Co-evolution of bidding strategies in power markets employing nodal pricing

Thomas Trötscher*

SINTEF Energy Research, Sem Sælandsvei 11, NO-7564 Trondheim, Norway

Abstract

This paper describes an agent-based simulation method for studying market power in network constrained, spot electricity markets. A genetic algorithm is used to evolve profit maximizing bidding strategies for the agents. Nodal pricing is used to clear the market; it is an approach which enables realistic modelling of the underlying power grid. The method is especially useful in identifying local market power due to transmission constraints which is difficult with conventional economic equilibrium models. Simulation outcomes are compared to Nash equilibrium analysis and previous research to give a better understanding of the results and build confidence in agent-based modelling.

Keywords: genetic algorithm, market power, electricity market, game theory, nodal pricing

*Tel.: +47 99245340; Email address:

Email address: thomas.trotscher@sintef.no (Thomas Trötscher)

Preprint submitted to Electrical Power and Energy Systems *December 15, 2009*

1. Introduction

Electricity markets around the world are moving from a centrally planned monopolistic structure to competitive markets. However, many previously state controlled generation companies (GENCOs) retained their large market shares in the deregulated market spurring concerns about market power. In addition to an oligopolist ownership structure, network constraints may give rise to market power specific to electricity markets. Assessment of market power potential has received much attention in the literature. [1] and [2] both give extensive reviews of methods for market power analysis in electricity supply. They distinguish between oligopoly equilibrium methods, production simulation models, and statistical methods. Additionally, experimental methods are reported in [3]. As both [1] and [2] conclude, the methods are complimentary and no single method has been accepted as the *de facto* assessment method for market power in electricity supply.

The plethora of analysis methods in use may arise from the difficulty of deriving analytical expressions for market power. Even for relatively small electrical networks with few participants it can be difficult to calculate a Nash Equilibrium (NE) for the power price. Also, for those simplified situations where analytical solutions can be computed, there may not exist a NE in pure strategies, leaving the interpretation of an equilibrium price open. The use of market concentration based methods also becomes difficult when the underlying network becomes congested such that several price areas arise.

Parallel to the more "traditional" methods presented above, several authors have focused on the use of agent based modelling and Genetic Algorithms (GAs) when modelling the power market. Taking a bottom up approach using evolving agents makes it possible to simulate strategic trading in complex markets. This has obvious advantages as realistic markets can be evaluated; the only requirement is that the economic dispatch of generators can be calculated – usually an easy exercise even for large power systems.

In the following, it is assumed that the reader is familiar with GAs[1]; for a good and plain introduction to the application of GAs in the simulation of strategic behavior in markets see [5]. [6] gives a theoretical treatment of *economical* GAs, that is a GA whose fitness function depends on the strategies of all players in the model. In particular, he showed that the population in an economical canonical[2] GA "[...] tends, over time, to move toward NE without fully reaching it."

Among the first to use GAs to evolve bidding strategies in power markets were [7]. Later applications of GAs to evolve strategically behaving agents in power markets include [8, 9, 10]. So far, focus has been on power markets with one price area. Experience from existing work seems to be that agents evolve plausible bidding strategies that can be observed in the real market. Closely related to the use of GAs is the use of agent based modelling with reinforcement learning to model market participants, see e.g. [11, 12]. The main distinction being that the learning process of the latter is not biologically inspired like a GA.

This paper will focus on the use of a real valued GA to evolve bidding

[1]For an introductory book on the subject see [4]. The term genetic algorithm is used for both binary and real valued encoding in this document.

[2]A canonical GA is the most basic GA and it is thoroughly described in the GA literature e.g. [4, Ch. 1, p. 10–14]

strategies for multiple players in electricity spot markets with a realistic description of the underlying power system. The objective of this paper is twofold:

1. *Establish a procedure to evolve bidding strategies for an arbitrary number of agents in power markets that employ nodal pricing.* Such a tool can enable authorities to identify market inefficiencies and investigate measures to improve competition. Among the possible measures that can be considered is the building of new transmission capacity.
2. *Compare simulation outcomes to classical NE analysis where possible.* The hope is to build confidence in agent based methods for the simulation of bidding strategies in power markets.

The remainder of this paper is organized as follows: Section 2 outlines the power market, the market participants, and the nodal pricing scheme. Section 3 describes the GA and the parametrized strategy function. Section 4 compares the results of GA simulation with NE analysis for several market set-ups ranging from simple well known games like Bertrand douploy to slightly more complex games. Finally, section 5 draws up conclusions and suggests model extensions.

2. The power market

This section describes the hypothetical power market, its participants, and the clearing mechanism used in the simulations. The market type presented is relatively realistic and has real life counterparts[3].

Only hourly contracts for physical delivery of electrical energy in the day-ahead (spot) market are considered. Spot trade is conducted at a Power Exchange (PX) and is obligatory for all buyers and sellers such that no bilateral trade can influence the power flow.

2.1. Participants

Each participant has one generator or one load. Generators are defined by their linear marginal cost (MC) functions; conversely, loads are defined by their marginal willingness to pay (WTP). The two expressions are mathematically identical, the formula for the MC is given in (1).

$$\begin{aligned} &MC_u(q_u) = i_u + s_u \cdot q_u \\ &\text{s.t.} \quad q_{u,min} \le q_u \le q_{u,max} \end{aligned} \tag{1}$$

Where, for a generator, $MC_u(q_u)$ is the cost of generating the q_uth MW with unit u during the hour in question. For a load, $WTP_u(q_u)$ is the willingness to pay for the q_uth MW of electricity. In general, the lower limit $q_{u,min}$ is zero for a generator and the upper limit $q_{u,max}$ is zero for a load such that negative q is consumption and positive q is generation. In the following discussion, all units will be regarded as generators.

The participants must auction their units into the market with linear bids. The bids may differ from the MC thus allowing the participant to bid strategically. A convenient and simple way of representing strategic bidding

[3]California, Texas, New York and New England all employ nodal pricing similar to the scheme presented here

is then to add a markup m to the marginal cost. The markup is a real value on some interval $[m_{min}, m_{max}]$ that should be selected such that the highest possible price in the market remains realistic[4]. The bid is then given by (2).

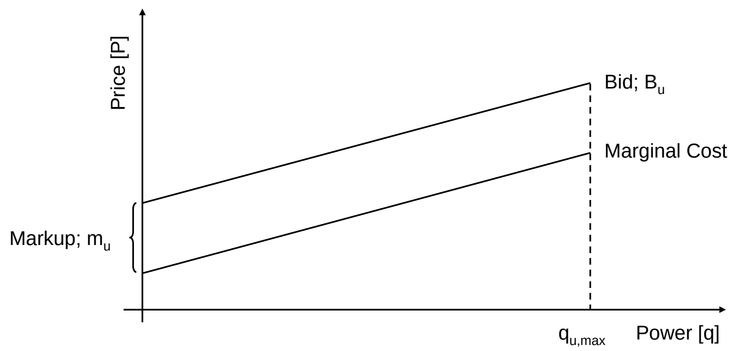

Figure 1: Marginal cost and markup of a generator unit

$$B_u(q_u) = MC_u(q_u) + m_u \tag{2}$$

Where $B_u(q_u)$ is the price per MWh its owner demands to generate q_u MW with unit u during the hour in question. Bids may be partly accepted, but each unit cannot generate more than its rated power $q_{u,max}$ and each load cannot consume more than its limit $q_{u,min}$.

That each participant only can own one unit is the gravest restriction imposed on the model. Further work will remove this restriction; still, several interesting results can be obtained with the model.

[4]A realistic price cap could be interpreted as the highest observed price, the price where emergency peak power becomes profitable, or the price where authorities likely will intervene.

2.2. Nodal pricing and market clearing

The PX collects all the market participant's bids as specified by (2). It is then the system operator (SO) who has to clear the market by finding the market cross point while taking transmission constraints into account. In other words, the SO must find the generation q_u by all generators that minimizes the system cost. With linear bids, this is a quadratic optimization problem as defined in (3) identical to the one used by [11].

$$\min \sum_{u} \left(i_u + m_u\right) q_u + \frac{1}{2} s_u q_u^2 \tag{3}$$

subject to

$$\begin{aligned}
& q_{u,min} \leq q_u \leq q_{u,max} && \forall\ u \in \Gamma \\
& \sum_{u \in \Gamma_k} q_u = \sum_{m \in \Omega_k} y_{km} \left(\theta_k - \theta_m\right) && \forall\ k \in \Omega \\
& y_{km} \left|\theta_k - \theta_m\right| \leq Q_{km,max} && \forall\ \{k, m \in \Omega | k \neq m\}
\end{aligned}$$

where: Indices k and m are nodal indices; Ω is the set of all nodes; Ω_k is the set of all nodes connected to k; Γ is the set of all units; Γ_k is the set of all units at node k; y_{km} is the admittance of the line connecting k with m; θ_m and θ_k are the nodal voltage angles; and, $Q_{km,max}$ is the flow constraint on the line from k to m.

The constraints are linearized approximations of the power flow equations[5] available in any standard textbook on power system analysis [6]. They represent a lossless electrical network and solutions are only approximate.

[5]Often called the DC power flow equations

[6]See for instance [13, Chapter 6.5] which is publicly available online

On the other hand, the linearized equations are faster to solve than the complete equations which is important because the problem is solved thousands of times in the GA. The power system simulation package Matpower [14] with the interior point solver BPMPD [15] was used to solve the power flow equations.

Solution of the problem in (3) gives the optimal dispatch of generators, i.e., generation q_u for all generators. Additionally, the nodal power prices, P_k, are inferred from the Lagrangian multipliers of the equality constraints; the Lagrangians are the costs of producing one additional MW of energy at node k.

2.3. Profit calculation

Generation units are paid the price at the physical point (the node/bus) at which power was fed into the network even if it was consumed at a node with a different price. Profit from trade between differently priced nodes – congestion rent – devolves to the SO. The profit of a unit $\Pi_u(q_u)$ is then:

$$\Pi_u(q_u) = P_k q_u - \int_0^{q_u} MC_u(q)dq = P_k q_u - \left(i_u q_u + \frac{1}{2} s_u q_u^2 \right) \tag{4}$$

Where P_k is the nodal price. Start and stop costs are not considered. The above reasoning also applies to loads as they can be regarded as negative generators.

3. Genetic Algorithm

The GA is used to evolve strategies for the market participants. Market participants are, in this context, those units that behave strategically in the market. In many cases, much of the load is completely inelastic and does not behave strategically.

Each market participant has its own pool of n strategies that are evolved using the GA. The aim of each participant is to maximize its profits Π_u. At each *generation* g of the simulation the participants meet several times, these are the *iterations*. At each iteration, each participant randomly selects a strategy from its pool of strategies. This strategy is evaluated to yield a bid B_u that is submitted to the market clearing algorithm. The market clearing returns participants profits that are stored in memory. After the desired number of iterations have passed, the fitness of a strategy, s, for unit, u, at generation, g, $F_{u,s,g}$, is calculated as the sum of profits it received divided by the number of times it was played. Thereafter, selection, crossover, and mutation are applied to produce the new generation of strategies and a new generation is started.

In pseudocode:

1. Initialize and load power system structure, $g = 0$
2. $g = g + 1$, $i = 0$
3. $i = i + 1$
4. For each participant u, select a random strategy s from the pool
5. Find the bid of each participant B_u, given the random strategy
6. Evaluate the market clearing algorithm given the bids of all players and store the achieved profits Π_u, s
7. If $i < numIterations$ goto line (3), else:
8. Calculate the fitness of each strategy $F_{u,s,g}$ as the average of all profits received this generation

9. Perform selection, crossover and mutation of all strategies
10. If $g < numGenerations$ goto line (2)
11. Simulation ends

The number of iterations is selected such that all strategies are played on average ten times each generation. In this paper, the standard number of strategies for each participant was $n = 50$. In other words, there were 500 iterations per generation. If a strategy – by chance – is not played at all, it receives a fitness of zero.

3.1. Decision space

The decision space is the set of possible decisions a strategy could turn up with. Possible dimensions in the decision space could be: the bid intercept i_u; the bid slope s_u; and, the bid quantity.

In this paper, a markup on the bid intercept, i_u, is used as the only dimension in the decision space. It is possible that other dimensions could give different results. At the same time, some experiments with alteration of the slope instead of the intercept have also been conducted. They suggest that simulation outcomes approach the same solution, but are less stable when the decision space is the slope.

The markup can take on any real value on a finite interval; another frequently used option is to let the agents select a markup from a set of discrete choices. As will be shown in section 4.1.3 this can have great implications for the simulation outcomes.

If only generators behave strategically, the markup can be restricted to positive values. However, if loads behave strategically they would want to bid below their willingness to pay; consequently, markup must be allowed to attain negative values.

3.2. Representation of one strategy

A strategy is the rule by which the participants determine their bid, in this case the markup. In mathematical terms, a strategy is a function that takes a set of inputs and returns a decision (markup). This function has parameters that determine the nature of the relationship between inputs and output; it is these parameters that the GA work upon.

Inspired by research on the Iterated Prisoner's Dilemma[7] (IPD), the inputs to the strategy are taken to be the markup and the price attained when the strategy was last played. The loose resemblance to an IPD is that the markup is ones own level of cooperation and the price corresponds to the cooperative level of the competitor(s). Using these two inputs, the strategy essentially has a memory of one round and is able to evolve tit-for-tat strategies and similar, i.e. *if you bid high last turn I also continue to bid high this turn.*

A convenient way of creating a parametrized function, that can represent a strategy as described above, is a multilayer feed-forward neural network. It can approximate any function arbitrarily well, given sufficient neurons.

[7]In the classical prisoners dilemma, there are two players each facing a choice between cooperation and defection. If both players defect, they both get a bad payoff; if both cooperate they get a good payoff; if one player defects and one cooperates then the defector is better off than if both cooperated and the cooperator is worse off than if both defected. The NE of this one-shot game is mutual defection, however, if the game is played iteratively mutual cooperation can be the outcome for many periods.

This approach of representing strategies has been used on the IPD, see e.g. [16, 17]. Figure 2 shows the structure of the neural network in abbreviated matrix notation. The input vector **P** consists of the price and markup last time the strategy was played, for convenience they are scaled to be in the range $[-1, 1]$; y is the output, also in $[-1, 1]$; S is the number of neurons in the hidden layer; **IW**, **LW**, **IB**, and **LB** are the weights and biases of the network. Matrix sizes are denoted below the corresponding matrices in italics.

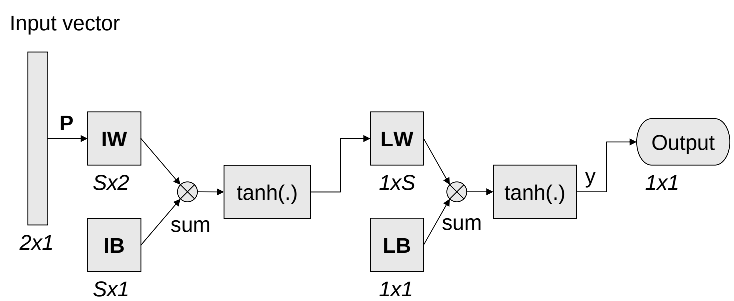

Figure 2: Feed-forward neural network structure

The neurons take their input, multiply it by a weight, add a bias and pass it through a hyperbolic tangent filter. The filter operates elementwise on matrices. Mathematically, the whole neural network can be represented by (5) which only requires a few arithmetic operations to be evaluated. To yield a markup, the output y is scaled according to (6)

$$y = \tanh\left(\mathbf{LW} \cdot \tanh\left(\mathbf{IW} \cdot \mathbf{P} + \mathbf{IB}\right) + \mathbf{LB}\right) \tag{5}$$

$$m = m_{min} + \frac{(y+1)\left(m_{max} - m_{min}\right)}{2} \tag{6}$$

The number of neurons in the hidden layer S is set to 10 in accordance with conclusions from [16]. The total number of weights and biases in the network are then 41[8]. These 41 floating-point numbers constitute the genes in one member of the population. At the beginning of a simulation the weights and biases are initialized randomly in the interval $[-1, 1]$.

The amount of information taken as input to the neural network could easily be extended to allow for more sophisticated strategies. One apparent extension would be to include information about results prior to the previous turn. However, care should be taken that all inputs to the strategy is information available to participants in the real market; inclusion of other participants bids or marginal costs would be unrealistic. As pointed out by [17], a strategy with two inputs and one output is easy to visualize. This is convenient when interpreting the strategies evolved by the agents and is part of the reason why only two inputs are used in this paper.

3.3. Crossover and mutation

This section briefly summarises the crossover and mutation operators that were used.

Crossover between strategies was performed within each pool, i.e. strategies from two different participants could not mate. Tournament selection was used to select the mating population. Crossover was arithmetic with a probability of 40% that parents were copied unmodified. The whole population went through mutation after crossover. Mutation was gaussian with a standard deviation of 0.5 and a probability such that each strategy on average suffered mutations in 2 of the 41 genes.

[8] 20 input weights, 10 input biases, 10 output weights and 1 output bias

3.4. Performance of participants

Whether participants perform better than in a perfectly competitive situation is a measure of the potential to exercise market power. A good performance measure is therefore needed to evaluate the simulation results. The performance measure used in this paper is the profit achieved by a participant divided by the highest profit, $\Pi_{u,best}$, that he or she could have achieved. This sounds intuitive, but requires some explanation as to what the highest possible profit of a player is.

The highest possible profit for a generator is the best profit it can achieve given that all other generators bid with the highest markup, m_{max}, and all loads bid according to their true willingness to pay.

The highest possible profit for a load is the best profit it can achieve given that all other loads bid with the lowest markup, m_{min}, and all generators bid according to their true marginal costs.[9]

The profit achieved by a participant is simply the average fitness of the strategies in its pool at the current generation. Thus, the performance measure, $\Psi_{u,g}$, of a unit is calculated as in (7):

$$\Psi_{u,g} = \frac{\frac{1}{n}\sum_{s=1}^{n} F_{u,s,g}}{\Pi_{u,best}} \tag{7}$$

[9]Higher or lower profits can arise if generators bid below their marginal cost or loads bid higher than their willingness to pay; however, profit seeking behaviour ensures that this never happens.

4. Simulation results

4.1. Single area

This section describes some examples of one bus markets, i.e., markets where all generators and loads are located at the same bus. This is equivalent to an uncongested network.

These simulations serve to show that the market simulator performs as expected in well-known situations like perfect competition and Bertrand duopoly.

4.1.1. Co-operative equilibrium

The first situation to be analysed is where two generators alone can serve exactly one half of a completely inelastic load due to capacity constraints. If one participant increases his markup, he or she improves the profit of both without chance of losing market shares to the other. Thus, the obvious best strategy is to bid with the highest markup possible.

As can be seen from figure 3, the generators quickly discover that they can bid as high as they want without loosing market shares; equilibrium is soon reached where at least one generator bids with the highest markup and thus fixes the price and profits of both participants at the maximum.

4.1.2. Perfect competition

Figure 4 depicts a situation that resembles a perfectly competitive market. Ten equal generators compete to serve the load which is completely inelastic. The capacity of the generators relative to the load is such that if one generator bids high it looses it's entire market share to the other participants. As can be expected the market price converges to competitive equilibrium.

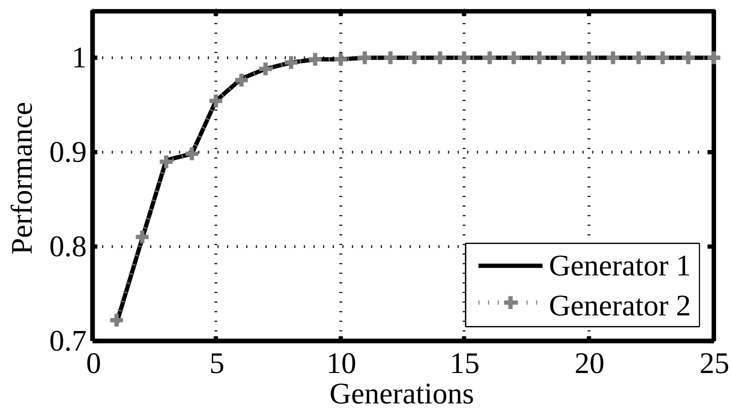

Figure 3: Performance of two generators that serve exactly half the load each

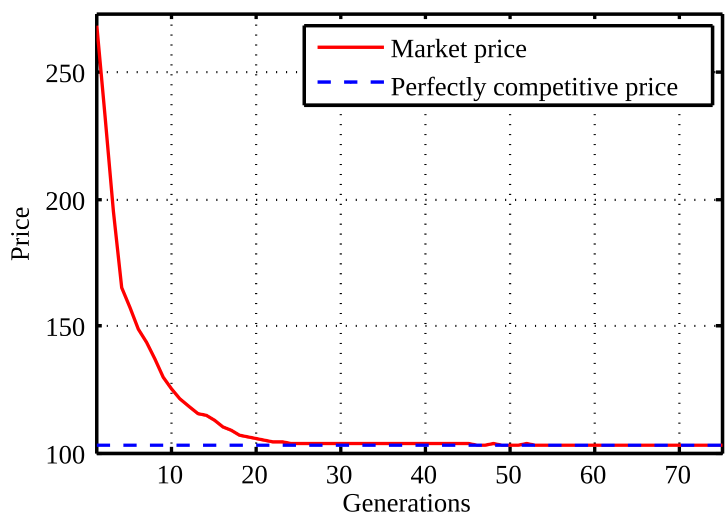

Figure 4: Market price compared to the competitive equilibrium price

4.1.3. Bertrand competition

In Bertrand duopoly, two generators with equal constant marginal costs compete to serve a load. The load does not behave strategically and the capacities of the generators are such that each one alone can serve the whole load at its zero profit price. In this game there is one NE where both generators bid according to their true marginal costs and receive zero profit just as in perfect competition.

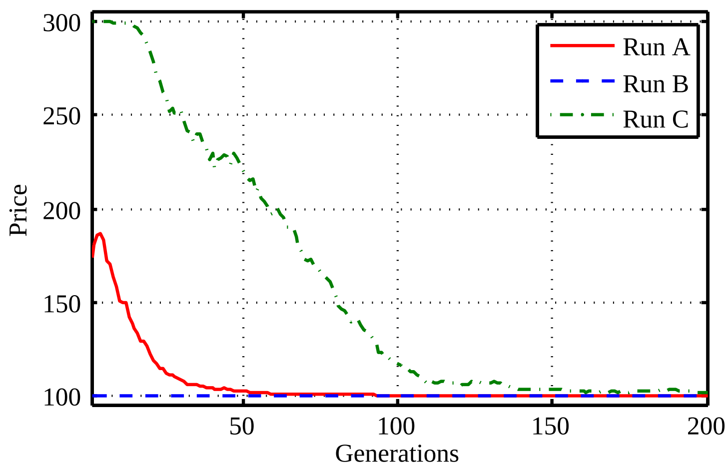

Figure 5: Market price in Bertrand competition for an initially random population (A), defecting population (B), and co-operative population (C)

Since there is no potential for market power in the Bertrand game, supernormal profit can only arise if the players cooperate by coordinating their prices. Figure 5 shows results of three simulation runs with different initial population. Run A has the weights and biases initialized randomly as normal; run B is initialized such that all members of the population start by defecting, i.e., bid at their true marginal costs; run C is initialized with a population of cooperators, i.e., all members bid with the highest possible markup. The generators have marginal costs of 100 which is also the perfectly competitive price level. In all three cases shown, competition eventually ensures that both generators bid at marginal costs even though convergence is slow when both generators start by co-operating. The simulations show that the players are unable to coordinate their prices.

Conceptually, Bertrand competition bears some resemblance to the Prisoners Dilemma: If both generators co-operate by bidding high they get a good profit, however, it is always in the interest of the other to defect and bid marginally lower, thus mutual defection is ensured. In the IPD game it is known that players can evolve strategies that maintain a mutually co-operative level even though the equilibrium for both players is to defect if the game is played only once. In the normal IPD, there are only two choices, either full cooperation or full defection. A common extension of the game is to include several discrete levels of cooperation. [17] noted that the chance of mutual cooperation dwindled as the possible choices in the IPD increased. Considering that each player in the Bertrand game can select any real positive number as her price (which can be thought of as the levels of cooperation), it is simply very unlikely to remain at a mutually cooperative level. This effect can explain why mutual cooperation is likely if the choice of each player is limited to only a few discrete price levels, but highly unlikely if the price can be selected from a continuum.

Other papers on agent based modelling [10, 11] limit the agents to select from a discrete set of choices. If the market situation resembles Bertrand competition, the above discussion might lead to think that mutual co-operation can occur, albeit with a small probability if the number of choices is large.

However, it turns out that the similarity of the Bertrand and Prisoners Dilemma game disappears if the number of choices is limited. In fact, the outcome of simulated "Bertrand" competition with discrete levels of cooperation is quite arbitrary and depends on the exact markups the players can choose from. An example serves to show this: Let $P = \{0, 2, 4, 6\}$ be the set of markups both players can select from. In this case there are three NE in pure strategies[10], but if instead $P = \{0, 2, 4, 8\}$ then there are four NE in pure strategies and one in mixed strategies. Furthermore, if $P = \{0, 2, 4, 9\}$ there are four NE in pure strategies and 3 NE in mixed strategies; finally if $P = \{0, 10, 15, 20\}$ there are two pure NE. Therefore, one can expect to see quite different outcomes from simulations of Bertrand competition for different sets of choices. Whether this also is the result for other types of competition has not been investigated. In any case one should be aware that the representation of the decision space can have wide-reaching implications for the results.

4.1.4. Bertrand with capacity constraints

With capacity constraints, one generator alone can not serve the entire market like in Bertrand competition. The participant that bids higher than its opponent still "captures" some residual demand. Suppose the same two generators as in the Bertrand example above, with constant marginal costs $MC_1 = MC_2 = 100$, serve a load with inverse demand function given by (8):

$$p(q) = 1000 - q \tag{8}$$

[10]Calculated using Gambit [18]

In contrast to the pure Bertrand game, both generators are constrained by maximum limits on generation $q_{1,max} = q_{2,max} = 540$. The participant that bids higher than the other then faces a residual demand given by (9):

$$q_u(p_u) = 1000 - 540 - p_u \tag{9}$$

In this one bus market, the price is set by the marginal unit accepted at the market. The profit of participant 1 can thus be expressed by (10) taking into account that the bid price $B_u = m_u + MC_u = m_u + 100$.

$$\Pi_1 = \begin{cases} (B_2 - MC_1) \cdot \min\{1000 - B_1, 540\} & \text{if} \quad B_1 < B_2 \\ (B_1 - MC_1) \cdot \frac{(1000 - B_1)}{2} & \text{if} \quad B_1 = B_2 \\ (B_1 - MC_1) \cdot \max\{1000 - 540 - B_1, 0\} & \text{if} \quad B_1 > B_2, B_2 < 460 \\ 0 & \text{if} \quad B_1 > B_2, B_2 \geq 460 \end{cases} \tag{10}$$

It is now possible to construct the reaction function[11] for participant 1 based on the cases in (10). As a starting point, consider the situation where $B_1 = B_2$: By inspection this is not a NE, both generators can benefit by diverging from this price. If the price where $B_1 = B_2$ is below some threshold, it is possible to realise a better profit by bidding higher and vice versa if the price is above the same threshold. For example, if $B_1 = B_2 = 500$ then profits $\Pi_1 = \Pi_2 = 400 \cdot 250 = 100\,000$, however if participant 1 bids slightly lower (by a margin δP) it captures its entire capacity and receives a profit

[11]The reaction function of participant 1 is how the optimal bid of participant 1 varies as a function of the bid by participant 2

$\Pi_1 = (400 - \delta P) \cdot 540 \approx 216\,000$. Similarly, if $B_1 = B_2 = 100$ then profits $\Pi_1 = \Pi_2 = 0$, however if participant 1 bids higher it can capture the residual demand at a profit $\Pi_1 > 0$, thereby also increasing the profit of player 2.

The first step in constructing the reaction function is to find the price threshold where a generator is indifferent as to which direction it should diverge as long as the opponent sticks to its price. The price of indifference occurs when undercutting the opponent gives the exact same profit as bidding higher than the opponent and capturing the residual demand. When bidding higher than the opponent, the best profit that can be achieved can be found by maximizing (11):

$$(B_1 - MC_1) \cdot \max\{1000 - 540 - B_1, 0\} = 560B_1 - B_1^2 \tag{11}$$

which occurs when:

$$B_1 = \frac{560}{2} = 280 \tag{12}$$

equating the profit of undercutting the opponent and the profit by bidding high we find:

$$\begin{aligned}(B_2 - MC_1)540 &= (B_1 - MC_1)(1000 - 540 - B_1) \\ B_2 &= \frac{(B_1 - 100)(460 - B_1)}{540} + 100\end{aligned}$$

substituting for $B_1 = 280$ yields the price threshold:

$$B_2 = 160 \tag{13}$$

When participant 2's bid is below the price threshold, the best action of participant 1 is to place a bid of $B_1 = 280$. When participant 2's bid is above the threshold, participant 1 should bid below participant 2 which can be ensured by bidding with a markup of zero. These two rules constitute the reaction function of player 1 and can be visualized as in figure 6. The reaction function of player two equals that of player 1 because costs and capacity limits are identical.

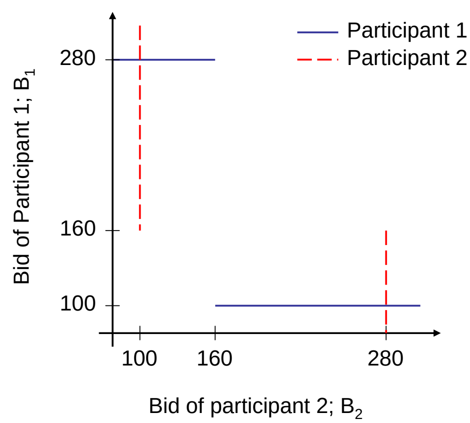

Figure 6: Reaction functions of participant 1 and 2. NE occurs where the two reaction functions intersect.

There are two symmetrical NE in this game which occur when one player bids 280 and the other bids below the threshold. The participant that ends up bidding high is unable to increase its profit by cutting the price and is worse off than the generator that bids low.

Figure 7 depicts the results of a simulation run for the example analyzed in this section. The bid of player 1 bumps around 280 while the bid of player 2 stays below the threshold of 160. In other words the simulation results are in accord with the NE analysis. Due to mutation, the values don't stabilize entirely; it is possible to achieve complete stabilization by utilizing a mutation step size reduction algorithm.

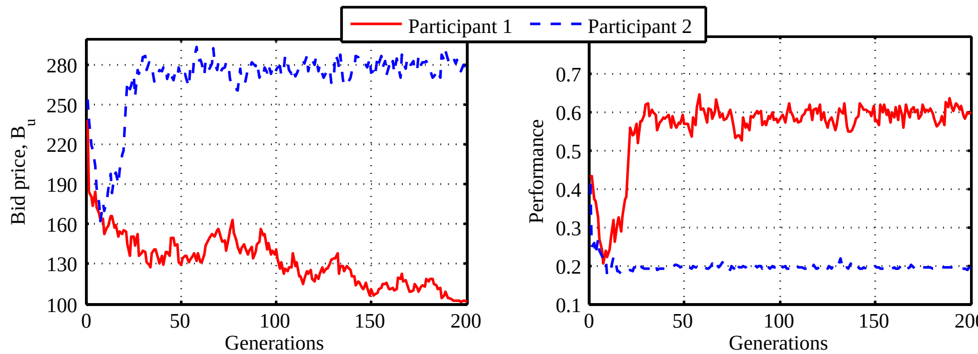

Figure 7: Simulation results for Bertrand competition with capacity constraints

Who ends up as the high bidder, and thus worse of, is entirely random. From figure 7 it can be seen that the two players undercut one another until, at the threshold price, player 1 starts increasing its markup.

4.1.5. Monopolistic load vs. monopolistic GENCO

So far, only participants owning generators have behaved strategically. This section considers a situation with a monopolistic GENCO that supplies a monopolistic load, e.g., a distribution company or an industrial customer. Let the true WTP and the true MC be defined by (14–15). Note that $s_d < 0$, $s_s \geq 0$, and $i_d > i_s$ to ensure a market cross point.

$$WTP_d(q) = i_d + s_d q \qquad \text{true WTP of the load} \qquad (14)$$

$$MC_s(q) = i_s + s_s q \qquad \text{true MC of the GENCO} \qquad (15)$$

Both the load and the GENCO behave strategically and want to maximize their profits by bidding strategically. Their bids, $B_d(q)$ and $B_s(q)$, are defined by (16–17):

$$B_d(q) = WTP_d(q) + m_d \qquad \text{bid of the load} \tag{16}$$

$$B_s(q) = MC_s(q) + m_s \qquad \text{bid of the GENCO} \tag{17}$$

The objective of the GENCO is to find the markup m_s that maximizes its profit. Since the optimal m_s is a function of the load markup, m_d, this process can be thought of as finding the reaction function of the GENCO. Being a monopolist on the supply side, the GENCO equates marginal revenue (MR) and MC to find the quantity and corresponding price and markup that maximizes profits:

$$MR = MC$$
$$\frac{B_d(q)q}{\mathrm{d}q} = MC_s(q)$$
$$2s_dq + i_d + m_d = s_sq + i_s$$

solving for q gives the monopolist quantity q^*:

$$q^* = \frac{i_s - i_d - m_d}{2s_d - s_s} \tag{18}$$

Next, finding the markup m_s that corresponds to the monopolist quantity q^*:

$$B_d(q^*) = B_s(q^*)$$
$$s_d\frac{i_s - i_d - m_d}{2s_d - s_s} + i_d + m_d = s_s\frac{i_s - i_d - m_d}{2s_d - s_s} + i_s + m_s$$

solving for m_s gives the reaction function of the GENCO:

$$m_s(m_d) = s_d\frac{i_d + m_d - i_s}{2s_d - s_s} \tag{19}$$

For the load, the derivation follows the same lines; the corresponding reaction function is the same as in (19), but with opposite subscripts. Since the reaction functions are straight unbroken lines there is one NE where the lines intersect.

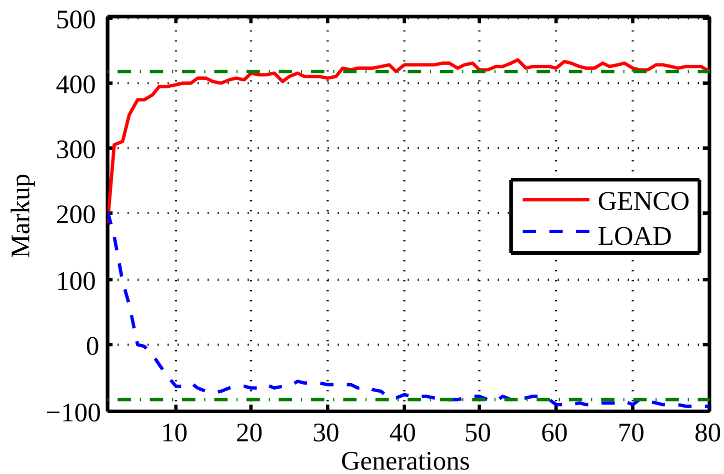

Figure 8: Markup for a monopolistic GENCO supplying a monopolistic load. The NE markups are shown with dotted lines.

A simulation was conducted for the following numerical example:

$$WTP_d d(q) = -q + 1000 \qquad \text{true WTP of the load}$$

$$MC_s(q) = 0.2q \qquad \text{true MC of the GENCO}$$

Plotting of the reaction functions reveal that there is one NE in pure strategies at $m_s = 416.67$ and $m_d = -83.33$. The simulated markups are shown in figure 8; apart from some noise from mutation the algorithm converges quickly to the NE. This example also shows how the monopolistic GENCO with relatively elastic MC is better of than the monopolistic load with relatively inelastic WTP. The result is a transfer of welfare from the load/consumer to the GENCO/producer.

4.2. Duopoly in a congested two bus network

So far, the underlying electrical network has been kept uncongested to enable analysis of some well known economic games. A congested network complicates the NE analysis and makes results less intuitive. In this section, the simplest possible constrained network will be analysed theoretically and simulated with the developed model. The network is depicted in figure 9. It consists of two buses with a generator and load each; the buses are connected by a transmission line capable of transferring K MW in either direction.

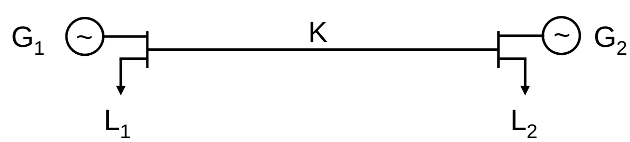

Figure 9: Two bus network with constrained transmission capacity

The participants owning the generators are assumed to be the only strategically behaving agents. Moreover, the transmission capacity K is less than the demand at either of the buses and the generators are capable of serving the load at their "own" bus and at the same time utilize the entire transmission capacity.

Accordingly, the generators face a strategic choice: (1) bid high to raise the price at its bus and serve the load that the opponent cannot capture because of the limited transmission capacity; or, (2) bid low and serve all the load it possibly can cover.

At first, this might seem like the Bertrand game with capacity constraints. However, there is one important difference: since the generators are paid the nodal price (and not a common system price), the low bidder can profitably raise its price towards that of the high bidder. As will be shown, the consequence is that there is no NE in pure strategies. In fact, the situation is very similar to Bertrand competition with capacity constraints in a *pay-as-bid* auction – a situation first analysed by [19], now known as the Edgeworth Paradox.

It is the intent to find the reaction functions of both participants in terms of markup. Let's begin with the supposition that the load is completely inelastic and that the generators have constant and equal MC and no capacity constraints[12]. The profit of player 1 can be expressed by (20):

$$\Pi_1 = \begin{cases} m_1(L_1 - K) & \text{if} \quad m_1 > m_2 \\ m_1 \min\{\frac{(L_1+L_2)}{2}, L_1 + K\} & \text{if} \quad m_1 = m_2 \\ m_1(L_1 + K) & \text{if} \quad m_1 < m_2 \end{cases} \tag{20}$$

where L_i is the load on bus i. If the transmission capacity K was zero, each participant would be a monopolist at its bus and fix the markup m_u at the highest possible value m_{max}. But with a positive transmission capacity, it is in the interest of both participants to lower their price by a little to capture some of the load at the other bus. Just as in Bertrand competition, the competitors undercut one another in steps, always with the aim of receiving the larger portion of the demand. This process may continue until the limiting markup m_u^* is reached where the profit of undercutting the rival equals the profit of

[12]The following deduction of the reaction function is less cumbersome with these assumptions, but ends up with qualitatively the same results as with relaxed assumptions

bidding with the highest markup. For player 1, the following holds:

$$m_{max}(L_1 - K) = m_1^*(L_1 + K)$$

solving for m_1^* yields:

$$m_1^* = m_{max}\frac{L_1 - K}{L_1 + K} \tag{21}$$

The reaction functions are illustrated in figure 10. Starting at point O, the players move down the diagonal as they undercut one another, until, at the point U, it is in the interest of one of the players to swap strategy and bid with m_{max}. The rival then follows suit, bids high and they are back at point O and the whole process starts anew.

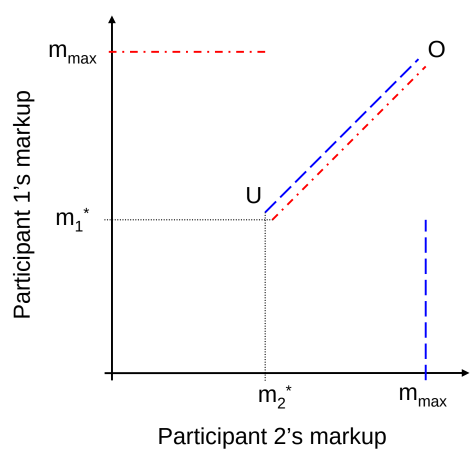

Figure 10: Reaction function of players in the two-bus market, adapted from [19].

Varying the capacity K of the transmission line can create game situations ranging from pure Bertrand competition to pure monopoly. Not counting the extremes, there is no NE in pure strategies regardless of K. However, the price will be confined to a small interval near the maximum price for small transmission capacities. For large capacities, m_u^* will be close to zero and the price will fluctuate in a larger interval.

Four simulations were conducted with the following parameters:

$$L_1 = L_2 = 300 \qquad \text{Load at the buses}$$

$$m_{max} = 200 \qquad \text{Highest allowable markup}$$

$$K = \{10, 50, 150, 300\} \qquad \text{Transmission capacity}$$

The four simulations were identical except for the transmission capacity. Results are shown in figure 11; from top to bottom in the diagram the simulation runs range from low to high transmission capacity. The corresponding limiting markups for the four cases, calculated from (21), are $m_u^* = \{187.1, 142.9, 66.0, 0\}$.

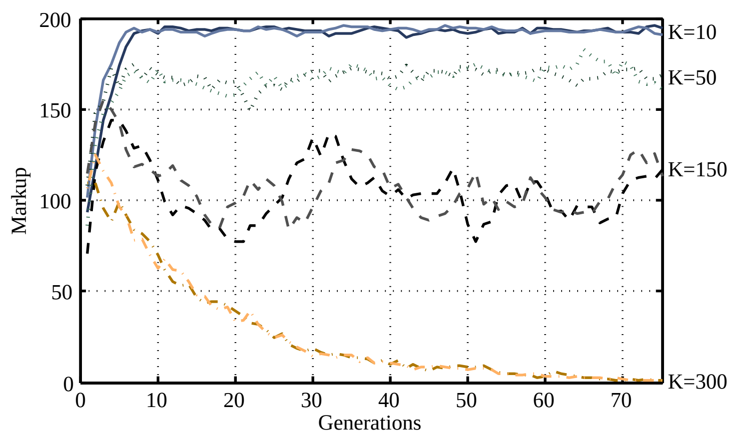

Figure 11: Simulated markups in the two-bus network for 4 different transmission capacity levels K

Contrary to what might be expected, the simulated markups do not span the whole range from m_u^* to m_{max}. This may partly be attributed to the averaging that occurs when the participant markup is calculated from the whole population of markups. However, part of the reason is also due to the mutation scheme used; when two competing generators reach point U in figure 10 it is simply very unlikely that a strategy is mutated to accommodate a “jump” up to point O. Instead, the participants slowly move up towards O along the diagonal, but before reaching that far they enter a new cycle of undercutting one another – again approaching U. The resulting behavior is the cyclical motion seen in figure 11.

4.3. Duopoly in a meshed grid

This section presents results from an analysis of a small benchmark power system (figure 12) developed by [11] for comparison of NE analysis and agent-based modelling. There are 3 strategically behaving generators in the system and all load is assumed to to be completely inelastic. All transmission lines are assumed to have sufficient transmission capacity with the exception of the line from bus 2 to bus 5 which is limited to 100MW. Table 1 gives the generator parameters.

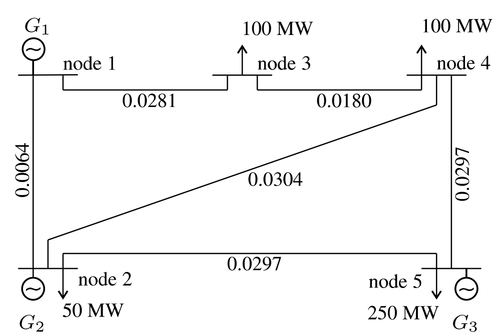

Figure 12: Power system diagram. The numbers close to the lines are reactances in per unit. Source: [11]

Table 1: Generator data

u	$q_{u,max}$	i_u	s_u
Generator 1	300	10	0.02
Generator 2	300	10	0.02
Generator 3	250	20	0.04

Krause et. al. [11] analytically computed the NE and simulated the system using adaptive agents with decisions restricted to a discrete set of markups. Two symmetrical NE were found; generator 3 should always bid its highest markup while either generator 1 or 2 bids with the highest markup and the other bids its marginal cost. The agent-based approach did discover the NE, but alternated between the two equilibria in a cyclical fashion. Here, the markups can be any real number on the interval $[0, 30]$, otherwise the simulation setup is identical to that of [11].

Simulations with the model presented in this paper confirm the results found by [11], however no cyclical behavior occurs. As soon as the agents have locked in on one NE the game stays in that equilibrium for the rest of the simulation. To discover several equilibria, the model must be rerun. The random nature of the model ensures that both NE are found after a couple of reruns. Experience has shown that this is the case for all systems with multiple NE that have been analysed, at the same time this has not been proved or confirmed in general.

Figure 13 depicts the simulation results. On the left are the simulated markups, on the right the corresponding consumer benefit, producer benefit, and congestion rent is shown. Since the load is completely inelastic, a maximum WTP must be assumed in the calculation to avoid infinite consumer benefit[13]. As can be seen from figure 13, the producers (generator 1, 2, and 3) take the lions share of the total benefits. The result can be compared to row one of table 2 where the socially optimal solution is reported, i.e. when all markups are zero. The total social benefit is approximately equal, but there has been a huge transfer from consumers to producers due to market power.

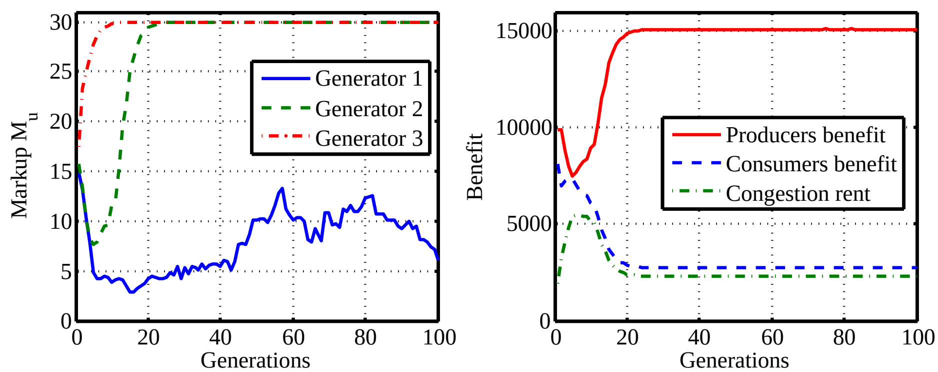

Figure 13: Results with 100 MW capacity on the line from bus 2 to bus 5.

In this case, it is evident that the producers exert market power, hence authorities might intervene to try and remedy the situation. One of the possible measures could be to double the transmission capacity on the congested line between bus 2 and 5. Figure 14 shows the situation with 200MW capacity on line 2–5; market power is not completely eradicated, but the new

[13]As the demand schedule becomes more and more inelastic the area underneath approaches infinity. The maximum WTP used was 55.811, corresponding to the highest possible market price

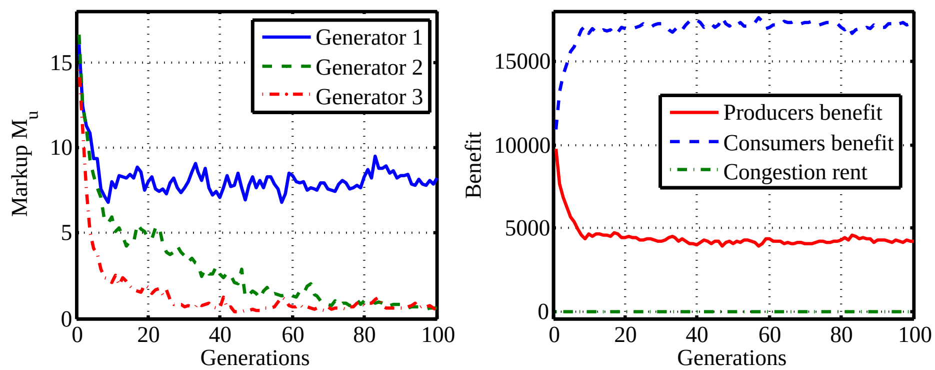

Figure 14: Results with 200 MW capacity on the line from bus 2 to bus 5.

equilibrium is much closer to the social optimal equilibrium which is shown in row 2 of table 2.

Table 2: Social optimal congestion rent, consumer benefit, and producer benefit

Capacity line 2–5	Prod. benef.	Cons. benef.	Cong. rent	Total
100MW	1056	17322	2019	20397
200MW	1250	20406	0	21656

5. Conclusion

In this paper, a method for the simulation of strategic behaviour in power markets employing nodal pricing has been developed. The method employs adaptive agents that try to maximize their profits by bidding strategically in the spot market. The agent's strategies are represented by neural networks that translate information from the previous turn of the game into a markup. The markup is added to the marginal cost function of the agent to yield a bid. Over the course of a game, the strategies are evolved using a genetic algorithm. This approach has the advantage that complex market situations – with many strategically behaving agents and a congested power grid – can be analysed.

To build confidence in the model, several well known games were analysed. Nash equilibria (NE) in pure strategies were derived using reaction functions and compared to simulated outcomes. The method proved to converge to the NE in pure strategies in all cases. If several pure strategy NE existed, the algorithm randomly converged to one of the NE.

The Bertrand duopoly in particular was analysed to see if co-operative behaviour can occur when there is no market power, i.e. as a result of the players co-ordinating their behaviour. This was disproved as simulations with populations of only co-operators eventually converged to the competitive Bertrand equilibrium. Furthermore it was found that restricting the strategies to a discrete set of markups creates game situations with quite different NE from the original Bertrand game. Thus, the simulation outcomes can be highly dependent on the choice of strategy representation. Whether this is the case for other competitive game situations was not analysed.

In a game set-up with no pure strategy NE the simulated prices exhibited cyclical behavior. The situation was caused by a transmission line with limited capacity between two buses. If the capacity limit was removed, the situation was reduced to Bertrand competition; if the capacity was set to zero, the situation was reduced to pure monopoly. In between these two extremes the price fluctuated in a interval given by the capacity of the line. The observed behavior is also found in other markets; the absence of a pure strategy NE in these market situations is called the Edgeworth Paradox.

As a final exercise, a small benchmark power system from a previous paper on adaptive agents in electricity markets was analysed. Simulation results were in line with the previous study and showed that the suppliers in the benchmark system exerted market power. Moreover, the market power was due to transmission constraints on one of the lines and was only detected because a proper model of the underlying network was used. Further simulations showed that market power could be mitigated to a great extent by building new transmission capacity.

Further work on the model will aim to enable each market participant to own several generation and load units dispersed at different nodes in the network. Additionally, it would be of interest to investigate whether different strategies are evolved if more information is available to the participants. With these improvements the method can give insights into the effect of various measures aimed at mitigating market power in electricity markets.

[1] E. Kahn, Numerical Techniques for Analyzing Market Power in Electricity, The Electricity Journal 11 (6) (1998) 34–43.

[2] K. David, A. F. Wen, Market power in electricity supply, Energy Conversion, IEEE Transactions on 16 (4) (2001) 352–360.

[3] J. Bernard, R. Zimmerman, W. Schulze, R. Thomas, T. Mount, R. Schuler, Alternative Auction Institutions for Purchasing Electric Power: An Experimental Examination, in: Proceedings of Bulk Power Systems Dynamics and Control IV, 1998.

[4] D. E. Goldberg, Genetic Algorithms in Search, Optimization and Machine Learning, Addison-Wesley Longman Publishing Co., Inc. Boston, MA, USA, 1989.

[5] T. C. Price, Using co-evolutionary programming to simulate strategic behaviour in markets, Journal of Evolutionary Economics 7 (3) (1997) 219–254.

[6] T. Riechmann, Genetic algorithm learning and evolutionary games, Journal of Economic Dynamics and Control 25 (6-7) (2001) 1019–1037.

[7] C. W. Richter, Jr, G. B. Sheble, Genetic Algorithm Evolution of Utility Bidding Strategies for the Competitive Marketplace, Power Systems, IEEE Transactions on 13 (1) (1998) 256–261.

[8] D. Menniti, A. Pinnarelli, N. Sorrentino, Simulation of producers behaviour in the electricity market by evolutionary games, Electric Power Systems ResearchArticle in press: doi:10.1016/j.epsr.2007.04.005.

[9] P. Attaviriyanupap, H. Kita, E. Tanaka, J. Hasegawa, New bidding strategy formulation for day-ahead energy and reserve markets based on evolutionary programming, International Journal of Electrical Power and Energy Systems 27 (3) (2005) 157–167.

[10] A. J. Bagnall, G. D. Smith, A multiagent model of the UK market in electricity generation, Evolutionary Computation, IEEE Transactions on 9 (5) (2005) 522–536.

[11] T. Krause, E. V. Beck, R. Cherkaoui, A. Germond, G. Andersson, D. Ernst, A comparison of Nash equilibria analysis and agent-based

modelling for power markets, International Journal of Electrical Power and Energy Systems 28 (9) (2006) 599–607.

[12] J. Bower, D. Bunn, Experimental analysis of the efficiency of uniform-price versus discriminatory auctions in the England and Wales electricity market, Journal of Economic Dynamics and Control 25 (3-4) (2001) 561–592.

[13] G. Andersson, Modelling and analysis of electric power systems, Lecture 227-0526-00 Swiss Federal Institute of Technology Zürich, available online at: `http://www.eeh.ee.ethz.ch/downloads/academics/courses/227-0526-00.pdf` (2006).

[14] R. D. Zimmerman, C. E. Murrillo-Sanchez, Matpower, Version 3.2, User's Manual, Power System Engineering Research Center, Cornell University, Available online at: `http://www.pserc.cornell.edu/matpower/` (2007).

[15] C. Mészáros, Fast Cholesky factorization for interior point methods of linear programming, Computers and Mathematics with Applications 31 (4-5) (1996) 49–54.

[16] P. G. Harrald, D. B. Fogel, Evolving continuous behaviors in the Iterated Prisoner's Dilemma, BioSystems 37 (1996) 135–145.

[17] P. Darwen, X. Yao, Co-evolution in iterated prisoner's dilemma with intermediate levels of cooperation: Application to missile defense, International Journal of Computational Intelligence and Applications 2 (1) (2002) 83–107.

[18] R. McKelvey, A. M. McLennan, T. L. Turocy, Gambit: Software tools for game theory, version 0.2007.01.30, Available online at: http://gambit.sourceforge.net (2007).

[19] F. Y. Edgeworth, La teoria pura del monopolio, Giornale Degli Economisti 40 (1897) 13–31.