Revenue Equivalence Revisited: Bounded Rationality in Auctions

Konrad Richter
University of Kiel
e-mail: konrad_richter@mckinsey.com

April 30, 2004

Abstract

The Revenue Equivalence Theorem states that for a variety of auction formats, expected seller revenue and the final allocation of goods are the same. The proof relies critically on perfect rationality of bidders. This paper considers repeated open bid first and second price auctions with heterogenous bounded rational bidders. It shows by mathematical and computational methods that the equivalence breaks down when bidders play their best responses to past play. In this case, expected seller revenue differs between auction formats and revenue volatility is persistently higher in first than in second price auctions. The consequences are twofold: First, the literature on auction theory should be revised to integrate the consequences of bounded rationality. Second, many auction formats in the real world could be replaced by auctions like second price auctions where bidding the true value is the dominant strategy. This would significantly reduce risk in the economy.

1 Introduction and Outline

1.1 Introduction

Auctions, though existent in the economy since thousands of years, have increasingly gained importance throughout the last decades. Nowadays, goods are allocated via auctions in a multitude of economic settings.
Most of current auction theory assumes that each bidder makes her decisions perfectly rational and knows that each other agent also does so. Under this assumption, bidders bid according to their Nash Equilibrium (NE) bidding functions. The most important theorem in auction theory is the Revenue Equivalence Theorem (RET).

The RET basically predicts that in the NE all auction formats are equivalent in two major ways:[1]

- First, the auctions result in the same allocation of goods, independent of the specific format. For auctions where the seller doesn't post a reserve price, the allocation is pareto-optimal.
- Second, expected seller revenue is the same for a wide range of formats.

According to the RET, first price auctions, higher price auctions like second or third price auctions and all-pay auctions are equivalent. Thereby, the theory justifies the observeable coexistence of different auction formats in the economy.
Today, most of current research still focuses on the investigation of Nash Equilibria in auctions [14], [16]. The main interest is to assess the effects of relaxing the assumptions of the RET.

The NE concept is a great tool to assess the outcome of games. However, it has several weaknesses [4]. First, there is experimental and empricial evidence that players deviate in many situations significantly from NE play. Second, the concept gets logically inconsistent if the thinking cost for finding the optimal solution to a game are properly taken into account. Third, we can not use the NE concept to assess the outcome of games if some players incidentally deviate from perfect rational play. In particular, it is not clear, whether players with initially heterogenous strategies will ever converge to NE play. Therefore, it is a much more natural assumption to investigate games under bounded - or limited - rationality. There, bidders use rules to select their strategies and update these rules according to some prespecified learning scheme. This approach allows us to model learning in games.[8]

Under best response dynamics, players try to maximize their expected payoff by playing the strategies that would have generated the highest payoff in the past. Considering the times at which players update their strategies, the updating can be myopic when bidders consider only information from the last round to update their strategies. Alternatively, they can have a long memory so that they base their strategy updating on more than one round or even the entire game.

My paper focuses on repeated open bid auctions under best response dynamics (fictituous play). It shows with mathematical and computational methods that bidders under best response play fail to find the NE bidding functions in first price auctions. In contrast, bidders in second price auctions converge easily to NE play.

What is the intuition behind investigating best response dynamics? To which economic situations can we apply our findings - and to which not? To start with, best response

[1] In detail, the assumptions of the RET are:

- For the bidders: Private independent values, identical commonly known value distributions, risk neutrality, no budget constraint, perfect rationality
- For the different auction formats: Equal participation cost, equal probability of winning with the highest bid

dynamics is *not* a good model to describe experimental auctions. It is observed there that experimental subjects significantly oberbid in first price auctions [5], [9], [13]. This behavior is commonly explained by assuming a quantal response model of learning. There, players play strategies in their strategy set with a probability that is proportional to the payoff that these strategies would have generated in past auctions taken to the power of some exponent γ. Like the best response dynamics, also quantal response can be myopic or with long memory.[2]
Quantal response play with low γ seems to be a good model to describe experiments with inexperienced bidders where only little profit is at stake. However, as the stakes increase, also the significance of the gains and losses - modeled by the parameter γ - increases. Therefore I claim that best response play - which is the limit case of $\gamma \to \infty$ - is a good model to describe economic environments where huge profits at stake. Professional economic agents like firms or traders in markets try to outperform their competitiors by the use of research departments, consultants and expert tools. This behavior can be modeled best by assuming best response dynamics.

The mathematical analysis shows that under best response dynamics, myopic bidders in first price auctions bid significantly below the NE prediction. As the memory strength of players increases from myopic to long memory play, the seller revenue in first price auctions also increases. I explicitely show how NE play emerges for players with perfect memory or infinite sophistication.
Bounded rational players with limited memory capabilities, however will deviate from NE play. A persistent volatility of bidding strategies accompanies this deviation. It comes from the permanent mutual adaptation of bidders and results in riskier asset prices than necessary. In contrast to first price auctions, players in second price auctions learn to bid according to their NE functions even under myopic play. The main conclusion therefore is that the prices of goods that are allocated via second price auctions are less volatile than in first price auctions and that the goods always go to the bidder who values it the most. Therefore the allocation of goods is maximal efficient in second price auctions.

The quantitative predictions of the mathematical analysis are replicated by computational model runs. The Auction Simulator is an agent based model for the investigation of different auction formats under various learning schemes. The double-check of computational results with mathematical predictions shows that the model works correct. Therefore we can use it as a starting base to investigate questions that can not so easily be tracked by mathematical methods. Generalizations are for instance the mutual adaption of agents, the participation of more than 2 bidders, arbitrary value distributions for bidders or statistical analyses of price and strategy time series. The comparison of simulated data with statistical features of empricial data like heteroscedasticity or leptokurtosis will allow the 'reality check' of different models of bounded rationality for different economic systems in the future.

In summary, mathematics and computation show that revenue volatility under best response play is persistently higher in first than in second price auctions. In generalization, the results

[2] In another paper I have modeled this behavior for an exponent of $\gamma = 1$ and shown analytically and computationally that myopic quantal response leads to the experimentally observed overbidding while for sufficiently long memory it leads to underbidding.

of this paper indicate that excess volatility of revenue is a property of any auction format that gives bidders incentive to shade their true values. Such value shading designs of repeated auctions can be found in many economic systems: Order books of financial or electricity markets, the emerging market for greenhouse certificates or supply chains to name but a few.

In conclusion, the paper suggests to take a closer look at the auction mechanisms that are implemented in the economy. Since there are a lot of additional possible constraints on the efficiency of auctions such as for instance the prevention of collusion, recommendations on the optimal auction mechanism must depend on the specific system. However, I claim that volatility in many markets could be substantially lowered by replacing value shading auction designs by equivalents of second price auctions. This would result in a better allocation of goods and in a decrease of risk for the economy as a whole.

1.2 Outline

The paper is structured as follows:
Chapter 2 mathematically analyzes the bids of 2 payoff-maximizing bidders who try to find their best responses in open bid auctions. It is explicitely derived how a bidder can calculate her best response against her opponents play if she is myopic or has perfect memory.
Chapter 3 introduces the Auction Simulator(AS). This agent based simulation correctly reproduces the mathematical findings of chapter 2. I use the AS to assess the impact of relaxing assumptions of the mathematical analysis.
Chapter 4 analyzes the implications of our findings for the real economy and gives an outlook for future research directions.

The paper is completed by several appendices and a bibliography.

2 Best Responses in Open Bid Auctions

2.1 Introduction to Auction Theory

2.1.1 The Standard Model

In the symmetric independent private values framework (SIPV) of auction theory, two standard auctions are distinguished: the 1^{st} price auction and the 2^{nd} price auction.[3] In the theoretical standard setting, each bidder has a private value. Values are drawn from a common random distribution that is known to each bidder. Each bidder knows her own but not

[3] A range of other formats is also considered in the literature: 3^{rd} and higher price auctions, all-pay auctions, etc. However, for our theoretical analysis we will focus on first and second price auctions.

the other players' values. Dependending on her value, each bidder decides on a bid which she secretly reports to the auctioneer. The assignment of a bid to each possible value is called the bidding strategy. The object under auction goes to the bidder who submits the highest bid. In the first price auction, the winner has to pay her own bid, in the second price auction the second highest bid.
In sealed bid auctions, only the winning bid is published after the auction while in open bid auctions every bid, also the loosing ones are published. In the one-shot auction, the publishing of information after the auction does not influence the bidding strategies of players who have to decide on their bids before the auction. However, in a repeated setting, open and sealed bid auctions are not equivalent because they induce a different informational structure. We will focus in this paper on repeated open bid auctions.[4]

2.1.2 Revenue Equivalence in the Standard Model

The revenue eqivalence theorem (RET) was first formulated by Vickrey [19] and then generalized independently by Myerson [17] and Riley and Samuelson [18]. It states that all auction forms yield the same expected seller revenue and allocation of goods as long as the following conditions are met:

- The item goes to the bidder who submits the highest bid
- The cost of submitting the lowest feasible bid is the same for the different auction formats
- Bidders have private independent values
- Bidders' values are drawn from the same distribution and this distribution is known to each bidder
- Bidders are risk neutral
- Bidders have no budget constraint
- Bidders are perfectly rational, they know that all the other players are also rational, and they know that all bidders are symmetric to themselves. This implies that all bidders bid according to their Nash Equilibrium bidding functions

The last point is often not explicitely mentioned because it is implicitely assumed in most of game theory. However, as I will show, it is crucial for revenue equivalence to hold: If we drop the assumption that each bidder knows that each other bidder is perfectly rational, we arrive at the best response dynamics. In this chapter I show that then the RET breaks down.

[4] I claim that under best response dynamics, repeated sealed bid auctions have even worse convergence to the NE than their open bid equivalents. The reason is that each bidder has less information available to calculate her best response.

2.1.3 Deviating from the Standard Model

The derivation of Nash Equilibria implies that players are a priori homogenous in that they play symmetric strategies. But can bidders find the NE bidding strategies if we don't restrict them to a priori symmetric solutions? What is the phase space dyamics of the system? Will it eventually converge to the symmetric Nash Equilibrium?

To answer these questions, I consider a repeated two player open bid auction. My analysis builds on the SIVP: I consider 2 risk neutral bidders without budget constraints. In each round, bidders' private values are drawn from a common distribution. To be specific, I assume a uniform distribution $v_i \sim U(0,1)$.The NE in first price auctions is then bidding half the value: $b_i = v_i/2$. In second price auctions, the NE is bidding the true value, $b_i = v_i$. These assumptions are all in accordance with the assumptions of the standard model, however I deviate in two aspects:

The main difference to the SIPV is that I consider heterogenous bidders under best response dynamics: Bidders stay rational in that they try to optimize their payoff by playing the best response to their opponents play. However, in contrast to the standard assumption of game theory, bidders do not necessarily believe that their opponents are perfectly rational. This implies that bidders can no more follow the reasoning that leads them to play NE strategies.

The second difference is that I restrict players to linear bidding functions $b_i(v_i) = \beta_i \cdot v_i$ where $\beta_i \in (0,1)$. This means that players have to decide on their strategy $\beta_i^{(\mathrm{t})}$ upfront before their value $v_i^{(\mathrm{t})}$ is drawn. This prevents the (more realistic) modeling of players who bid more aggressively for small values $v_i^{(\mathrm{t})}$ than for large ones. However, the purpose of this paper is to show that in first price auctions players fail to converge to NE play. In this sense I reduce the infinite-dimensional search space of bidding functions $b_i(v_i)$ for each agent to the one-dimensional space $\beta_i \in (0,1)$ that includes the NE bidding function $\beta_i = \frac{1}{2}$. If agents can not even find the NE in this space, they will certainly be unable to find it in the infinite-dimensional original one. Therefore, this restriction does not weaken the conclusions on exess volatility. On the contrary, it makes them even more pronounced.

2.1.4 Plan of this Chapter

In this chapter I analyze the best response dynamics in 2 player open auctions along the following steps:

- 2.2 identifies the best response function of player 0 in first price open bid auctions (1POBAs)[5] $\beta_0^{\mathrm{br}}(\beta_1, v_0, v_1)$ if all parameters are known to her. I use the result to quantify the average strategy of player 0 under myopic best response play.
- 2.3 identifies the best response function $\beta_0^{\mathrm{br}}(\beta_1)$ for v_0 and v_1 both $\sim U(0,1)$; β_1 fixed. This is the best response against a known strategy if neither the own nor the opponent's value are known upfront. I use the result to analyze the average strategy of player 0 under long memory best response play.

[5] All abbreviations are found in the appendix

- 2.4 analyzes the best response dynamics in second price auctions. I show that under best response dynamics each player immediately learns the NE strategy of bidding her true value.
- In 2.5 I conclude that the RET emerges from the best response dynamics if bidders have infinite memory strength or are infinitely spohisticated. Under realistic assumptions however the RET breaks down.

2.2 Unsophisticated Myopic Best Response in 1POBAs

2.2.1 The Payoff Function

The probability for player 0 to win a 1POBA is given by

$$p_0^{win} = \begin{cases} 1 \text{ if } \beta_0 v_0 > \beta_1 v_1 \Leftrightarrow \text{ if } \frac{\beta_0 v_0}{\beta_1 v_1} > 1 \\ 0 \text{ if } \beta_0 v_0 < \beta_1 v_1 \Leftrightarrow \text{ if } \frac{\beta_0 v_0}{\beta_1 v_1} < 1 \end{cases}$$

So, the winning probability can be rewritten as

$$p_0^{\text{win}} = \Theta(\beta_0 v_0 - \beta_1 v_1)$$

For details on the Θ function please check Appendix A.
Player 0's payoff is therefore given by

$$PO_0(\beta_0, \beta_1, v_0, v_1) = v_0(1 - \beta_0)\Theta(\beta_0 v_0 - \beta_1 v_1) \tag{2.1}$$

The following graph shows the payoff of player 0 for different choices of β_0 when β_1, v_0 and v_1 are fixed.

[INSERT GRAPH 'PAYOFF OF STRATEGIES']

The best response function of player 0 if she knows β_1, v_0 and v_1 is therefore given by

$$\beta_0^{\text{br}}(\beta_1, v_0, v_1) = \begin{cases} \lim_{\epsilon \to 0} \frac{\beta_1 v_1}{v_0} + \epsilon = \frac{\beta_1 v_1}{v_0} \text{ if } \frac{\beta_1 v_1}{v_0} < 1 \\ \text{arbitrary} \in (0, 1) \text{ if } \frac{\beta_1 v_1}{v_0} > 1 \end{cases}$$

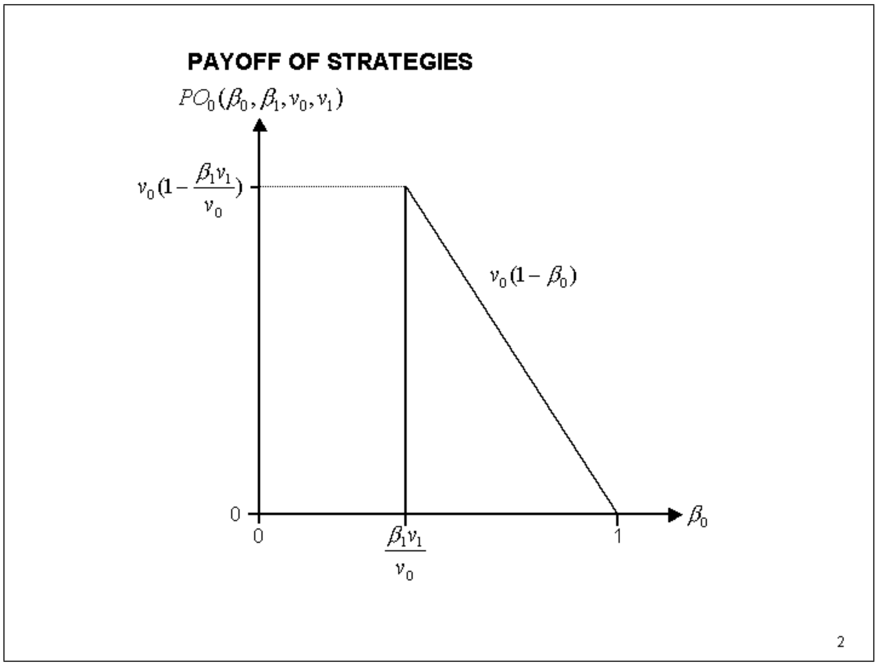

Figure 1: Payoff of Strategies

2.2.2 Unsophisticated Myopic Best Response Play

The best response function determines the average strategy of player 0 if she plays in each round the β_0 that would have maximized her last round's payoff. I stick to the convention that, if no feasible strategy[6] of player 0 could have won the auction (because her value was too low resp. player 1s value was unusually high) then she continues to use her previous rounds strategy β_0^{old}.
Therefore player 0s next rounds strategy is given as

$$\beta_0^{\text{new}}(\beta_1, v_0, v_1) = \frac{\beta_1 v_1}{v_0}\Theta(v_0 - \beta_1 v_1) + \beta_0^{\text{old}}\Theta(\beta_1 v_1 - v_0)$$

$\beta_0^{\text{new}}(\beta_1, v_1)$ can be found by integrating over v_0:

$$\beta_0^{\text{new}}(\beta_1, v_1) = \int_0^1 dv_0 \frac{\beta_1 v_1}{v_0}\Theta(v_0 - \beta_1 v_1) + \beta_0^{\text{old}} \int_0^1 dv_0 \Theta(\beta_1 v_1 - v_0) =$$

$$= \int_{\beta_1 v_1}^1 dv_0 \frac{\beta_1 v_1}{v_0} + \beta_0^{\text{old}} \int_0^{\beta_1 v_1} dv_0 = \beta_1 v_1(\beta_0^{\text{old}} - \ln(\beta_1 v_1))$$

[6] The set of feasible strategies for each player is the interval $[0, 1]$

Integration over v_1 yields

$$\beta_0^{\text{new}}(\beta_1) = \int_0^1 dv_1 \beta_0^{\text{new}}(\beta_1, v_1) = \frac{\beta_1}{4}(1 + 2\beta_0^{\text{old}} - 2\ln\beta_1)$$

Now, since on average $\beta_0^{\text{new}} = \beta_0^{\text{old}} \stackrel{!}{=} \beta_0^{\text{mbr}}$, we arrive at the following

Theorem 2.1 *Consider a 2 player first price open bid auction with linear bidding strategies. Assume that player 1 plays a fixed strategy and player 0 plays in each round the strategy that would have maximized her payoff in the previous round. If she could not have won the last round with any feasible strategy, she continues to use the previous round's strategy. Then, her average strategy satisfies*

$$\beta_0^{mbr}(\beta_1) = \frac{\beta_1 - 2\beta_1 ln\beta_1}{4 - 2\beta_1} \tag{2.2}$$

2.3 Long Memory Best Response in 1POBAs

2.3.1 Best Responses to Known Pure Strategies

The expected payoff $EPO_0(\beta_0, \beta_1, v_0)$ for $v_1 \sim U(0,1)$ is obtained by integrating (2.1) over v_1:

$$EPO_0(\beta_0, \beta_1, v_0) = \int_0^1 dv_1 \Theta(\frac{\beta_0 v_0}{\beta_1} - v_1) v_0 (1 - \beta_0)$$

Now note that the Θ-function is 0 for $v_1 > \frac{\beta_0 v_0}{\beta_1}$, so the integrand takes non-zero values only for $v_1 < \frac{\beta_0 v_0}{\beta_1}$. Additionally, the restriction $v_1 < 1$ is imposed from the upper limit of the integral. So, the integrand only takes non-zero values for

$$v_1 < \min(1, \frac{\beta_0 v_0}{\beta_1}) = 1 - \Theta(1 - \frac{\beta_0 v_0}{\beta_1})(1 - \frac{\beta_0 v_0}{\beta_1})$$

and therefore the integral transforms into

$$EPO_0(\beta_0, \beta_1, v_0) = \int_0^{1-\Theta(1-\frac{\beta_0 v_0}{\beta_1})(1-\frac{\beta_0 v_0}{\beta_1})} dv_1 v_0 (1 - \beta_0) = [1 - \Theta(1 - \frac{\beta_0 v_0}{\beta_1})(1 - \frac{\beta_0 v_0}{\beta_1})] v_0 (1 - \beta_0) \tag{2.3}$$

The following graph shows the expected payoff for player 0 for a fixed value $v_0 = 0.8$ and different strategies β_1 of player 1. Note that for all $\beta_1 < 0.4 = \frac{v_0}{2}$, player 0 maximizes her payoff by playing less than the NE strategy $\beta_0 = \frac{1}{2}$.

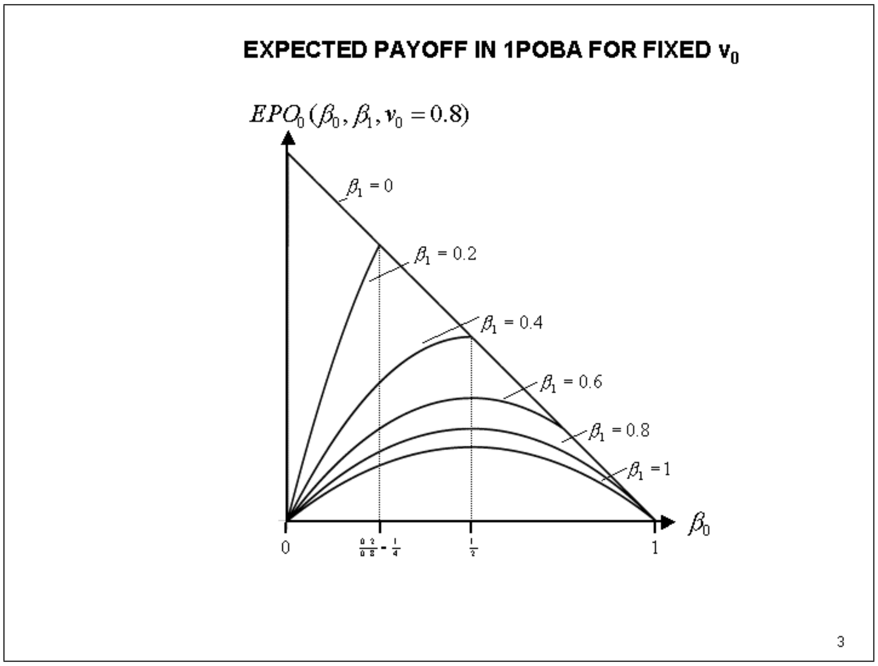

Figure 2: Expected Payoff in 1POBA for fixed v_0

[INSERT GRAPH 'EXPECTED PAYOFF IN 1POBA FOR FIXED v_0']

Now let us determine the optimal response function $\beta_0^{\text{br}}(\beta_1)$, if both values v_0 and v_1 are drawn from a uniform random distribution. This tells us how player 0 would play if she knows player 1s strategy but not next rounds realizations of v_0 and v_1.
Using (2.3), the expected payoff is given by

$$EPO_0(\beta_0, \beta_1) =$$

$$= \int_0^1 dv_0 EPO(\beta_0, \beta_1, v_0) = \int_0^1 dv_0 [1 - \Theta(1 - \frac{\beta_0 v_0}{\beta_1})(1 - \frac{\beta_0 v_0}{\beta_1})] v_0 (1 - \beta_0)$$

$$= \int_0^1 dv_0 v_0 (1 - \beta_0) - \int_0^1 dv_0 \Theta(\frac{\beta_1}{\beta_0} - v_0)(1 - \frac{\beta_0 v_0}{\beta_1}) v_0 (1 - \beta_0) \tag{2.4}$$

The integrand of the second integral only takes on positive values if $v_0 < \frac{\beta_1}{\beta_0}$ and if $v_0 < 1$. Therefore, the condition

$$v_0 < \min(1, \frac{\beta_1}{\beta_0}) = 1 - \Theta(1 - \frac{\beta_1}{\beta_0})(1 - \frac{\beta_1}{\beta_0}) = 1 - \Theta(\beta_0 - \beta_1)(1 - \frac{\beta_1}{\beta_0})$$

yields

$$EPO_0(\beta_0,\beta_1) = \frac{1}{2}(1-\beta_0) - \int_0^{1-\Theta(\beta_0-\beta_1)(1-\frac{\beta_1}{\beta_0})} dv_0(1-\frac{\beta_0 v_0}{\beta_1})v_0(1-\beta_0) =$$

$$= \frac{1}{2}(1-\beta_0) - (1-\beta_0)\{\frac{1}{2}[1-\Theta(\beta_0-\beta_1)(1-\frac{\beta_1}{\beta_0})]^2 - \frac{\beta_0}{3\beta_1}[1-\Theta(\beta_0-\beta_1)(1-\frac{\beta_1}{\beta_0})]^3\} =$$

$$= (1-\beta_0)\{\Theta(\beta_0-\beta_1)(\frac{1}{2}-\frac{\beta_1^2}{6\beta_0^2}-\frac{\beta_0}{3\beta_1}) + \frac{\beta_0}{3\beta_1}\} \qquad (2.5)$$

The next graph shows the expected payoff of player 0 for different strategies β_1 of player 1. Note that for all $\beta_1 \geq 0.5$, player 0 maximizes her payoff by playing her NE strategy $\beta_0 = \frac{1}{2}$. For all $\beta_1 < 0.5$, player 0s best response is also smaller than the NE.

[INSERT GRAPH 'EXPECTED PAYOFF FOR $v_0, v_1 \sim U(0,1)$']

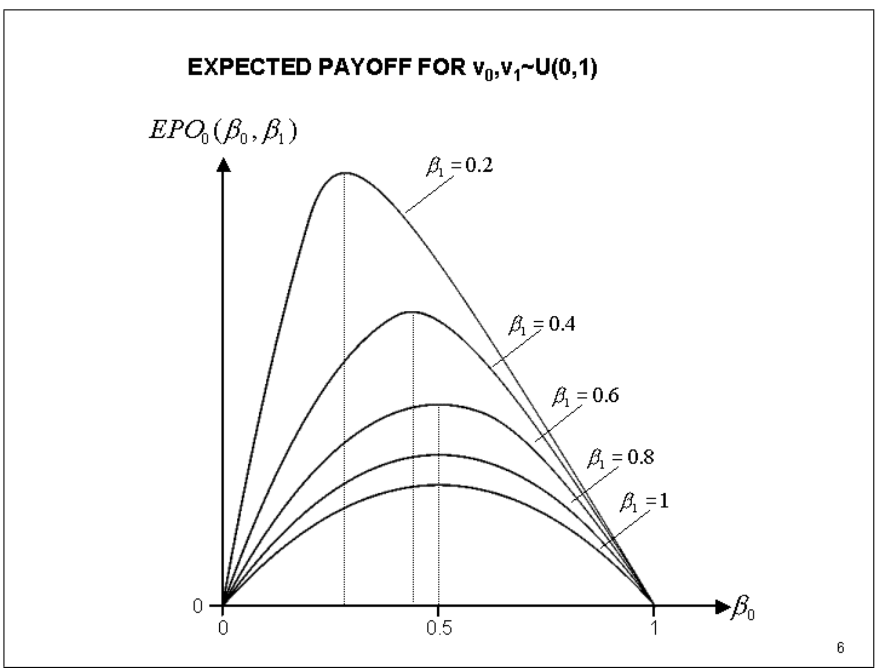

Figure 3: Expected Payoff with v_0, $v_1 \sim U(0,1)$

To analytically derive the optimal response function $\beta_0^{br}(\beta_1)$, we have to discern the cases

$$\beta_0 \begin{cases} > \beta_1 \\ < \beta_1 \\ = \beta_1 \end{cases}$$

For $\beta_0 > \beta_1$, $\Theta(\beta_0 - \beta_1) = 1$, therefore

$$EPO_0(\beta_0, \beta_1) = (1 - \beta_0)\{\frac{1}{2} - \frac{\beta_1^2}{6\beta_0^2}\}$$

so that the first order condition

$$\frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1) \stackrel{!}{=} 0 = \frac{\beta_1^2}{3\beta_0^3} - \frac{\beta_1^2}{6\beta_0^2} - \frac{1}{2}$$

yields

$$\beta_0^{br}(\beta_1) = \frac{\beta_1^{\frac{4}{3}} - \beta_1^{\frac{2}{3}}(\sqrt{81 + \beta_1^2} - 9)^{\frac{2}{3}}}{3(\sqrt{81 + \beta_1^2} - 9)^{\frac{1}{3}}} \tag{2.6}$$

as the only real solution. Plugging the solution back into the condition $\beta_0 > \beta_1$ shows that this solution is valid as long as $\beta_1 \leq \frac{1}{2}$.

For $\beta_0 < \beta_1$, $\Theta(\beta_0 - \beta_1) = 0$ and therefore

$$EPO_0(\beta_0, \beta_1) = (1 - \beta_0)\frac{\beta_0}{3\beta_1}$$

The first order condition yields the solution $\beta_0^{br} = \frac{1}{2}$. Plugging this back into the condition $\beta_0 < \beta_1$ shows that this solution is valid for $\beta_1 > \frac{1}{2}$.

For the symmetric solution $\beta_0 = \beta_1$ we can't proceed straight forward because the first derivative of $EPO_0(\beta_0, \beta_1)$ is discontinuous in $\beta_0 = \beta_1$. Instead, we use the fact that the first derivative to the left of a maximum of a continuous function has to positive whereas right of the maximum it has to be negative. Therefore we look for the solution to the two inequalities

$$\frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1) \begin{cases} > 0 \text{ for } \beta_0 = \beta_1 - \epsilon \\ < 0 \text{ for } \beta_0 = \beta_1 + \epsilon \end{cases}$$

For $\beta_0 = \beta_1 - \epsilon$ we obtain

$$\frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1) = \frac{1 - 2\beta_0}{\beta_1} \stackrel{!}{>} 0$$

The expression is > 0 if $2\beta_0 < 1$. Resubstituting $\beta_0 = \beta_1 - \epsilon$ and taking $\lim\limits_{\epsilon \to 0}$ we arrive at

$$\beta_1 \overset{!}{\leq} \frac{1}{2}.$$

For $\beta_0 = \beta_1 + \epsilon$ we obtain

$$\frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1) = \frac{\beta_1^2}{3\beta_0^3} - \frac{\beta_1^2}{6\beta_0^2} - \frac{1}{2} \overset{!}{<} 0$$

Resubstituting $\beta_0 = \beta_1 + \epsilon$ and taking $\lim\limits_{\epsilon \to 0}$, we arrive at

$$\beta_1 \overset{!}{\geq} \frac{1}{2}.$$

In conclusion, we see that $\beta_0 = \beta_1 = \frac{1}{2}$ is the only symmetric solution. Our results can be summarized as

Theorem 2.2 *Consider a two bidder first price open bid auction where bidders bid according to the linear bidding functions $b_i = \beta_i v_i$. Assume that*

- $\beta_1 = const \in (0, 1)$
- $v_0 \in U(0, 1)$
- $v_1 \in U(0, 1)$

Then, player 0 maximizes her payoff by choosing her best response bidding strategy $\beta_0^{br}(\beta_1)$ as

$$\beta_0^{br}(\beta_1) = \frac{\beta_1^{\frac{4}{3}} - \beta_1^{\frac{2}{3}}(\sqrt{81 + \beta_1^2} - 9)^{\frac{2}{3}}}{3(\sqrt{81 + \beta_1^2} - 9)^{\frac{1}{3}}} \ for\ \beta_1 < \frac{1}{2} \tag{2.7}$$

$$\beta_0^{br}(\beta_1) = \frac{1}{2} \ else \tag{2.8}$$

In particular, $\beta_0^{br}(\beta_1)$ satisfies

$$\beta_1 < \beta_0^{br}(\beta_1) < \frac{1}{2} \ for\ \beta_1 < \frac{1}{2}$$

$$\beta_0^{br}(\beta_1) = \frac{1}{2} \leq \beta_1 \ for\ \beta_1 \geq \frac{1}{2}$$

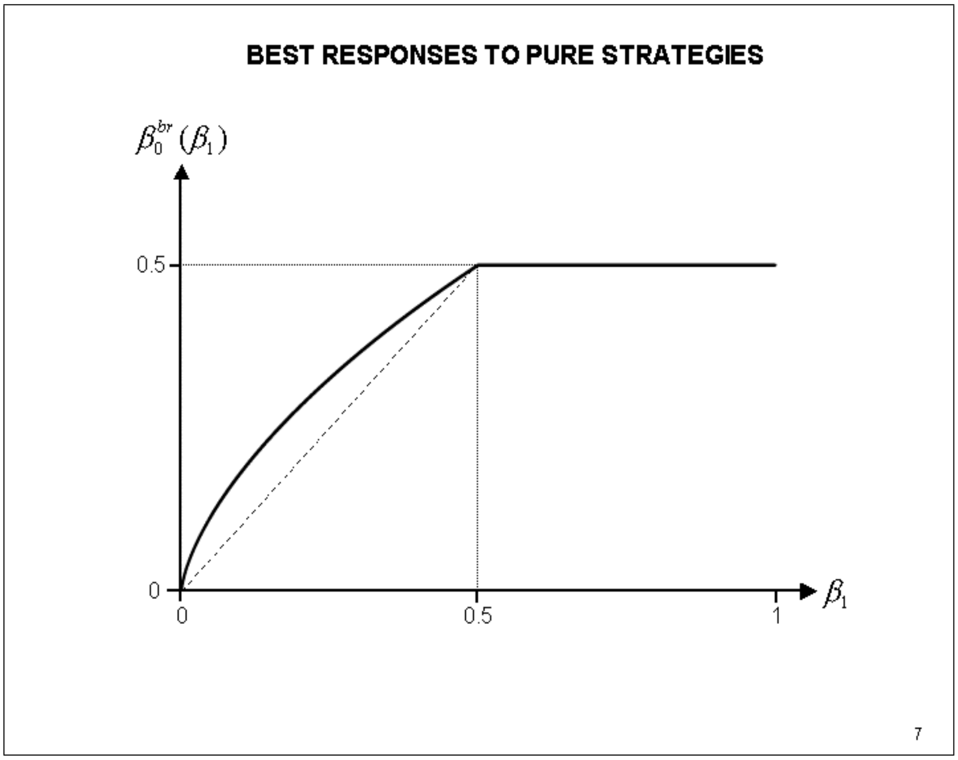

Figure 4: Best Responses to Pure Strategies

The following graph plots the best response function $\beta_0^{\mathrm{br}}(\beta_1)$ against β_1. As stated in Theorem 2.2, for $\beta_1 < \frac{1}{2}$, the best response function satisfies $\beta_1 < \beta_0^{\mathrm{br}}(\beta_1) < \frac{1}{2}$ - as can easily be seen, since the function plot in this interval is always above the first median - the dashed line - and below $\frac{1}{2}$.

[INSERT GRAPHICS 'BEST RESPONSES TO PURE STRATEGIES']

It can be seen immediately from the graph that under PM, $\beta_0^* = \beta_1^* = \frac{1}{2}$ is an NE for the 1POBA: If player 1 plays $\beta_1 = \frac{1}{2}$, player 0s best response is $\beta_0^{\mathrm{br}}(\beta_1^*) = \frac{1}{2}$.
However, the NE is not trembling-hand perfect: If player 1 starts to 'tremble', the best response to β_1s that lie to the right of $\frac{1}{2}$ is $\frac{1}{2}$ while the best response to the β_1s left of $\frac{1}{2}$ is smaller than $\frac{1}{2}$. So, 0s best response against 1s play will be smaller than $\frac{1}{2}$, showing that the NE is not trembling-hand perfect. The next subsection explores this property in more detail.

2.3.2 Long Memory Best Response Play

Let us now investigate 2 player 1POBAs where players don't know their opponents strategy for sure. Instead, they use previously observed bids to estimate their opponents strategy and themselves play their best response against this estimate.

For a start let us imagine that player 1 always plays the same strategy β_1. However, player 0 doesn't know β_1. What wil player 0s best response be?
Assume that the game has already been going on for say $1,000$ rounds. Then player 0 can use the $1,000$ observed bids of player 1 to estimate her strategy β_1 as

$$\tilde{\beta}_1 = \sum_{t=0}^{999} \frac{b_1^{(t)}}{500}$$

since she knows that the values are distributed according to $v_1^{(t)} \sim U(01)$. Her estimate $\tilde{\beta}_1$ will be distributed symmetrically around the true value β_1. Therefore, her best response will in general deviate from the true best response $\beta_0^{\text{br}}(\beta_1)$.

As we increase the number of rounds that player 0 uses for estimating β_1, her estimate of β_1 gets ever more accurate. In the limit of infinitely many rounds, her play therefore converges to the best response function $\beta_0^{\text{br}}(\beta_1)$.

Let us now drop the assumption that player 1 always sticks to the same strategy β_1 but let us consider mutually adapting players. Players use a strategy for a certain number of rounds, after which they simulataneously update their strategy as a best response to their opponents play. Assume for the moment that the number of rounds is sufficiently high so that players can estimate their opponents true strategy accurately.[7] Then, the sequence of strategies at the strategy revision times will be:

$$\beta_0^{(0)}, \beta_0^{(1)} = \beta_0^{\text{br}}(\beta_1^{(0)}), \beta_0^{(2)} = \beta_0^{\text{br}}(\beta_1^{(1)}) = \beta_0^{\text{br}}(\beta_1^{\text{br}}(\beta_0^{(0)})), \cdots$$

$$\beta_1^{(0)}, \beta_1^{(1)} = \beta_1^{\text{br}}(\beta_0^{(0)}), \beta_1^{(2)} = \beta_1^{\text{br}}(\beta_0^{(1)}) = \beta_1^{\text{br}}(\beta_0^{\text{br}}(\beta_1^{(0)})), \cdots$$

Now, Theorem 2.2 shows that for $\beta_0^{(0)} < \frac{1}{2}$,

$$\beta_0^{(0)} < \beta_1^{\text{br}}(\beta_0) < \beta_0^{(2)} = \beta_0^{\text{br}}(\beta_1^{\text{br}}(\beta_0^{(0)})) < \beta_0^{(4)} < \beta_0^{(6)} < \cdots < \frac{1}{2}$$

So, if $\beta_0^{(0)} < \frac{1}{2}$, player 0s strategies constitute a monotonically increasing series and converge to the NE of $\frac{1}{2}$.
If $\beta_0^{(0)} \geq \frac{1}{2}$, $\beta_1^{(1)} = \frac{1}{2}$ and therefore $\beta_0^{(2)} = \beta_0^{(3)} = \cdots = \frac{1}{2}$.
An analogous consideration holds for player 1 so that mutually adaptive play converges to the NE - under the assumption that players update their strategies after so many rounds that they can estimate their oponent's strategy absolutely accurate.

But in reality, if players update their strategies after finitely many rounds, they necessarily have some uncertainty about their opponent's true strategy. Therefore players underbid

[7] This is of course a pure Gedankenexperiment: Players can identify their opponent's true strategy only after infinitely many rounds; however, this would mean that players in fact never update their strategies.

when compared to the NE:
Assume that one player plays the NE strategy but the opponent is not absolutely certain about that. To account for the possibiliy that her opponent plays less than the NE strategy, she will adjust her best response downwards. But the corresponding possibility that the opponent plays above the NE doesn't lead to an upwards shift since the best response against all $\beta_i \geq \frac{1}{2}$ is $\frac{1}{2}$.

In summary we arrive at

Theorem 2.3 *Consider a 2 player first price open bid auction with linear bidding functions. Players infer their opponents strategy in regular intervals from the last* `R` *observed bids and revise their strategies according to the best response dynamics.*
For infinite `R` *the play converges to the NE. However it reaches it only after an infinitely long time. For finite* `R` *the process leads to underbidding when compared with NE play.*

2.4 Second Price Auctions

As already noted in the introduction, the payoff-maximizing bidding strategy in second price open and closed bid auctions is always bidding the true value: Higher bids run the risk of making negative payoff; lower bids might forego positive payoff. This result is well established in the literature and holds for any number of participants and all kinds of risk aversion.
This result can be rederived by using the methodology that we developed in the previous sections. By following the analogous calculations we arrive at an expected payoff of

$$EPO_0(\beta_0, \beta_1) = \Theta(\beta_0 - \beta_1)(\frac{\beta_1^2}{3\beta_0} - \frac{\beta_1^2}{6\beta_0^2} + \frac{1}{2} - \frac{\beta_1}{2}) + \Theta(\beta_1 - \beta_0)(\frac{\beta_0}{3\beta_1} - \frac{\beta_0^2}{6\beta_1}) \tag{2.9}$$

[INSERT GRAPHICS 'EXPECTED PAYOFF IN 2POBAs']

Differentiation of the expected payoff gives the closed form solution for the best response function $\beta_0^{\text{br}}(\beta_1) = 1\ \forall \beta_1 \in [0, \infty)$. Alternatively, also graphical inspection shows that the expected payoff always has its maximum in $\beta_0 = 1$. Either way we arrive at

Theorem 2.4 *The best response $\beta_0^{br}(\beta_1)$ against any bidding strategy β_1 in a second price open bid auction with risk neutral bidders and private values v_0, v_1 drawn from $U(0,1)$ is given by*

$$\beta_0 = 1\ \forall \beta_1.$$

Therefore, the best response dynamics in 2POBAs instantaneously leads to NE play for myopic and long memory play.

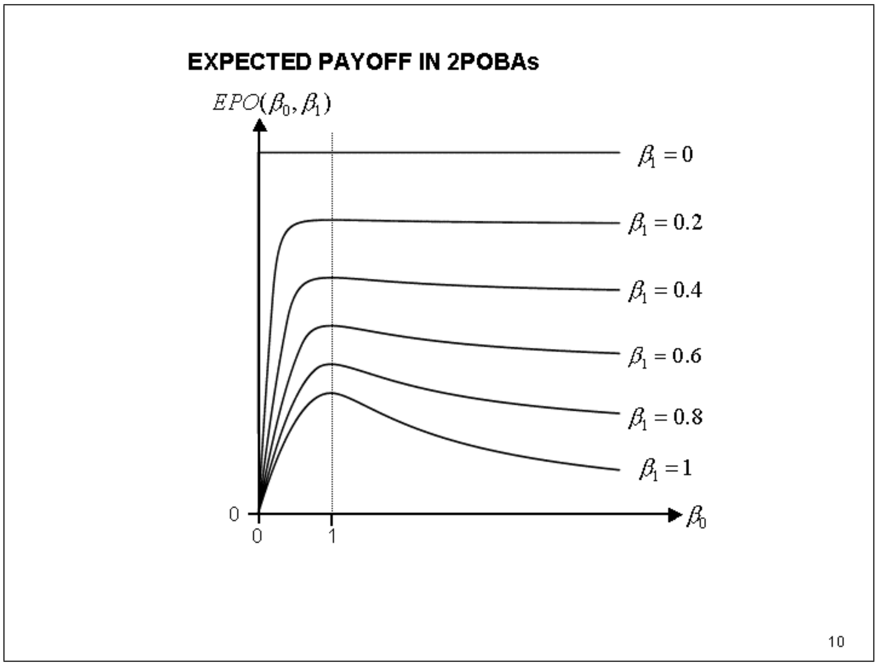

Figure 5: Expected Payoff in 2POBAs

2.5 Revenue Equivalence Revisited

In all considerations up till now we have implicitely assumed that each player believes that her opponent will keep her strategy unchanged for the next `numRounds` auctions. This assumption is crucial. A player that knows that the other player may change her strategy plays different than a player that assumes that her opponent will stick to her strategy.

Consider a payoff-maximizing player 0 that starts with some bid $\beta_0^{(0)}$. Instead of assuming after the first `numRounds` auctions that player 1 will stick to her strategy $\beta_1^{(0)}$, player 0 now assumes her to switch to the best response against her estimate of player 0s strategy $\beta_1^{(1)} = \beta_1^{\mathrm{br}}(E[\beta_0^{(0)}])$. Player 0s best response is now to play $\beta_0^{\mathrm{br}}(\beta_1^{(1)}) = \beta_0^{\mathrm{br}}(\beta_1^{\mathrm{br}}(E[\beta_0^{(0)}]))$.
Taking this iteration a step further, player 1 would also foresee this behavior of player 0 and therefore herself play her best response against this new β_0. This process can be taken further and by induction like in 2.3.2., we see that it converges to $\beta_0 = \beta_1 = \frac{1}{2}$.
So, if payoff-maximizing agents have an infinite foresight horizon H, the process converges to the NE within one round:

$$\lim_{H\to\infty} \beta_0^{(\mathrm{t})} = \lim_{H\to\infty} \beta_1^{(\mathrm{t})} = \frac{1}{2} \; \forall t$$

The infinitely iterated process describes perfectly rational bidders wo assume that their opponents are also perfectly rational. Therefore they bid according to the NE. However, if bidders have limited foresight, i.e., they think this process through for a finite number of rounds only, repeated 1POBAs will show persistent deviations from the NE strategies.

Therefore we arrive at

Theorem 2.5 *Consider a two bidder open bid auction with private values drawn from a common random distribution. Bidders with a limited foresight horizon H bid in consecutive auctions. Players update their strategies to the the best response against their opponent's play every R rounds.*
Then, a repeated first price open bid auction gives lower expected seller revenue and higher earnings volatiliy than a second price open bid auction. The Revenue Equivalence Theorem emerges under the best response dynamics as $\lim_{H\to\infty}$ *or* $\lim_{R\to\infty}$ *and* $\lim_{t\to\infty}$.

3 The Auction Simulator

3.1 Motivation for the Use of Simulation

The idea of simulating auctions is not new [1], [3], [6]. Still, there are various reasons why I developed my own computational auction model, the Auction Simulator(AS), in parallel to the mathematical models about bounded rational bidding behavior:

- Maybe most important, the development of the model provided guidance in identifying the drivers of bounded rational bidding in 1POBAs. Moreover, the model gave significant help in structuring the analysis and the proofs in the mathematical part of the paper.
- Even if several people have worked through mathematical proofs, this is still not a guarantee that they have no flaws. Therefore it is a great help to be able to double-check the quantitative results against a computational model.
- Conversely, the mathematics also provides a quality check for the computer program. It is thereby easier to guarantee a bug-free program. As I will show, the Auction Simulator correctly reproduces the results of the mathematical analysis. Therefore, the AS is a good starting place to investigate topics that are mathematically nasty or simply not trackable. To name a few:
  - I have only analyzed the two player case. In many auctions there is a multitude of participants. Moreover, the number of agents may fluctuate and bidders may not know the number of their competitors in a specific auction.[8]
  - Value distributions need not be uniform but may be normal, log-normal etc.
  - The statistical properties of timeseries that are generated by interacting bounded rational agents are mathematically very difficult if not impossible to track. However, I can use computationally generated timeseries to investigate advanced statistical properties of the process like heteroscedasticity or leptokurtosis.

Because of all these reasons I programmed the Auction Simulator. At the current stage, the program can simulate single-unit open-bid first and second price auctions under best response and quantal response dynamics. In this paper I focus only on the simulation of best response dynamics.

[8] This setup of stochastic entry is not merely of academic interest. For instance in procurement auctions, the suppliers often do not know the number of their competitors.

3.2 Basic Simulation Setup

The Auction Simulator (AS) is programmed in SWARM. SWARM is a programming language that was designed at the Santa Fe Institute for facilitating agent based modelling. In principle it is a library of Objective C objects. It was published under the open GNU license and is downloadable from www.swarm.org. The sourcecode of the AS can be obtained from the author of this paper. Mail to: konrad_richter@mckinsey.com

The AS simulates bidding behavior in repeated auctions. In the investigated setting of open bid auctions, bidders maximize their expected payoff by trying to acquire as often an asset as possible for as little a price as possible. The assets value for the seller is always 0 and there is no reserve price.

The program simulates an arbitrary number of bidders, `numPlayers`. Each bidder i has access to a private set of `numStrategies` strategies. Strategies are real numbers β_i. Each strategy assigns to the players private value a bid according to `bid` $= \beta \cdot$ `value`. Each player chooses an active strategy that determines the bid that the bidder is actually placing. All other strategies are evaluated as well to see how they would have performed if they would have been the active ones. Changes of the active strategy are possible only every `numRounds` ronds. In the GA case, at these times, the strategy population is replaced by a new generation. In the FS case, a new strategy is chosen as active.

A model run consists of the following steps:

1. bidders randomly initialize their active and `numStrategies-1` non-active strategies
2. for `numGenerations` generations of strategy sets
   (a) bidders and seller reset their current payoffs
   (b) for `numRounds` auctions
      i. bidders' values are chosen from a uniform distribution on $[0, 1]$
      ii. bidders submit their bids
      iii. the auction module determines the winner
      iv. bidders update the current payoffs of their active and passive strategies
   (c) bidders evaluate the payoff, their strategies (would) have generated during the `numRounds` auctions and choose the best one as the active one for the next `numRounds` auctions
   (d) if the GA is used, the strategy population is updated
      i. the best `numElite` strategies are kept unchanged for the next generation
      ii. the best `numParents` strategies are taken as parents to create offspring
      iii. bidders update their worst `numStrategies-numElite` strategies using the genetic operators mutation and crossover
   (e) proceed with (a)

For further reference and to get a feeling for the program capabilities, take a look at the model parameters in Appendix B.

3.3 Simulating Learning

3.3.1 Genetic Algorithms

In the Auction Simulator we are using real GAs instead of ones with bitstring populations. Real GAs are faster than bitstring GAs, note however that the theoretical properties are less well understood than for bitstring GAs [11], [15].

The implementation of the real GA in the AS is as follows:
In the beginning, `numStrategies` real numbers are randomly generated for each bidder and used as bidding strategies. Strategy number 0 is used as the active strategy that determines the playing behavior of the bidder. Subsequently, bidders use their active strategy for bidding in `numRounds` consecutive auctions.
For each strategy, the payoff is added up for the `numRounds` rounds.[9] After that, the population of bidding strategies is updated:

The first step in the updating process is to rank the strategies according to their fitness:
The best strategy is chosen as the active one for the next generation.
The best `numElite` strategies are left unchanged for the next round. This elitism reflects the assumption that a bidder would like to evaluate her most successful strategies also in the next auction without any change.
The best `numParents` different strategies are collected in a breeding list. Two strategies β_i and β_j are considered as different if $\mid \beta_i - \beta_j \mid >$`strategyDistance`.[10]. If there are less than `numParents` different strategies in the strategy set, the missing positions are filled up by randomly generated strategies.

The second step is the creation of `numStrategies-numElite` new strategies by application of the crossover operator:
The crossover operator randomly selects two different parent strategies from the `breedingList`. The selection probabilities are assigned according to their rank: The best strategy is selected with relative probability `numParents`, the next with relative probability `numParents-1` and so on. The last strategy in the `breedingList` has a relative selection probability of 1.
Rank proportional selection is better suited than fitness proportional selection to maximize expected payoff. The reason is that in fitness proportional selection, the selection pressure rapidly declines if all strategies are near the global optimum. With rank-proportional selection, the best solutions in the `breedingList` are always much more likely to create offspring than worse ones - even if their absolute fitness advantage is arbitrarily small.
Having selected the two strategies, denote the lower by $\beta^{\min}$ and the larger by $\beta^{\max}$. A new strategy is constructed by selecting a number between $\beta^{\min}(1 - \texttt{crossoverPar}(\beta^{\max} - \beta^{\min}))$ and $\beta^{\max}(1 + \texttt{crossoverPar}(\beta^{\max} - \beta^{\min}))$ with uniform probability.[11].

[9] For the inactive strategies this is the payoff, they *would* have generated if they had been the active one.
[10] This prevents the GA from getting stuck in too homogenous populations
[11] The extension of the crossover interval offsets the tendency of crossover to equalize all strategies. Usually the literature assumes a normal distribution for the mutation. However, in this simulation a uniform distribution yields better results.

The third step is the application of the mutation operator to the offspring:
With probability `mutationProb` strategy β is changed into another value according to a normal distribution with mean β and a variance of `mutationPar` percent. The mutation operator is a further mechanism that prevents the population from getting too homogenous. Together with the `numElite` unchanged strategies, the newly constructed strategies form the next generation strategy set.

3.3.2 Fixed Strategies

As an alternative simulation tool we employ fixed strategies:
In the beginning we partition the strategy space $(0, 1)$ into a grid of equidistant strategies, e.g. `numStrategies`=100 ranging from 0.00 to 0.99. One strategy is randomly selected as active and determines the initial bidding behavior of the agent.
For each auction the current payoff that each strategy β_0 would have generated in round i is calculated by $PO_0^{\mathrm{curr};(i)}(\beta_0) = v_0(1-\beta_0)\Theta(\beta_0 v_0 - \beta_1 v_1)$.
After the first `numRounds` auctions, each strategies' payoffs are added up and yield the payoff, the strategy generated in the first `numRounds` auctions:

$$PO_0^{\mathrm{numRounds};(1)}(\beta_0) = \sum_{i=1}^{\mathrm{numRounds}} PO_0^{\mathrm{curr};(i)}(\beta_0)$$

After, the strategy that generated the highest payoff is selected as the active one. If no strategy has positive payoff, the currently active strategy remains active.[12] Denote

$$PO_0^{\mathrm{cum};(1)}(\beta_0) = PO_0^{\mathrm{numRounds};(1)}(\beta_0)$$

After the next `numRounds` auctions the cumulated payoff is recalculated as a weighted sum of the old cumulated payoff and the new payoff generated in the last `numRounds` auctions:

$$PO_0^{\mathrm{cum};(2)}(\beta_0) = \mathtt{memoryStrength} PO_0^{\mathrm{cum};(1)}(\beta_0) + PO_0^{\mathrm{numRounds};(2)}(\beta_0)$$

Again, the strategy that generated the highest payoff is selected as the active one and the process starts anew.

3.4 Simulation Results

For details on the parameter settings please consult appendix B

[12]The case that no strategy has positive payoff happens if $v_0 < \beta_1 v_1$ for all auctions witin the last generation

3.4.1 Responses against Pure Strategies

The first experiment checks the predictions of the Theorems in chapter 2 against the computer model. I show that the Auction Simulator correctly replicates the playing behavior of boundedly rational players under best response dynamics.

I fix the values of β_1 at numbers between 0.1 and 0.9 in steps of 0.1.
v_0 and v_1 are drawn from a uniform random distribution on $(0, 1)$.
Each experiments consists of 4 runs. The random seeds for the runs are 10, 11, 12 and 13. Each run consists of 6.000 generations after a phase-in of 4.000 generations.

For myopic best response play in 1POBAs, the theory predicts $\beta_0^{\text{mbr}}(\beta_1)$ as given by Theorem 2.1. For the simulation I use fixed strategies with the following parameter settings:

`numRounds`	1	`numStrategies`	100	`selectionType`	2	`memoryStrength`	0

Since I use only 100 strategies, I have to correct the result for the overestimation by 0.005. I do not simulate Myopic Play with GAs since GAs are known to perform poorly when they are confronted with a frequently varying fitness landscape.

For long memory best response, $\beta_0^{\text{lbr}}(\beta_1)$ is given according to Theorem 2.2. For the simulation I use a GA with the following parameter settings:

`numRounds`	1000	`numStrategies`	10	`numParents`	5
`numElite`	2	`strategyDistance`	0.01	`crossoverPar`	0.1
`mutationProb`	7.5%	`mutationPar`	0.1		

Alternatively I use fixed strategies on $(0, 1)$ with the following parameters:

`numRounds`	1000	`numStrategies`	100	`selectionType`	2	`memoryStrength`	1

Results are as follows:

$\beta_0^{\mathbf{mbr}}(\beta_1)$ via FSs

β_1	$\beta_0^{th}(\beta_1)$	$\beta_0^{ex}(\beta_1)$	$\beta_0^{\mathbf{corr}}(\beta_1)$	**dev**$(\beta_0^{th}, \beta_0^{ex})$
0.1	0.14750	0.15241	0.14741	-0.06%
0.2	0.23438	0.24044	0.23544	0.45%
0.3	0.30070	0.30420	0.29930	-0.50%
0.4	0.35407	0.35713	0.35213	-0.54%
0.5	0.39772	0.40256	0.39756	-0.04%
0.6	0.43321	0.43720	0.43220	-0.23%
0.7	0.46129	0.46643	0.46143	0.03
0.8	0.48210	0.48780	0.48280	0.15%
0.9	0.49530	0.49898	0.49398	-0.27%

$\beta_0^{\mathbf{lbr}}(\beta_1)$ **via GAs and FSs**

β_1	$\beta_0^{th}(\beta_1)$	$\beta_0^{GA}(\beta_1)$	**dev**(β_0^{th},β_0^{ex})	$\beta_0^{FS}(\beta_1)$	**dev**(β_0^{th},β_0^{ex})
0.1	0.18231	0.18393	0.89%	0.18200	-0.17%
0.2	0.28390	0.28490	0.35	0.28320	-0.25%
0.3	0.36598	0.36673	0.20%	0.36507	-0.25%
0.4	0.43685	0.43665	-0.05%	0.43554	-0.30%
0.5	0.5	0.48909	-2.18%	0.48838	-2.32%
0.6	0.5	0.49789	-0.42%	0.49616	-0.76%
0.7	0.5	0.49659	-0.68%	0.49582	-0.84%
0.8	0.5	0.49709	-0.58%	0.49552	-0.90%
0.9	0.5	0.49273	-1.45%	0.49547	-0.91%

The table's first column gives player 1s fixed strategy β_1, the second column the optimal response β_0 predicted by the according Theorems. The third column shows the mean of the simulated β_0 time series, averaged over the four runs. The last column gives the deviation between analytical prediction and simulation. In the first table, I use an additional column to correct for the finite size of the strategy set.
In all cases, we see that the Auction Simulator replicates the analytical reults very well. The simulation results lie within 1% of the theoretical values for nearly all parameter settings. The persistent slight underbidding - that is especially pronounced $\beta_1 = 0.5$ - stems from the convex shape of the best response function $\beta_0^{\text{lbr}}(\beta_1)$.
For second price auctions I performed the analogous experiments with GAs and FSs. As theoretically predicted, strategies in the experiments always converged to 1.[13]

From now on I will focus on simulations with FSs since they yield the same results as GAs but allow for more flexiblity.

3.4.2 Mutual Adaptation

This experiment investigates the bidding strategies of two mutually adapting bidders in 1POBAs via FSs. In particular I examine the influence on the mean and the volatility of the bidding strategies if bidders use more and more past information for revising their strategies. I consider a two bidder first price auction with the NE of 0.5. Starting from a default setting of completely myopic play (`numRounds`=1 and `memoryStrength`=1) I assess then the impact of increasing `numRounds` and `memoryStrength`.

Results are as follows:

$\beta_0^{\mathbf{br;ad}}(\beta_1)$ **in dependence of using past information**

numR	β_0^{ex}	$\sigma^2(\beta_0^{ex})$	**memS**	β_0^{ex}	$\sigma^2(\beta_0^{ex})$
1	0.11933	0.17788	0	0.11930	0.17788
10	0.33264	0.11814	0.3	0.1918	0.14774
100	0.44303	0.05888	0.6	0.26193	0.12673
1000	0.47851	0.02727	1.0	0.46072	0.007531.0

[13]For FS, a positive but arbitrarily small memoryStrength is needed for convergence.

We see that as predicted, under best response learning rules the bids are always below the Nash Equilibrium. Bids get closer to the NE and volatility of bidding strategies drops as the intervals for strategy updating get larger or the memory strength of bidders increases. This supports the prediction of Theorem 2.5 that bidding behavior converges under best response to the NE as more and more past information is used for strategy updating.
Note that this result is in line with others from the literature. Many models of bounded rationality show supoptimality if strategies are revised too frequently [7], [12].

3.4.3 Increasing Number of Bidders

This experiment investigates the influence of the number of participants on the outcome of first price auctions. In all of the experiments, `numRounds` was set to 1 and memoryStrength was set to 0 respective 1.

First price auctions with more than two participants

numBidders	β_0^{NE}	$\beta_0^{ex}(0.0)$	**dev**$(\beta_0^{NE}, \beta_0^{ex})$	$\beta_0^{ex}(1.0)$	**dev**$(\beta_0^{NE}, \beta_0^{ex})$
2	0.5	0.11933	-76.1%	0.46072	-7.9%
5	0.8	0.70308	-12.1%	0.77881	-2.6%
10	0.9	0.85418	-5.1%	0.89181	-0.9%
100	0.99	0.98264	-0.7%	0.98467	-0.5%

The results show that best response leads to persistent underbidding for all numbers of bidders. However, as the number of bidders increases, the bidding strategies converge towards the NE. So we can conclude that deviations from NE play are especially pronounced if there are only few participants. Additionally we see again the effect that perfect memory drives the bids more towards the NE than myopic play.

3.4.4 Price volatility for normal distributed values

The following experiment deviates from the hitherto assumed uniform value distributions for players. Instead I assume normal value distributions with varying standard deviations. I am especially interested in how much first and second price auctions 'blow up' the volatility of the underlying value distribution. This gives us a good proxy for comparing the 'riskiness' of first and second price auctions. I investigated here 20 players with 1.500 fixed strategies each. The value of `memoryStrength` was set to 0.5.

Mutually adapting players with normal distributed values

$\mu(v)$	$\sigma(v)$	$\sigma(v)$ in %	$\mu(PO_{1PA}^{sell})$	$\sigma(PO_{1PA}^{sell})$	σ in %	$\mu(PO_{2PA}^{sell})$	$\sigma(PO_{2PA}^{sell})$	σ in %
5	1	20%	6.31705	0.52755	8.4%	6.39466	0.39752	6.2%
5	0.1	2%	5.41903	0.17304	3.2%	5.4222	0.12599	2.3%
5	0.01	0.2%	5.1328	0.05569	1.1%	5.14062	0.03986	0.8%
5	0.001	0.02%	5.04195	0.01769	0.4%	5.04456	0.01260	0.2%

The experiment shows that first price open bid auctions yield for all parameter settings a higher price volatility than 2POBAs.[14] When we compare the volatility of first and second price auctions, we see that the differences between the auction formats become more pronounced as the underlying value distribution gets more peaked. Graphical inspection of graphs and histograms also suggests heteroscedasticity and leptokurtosis in the seller returns in 1POBAs. This suggests that the model is a good starting point to investigate high frequency data in financial markets. However, this is a topic for future research.

4 Conclusion and Outlook

In this paper I have used mathematical and computational methods to investigate repeated auctions. 2 bidders bid in repeated open auctions for assets. They have different valuations of the asset currently under auction which change for every auction. Over time, each bidder tries to learn her optimal bidding strategy that maximizes her payoff.
The results of mathematical and computational analysis showed that 1POBAs under best response dynamics show excess volatility and therefore a suboptimal allocation of goods. In contrast, players in 2POBAs find the NE easily - either upfront by simple reasoning or in the course of time by following the best response strategy updating. Similar arguments show that in the sealed bid case second price auctions are more efficient than first price auctions.

The underlying reason of the optimal properties of second price auctions is that in these auctions bidding the true value is always a (weakly) dominant strategy. Therefore, the optimal strategy doesn't depend on the other players' strategies. In contrast, the optimal bid in value shading auctions -like first price auctions, all-pay auctions etc. - depends on the strategies of all other players. An analysis parallel to the one we performed for single unit auctions shows that *any value shading auction design must show excess volatility* when compared to an auction design where bidding the true value is dominant.
For single unit auctions this showed that the second price auction allocates goods optimal and minimizes volatility. For multi-unit auctions the corresponding optimal auction would be the Ausubel auction [2]. For double auctions, I currently know of no auction mechanism where bidding the true value is the dominant strategy. The development of such a format is a question for further research.

What do these results imply for the economy?
A first obvious application is to stock and derivative markets. Floor-based interaction between brokers is a repeated open double-sided multi-unit auction. Current orderbook designs are uniform price auctions that incentivize for bid shading. This is especially true for bids of large players. My aim is to redesign order books to an auction format where bidding the true value is the dominant strategy.
The increased stability of stock prices would have clear advantages for the productive economy: First, it would reduce the need for firms to purchase financial derivatives for risk

[14] Strictly speaking, the price distribution in 1POBAs is not normal, so σ is not the correct measure for the volatility. However, here I use this measure as a proxy to compare the riskiness of first and second price auctions.

hedging purposes; therefore more money could be invested in the productive sector. Second, the allocation of capital would be more efficient since stock and derivative prices would better reflect the underlying fundamental value. Third it would reduce the occurence of bubbles and crashes since there would be less random price movements that can be exploited by chartists; therefore it would help to stabilize the world economy.

Other markets where order book redesign would have clear advantages are electricity markets and the currently emerging global market for emission certificates. Especially the latter would be the ideal place for a first implementation of revised order books: It is a new market without strong incumbents and the declared aim is -according to the Kyoto protocol - to use the certificates as efficiently as possible.

Another important application is with respect to supply chains and -networks. In experiment 4 we saw the destabilizing power of first price auctions to blow up the volatility of values to a much higher volatility of prices. Supply chains can be viewed as chains of interconnected double auctions. If these auctions incentivize for value shading, even a very small volatility of production cost could result in highly exeggerated price volatility on the consumer side. However, the introduction of second price auctions and equivalents would have to be accompanied by measures to prevent collusion among the participants.
Further possible applications of auction redesign include landconservation auctions [10], IPOs, treasury bill auctions, procurement auctions and many more.

As an intermediary step of my future research I want to computationally and mathematically investigate multi-unit double auctions. The tools needed for this are in principle the ones I developed in this paper. I hope that the improved understanding of double sided multi-unit auctions will allow me to develop specific improvement recommendations that help to stabilize the economic system.

A The Θ function

The Θ function is defined by

$$\Theta(x) = \begin{cases} 0 \text{ for } x < 0 \\ 1 \text{ for } x \geq 0 \end{cases}$$

The minimum of two values x and y can be written as

$$\min(x, y) = x - \Theta(x - y)(x - y)$$

B Simulation Parameters

Parameters in model.setup

parameter	**explanation**
`numRounds`	number of auctions played between strategy updatings
`auctionType`	if ==1: first price sealed bid auction if ==2: second price sealed bid auction
`transactionFee`	transaction fee payed in each round by winning bidder and seller
`numPlayers`	number of bidders participating in the auction
`startOfAverageCalculation`	number of initial strategy updatings after which the calculation of averages for graphical output starts
`reportingPlayerID`	ID of player that reports her strategy list
`randomSeed`	sets the initial state of the random generator

Parameters in player.setup for GAs and Fixed Strategies

parameter	explanation
`learningType`	if==1, FS are used; if==2, GA is used
`selectionType`	if==1, fitness proportional selection; if==2, rank proportional selection
`fixedStrategyMarker`	if ==0: each player updates her strategy set; mutual adaptation if ==1: player 1 plays a fixed strategy, all others update their strategy sets if ==2: player 0 updates her strategy set, all others play fixed strategies
`fixedBeta`	this parameter is only important if `fixedStrategyMarker` $\neq 0$ if $\in (0,1)$: β=`fixedBeta`; if $\in (-1,0)$: $\beta \sim U(-fixedBeta, 1)$
`fixedValueMarker`	if==0: each player has random values if==1: player 1 alwas has fixed value; all other values are random if==2: player 0 has random value; all other players have fixed value
`fixedValue`	determines respective values if `fixedValueMarker`$\neq 0$
`player0FixedValueMarker`	if ==0: random v_0; if ==1: fixed v_0
`player0FixedValue`	if `player0FixedValueMarker`==1, then v_0= `player0FixedValue`
`numStrategies`	number of strategies in each players strategy set
`valDistShape`	if==0: values from normal Dist; if==1: values from uniform Dist
`valDet0`	lower bound for uniform Dist, resp. variance for normal Dist
`valDet1`	upper bound for uniform Dist, resp. mean for normal Dist

Parameters in player.setup for Fixed Strategies only

parameter	explanation
`memoryStrength`	multiplicative weigthing factor for previous rounds' payoffs
`minStrategy`	value of the minimal strategy
`maxStrategy`	value of the maximal strategy

Parameters in player.setup for GAs only

parameter	explanation
`strategyDistance`	minimum distance between strategies in the breedingList
`numElite`	number of fittest strategies that stay unchanged in the population
`numParents`	number of fittest strategies that are used for creating offspring by crossover
`crossoverPar`	fraction of the interval between two parent strategies by which offspring is allowed to lie outside the interval
`mutationType`	if==0: uniform Dist between $\beta(1\pm$ `mutationPar`$)$ if==1: normal Dist with variance $\beta\cdot$ `mutationPar`
`mutationProb`	probability of mutation of a strategy
`mutationPar`	determines extent of mutation; see `mutationType`

C Abbreviations

abbreviation	explanation
1POBA, 1PSBA	First Price Open Bid Auction resp. Sealed Bid Auction
2POBA, 2PSBA	Second Price Open Bid Auction resp. Sealed Bid Auction
AS	Auction Simulator
BR	Best Response
FS	Fixed Strategy
GA	Genetic Algorithm
LBR	Long Memory Best Response
MBR	Myopic Best Response
NE	Nash Equilibrium
RET	Revenue Equivalence Theorem
SIPV	Symmetric Independend Private Values Framework

References

[1] Andreoni J. and Miller J., 1995, *Auctions with Artificial Adaptive Agents*, Games and Economic Behavior, *10*, 39-64

[2] Ausubel L. M., 1997, *An efficient Ascending-Bid Auction for Multiple Objects*, Working Paper 97-06, University of Maryland

[3] Byde A., 2002, *Applying Ecolutionary Game Theory to Auction Mechanism Design*, Hewlett-Packard Workingpaper, http://www.hpl.hp.com/techreports/2002/HPL-2002-321.pdf

[4] Conlisk J., 1996, *Why Bounded Rationality?*, Journal of Economic Literature, *XXXIV*, 669-700

[5] Cox J., Roberson B. and Smith V. L., 1982, *Theory and Behavior of Single Object Auctions*, in Vernon L. Smith, ed., *Research in Experimental Economics*, Greenwich, JAI Press

[6] Dawid H., 1999, *On the convergence of genetic learning in a double auction market*, Journal of Economic Dynamics & Control, 1545-1567

[7] Ellison G. and Fudenberg D., 1995, *Word-of-Mouth Communication and Social Learning*, Quarterly Journal of Economics, *110*, 93-125

[8] Fudenberg D., and Levine D., 1999, *The Theory of Learning in Games*, MIT Press, Cambridge, London, second edition

[9] Harrison G. W., 1989, *Theory and Misbehavior of First-Price Auctions*, American Economic Review, *79*, 749-62

[10] Hailu A. and Schilizzi S., 2002, *Learning in a 'Basket of Crabs': An Agent-Based Computational Model of Repeated Conservation Auctions*, www.bwl.uni-kiel.de/vwlinstitute/gwrp/wehia/papers/hailu.pdf

[11] Holland J., 1992, *Adaptation in Natural and Artificial Systems*, Massachusetts, MIT Press, 2^{nd} edition

[12] Joshi S., Parker J., Bedau M., 1999, *Financial Markets can be at Sub-Optimal Equilibria*, SFI Workingpaper, www.santafe.edu/sfi/publications/Working-Papers/99-03-023.pdf

[13] Kagel J. H., 1995, *Auctions: A survey of Experimental Research*, in Kagel and Roth ed., *The Handbook of Experimental Economics*, Princeton, Princeton University Press

[14] Lebrun B., 1999, *First Price Auctions in the Asymmetric N Bidder Case*, International Economic Review, *40, No.1*, 125-142

[15] Lux T. and Schornstein S., 2002, *Genetic Learning as an Explanation of Stylized Facts of Foreign Exchange Markets*, Workingpaper, www.bwl.uni-kiel.de/vwlinstitute/gwrp/publications/lux_gen_learning.pdf

[16] Maskin E. S. and Riley J. G., 2000, *Asymmetric Auctions*, Review of Economic Studies, *67*, 413-438

[17] Myerson, R. B., 1981, *Optimal Auction Design*, Mathematics of Operations Research, *6*, 58-73

[18] Riley J. G. and Samuelson, W. F., 1981, *Optimal Auctions*, American Economic Review, *71*, 381-92

[19] Vickrey W., 1961, *Counterspeculation, Auctions and Competitive Sealed Tenders*, J. Finance, *16*, 8-37