy that maximises the expected profit for varying demand. In case demand of the previous auction is known, players still face


²³This restriction is important. If demand changes were completely random in the whole interval $[-10;10]$, information about the previous demand would be entirely worthless.

the uncertainty of demand changes from one auction to the other within the range $\{-1, 0, 1\}$. However, knowing the previous demand reduces uncertainty about the demand in the next auction and gives an indication about scarcity in the market and whether a player might be pivotal or not.


We are aware of the fact that in most European balancing markets, the demand for balancing capacity is rather constant and the supply is varying instead. Even though one might think that it is irrelevant whether there is uncertainty about demand or supply, there are some structural differences. This is mainly due to the fact that if supply is varying, the individual players possess additional private information as they are aware of the change of their own supply. Hence, even in the Info 1 regime in which players do not know the total supply, they are aware of their own supply changes. The larger a player's capacity is, the more conclusions he can draw from its own to the total supply changes. This higher level of information in the Info 1 case diminishes the differential between the two information regimes. For this reason, we decided to show demand instead of supply variation in this section.

The results of the 4 player asymmetric case with one large player being endowed with half of the capacity and the remaining players with one third of 50 each is shown in Figure 7(a). We choose the asymmetric case as differences between the two info levels are much more pronounced compared to symmetric player case. As seen in the previous sections, UPA prices always exceed those of DPAs and prices rise with increasing demand. The effect of the information level on prices is significant around the pivotal demand level of the large player (37.5 to 62.5). If players know about the total demand of the previous auction (Info 2), prices exceed those of the Info 1 regime for both UPA and DPA. The difference between the two information regimes is largest at a demand level of 37.5 for UPAs and 50 for DPAs. However, at demand levels of 75 or higher the differential between the two information regimes becomes very small, with Info 1 prices even slightly exceeding those of Info 2.

A first explanation for the price pattern observed can be found in Figure 7(b). Here, each observation point of Figure 7(a) is further decomposed into three sub demand levels. For example, the average Info 2 UPA price at an average demand of 37.5 amounts to 4.6 (see Figure 7(a)). However, the actual underlying demand varies between 27.5 and 47.5. We divide this range into three parts of equal size and then report their average prices in Figure 7(b). In this case, the average price for low demand levels (interval [27.5;34.2]) amounts to 2.4, for medium demand levels [34.2;40.8] to 4.8 and for high demand [40.8;47.5] to 6.7. The Info 2 gradient of the price increase exceeds the one of Info 1 at all demand levels and for both auction types. This indicates that players use the additional information about the demand level of the previous auction to actively differentiate their bidding strategies, which is not possible in the Info 1 regime. This effect is largest around the pivotal demand level of 50 of the large player. However, at high average demand levels, Info 1 prices exceed those of Info 2 in the low demand sub-interval, as players in Info 2 actively push down prices while the slope of Info 1 is rather flat. This effect results in overall slightly higher Info 1 prices for demand levels of 75 or higher as seen in Figure 7(a).




(a) overview




(b) grouped by demand level

Figure 7: Avg. price in UPA and DPA with demand uncertainty and information levels with 4 asymmetric players

As already assumed, the large player can be identified as the main driver for higher prices and steeper price gradients observed in the Info 2 regime. The average bid prices²⁴ of the large and small players are displayed in Figure 8. At average demand levels of 37.5 to 62.5, the large player aggressively differentiates its bidding strategy with respect to the sub demand levels (see Figure 8(a)). This is intuitive because if the large players know that he is pivotal, he can exploit this situation more strongly compared to the case in which he has to guess its strategic position (Info 1).



(a) large player



(b) avg. of small players

Figure 8: Avg. bid prices in UPA and DPA with demand uncertainty and information levels with 4 asymmetric players

Small players on the other side exhibit different bidding strategies as displayed in Figure 8(b). In UPAs, average bid prices are significantly lower when compared to the large player and the differential between Info 1 and Info 2 is smaller as well. This is mainly due to the fact that small players free ride, independent of the information regime. In DPAs, differences compared to the large player are smaller, but absolute bid

²⁴To simplify the graphic, we show the combined average price of the first and second bid in Figure 8.

prices are still at a generally lower level. It is very interesting to observe, that the largest price increase in DPA Info 2 occurs at an average demand level of 87.5, the level at which the small players become pivotal.

We choose to exclusively show the 4 player asymmetric case in this section which features the most significant differences between the information regimes. However, we have conducted several test runs with symmetric and with 8 players. In both cases, the price differentials between the two information regimes diminish. This is mainly due to an effect that we have observed in our 4 player asymmetric case at high demand levels of 62.5 and 75 (see Figure 7(a)). In the absence of a dominant player, players still actively differentiate their bidding strategies with Info 2. However, their higher prices in the high demand subinterval are mostly offset by lower prices in the low demand subinterval when compared to the Info 1 prices. Hence, average prices turn out to be very similar.

5. Conclusion

In this paper, we develop an agent-based Q-learning model to simulate strategic bidding behaviour in repeated auctions under varying market clearing schemes. To our knowledge, we are the first to extend the Q-learning algorithm to a MUMB set-up. In addition, we analyse the influence of uncertainty and different information regimes regarding previous auctions provided to the players. Our findings are manifold:

First, we find that with an increasing number of players and increasing supply to demand ratio prices are decreasing. This corresponds to common expectations and indicates reasonable bidding behaviour of the modelled players. Furthermore, we observe higher average prices under UPA than with DPA. This result is valid for all levels of demand and number of players. Uniform pricing seems to facilitate strategic bidding in repeated auctions even if the available supply significantly exceeds demand and no single player is pivotal. This is mainly due to the fact that players aggressively engage in bid shading by rising the price of their second bids as described by Krishna (2002). Even with low bid acceptance rates, their first bid can profit disproportionately from an elevated second bid. This is not the case in DPAs. Our analysis shows that the difference between the first and second bid is significantly smaller when compared to UPA and the same bid price strategy is applied more often. As the number of players increases and as prices approach the cap or the floor price, price differences between UPA and DPA diminish. With regard to asymmetric players, our findings confirm Kahn et al. (2001) who argue that smaller bidders are disadvantaged in DPAs as they are likely to benefit less from the exertion of market power by bigger players. In UPAs on the other side, small players can free ride on the large players' bidding strategy and obtain significantly higher profits per capacity owned.

Second, we compare the auction types with regard to their efficiency. We find, that the results are ambiguous and that the argument of Kahn et al. (2001) who claim that UPAs are always more efficient as in DPAs, some low cost bids might be rejected as bidders overestimate the market clearing price, cannot be

confirmed. While in the majority of cases their statement can be confirmed, we have shown several examples in which the efficiency of UPAs is lower than in DPAs. This is particularly true for low demand to supply ratios and for certain cost allocations among the players. Our results indicate that the common expectations about efficiency of MUMB UPAs might not be generally true. As bidders engage in bid shading and increase prices of their second bids, UPA results are generally inefficient as well.

Third, we are able to analyse the influence of published information concerning previous auctions on average prices. For this purpose we introduce demand uncertainty in our model. Although the effect of providing more information about the demand level of previous auctions is ambiguous with symmetric players, prices tend to increase with asymmetric players with additional information. This is particularly true if the demand levels are close to the pivotal level of the large player. This is due to the fact that the large player aggressively differentiates its bidding strategy with respect to the sub demand levels as he knows whether he is pivotal or not. Without additional information, the large player bids at lower prices as he has to guess its strategic position. Again, differences get smaller when supply to demand ratios, the number of players or symmetry among players increases.

Based on our simulation results we conclude the following: For markets with many participants and limited market concentration, UPAs may be favourable compared to DPAs even though they might result in slightly increased market prices. Typical issues with DPAs can be avoided when choosing UPAs. With UPAs, players have lower transaction costs (in DPAs they need to forecast the marginal price) and smaller players are disadvantaged as they benefit less from the exertion of market power by bigger players. This also leads to the fact that market entry of new players is less likely. Plus in general, efficiency of UPAs is higher, especially if there are many symmetric players.

Concerning markets with a small number of players and potentially few large players – which might be the case in some balancing markets – our results indicate that DPAs are advisable if low prices are the main objective. They limit the potential for the exertion of market power and result in lower average market prices. The publishing of previous supply to demand ratios should also be handled with care, as our results indicate that additional information may facilitate strategic bidding behaviour. Coming back to our introductory example concerning the changes in the German balancing markets, our analysis provides some support to the choice by BNetzA not to disclose information about the supply to demand ratio of past auctions. However, as the number of players increases²⁵ and pivotal players disappear, a switch to UPA may be advisable and the amount of information about past auction results should be increased.

²⁵See Viehmann (2017) for the recent development of the number of balancing capacity providers in Germany.

6. Acknowledgements

The views expressed herein are strictly those of the authors and do not necessarily represent those of EWI or IEA. Stefan Lorenczik and Raimund Malischek work as energy analysts at the International Energy Agency (IEA). The content of this article reflects the authors' personal views and does not necessarily represent the views of the IEA Secretariat or the IEA Member countries. We would like to thank Professor Felix Höffler for helpful comments and suggestions. All errors and omissions are the responsibility of the authors.

References

- Bakirtzis, A., Tellidou, A., 2006. Agent-Based Simulation of Power Markets under Uniform and Pay-as-Bid Pricing Rules using Reinforcement Learning. IEEE PES Power System Conference and Exposition.
- BNetzA, April 2011. Beschluss in dem Verwaltungsverfahren wegen der Festlegung zu Verfahren zur Ausschreibung von Regelleistung in Gestalt der Sekundärregelung, Az: BK6-10-098.
- 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 (3-4), 561–592.
- Bunn, D. W., Oliveira, F., 2007. Agent-based analysis of technological diversification and specialisation in electricity markets. *European Journal of Operational Research* 181 (3), 1265–1278.
- ENTSO-E, 29th of May 2013. Entso-e draft network code on electricity balancing. Version 1.22.
- ENTSO-E, August 2014. ENTSO-E Network Code on Electricity Balancing. Version 3.0.
URL https://www.entsoe.eu/Documents/Network%20codes%20documents/NC%20EB/140806_NCEB_Resubmission_to_ACER_v.03.PDF
- ENTSO-E, May 2017a. Consultation Report "FCR Cooperation".
URL https://www.entsoe.eu/Documents/Consultations/20170601_FCR_Consultation_Report.pdf
- ENTSO-E, Mar. 2017b. ENTSO-E Network Code on Electricity Balancing.
URL https://www.entsoe.eu/Documents/Network%20codes%20documents/NC%20EB/Informal_Service_Level_EBGL_16-03-2017_Final.pdf
- ENTSO-E, Mar. 2017. Survey on Ancillary Services Procurement, Balancing Market Design 2016.
URL https://www.entsoe.eu/Documents/Publications/Market%20Committee%20publications/WGAS_Survey_final_10.03.2017.pdf
- Fabra, N., Fehr, N.-H. v. d., Harbord, D., 2006. Designing electricity auctions. *RAND Journal of Economics* 37 (1), 23–46.
- Frederico, G., Rahman, D., 2003. Bidding in an electricity pay-as-bid auction. *Journal of Regulatory Economics* 24 (2), 175–211.
- Geanakoplos, J., Robert, A., Farmer, J. D., Howitt, P., Conlee, B., Goldstein, J., Hendrey, M., Palmer, N., Yang, C.-Y., 2012. Getting at systemic risk via an agent-based model of the housing market. *American Economic Review* 102 (3), 53–58.
- Hailu, A., Thoyer, S., 2007. Designing Multi-unit Multiple-Bid Auctions: An Agent-based Computational Model of Uniform, Discriminatory and Generalised Vickrey Auctions. *The Economic Record* 83 (1), 57–72.
- Heim, S., Götz, G., 2013. Do Pay-as-Bid Auctions Favor Collusion? Evidence from Germany's Market for Reserve Power. ZEW Discussion Papers No. 13-035 No. 13-035.
- Kahn, A. E., Cramton, P. C., Porter, R. H., Tabors, R. D., 2001. Uniform Pricing or Pay-as-Bid Pricing: A Dilemma for California and Beyond. *The Electricity Journal* 14 (6), 70–79.
- Kimbrough, S. O., Murphy, F. H., 2013. Strategic bidding of offer curves: An agent-based approach to exploring supply curve equilibria. *European Journal of Operational Research* 229 (1), 165–178.

- Krishna, V., 2002. Auction Theory. Elsevier Science (USA).
- Kutschinski, E., Uthmann, T., Polani, D., 2003. Learning competitive pricing strategies by multi-agent reinforcement learning. *Journal of Economic Dynamics & Control* 27 (11-12), 2207–2218.
- Müsgens, F., Ockenfels, A., 2011. Design von Informationsfeedback in Regelenenergiemärkten. *Zeitschrift für Energiewirtschaft* 35 (4), 249–256.
- Müsgens, F., Ockenfels, A., Peek, M., 2012. Balancing Power Markets in Germany: Timing Matters. *Zeitschrift für Energiewirtschaft* 36 (1), 1–7.
- Naghbi-Sistani, M. B., Akbarzadeh-Tootoonchi, M. R., Javidi-Dasht Bayaz, M. H., Rajabi-Mashhadi, H., 2006. Application of Q-learning with temperature variation for bidding strategies in market based power systems. *Energy Conversion & Management* 47 (11-12), 1529–1538.
- Rassenti, S., Smith, L., Wilson, B., 2003. Discriminatory Price Auctions in Electricity Markets: Low Volatility at the Expense of High Price Levels. *Journal of Regulatory Economics* 23 (2), 109–123.
- Vickrey, W., 1961. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance* 16 (1), 8–37.
- Viehmann, J., 2017. State of the German Short-Term Power Market. *Zeitschrift für Energiewirtschaft* 41 (2), 87–103.
- Waltman, U., Kaymak, U., 2008. Q-learning agents in a Cournot oligopoly model. *Journal of Economic Dynamics & Control* 32 (10), 3275–3293.
- Weidlich, A., Veit, D., 2008. A critical survey of agent-based wholesale electricity market models. *Energy Economics* 30 (4), 1728–1759.
- Xiong, G., Okuma, S., Fujita, H., 2004. Multi-agent Based Experiments on Uniform Price and Pay-as-Bid Electricity Auction Markets. *IEEE Conference on Electric Utility Deregulation, Restructuring and Power Technologies*.