

Evolutionary Learning in Auctions

Inaugural-Dissertation zur Erlangung des Grades eines
Doktors der Wirtschafts- und Sozialwissenschaften der
Wirtschafts- und Sozialwissenschaftlichen Fakultät der
Christian-Albrechts-Universität zu Kiel

vorgelegt von

Mag. rer. nat. Konrad Richter
aus Wien

gedruckt in Wien, 2005

Before starting with the thesis I want to thank everybody who made it possible that I could write the book that you have in hands. This is first of all Katharina who supported me in each phase of the work - including a one year long-distance relationship. Thomas Lux gave me great freedom in choosing my research area and guided my efforts with stimulating questions. Numerous colleagues whom I met at university, at conferences and at McKinsey were valuable spare-ring partners for developing the ideas presented here. A special thanks to Wolfgang Gatterbauer whose curiosity and straight-to-the-point questions often forced me to think one step further than I previously had. My parents and my friends provided me kindly with opportunity to molest them with my enthusiasm for certain aspects of my research - and also to regenerate from work.

The research would not have been possible in this form without the generous financial support that McKinsey&Co. provides as part of their Fellow-Program. And finally thanks to the DAAD for providing me with a scholarship and for funding the PhD-program 'Quantitative Economics' in Kiel.

Contents

| | | |
|----------|--|-----------|
| 1 | Motivation and Outline | 6 |
| 1.1 | Motivation | 6 |
| 1.2 | Outline | 7 |
| 2 | Literature Background on Auctions | 9 |
| 2.1 | Private-Value Single-Unit Pure-Bidder Auctions | 9 |
| 2.1.1 | Standard Formats | 9 |
| 2.1.2 | Revenue Equivalence | 10 |
| 2.2 | A Quick Tour through Auction Theory | 12 |
| 2.3 | Auction Experiments | 15 |
| 3 | Literature Background on Game Theory | 17 |
| 3.1 | Motivation | 17 |
| 3.2 | Bounded Rationality | 17 |
| 3.2.1 | Reasons for Bounded Rationality | 17 |
| 3.2.2 | Concepts from Evolutionary Game Theory | 19 |
| 3.3 | Auction Theory under Bounded Rationality | 21 |
| 3.4 | Simulations | 22 |
| 3.4.1 | Genetic Algorithms | 22 |
| 3.4.2 | Auction Simulations | 24 |
| 3.4.3 | Financial Market Simulations | 24 |
| 4 | Best Response in Open Bid Auctions | 26 |
| 4.1 | Structure of this Chapter | 26 |
| 4.2 | Myopic Best Response to Fixed Strategies in IPOBAs | 27 |
| 4.2.1 | Payoff Function | 27 |
| 4.2.2 | Myopic Best Response to β_1 | 27 |
| 4.2.3 | Myopic Best Response with Seller Reserve Price | 30 |
| 4.3 | Mutual Myopic Best Response Adaptation | 32 |

| | | |
|----------|---|-----------|
| 4.3.1 | Strategy Distribution under Myopic Best Response | 32 |
| 4.3.2 | Convergence under Myopic Mutual Adaptation | 35 |
| 4.4 | Perfect Memory Best Response to Fixed Strategies in 1POBAs | 37 |
| 4.4.1 | Expected Payoff Function | 38 |
| 4.4.2 | Perfect Memory Best Response to β_1 | 40 |
| 4.5 | Mutual Perfect Memory Best Response Adaptation | 41 |
| 4.6 | Second-Price Auctions | 43 |
| 4.7 | Revenue Equivalence Revisited | 45 |
| 5 | Addenda to the Best Response Dynamics | 47 |
| 5.1 | Value Dependent Perfect Memory Best Response | 47 |
| 5.2 | Probability-Weighted Myopic Best Response Play | 50 |
| 5.2.1 | Strategy Estimation from Observed Bids | 50 |
| 5.2.2 | Probability-Weighted Myopic Best Response Play | 51 |
| 5.3 | Sealed Bid Auctions | 54 |
| 6 | Quantal Response in Open Bid Auctions | 55 |
| 6.1 | Myopic Quantal Response to Fixed Strategies in 1POBAs | 55 |
| 6.1.1 | Payoff-Linear Myopic Quantal Response | 55 |
| 6.1.2 | Payoff-Exponential Myopic Quantal Response | 58 |
| 6.2 | Perfect Memory Quantal Response to Fixed Strategies in 1POBAs | 59 |
| 6.2.1 | Payoff-Linear Perfect Memory Quantal Response | 59 |
| 6.2.2 | Payoff-Exponential Perfect Memory Quantal Response | 60 |
| 6.3 | Quantal Response in 2POBAs | 61 |
| 7 | Best Response in Pure Seller Auctions | 63 |
| 7.1 | Payoff Function for Sellers | 63 |
| 7.2 | Myopic Best Response to Fixed σ_1 | 64 |
| 7.3 | Perfect Memory Best Response to Fixed σ_1 | 65 |
| 8 | Best Response in Double Auctions | 68 |
| 8.1 | Double Auctions with Linear Bidding Strategies | 68 |
| 8.1.1 | Myopic Best Response | 69 |
| 8.1.2 | Perfect Memory Best Response | 70 |
| 8.2 | Outlook to Future Research: Dynamics in Repeated DAs with Affine Bidding Strategies | 72 |
| 8.2.1 | Expected Payoffs under Perfect Memory Play | 73 |
| 8.2.2 | The NE Region | 75 |
| 8.2.3 | Dynamics around the NE | 77 |
| 8.2.4 | Phase Space Regions off the NE | 78 |

| | | |
|-----------|--|------------|
| 9 | The Auction Simulator | 81 |
| 9.1 | Motivation for Simulation | 81 |
| 9.2 | Program Architecture | 81 |
| 9.3 | Main Aspects of the Program | 83 |
| 9.3.1 | Simulation Flow | 83 |
| 9.3.2 | Determination of Payoffs | 84 |
| 9.4 | Learning in the Auction Simulator | 85 |
| 9.4.1 | Real GAs | 85 |
| 9.4.2 | Fixed Strategies | 86 |
| 9.4.3 | Interpretation of GA and FS Results | 86 |
| 10 | Double-check of Analytic Results | 88 |
| 10.1 | Bidder Response against Fixed Bidding Strategies | 88 |
| 10.2 | Seller Best Response against Fixed Asking Strategies | 92 |
| 10.3 | Linear Strategy Best Response in Double Auctions | 94 |
| 10.4 | Myopic Best Response with Reserve Price | 96 |
| 10.5 | Strategy Distribution under Myopic Best Response | 96 |
| 10.6 | Convergence under Mutual Best Response Adaptation | 97 |
| 10.7 | Payoff-Exponential Quantal Response | 101 |
| 10.8 | Comments on Chapter 5 | 102 |
| 10.9 | Conclusions from this Chapter | 102 |
| 11 | Computational Analysis of Mutually Adapting Bidders | 104 |
| 11.1 | Auction Efficiency | 104 |
| 11.2 | Myopic Best Response with Normal Value Distributions | 106 |
| 11.2.1 | Influence of the Variance | 106 |
| 11.2.2 | Myopic Best Response | 109 |
| 11.3 | Imperfect Memory Best Response with Normal Value Distributions | 113 |
| 11.4 | Truncated Fictitious Play with Uniform Value Distributions | 114 |
| 11.5 | Computational Assessment of Nash Equilibria | 120 |
| 11.6 | Conclusion from this Chapter | 121 |
| 12 | Conclusion and Areas for Future Research | 122 |
| A | Θ- and δ-Function | 131 |
| B | Simulation Parameters | 133 |

| | |
|--|------------|
| C Abbreviations | 135 |
| D PhD-related Activities besides Writing the Thesis | 136 |
| E Lebenslauf | 137 |

Chapter 1

Motivation and Outline

1.1 Motivation

Auctions, though existing in the economy since thousands of years, have increasingly gained importance throughout the last decades. Nowadays, goods in the value of billions of Euros are allocated each day via auctions in a variety of economic settings. Accordingly, the challenge is for economists to provide a comprehensive theory of auctions.

Most of current auction theory rests heavily on the assumption that all participants play according to the Nash Equilibrium. But it is silent when it comes to explaining how bidders arrive there. The assumption of Nash Equilibrium play in a game can be motivated by the assumption that the participating players are all perfectly rational and that each player knows with absolute certainty that all other participating players are perfectly rational as well. In this thesis I assess the impact of dropping this assumption. As an alternative, I consider evolutionary learning: Either I still model each player as perfectly rational but assume that she doesn't necessarily think that all the other participants are perfectly rational as well. Instead she uses the strategies that would have generated the highest payoff in past play. This gives rise to best response dynamics. As an alternative I assume that the each bidder also makes errors in the individual decision process. This gives rise to quantal response. With these models I want to answer questions like the following: Can evolutionary learning justify the assumption of Nash Equilibrium play in repeated auctions? Will agents with initially heterogeneous strategies coordinate on Nash Equilibrium play? And if yes, which additional assumptions do we have to impose for the convergence and how realistic are they?

The thesis aims at unifying two research fields in economics that have so far been largely separated, namely auction theory and the theory of bounded rationality. Since there exists very few literature in this field, I have to develop the tools for this investigation in the course of this thesis. I use two complementary approaches:

On the one hand I develop a mathematical formalism that quantifies the average play under different assumptions about the bounded rationality of the bidders. This is a top-down approach that only works if the setup of the auctions – such as the number of bidders and the underlying value distribution – are known. From a methodological point of view, the mathematical analysis has to overcome two particular complications: On the one hand, bidders can't observe their opponents' strategies but only their bids. On the other hand I have to develop a formalism that allows the integration of discontinuous functions. In the thesis I apply this mathematical methodology to the analysis of several simple auction setups. This allows me to quantify bidding strategies under evolutionary learning.

As a complementary approach I program and employ an agent based simulation, the Auction Simulator. Agent based simulations are a bottom-up approach where the individual reaction of agents to their environment is accurately specified. The simulation then calculates the interactions of a number of agents in a common environment. By comparing the bottom-up simulation results with the top-down mathematical predictions that I derived for simple setups, I show that the simulation works correctly. After this quality-check I use it to reliably assess more general auction setups where the mathematical methodology would become too complicated

for the analysis. The simulation results are then analyzed by statistical methods. Simulation doesn't have to make the simplifications that mathematical analysis requires and therefore has the potential to describe auctions more realistically than currently possible. While allowing already for a variety of setups, the program is designed for easy extendability in the future. Thereby the Auction Simulator aims at providing a tool for a comprehensive analysis of auctions.

The main question that I set out to answer in this thesis is whether initially heterogeneous bidders in repeated auctions converge to the Nash Equilibrium. In order to give bidders the maximum chance of converging, I consider open bid auctions – i.e., auctions where each bidder gets to know the bids of all other bidders after the auction is over. This assumption is not necessarily satisfied in real-life auctions: Many auctions in the economy just publish the winner's bid. This makes no difference for the strategy updating of losing bidders, however the winner gets in such an auction no new information and therefore can not learn. This shows that open bid auctions give bidders the maximum chance to converge to the NE. The results of my analysis show that even in the open bid case best response bidders in general don't converge to the Nash Equilibrium in first-price auctions. However, they do converge in second-price auctions, no matter, whether an open or a sealed bid format is implemented. The intuitive reason for this is that in first-price auctions a bidder's strategy influences the opponents' optimal strategy choice so that bidders' strategies fluctuate in permanent mutual adaptation. In contrast, in second-price auctions it is a weakly dominant strategy to bid the true value independently of the opponents' strategies. Therefore bidders play their NE strategies, either by simple upfront-reasoning or by best response learning. In consequence, the allocation in first-price auctions is less efficient than in second-price auctions and the seller revenue is more volatile.

In generalization, the thesis suggests therefore that in a repeated setting auction formats with weakly dominant strategies are better suited to reach an optimal allocation of goods. Therefore, the thesis suggests to take a closer look at the auction mechanisms that are implemented in the economy and occasionally replace them. Since there are a lot of additional possible constraints on the efficiency of auctions such as for instance the prevention of bidding rings, recommendations on the optimal auction mechanism must depend on the specific economic environment. However, I conjecture that volatility in many markets could be substantially lowered by replacing value shading auction designs by auctions with weakly dominant strategies. 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 thesis is structured as follows:

Chapter 2 reviews results about auction theory and auction experiments from the literature.

Chapter 3 reviews concepts from game theory and surveys the usage of simulations to investigate markets in the economy.

Chapter 4 mathematically shows that two best response bidders with linear bidding strategies and values from a uniform distribution don't settle down in the NE in first price auctions (1PAs) while they do in second-price auctions (2PAs). I derive analytically how the bidders can calculate their best responses. Bidders can be modeled as being myopic in that they consider only last period's payoff. Alternatively they can be modeled as having perfect memory, where they use the strategy's aggregated payoff of the whole game. The quantitative analysis of these two extremes allows me also to predict the behavior of bidders with imperfect memory who give more weight to recent information than to information from the beginning of the game.

Chapter 5 discusses the consequences of relaxing certain assumptions in chapter 4. The qualitative results are shown to remain the same.

Chapter 6 mathematically analyzes the average bidding behavior of two quantal response bidders. Again, bidders can have various degrees of memory strength.

Chapter 7 shows how the newly developed mathematical methods can be applied to pure seller auctions.

Chapter 8 analyzes the stability of the NE in double auctions. To this end, it extends the mathematical analysis of linear bidding strategies for perfect memory best response dynamics to affine linear bidding functions.

Chapter 9 introduces the Auction Simulator. I discuss the program flow, the key parameters of the model and the alternative strategy updating algorithms that are used in the simulation.

Chapter 10 reports the simulation results for the setups where I have derived mathematical predictions in the previous chapters. It shows that the AS quantitatively reproduces these predictions and thus both, theory and simulation are double-checked for their correctness.

Chapter 11 analyzes the statistics of more complicated simulation setups that are too difficult to deal with analytically.

Chapter 12 concludes and gives an outlook to future research. It discusses some implications of the theory for the real economy. In particular, it argues that auction designs that give participants incentives to shade their values should be carefully revised and occasionally be replaced by auction designs with dominant strategies.

The thesis is completed by four appendices and a bibliography.

Chapter 2

Literature Background on Auctions

2.1 Private-Value Single-Unit Pure-Bidder Auctions

2.1.1 Standard Formats

In the symmetric independent private values framework (SIPV) of auction theory, two standard auctions are distinguished: the 1st price auction and the 2nd price auction.¹ In these auction formats, one seller faces several bidders. Each bidder has a private value that is not known to the other participants. Values are drawn from a common random distribution that is known to each bidder. Each bidder knows her own but not the other players' values. Depending 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.²

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 losing ones are published. In the one-shot auction, the publishing of information after the auction of course 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. In this thesis I am particularly interested whether repeated auctions converge to the NE. Therefore I will focus on repeated open bid auctions since they give bidders the maximum amount of possible information. I conjecture 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 the winning bidder has less information available to calculate her best response.

The first price sealed bid auction (1PSBA)³ is equivalent to the Dutch auction. There, the price of the object for sale is initially set to a very high level and then continuously lowered, e.g. by a clock that counts down the price. When the first bidder raises her hand, she acquires the object at the price that is currently shown by the clock. 1PSBAs and Dutch auctions are widely used throughout the economy: Procurement of companies and governments, the sale of agricultural products, land conservation auctions, IPOs and the privatization of government owned companies are all done via these auction formats.⁴

The 2POBA is roughly equivalent to the English auction where the price is increased until

¹A range of other formats is also considered in the literature: 3rd and higher price auctions where the winner has to pay the n^{th} highest bid, all-pay auctions where all bidders have to pay their bids - also the losing ones, etc. However, for most of the analysis I will focus on first and second-price auctions.

²In principle, also other allocation rules are possible such as the allocation to a certain bidder with a probability proportional to the bid of that bidder.

³All abbreviations can be found in the appendix

⁴Or in the multi unit equivalents of first-price auctions such as discriminatory or uniform-price auctions.

all but one bidder have dropped out.⁵ This auction type is typically used in art auctions or auctions of rare wines.

2.1.2 Revenue Equivalence

Friedman (1955) was the first to analyze bidding in 1PAs by mathematical methods. He analyzed repeated auction settings and proposed to use simple statistical methods to extract information on the bidding behavior of competitors. This information can be used by bidders to increase their expected profit in future auctions.

A theoretical break-through in the analysis of auctions was the Revenue Equivalence Theorem (RET). It was first published in Vickrey (1961) and then generalized independently by Myerson (1981) and Riley and Samuelson (1981).

The RET states that all private value single sided auction formats 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 bid⁶
- The cost of submitting the lowest feasible bid is equal 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 that all of these other players know that all the other players are perfectly rational, etc. ad infinitum. This implies that all bidders bid according to their Nash Equilibrium bidding functions

The last point is often not explicitly mentioned because it is implicitly 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 opponent is perfectly rational, the RET breaks down.

In the specific cases of 1PAs and 2PAs with n bidders and values drawn from a uniform distribution, the Nash Equilibrium (NE) bidding functions are given as follows:

$$\begin{aligned} \text{for 1PAs: } \quad \text{bid} &= \frac{n-1}{n} \cdot \text{value} \\ \text{for 2PAs: } \quad \text{bid} &= \text{value} \end{aligned} \tag{2.1}$$

I will prove this well known result subsequently.

The expected seller revenue is given by

$$\begin{aligned} \text{for 1PAs: } \quad \frac{n-1}{n} \cdot \text{E}[\text{highest value}] &= \frac{n-1}{n} \frac{n}{n+1} = \frac{n-1}{n+1} \\ \text{for 2PAs: } \quad 1 \cdot \text{E}[\text{highest value}] &= \frac{n-1}{n+1} \end{aligned} \tag{2.2}$$

which is a special case of the RET. The order statistics can be found in any statistics textbook.

A more general proof of the RET for arbitrary value distributions was first developed by Vickrey (1961). The exposition here closely follows Jehle and Reny (2001). They restrict bidders to symmetric bidding functions $b(v)$ that increase monotonically in v . Denote the probability density of the value by $f(v)$ and the CDF by $F(v)$.

⁵2PSBA, Japanese and English auction with private values are equivalent in that the unique equilibrium bidding strategy is bidding the true value of the object. However, they differ in the information dynamics in the case of affiliated values: In the English and the Japanese auction each bidder learns the bids of all losing bidders. This can change the outcome even in a one-shot auction in the case of affiliated or common values. In a repeated setting this holds true even more.

⁶The RET even holds for the more general assumption that the probability of winning with a certain bid is the same among the different auction formats.

Now, imagine that bidders do not report their bids but their values to the auctioneer who uses a publicly known payment rule that determines the allocation and the payments that the bidders have to make. The central part of Vickrey's idea is to find an incentive-compatible payment rule that induces truthful reporting of values by the participating bidders.

Imagine that a bidder in a first-price auction with n participating bidders reports a value of r while having a true value of v . Then her expected payoff is given as

$$EPO(r, v) = \underbrace{F^{n-1}(r)}_{\text{prob of winning}} \cdot \underbrace{\{v - b(r)\}}_{\text{payoff if she wins}} \quad (2.3)$$

Note the implicit assumption of homogeneous bidding functions. Only under this assumption can the probability of winning be written as a function of the reported value. If we allow for heterogeneous bidding functions, Vickrey's argument doesn't hold anymore. A great part of the remainder of this thesis will focus on the implications of allowing heterogeneous bidding functions in auctions.

Setting the derivative of equation (2.3) with respect to r equal to 0 at $r = v$ (because it shall be optimal to report truthfully) and solving for $b(r)$ yields the NE bidding function

$$b(v) = \frac{1}{F^{n-1}(v)} \int_0^v x dF^{n-1}(x) \quad (2.4)$$

In particular for a uniform value distribution with $f(v) = 1$ and $F(v) = v$ we obtain

$$b(v) = \frac{1}{v^{n-1}} \int_0^v x dx^{n-1} = \frac{1}{v^{n-1}} \int_0^v x(n-1)x^{n-2} dx = \frac{n-1}{v^{n-1}} \frac{v^n}{n} = \frac{n-1}{n} v \quad (2.5)$$

which proves the claim that I made in equation (2.1). For two participating bidders this amounts to bidding half their value in a first-price auction.

In contrast, bidders in second-price auctions have a (weakly) dominant strategy of bidding their true value: Bidding more than their value induces an unnecessary risk of negative payoff. Bidding less than the value on the other hand is suboptimal since then the bidder risks foregoing positive payoff.

The comparison of payoffs between first and second-price auctions is straight forward. In first-price auctions, the density of the highest value of n independent draws from a value distribution is given by $dF^n(v) = nF(v)^{n-1}f(v)$.

Therefore the expected seller revenue in a first-price auction with an arbitrary value distribution is given as

$$EPO_{1PA}^{\text{sell}} = \int_0^1 dv b(v) n F(v)^{n-1} f(v) = n \int_0^1 \int_0^v x dF(x)^{n-1} f(v) \quad (2.6)$$

since

$$b(v) = \frac{1}{F(v)^{n-1}} \int_0^v x dF(x)^{n-1} \quad (2.7)$$

In second-price auctions the crucial quantity is the density of the second highest value. The probability density that the value of a specific bidder i is v is given by $f(v)$. The probability that a specific bidder $j \neq i$ of the remaining $(n-1)$ bidders has a value above v while all the other $(n-2)$ bidders have values below v is given by $(1-F(v))F(v)^{n-2}$. Therefore the probability that any bidder $\neq i$ has a value above v while all the others have values below v is given as $(n-1)F(v)^{n-2}(1-F(v))$. Since there are n bidders, the probability density of the second highest value is therefore given by

$$g(v) = n(n-1)f(v)F(v)^{n-2}(1-F(v)) \quad (2.8)$$

Therefore the expected seller revenue in a second-price auction is given as

$$EPO_{2PA}^{\text{sell}} = \int_0^1 dv v n(n-1)f(v)F(v)^{n-2}(1-F(v)) \quad (2.9)$$

since in the Nash Equilibrium $b(v) = v$.

To show the equivalence of expected seller revenue in first- and second-price auctions consider the following series of transformations:

$$\begin{aligned}
 EPO_{1PA}^{\text{sell}} &= n \int_0^1 dv \int_0^v x dF(x)^{n-1} f(v) = n \int_0^1 dv f(v) \int_0^v x(n-1)F(x)^{n-2} f(x) dx = \\
 &= n(n-1) \int_0^1 dx \left(\int_x^1 dv f(v) \right) x f(x) F(x)^{n-2} = \\
 &= n(n-2) \int_0^1 dx (1-F(x)) x f(x) F(x)^{n-2} = EPO_{2PA}^{\text{sell}} \tag{2.10}
 \end{aligned}$$

Riley and Samuelson (1981) and Myerson (1981) independently generalize the result and show that in the independent private values framework all auction formats with the same allocation rule - which might be probabilistic - and equal participation cost yield the same expected seller revenue. Myerson (1981) in addition develops an auction that maximizes expected seller revenue. Interestingly this auction does not necessarily allocate the good to the bidder who values it most. Moreover it requires the seller to set a reserve price above her value so that possible mutually beneficial trades are prevented. So, even under the assumption of perfect rationality, there is a tradeoff between the efficiency of an auction and its ability to maximize expected seller revenue.

2.2 A Quick Tour through Auction Theory

Besides private-value single-unit pure-bidder auctions, there are various other auction types:

Seller Participation: In many auctions not only bidders but also several sellers might act strategically.

If only one bidder who wants to buy some good participates, but several potential sellers compete to sell the good, we face a pure seller auction. An example of this is are procurement auctions where several suppliers compete for a contract.⁷

In a double auction, a bunch of bidders faces a bunch of sellers. There, in addition to bidders and sellers competing against their peers on the same side of the market, they also have to act strategically against the opposite side of the market. A profitable deal requires the match of a bidder with a seller who asks less than the bidder bids.

Like in the pure bidder case, also the second-price pure seller auction has a dominant strategy of bidding the true value. For double auctions, McAfee (1992) designs a dominant strategy format. The price is set in the middle of the highest bid and the lowest ask that could not be matched. Note that this auction format might prevent the smallest profitable trade, so it is not fully efficient.

Multi-unit Auctions: Instead of bidders/ sellers who want to trade exactly one unit, bidders might want to buy/ sell multiple units. For the exposition here I focus on the pure bidder case, auctions with seller participation are straight forward generalizations of this setup.

In the discriminatory auction (DA) bids are ordered from the highest to the lowest. Starting from the highest bid, the demand of bidders is satisfied until the supply is exhausted. Each bidder has to pay her own bid. The DA is the equivalent to the first-price auction in the single-unit case.

In the uniform price auction (UPA), each successful bidder pays the highest unsuccessful bid.⁸

⁷Of course, a model that is only based on the contract value fails to capture other important determinants such as the expected quality of fulfilment

⁸The market clearing price in the UPA can in principle be set anywhere between the highest losing and the lowest winning bid. I will stick to the particular convention that the highest losing bid determines the payment.

The UPA is sometimes sloppily denoted as the multi-unit equivalent of the second-price auction in that bidders have a dominant strategy of bidding their true value. However, this is only true in the limit of infinitely many participating bidders or if only one unit is for sale. If only few bidders participate or if some bidders dominate the market by demanding the majority of units, the UPA loses its dominant strategy property. This is intuitively clear: A big player has an incentive to bid below her value. Even if she foregoes some profitable trades by bidding less than her opponents, she might in total still profit from the lower price she has to pay for the other units.

Ausubel (2004) designs an open bid format for multi-unit auctions with a weakly dominant strategy⁹. The seller continuously raises the price and bidders indicate their demand at the prevailing price. Whenever there is excess supply because a bidder reduced her demand, the other bidders get their residual supply¹⁰ at that price. A sealed bid dominant strategy multi-unit auction is the Vickrey multi-unit auction. In this format, a winner who wins k^i units pays the k^i highest losing bids of the other bidders.

Note that multi-unit auctions allow for a much richer structure than single-unit auctions since the demand for units might be interdependent: In the case of homogeneous goods, the willingness to pay for an additional unit might increase or decrease with the number of units the bidder has obtained so far. In the case of heterogeneous goods, goods can be substitutes or complements so that the willingness to pay for a certain unit depends not only on the number but also on the composition of goods the bidder has obtained so far. Much of current research focuses on the assessment of the impact of these dependency structures on the outcome of auctions.

Also for multi-unit auctions, there exists a Revenue Equivalence Theorem: In equilibrium, the expected payoffs (and payments) in multi-unit auctions with the same allocation rule differ at most by an additive constant

In addition to investigating these auction structures, considerable effort has been made to relax the assumptions of the RET one by one. Each of these extensions causes the RET to break down.

Risk Aversion: It is easy to see that risk aversion of bidders destroys revenue equivalence: In first-price auctions risk averse bidders weight the risk of not obtaining the asset heavier than if they were risk neutral. Therefore they bid higher and the expected seller revenue is raised in first-price auctions by introducing risk aversion of bidders. In second-price auctions, bidding the true value involves no risk, therefore there is no change in bidding behavior. Consequently, if we assume risk averse bidders, first-price auctions give more expected seller revenue than second-price auctions - the RET breaks down.

Asymmetric Value Distributions: Instead of having values drawn from the same distribution for all bidders/sellers, auctions behave significantly differently if this assumption is dropped. If values are drawn from different distributions, then Maskin and Riley (1999) show that strong¹¹ buyers prefer 2PAs, whereas weak buyers prefer 1PAs. Lebrun (1999) shows that for 1PAs with asymmetric value distributions there always exists an equilibrium. Fibich and Gaviols (2003) explicitly calculate such equilibrium bids in asymmetric 1PAs by using perturbation theory.

Value Dependencies: Instead of private independent values, values can be affiliated or common. Affiliated values can be positively or negatively correlated. To model affiliated values, each bidder i is assumed to obtain a signal x_i . Her true value depends then on all the other bidder's signals as well,

$$v_i = v_i(x_1, x_2, \dots, x_n) \quad (2.11)$$

If the value is the same for each bidder,

$$v_i = v(x_1, \dots, x_n) \quad (2.12)$$

we speak of a common value setting.

Common values give rise to the winners curse: Assume a jar full of coins is auctioned off and

⁹i.e., an auction format where the bidders' optimal strategy does not depend on the opponents' strategies

¹⁰i.e., the total supply minus the sum of amounts demanded by the other bidders, provided this is nonnegative

¹¹In the sense that a strong buyer's value distribution stochastically dominates the weaker buyer's distribution

bidders have a normally distributed estimate of the value. In each auction type, the highest bidder is likely to have estimated too much for the jar and therefore to make a loss - the winners curse. Consequently, bidders in all auction types should reduce their bids because they take this risk of overbidding into account. An interesting possibility to overcome the winners curse is given by the linkage principle. It states that sellers should commit to publish all information they have about the value of a good -*even the bad news*. The resulting narrowing of value estimates leads to a reduction in the winners curse and hence to higher bidding. Therefore expected seller revenue is raised by the release of information.

Obviously, in affiliated and common value auctions, the revenue equivalence between second-price and English auctions breaks down. This is because in open bid formats, bidders get information about their opponents' signals during the bidding process that influence their own value. Moreover, also revenue equivalence between first- and second-price auctions breaks down: Expected seller revenue from second-price auctions is at least as high as that of first-price auctions. For more information on auctions with interdependent values see, e.g., Krishna (2002).

Repeated and sequential auctions: Another current research topic is equilibrium bidding in sequential auctions.

Jeitschko (1998) analyzes equilibrium strategies in 1PSBAs with 2 rounds and 3 bidders. In each round there is one unit of a homogeneous good up for auction. Bidders can have either a high or a low valuation. Each bidder demands exactly one unit, so the winner of the first round drops out. Two effects are considered: Bidders can estimate their opponents' types from the first auction and condition their second round bidding strategy on this information. Additionally, bidders know that their first round bids influence their opponents' bids of the second round (low bids in the first round lead to lower bids in the second round) and therefore bid strategically. This induces a tradeoff between raising and lowering bids. In effect, prices may fluctuate between the two rounds but there is no unique result whether prices increase or decrease.

Hausch (1986) shows similarly for a multi-unit setting that the strategic anticipation of informational effects by bidders can lead to higher or lower expected seller revenue in sequential vs simultaneous auctions. This result is driven by the following tradeoff: The release of information about bidders' valuations in sequential auctions reduces via the linkage principle the winners curse and thereby increases seller revenue. On the other hand - as in Jeitschko (1998) - bidders in sequential auctions have an incentive to reduce their bid in order to reduce the others' bids in later rounds. This reduces seller revenue.

Jeitschko and Wolfstetter (2002) consider a two round auction where bidders demand more than one unit. The essential feature of this model is the investigation of the impact of (dis)economies of scale on seller revenue: A bidder's value for the second item depends on whether she has won the first round. The paper shows that under economies of scale, 2PAs are more profitable for the seller while the converse is not necessarily true.

McAfee and Vincent (1997) investigate sequential auctions where the seller posts a reserve price. If the seller does not commit not to auction off the good in a later round, bidders might react strategically and lower their bids below the reserve price in order to force the seller to lower the reserve price in a later round. The paper shows that as the time between auctions goes to zero, expected seller revenues converge to those of a static auction without reserve price. Note that in contrast to the other papers, here the same good is repeatedly auctioned in all rounds as might be the case for instance in the auctioning of rare wines.

Theoretical work on auction theory is virtually countless. A good informal overview is given by Klemperer (1999). More advanced introductions are given in Jehle and Reny (2001) and Wolfstetter (1999). A handbook that summarizes much of the current state of auction theory is Krishna (2002).

In summary, theoretical auction literature focuses on deriving the NE of specific auction setups. So bidders in these models are implicitly assumed to be perfectly rational and to know with certainty that each of their opponents is perfectly rational as well. In addition, bidders are assumed to know exactly the details of the auction setup such as the number of participating bidders and the underlying values. This assumption will be dropped in my thesis: Best response play corresponds to dropping the assumption that each bidder knows with certainty that each opponent is perfectly rational while she is still assumed to be perfectly rational herself. Quantal response in addition assumes deviations from perfect rationality for the individual

bidder. Therefore, the literature doesn't provide much insight into the questions that I want to investigate.

2.3 Auction Experiments

There also exists a vast literature on auction experiments. The bottom line of this literature is that the NE is a good proxy for describing human behavior in auctions. However, there are significant deviations from the predictions which force the experimenters to refrain from the assumption of perfectly rational bidders.

Among the oldest and the most influential papers are the experiments described in Cox, Roberson, and Smith (1982), their criticism in Harrison (1989) and the following discussion in the *American Economic Review*, see Friedman (1992), Kagel and Roth (1992), Cox, Smith, and Walker (1992), Merlo and Schotter (1992) and Harrison (1992).

Cox, Roberson, and Smith (1982) describe experiments where bidders in 1PSBAs tend not to bid according to their NE bidding functions but to overbid. This result also held in an experimental setup that compensated for risk aversion. Harrison (1989) argues that in the specific experimental setup the expected payoff losses when not bidding according to the equilibrium bidding function were too small to compensate for the increased deliberation cost. This is termed as the "flat maximum critique". Accordingly, bidders in these experiments act boundedly rational. When evaluated in payoff-space instead of the space of bidding strategies, the deviations from theory were not significant. This comment triggered one of the fiercest discussions in experimental economics up to date because many experimentalists saw the article as a severe criticism to their methods from an outsider to the field.

A more comprehensive survey of auction experiments up to 1995 is given in Kagel and Roth (1995). They show that auction experiments show in general persistent deviations from equilibrium predictions. In particular, in contradiction to the predictions of the RET, all auction formats yield different expected seller revenue:

1PSBAs yield more revenue than Dutch auctions. In both formats bidding is significantly above the NE prediction. Bids in 2PSBAs are also significantly above the NE prediction whereas repeated English auctions quickly converge to the NE of bidding the true values.

Kagel explains this finding by assuming that in 1PSBAs and 2PSBAs bidders focus on the price dimension of the auction whereas the accept/reject problem in Dutch and English auctions leads them to focus on the profitability dimension thus pushing them closer to the NE. For the efficiencies they report that 88% of 1PAs and 2PAs are efficient compared to 80% of Dutch auctions. English auctions yield the optimal efficiency of $\sim 100\%$.

Similarly, in experiments for multi-unit auctions also persistent deviations from the NE were observed, see e.g., Cox, Smith, and Walker (1984).

Kagel and Roth (1992) report that in 1PSBAs, the outcome depends on the number of participants: For relatively few players there is overbidding and for relatively many players underbidding when compared to the NE strategies.

Goeree, Holt, and Palfrey (1999) explain the experimentally observed overbidding in first-price auctions by a quantal response equilibrium under constant relative risk aversion. They investigate also two alternative explanations, namely nonlinear probability weighting and a model based on the "joy of winning". While the first alternative provides a reasonable good fit to the data, it does not show the expected "inverted S-shape" corresponding to the overweighing of low risk probabilities and underweighing of high probabilities. Therefore, the authors prefer the QR/CRRRA model. The "joy of winning" model performs significantly worse than the other two.

Engelmann and Grimm (2003) report on experiments for multi-unit auctions. While value shading auction designs like discriminatory or uniform price auctions show significant deviations from the NE, bidding behavior in dominant strategy auction formats like Vickrey or Ausubel auctions are close to the theoretical predictions.

Pezanis-Christou and Sadrieh (2004) report on experimental evidence that bidders in symmetric first-price auctions follow a best reply dynamics. The authors let bidders specify in each round a stepwise linear bidding function with two value segments before their value is drawn: Bidders decide on the slope β of the bidding function separately for the two value segments. They

consider 2 asymmetric and one symmetric treatment. For each bidder in each treatment they determine the empirical best reply function. By this they denote the risk-neutral best reply bid function to the distribution of the rivals' actual bid functions. Then they investigate whether the shapes of observed bid functions (concave, convex, humped or linear) can best be matched by empirical best reply functions, by risk-neutral NE bidding functions or by risk-averse NE bidding functions. Their result is that the shapes of the bid functions in the different treatments are best matched by assuming best reply behavior of bidders.

Garrat and Wooders (2004) reinvestigate data that were originally reported by Kagel. They show that in experimental second-price auctions, bidders tend to overbid, but auction efficiency is rising over time since players learn to coordinate on the same strategy above the NE.

In conclusion, the literature on auction experiments shows that the NE is a good proxy for behavior of bidders in auctions. However, there are significant systematic deviations from it. This indicates that bidders in experimental auctions act boundedly rational. In the next chapter I will therefore briefly survey concepts from bounded rationality.

Chapter 3

Literature Background on Game Theory

3.1 Motivation

As already noted before, this thesis is an attempt to construct mechanisms that lead to 'optimal' economic systems. In particular I look for mechanisms that increase the allocative efficiency of auctions and markets. This property goes hand in hand with a decrease of risk and instability in these economic systems. Of course, these claims rest on several assumptions.

The most important assumption is about the nature of rationality of the participants. Classic economic theory assumes each agent to be perfectly rational and to know that all the other agents are also perfectly rational. Proofs are literally made under the assumption that each person in the world, each decision maker in a firm would act in exactly the same way if you ask them to solve the identical problem and that everybody knows this fact.

But real economic agents are not perfectly rational. They have different experience, different aims and pursue different strategies. The concept of bounded rationality tries to take this into account by relaxing some but not all the assumptions of standard game theory.¹

In this thesis I investigate the efficiency of auctions under different relaxations of agents' rationality: Agents might stay perfectly rational (best response or fictitious play) or make errors in their responses (quantal response). However, no agent can be sure that the other participants are necessarily perfectly rational. Therefore, agents in auctions can *not derive* how their opponents will act in the *future*. Instead, I will focus as an alternative concept on evolutionary learning where players *observe* the success of their strategies in the *past* in order to choose their strategy for future play.

3.2 Bounded Rationality

3.2.1 Reasons for Bounded Rationality

Besides the observation that in reality people do not act perfectly rational, there are also a number of more formal arguments that cast doubt on the underlying concepts of standard game theory:

Rubinstein (1991) gives a wonderful example that shows the problem of merely identifying and including all relevant information into the game solving process: He considers the battle of sexes with payoffs of 3 and 1 if coordination is reached. This game has three possible equilibria: There exist two pure strategy equilibria where players coordinate and one mixed equilibrium. However, including the possibility that player 1 throws away one dollar at the start of the game

¹Obviously, if we relax each and every assumption on agent rationality we can not predict anything interesting.

leads to the unique selection of the most favorable equilibrium for player 1 by iterated weak dominance arguments. In the final outcome player 1 actually doesn't dispose of the dollar, so the change in the outcome just comes from assumption that player 2 has thought of the possibility that player 1 disposes of one dollar. But one can never be sure to have thought of all relevant influence factors for the outcome of a game and this severely undermines the power of the NE solution concept.

A battery of arguments are given by Conlisk (1996). He structures his criticism around two main points:

First, it has repeatedly been seen that subjects in economic experiments fail to behave perfectly rational. Instead, it is often observed that the subjects use rules of thumb and avoid thinking through the consequences of their decision in all detail. Deviations from perfect rationality can also be seen in 'real' economic systems.

Second, the assumption of perfect rationality fails to take deliberation costs into account and therefore is logically inconsistent: The time, computational power and research needed to find the perfectly rational solution are not for free. By incorporating deliberation cost into perfectly rational behavior, one runs into an infinite regress problem: Imagine, players try to determine how deeply they should think through some problem before they make their next move. Then in effect, this is an optimization problem where players try to balance the trade-off between expected payoff and deliberation cost. But the deliberation cost of solving this optimization problem is no yet taken into account. Taking it into account results in a higher-order optimization problem which again has to be solved in some resource intensive way.

These shortcomings of the concept of perfect rationality of course haven't gone unnoticed by the developers of game theory as reported in Goeree and Holt (1999), p. 10.564:

There is, however, widespread criticism of theories based on the "rational choice" assumptions of perfect decision making (no errors) and perfect foresight (no surprises). This skepticism is reinforced by evidence from laboratory experiments with financially motivated subjects. Nash participated in such experiments as a subject and later designed similar experiments of his own but lost whatever confidence he had in game theory when he saw how poorly it predicted human behavior. And Reinhard Selten, who shared the 1995 economics Nobel Prize with Nash and Harsanyi remarked that "game theory is for proving theorems not for playing games."

In Goeree and Holt (2001), the authors formulate their criticism on game theory even sharper:

"The rationality assumptions that underlie this analysis are often preceded by persuasive adjectives like "perfect", "intuitive" and "divine". If any noise in decision-making is admitted, it is eliminated in the limit in a process of "purification". It is hard not to notice parallels with theology, and the highly mathematical nature of the developments makes this work about as inaccessible to mainstream economists as medieval treatises on theology would have been to the general public."

I agree with the criticism that the NE concept is somewhat fishy. Certainly, the NE concept is a powerful tool to determine the general direction of the outcome of games. But it is inappropriate to assess the finer subtleties of player behavior. Only in one respect I disagree with the authors: Accounting for bounded rationality in games surely doesn't make the mathematics easier.

As an alternative to the assumption of perfect rationality, theory has in the last 15 years developed a theory of evolutionary learning. This research area has its roots in two quite distinct scientific fields: Economics and theoretical biology. The main idea in evolutionary learning is that players in repeated games use simple rules to update their strategies. The resulting dynamics can be amazingly complex. One of the main questions in that context is whether evolutionary learning converges to the NE (convergence) and - if there exists more than one NE - to which one (equilibrium selection).

3.2.2 Concepts from Evolutionary Game Theory

The following section closely follows Fudenberg and Levine (1999). Short overviews of evolutionary game theory are also given in Sigmund and Young (1995) and Fudenberg and Levine (1996).

Myopic best response: Under myopic best response, in each period each player plays her best response against the observed behavior of her opponent in the last round.² The underlying story is that once a player has formed an opinion about the strategy of the other player in a particular round, she believes that her opponent will stick to that strategy also in the next round. So, agents are boundedly rational with limited foresight.

The conceptual problem with this story is that each player implicitly supposes that she is more intelligent than her opponent, since she assumes, she can change her strategy while the opponent doesn't. Even if a sophisticated player 0 assumes an intelligent opponent who will play a best response against her own last move and plays a best response against this strategy, she implicitly assumes that she can think through more iterations than her opponent.

Nevertheless, (near to) myopic strategy updating has been experimentally observed in various settings.

Fictitious Play: The basic idea of fictitious play (FP) is that players keep track of the whole game and play their best response against the historical frequency distribution of opponents' strategies. Originally fictitious play has been proposed as a method for NE calculation. However more recently it has come into its own right as a learning model.

Formally, each player i has an exogenous initial weighting function, mapping the opponents' strategy space onto the real numbers, $\kappa_i^0 : S^{-i} \rightarrow \mathcal{R}_+$. κ_i^0 indicates player i 's prior beliefs that her opponents play a particular set of strategies. The weight is updated each time, the opponents make a particular move by

$$\kappa_i^t(s^{-i}) = \kappa_i^{t-1}(s^{-i}) + \begin{cases} 1 & \text{if } s_{-i}^{t-1} = s^{-i} \\ 0 & \text{if } s_{-i}^{t-1} \neq s^{-i} \end{cases} \quad (3.1)$$

The probability that player i assigns to the event that the players $-i$ play s^{-i} at date t is given by

$$\gamma_i^t(s^{-i}) = \frac{\kappa_i^t(s^{-i})}{\sum_{\tilde{s}^{-i} \in S^{-i}} \kappa_i^t(\tilde{s}^{-i})} \quad (3.2)$$

Fictitious play is defined as a rule that assigns a best response to the assumed probability distribution of opponents' strategies.

Later in this thesis I will analyze fictitious play in auctions. In auctions, no player knows her opponents' values, therefore no player can observe her opponents' bidding strategies directly. I will assume then - in analogy to FP - that players at each strategy updating choose their own strategy that performed best against the historical bid distribution of her opponents. So, she plays the strategy that would have maximized her payoff if she had played it for the whole game up to the current round.

In the thesis I will alternatively use the terms fictitious play and perfect memory best response to describe the same behavior.

Truncated Fictitious Play and Imperfect Memory Play: Obviously, myopic play and fictitious play are two extremes that do not capture the behavior of real people. In reality, people will mainly act based on information from the last few rounds. To model this behavior, I consider two possibilities:

As a first possibility, players could consider only the last n rounds and update their strategies according to fictitious play over this period. In other words, they weight information from the last n rounds equally and neglect all information from rounds longer ago. This is called truncated fictitious play. Note that $n = 1$ corresponds to myopic play. If, on the other hand,

²Note that the term "myopic play" is often also used to describe behavior that tries to maximize just next round's payoff and doesn't allow for more sophisticated strategic behavior. However, when I use the term "myopic" in the thesis, I refer to agents who only look at information from the last round to update their strategies.

$n = N$ where N is the number of rounds played in the game so far, then we rediscover fictitious play.

As an alternative I offer the model of imperfect memory play. I have not found a comparable model in the literature, however it seems to me at least as appropriate as truncated fictitious play for interpolation between myopic and fictitious play (=perfect memory play).

The main idea of imperfect memory play is that the importance that players give to past information gradually dies out³. In auctions, the information of interest is the payoff that each strategy generated in the past. I denote the payoff that strategy j would have generated t as $PO^{(t)}(s_j)$. The cumulated payoff under imperfect memory is then given as

$$PO_{\text{imp}}^{(t)}(s_j) = PO^{(t)}(s_j) + \lambda PO^{(t-1)}(s_j) + \lambda^2 PO^{(t-2)}(s_j) + \dots \quad (3.3)$$

Updating to the strategy that maximizes $PO_{\text{imp}}^{(t)}$ yields myopic best response for $\lambda = 0$, fictitious play for $\lambda = 1$ and a smooth interpolation for $\lambda \in (0, 1)$. Under imperfect memory play memory gradually dies out according to an exponential function.⁴ As much as I know from the psychological literature, this captures the decline of players' memory more realistically than truncated fictitious play.

Quantal Response: Under quantal response (QR) each player chooses next round's strategy with a probability proportional to past payoffs of that strategy. Mathematically, this behavior can be modeled as a probit rule

$$\text{selection probability of strategy } s_i = \frac{PO(s_i)^\xi}{\sum_j PO(s_j)^\xi} \quad (3.4)$$

or as a logit rule

$$\text{selection probability of strategy } s_i = \frac{e^{\xi PO(s_i)}}{\sum_j e^{\xi PO(s_j)}} \quad (3.5)$$

In this thesis I will confine myself to probit rules. Quantal response can be thought of as noisy best response with a specific noise structure. Goeree and Holt (2001) report a number of experiments for one-shot games for which QR explains findings that are in contradiction to perfect rationality.

Like in the case of best response, one can model quantal response as being myopic, with perfect memory (like fictitious play) or - as an interpolation - with truncated fictitious play or imperfect memory play.

Replicator Dynamics: The concept of replicator dynamics (RD) originally stems from the field of evolutionary biology. There, RD is used to model the evolution of species with different - in general mutually interdependent - fitnesses. RD assumes that on average the population of a species grows according to its fitness.

In learning theory, RD is used to model social learning where agents learn by observing other agents. Alternatively, RD can also be seen as the outcome of individual learning by a stimulus-response model. In such models, species are interpreted as specific behaviors that a player can adopt and that give rise to different payoffs. Well performing actions are reinforced and so are more likely to be used again in the future.

It can be shown that under the RD the probability that a specific strategy is played grows or shrinks exponentially over time. The exponent is proportional to the difference between the fitness of that strategy and the mean fitness of the population of all strategies. Holland's Schema Theorem (see Holland (1992)) proofs that Genetic Algorithms (GAs) give rise to the same dynamics: Schemata grow exponentially to their relative fitness. Therefore GAs are an apt tool to simulate RD. I will say a little bit more on GAs and the Schema Theorem in section 3.4.1. RD can be generalized to various other classes of processes that use some other payoff-monotone measure to determine the success of individuals. These models qualitatively all exhibit the same dynamics.

Under RD, any stable state that is the limit of a path that originates in the interior of the

³As opposed to truncated fictitious play that leads to abrupt changes in the information sets that the players use for strategy updating.

⁴since the weight of $PO^{(t-i)}$ is λ^i

strategy space is a NE.

A concept closely related to RD is that of an evolutionary stable strategy (ESS). Intuitively an ESS is an equilibrium that is able to repel invaders. This means that if the population is in an ESS and a small part of the population starts to play a deviating strategy, then the payoff of the old strategy against the new resulting mixture of strategies is higher than the payoff of the new strategy against this mixture.

Introspection: None of the above models allows for the fact that players use introspection to come up with better strategies just to play some response to past play. By introspection I denote the chain of thinking "I play the best response to your best response to my best response to \dots ". In the context of repeated games the assumption of no introspection can be justified by assuming that bidders act under severe time constraints (like e.g., traders in Financial Markets) or that they do not believe in the intelligence of their opponents. Goeree and Holt (2001) propose to model introspection in one-shot games by assuming iterated quantal response (since humans make errors) with increasing noise (lower ξ) for higher iterations.

Strategic Reasoning: Most papers in evolutionary game theory don't allow for strategically acting players who take into account the consequences of their own actions. Evolutionary learning focuses only on the *past* as opposed to strategic reasoning that focuses on the *future*. In repeated first-price auctions, for instance, bidding low in the first few rounds - no matter what the opponent does - could be a feasible strategy since this could lead the opponent to also bid lower and therefore to larger profits for both in later rounds. Second-price auctions are even more prone to collusion as it is a NE if the bidder with the highest valuation bids truthfully and all the others bid 0. However, the highest bidder would have to be determined in a (presumably illegal) pre-auction. In frequently repeated auctions, there presumably is no time for meeting secretly and running these pre-auctions.

Evolutionary game theory mostly neglects strategic reasoning by considering models where players come from a large population and are matched randomly with each other. As the population size goes to infinity the chance that the same two players will be matched again goes to zero and therefore this kind of strategic reasoning is not taken into account. Models where the same players are matched over and over again are referred to as "repeated matching".

Evolutionary Learning and Auctions: Most of the literature on evolutionary learning deviates in two important aspects from setups that would allow conclusions on repeated auctions.

- First, the literature focuses on games where players can observe their opponents' strategies exactly. In auctions, however, bidders can only observe their opponents' bids but not their opponents' strategies, since they have no information on the opponents' (private) values. This implies that in repeated auctions, bidders can base their strategy updating only on the payoffs that their own strategies have reached in the past.
- Second, the literature focuses on games with finitely many strategies (usually 2). In auctions, however, there are infinitely many possible bidding strategies for each bidder. Even in the restricted case of linear bidding functions that I will focus on in the thesis, there is a continuum of feasible bidding strategies namely the interval $[0, 1]$ for 1PAs and $[0, \infty)$ for 2PAs.

In the next section I discuss the sparse literature that exists on auctions under bounded rationality. Note that most of it doesn't deal with evolutionary learning.

3.3 Auction Theory under Bounded Rationality

To reconcile auction theory with experimental observations, theory has to be extended to describe bidders that do not bid according to their equilibrium bidding functions. However, literature about non-equilibrium bidding in 1PSBAs is sparse.

Most of the literature focuses on deviations from perfect rationality in one-shot auctions. Battigali and Siniscalchi (2003) investigate the general properties of non-equilibrium bidding in 1PSBAs. They show that for private or affiliated values every bid below the equilibrium and some bids above the equilibrium are interim-rationalizable (=consistent with thinking through

finitely many best responses). The paper argues that bounded rationality may explain the experimentally observed deviations from NE bidding. The paper makes no attempt to explicitly formulate strategy updating rules and investigate their dynamics. The contribution of the paper is an existence proof that can not be used for quantitative calculations.

Deltas and Engelbrecht-Wiggans (2001) investigate common value ascending price auctions – i.e., English auctions – with entry costs. There are infinitely many fully rational expert bidders and one non-expert⁵ bidder. The expert bidders know which questions to ask from the seller and thereby get to know the exact value of the auctioned asset. The non-expert does not know how to accurately assess the value and therefore gets her value from a random distribution with mean v . The authors show that the presence of the non-expert can deter all the infinitely many rational bidders from entering the auction and leaves positive expected payoff only to the non-expert. So, in equilibrium, the population of bidders in repeated auctions will partly consist of non-expert bidders.

In a related paper, Deltas and Engelbrecht-Wiggans (2000), the authors investigate again a common value setting. They show that bidders who don't fully understand the winners curse might have more expected payoff than fully rational bidders who do understand it.

The only paper I found about evolutionary learning in repeated auctions is Hon-Snir and Sela (1998). They investigate repeated 1PSBAs where bidders' values are determined before the first round and held constant thereafter. Under a variety of learning rules the outcome is the same as in a one-shot 1PA where the bidders' types are commonly known. The results, however, rely critically on the fact that bidders' values don't change between auctions. No result for repeated auction where valuations change between auctions is derived.

3.4 Simulations

As we saw above, the literature on auctions focuses mainly on equilibrium bidding. This is in accordance with standard game theoretic assumptions on perfectly rational players. Once we abandon these assumptions, the analysis becomes cumbersome or simply intractable.

As an alternative to mathematics, researchers have in the last 20 years increasingly employed computational methods for the investigation of heterogeneous agents. Computers have turned out as an indispensable tool to investigate systems that are too complicated to analyze them analytically (see, e.g., Casti (1997)). In this section I want to give a quick introduction into market simulations.

3.4.1 Genetic Algorithms

Genetic algorithms (GAs) are a computational tool to determine the solution to optimization problems. A good introduction is given in Mitchell (1999). Other introductory books are, e.g., Goldberg (1989) and Michalewicz (1999). In the standard setting proposed in Holland (1992), a genetic algorithm (GA) is a mechanism that determines the evolution of a population of bit strings.

A bit string is a representation of a number in the binary system. Consider for instance the 16-digit bit string 0110001001001101 and assume, the value of the i th position is $\frac{1}{2^i}$ where positions are numbered from 0 to 15. Let $j(i)$ denote the i th bit, i.e., $j(i) = 0$ or $j(i) = 1$. Then the value of this string in real numbers is $\sum_{i=0}^{15} j(i) \frac{1}{2^i} \sim 0.7697$.

The real values encoded in the bit strings are tried as potential solutions to the optimization problem under consideration. The relative success of the bit string in solving the problem determines its fitness.⁶ The fitness is used by the GA to construct a new generation of strategies from the old one. For this, usually the crossover and the mutation operator are used.

The crossover operator selects two bit strings as parents with probability proportional to their fitness. Then it constructs a new bit string by cutting both bit strings at the same random

⁵The authors actually use the term "inexpert", however I did not find it in any dictionary.

⁶E.g., if the problem is to find the argument that maximizes a function, the fitness of a bit string could simply be defined as the value the function takes at the real value encoded by the string. In computational auctions bit strings are interpreted as bidding strategies. The fitness of the bit string is then defined as the payoff, the strategy would have generated in past play.

position and joining the first part of the first bit string with the second part of the second bit string. The mutation operator subsequently flips each position in the bit string with a low probability.

The good performance of bit string GAs is theoretically justified by Holland’s Schema Theorem, see Holland (1992). The Theorem shows that the success of GAs in solving complex problems lies in its ability of massive parallel computation:

A schema is a template made up of 0s, 1s and don’t care-symbols ‘*’. Each bit string of length l belongs to 2^l schemata.⁷ By explicitly evaluating the fitness of one bit string, the GA implicitly evaluates 2^l schemata. Given the number of instances of a certain schema in the current population⁸, Hollands’ schema Theorem gives a lower boundary on the expected number of instances in the next generation. The Theorem accounts for the genetic operators that influence the number of instances. It proves that if a Genetic Algorithm follows a fitness proportional selection rule, then, the dynamics of instances of a certain schema in the population of a GA satisfies the following inequality:

$$E[m(H, t + 1)] \geq m(H, t) \cdot \underbrace{\frac{f(H, t)}{\bar{f}(t)}}_{\text{Selection}} \cdot \underbrace{(1 - p_m)^{\mathcal{O}(H)}}_{\text{Mutation}} \cdot \underbrace{\left[1 - p_c \frac{\mathcal{L}(H)}{N - 1} \left(1 - \frac{m(H, t)f(H, t)}{M\bar{f}(t)}\right)\right]}_{\text{Crossover}} \cdot p_d(H, t) \quad (3.6)$$

The variables in (3.6) have the following meaning:

| | |
|------------------|---|
| $E[m(H, t + 1)]$ | expected number of strings matching the schema H at generation $t + 1$ |
| $m(H, t)$ | number of strings matching H at generation t |
| $f(H, t)$ | mean fitness of instances of H at generation t |
| $\bar{f}(t)$ | mean fitness of strings in the population at generation t |
| p_m | mutation probability |
| $\mathcal{O}(H)$ | order of H (the number of defining bits $\neq *$ in the schema) |
| $\mathcal{L}(H)$ | defining length of H (length between the outermost defining bits in the schema) |
| p_c | crossover probability |
| N | length of bit strings in the population |
| M | population size |
| p_d | probability of disruption if crossover is applied |

The intuition for (3.6) is straightforward: The number of instances of a certain schema H rises with the relative fitness of its instances, because then the schema is more likely to be selected for cloning or crossover. The higher the mutation rate is and the more defining bits H has, the more likely it is to be destroyed under mutation. The probability that the schema gets destroyed by crossover rises with the crossover probability and also with the defining length. The longer the bit string is, the more likely it is that crossover happens outside the schema so that it doesn’t get destroyed. And finally, if crossover happens between two instances of a schema, the offspring will also be instances of the schema.

Note that (3.6) only gives a lower boundary for $E[m(H, t + 1)]$. The reason is that it does not account for the schemas H that are *created* by the genetic operators. However, the recombination of low-order schemata to increasingly fitter high-order schemata by crossover is believed to be one of the major sources of GA power. This hypothesis is commonly known as the Building Block Hypothesis. Therefore, (3.6) fails to take into account of one of the main sources of the power of GAs. Still, the Schema Theorem serves as one of the main foundations for the analytic investigation of GAs.

One of the immediate predictions of the Theorem is that for infinite population sizes, $m(H, t)$ will grow at an exponential rate with t ⁹. Under certain assumptions it can be shown that this behavior implies a tradeoff between exploration and exploitation of the search space that leads

⁷For example, the bit string 110 belongs to the 8 schemata 110, *10, 1*0, 11*, **0, *1*, 1**, ***.

⁸i.e., the number of bit strings belonging to that particular schema

⁹To be precise, $m(H, t)$ grows exponentially as long as $\bar{f}(t)$ is constant.

to an optimal online performance of GAs.¹⁰ In other words, the Schema Theorem ensures us that the GA works in a way that maximizes the mean fitness of the population over time.

In summary, the instances of a particular schema in the population grow proportionally to the difference of the fitness of the schema to the average fitness of the population. This property makes GAs a good tool to simulate the replicator dynamics.

3.4.2 Auction Simulations

Andreoni and Miller (1995) report the outcome of simulations of repeated pure bidder auctions. They investigate auctions with 4-8 bidders and 40 strategies each. Learning is simulated by GAs with fitness-proportional selection. This corresponds to QR learning. Each thousand auctions the GA updates the strategy population according to the cumulated payoffs. So, the simulation models QR with near to perfect memory. The paper analyzes auctions with private, affiliated and common values. In all cases, the simulation qualitatively replicates the experimentally observed behavior of human players. In particular, the simulation results in overbidding for first- and second-price auctions in an independent private values setting. The paper argues that usually experimental auctions consist of too few auctions¹¹ to assess many of the interesting aspects of the long-term behavior.

Dawid (1999) investigates a double auction market (repeated double auctions) where bidders and sellers have fixed deterministic values/cost that are the same in each auction. This setup possesses a continuum of equilibria. The question is to which equilibrium the game converges - if it converges at all. Dawid attacks this problem complementarily: He simulates the market via GAs and analyzes their convergence properties mathematically by a kind of "mean field approach" based on the analysis of the replicator dynamics.

The simulation results show convergence to an equilibrium where bidders' bids are slightly above sellers asks. However, the equilibrium selection¹² is not unique but depends on the random seed of the simulation: The price may be in favor of the sellers or the bidders in different simulation runs.

Note that in this simulation, each bidder is in each round randomly matched with a seller. Though the title of his paper suggests so, I would say that Dawid does not consider a standard double auction since in double auctions, bids and asks are ordered by size and then matched. Therefore, each participants' optimal strategy depends on all the other participants' strategies as well. In Dawid's model, however, each bidder's optimal strategy depends only on the population of sellers but not on the other bidders' strategies. The same holds true for the sellers. I would rather interpret his simulation as repeated two-sided ultimatum bargaining than as a model of a double auction. Note moreover that the simulation does not work with value distributions but with fixed values that are same among all bidders and all sellers for all rounds.

Another stream of research focuses on the simulation of new auction designs to design them "optimally" in some specified way. An example is Byde (2002) who proposes a mixture of 1PAs and 2PAs in order to maximize seller revenue.

3.4.3 Financial Market Simulations

Though financial market simulations are a little bit off-topic in the context of auctions I want to give here a short overview over this field. The first reason is that this financial markets are one of the major playgrounds for agent based modeling (ABM) in economics. Therefore it might give readers who are not so acquainted with ABM an impression of how these models work and how they are used. Since the Auction Simulator, that I introduce later in the thesis, is also an ABM this might simplify the understanding of the program philosophy. The second reason is that in the future I want to couple my simulation of auctions to a financial market simulation in order to investigate the feedbacks between order book design, technical trading and market volatility. I will say more on this in the last chapter. Therefore I want to shortly explain the

¹⁰Holland derived an analytic solution to the optimal playing strategy for multi-armed bandit problems and showed that the GA replicates this strategy.

¹¹Usually in the order of 18-30 consecutive auctions

¹²i.e., the mean of bids and asks

relevant concepts of financial market simulation here for later reference. The part can be left out upon first reading.

FM simulations mostly focus on the interaction of few different types of investors. In most models, there are two flavors of investors: Fundamentalists assess the intrinsic asset value by some method such as discounted future dividends. Fundamentalists believe that the price will move into the direction of the fundamental value, so they buy if the price is below the fundamental value and sell if it is above. In this sense they help to stabilize the market.¹³ On the other hand, chartists use technical trading rules to exploit past information on price and volume movements for the prediction of future prices. These traders generate bubbles and crashes by herding and thus introduce instability into the market. The assumption of traders who base their trading activities not only on the fundamental value but also on the strategies of the other traders is empirically justified by the observation that price movements in markets are often governed by fads which lack any fundamental reason (see, e.g., Shiller (1999)). The interaction between fundamentalists and chartists generates in these simulations often realistic features such as leptokurtosis and heteroscedasticity of the time series of returns or cross-correlation of volatility and volume. The first model of this kind was the Santa Fe Artificial Stock Market Model (SFASM), see e.g., Arthur, Durlauf, and Lane (1997). Technical trading is simulated there by classifier systems. Classifiers are descriptors of past states that consist of sequences of *s, 0s and 1s. Their evolution is closely related to that of a population of schemata which makes this model a playground for GA modeling. Joshi, Parker, and Bedau (2002) show that the SFASM operates in two states. The chaotic state qualitatively reproduces stylized facts of real markets in the return series. It is reached as updating frequencies of GAs are increased. This is intuitively clear since frequent updating of trading strategies implies frequent reallocation of capital and thereby more pronounced price movements.

The last ten years have seen the growth of a rich ecology of different financial market simulations. An overview of some of the most important models is given in Lux, Samaindou, Stauffer, and Zschischang (2004).

Most financial market simulations do not focus too much on the details of the order book. A notable exception is Chiarella and Iori (2002). They simulate a limit order market where orders arrive at random times. The matching of orders is explicitly modeled. Demands of traders are modeled as consisting of three components which are heterogeneous among agents: A fundamentalist component (with fixed fundamental value), a chartist component (extrapolating past returns) and a noise component. The authors are able to replicate realistic market features such as clustered volatility or cross-correlation between volatility, trading volume and bid-ask spread.

Another interesting paper is Daniels, Farmer, Gillemot, Iori, and Smith (2003). They focus on deriving properties of the order book without allowing for strategic acting of agents. Orders are coming randomly to both sides of the market. The authors define dimensionless characteristic quantities for the number of shares, the price interval and time. By complementary mathematical analysis for simplified cases (logarithmic prices for limit orders are uniformly distributed on the interval $(-\infty, \infty)$) and simulation for more general setups they are able to replicate properties of real financial market order books like the concavity of the price impact function. However, the simulation results also show different power laws for short- and long term volatility which the authors can not explain by their mathematical model.

¹³Note that most models assume one fundamental value that is in each time step the same for all investors. It would probably be more realistic to model them in each time step as value distributions so that fundamental values differ among investors in each round.

Chapter 4

Best Response in Open Bid Auctions

We saw in the last chapters that current auction theory largely fails to capture the effects of bounded rationality. The main contribution of this thesis is to marry these two research fields. In the first part of my own research, that starts with this chapter, I analyze mathematically the behavior of boundedly rational bidders in repeated auctions.

4.1 Structure of this Chapter

In this chapter I investigate the best response dynamics in 2 player open bid auction where values are drawn anew in each round and come from a uniform value distribution, $v_i \sim U(0, 1)$. I am in particular interested, whether bidders with initially heterogeneous bidding functions converge to the Nash Equilibrium in first- and second-price auctions. Note that the analysis in this chapter focuses on linear bidding strategies $b_i(v_i) = \beta_i v_i$. Near the NE - which is given in 1PAs and 2PAs by linear strategies - this might be a realistic assumption. In principle we make it easier for bidders to find the NE by confining their search from the infinite-dimensional function space of all possible bidding functions $b(v)$ to the one-dimensional space of linear bidding strategies $b(v) = \beta v$ that entails the NE. However, it is a simplification. In the next chapter I will investigate the more general case of nonlinear bidding functions. The current chapter is organized as follows:

- **4.2** focuses on myopic best response play in first price open bid auctions¹ (1POBAs)². I quantify the average strategy $\bar{\beta}_0^{\text{mBR}}(\beta_1)$ of player 0 under myopic best response play against a fixed strategy β_1 of bidder 1. As a finger exercise I additionally quantify how a fixed seller reserve price influences the average bid of a bidder who plays a myopic best response to a fixed strategy of a second bidder.
- **4.3** quantifies the strategy distributions of bidder 0 against a fixed strategy. Subsequently I use the result to show that two mutually adapting myopic best response bidders converge to bidding 0.
- **4.4** identifies the perfect memory best response function $\beta_0^{\text{pBR}}(\beta_1)$ against a fixed strategy β_1 . This is the best response against a known strategy if neither the own nor the opponent's value are known upfront.
- **4.5** analyzes the dynamics of mutually adapting bidders under perfect memory best response. It shows that bidders converge to the NE but that it takes them infinitely long to arrive there.

¹By open bid auction I mean that each bidder makes her bid simultaneously and all bidders can observe all the other participants' bids.

²All abbreviations are found in the appendix

- **4.6** 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.
- **4.7** concludes that the RET emerges after infinite time from the best response dynamics if bidders have infinite memory strength or instantaneously if bidders are infinitely sophisticated. Under realistic assumptions however the RET breaks down.

4.2 Myopic Best Response to Fixed Strategies in 1POBAs

4.2.1 Payoff Function

Unless explicitly noted, I focus in this thesis on the analysis of linear bidding strategies β_i . If bidder i has a private value v_i , her bid is given as $b_i(v_i) = \beta_i v_i$.

Bidder 0 wins an auction, if her bid is higher than that of the opponent. Otherwise, she loses the auction. Therefore, the probability for player 0 to win a 1POBA is given by

$$p_0^{win} = \begin{cases} 1 & \text{if } \beta_0 v_0 > \beta_1 v_1 \\ 0 & \text{if } \beta_0 v_0 < \beta_1 v_1 \end{cases} \quad (4.1)$$

In closed form, the winning probability can be rewritten as

$$p_0^{win} = \Theta(\beta_0 v_0 - \beta_1 v_1) \quad (4.2)$$

where the Θ - or unitstep-function $\Theta(x)$ assigns a value of 1 to every positive number x and 0 to every negative number. For the value of $\Theta(0)$, there exist different conventions such as $\Theta(0) = 0$, $\Theta(0) = 1$ or $\Theta(0) = 0.5$. For defining the winning probability in auctions with continuous value distributions it doesn't matter which particular convention I apply, since the equality of two bids occurs with probability 0. For details on the Θ function please check Appendix A.

In a first-price auction, a winning bidder has a payoff of $v_i - \beta_i v_i$ whereas her payoff is 0 if she loses. Therefore, player 0's payoff in a 1PA is given as

$$PO_0(\beta_0, \beta_1, v_0, v_1) = v_0(1 - \beta_0)\Theta(\beta_0 v_0 - \beta_1 v_1) \quad (4.3)$$

Figure 4.1 shows the payoff of player 0 for different choices of β_0 when β_1 , v_0 and v_1 are fixed.

A bidder maximizes her payoff in a 1PA by bidding just slightly more than her competitor. If the opponent bids more than her value then all of her feasible strategies³ would have generated 0 payoff and therefore it doesn't matter, which strategy she chooses. Therefore, the best response function of player 0 if she knows β_1 , v_0 and v_1 is given by

$$\beta_0^{br}(\beta_1, v_0, v_1) = \begin{cases} \lim_{\epsilon \rightarrow 0} \frac{\beta_1 v_1}{v_0} + \epsilon = \frac{\beta_1 v_1}{v_0} & \text{if } \frac{\beta_1 v_1}{v_0} < 1 \\ \text{arbitrary } \in (0, 1) & \text{if } \frac{\beta_1 v_1}{v_0} > 1 \end{cases} \quad (4.4)$$

4.2.2 Myopic Best Response to β_1

The best response function determines the average strategy of player 0 if she plays in each round the strategy β_0 that would have maximized her last round's payoff. Here, I stick to the convention that, if at time t no feasible strategy of player 0 could have won the auction (i.e., if

³The set of feasible strategies in 1PAs is for each player the interval $[0, 1]$. All strategies > 1 yield negative expected payoff.




Figure 4.1: Payoff of Strategies

$v_0^{(t)} < \beta_1 v_1^{(t)}$) then she continues to use her previous rounds strategy $\beta_0^{(t)}$ also at time $t + 1$. Therefore player 0's next rounds strategy is given as

$$\beta_0^{(t+1)}(\beta_0^{(t)}, \beta_1, v_0^{(t)}, v_1^{(t)}) = \frac{\beta_1 v_1^{(t)}}{v_0^{(t)}} \Theta(v_0^{(t)} - \beta_1 v_1^{(t)}) + \beta_0^{(t)} \Theta(\beta_1 v_1^{(t)} - v_0^{(t)}) \quad (4.5)$$

Note that this is just a short way for writing

$$\beta_0^{(t+1)} = \begin{cases} \frac{\beta_1 v_1^{(t)}}{v_0^{(t)}} & \text{if } v_0^{(t)} > \beta_1 v_1^{(t)} \\ \beta_0^{(t)} & \text{if } v_0^{(t)} < \beta_1 v_1^{(t)} \end{cases} \quad (4.6)$$

The expected value of $\beta_0^{(t+1)}$, $\bar{\beta}_0^{(t+1)}(\beta_1, v_1^{(t)})$ can be found by integrating over $v_0^{(t)}$:

$$\begin{aligned} \bar{\beta}_0^{(t+1)}(\beta_1, v_1^{(t)}) &= \int_0^1 dv_0^{(t)} \frac{\beta_1 v_1^{(t)}}{v_0^{(t)}} \Theta(v_0^{(t)} - \beta_1 v_1^{(t)}) + \beta_0^{(t)} \int_0^1 dv_0^{(t)} \Theta(\beta_1 v_1^{(t)} - v_0^{(t)}) = \\ &= \int_{\beta_1 v_1^{(t)}}^1 dv_0^{(t)} \frac{\beta_1 v_1^{(t)}}{v_0^{(t)}} + \beta_0^{(t)} \int_0^{\beta_1 v_1^{(t)}} dv_0^{(t)} = \beta_1 v_1^{(t)} (\beta_0^{(t)} - \ln(\beta_1 v_1^{(t)})) \end{aligned} \quad (4.7)$$

To understand how I get from the second to the third expression, think through what happens for different value realizations of $v_0^{(t)}$. The first integral gets only a positive contribution from the integrand, if $v_0^{(t)} > \beta_1 v_1^{(t)}$. Additionally, the values of $v_0^{(t)}$ are restricted between 0 and 1 because of the integration borders. Since $0 < \beta_1 < 1$ and $0 < v_1^{(t)} < 1$, also $0 < \beta_1 v_1^{(t)} < 1$. In particular, $\beta_1 v_1^{(t)} > 0$. Therefore, the values of $v_0^{(t)}$ for which the integrand gives a positive contribution are the ones between $\beta_1 v_1^{(t)}$ and 1. Similarly, in the second integral, the conditions $v_0^{(t)} < \beta_1 v_1^{(t)}$ and $0 < v_0^{(t)} < 1$ lead to the integration between 0 and $\beta_1 v_1^{(t)}$ because $\beta_1 v_1^{(t)} < 1$. Subsequent integration of (4.7) over $v_1^{(t)}$ yields the expected value of bidder 0's strategy in dependence of bidder 1's strategy:

$$\bar{\beta}_0^{(t+1)}(\beta_1) = \int_0^1 dv_1^{(t)} \bar{\beta}_0^{(t+1)}(\beta_1, v_1^{(t)}) = \frac{\beta_1}{4} (1 + 2\beta_0^{(t)} - 2\ln\beta_1) \quad (4.8)$$

Now, since in the long-term on average the relation $\bar{\beta}_0^{(t+1)} = \bar{\beta}_0^{(t)} \stackrel{!}{=} \bar{\beta}_0^{\text{mBR}}$ must hold, I obtain for the average strategy of bidder 0, $\bar{\beta}_0^{\text{mBR}}(\beta_1)$, the condition

$$\bar{\beta}_0^{\text{mBR}}(\beta_1) = \frac{\beta_1}{4}(1 + 2\bar{\beta}_0^{\text{mBR}}(\beta_1) - 2\ln \beta_1) \quad (4.9)$$

By solving for $\bar{\beta}_0^{\text{mBR}}$, I arrive at the following

Theorem 4.1 *Consider a two bidder first price open bid auction with linear bidding strategies and $v_i \sim U(0,1)$. Assume that bidder 1 plays a fixed strategy β_1 and bidder 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 $\beta_0 \in [0,1]$, she continues to use the previous round's strategy. Then, in the long term her average strategy is given as*

$$\bar{\beta}_0^{\text{mBR}}(\beta_1) = \frac{\beta_1(1 - 2\ln \beta_1)}{2(2 - \beta_1)} \quad (4.10)$$

In particular, $\bar{\beta}_0^{\text{mBR}}(\beta_1) < \frac{1}{2} \forall \beta_1 \in [0,1)$.

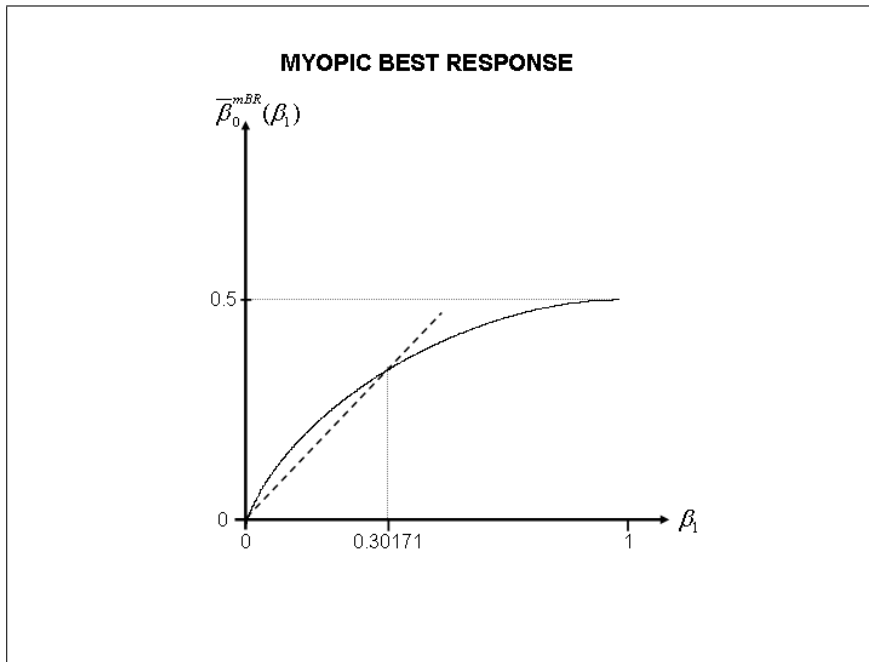


Figure 4.2: Myopic Best Response

Figure 4.2 plots $\bar{\beta}_0^{\text{mBR}}(\beta_1)$. We see that the average myopic best response against any feasible strategy $\beta_1 \in (0,1)$ is below the NE of $\frac{1}{2}$. As $\beta_1 \rightarrow 1$, $\bar{\beta}_0^{\text{mBR}}(\beta_1)$ approaches the NE from below.

The dashed line shows the first median $\beta_0 = \beta_1$. Since there are two points where the first median crosses the best response function, (4.10) has two points with $\bar{\beta}_0^{\text{mBR}}(\beta_1) = \beta_1$. The first point is $\beta_1 = 0$ and the second point is $\beta_1 \sim 0.30171$. For low strategies of bidder 1, $0 < \beta_1 < 0.30171$, bidder 0 bids under myopic best response on average more than β_1 , since the best response function is above the first median. For $\beta_1 > 0.30171$ she bids on average less than β_1 .

How can we understand this intuitively? The lower β_1 , the lower also $\beta_1 v_1^{(t)}$ and therefore the smaller the chance that $v_0^{(t)} < \beta_1 v_1^{(t)}$. Therefore, the rule $\beta_0^{(t+1)} = \beta_0^{(t)}$ is only seldom applied and the value of next rounds strategy is basically given as $\beta_0^{(t+1)} = \frac{\beta_1 v_1^{(t)}}{v_0^{(t)}}$. This is on average bigger than β_1^4 .

⁴If the rule $\beta_0^{(t+1)} = \beta_0^{(t)}$ would never be applied, the expected value of $\beta_0^{(t)}$ would be given as $\int_0^1 dv_0^{(t)} \int_0^1 dv_1^{(t)} \frac{\beta_1 v_1^{(t)}}{v_0^{(t)}} = \text{"}\infty\text{"} > \beta_1$.

As β_1 increases, the rule $\beta_0^{(t+1)} = \beta_0^{(t)}$ gets applied increasingly often. This prevents some strategy increases of player 0 but puts no restriction on decreases of her strategy. Therefore this rule lowers the average strategy of bidder 0.

If (4.10) was of the form $\beta_0 = f(\beta_1)$, such a constellation where the best response to a low strategy is higher than the opponents' strategy and but lower for a high strategy of the opponent would mean that the strategy combination $(\beta_0 = 0.30171, \beta_1 = 0.30171)$ is an attracting fixed point where mutually adapting bidders converge to. However, (4.10) just quantifies the *average* strategy of bidder 0 and not the strategy itself. For a fixed strategy of the opponent, the strategy of the respondent permanently fluctuates around the mean that is determined by (4.10). Therefore, to assess the long term behavior of bidders under mutual adaptation, one needs to determine the distribution of these fluctuations and then quantify the interplay between the strategy distributions of two players. In section 4.3 I will follow this route to assess the long term behavior of mutually adapting bidders under myopic best response.

4.2.3 Myopic Best Response with Seller Reserve Price

This subsection is not necessary to understand the remainder of the thesis. It can be left out upon first reading. I extend the above auction setup by allowing the seller to participate in the auction:

Still, bidder 1 plays the fixed strategy β_1 whereas bidder 0 plays in each round the strategy that would have maximized her last round's payoff. In addition, the seller posts a fixed reserve price $\rho \in (0, 1)$ such that the item is only sold to bidder 0 if her bid is – in addition to being bigger than bidder 1s bid – also bigger than ρ . Player 0s next round strategy is therefore given as

$$\beta_0^{(t+1)}(\beta_0^{(t)}, \beta_1, v_0^{(t)}, v_1^{(t)}, \rho) = \begin{cases} \frac{\max(\beta_1 v_1^{(t)}, \rho)}{v_0^{(t)}} & \text{if } v_0^{(t)} > \max(\beta_1 v_1^{(t)}, \rho) \\ \beta_0^{(t)} & \text{if } v_0^{(t)} < \max(\beta_1 v_1^{(t)}, \rho) \end{cases} \quad (4.11)$$

where $\max(\beta_1 v_1^{(t)}, \rho) = \rho + \Theta(\beta_1 v_1^{(t)} - \rho)(\beta_1 v_1^{(t)} - \rho) < 1$. (see Appendix A)
Integration over $v_0^{(t)}$ yields the expected strategy of bidder 0, still dependent on $v_1^{(t)}$,

$$\begin{aligned} \bar{\beta}_0^{(t+1)}(\beta_0^{(t)}, \beta_1, v_1^{(t)}, \rho) &= \\ & \int_{\rho + \Theta(\beta_1 v_1^{(t)} - \rho)(\beta_1 v_1^{(t)} - \rho)}^1 dv_0^{(t)} (\rho + \Theta(\beta_1 v_1^{(t)} - \rho)(\beta_1 v_1^{(t)} - \rho)) \frac{1}{v_0^{(t)}} + \\ & + \beta_0^{(t)} \int_0^{\rho + \Theta(\beta_1 v_1^{(t)} - \rho)(\beta_1 v_1^{(t)} - \rho)} dv_0^{(t)} = \\ & = \dots = \rho(\beta_0^{(t)} - \ln \rho) + \Theta(\beta_1 v_1^{(t)} - \rho)(\beta_0^{(t)} - \rho - \ln \beta_1 v_1^{(t)})(\beta_1 v_1^{(t)} - \rho) \end{aligned} \quad (4.12)$$

Regrouping and integration over $v_1^{(t)}$ yields the expected strategy

$$\begin{aligned} \bar{\beta}_0^{(t+1)}(\beta_0^{(t)}, \beta_1, \rho) &= \underbrace{\int_0^1 dv_1^{(t)} \Theta\left(\frac{\rho}{\beta_1} - v_1^{(t)}\right) (\beta_0^{(t)} \rho - \rho \ln \rho)}_{(I)} + \\ & \underbrace{\int_0^1 dv_1^{(t)} \Theta\left(v_1^{(t)} - \frac{\rho}{\beta_1}\right) (\beta_0^{(t)} \beta_1 v_1^{(t)} - \beta_1 v_1^{(t)} \ln \beta_1 v_1^{(t)})}_{(II)} \end{aligned} \quad (4.13)$$

In expression (I), the integrand is only non-zero if $v_1^{(t)} < \min(1, \frac{\rho}{\beta_1}) = 1 - \Theta(1 - \frac{\rho}{\beta_1})(1 - \frac{\rho}{\beta_1}) < 1^5$. Therefore, the integral is given as

$$(I) = \int_0^{1 - \Theta(1 - \frac{\rho}{\beta_1})(1 - \frac{\rho}{\beta_1})} dv_1^{(t)} (\beta_0^{(t)} \rho - \rho \ln \rho) = \beta_0^{(t)} \rho - \rho \ln \rho - \Theta(\beta_1 - \rho)(1 - \frac{\rho}{\beta_1})(\beta_0^{(t)} \rho - \rho \ln \rho) \quad (4.14)$$

In expression (II), the integrand is only non-zero if $v_1^{(t)} > \frac{\rho}{\beta_1}$ and if $v_1^{(t)} < 1$. Therefore, the integral is given as⁶

$$(II) = \Theta(1 - \frac{\rho}{\beta_1}) \int_{\frac{\rho}{\beta_1}}^1 dv_1^{(t)} (\beta_0^{(t)} \beta_1 v_1^{(t)} - \beta_1 v_1^{(t)} \ln \beta_1 v_1^{(t)}) = \\ = \dots = \Theta(\beta_1 - \rho) \frac{\beta_1^2 + 2\beta_0^{(t)} \beta_1^2 - \rho^2 - 2\beta_0^{(t)} \rho^2 - 2\beta_1^2 \ln \beta_1 + 2\rho^2 \ln \rho}{4\beta_1} \quad (4.15)$$

The prefactor $\Theta(1 - \frac{\rho}{\beta_1})$ guarantees that for $\frac{\rho}{\beta_1} > 1$, the integral doesn't give a negative contribution.

In summary, I obtain the expected strategy of bidder 0 as

$$\bar{\beta}_0^{(t+1)}(\beta_0^{(t)}, \beta_1, \rho) = (I) + (II) = \rho(\beta_0^{(t)} - \ln \rho) - \Theta(\beta_1 - \rho)(1 - \frac{\rho}{\beta_1})(\beta_0^{(t)} \rho - \rho \ln \rho) + \\ + \Theta(\beta_1 - \rho) \frac{\beta_1^2 + 2\beta_0^{(t)} \beta_1^2 - \rho^2 - 2\beta_0^{(t)} \rho^2 - 2\beta_1^2 \ln \beta_1 + 2\rho^2 \ln \rho}{4\beta_1} \quad (4.16)$$

To determine the average behavior, I set - like in the previous section - $\bar{\beta}_0^{(t+1)} = \bar{\beta}_0^{(t)} \stackrel{!}{=} \bar{\beta}_0^{\text{mBR}}$ and solve for $\bar{\beta}_0^{\text{mBR}}$. For $\beta_1 < \rho$ this yields the solution

$$\bar{\beta}_0^{\text{mBR}}(\beta_1, \rho) = \frac{\rho}{\rho - 1} \ln \rho \quad (4.17)$$

For $\beta_1 > \rho$ the solution is

$$\bar{\beta}_0^{\text{mBR}}(\beta_1, \rho) = \frac{\rho^2 - \beta_1^2 + 2\beta_1^2 \ln \beta_1 + 2\rho^2 \ln \rho}{2(\rho^2 + \beta_1^2 - 2\beta_1)} \quad (4.18)$$

In summary I arrive at

Theorem 4.2 *Consider a first price open bid auction with linear bidding strategies and $v_i \sim U(0, 1)$. Two bidders and one seller participate in the auction. Assume that bidder 1 plays in each round a fixed strategy β_1 and the seller has a fixed reserve price ρ . Bidder 0 follows a myopic best response strategy updating rule. Then, in the long term, her average strategy is given as*

$$\bar{\beta}_0^{\text{mBR}}(\beta_1, \rho) = \Theta(\rho - \beta_1) \frac{\rho \ln \rho}{\rho - 1} + \Theta(\beta_1 - \rho) \frac{\rho^2 - \beta_1^2 + 2\beta_1^2 \ln \beta_1 + 2\rho^2 \ln \rho}{2(\beta_1^2 + \rho^2 - 2\beta_1)} + \quad (4.19)$$

Figure 4.3 plots the average strategy of bidder 0 in dependence of bidder 1s strategy for different reservation prices of the seller as given by equation (4.19). What is the intuition behind the behavior of bidder 0?

(4.19) consists of two parts. The first part describes the average strategy of bidder 0 if the

⁵ $v_1^{(t)}$ has to be < 1 because of the upper integration border. On the other hand, the Θ -function is nonzero only for $v_1^{(t)} < \frac{\rho}{\beta_1}$.

⁶The term $\Theta(1 - \frac{\rho}{\beta_1})$ in front of the integral ensures that in the case of $\frac{\rho}{\beta_1} > 1$ the integral doesn't give a negative contribution.




Figure 4.3: Myopic Best Response with Reserve Price

sellers reservation price is bigger than bidder 1s strategy. Then, the strategy depends only on the sellers reservation price. This is intuitively clear. Bidder 0 has to bid above both, the seller reservation price and bidder 1s bid to make positive payoff. If $\beta_1 < \rho$ than because of $b_1 = \beta_1 v_1^{(t)} < \beta_1 < \rho$, the seller reservation price is always above bidder 1s bid. Therefore bidder 0s strategy must be completely independent of bidder 1s strategy. In Figure 4.2, the parameter combinations where $\rho > \beta_1$ are represented by the horizontal line at the left hand side of graphs with $\rho > 0$.

If, however, bidder 1s strategy is higher than the seller reservation price, then, at least for some value realizations v_1 , bidder 1s bid is above the reserve price. Therefore, bidder 1s bids have an impact on bidder 0s best response. The higher bidder 1s strategy is, the more will bidder 0 bid on average. This is the second term of (4.19) which corresponds to the monotonic increase at the right side of the graphs.

Note that for $\rho = 0$, equation (4.19) reduces to the equation for myopic best response without reserve price, (4.10), just as it has to be.

4.3 Mutual Myopic Best Response Adaptation

In this section I want to investigate the dynamics of 1POBAs with two myopic best response bidders. In the first part I derive the strategy distribution of 1 bidder who plays a myopic best response against a fixed strategy of her opponent. Subsequently I use this result to derive the strategy distributions of two mutually adapting bidders. Finally I derive the outcome of 1POBAs with two mutually adapting bidders under myopic best response dynamics. The surprising result is that this learning rule leads both bidders to bid 0.

4.3.1 Strategy Distribution under Myopic Best Response

I assume that the strategy fluctuations follow a stationary distribution if a bidder uses myopic best response strategy updating against a fixed strategy of her opponent. In the end of this derivation we will see that the distribution indeed exists. Without loss of generality I assume that bidder 1 always bids according to the strategy β_1 and bidder 0 is the adapting bidder.

The updating rule of bidder 0 is

$$\beta_0^{(t+1)} = \begin{cases} \frac{\beta_1 v_1^{(t)}}{v_0^{(t)}} & \text{if } \frac{\beta_1 v_1^{(t)}}{v_0^{(t)}} < 1 \\ \beta_0^{(t)} & \text{if } \frac{\beta_1 v_1^{(t)}}{v_0^{(t)}} > 1 \end{cases} \quad (4.20)$$

For the sake of notational compactness, I drop from now on the time indices and write v_0 for $v_0^{(t)}$ and v_1 for $v_1^{(t)}$. The cumulated distribution function (CDF) of bidder 0's new strategy, i.e., the probability that bidder 0's strategy is between 0 and $\tilde{\beta}$ after the updating, is given as

$$\begin{aligned} F(\tilde{\beta}) &= \\ &= P(0 < \beta_0^{\text{new}} < \tilde{\beta}) = P\left(\frac{\beta_1 v_1}{v_0} < 1\right) P\left(\frac{\beta_1 v_1}{v_0} < \tilde{\beta} \mid \frac{\beta_1 v_1}{v_0} < 1\right) + P\left(\frac{\beta_1 v_1}{v_0} > 1\right) P(0 < \beta_0^{\text{old}} < \tilde{\beta}) = \\ &= P\left(\frac{\beta_1 v_1}{v_0} < 1\right) \frac{P\left(\frac{\beta_1 v_1}{v_0} < \tilde{\beta}\right)}{P\left(\frac{\beta_1 v_1}{v_0} < 1\right)} + (1 - P\left(\frac{\beta_1 v_1}{v_0} < 1\right)) F(\tilde{\beta}) = \\ &= P\left(v_1 < \frac{v_0 \tilde{\beta}}{\beta_1}\right) + F(\tilde{\beta}) - P\left(v_1 < \frac{v_0}{\beta_1}\right) F(\tilde{\beta}). \end{aligned} \quad (4.21)$$

By canceling the term $F(\tilde{\beta})$ from both sides I obtain the equation for the CDF

$$P\left(v_1 < \frac{v_0}{\beta_1}\right) F(\tilde{\beta}) = P\left(v_1 < \frac{v_0 \tilde{\beta}}{\beta_1}\right). \quad (4.22)$$

For uniform value distributions, the pertinent probabilities can be calculated as

$$\begin{aligned} P\left(v_1 < \frac{v_0}{\beta_1}\right) &= \min\left(1, \frac{v_0}{\beta_1}\right) = 1 - \Theta\left(1 - \frac{v_0}{\beta_1}\right) \left(1 - \frac{v_0}{\beta_1}\right) = 1 - \Theta(\beta_1 - v_0) \left(1 - \frac{v_0}{\beta_1}\right) \\ P\left(v_1 < \frac{v_0 \tilde{\beta}}{\beta_1}\right) &= \min\left(1, \frac{v_0 \tilde{\beta}}{\beta_1}\right) = 1 - \Theta\left(1 - \frac{v_0 \tilde{\beta}}{\beta_1}\right) \left(1 - \frac{v_0 \tilde{\beta}}{\beta_1}\right) = 1 - \Theta\left(\frac{\beta_1}{\tilde{\beta}} - v_0\right) \left(1 - \frac{v_0 \tilde{\beta}}{\beta_1}\right). \end{aligned} \quad (4.23)$$

Therefore I obtain by integration of (4.22) over v_0 the condition for the CDF

$$F(\tilde{\beta}) \int_0^1 dv_0 \left\{1 - \Theta(\beta_1 - v_0) \left(1 - \frac{v_0}{\beta_1}\right)\right\} = \int_0^1 dv_0 \left\{1 - \Theta\left(\frac{\beta_1}{\tilde{\beta}} - v_0\right) \left(1 - \frac{v_0 \tilde{\beta}}{\beta_1}\right)\right\} \quad (4.24)$$

This is equivalent to

$$F(\tilde{\beta}) \left\{ \int_0^{\beta_1} dv_0 \frac{v_0}{\beta_1} + \int_{\beta_1}^1 dv_0 \right\} = \int_0^{1 - \Theta\left(1 - \frac{\beta_1}{\tilde{\beta}}\right) \left(1 - \frac{\beta_1}{\tilde{\beta}}\right)} dv_0 \frac{v_0 \tilde{\beta}}{\beta_1} + \int_{1 - \Theta\left(1 - \frac{\beta_1}{\tilde{\beta}}\right) \left(1 - \frac{\beta_1}{\tilde{\beta}}\right)}^1 dv_0 \quad (4.25)$$

because $\min\left(1, \frac{\beta_1}{\tilde{\beta}}\right) = 1 - \Theta\left(1 - \frac{\beta_1}{\tilde{\beta}}\right) \left(1 - \frac{\beta_1}{\tilde{\beta}}\right)$. Integrating this out yields

$$F(\tilde{\beta}) \left\{ \frac{v_0^2}{2\beta_1} \Big|_0^{\beta_1} + (1 - \beta_1) \right\} = \frac{v_0^2 \tilde{\beta}}{2\beta_1} \Big|_0^{1 - \Theta(\tilde{\beta} - \beta_1) \left(1 - \frac{\beta_1}{\tilde{\beta}}\right)} + \Theta(\tilde{\beta} - \beta_1) \left(1 - \frac{\beta_1}{\tilde{\beta}}\right) \quad (4.26)$$

and therefore

$$\left(1 - \frac{\beta_1}{2}\right) F(\tilde{\beta}) = \frac{\tilde{\beta}}{2\beta_1} \left[1 - \Theta(\tilde{\beta} - \beta_1) \left(1 - \frac{\beta_1}{\tilde{\beta}}\right)\right]^2 + \Theta(\tilde{\beta} - \beta_1) \left(1 - \frac{\beta_1}{\tilde{\beta}}\right). \quad (4.27)$$

Solving this for $F(\tilde{\beta})$ yields

$$F(\tilde{\beta}) = \frac{2 \left\{ \frac{\tilde{\beta}}{2\beta_1} \left[1 - \Theta(\tilde{\beta} - \beta_1) \left(1 - \frac{\beta_1}{\tilde{\beta}}\right)\right]^2 + \Theta(\tilde{\beta} - \beta_1) \left(1 - \frac{\beta_1}{\tilde{\beta}}\right) \right\}}{2 - \beta_1} \quad (4.28)$$

so

$$F(\tilde{\beta}) = \begin{cases} \frac{2\tilde{\beta}-\beta_1}{\tilde{\beta}(2-\beta_1)} & \text{if } \beta_1 < \tilde{\beta} \\ \frac{\tilde{\beta}}{\beta_1(2-\beta_1)} & \text{if } \beta_1 > \tilde{\beta} \end{cases} \quad (4.29)$$

The PDF $f(\tilde{\beta})$ is the derivative of the CDF $F(\tilde{\beta})$ with respect to $\tilde{\beta}$. Therefore it is given by

$$f(\tilde{\beta}) = \begin{cases} \frac{\beta_1}{(2-\beta_1)\tilde{\beta}^2} & \text{if } \beta_1 < \tilde{\beta} \\ \frac{1}{\beta_1(2-\beta_1)} & \text{if } \beta_1 > \tilde{\beta} \end{cases} \quad (4.30)$$

So I arrive at the following

Theorem 4.3 Consider a two bidder first price open bid auction with linear bidding strategies and $v_i \sim U(0, 1)$. Assume that bidder 1 plays a fixed strategy β_1 and bidder 0 follows a myopic best response strategy updating rule. Then, in the long term, her strategy distribution is given as

$$f(\beta_0) = \frac{\beta_1}{(2-\beta_1)\beta_0^2} \Theta(\beta_0 - \beta_1) + \frac{1}{\beta_1(2-\beta_1)} \Theta(\beta_1 - \beta_0) \quad (4.31)$$

Figure 4.4 shows the strategy distribution of player 0 against a fixed $\beta_1 = 0.3$. The figure shows

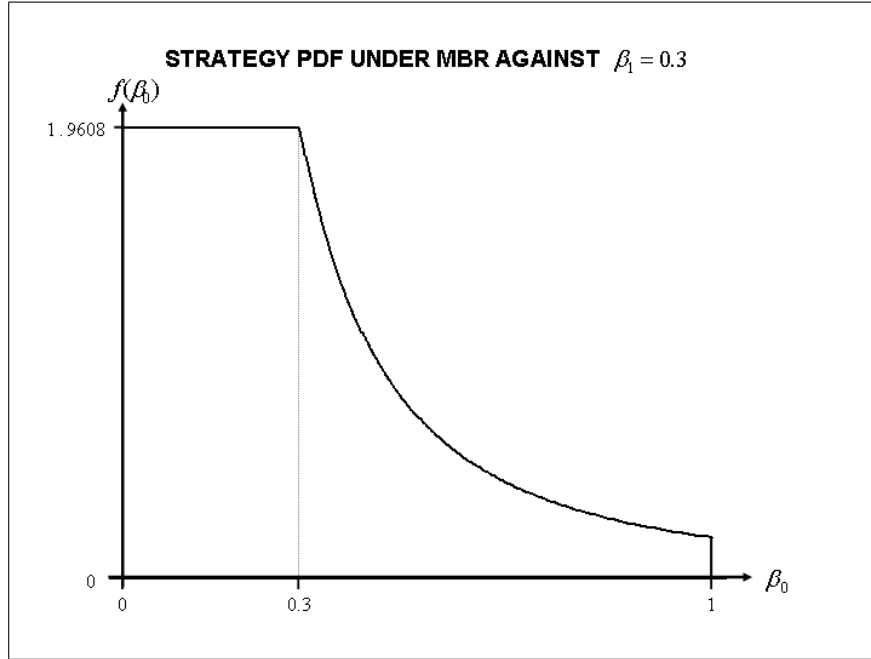


Figure 4.4: Strategy PDF under mBR against $\beta_1 = 0.3$

that, on average, bidder 0 plays all strategies that are lower than the strategy of her opponent with equal probability. Strategies above the opponent's strategy are played less often and the higher the strategy, the less frequently it is played.

As a short double-check of the correctness of this result I calculate the average bid:

$$\int_0^1 d\tilde{\beta} \tilde{\beta} f(\tilde{\beta}) = \int_0^{\beta_1} \frac{\tilde{\beta}}{\beta_1(2-\beta_1)} d\tilde{\beta} + \int_{\beta_1}^1 \frac{\beta_1}{(2-\beta_1)\tilde{\beta}} d\tilde{\beta} = \quad (4.32)$$

$$= \frac{1}{\beta_1(2-\beta_1)} \frac{\beta_1^2}{2} - \frac{\beta_1}{(2-\beta_1)} \ln \beta_1 = \frac{\beta_1(1-2\ln \beta_1)}{2(2-\beta_1)} \quad (4.33)$$

which is in accordance with Theorem 4.1.

4.3.2 Convergence under Myopic Mutual Adaptation

In the next step I want to derive the strategy distribution of two mutually adapting bidders who follow a myopic best response strategy updating rule. Again I assume initially that the strategy distribution exists. However, we will see in the course of the derivation that this is only true if bidders have a lowest strategy ϵ in their strategy set that is positive. The degenerate distribution in the case of $\beta_i \in [0, 1]$ is then found by taking $\lim \epsilon \rightarrow 0$.

We know from the last theorem the strategy distribution of bidder 0 if bidder 1 plays a fixed strategy β_1 . If bidder 1 plays instead a mixed strategy with PDF $f_1(\beta_1)$ then the probability that bidder 0 plays a certain strategy $\tilde{\beta}_0$ is given by the integral over all probabilities that player 1 plays a certain strategy β_1 , $f_1(\beta_1)$, times the probability that bidder 0 responds to β_1 with the strategy $\tilde{\beta}_0$ which is known from Theorem 4.3. Therefore the strategy distribution of bidder 0 against a strategy distribution $f_1(\beta_1)$ is given by

$$f_0(\tilde{\beta}_0) = \int_0^1 d\beta_1 f_1(\beta_1) (f_0(\tilde{\beta}_0) | \beta_1) = \quad (4.34)$$

$$= \int_0^{\tilde{\beta}_0} d\beta_1 f_1(\beta_1) \frac{\beta_1}{(2 - \beta_1)\tilde{\beta}_0^2} + \int_{\tilde{\beta}_0}^1 d\beta_1 f_1(\beta_1) \frac{1}{\beta_1(2 - \beta_1)}. \quad (4.35)$$

For mutually adapting bidders the situation is symmetric. I assume that there exists a common, differentiable strategy distribution. Denoting this common strategy distribution by $f_0 = f_1 \stackrel{!}{=} f$ and relabeling the variables, f must satisfy

$$f(\tilde{\beta}) = \int_0^{\tilde{\beta}} d\beta f(\beta) \frac{\beta}{(2 - \beta)\tilde{\beta}^2} + \int_{\tilde{\beta}}^1 d\beta f(\beta) \frac{1}{\beta(2 - \beta)}. \quad (4.36)$$

To simplify the notation, I denote by $I_1(\beta)$ the primitive (=antiderivative) of $f(\beta) \frac{\beta}{(2 - \beta)}$ and by $I_2(\beta)$ the primitive of $f(\beta) \frac{1}{\beta(2 - \beta)}$. Then the condition for the symmetric distribution can be written as

$$f(\tilde{\beta}) = \frac{1}{\tilde{\beta}^2} [I_1(\tilde{\beta}) - I_1(0)] + [I_2(1) - I_2(\tilde{\beta})]. \quad (4.37)$$

Differentiation yields

$$\begin{aligned} f'(\tilde{\beta}) &= -\frac{2}{\tilde{\beta}^3} [I_1(\tilde{\beta}) - I_1(0)] + \frac{1}{\tilde{\beta}^2} \left\{ f(\tilde{\beta}) \frac{\tilde{\beta}}{2 - \tilde{\beta}} - 0 \right\} + f(1) - f(\tilde{\beta}) \frac{1}{\tilde{\beta}(2 - \tilde{\beta})} = \\ &= -\frac{2}{\tilde{\beta}^3} [I_1(\tilde{\beta}) - I_1(0)] + f(1). \end{aligned} \quad (4.38)$$

This equation can be transformed so that the primitives stand on one side:

$$-\frac{\tilde{\beta}^3}{2} f'(\tilde{\beta}) + \frac{\tilde{\beta}^3}{2} f(1) = I_1(\tilde{\beta}) - I_1(0). \quad (4.39)$$

Differentiating again and using the fact that the derivative of the primitive I_1 at $\tilde{\beta}$ gives the original function $f(\tilde{\beta}) \frac{\tilde{\beta}}{2 - \tilde{\beta}}$ yields the second order differential equation

$$-\frac{3\tilde{\beta}^2}{2} f'(\tilde{\beta}) - \frac{\tilde{\beta}^3}{2} f''(\tilde{\beta}) + \frac{3\tilde{\beta}^2}{2} f(1) = f(\tilde{\beta}) \frac{\tilde{\beta}}{2 - \tilde{\beta}} \quad (4.40)$$

The solution to this differential equation is given by Mathematica as

$$f(\tilde{\beta}) = \frac{-8C_1 + 4\tilde{\beta}C_1 - 2C_2 - 4\tilde{\beta}f(1) + 4\tilde{\beta}^2f(1) - (\tilde{\beta} - 2)(C_2 - 4f(1)) \ln(\tilde{\beta} - 2) + (\tilde{\beta} - 2)C_2 \ln \tilde{\beta}}{4\tilde{\beta}} \quad (4.41)$$

What remains is to determine the three constants C_1 , C_2 and $f(1)$.

First notice that $\tilde{\beta} - 2 < 0 \forall \tilde{\beta} \in (0, 1)$ and therefore $\ln(\tilde{\beta} - 2)$ in the next to last term is a complex number. Since $f(\tilde{\beta})$ is a (real valued) probability density, the prefactor of $\ln(\tilde{\beta} - 2)$ has to be equal to 0. This yields the condition $C_2 = 4f(1)$. Therefore the expression for $f(\tilde{\beta})$ reduces to

$$f(\tilde{\beta}) = \frac{-8C_1 + 4\tilde{\beta}C_1 - 8f(1) - 4\tilde{\beta}f(1) + 4\tilde{\beta}^2f(1) + 4f(1)(\tilde{\beta} - 2) \ln \tilde{\beta}}{4\tilde{\beta}} \quad (4.42)$$

Since the equation has to be satisfied in particular for $\tilde{\beta} = 1$, I obtain the condition

$$f(1) = \frac{4C_1 - 8C_1 - 8f(1)}{4} = -C_1 - 2f(1) \quad (4.43)$$

and therefore

$$f(1) = -\frac{C_1}{3} \quad (4.44)$$

Substituting this back yields

$$f(\tilde{\beta}) = \dots = \frac{4\tilde{\beta}C_1 - 4C_1 - \tilde{\beta}^2C_1 - C_1(\tilde{\beta} - 2) \ln \tilde{\beta}}{3\tilde{\beta}} \quad (4.45)$$

To obtain C_1 I use the scaling condition that the integral over a probability density must yield 1. Therefore

$$1 \stackrel{!}{=} \int_0^1 d\tilde{\beta} f(\tilde{\beta}) = C_1 \int_0^1 d\tilde{\beta} \frac{4\tilde{\beta} - 4 - \tilde{\beta}^2 - (\tilde{\beta} - 2) \ln \tilde{\beta}}{3\tilde{\beta}} \quad (4.46)$$

This integral doesn't converge. In effect this means that there doesn't exist a differentiable strategy distribution for bidders under myopic best response. Therefore I construct now a series of differentiable distribution functions that converges towards this distribution: Assume that the lowest strategy in the players' strategy set is $\epsilon > 0$. Then, the condition

$$\int_{\epsilon}^1 d\tilde{\beta} f(\tilde{\beta}) = C_1 \int_{\epsilon}^1 d\tilde{\beta} \frac{4\tilde{\beta} - \tilde{\beta}^2 - (\tilde{\beta} - 2) \ln \tilde{\beta}}{3\tilde{\beta}} = C_1 \frac{1}{6} (9 - 10\epsilon + \epsilon^2 + 8 \ln \epsilon + 2\epsilon \ln \epsilon - 2 \ln^2 \epsilon) \quad (4.47)$$

yields the constant C_1 as

$$C_1 = \frac{6}{9 - 10\epsilon + \epsilon^2 + 8 \ln \epsilon + 2\epsilon \ln \epsilon - 2 \ln^2 \epsilon} . \quad (4.48)$$

In conclusion, mutually adapting bidders with a strategy set $\beta_i \in [\epsilon, 1]$ bid according to a continuous, differentiable strategy distribution that is defined by equations (4.45) and (4.48). This distribution puts most of the mass to the lowest strategy and dies out exponentially for higher strategies. As $\epsilon \rightarrow 0$, i.e., if bidders can use all strategies in the interval $[0, 1]$, bidders end up using the degenerate distribution $f(\tilde{\beta}) = \delta(\tilde{\beta})$ that puts all the weight on the strategy 0⁷. In other words, the only strategy that bidders use is $\beta_0 = \beta_1 = 0$.

The average strategy is given by

$$\begin{aligned} \bar{\beta} &= \int_{\epsilon}^1 d\tilde{\beta} \tilde{\beta} f(\tilde{\beta}) = \\ &= \frac{6}{9 - 10\epsilon + \epsilon^2 + 8 \ln \epsilon + 2\epsilon \ln \epsilon - 2 \ln^2 \epsilon} \int_{\epsilon}^1 d\tilde{\beta} \tilde{\beta} \frac{4\tilde{\beta} - 4 - \tilde{\beta}^2 - (\tilde{\beta} - 2) \ln \tilde{\beta}}{3} = \\ &= \frac{72\epsilon - 27\epsilon^2 + 4\epsilon^3 - 24\epsilon \ln \epsilon + 6\epsilon^2 \ln \epsilon - 49}{6(9 - 10\epsilon + \epsilon^2 + 8 \ln \epsilon + 2\epsilon \ln \epsilon - 2 \ln^2 \epsilon)} \end{aligned} \quad (4.49)$$

Letting $\epsilon \rightarrow 0$, the denominator grows because of the term $\ln^2 \epsilon$ faster towards ∞ than the numerator, so the average strategy converges to 0. Formally I obtain this result by applying de'l Hospital's rule.

Summarizing these results gives

⁷For the definition of the δ -function see Appendix A

Theorem 4.4 Consider a two bidder first price open bid auction with linear bidding strategies and $v_i \sim U(0, 1)$. Assume that both bidders mutually adapt by following a myopic best response updating rule. Then, in the long term, strategies and bids converge to 0.

At first sight, this result seems to be in contradiction to Theorem 4.1. There, equation 4.10 predicted that if one bidder plays a strategy between 0 and 0.30171, the opponent answers on average with a strategy that is *higher* than her strategy. So, how can it be that strategies still converge to 0?

The answer can be found by looking at the CDF of the strategy distribution that is given by equation (4.29). The probability that the updating bidder's strategy is below the opponent's strategy β is given by $F(\beta) = \frac{\beta}{\beta(2-\beta)} = \frac{1}{2-\beta} > \frac{1}{2} \quad \forall \beta \in [0, 1]$. So, the updating bidder's strategy will more often be below the opponent's strategy than above it. In other words, the *median* of the updating bidder's strategy is below the opponent's strategy. But the strategies above the opponent's strategy are sufficiently high to raise the *mean* above the opponent's strategy.

So, the dynamics of mutually adapting myopic best response bidders is governed by two forces. On the one hand, strategies can in each time step get arbitrarily large – as long as they stay below one. This is since – considering for the sake of clarity bidder 0 who updates against the strategy of bidder 1 – bidder 0's strategy in the next round is given by

$$\beta_0^{(t+1)} = \frac{\beta_1^{(t)} v_1^{(t)}}{v_0^{(t)}} \quad (4.50)$$

and as long as bidder 1 plays a strategy $\beta_1^{(t)} > 0$, there can always be value combinations that make bidder 0's strategy arbitrarily large: Any strategy $\beta_0^{(t+1)} \in (0, 1)$ can be reached from any strategy $\beta_1^{(t)}$ and any value $v_1^{(t)}$, if bidder 0's value is $\beta_1^{(t)} v_1^{(t)} / \beta_0^{(t+1)}$. However, the lower bidder 1's strategy, the smaller is the chance that bidder 0's value is such that a given strategy magnitude is reached.

On the other hand, as noted above, the updating bidder's strategy is more often below the opponent's strategy than above it. This induces in general a decrease of strategies. The speed of decrease gets ever lower, the nearer strategies are to 0 since then $F(\beta) \rightarrow 0.5$. So, what we can expect for the dynamics is a general drift of strategies towards 0 that is interrupted by occasional breakouts towards higher strategies. The lower the strategies of the bidders are, the lower is the chance for such breakouts. However, the lower the strategies are, the slower is also the speed of their decrease. So, the question is which of the two forces dominates in the long run: Will the breakouts be sufficiently strong and sufficiently often that the strategies never converge to 0? Or will strategies converge sufficiently fast to 0 such that they don't break away anymore?

What Theorem 4.4 tells us is that in the case of continuous strategies, the downward drift of the strategies is the dominating force. It is sufficient to compensate for the occasional break-outs of strategies such that the strategies die out over time. Below we will see in the simulations that the existence of a smallest strategy $\epsilon > 0$ changes the dynamics fundamentally. Then, strategies involve in breakouts to arbitrarily high strategies (compare Figure 10.3). These breakouts occur with a fixed frequency that depends on the lowest strategy and doesn't decrease over time.

4.4 Perfect Memory Best Response to Fixed Strategies in 1POBAs

In the previous section I assumed that players update in each round their strategy as the best response to the previous round. In this section, in contrast, I consider fictitious play (perfect memory best response): Bidders choose in each round the strategy that would have obtained the highest payoff in the whole game up to that particular round. The basic idea in the analysis of perfect memory best response is that - if the opponent plays a fixed strategy and the game has been going on infinitely long - choosing the strategy that would have maximized past payoff is the same as maximizing expected payoff against this fixed strategy: If bidder 1 plays a fixed

strategy, bidder 0 is able to estimate bidder 1s strategy accurately after infinitely many rounds. Therefore, the best response based on expected payoffs that I derive here is just the same as perfect memory best response.

4.4.1 Expected Payoff Function

The expected payoff $EPO_0(\beta_0, \beta_1, v_0)$ for $v_1 \sim U(0, 1)$ is obtained by integrating (4.3) over v_1 :

$$EPO_0(\beta_0, \beta_1, v_0) = \int_0^1 dv_1 PO_0(\beta_0, \beta_1, v_0, v_1) = \int_0^1 dv_1 \Theta\left(\frac{\beta_0 v_0}{\beta_1} - v_1\right) v_0 (1 - \beta_0). \quad (4.51)$$

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

$$v_1 < \min\left(1, \frac{\beta_0 v_0}{\beta_1}\right) = 1 - \Theta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right) \left(1 - \frac{\beta_0 v_0}{\beta_1}\right) \quad (4.52)$$

and therefore the integral transforms into

$$\begin{aligned} EPO_0(\beta_0, \beta_1, v_0) &= \\ &= \int_0^{1 - \Theta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right) \left(1 - \frac{\beta_0 v_0}{\beta_1}\right)} dv_1 v_0 (1 - \beta_0) = \left[1 - \Theta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right) \left(1 - \frac{\beta_0 v_0}{\beta_1}\right)\right] v_0 (1 - \beta_0). \end{aligned} \quad (4.53)$$

Figure 4.5 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 expected payoff by playing less than the NE strategy $\beta_0 = \frac{1}{2}$. For all $\beta_1 > 0.4$, player 0 maximizes her expected payoff by playing the NE strategy. In the following I will determine the optimal response

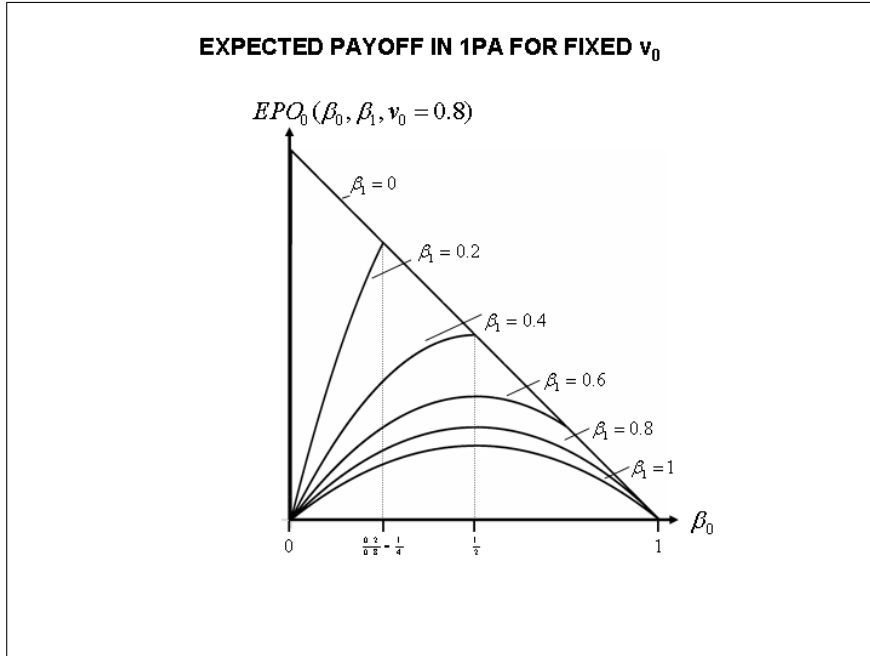


Figure 4.5: Expected Payoff in First-Price Auction for Fixed v_0

function $\beta_0^{\text{pBR}}(\beta_1)$, if both values v_0 and v_1 are drawn from a uniform random distribution. This tells us how player 0 would optimally play if she knows player 1s strategy but not next

round's realizations of v_0 and v_1 . For this I have to determine the expected payoff $EPO_0(\beta_0, \beta_1)$ if both values are unknown.

Using (4.53), the expected payoff is given by

$$\begin{aligned} EPO_0(\beta_0, \beta_1) &= \\ &= \int_0^1 dv_0 EPO(\beta_0, \beta_1, v_0) = \int_0^1 dv_0 [1 - \Theta(1 - \frac{\beta_0 v_0}{\beta_1})(1 - \frac{\beta_0 v_0}{\beta_1})] v_0 (1 - \beta_0) \\ &= \int_0^1 dv_0 v_0 (1 - \beta_0) - \int_0^1 dv_0 \Theta(\frac{\beta_1}{\beta_0} - v_0) (1 - \frac{\beta_0 v_0}{\beta_1}) v_0 (1 - \beta_0). \end{aligned} \quad (4.54)$$

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, rewriting the integration border

$$v_0 < \min(1, \frac{\beta_1}{\beta_0}) = 1 - \Theta(1 - \frac{\beta_1}{\beta_0})(1 - \frac{\beta_1}{\beta_0}) = 1 - \Theta(\beta_0 - \beta_1)(1 - \frac{\beta_1}{\beta_0}) \quad (4.55)$$

yields

$$\begin{aligned} EPO_0(\beta_0, \beta_1) &= \frac{1}{2}(1 - \beta_0) - \int_0^{1 - \Theta(\beta_0 - \beta_1)(1 - \frac{\beta_1}{\beta_0})} dv_0 (1 - \frac{\beta_0 v_0}{\beta_1}) v_0 (1 - \beta_0) = \\ &= \frac{1}{2}(1 - \beta_0) - (1 - \beta_0) \left\{ \frac{1}{2} [1 - \Theta(\beta_0 - \beta_1)(1 - \frac{\beta_1}{\beta_0})]^2 - \frac{\beta_0}{3\beta_1} [1 - \Theta(\beta_0 - \beta_1)(1 - \frac{\beta_1}{\beta_0})]^3 \right\} = \\ &= (1 - \beta_0) \left\{ \Theta(\beta_0 - \beta_1) \left(\frac{1}{2} - \frac{\beta_1^2}{6\beta_0^2} - \frac{\beta_0}{3\beta_1} \right) + \frac{\beta_0}{3\beta_1} \right\}. \end{aligned} \quad (4.56)$$

Figure 4.6 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.

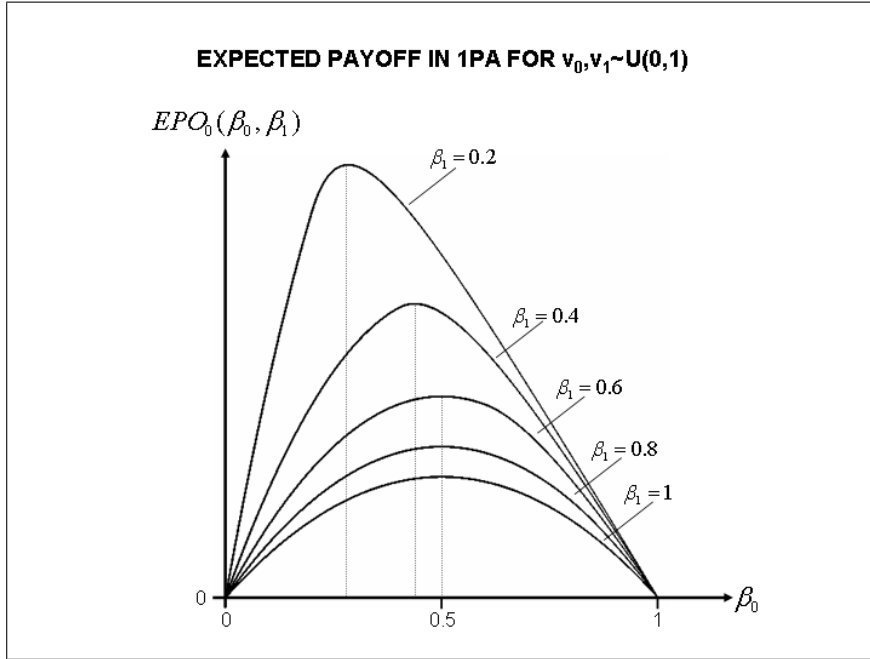


Figure 4.6: Expected Payoff in First-Price Auction for $v_0, v_1 \sim U(0, 1)$

4.4.2 Perfect Memory Best Response to β_1

To analytically derive the optimal response function $\beta_0^{pBR}(\beta_1)$, I have to distinguish the cases

$$\beta_0 \begin{cases} > \beta_1 \\ < \beta_1 \\ = \beta_1 \end{cases} \quad (4.57)$$

For $\beta_0 > \beta_1$, $\Theta(\beta_0 - \beta_1) = 1$, therefore (4.56) reduces to

$$EPO_0(\beta_0, \beta_1) = (1 - \beta_0) \left\{ \frac{1}{2} - \frac{\beta_1^2}{6\beta_0^2} \right\} \quad (4.58)$$

so that the first order condition

$$\frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1) \stackrel{!}{=} 0 = \frac{\beta_1^2}{3\beta_0^3} - \frac{\beta_1^2}{6\beta_0^2} - \frac{1}{2} \quad (4.59)$$

yields

$$\beta_0^{pBR}(\beta_1) = \frac{\beta_1^{\frac{4}{3}} - \beta_1^{\frac{2}{3}} (\sqrt{81 + \beta_1^2} - 9)^{\frac{2}{3}}}{3(\sqrt{81 + \beta_1^2} - 9)^{\frac{1}{3}}} \quad (4.60)$$

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

For $\beta_0 < \beta_1$, $\Theta(\beta_0 - \beta_1) = 0$ and therefore (4.56) reduces to

$$EPO_0(\beta_0, \beta_1) = (1 - \beta_0) \frac{\beta_0}{3\beta_1} \quad (4.61)$$

The first order condition yields the solution $\beta_0^{pBR} = \frac{1}{2}$. 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 $\beta_0 = \beta_1$ I can't proceed straight forward because the first derivative of $EPO_0(\beta_0, \beta_1)$ is discontinuous in $\beta_0 = \beta_1$. Instead, I use the fact that the first derivative to the left of a maximum of a continuous function has to be positive whereas right of the maximum it has to be negative. Therefore I look for the solution to the two inequalities

$$\frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1) \begin{cases} > 0 \text{ for } \beta_0 = \beta_1 - \epsilon \\ < 0 \text{ for } \beta_0 = \beta_1 + \epsilon \end{cases} \quad (4.62)$$

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

$$\frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1) = \frac{1 - 2\beta_0}{\beta_1} \stackrel{!}{>} 0. \quad (4.63)$$

The expression is > 0 if $2\beta_0 < 1$. Resubstituting $\beta_0 = \beta_1 - \epsilon$ and taking $\lim_{\epsilon \rightarrow 0}$ I arrive at

$$\beta_1 \stackrel{!}{\leq} \frac{1}{2}. \quad (4.64)$$

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

$$\frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1) = \frac{\beta_1^2}{3\beta_0^3} - \frac{\beta_1^2}{6\beta_0^2} - \frac{1}{2} \stackrel{!}{<} 0. \quad (4.65)$$

Resubstituting $\beta_0 = \beta_1 + \epsilon$ and taking $\lim_{\epsilon \rightarrow 0}$, I arrive at

$$\beta_1 \stackrel{!}{\geq} \frac{1}{2}. \quad (4.66)$$

In conclusion, the NE $\beta_0 = \beta_1 = \frac{1}{2}$ is the only symmetric solution. The results are summarized in

Theorem 4.5 Consider a two bidder first price open bid auction with linear bidding strategies and $v_i \sim U(0, 1)$. Assume that bidder 1 plays a fixed strategy β_1 and bidder 0 follows a perfect memory best response updating rule. Then, in the long term her strategy is given as

$$\beta_0^{pBR}(\beta_1) = \frac{\beta_1^{\frac{4}{3}} - \beta_1^{\frac{2}{3}}(\sqrt{81 + \beta_1^2} - 9)^{\frac{2}{3}}}{3(\sqrt{81 + \beta_1^2} - 9)^{\frac{1}{3}}} \quad \text{for } \beta_1 < \frac{1}{2}$$

$$\beta_0^{pBR}(\beta_1) = \frac{1}{2} \quad \text{for } \beta_1 > \frac{1}{2}$$
(4.67)

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

$$\beta_1 < \beta_0^{pBR}(\beta_1) < \frac{1}{2} \quad \text{for } \beta_1 < \frac{1}{2}$$

$$\beta_0^{pBR}(\beta_1) = \frac{1}{2} \leq \beta_1 \quad \text{for } \beta_1 \geq \frac{1}{2}$$
(4.68)

Figure 4.7 plots the best response function $\beta_0^{pBR}(\beta_1)$ against β_1 . As stated in Theorem 4.5, for $\beta_1 < \frac{1}{2}$, the best response function satisfies $\beta_1 < \beta_0^{pBR}(\beta_1) < \frac{1}{2}$ - as can easily be seen, since the function plot in this interval is always above the first median - the dashed line - and below $\frac{1}{2}$. It can be seen immediately from the Figure that $\beta_0^* = \beta_1^* = \frac{1}{2}$ is an NE for the 1POBA: If

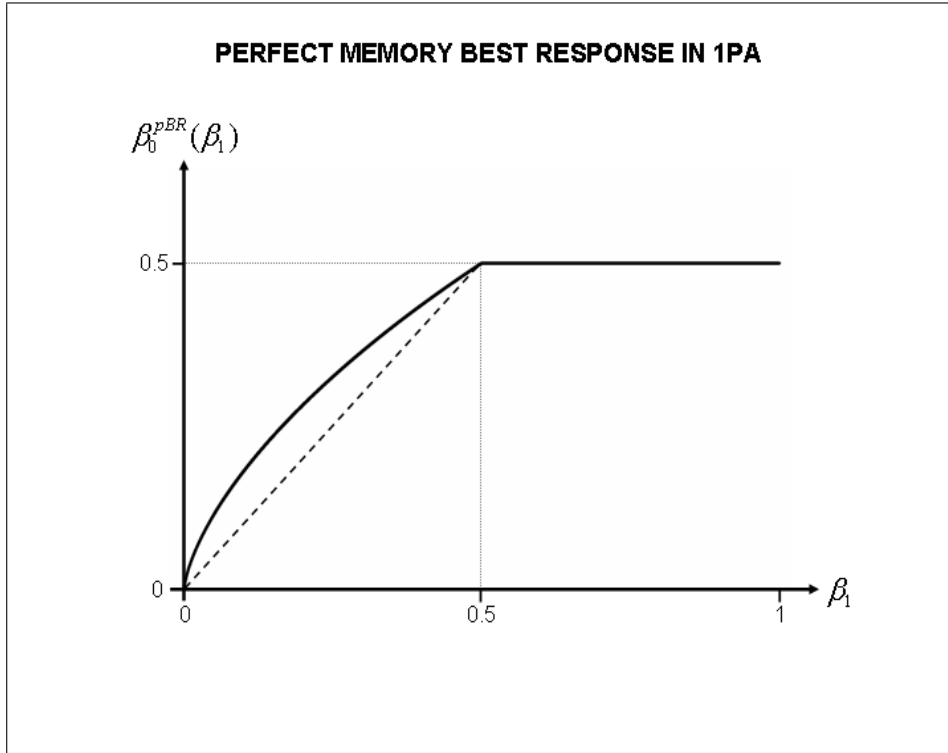


Figure 4.7: Perfect Memory Best Response in First-Price Auction

player 1 plays $\beta_1 = \frac{1}{2}$, player 0's best response is $\beta_0^{pBR}(\beta_1^*) = \frac{1}{2}$ and vice versa. However, if player 0 is unsure whether her opponent plays according to the NE, she will play the NE only if she estimates the opponent's strategy too high. If she estimates it too low, she plays below the NE. So, uncertainty reduces the average bid. The next section explores behavior around the NE in more detail.

4.5 Mutual Perfect Memory Best Response Adaptation

Let me now investigate 2 player 1POBAs under mutual adaptation. Players don't know their opponents strategy for sure. Instead, they have to use previously observed bids to estimate

their opponents strategy and play their best response against this estimate.

For a start let us imagine that player 1 always plays the fixed strategy β_1 . But player 0 doesn't know β_1 . How would bidder 0 come up with an estimate of β_1 ?

Assume that bidder 0 updates her strategy every R rounds. Then she can use the R last bids of player 1 to estimate β_1 as

$$\tilde{\beta}_1 = \frac{E[b_1]}{E[v_1]} = \frac{1}{R} \sum_{t=1}^R \frac{b_1^{(t)}}{0.5} \quad (4.69)$$

since she knows that the values are distributed according to $v_1^{(t)} \sim U(0, 1)$. Her estimate $\tilde{\beta}_1$ from different blocks á R rounds will be distributed normally around the true value β_1 . Therefore, her best response $\beta_0^{\text{pBR}}(\tilde{\beta}_1)$ will in general deviate from the true best response $\beta_0^{\text{pBR}}(\beta_1)$.

Note that even if bidder 1 always plays the NE strategy $\beta_1 = \frac{1}{2}$, bidder 0 plays according to the NE only half the time: If her estimate of bidder 1s strategy is above the NE, $\tilde{\beta}_1 > \frac{1}{2}$, her best response is according to Theorem 4.5 also to play the NE strategy $\beta_0 = \frac{1}{2}$. If she estimates however $\tilde{\beta}_1 < \frac{1}{2}$, her best response is according to Theorem 4.5 below the NE.

By increasing the number of rounds that player 0 uses for estimating β_1 , her estimate of β_1 gets ever more accurate - provided that bidder 1 indeed doesn't change her strategy. In the limit of infinitely many rounds between updating, her play therefore converges to the best response function $\beta_0^{\text{pBR}}(\beta_1)$.

Now I drop the assumption that player 1 always sticks to the same strategy β_1 but instead consider mutually adapting players. To investigate whether heterogeneous bidders converge to the NE, I first consider truncated fictitious play. Players use a strategy for a certain number of rounds, after which they simultaneously update their strategy to the best response against their estimate of the opponent's play. As argued above, even if both bidders start in the NE, each bidder has only a chance of $\frac{1}{2}$ that she also bids according to the NE after the first updating. So, after one updating, there is a $\frac{3}{4}$ chance that at least one bidder bids less than the NE⁸. And if one bidder bids less than the NE, then the chance that the opponent bids less than the NE after the next updating is even higher than $\frac{1}{2}$. So, bidders under truncated fictitious play will not stay in the NE but underbid on average. This underbidding is the less severe, the more rounds bidders consider for strategy updating.

What, if bidders didn't start in the NE? If a bidder - say bidder 0 - starts sufficiently above the NE, the opponent's strategy estimate will be above the NE and therefore, her best response will be - according to Theorem 4.5 - the NE strategy. With the opponent playing the NE strategy after the first updating, the strategy of bidder 0 will move below the NE in subsequent rounds as argued above and both strategies will permanently fluctuate.

If a bidder starts sufficiently below the NE, then her opponent will typically use a higher strategy after the updating:

She will use a higher strategy than β_1 if her estimate $\tilde{\beta}_1$ is $\geq \beta_1$ or even below β_1 but satisfies $\tilde{\beta}_1 \geq (\beta_0^{\text{pBR}})^{-1}(\beta_1)$ where $(\beta_0^{\text{pBR}})^{-1}(\beta_1)$ denotes the inverse function of β_0^{pBR} . Since the estimate $\tilde{\beta}_1$ of β_1 is symmetric around the true value β_1 , the opponent will therefore use in general a higher strategy. So, if both bidders start below the NE, strategies will increase towards the NE. However, the strategies do not converge into the NE but instead permanently fluctuate below it.

In summary, no matter where the bidders start under truncated fictitious play, they end up in permanent mutual strategy updating below the NE. In the long term, the strategy distribution of these fluctuations depends merely on the number of rounds between strategy updating: The more auctions are conducted between the strategy updating, the closer the strategies move to the NE and the narrower the strategy distribution gets. However, for finite time, the strategies do not settle down in the NE and if we allow infinite time between updating, strategies do not change at all. So we see that under truncated fictitious play, strategies don't settle down in the NE.

Fictitious play (perfect memory best response), on the other hand, also doesn't settle down in the NE: Under FP, bidders use all the bids played in the game so far to assess the opponent's strategy and play their best response against the estimated average strategy. Therefore, as soon as a player uses a strategy below the NE⁹, this is never forgotten by the opponent who

⁸Note that this probability of $\frac{3}{4}$ to move out of the NE is independent of the number of rounds between updating.

⁹In the beginning of the game, players underbid since the dynamics there is still very close to myopic play.

calculates the average of all opponent strategies. Since each bidders best response is above the opponent's average strategy, strategies gradually approach the NE from below. However, they never reach the NE, since strategies below the NE are never forgotten.

Up to now I assumed that bidders use their estimate of the opponent's strategy and calculate their best response to it. This requires bidders to know the number of bidders in the auction and the underlying value distribution. Note however, that the same result would hold true if bidders only observe their own strategies and play the strategy that would have generated the highest payoff, had it been the active strategy in the last R rounds. If R is sufficiently high, the strategy that generated the highest payoff will be very close to the true best response against the opponent's strategy. However, for finite R , the random value realizations might not lead the bidder to play her true best response. Due to the concave shape of the best response function, deviations to lower strategies are more severe than deviations to high strategies so that – by a similar argumentation as above – bidders involve in permanent mutual adaptation below the NE. Note that this result is somewhat more general since it describes the behavior of bidders who don't have to know the underlying value distribution or the number of bidders who participate in the auction. The only information that each bidder uses is how much payoff her strategies generate in response to the environment that she faces.

In summary, I arrive at

Theorem 4.6 *Consider a two bidder first price open bid auction with linear bidding strategies and $v_i \sim U(0, 1)$. Assume that each bidder estimates her opponent's strategy in regular intervals from the last R observed bids and revises her strategy as the best response against this estimate. Alternatively, assume that bidders play the strategy that would have generated the highest payoff, had it been the active one in the last R rounds.*

Both models yield qualitatively the same dynamics: For infinite R , strategies converge to the NE from below. However they reach it only after an infinitely long time. For finite R the process leads to underbidding when compared with NE play. This deviation from the NE is accompanied by a permanent fluctuation of bidders' strategies. The lower R , the more pronounced are underbidding and strategy fluctuations.

4.6 Second-Price Auctions

The payoff-maximizing bidding strategy in second-price open and closed bid auctions is always bidding the true value. This result is well established in the literature and holds for any number of participants and all kinds of risk aversion.

This section rederives this result using the methodology of the previous sections. It serves mainly as a finger-exercise and a double-check for the mathematical methodology.

In a second-price auction the payoff for player 0 is given by

$$PO_0(\beta_0, \beta_1, v_0, v_1) = (v_0 - \beta_1 v_1) \Theta\left(\frac{\beta_0 v_0}{\beta_1} - v_1\right) \quad (4.70)$$

The expected payoff for a fixed v_0 and $v_1 \sim U(0, 1)$ can be calculated as

$$EPO_0(\beta_0, \beta_1, v_0) = \int_0^1 dv_1 (v_0 - \beta_1 v_1) \Theta\left(\frac{\beta_0 v_0}{\beta_1} - v_1\right). \quad (4.71)$$

The integrand takes positive values if

$$v_1 < \min\left(1, \frac{\beta_0 v_0}{\beta_1}\right) = 1 - \Theta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right) \left(1 - \frac{\beta_0 v_0}{\beta_1}\right) = 1 - \Theta(\beta_1 - \beta_0 v_0) \left(1 - \frac{\beta_0 v_0}{\beta_1}\right). \quad (4.72)$$

Therefore the integral transforms to

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

There, as I have shown above, on average the myopic best response against *any* strategy of the opponent is below the NE.

$$= v_0[1 - \Theta(\beta_1 - \beta_0 v_0)(1 - \frac{\beta_0 v_0}{\beta_1})] - \frac{\beta_1}{2}[1 - \Theta(\beta_1 - \beta_0 v_0)(1 - \frac{\beta_0 v_0}{\beta_1})]^2. \quad (4.73)$$

To obtain the expected payoff for $v_0 \sim U(0, 1)$, I integrate this expression over v_0 :

$$EPO_0(\beta_0, \beta_1) = \int_0^1 dv_0 \{v_0[1 - \Theta(\frac{\beta_1}{\beta_0} - v_0)(1 - \frac{\beta_0 v_0}{\beta_1})] - \frac{\beta_1}{2}[1 - \Theta(\frac{\beta_1}{\beta_0} - v_0)(1 - \frac{\beta_0 v_0}{\beta_1})]^2\}. \quad (4.74)$$

There are two cases for the integral that I have to distinguish: Either $\beta_0 > \beta_1$ or $\beta_0 < \beta_1$.

Case 1:

$$\beta_1 > \beta_0 \Rightarrow \frac{\beta_1}{\beta_0} > 1 \Rightarrow \Theta(\frac{\beta_1}{\beta_0} - v_0) = 1 \quad \forall v_0 \in (0, 1) \quad (4.75)$$

Then, the integral transforms into

$$EPO_0(\beta_0, \beta_1) = \int_0^1 dv_0 v_0^2 \left\{ \frac{\beta_0}{\beta_1} - \frac{\beta_0^2}{2\beta_1} \right\} = \frac{\beta_0}{3\beta_1} - \frac{\beta_0^2}{6\beta_1} \quad (4.76)$$

Case2:

$$\beta_1 < \beta_0 \Rightarrow \frac{\beta_1}{\beta_0} < 1 \Rightarrow \Theta(\frac{\beta_1}{\beta_0} - v_0) = \begin{cases} 1 & \forall v_0 \in (0, \frac{\beta_1}{\beta_0}) \\ 0 & \forall v_0 \in (\frac{\beta_1}{\beta_0}, 1) \end{cases}. \quad (4.77)$$

Then,

$$EPO_0(\beta_0, \beta_1) = \int_0^{\frac{\beta_1}{\beta_0}} dv_0 v_0^2 \left(\frac{\beta_0}{\beta_1} - \frac{\beta_0^2}{2\beta_1} \right) + \int_{\frac{\beta_1}{\beta_0}}^1 dv_0 v_0 - \frac{\beta_1}{2} = \frac{\beta_1^2}{3\beta_0} - \frac{\beta_1^2}{6\beta_0^2} + \frac{1}{2} - \frac{\beta_1}{2}. \quad (4.78)$$

Combining the two cases yields the general expression for the expected payoff in second-price auctions:

$$EPO_0(\beta_0, \beta_1) = \Theta(\beta_0 - \beta_1) \left(\frac{\beta_1^2}{3\beta_0} - \frac{\beta_1^2}{6\beta_0^2} + \frac{1}{2} - \frac{\beta_1}{2} \right) + \Theta(\beta_1 - \beta_0) \left(\frac{\beta_0}{3\beta_1} - \frac{\beta_0^2}{6\beta_1} \right). \quad (4.79)$$

Figure 4.8 shows the expected payoff of bidder 0 for different strategies of bidder 1. Graphical inspection shows that this payoff function has its maximum in $\beta_0 = 1$ for all values of β_1 in the interval $(0, 1)$. For low values of β_1 the payoff-function is quite flat to the right of the maximum, so the reduction in expected payoff when bidding above the optimal strategy $\beta_0 = 1$ is small for even large deviations.¹⁰

This intuitively clear. Imagine that bidder 1 plays according to the NE, $\beta_1 = 1$. Expected payoff is for every nonnegative bid $\beta_0 > 0$ positive: Even $\lim_{\beta_0 \rightarrow \infty} EPO(\beta_0, \beta_1) = 0+$, since if player 0 bids infinitely much, she obtains in each round the payoff $v_0 - v_1$. Since $v_0, v_1 \sim U(0, 1)$ the probabilities of winning and losing sum up to 0.

To analytically calculate the maximum of this function, I derive the FOCs for (4.79):¹¹

$$\begin{aligned} 0 &\stackrel{!}{=} \frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1) = \delta(\beta_0 - \beta_1) \left(\frac{\beta_1^2}{3\beta_0} - \frac{\beta_1^2}{6\beta_0^2} + \frac{1}{2} - \frac{\beta_1}{2} \right) + \\ &+ \Theta(\beta_0 - \beta_1) \left(\frac{\beta_1^2}{3\beta_0^3} - \frac{\beta_1^2}{3\beta_0^2} \right) - \delta(\beta_1 - \beta_0) \left(\frac{\beta_0}{3\beta_1} - \frac{\beta_0^2}{6\beta_1} \right) + \Theta(\beta_1 - \beta_0) \left(\frac{1}{3\beta_1} - \frac{\beta_0}{3\beta_1} \right) \end{aligned} \quad (4.80)$$

Like in the case of the 1POBA, I distinguish the cases $\beta_0 > \beta_1$, $\beta_0 < \beta_1$ and $\beta_0 = \beta_1$.

For $\beta_0 > \beta_1$, 4.6 transforms to

$$\frac{\beta_1^2}{3\beta_0^3} = \frac{\beta_1^2}{3\beta_0^2} \quad (4.81)$$

¹⁰Compare this to the experimental observation that inexperienced bidders bid above the NE also in 2PAs.

¹¹Remember that the derivative of the Θ -function is the δ -function; see Appendix A.

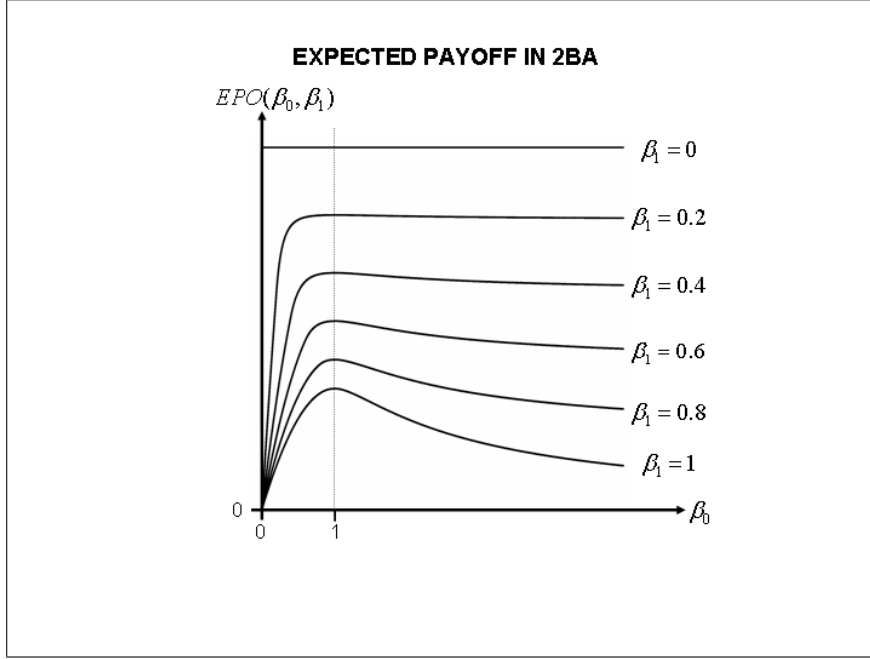


Figure 4.8: Expected Payoff in Second-Price Auction

with the solution $\beta_0 = 1$.

For $\beta_0 < \beta_1$, equation 4.6 transforms to

$$\frac{1}{3\beta_1} = \frac{\beta_0}{3\beta_1} \quad (4.82)$$

also with the solution $\beta_0 = 1$.

The case $\beta_0 = \beta_1$ has to be investigated by looking at the sign of the first derivative left and right of the point $\beta_0 = \beta_1$:

The condition

$$\frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1) \Big|_{\beta_0 = \beta_1 - \epsilon} = \frac{1}{3\beta_1} - \frac{1}{3} + \frac{\epsilon}{\beta_1} \stackrel{!}{>} 0 \quad (4.83)$$

transforms by taking $\lim_{\epsilon \rightarrow 0}$ into the condition $\beta_1 \stackrel{!}{\leq} 1$.

The condition

$$\frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1) \Big|_{\beta_0 = \beta_1 + \epsilon} = \frac{\beta_1^2}{3(\beta_1 + \epsilon)^3} - \frac{\beta_1^2}{3(\beta_1 + \epsilon)^2} \stackrel{!}{<} 0 \quad (4.84)$$

likewise transforms by taking $\lim_{\epsilon \rightarrow 0}$ into the condition $\beta_1 \stackrel{!}{\geq} 1$.

So we see that the only symmetric solution is $\beta_0 = \beta_1 = 1$. The findings of this section therefore reestablish the following

Theorem 4.7 *The perfect memory best response $\beta_0^{pBR}(\beta_1)$ against any bidding strategy β_1 in a private value second-price auction with risk neutral bidders and private values $v_0, v_1 \sim U(0, 1)$ is given by*

$$\beta_0 = 1 \quad \forall \beta_1. \quad (4.85)$$

4.7 Revenue Equivalence Revisited

In all considerations up to now I have implicitly assumed that each player believes that her opponent will keep her strategy unchanged for the next R auctions. This assumption is crucial.

A player who knows that the other player may change her strategy plays different than a player who assumes that her opponent will stick to her strategy.

To investigate the consequences of strategic sophistication, consider player 0 in a 1PA who starts with some bid $\beta_0^{(0)}$. Instead of assuming after the first R auctions that player 1 will stick to her initial strategy $\beta_1^{(0)}$, player 0 now assumes player 1 to switch to the best response against her estimate of player 0s strategy $\beta_1^{(1)} = \beta_1^{\text{pBR}}(E[\beta_0^{(0)}])$. Here, $E[\beta_0^{(0)}]$ denotes the estimate of player 0s strategy by player 1.¹² Player 0s best response is now to play $\beta_0^{\text{pBR}}(\beta_1^{(1)}) = \beta_0^{\text{pBR}}(\beta_1^{\text{pBR}}(E[\beta_0^{(0)}]))$.

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 further ad libitum. Since

$$\beta_0 < \beta_1^{\text{pBR}}(\beta_0^{(0)}) < \beta_0^{\text{pBR}}(\beta_1^{\text{pBR}}(\beta_0^{(0)})) < \dots < \frac{1}{2} \quad (4.86)$$

we see that this chain converges to the NE. So, if best response bidders have an infinite foresight horizon H , the process converges to the NE within one round:

$$\lim_{H \rightarrow \infty} \beta_0^{(t)} = \lim_{H \rightarrow \infty} \beta_1 = \frac{1}{2} \quad \forall t \quad (4.87)$$

This is intuitively clear: The infinitely iterated process describes perfectly rational bidders who assume that their opponents are also perfectly rational and that their opponents also assume that they are 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 IPOBAs will move closer to the NE than in the case of unsophisticated players but they will still persistently fluctuate below the NE.

In summary I arrive at

Theorem 4.8 *Consider a two bidder open bid auction with $v_i \sim U(0, 1)$. Bidders with a limited foresight horizon H adapt each R rounds their strategies to the best response to the H^{th} -order best response of their opponent.*

Then, a repeated first-price open bid auction gives lower expected seller revenue and higher earnings volatility than a second-price open bid auction. The Revenue Equivalence Theorem emerges under the best response dynamics as $\lim_{H \rightarrow \infty}$ or $\lim_{R \rightarrow \infty}$ and $\lim_{t \rightarrow \infty}$.

¹²In the case of $v_i \sim U(0, 1)$ this is simply $\sum_{i=1}^R \frac{2b_0}{R}$

Chapter 5

Addenda to the Best Response Dynamics

This chapter deals with additional aspects of the best response dynamics which complicate matters further but yield no significant new conclusions. The basic result that first-price auctions do not converge to the NE still holds in the extended models.

In the first section I investigate the consequences of allowing bidders under perfect memory best response to have non-linear bidding functions. It turns out that the best response function then is different. However, bidders in first-price auctions again don't converge to the NE.

In the second section, I consider a more sophisticated version of myopic best response where bidders use the last observed bid and their knowledge of the common value distribution to determine their opponent's strategy. Also under this assumption, play doesn't converge to the NE – even not on average.

In the third section I qualitatively argue that bidders in sealed bid first-price auctions will even less converge to the NE than in the open bid case because bidders get less information about their opponent.

5.1 Value Dependent Perfect Memory Best Response

So far I have assumed that bidders use linear bidding functions. This means that they always bid $b_i(v_i) = \beta_i v_i$ with the same β_i , no matter what their value is. Mathematically this corresponded to the integration of player 0's values over the interval $[0, 1]$ in equation 4.54 *before* we take the derivative to determine the maximum.

In general, this will be an unrealistic assumption: Bidders with a value close to 1 who bid against an unknown β_1 may bid less aggressively since their chance of acquiring the object is already very high and so they will try to increase their payoff by bidding less.

In this section I explore the consequences of allowing nonlinear bidding functions. I will look for best responses $\beta_0^{\text{pBR}}(\beta_1, v_0)$ where player 0 takes into account her own value v_0 . Note that I still assume that bidder 1 plays a linear β_1 strategy here.

How can I proceed mathematically? From equation 4.53 I know that

$$EPO_0(\beta_0, \beta_1, v_0) = [1 - \Theta(1 - \frac{\beta_0 v_0}{\beta_1})(1 - \frac{\beta_0 v_0}{\beta_1})]v_0(1 - \beta_0) \quad (5.1)$$

I determine the value-dependent best response function $\beta_0^{\text{pBR}}(\beta_1, v_0)$ by maximizing $EPO_0(\beta_0, \beta_1, v_0)$:

$$\begin{aligned} 0 &\stackrel{!}{=} \frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \beta_1, v_0) = \\ &= \left\{ -\delta \left(1 - \frac{\beta_0 v_0}{\beta_1}\right) \left(-\frac{v_0}{\beta_1}\right) \left(1 - \frac{\beta_0 v_0}{\beta_1}\right) - \Theta \left(1 - \frac{\beta_0 v_0}{\beta_1}\right) \left(-\frac{v_0}{\beta_1}\right) \right\} v_0 (1 - \beta_0) - \end{aligned}$$

$$-\left\{1 - \Theta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right)\left(1 - \frac{\beta_0 v_0}{\beta_1}\right)\right\} v_0, \quad (5.2)$$

which can be further transformed to

$$\frac{v_0}{\beta_1} \left\{ \delta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right)\left(1 - \frac{\beta_0 v_0}{\beta_1}\right) + \Theta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right) \right\} (1 - \beta_0) - \left\{ 1 - \Theta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right)\left(1 - \frac{\beta_0 v_0}{\beta_1}\right) \right\} \stackrel{!}{=} 0. \quad (5.3)$$

Figure 5.1 shows the first derivative of the expected payoff for a value of $v_0 = 0.8$. Note that for $\beta_1 \geq v_0/2$, the FOC yields a best response of 0.5, i.e., of the NE. For $\beta_1 < v_0/2$, however, the best response is below the NE.

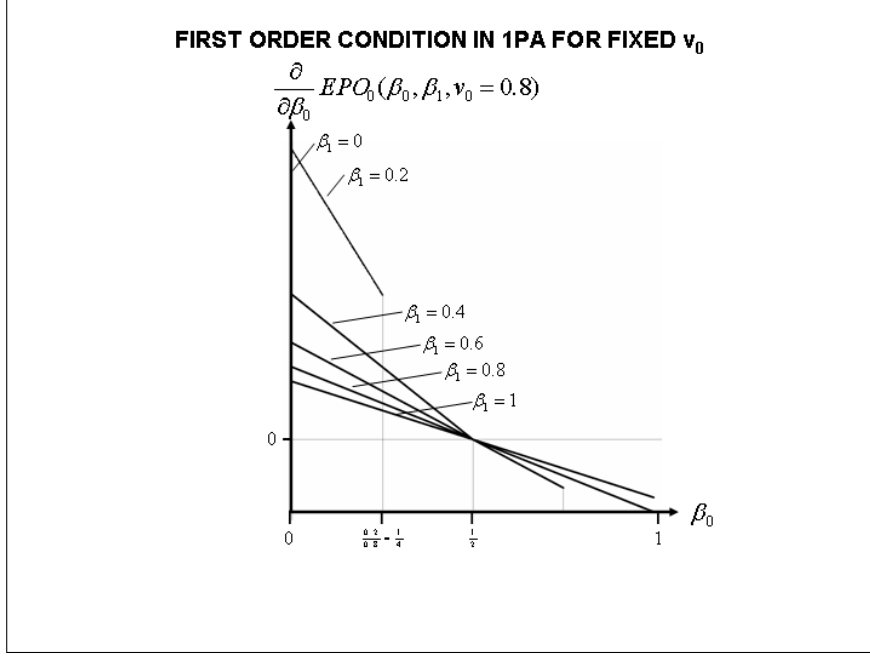


Figure 5.1: First Order Condition in First-Price Auction for Fixed v_0

To find the closed-form solution to (5.3), I have to distinguish three cases

$$\frac{\beta_0 v_0}{\beta_1} \begin{cases} < 1 \\ > 1 \\ = 1 \end{cases} \quad (5.4)$$

since the values of the Θ and the δ function depend critically on this quotient.

If $\frac{\beta_0 v_0}{\beta_1} < 1$, then $\Theta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right) = 1$ and $\delta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right) = 0$. Therefore, (5.3) is transformed into

$$\frac{v_0}{\beta_1} (1 - \beta_0) = \frac{\beta_0 v_0}{\beta_1} \quad (5.5)$$

with the unique solution

$$\beta_0 = \frac{1}{2}. \quad (5.6)$$

Plugging this back into the condition $\frac{\beta_0 v_0}{\beta_1} < 1$ shows that this is a valid solution only for $v_0 < 2\beta_1$.

If $\frac{\beta_0 v_0}{\beta_1} > 1$, then $\Theta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right) = \delta\left(1 - \frac{\beta_0 v_0}{\beta_1}\right) = 0$, so (5.3) reduces to

$$0 = 1, \quad (5.7)$$

so there exist no solutions with $\frac{\beta_0 v_0}{\beta_1} > 1$.

To check for solutions with $\frac{\beta_0 v_0}{\beta_1} = 1$, I can not directly solve equation (5.3), since the value of $\delta(0) \cdot 0$ is not defined. Intuitively, the reason is that $EPO_0(\beta_0, \beta_1, v_1)$ has a kink at $\beta_0 = \frac{\beta_1}{v_0}$ so

that its first derivative is not uniquely determined there. As a workaround, as above, I evaluate the derivative (5.3) at the points $\beta_0 = \frac{\beta_1}{v_0} \pm \epsilon$ and keep in mind that

$$\beta_0 = \frac{\beta_1}{v_0} \text{ is a maximum} \iff \begin{cases} \frac{\partial}{\partial \beta_0} EPO(\frac{\beta_1}{v_0} - \epsilon, \beta_1, v_0) > 0 \text{ (condition 1)} \\ \frac{\partial}{\partial \beta_0} EPO(\frac{\beta_1}{v_0} + \epsilon, \beta_1, v_0) < 0 \text{ (condition 2)} \end{cases} \quad (5.8)$$

If $\beta_0 = \frac{\beta_1}{v_0} - \epsilon$ then $\Theta(1 - \frac{\beta_0 v_0}{\beta_1}) = 1$ and $\delta(1 - \frac{\beta_0 v_0}{\beta_1} = 0)$, so condition 1 yields

$$\frac{v_0}{\beta_1} (1 - \frac{\beta_1}{v_0} + \epsilon) - 1 \stackrel{!}{>} 0 \quad (5.9)$$

With $\epsilon \rightarrow 0$ I obtain

$$\frac{v_0}{\beta_1} \stackrel{!}{\geq} 2, \quad (5.10)$$

so $\beta_0 = \frac{\beta_1}{v_0}$ is a solution for

$$v_0 \geq 2\beta_1, \quad (5.11)$$

Similarly, condition 2 leads to

$$\frac{v_0}{\beta_1} (1 - \frac{\beta_1}{v_0} + \epsilon) - 1 < 0 \quad (5.12)$$

which is always satisfied. Collecting these results, I arrive at the following

Theorem 5.1 Consider a two bidder first-price open bid auction with value-dependent strategies and $v_i \sim U(0, 1)$. Assume that bidder 1 plays a fixed linear strategy β_1 and bidder 0 follows a perfect memory best response strategy updating rule. Then, in the long term her strategy is given as

$$\beta_0(\beta_1, v_0) = \begin{cases} \frac{1}{2} & \text{for } v_0 \leq 2\beta_1 \\ \frac{\beta_1}{v_0} & \text{for } v_0 \geq 2\beta_1 \end{cases} \quad (5.13)$$

Figure 5.2 shows the best response $\beta_0^{\text{BR}}(\beta_1)$ in dependence of v_0 for different values of β_1 .

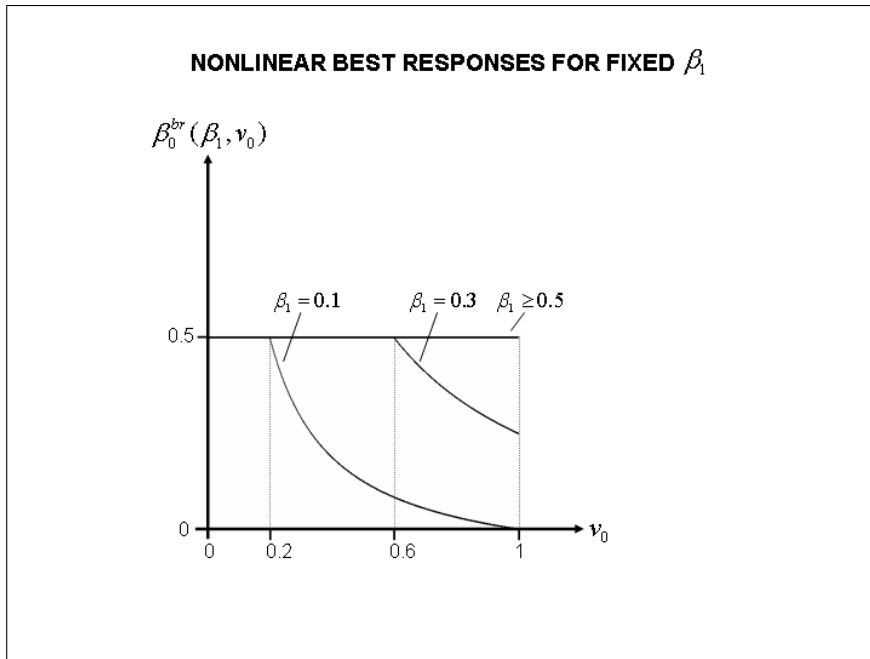


Figure 5.2: Nonlinear Best Responses for Fixed β_1

The main question of interest in this section is of course, whether the value-dependent bidding functions can lead mutually adapting bidders to NE bidding. Imagine, we start in the NE $\beta_0^{(0)} = \beta_1^{(0)} = \frac{1}{2}$ and we run sufficiently many auctions that bidders can identify their opponent's strategy correctly. After bidders choose their strategies for the next round. Then, for any value the best response is $\beta_0^{(1)} = \beta_1^{(1)} = \frac{1}{2}$. So bidders stay in the NE forever.

But – like in the last chapter – what happens if bidders didn't start in the NE or if bidders have some uncertainty about the strategy their opponent pursues¹? If $\beta_1 < \frac{1}{2}$ then there exist always values $v_0 \in (2\beta_1, 1)$ where the best response β_0^{BR} is also a strategy that is smaller than $\frac{1}{2}$. Consequently, the average best response strategy against any strategy $< \frac{1}{2}$ is $< \frac{1}{2}$ but bigger than the strategy of the opponent. The same is of course true for bidder 1. So, the situation is qualitatively completely analogous to the case of linear bidding strategies. In summary I arrive at the following

Theorem 5.2 *Consider a 1POBA where bidders bid according to value-dependent bidding functions. Then, the same conclusions as in the case of linear bidding functions hold. In particular, it takes bidders infinitely long to reach the NE if they pursue fictitious play and they involve in permanent mutual adaptation below the NE if they pursue truncated fictitious play.*

5.2 Probability-Weighted Myopic Best Response Play

My previous model of myopic play was a bit unsatisfactory. After all, player 0 should be sophisticated enough to take into account that the best response in last round's auction depended on the last rounds realization of $v_0^{(t-1)}$ and $v_1^{(t-1)}$. But by playing at time t the last rounds best response $\beta_0^{\text{pBR}; (t-1)}(\beta_1, v_0^{(t-1)}, v_1^{(t-1)})$ she implicitly would assume that next rounds' realizations of $v_0^{(t)}$ and $v_1^{(t)}$ are the same as at time $(t-1)$.

As a more sophisticated alternative I assume now that player 0 calculates her strategy $\beta_0^{(t)}$ from last rounds' observed play $b_1^{(t-1)} = \beta_1^{(t-1)} v_1^{(t-1)}$ as follows:

- In the first step, player 0 estimates from $b_1^{(t-1)}$ the strategies that player 1 might have used at time $(t-1)$. This gives rise to a probability distribution $f^{b_1 = \beta_1 v_1^{\text{old}}}(\tilde{\beta}_1)$ for player 0s estimate of player 1s active strategy given the observed bid. Here, β_1 is the strategy that player 1 actually uses, $\tilde{\beta}_1$ runs through player 1s strategy set, $f^{b_1}(\tilde{\beta}_1)$ is the probability density that player 0 assigns to each strategy $\tilde{\beta}_1$ given the observed bid b_1 and v_1^{old} is the last rounds' realization of player 1s value
- In the second step, player 0 uses the probability distribution $f(\tilde{\beta})$ to maximize her expected payoff at time t by choosing an appropriate probability-weighted myopic best response $\beta_0^{\text{pmBR}}(\beta_1)$.

Note that the rather complicated estimation of the opponent's strategy is only necessary under (near to) myopic play. Under perfect memory play (or truncated fictitious play with sufficiently long updating intervals), as players observe bids from many rounds they can directly derive the opponent's strategy as $\tilde{\beta}_1 = \frac{\text{numRounds}}{\sum_{t=1}^{\text{numRounds}} \frac{2b_1^{(t)}}{\text{numRounds}}}$ if $v_1 \sim U(0, 1)$. This estimate is asymptotically normally distributed around the true value of β_1 and gets more accurate, the higher the number of considered rounds.

5.2.1 Strategy Estimation from Observed Bids

From $b_1 = \beta_1 v_1$ it follows that $\beta_1 = \frac{b_1}{v_1}$. But since player 0 doesn't know v_1 , she can only derive a probability distribution $f^{b_1}(\tilde{\beta}_1)$ for the possible values of β_1 . In this section I derive this distribution via Bayesian updating.²

¹Such that they give positive probabilities to values above and below the true strategy of their opponent

²A special thanks to Prof. Lux for the hints for this section.

A priori, bidder 0 doesn't know anything about the value or the strategy of her opponent. Therefore, her priors are $\beta_1 \sim U[0, 1]$ and $v_1 \sim U[0, 1]$. To calculate the probability density $f^{b_1}(\beta_1)$ if bidder 0 observes the bid b_1 , I use the Bayesian formula for the posteriori distribution

$$f^{b_1}(\beta_1) = f(\beta_1 | b_1) = \frac{f(b_1 | \beta_1)f(\beta_1)}{\int d\tilde{\beta}_1 f(\tilde{\beta}_1)f(b_1 | \tilde{\beta}_1)}. \quad (5.14)$$

Since a priori bidder 0 assumes all strategies of her opponent to be equally likely on the interval $[0, 1]$, I set $f(\beta_1) = 1$.

Moreover, if player 1 plays strategy β_1 , her bids will be uniformly distributed on $(0, \beta_1)$ because $v_1 \sim U[0, 1]$. Therefore

$$f(b_1 | \beta_1) = \begin{cases} \frac{1}{\beta_1} & \text{for } b_1 \leq \beta_1 \\ 0 & \text{for } b_1 > \beta_1, \end{cases} \quad (5.15)$$

or in closed form

$$f(b_1 | \beta_1) = \frac{1}{\beta_1} \Theta(\beta_1 - b_1). \quad (5.16)$$

The integral in the denominator is given by

$$\begin{aligned} \int_0^1 d\tilde{\beta}_1 f(\tilde{\beta}_1)f(b_1 | \tilde{\beta}_1) &= \int_0^1 d\tilde{\beta}_1 1 \frac{1}{\tilde{\beta}_1} \Theta(\tilde{\beta}_1 - b_1) = \\ &= \int_{b_1}^1 d\tilde{\beta}_1 \frac{1}{\tilde{\beta}_1} = -\ln b_1. \end{aligned} \quad (5.17)$$

Putting all this together, the posteriori distribution is given by the following

Theorem 5.3 *Consider a two bidder first-price open bid auction with linear bidding strategies and $v_i \sim U(0, 1)$. Assume that bidder 1 plays a fixed strategy β_1 and bidder 0 observes in each round the bid $\beta_1 v_1$ that bidder 1 is placing. Then, player 0 derives the PDF that player 1 strategy is $\tilde{\beta}_1$ if she observes the bid $\beta_1 v_1$ as*

$$f^{\beta_1 v_1}(\tilde{\beta}_1) = -\frac{1}{\beta_1 \ln b_1} \Theta(\tilde{\beta}_1 - b_1). \quad (5.18)$$

Figure 5.3 shows the strategy distribution that bidder 0 derives from different bids of bidder 1.

5.2.2 Probability-Weighted Myopic Best Response Play

In this subsection I analyze a more sophisticated version of myopic best response than in the last chapter. In the derivation of Theorem 4.1 I assumed that player 0 plays in each round the strategy that would have maximized her payoff in the last round. This neglected the influence of the randomness bidder 1s values v_1 .

Instead, I assume now that player 0 uses her observation of player 1s last rounds bid to estimate player 1s strategy and maximizes based on this distribution her expected payoff. According to Theorem 5.3 she can calculate the distribution $f^{\beta_1 v_1^{\text{old}}}(\tilde{\beta}_1)$ of player 1s true strategy β_1 that depends on the realization of player 1s value in the previous round, v_1^{old} . If player 0 knew player 1s strategy β_1 for sure, she could calculate her expected payoff $EPO_0(\beta_0, \beta_1)$ by equation (4.56). However, player 0 doesn't know β_1 for sure but instead has to weight the expected payoffs with the according probabilities $f^{\beta_1 v_1^{\text{old}}}(\tilde{\beta}_1)$ for bidder 1's expected strategy,

$$EPO_0^{\beta_1 v_1^{\text{old}}}(\beta_0, b_1) = \int_0^1 d\tilde{\beta}_1 f^{\beta_1 v_1^{\text{old}}}(\tilde{\beta}_1) EPO_0(\beta_0, \tilde{\beta}_1), \quad (5.19)$$

where

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

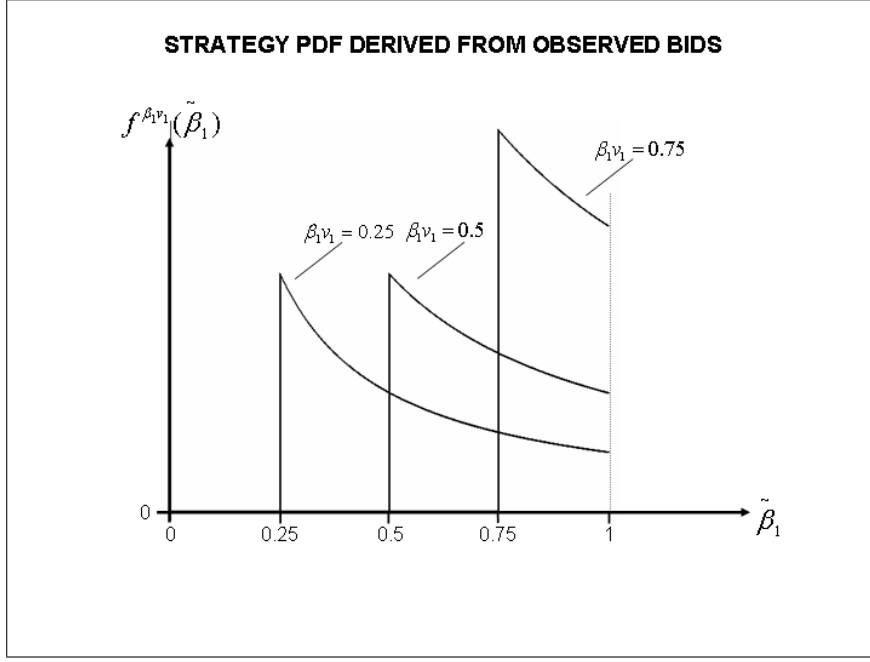


Figure 5.3: Strategy PDF Derived from Observed Bids

$$= \begin{cases} (1 - \beta_0) \frac{\beta_0}{3\beta_1} & \text{for } \beta_0 < \beta_1 \\ (1 - \beta_0) \left(\frac{1}{2} - \frac{\beta_1^2}{6\beta_0^2} \right) & \text{for } \beta_0 > \beta_1 \end{cases} \quad (5.20)$$

and

$$f^{\beta_1 v_1^{\text{old}}}(\tilde{\beta}_1) = -\frac{1}{\tilde{\beta}_1 \ln b_1} \Theta(\tilde{\beta}_1 - b_1). \quad (5.21)$$

For the analysis I treat the two cases $\beta_0 < b_1$ and $\beta_0 > b_1$ separately:

If $\beta_0 < b_1$ then certainly also $\beta_0 < \beta_1$ and therefore the Θ -function in the expression for $EPO_0(\beta_0, \beta_1)$ is zero. Therefore the probability-weighted expected payoff is given as

$$EPO_0^{b_1}(\beta_0, b_1) = \int_{b_1}^1 d\tilde{\beta}_1 \left(-\frac{1}{\tilde{\beta}_1 \ln b_1} \right) (1 - \beta_0) \frac{\beta_0}{3\tilde{\beta}_1} = \frac{(1 - \beta_0)\beta_0(1 - b_1)}{3b_1 \ln b_1} \quad (5.22)$$

for $\beta_0 < b_1$.

If $\beta_0 > b_1$ then β_0 might be bigger or smaller than β_1 . Therefore, the expression for the expected payoff consists in this case of two parts,

$$\begin{aligned} EPO_0^{b_1}(\beta_0, b_1) &= \int_{b_1}^{\beta_0} d\tilde{\beta}_1 \left(-\frac{1}{\tilde{\beta}_1 \ln b_1} \right) (1 - \beta_0) \left(\frac{1}{2} - \frac{\beta_1^2}{6\beta_0^2} \right) + \int_{\beta_0}^1 d\tilde{\beta}_1 \left(-\frac{1}{\tilde{\beta}_1 \ln b_1} \right) (1 - \beta_0) \frac{\beta_0}{3\tilde{\beta}_1} = \\ &= \frac{(1 - \beta_0)(4\beta_0^3 + 6\beta_0^2 \ln b_1 - 3\beta_0^2 - b_1^2 - 6\beta_0^2 \ln \beta_0)}{12\beta_0^2 \ln b_1}. \end{aligned} \quad (5.23)$$

In summary, the expected payoff of bidder 0 if she observes bidder 1 bidding b_1 and answers with β_0 is given as

$$\begin{aligned} EPO_0^{b_1}(\beta_0, b_1) &= \frac{(1 - \beta_0)\beta_0(1 - b_1)}{3b_1 \ln b_1} \Theta(b_1 - \beta_0) + \\ &+ \frac{(1 - \beta_0)(4\beta_0^3 + 6\beta_0^2 \ln b_1 - 3\beta_0^2 - b_1^2 - 6\beta_0^2 \ln \beta_0)}{12\beta_0^2 \ln b_1} \Theta(\beta_0 - b_1). \end{aligned} \quad (5.24)$$

Next I derive the best response function by maximizing the expected payoff. If $\beta_0 < b_1$ then

$$EPO_0^{b_1}(\beta_0, b_1) = \frac{(1 - \beta_0)\beta_0(1 - b_1)}{3b_1 \ln b_1} \quad (5.25)$$

The FOC is given by $\frac{\partial}{\partial \beta_0} EPO_0 \stackrel{!}{=} 0$ and yields

$$\frac{1 - 2\beta_0(1 - b_1) - b_1}{3b_1 \ln b_1} \stackrel{!}{=} 0 \Rightarrow \beta_0 = \frac{1}{2}. \quad (5.26)$$

Plugging this back into the condition $\beta_0 < b_1$ shows that this solution is valid for all $b_1 > \frac{1}{2}$. This is intuitively clear: If the bid is bigger than $\frac{1}{2}$, then certainly also the strategy is bigger than $\frac{1}{2}$ and therefore the estimated strategy distribution has support only above the NE. Since the best response against any strategy $> \frac{1}{2}$ is $\frac{1}{2}$, the best response against a distribution that consists only of strategies bigger than $\frac{1}{2}$ must also be $\frac{1}{2}$.

If $\beta_0 > b_1$, then

$$EPO_0^{b_1}(\beta_0, b_1) = \frac{(1 - \beta_0)(4\beta_0^3 + 6\beta_0^2 \ln b_1 - 3\beta_0^2 - b_1^2 - 6\beta_0^2 \ln \beta_0)}{12\beta_0^2 \ln b_1}. \quad (5.27)$$

The FOC yields

$$6\beta_0^2 - 13\beta_0^3 + 8\beta_0^4 - 2b_1^2 + \beta_0 b_1^2 - 6\beta_0^3 \ln \beta_0 + 6\beta_0^3 \ln b_1 \stackrel{!}{=} 0. \quad (5.28)$$

Unfortunately there is no closed-form solution to this logarithmic equation. Instead I use Mathematica to solve the FOC numerically. Figure 5.4 shows the probability weighted myopic best response of bidder 0 if she observes the bid b_1 of her opponent. The graph is qualitatively




Figure 5.4: Probability Weighted Myopic Best Response Against Observed Bids

the same as Figure 4.7, only that now the x -axis consists of bidder 1s bid instead of her strategy. Note that the left part of the plot is higher than the corresponding part in Figure 4.7. This is intuitively clear: Bidder 0s estimate of bidder 1s strategy based on the bid b_1 puts positive probability on every strategy between b_1 and 1. Since the best response against higher strategies is also higher, this pushes the best response function up. Completely analogous to the discussion in the case of perfect memory best response, follows

Theorem 5.4 *Consider a two bidder first-price open bid auction with linear bidding strategies and $v_i \sim U(0, 1)$. Assume that both bidders revise their strategies according to probability-weighted myopic best response: After each round they calculate from their opponent's bid a probability distribution of her strategy. Then they use this distribution to calculate their own best strategy to maximize next round's expected payoff.*

The dynamics of this game leads to persistent underbidding and does not converge to the NE.

Note that it is crucial for probability-weighted myopic best response that bidders know exactly, how many bidders participate in the auction and the underlying value distribution. Otherwise the bidders could not calculate the expected payoff. There is no simple learning rule that gives rise to probability-weighted myopic best response. Therefore it is mainly of academic interest but will in general not be very relevant for real auctions.

5.3 Sealed Bid Auctions

So far, I have only considered open bid auctions where each bidder can observe her opponent's bid, no matter whether she has won or lost.

Many auction formats, however are sealed bid. This means that only the winning bid but not any losing bid is published. How would this influence the outcome of repeated auctions? The main problem here is that I have to define a rule that specifies, how winning bidders behave that get no new information about their opponent's strategies.

The first intuition could be to define a rule analogous to myopic best response that says: 'If the bidder gets no new information, she sticks with the same strategy. If she gets new information, she plays her best response.'

This rule sounds straight forward but it is completely infeasible: A losing bidder who observes the winning bid and updates her strategy always increases her strategy to a higher level. However, winning bidders get no new information. Therefore, winning bidders stick to their strategies and strategies will never move down again. Consequently, this updating rule will lead over time both bidders to bid exactly their true value. This behavior would be the outcome of all updating schemes considered in this thesis: Myopic and Perfect Memory Best or quantal response³ all lead to bidding the true value in 1PSBAs, leaving no payoff to the bidders.

Clearly, this is not a realistic result.

Consequently, the only plausible assumption is that winning bidders decrease their strategy in order to try to get more payoff in the future. But by how much? There is no endogenous rule that tells us whether winning bidders discount each time 1%, 5% or 10%. Or do bidders decrease their strategies the more, the more often the strategy has consecutively won? Another possibility would be that winning bidders calculate a probability distribution for their opponents' strategies that takes into account the information that they have not won the auction. Again, like in the case of probability-weighted myopic best response, this rule would depend crucially on the assumption that the participants know the underlying value distribution and the number of participants.

These considerations show that in first-price sealed bid auctions bidders will have even less chance to find the NE than in first-price open bid auctions. In second-price auctions, however, also in the sealed bid case, bidding the true value is the dominant strategy and the best response dynamics converges to that equilibrium in dominant strategies.

Note that the differences between information structures of open and sealed bid auctions diminish as the number of participants in the auction increases: The only difference between these two formats lies in the information that the winning bidder gets. The relevant information for the remaining $(N - 1)$ bidders – namely the highest bid – is in both cases published. As $N \rightarrow \infty$ therefore the information structures converge. Of course, this is just a rough argument because it assumes that the losing bidders are completely uninfluenced by the other losing bidders's bids and it implicitly also assumes that as $N \rightarrow \infty$, in each auction another bidder wins. But the underlying intuition should be correct.

³I investigate QR in detail in the next chapter.

Chapter 6

Quantal Response in Open Bid Auctions

Best response play is an arguable model for companies with specialized analytic departments like banks or insurances: Agents who face decisions where particularly much money is at stake might devote a lot of effort to collecting data about competitors and evaluating historical data. However, best response is in general not an apt model for the investigation of auctions with inexperienced bidders and low payoffs like they are typically found in auction experiments. Moreover, if bidders face time constraints, even experienced bidders might not be able to find the best response in the time given to them.

In this chapter I investigate a one-parameter family of alternative models for the formation of bidders beliefs. As I will show, this family encloses the best response dynamics as an extreme case. Under quantal response dynamics players keep track of how much payoff each of their own strategies could have generated in past play. The payoff of each strategy determines the probability of choosing the strategy for play in the next round. Like under best response, play can be myopic where bidders consider only last round's payoffs or it can be with perfect memory where bidders consider payoff information from the whole game.

6.1 Myopic Quantal Response to Fixed Strategies in 1POBAs

6.1.1 Payoff-Linear Myopic Quantal Response

For a start, I assume that players have only a finite set¹ of strategies β_0^i and that they choose strategies for their play by observing how much payoff each of them has generated in the last round. Myopic quantal response (mQR) in its simplest version is defined in the following way:

- Each bidder starts with some initial bid
- Bidder i observes the outcome of the auction
- Denote the set of feasible strategies of player 0 that would have won last round's auction by $\mathcal{B}_0 = \{\beta_0^{\text{old}} \mid \beta_0^{\text{old}} v_0^{\text{old}} > \beta_1^{\text{old}} v_1^{\text{old}}, 0 < \beta_0^{\text{old}} < 1\}$
- Bidders select a strategy for the next round from \mathcal{B} with a probability that is linearly proportional² to that strategy's payoff in the last round's auction: The selection probability for a specific strategy β is given as

$$\frac{PO(\beta)}{\int_0^1 d\tilde{\beta} PO(\tilde{\beta})} \tag{6.1}$$

¹e.g., one could think of the 100 strategies $\beta = 0.00, 0.01, \dots, 0.99$

²Below, I will generalize this assumption.

- If $\mathcal{B}_0 = \emptyset$, the bidder keeps her current strategy unchanged in the next round

Mathematically I can capture mQR in the following way:

The payoff of bidder 0 *conditional on that she could have won the auction with some feasible strategy, i.e., conditional on $\mathcal{B}_0 \neq \emptyset$* can be calculated according to the Figure 6.1: If there is




Figure 6.1: Payoff of Strategies

some feasible strategy β_0 that would have won the auction (and therefore would have generated positive payoff) that strategy β_0 satisfies $\beta_0 v_0 > \beta_1 v_1$ and therefore $\beta_0 > \frac{\beta_1 v_1}{v_0}$. The payoff at $\beta_0 = \frac{\beta_1 v_1}{v_0} + \epsilon$ is given as $v_0 - b_0 = v_0(1 - \beta_0) = v_0(1 - \frac{\beta_1 v_1}{v_0})$.

The selection density for some strategy $\beta_0 > \frac{\beta_1 v_1}{v_0}$ conditioned on $\mathcal{B}_0 \neq \emptyset$, i.e. on $v_0 > \beta_1 v_1$, is given by

$$f(\beta_0, \beta_1, v_0, v_1) = \frac{PO_0(\beta_0, \beta_1, v_0, v_1)}{\int_{\frac{\beta_1 v_1}{v_0}}^1 d\tilde{\beta}_0 PO_0(\tilde{\beta}_0, \beta_1, v_0, v_1)} = \frac{(1 - \beta_0)v_0}{\int_{\frac{\beta_1 v_1}{v_0}}^1 d\tilde{\beta}_0 (1 - \tilde{\beta}_0)v_0} \quad (6.2)$$

The expected value of β_0 , denoted by $\bar{\beta}_0^{\text{mQR}}$, still conditioned on $\mathcal{B}_0 \neq \emptyset$ is given by

$$\bar{\beta}_0^{\text{mQR}}(\beta_1, v_0, v_1) = \int_{\frac{\beta_1 v_1}{v_0}}^1 d\tilde{\beta}_0 f(\tilde{\beta}_0, \beta_1, v_0, v_1) \tilde{\beta}_0 = \frac{\int_{\frac{\beta_1 v_1}{v_0}}^1 d\tilde{\beta}_0 \tilde{\beta}_0 (1 - \tilde{\beta}_0)v_0}{\int_{\frac{\beta_1 v_1}{v_0}}^1 d\tilde{\beta}_0 (1 - \tilde{\beta}_0)v_0} = \quad (6.3)$$

$$\frac{\frac{\beta_0^2}{2} - \frac{\beta_0^3}{3} \Big|_{\frac{\beta_1 v_1}{v_0}}^1}{\beta_0 - \frac{\beta_0^2}{2} \Big|_{\frac{\beta_1 v_1}{v_0}}^1} = \frac{v_0 + 2\beta_1 v_1}{3v_0} \quad (6.4)$$

Next I calculate for illustrative purposes the average bid for all possible values v_0 *conditioned on $\mathcal{B}_0 \neq \emptyset$* .

A strategy $\beta_0 \in \mathcal{B}_0$ satisfies $\beta_0 v_0 > \beta_1 v_1$. Since $v_0 > \beta_0 v_0$ for $\beta_0 \in (0, 1)$, \mathcal{B}_0 can only be non empty if $v_0 > \beta_1 v_1$. Therefore I integrate v_0 not from 0 to 1 but only from $\beta_1 v_1$ to 1. Simultaneously I divide by the length of the integration interval, $(1 - \beta_1 v_1)$, to obtain the

expected value of β_0 :

$$\bar{\beta}_0^{\text{mQR}}(\beta_1, v_1) = \frac{1}{1 - \beta_1 v_1} \int_{\beta_1 v_1}^1 dv_0 \frac{v_0 + 2\beta_1 v_1}{3v_0} = \frac{\beta_1 v_1 + 2\beta_1 v_1 \ln(\beta_1 v_1) - 1}{3(\beta_1 v_1 - 1)} \quad (6.5)$$

To obtain $\bar{\beta}_0^{\text{mQR}}(\beta_1)$ I can *not* simply integrate $\bar{\beta}_0^{\text{mQR}}(\beta_1, v_1)$ over v_1 from 0 to 1. The reason is that taking the average to obtain the expected value of β_0 - as I did when I divided by $1 - \beta_1 v_1$ - is an operation that may only be performed in the end of a calculation.³

Therefore, I have to go back to equation (6.4) and calculate

$$\bar{\beta}_0^{\text{mQR}}(\beta_1) = \frac{1}{1 - \frac{\beta_1}{2}} \int_0^1 dv_1 \int_{\beta_1 v_1}^1 dv_0 \frac{v_0 + 2\beta_1 v_1}{3v_0} = \frac{2}{3} \frac{1 - \beta_1 \ln \beta_1}{2 - \beta_1} \quad (6.6)$$

The integration borders in the integration over v_0 are justified again since I condition on the event that at least some strategies in player 0's strategy set could have won against player 1's bid of $\beta_1 v_1$.

The prefactor $\frac{1}{1 - \frac{\beta_1}{2}}$ is the equivalent of the prefactor $\frac{1}{1 - \beta_1 v_1}$ that I used in the derivation of $\bar{\beta}_0^{\text{mQR}}(\beta_1, v_1)$. Here I additionally I used $v_1 \sim U(0, 1)$ which gives an average value of $\frac{1}{2}$. Figure 6.2 shows the average strategy of bidder 0 under linear myopic quantal response play in response to a fixed strategy of bidder 1. Note that the myopic quantal response against the NE is above the NE: $\bar{\beta}_0^{\text{mQR}}(\frac{1}{2}) \sim 0.6$. So far I have considered only the events where bidder 0




Figure 6.2: Myopic Quantal Response Play against fixed β_1

has at least one feasible strategy in \mathcal{B}_0 that could have won the auction. However, the above calculated $\bar{\beta}_0^{\text{mQR}}(\beta_1)$ is the same if bidders keep using the same strategy as long as they get no new information. The result is even the same if we allow bidders to select new strategies each round according to the payoff-proportional distribution of β_0 s they calculated in the last round where their value satisfied $v_0 > \beta_1 v_1$.

³For a more intuitive example consider the computation of the expected value of the squares of three numbers, e.g., 2, 3 and 4. It is given by $\frac{1}{3}(2^2 + 3^2 + 4^2) = 9\frac{2}{3}$. This is different from first taking the average $\frac{1}{3}(2 + 3 + 4) = 3$ and then squaring it to a result of 9.

The trick is to see that the probability that a strategy stays another round active is independent of the strategy value:

The active strategy β_0 is left unchanged if $v_0 \leq \beta_1 v_1$ because then no new payoff distribution on the strategy space $(0, 1)$ is generated. This condition is independent of the currently active β_0 . So, the strategies that stay another round on average have exactly the same distribution as the strategies that are selected in rounds where new information is generated.

Therefore the above calculated $\bar{\beta}_0^{mQR}(\beta_1)$ gives the expected value of β_0 under myopic quantal response for a wide range of possible updating rules if the last auction generated no payoff-distribution on the strategy space $(0, 1)$. This consideration gives rise to the following

Theorem 6.1 *Consider a two bidder first-price open bid auction with linear bidding strategies and $v_i \sim U(0, 1)$. Assume that bidder 1 plays a fixed strategy β_1 and bidder 0 follows a myopic payoff-linear quantal response updating rule:*

- if one or more strategies in her set could have generated positive payoff at time t , she selects one among them with probability proportional to their payoff
- if no strategy could have generated positive payoff at time t , she either
 - keeps the strategy she used at time t or
 - selects another strategy with probability proportional to the payoff the strategy would have obtained in the last round

Then, in the long term, her average strategy is given by

$$\bar{\beta}_0^{mQR}(\beta_1) = \frac{2}{3} \frac{1 - \beta_1 \ln \beta_1}{2 - \beta_1} \quad (6.7)$$

6.1.2 Payoff-Exponential Myopic Quantal Response

In the last section I assumed that players play a strategy in the next round with a probability that is directly proportional to the payoff that strategy generated in the last round. In the following I generalize this rule. I investigate what happens if player select a strategy β_i with a probability that is proportional to $[PO(\beta_i)]^\xi$ where ξ is a parameter that can take values between 0 and ∞ . For $\xi \rightarrow \infty$, myopic quantal response converges to myopic best response since then the strategy that generated the highest payoff in the last round is chosen in the next round with probability 1. For $\xi = 0$ we obtain essentially random play. $\xi = 1$ corresponds to payoff-linear QR considered above.

The probability that β_0 is chosen for next rounds play if the parameters of the last round were β_1, v_0 resp v_1 , is given by

$$f(\beta_0, \beta_1, v_0, v_1, \xi) = \frac{PO_0^\xi(\beta_0, \beta_1, v_0, v_1)}{\int_{\frac{\beta_1 v_1}{v_0}}^1 d\tilde{\beta}_0 PO_0^\xi(\tilde{\beta}_0, \beta_1, v_0, v_1)} = \frac{[(1 - \beta_0)v_0]^\xi}{\int_{\frac{\beta_1 v_1}{v_0}}^1 d\tilde{\beta}_0 [(1 - \tilde{\beta}_0)v_0]^\xi} \quad (6.8)$$

Therefore, the average strategy of player 0 conditional on v_0 and v_1 is given as

$$\begin{aligned} \bar{\beta}_0^{mQR}(\beta_1, v_0, v_1, \xi) &= \int_{\frac{\beta_1 v_1}{v_0}}^1 d\tilde{\beta}_0 \tilde{\beta}_0 f(\tilde{\beta}_0, \beta_1, v_0, v_1) = \frac{\int_{\frac{\beta_1 v_1}{v_0}}^1 d\tilde{\beta}_0 \tilde{\beta}_0 [(1 - \tilde{\beta}_0)v_0]^\xi}{\int_{\frac{\beta_1 v_1}{v_0}}^1 d\tilde{\beta}_0 [(1 - \tilde{\beta}_0)v_0]^\xi} = \\ &= \frac{(v_0 - \beta_1 v_1)^{\xi+1} (v_0 + \beta_1 v_1 (\xi+1))}{v_0^2 (2+3\xi+\xi^2)} = \frac{v_0 + \beta_1 v_1 (\xi + 1)}{v_0 (\xi + 2)}. \end{aligned} \quad (6.9)$$

By setting $\xi = 1$ this reduces to the pertinent equation (6.4) of the last subsection.

To calculate $\bar{\beta}_0^{mQR}(\beta_1, \xi)$ I proceed exactly as in the last subsection.

$$\bar{\beta}_0^{mQR}(\beta_1, \xi) = \frac{1}{1 - \frac{\beta_1}{2}} \int_0^1 dv_1 \int_{\beta_1 v_1}^1 dv_0 \frac{v_0 + \beta_1 v_1 (\xi + 1)}{v_0 (\xi + 2)} =$$

$$\frac{2}{2 - \beta_1} \int_0^1 dv_1 \frac{-1 + \beta_1 v_1 + \beta_1(\xi + 1)v_1 \ln(\beta_1 v_1)}{\xi + 2} = \dots = \frac{\beta_1(\xi - 1) - 2\beta_1(\xi + 1)\ln\beta_1 + 4}{2(2 - \beta_1)(\xi + 2)} \quad (6.10)$$

So I arrive at the following

Theorem 6.2 *Consider a two bidder first-price open bid auction with the same assumptions as in Theorem 6.1. The only difference is that the selection of new strategies is with a probability proportional to PO^ξ where the exponential factor ξ is a measure for the importance that bidders give to the payoff differences they obtain by using different strategies. Then, in the long term, the average strategy of player 0 is given as*

$$\bar{\beta}_0^{mQR}(\beta_1, \xi) = \frac{\beta_1(\xi - 1) - 2\beta_1(\xi + 1)\ln\beta_1 + 4}{2(2 - \beta_1)(\xi + 2)} \quad (6.11)$$

For $\xi \rightarrow \infty$, $\bar{\beta}_0^{mQR}(\beta_1, \xi)$ converges to myopic best response play.

Note that substitution of $\xi = 1$ rederives the prediction of Theorem 6.6 for $\bar{\beta}_0^{mQR}(\beta_1)$. For $\xi \rightarrow \infty$ I use that $\xi - 1 \sim \xi + 1 \sim \xi + 2$ and therefore can cancel the ξ terms from the fraction. This yields the expression

$$\bar{\beta}_0^{mQR}(\beta_1, \xi = \infty) = \frac{\beta_1(1 - 2\ln\beta_1)}{2(2 - \beta_1)} \quad (6.12)$$

This is exactly the expression for the average myopic best response strategy against a fixed strategy from Theorem 4.1. This is no coincidence as noted above since $\xi \rightarrow \infty$ means that we always take for sure the strategy with the best payoff. So the parameter ξ can be interpreted as an indicator of the importance that players give to the payoff.

6.2 Perfect Memory Quantal Response to Fixed Strategies in 1POBAs

6.2.1 Payoff-Linear Perfect Memory Quantal Response

Under perfect memory quantal response, like with mQR, bidders choose the strategies for their play by observing how well their individual strategies have behaved in the past. However, they base their selection on the accumulated payoff the strategy generated in the entire game as opposed to looking only at the last round.

Under the assumption of $v_0, v_1 \sim U(0, 1)$, the expected payoff of any strategy β_0 against a strategy β_1 is given by equation 4.53:

$$EPO_0(\beta_0, \beta_1) = (1 - \beta_0) \left\{ \Theta(\beta_0 - \beta_1) \left(\frac{1}{2} - \frac{\beta_1^2}{6\beta_0^2} - \frac{\beta_0}{3\beta_1} \right) + \frac{\beta_0}{3\beta_1} \right\} \quad (6.13)$$

If the game has been going on for long enough, the cumulated payoff of that strategy since the start of the game converges to the expected payoff.

With payoff-proportional strategy selection, the expected value of β_0 is the weighted mean of all possible β_0 s. The weights are proportional to the payoff that the pertinent strategies have generated in the past. The expected value $\bar{\beta}_0^{pQR}(\beta_1)$ is therefore given by:

$$\bar{\beta}_0^{pQR}(\beta_1) = \int_0^1 d\tilde{\beta}_0 f(\tilde{\beta}_0, \beta_1) \tilde{\beta}_0 = \int_0^1 d\tilde{\beta}_0 \frac{EPO_0(\tilde{\beta}_0, \beta_1)}{\int_0^1 d\tilde{\beta}_0 EPO_0(\tilde{\beta}_0, \beta_1)} \tilde{\beta}_0 = \frac{\int_0^1 d\tilde{\beta}_0 EPO_0(\tilde{\beta}_0, \beta_1) \tilde{\beta}_0}{\int_0^1 d\tilde{\beta}_0 EPO_0(\tilde{\beta}_0, \beta_1)} \quad (6.14)$$

The evaluation of the integrals is straight forward:
The numeraire is given by:

$$\begin{aligned} \int_0^1 d\tilde{\beta}_0 EPO_0(\tilde{\beta}_0, \beta_1) \tilde{\beta}_0 &= \int_0^{\beta_1} d\tilde{\beta}_0 (1 - \tilde{\beta}_0) \frac{\tilde{\beta}_0^2}{3\beta_1} + \int_{\beta_1}^1 d\tilde{\beta}_0 (\tilde{\beta}_0 - \tilde{\beta}_0^2) \left(\frac{1}{2} - \frac{\beta_1^2}{6\tilde{\beta}_0^2}\right) = \\ &= \dots = \frac{1}{12} + \frac{\beta_1^2}{36} + \frac{\beta_1^2 \ln \beta_1}{6} - \frac{\beta_1^3}{12}. \end{aligned} \quad (6.15)$$

The denominator is given by

$$\begin{aligned} \int_0^1 d\tilde{\beta}_0 EPO_0(\tilde{\beta}_0, \beta_1) &= \int_0^{\beta_1} d\tilde{\beta}_0 (1 - \tilde{\beta}_0) \frac{\tilde{\beta}_0}{3\beta_1} + \int_{\beta_1}^1 d\tilde{\beta}_0 (1 - \tilde{\beta}_0) \left(\frac{1}{2} - \frac{\beta_1^2}{6\tilde{\beta}_0^2}\right) = \\ &= \dots = \frac{1}{4} - \frac{\beta_1}{2} + \frac{11\beta_1^2}{36} - \frac{\beta_1^2}{6} \ln \beta_1 \end{aligned} \quad (6.16)$$

Simplification yields the following

Theorem 6.3 Consider a two bidder first-price open bid auction with linear strategies and $v_i \sim U(0, 1)$. Assume that bidder 1 plays a fixed strategy and bidder 0 follows a perfect memory quantal response rule for strategy updating: Strategies are selected with probability proportional to the cumulated payoff the strategy would have generated in all auctions up to time t . Then, her average strategy is given by

$$\bar{\beta}_0^{pQR}(\beta_1) = \frac{3 + \beta_1^2 - 3\beta_1^3 + 6\beta_1^2 \ln \beta_1}{9 - 18\beta_1 + 11\beta_1^2 - 6\beta_1^2 \ln \beta_1} \quad (6.17)$$

In particular, $\bar{\beta}_0^{pQR} < \frac{1}{2} \forall \beta_1 \in U(0, 1)$.

Average perfect memory quantal response against a fixed strategy of bidder 1 is plotted in Figure 6.3. Note from Figs. 6.2 and 6.3 that linear quantal response leads under both, myopic and perfect memory play, bidders to use the (average) strategy $\beta_i = \frac{1}{3}$ if the opponent bids nothing, $\beta_{-i} = 0$. This shows that bidders overbid significantly for low strategies when compared to their best response. For higher strategies, myopic quantal response leads to overbidding whereas perfect memory quantal response stays below the NE for all strategies of the opponent.

6.2.2 Payoff-Exponential Perfect Memory Quantal Response

Like for the myopic case I tried to investigate perfect memory payoff-exponential quantal response play where the probability that a certain strategy is played is proportional to PO_0^ξ . By proceeding exactly like for the $\xi = 1$ -case we can write

$$\bar{\beta}_0^{pQR}(\beta_1, \xi) = \frac{\int_0^1 d\tilde{\beta}_0 \frac{EPO_0^\xi(\tilde{\beta}_0, \beta_1) \tilde{\beta}_0}{\int_0^1 d\tilde{\beta}_0 EPO_0^\xi(\tilde{\beta}_0, \beta_1)}} \quad (6.18)$$

where the expected payoff is given by

$$EPO_0(\beta_0, \beta_1) = \begin{cases} (1 - \beta_0) \frac{\beta_0}{3\beta_1} & \text{for } \beta_0 < \beta_1 \\ (1 - \beta_0) \left(\frac{1}{2} - \frac{\beta_1^2}{6\beta_0^2}\right) & \text{for } \beta_0 > \beta_1 \end{cases} \quad (6.19)$$

The problem is however that I could not derive closed-form expressions for these integrals. In the parts of the integrals that I could solve there occur hypergeometric functions.

Of course it would have been nice to have a closed-form solution for each value of ξ . The Auction Simulator gives the possibility to investigate payoff-exponential perfect memory quantal response. I will show with the Auction Simulator in a later chapter that analogous to myopic quantal response, perfect memory payoff-exponential quantal response converges to best response play for $\xi \rightarrow \infty$.




Figure 6.3: Perfect Memory Quantal Response Play against Fixed β_1

6.3 Quantal Response in 2POBAs

Under quantal response strategy updating, the equilibrium strategy of bidding the true value in 2PAs is in general not found. This is intuitively clear by looking back at Figure 4.8. The average play depends on the size of the strategy set:

In 1PAs it is intuitively clear that bidders have to confine themselves to the strategy set $\beta_i \in (0, 1)$. Strategies have to be bigger than 0 to be able to obtain the object and they have to be smaller than 1 to avoid losses. In 2PAs, all strategies between 0 and ∞ generate positive expected payoff.⁴ Therefore, the average play against a fixed strategy depends on the cutoff-level that human bidders assume for their strategy set. In conclusion, quantal response in 2PAs does not give rise to a unique analytical solution.

Note however that the reasoning, why in 2PAs it is optimal to bid exactly the own valuation is quite simple. Therefore, even if we assume that human bidders use quantal response strategy updating in 1PAs, it might be a reasonable assumption that most bidders in 2PAs do discover their NE bidding functions or at least play very close to it.

Experimental evidence shows that bidders in repeated second price auctions move very closely to the NE. However, some experiments show small but significant deviations from NE play (see, e.g., Kagel and Roth (1995), Garrat and Wooders (2004)). This can be explained by a perfect memory quantal response model with very high ξ . Intuitively, the interpretation of this model is that bidders in principle can determine their best strategy, but they can make small errors in this determination. They play good strategies more often than bad ones. Now, take a look at Figure 4.8. If ξ is very high, this means that bidders' selection criteria are very strict. Strategies with a payoff below some threshold level are virtually never played. Therefore, the integration to determine the average strategy can be approximated by integrating only over a certain region around the NE of $\beta = 1$.

$$\bar{\beta} = \int_0^{\infty} d\beta \beta PO^{\xi}(\beta) \sim \int_{1-\epsilon}^{1+\epsilon} d\beta \beta PO^{\xi}(\beta) \quad (6.20)$$

where ϵ is such that strategies outside the interval $1 \pm \epsilon$ are virtually never played. A look at Figure 4.8 shows that the payoff function is flatter to the right than to the left. This means that – no matter, which strategy the opponent plays – strategies above the NE of $\beta_i = 1$ are more

⁴See the section about perfect memory best response in 2PAs.

often played than strategies below it. Therefore a perfect memory payoff-exponential quantal response model with high ξ could explain the experimentally observed overbidding in 2PAs.

Chapter 7

Best Response in Pure Seller Auctions

To investigate auctions where only sellers and no bidders participate, I will in principle make the same analyses as in the case of pure bidder auctions. Again I consider a formalized setting where the sellers' private values are in each round drawn anew from $U(0, 1)$.

An example for such an auction in a repeated setting is e.g. procurement. In a specific industry sector, big suppliers repeatedly compete for contracts. They want to sell their products or services to a specific firm. The firm takes the role of the bidder. In this chapter I assume that the bidder posts no reservation price – so her willingness to pay is assumed to be infinitely high. Though this might not be very realistic, it helps to focus only on the interaction of the sellers. Extensions that allow for a bidder reservation price can be calculated exactly in the same way as in the case of pure bidder auctions that were considered in section 4.2.3. Pure seller auctions can – like in the pure bidder case – be organized as first- second- or higher-price auctions. The question is here again, which of these auction formats converge under which assumptions.

As in the case of pure bidder auctions I allow only for linear asking strategies. I stick to the convention that a seller with value v and strategy σ asks $\frac{v}{\sigma}$. This convention guarantees that the set of feasible asking strategies is confined to the interval $(0, 1]$.

Note however that the seller perspective is not completely symmetric to the bidder perspective: Assuming a value distribution $U(0, 1)$, a bidder's maximal possible payoff in a pure bidder auction is 1¹. A seller in a pure seller auction can obtain an infinitely high profit if the bidder has no reservation price, i.e., if the bidders willingness to pay is infinitely high. Due to this asymmetry between the bidder- and the seller-format, it is a priori not clear, whether the convergence behavior of the auction formats will be similar.

7.1 Payoff Function for Sellers

The probability that seller 0 wins by asking v_0/σ_0 against a second seller who asks v_1/σ_1 , is given by

$$p_0^{\text{win}} = \begin{cases} 1 & \text{if } \frac{v_0}{\sigma_0} < \frac{v_1}{\sigma_1} \Leftrightarrow \text{if } v_0\sigma_1 < v_1\sigma_0 \\ 0 & \text{if } \frac{v_0}{\sigma_0} > \frac{v_1}{\sigma_1} \Leftrightarrow \text{if } v_0\sigma_1 > v_1\sigma_0 . \end{cases} \quad (7.1)$$

In closed form, the winning probability for seller 0 can therefore be written as

$$p_0^{\text{win}} = \Theta(v_1\sigma_0 - v_0\sigma_1) \quad (7.2)$$

and the payoff is given as

$$PO_0(\sigma_0, \sigma_1, v_0, v_1) = v_0 \left(\frac{1}{\sigma_0} - 1 \right) \Theta(v_1\sigma_0 - v_0\sigma_1) \quad (7.3)$$

¹Assuming that the bidder's value is 1 and she wins the auction with a bid of $\epsilon \rightarrow 0$

because if seller 0 can sell the item that she values at v_0 for a price of v_0/σ_0 , her payoff is $v_0/\sigma_0 - v_0$. Figure 7.1 shows the payoff of seller 0 if she uses different strategies against a fixed combination of σ_1 , v_0 and v_1 .




Figure 7.1: Payoff of Ask Strategies

7.2 Myopic Best Response to Fixed σ_1

In this section I assume that seller 1 uses a fixed strategy σ_1 that is the same in all rounds. I want to assess seller 0's average myopic best response to σ_1 , i.e., I want to assess the long term average of her strategy if she uses myopic best response. Like in the pure bidder case I stick to the convention that seller 0 keeps last round's strategy if no feasible strategy of her could have won the auction. Seller 0's best response strategy against last rounds' play is therefore given as

$$\sigma_0^{\text{new}}(\sigma_0^{\text{old}}, \sigma_1, v_0, v_1) \stackrel{!}{=} \sigma_0^{\text{br}}(\sigma_0^{\text{old}}, \sigma_1, v_0, v_1) = \begin{cases} \lim_{\epsilon \rightarrow 0} \frac{v_0\sigma_1}{v_1} + \epsilon = \frac{v_0\sigma_1}{v_1} & \text{if } \frac{v_0\sigma_1}{v_1} < 1 \\ \sigma_0^{\text{old}} & \text{if } \frac{v_0\sigma_1}{v_1} > 1 \end{cases} \quad (7.4)$$

where σ_0^{new} denotes the respondent's strategy in round $t+1$ whereas the values of σ_0^{old} , v_0 and v_1 are taken at time t . In short term notation I can write this as

$$\sigma_0^{\text{new}}(\sigma_0^{\text{old}}, \sigma_1, v_0, v_1) = \frac{v_0\sigma_1}{v_1} \Theta(v_1 - v_0\sigma_1) + \sigma_0^{\text{old}} \Theta(v_0\sigma_1 - v_1) \quad (7.5)$$

Integration over v_1 yields the expected value of the myopic best response strategy in dependence of the opponent's strategy and the respondent's value v_0 ,

$$\bar{\sigma}_0^{\text{new}}(\sigma_0^{\text{old}}, \sigma_1, v_0) = v_0\sigma_1(\sigma_0^{\text{old}} - \ln(v_0\sigma_1)) \quad (7.6)$$

and finally integration over v_0 and setting $\bar{\sigma}_0^{\text{new}} = \bar{\sigma}_0^{\text{old}} \stackrel{!}{=} \bar{\sigma}_0^{\text{mBR}}$ yields the following

Theorem 7.1 Consider a two seller first-price open bid auction with linear ask strategies and $v_i \sim U(0, 1)$. Assume that seller 1 plays a fixed strategy σ_1 and seller 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 $\sigma_0 \in (0, 1]$, she continues to use the previous round's strategy. Then, in the long term, her average strategy satisfies

$$\bar{\sigma}_0^{\text{mBR}}(\sigma_1) = \frac{\sigma_1(1 - 2\ln\sigma_1)}{2(2 - \sigma_1)} \quad (7.7)$$

In particular, $\bar{\sigma}_0^{mBR}(\sigma_1) < \frac{1}{2} \forall \sigma_1 \in [0, 1]$.

Note that the result is the same as in Theorem 4.1 only that the β s are replaced by σ s. Therefore, the plot of the myopic best response function for seller auctions would look like in Figure 4.2. Consequently, also the behavior of mutually adapting sellers is completely analogous to the case of mutually adapting bidders. In particular, they would end up by asking infinitely much in the case of mutual adaptation. The underlying reason is that myopic best response in both, the seller and the bidder case, only depends on the place where the payoff function is discontinuous but not on the specific shape of the payoff function. The updating equations – given in (4.6) and (7.4) – have the same structure for both, the pure seller and the pure bidder case. Therefore, myopic best response yields the same equation structure in Theorems 4.1 and 7.1.

7.3 Perfect Memory Best Response to Fixed σ_1

Assume, seller 1 plays a fixed strategy σ_1 , then the payoff for seller 0 is given by

$$PO_0(\sigma_0, \sigma_1, v_0, v_1) = \begin{cases} \frac{v_0}{\sigma_0} - v_0 & \text{if } \frac{v_0}{\sigma_0} < \frac{v_1}{\sigma_1} \\ 0 & \text{if } \frac{v_0}{\sigma_0} > \frac{v_1}{\sigma_1} . \end{cases} \quad (7.8)$$

In closed form, the payoff can be written as

$$PO_0 = v_0 \left(\frac{1}{\sigma_0} - 1 \right) \Theta(\sigma_0 v_1 - \sigma_1 v_0) . \quad (7.9)$$

Like in the previous chapter, the best response function is found by integration over v_0 and v_1 and subsequent solving of the FOC.

$$\begin{aligned} EPO_0(\sigma_0, \sigma_1, v_1) &= \int_0^1 dv_0 v_0 \left(\frac{1}{\sigma_0} - 1 \right) \Theta\left(\frac{\sigma_0 v_1}{\sigma_1} - v_0\right) = \int_0^{\min(1, \frac{\sigma_0 v_1}{\sigma_1})} dv_0 v_0 \left(\frac{1}{\sigma_0} - 1 \right) = \\ &= \int_0^{1 - \Theta(1 - \frac{\sigma_0 v_1}{\sigma_1})(1 - \frac{\sigma_0 v_1}{\sigma_1})} dv_0 v_0 \left(\frac{1}{\sigma_0} - 1 \right) = \left(\frac{1}{\sigma_0} - 1 \right) \frac{1}{2} \left\{ 1 - \Theta\left(1 - \frac{\sigma_0 v_1}{\sigma_1}\right) \left(1 - \frac{\sigma_0 v_1}{\sigma_1}\right) \right\}^2 = \\ &= \begin{cases} \frac{1}{2} \left(\frac{1}{\sigma_0} - 1 \right) & \text{if } \sigma_1 < \sigma_0 v_1 \\ \frac{1}{2} \left(\frac{1}{\sigma_0} - 1 \right) \frac{\sigma_0^2 v_1^2}{\sigma_1^2} & \text{if } \sigma_1 > \sigma_0 v_1 . \end{cases} \end{aligned} \quad (7.10)$$

In closed form this can be rewritten as

$$EPO_0(\sigma_0, \sigma_1, v_1) = \frac{1}{2} \left(\frac{1}{\sigma_0} - 1 \right) \left\{ \Theta(v_1 - \frac{\sigma_1}{\sigma_0}) + \frac{\sigma_0^2 v_1^2}{\sigma_1^2} \Theta\left(\frac{\sigma_1}{\sigma_0} - v_1\right) \right\} . \quad (7.11)$$

Integration over v_1 yields

$$EPO_0(\sigma_0, \sigma_1) = \underbrace{\int_0^1 dv_1 \frac{1}{2} \left(\frac{1}{\sigma_0} - 1 \right) \Theta(v_1 - \frac{\sigma_1}{\sigma_0})}_{(I)} + \underbrace{\int_0^1 dv_1 \frac{1}{2} \left(\frac{1}{\sigma_0} - 1 \right) \frac{\sigma_0^2 v_1^2}{\sigma_1^2} \Theta\left(\frac{\sigma_1}{\sigma_0} - v_1\right)}_{(II)} . \quad (7.12)$$

The first integral is given as

$$(I) = \frac{1}{2} \left(\frac{1}{\sigma_0} - 1 \right) \Theta\left(1 - \frac{\sigma_1}{\sigma_0}\right) \int_{\frac{\sigma_1}{\sigma_0}}^1 dv_1 = \frac{1}{2} \left(\frac{1}{\sigma_0} - 1 \right) \left(1 - \frac{\sigma_1}{\sigma_0}\right) \Theta(\sigma_0 - \sigma_1) \quad (7.13)$$

where the Θ -function in front of the integral guarantees that the value of the integral is 0 if $\frac{\sigma_1}{\sigma_0} > 1$ since then all $v_1 \in (0, 1)$ give no contribution due to $\Theta(v_1 - \frac{\sigma_1}{\sigma_0}) = 0$. The second integral is given as

$$\begin{aligned} (II) &= \frac{1}{2} \left(\frac{1}{\sigma_0} - 1 \right) \int_0^{1 - \Theta(\sigma_0 - \sigma_1)(1 - \frac{\sigma_1}{\sigma_0})} dv_1 \frac{\sigma_0^2 v_1^2}{\sigma_1^2} = \\ &= \frac{1}{2} \left(\frac{1}{\sigma_0} - 1 \right) \frac{\sigma_0^2}{\sigma_1^2} \frac{v_1^3}{3} \Big|_0^{1 - \Theta(\sigma_0 - \sigma_1)(1 - \frac{\sigma_1}{\sigma_0})} = \\ &= \frac{1}{6} \left(\frac{1}{\sigma_0} - 1 \right) \frac{\sigma_0^2}{\sigma_1^2} \left\{ 1 - \Theta(\sigma_0 - \sigma_1)(1 - \frac{\sigma_1}{\sigma_0}) \right\}^3. \end{aligned} \quad (7.14)$$

Therefore the expected payoff is given by

$$EPO_0(\sigma_0, \sigma_1) = \frac{1}{2} \left(\frac{1}{\sigma_0} - 1 \right) (1 - \frac{\sigma_1}{\sigma_0}) \Theta(\sigma_0 - \sigma_1) + \frac{1}{6} \left(\frac{1}{\sigma_0} - 1 \right) \frac{\sigma_0^2}{\sigma_1^2} \left\{ 1 - \Theta(\sigma_0 - \sigma_1)(1 - \frac{\sigma_1}{\sigma_0}) \right\}^3. \quad (7.15)$$

To solve the first order condition, $\partial/\partial\sigma_0 EPO_0(\sigma_0, \sigma_1) \stackrel{!}{=} 0$, I have to distinguish two cases². For $\sigma_0 < \sigma_1$, the expected payoff is given as

$$EPO_0(\sigma_0, \sigma_1) = \frac{1}{6} \left(\frac{1}{\sigma_0} - 1 \right) \left(\frac{\sigma_0^2}{\sigma_1^2} \right) = \frac{1}{6\sigma_1^2} (\sigma_0 - \sigma_0^2) \quad (7.16)$$

The FOC is given by

$$\frac{\partial}{\partial\sigma_0} EPO_0(\sigma_0, \sigma_1) = \frac{1}{6\sigma_1^2} (1 - 2\sigma_0) \stackrel{!}{=} 0 \quad (7.17)$$

with the solution

$$\sigma_0^{pBR}(\sigma_1) = \frac{1}{2} \quad \forall \sigma_1 \geq \frac{1}{2} \quad (7.18)$$

where the condition $\sigma_1 \geq \frac{1}{2}$ is obtained by plugging the solution back into the condition $\sigma_0 < \sigma_1$. Similarly, if $\sigma_0 > \sigma_1$ the expected payoff is given as

$$\begin{aligned} EPO_0(\sigma_0, \sigma_1) &= \frac{1}{2} \left(\frac{1}{\sigma_0} - 1 \right) (1 - \frac{\sigma_1}{\sigma_0}) + \frac{1}{6} \left(\frac{1}{\sigma_0} - 1 \right) \frac{\sigma_0^2}{\sigma_1^2} \frac{\sigma_1^3}{\sigma_0^3} = \dots = \\ &= -\frac{1}{2} + \frac{1}{\sigma_0} \left(\frac{1}{2} + \frac{\sigma_1}{2} - \frac{\sigma_1}{6} \right) - \frac{1}{\sigma_0^2} \frac{\sigma_1}{3} \end{aligned} \quad (7.19)$$

The FOC is given by

$$\frac{\partial}{\partial\sigma_0} EPO_0(\sigma_0, \sigma_1) = -\frac{1}{\sigma_0^2} \left(\frac{1}{2} + \frac{\sigma_1}{3} \right) + \frac{2\sigma_1}{3\sigma_0^3} \stackrel{!}{=} 0 \quad (7.20)$$

The solution to this equation is easily determined as

$$\sigma_0^{pBR}(\sigma_1) = \frac{4\sigma_1}{3 + 2\sigma_1} \quad \forall \sigma_1 \leq \frac{1}{2} \quad (7.21)$$

Therefore I arrive at the following

Theorem 7.2 *Consider a two seller first-price auction with linear ask strategies and $v_i \sim U(0, 1)$. Assume that seller 1 plays a fixed strategy σ_1 and seller 0 follows a perfect memory best response strategy updating rule. Then, in the long term, her strategy is given as*

$$\sigma_0^{pBR}(\sigma_1) = \frac{4\sigma_1}{3 + 2\sigma_1} \Theta\left(\frac{1}{2} - \sigma_1\right) + \frac{1}{2} \Theta\left(\sigma_1 - \frac{1}{2}\right). \quad (7.22)$$

²The symmetric case follows from continuity arguments. Alternatively I could also apply the same methodology as in the case of pure bidder auctions.

Note that the structure of this equation is the same as the structure of equation (4.67) in the bidder case: If seller 1 uses a strategy above 0.5, the best response of the opponent is 0.5 as well. If seller 1 uses a strategy below 0.5, then the best response of the opponent is also smaller than 0.5 but above the strategy of seller 1. Therefore, the plot of the perfect memory best response of seller 0 to seller 1 looks qualitatively like in the bidder case which is shown in Figure 4.7. The only difference is that the left branch for strategies of seller 1 below 0.5 is higher than in the bidder case. The underlying reason is that the shape of the payoff-function and therefore also the shape of the expected payoff is different in the pure bidder and in the pure seller case. As a consequence of this, since the equations for pure bidder and pure seller perfect memory best response have qualitatively the same structure, also the discussion of the convergence properties of pure seller auctions under truncated fictitious play and perfect memory best response are like in the pure bidder case: Perfect memory best response slowly converges from below to the NE whereas truncated fictitious play involves in permanent strategy fluctuations around a mean below the NE.

Chapter 8

Best Response in Double Auctions

8.1 Double Auctions with Linear Bidding Strategies

In this chapter I consider a repeated double auction with one bidder and one seller. Private values for both, the bidder and the seller, are drawn in each round independently from a uniform random distribution, $v_i \sim U(0, 1)$. Let the bidder with private value v_0 bid $b_0(v_0)$ and let the seller with value v_1 ask $s_1(v_1)$. If $b_0(v_0) > s_1(v_1)$, the price is set at

$$p = \frac{1}{2}(b_0(v_0) + s_1(v_1)). \quad (8.1)$$

The bidder then obtains a payoff of

$$PO_0(v_0, b_0(v_0), s_1(v_1)) = v_0 - p \quad (8.2)$$

whereas the seller obtains a payoff of

$$PO_1(v_1, s_1(v_1), b_0(v_0)) = p - v_1 \quad (8.3)$$

If $b_0(v_0) < s_1(v_1)$, no trade takes place and both participants obtain zero payoff.

This setup corresponds to bargaining under incomplete information. It serves, e.g., as a model for union negotiations which may result in a compromise (a trade) or a strike (no trade takes place). It has first been investigated in Chatterjee and Samuelson (1983).

In the first section I consider the case of linear bidding strategies. Bidder 0 decides on her strategy $\beta_0 \in [0, 1]$ and offers for value v_0 the bid $b_0(v_0) = \beta_0 v_0$ where β_0 is the same for all v_0 . Analogously, the seller demands for value v_1 the ask v_1/σ_1 with $\sigma_1 \in [0, 1]$. This section shows that perfect memory best response with linear bidding strategies yields a dominant strategy pair ($\beta_0^*(\sigma_1) = \frac{2}{3}, \sigma_1^*(\beta_0) = \frac{3}{4}$). This means that no matter what the opponent does, it is always optimal to play the strategy $\frac{2}{3}$ for the bidder and $\frac{3}{4}$ for the seller. Note that this leaves exactly the same expected payoff for the bidder and the seller: For values v_0 resp. v_1 , the bidders payoff is

$$PO_0(v_0) = v_0 - \frac{2}{3}v_0 = \frac{v_0}{3} \quad (8.4)$$

and the sellers payoff is

$$PO_1(v_1) = \frac{4}{3}v_1 - v_1 = \frac{v_1}{3}. \quad (8.5)$$

Since values come from the same distribution, their expected payoff is equal.

In conclusion, a double auction under perfect memory best response comes to a remarkably robust equilibrium if the participants are confined to linear bidding strategies. Moreover, the expected gains from trade are split in this equilibrium equally between the seller and the bidder.

Though these insights are quite interesting, the model in this section suffers a fundamental flaw: The considered linear strategy space does not entail the NE of a double auction. Chatterjee and Samuelson (1983) have shown that double auctions with uniform value distributions and one participating bidder and seller each have affine NE bidding functions $b_0(v_0) = \alpha_0 + \beta_0 v_0$ and $s_1(v_1) = \alpha_1 + \frac{v_1}{\sigma_1}$ respectively. Therefore, the restriction to linear strategies is not innocuous anymore – even not near the NE. In the next section I will outline how to investigate this more complex case of affine bidding functions with four free parameters.

However, this section about linear strategies is not in vain: It gives additional opportunity to practice the mathematical methods. And – maybe more important – it gives additional data points to gauge the correctness of simulation results in the later chapters of this thesis.

8.1.1 Myopic Best Response

Bidder Response

The myopic best response function for bidder 0 is given as

$$\beta_0^{\text{new}}(\sigma_1, v_0, v_1, \beta_0^{\text{old}}) = \begin{cases} \frac{v_1}{\sigma_1 v_0} & \text{if } \frac{v_1}{\sigma_1 v_0} < 1 \\ \beta_0^{\text{old}} & \text{if } \frac{v_1}{\sigma_1 v_0} > 1 \end{cases} \quad (8.6)$$

which can be written in closed form by means of the Θ -function as

$$\beta_0^{\text{new}}(\sigma_1, v_0, v_1, \beta_0^{\text{old}}) = \frac{v_1}{\sigma_1 v_0} \Theta(\sigma_1 v_0 - v_1) + \beta_0^{\text{old}} \Theta(v_1 - \sigma_1 v_0). \quad (8.7)$$

Integration over v_1 yields the average strategy of bidder 0 in dependence of her own value v_0 ,

$$\begin{aligned} \bar{\beta}_0^{\text{new}}(\sigma_1, v_0, \beta_0^{\text{old}}) &= \int_0^1 dv_1 \beta_0^{\text{new}}(\sigma_1, v_0, v_1, \beta_0^{\text{old}}) = \frac{v_1^2}{2\sigma_1 v_0} \Big|_0^{\sigma_1 v_0} + \beta_0^{\text{old}} v_1 \Big|_{\sigma_1 v_0}^1 = \\ &= \frac{\sigma_1 v_0}{2} + \beta_0^{\text{old}} (1 - \sigma_1 v_0). \end{aligned} \quad (8.8)$$

Subsequent integration over v_0 yields

$$\begin{aligned} \bar{\beta}_0^{\text{new}}(\sigma_1, \beta_0^{\text{old}}) &= \int_0^1 dv_0 \left\{ \frac{\sigma_1 v_0}{2} + \beta_0^{\text{old}} (1 - \sigma_1 v_0) \right\} = \frac{\sigma_1 v_0^2}{4} + \beta_0^{\text{old}} v_0 - \beta_0^{\text{old}} \sigma_1 \frac{v_0^2}{2} \Big|_0^1 = \\ &= \frac{\sigma_1}{4} + \beta_0^{\text{old}} - \beta_0^{\text{old}} \frac{\sigma_1}{2} \end{aligned} \quad (8.9)$$

On average the relation $\bar{\beta}_0^{\text{new}} = \bar{\beta}_0^{\text{old}} \stackrel{!}{=} \bar{\beta}_0^{\text{mBR}}$ must hold and by solving for $\bar{\beta}_0^{\text{mBR}}$, the myopic best response of the bidder must satisfy the equation

$$\bar{\beta}_0^{\text{mBR}} = \frac{\sigma_1}{4} + \bar{\beta}_0^{\text{mBR}} - \bar{\beta}_0^{\text{mBR}} \frac{\sigma_1}{2} \quad (8.10)$$

with the unique solution

$$\bar{\beta}_0^{\text{mBR}} = \frac{1}{2}. \quad (8.11)$$

So I arrive at the following

Theorem 8.1 *Consider a repeated single-unit double auction with one bidder and one seller, linear strategies and $v_i \sim U(0, 1)$. If a trade occurs, the price is set in the middle of bid and ask. Then, in the long term, the average myopic best response function of the bidder is given by*

$$\bar{\beta}_0^{\text{mBR}}(\sigma_1) = \frac{1}{2} \quad \forall \sigma_1 \quad (8.12)$$

Note in particular that the average myopic best response function of the bidder does not depend on the strategy of the seller.

Seller Response

Analogously to the bidder case, the myopic best response function of the seller is given in a closed form as

$$\sigma_1^{\text{new}}(\beta_0, v_1, v_0, \sigma_1^{\text{old}}) = \frac{v_1}{\beta_0 v_0} \Theta(\beta_0 v_0 - v_1) + \sigma_1^{\text{old}} \Theta(v_1 - \beta_0 v_0) \quad (8.13)$$

Integration over v_1 yields

$$\bar{\sigma}_1^{\text{new}}(\beta_0, v_0, \sigma_1^{\text{old}}) = \frac{v_1^2}{2\beta_0 v_0} \Big|_0^{\beta_0 v_0} + \sigma_1^{\text{old}}(1 - \beta_0 v_0) = \frac{\beta_0 v_0}{2} + \sigma_1^{\text{old}}(1 - \beta_0 v_0) \quad (8.14)$$

By double integration over v_0 and v_1 and denoting $\bar{\sigma}_1^{\text{new}} = \bar{\sigma}_1^{\text{old}} \stackrel{!}{=} \bar{\sigma}_1^{\text{mBR}}$ I arrive at

$$\bar{\sigma}_1^{\text{mBR}} = \frac{\beta_0}{4} + \bar{\sigma}_1^{\text{mBR}} - \frac{\beta_0 \bar{\sigma}_1^{\text{mBR}}}{2} \quad (8.15)$$

with the solution

$$\bar{\sigma}_1^{\text{mBR}} = \frac{1}{2} \quad (8.16)$$

So I arrive at

Theorem 8.2 *Consider a repeated single-unit double auction with one bidder and one seller, linear strategies and $v_i \sim U(0, 1)$. Then, in the long term, the average myopic best response function of the seller is given by*

$$\bar{\sigma}_1^{\text{mBR}}(\beta_0) = \frac{1}{2} \quad \forall \beta_0. \quad (8.17)$$

Note in particular the average myopic best response function of the seller does not depend on the strategy of the bidder.

Note that $(\beta = 1/2, \sigma = 1/2)$ is *not* a dominant strategy pair. Theorems 8.1 and 8.2 describe the average strategy of the bidder and the seller, however under myopic best response, strategies constantly fluctuate around the mean and never settle down because the participants react to every new information.

8.1.2 Perfect Memory Best Response

Bidder Response

Bidder 0s payoff is given as

$$PO_0(\beta_0, \sigma_1, v_0, v_1) = \begin{cases} 0 & \text{for } \beta_0 v_0 < \frac{v_1}{\sigma_1} \\ v_0 - \frac{\beta_0 v_0 + \frac{v_1}{\sigma_1}}{2} & \text{for } \beta_0 v_0 > \frac{v_1}{\sigma_1}. \end{cases} \quad (8.18)$$

So the payoff can be written in a closed form as

$$PO_0(\beta_0, \sigma_1, v_0, v_1) = \left\{ v_0 - \frac{\beta_0 v_0 + \frac{v_1}{\sigma_1}}{2} \right\} \Theta(\beta_0 v_0 \sigma_1 - v_1). \quad (8.19)$$

Integration over v_1 yields

$$EPO_0(\beta_0, \sigma_1, v_0) = \int_0^{\beta_0 v_0 \sigma_1} dv_1 \left(v_0 - \frac{\beta_0 v_0}{2} - \frac{v_1}{2\sigma_1} \right) \quad (8.20)$$

since $\beta_0 v_0 \sigma_1 < 1$. Evaluation of the integral yields

$$\begin{aligned} EPO_0(\beta_0, \sigma_1, v_0) &= v_0 v_1 - \frac{\beta_0 v_0 v_1}{2} - \frac{v_1^2}{4\sigma_1} \Big|_0^{\beta_0 v_0 \sigma_1} = \\ &= v_0^2 \beta_0 \sigma_1 - \frac{\beta_0^2 v_0^2 \sigma_1}{2} - \frac{\beta_0^2 v_0^2 \sigma_1^2}{4\sigma_1} = v_0^2 \beta_0 \sigma_1 - \frac{3\beta_0^2 v_0^2 \sigma_1}{4}. \end{aligned} \quad (8.21)$$

Subsequent integration over v_0 yields

$$EPO_0(\beta_0, \sigma_1) = \frac{v_0^3}{3} \beta_0 \sigma_1 - \frac{3\beta_0^2 v_0^3 \sigma_1}{12} \Big|_0^1 = \frac{\beta_0 \sigma_1}{3} - \frac{\beta_0^2 \sigma_1}{4}. \quad (8.22)$$

Taking the FOC

$$\frac{\partial}{\partial \beta_0} EPO_0(\beta_0, \sigma_1) = \frac{\sigma_1}{3} - \frac{2\beta_0 \sigma_1}{4} \stackrel{!}{=} 0 \quad (8.23)$$

yields the best response function

$$\beta_0^{\text{pBR}}(\sigma_1) = \frac{2}{3}. \quad (8.24)$$

So we arrive at the following

Theorem 8.3 *Consider a repeated single-unit double auction with one bidder and one seller, linear strategies and $v_i \sim U(0, 1)$. If a trade occurs, the price is set in the middle of bid and ask. Then the perfect memory best response function of the bidder is given by*

$$\beta_0^{\text{pBR}}(\sigma_1) = \frac{2}{3} \quad \forall \sigma_1. \quad (8.25)$$

Note in particular that the perfect memory best response function of the bidder does not depend on the strategy of the seller, so it is a dominant strategy.

Seller Response

The payoff for the seller is given by

$$PO_1(\sigma_1, \beta_0, v_1, v_0) = \begin{cases} 0 & \text{for } \beta_0 v_0 < \frac{v_1}{\sigma_1} \\ \frac{\beta_0 v_0 + \frac{v_1}{\sigma_1}}{2} & \text{for } \beta_0 v_0 > \frac{v_1}{\sigma_1}. \end{cases} \quad (8.26)$$

Therefore

$$PO_1(\sigma_1, \beta_0, v_1, v_0) = \left\{ \frac{\beta_0 v_0 + \frac{v_1}{\sigma_1}}{2} - v_1 \right\} \Theta(\beta_0 v_0 \sigma_1 - v_1). \quad (8.27)$$

Integration over v_1 yields

$$\begin{aligned} EPO_1(\sigma_1, \beta_0, v_0) &= \int_0^{\beta_0 v_0 \sigma_1} dv_1 \left(\frac{\beta_0 v_0}{2} + \frac{v_1}{2\sigma_1} - v_1 \right) = \frac{\beta_0 v_0 v_1}{2} + \frac{v_1^2}{4\sigma_1} - \frac{v_1^2}{2} \Big|_0^{\beta_0 v_0 \sigma_1} = \\ &= \frac{\beta_0^2 v_0^2 \sigma_1}{2} + \frac{\beta_0^2 v_0^2 \sigma_1}{4} - \frac{\beta_0^2 v_0^2 \sigma_1^2}{2} = \frac{3\beta_0^2 v_0^2 \sigma_1}{4} - \frac{\beta_0^2 v_0^2 \sigma_1^2}{2}. \end{aligned} \quad (8.28)$$

Subsequent integration over v_0 yields

$$EPO_1(\sigma_1, \beta_0) = \int_0^1 dv_0 EPO_1(\sigma_1, \beta_0, v_0) = \frac{\beta_0^2 \sigma_1}{4} - \frac{\beta_0^2 \sigma_1^2}{6}. \quad (8.29)$$

Solving the FOC

$$\frac{\partial}{\partial \sigma_1} EPO_1(\sigma_1, \beta_0) = \frac{\beta_0^2}{4} - \frac{\sigma_1 \beta_0^2}{3} \stackrel{!}{=} 0 \quad (8.30)$$

yields the following

Theorem 8.4 Consider a repeated single-unit double auction with one bidder and one seller, linear strategies and $v_i \sim U(0, 1)$. If a trade occurs, the price is set in the middle of bid and ask. Then the perfect memory best response function of the seller is given by

$$\sigma_1^{pBR}(\beta_0) = \frac{3}{4} \quad \forall \beta_0. \quad (8.31)$$

Note in particular that the perfect memory best response of the seller does not depend on the strategy of the bidder, so it is a dominant strategy.

8.2 Outlook to Future Research: Dynamics in Repeated DAs with Affine Bidding Strategies

As noted in the previous section, in their seminal paper on double auctions, Chatterjee and Samuelson (1983) consider a very simple double auction: One bidder and one seller face each other, each with private value $\sim U(0, 1)$. The price is set exactly in the middle of bid and ask.¹ The authors interpret this setup as a toy-model for union negotiations and their breakdown such as strikes or work stoppages.

The authors derive the Nash Equilibrium solutions as affine functions of the form

$$b_0(v_0) = \alpha_0 + \beta_0 v_0 \quad (8.32)$$

$$s_1(v_1) = \alpha_1 + \frac{v_1}{\sigma_1} \quad (8.33)$$

with

$$\alpha_0 = \frac{1}{12}; \quad \alpha_1 = \frac{1}{4}; \quad \beta_0 = \frac{2}{3}; \quad \sigma_1 = \frac{3}{2}. \quad (8.34)$$

In this section I want to investigate whether best response bidders in repeated double auctions have any chance to find this NE under best response dynamics. In other words, I check whether the NE is attracting, repelling or is a saddle point with respect to the phase space dynamics that arises from best response play. Therefore I extend the methodology that I derived in the previous chapters. The analysis is tedious: Instead of 2 free parameters β_0, β_1 , respectively σ_0, σ_1 , I consider now four free parameters $\alpha_0, \alpha_1, \beta_0, \sigma_1$. Like in the case of first-price auctions, I derive the expected bidder payoff $EPO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1)$ and the expected seller payoff $EPO_1(\alpha_0, \beta_0, \alpha_1, \sigma_1)$. Then I use the first order conditions

$$\begin{aligned} \frac{\partial}{\partial \alpha_0} EPO_0 &\stackrel{!}{=} 0 & \frac{\partial}{\partial \beta_0} EPO_0 &\stackrel{!}{=} 0 \\ \frac{\partial}{\partial \alpha_1} EPO_1 &\stackrel{!}{=} 0 & \frac{\partial}{\partial \sigma_1} EPO_1 &\stackrel{!}{=} 0 \end{aligned}$$

to derive the best response functions of the bidder and the seller.

The necessary calculations involve some nasty expressions. I will in the worst cases refrain from replicating such formula-monsters which provide no intuitive insights. Instead I will focus on the logic of the derivation. I did the calculations in Mathematica and the notebooks are obtainable upon request.

Note that in this section – like in the case of first-price auctions – I again restrict the agents to the minimal functional space that entails the NE. This is motivated by the assumption that sufficiently close to the NE I can closely approximate the true non-linear best response functions by affine ones. By applying this methodology I show that the NE is not attracting. Moreover I discover distinct phase space regions where the parameters are coupled via different sets of equations. Since the other regions are not near the NE anymore, I probably can not accurately describe them by affine functions anymore. The unrestricted adaptation dynamics of non-linear best response functions might be richer than indicated in this section where I look only at affine strategies.

¹Actually, the authors investigate broader classes of double auctions with different price setting mechanisms k and more general value distributions $f(v)$. Their focus is on the comparative statics of the NE in dependence of k and $f(v)$. However, one of their examples is this particular setup.

The remarkable result in this section is that I find NE-repelling dynamics in double auctions even under perfect memory best response. This is in stark contrast to perfect memory best response in single sided 1PAs where players converged to the NE at least in infinite time. The observation that even perfect memory play does not converge to the NE in double auctions raises serious doubts on the efficiency of standard double auction mechanisms in the economy. In particular, it has drastic implications not only in repeated settings but also if we think about strategic sophistication in one-shot auctions where participants play the best response to the best response of their opponent to ... to their prior beliefs. If the participants are not completely sure that their opponent plays in the NE, they will in general play off-NE as a best response. If they are however strategically sophisticated, they will play the best response against the opponent's best response to this off-NE play. The results of this section show that she then is even further away from the NE than by just playing the best response against her believe about the opponent's play. So, instead of convergence to the NE by increasing the iterations of strategic sophistication, bidders who have any initial doubt on the rationality of her opponent or on the quality of information that she possesses are led away from the NE. And they are led away the further, the more sophisticated we assume them. Even more persistent than in the case of single sided auctions I conjecture therefore that double auctions with weakly dominant strategies like the one suggested in McAfee (1992) could significantly reduce risk in the real economy.

In this thesis my focus is on single-sided auctions and linear bidding strategies. The extension of my research to the rigorous mathematical analysis and simulation of double auctions with non-linear bidding functions is something I want to pursue in more detail in the future when I have obtained my PhD.

8.2.1 Expected Payoffs under Perfect Memory Play

I assume affine bidding- and asking functions of the form

$$b_0(v_0) = \alpha_0 + \beta_0 v_0, \quad (8.35)$$

$$s_1(v_1) = \alpha_1 + \frac{v_1}{\sigma_1}. \quad (8.36)$$

A trade is only made if $b_0(v_0) > s_1(v_1) \Leftrightarrow \alpha_0 + \beta_0 v_0 - \alpha_1 - \frac{v_1}{\sigma_1} > 0 \Leftrightarrow v_1 < \alpha_0 \sigma_1 + \beta_0 \sigma_1 v_0 - \alpha_1 \sigma_1$. The price is then given as

$$\frac{1}{2}(b_0(v_0) + s_1(v_1)) = \frac{1}{2}\left(\alpha_0 + \beta_0 v_0 + \alpha_1 + \frac{v_1}{\sigma_1}\right). \quad (8.37)$$

The bidder's payoff is $v_0 - p$ if a trade is made and 0 otherwise. In closed form it can therefore be written as

$$\begin{aligned} PO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1, v_0, v_1) &= (v_0 - p)\Theta(b_0(v_0) - s_1(v_1)) = \\ &= \left(v_0 - \frac{1}{2}\left(\alpha_0 + \beta_0 v_0 + \alpha_1 + \frac{v_1}{\sigma_1}\right)\right)\Theta(\alpha_0 \sigma_1 + \beta_0 v_0 \sigma_1 - \alpha_1 \sigma_1 - v_1). \end{aligned} \quad (8.38)$$

The expected bidder payoff for $v_0, v_1 \sim U(0, 1)$ is given by integrating the payoff over all possible value combinations,

$$EPO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1) = \int_0^1 dv_0 \int_0^1 dv_1 PO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1, v_0, v_1). \quad (8.39)$$

The problem is now to determine the correct integration borders since the values that provide a positive contribution to the integral are restricted by the Θ -function and by the integration borders.

The expected payoff of bidder 0 in dependence on v_0 is given as

$$EPO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1, v_0) = \int_0^1 dv_1 (v_0 - p)\Theta(\alpha_0 \sigma_1 + \beta_0 v_0 \sigma_1 - \alpha_1 \sigma_1 - v_1). \quad (8.40)$$

For later reference I define here the expression E_1 as

$$E_1 := \alpha_0 \sigma_1 + \beta_0 v_0 \sigma_1 - \alpha_1 \sigma_1 . \quad (8.41)$$

The integrand in the integral for EPO_0 gives a positive contribution only if

$$\begin{aligned} 0 &\leq v_1 \leq 1 \\ v_1 &\leq E_1 \end{aligned} \quad (8.42)$$

Therefore v_1 has to be smaller than $\min(1, E_1)$. However, if $E_1 \leq 0$, the expected payoff is not negative but 0. Therefore I multiply with the expression $\Theta(E_1)$ that guarantees that the expected payoff is zero for $E_1 < 0$. This simplifies equation (8.40) to

$$EPO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1, v_0) = \Theta(E_1) \int_0^{\min(1, E_1)} dv_1 (v_0 - p) . \quad (8.43)$$

The expected payoff in dependence only of the parameters is then obtained by integration over v_0 as

$$EPO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1) = \int_0^1 dv_0 \Theta(E_1) \int_0^{\min(1, E_1)} dv_1 (v_0 - p) . \quad (8.44)$$

Rewriting $\Theta(E_1)$ as

$$\Theta(E_1) = \Theta(\alpha_0 \sigma_1 + \beta_0 v_0 \sigma_1 - \alpha_1 \sigma_1) = \Theta\left(v_0 - \frac{\alpha_1 - \alpha_0}{\beta_0}\right) \quad (8.45)$$

shows that the integrand takes on non-zero values only for $v_0 > \frac{\alpha_1 - \alpha_0}{\beta_0}$ and the integral is 0 if $\frac{\alpha_1 - \alpha_0}{\beta_0} > 1$. Therefore I obtain a closed form solution for the expected payoff as

$$\begin{aligned} EPO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1) &= \Theta\left(1 - \frac{\alpha_1 - \alpha_0}{\beta_0}\right) \int_{\frac{\alpha_1 - \alpha_0}{\beta_0}}^1 dv_0 \int_0^{\min(1, E_1)} dv_1 (v_0 - p) = \\ &= \Theta(\alpha_0 + \beta_0 - \alpha_1) \int_{\frac{\alpha_1 - \alpha_0}{\beta_0}}^1 dv_0 \int_0^{\min(1, E_1)} dv_1 (v_0 - p) . \end{aligned} \quad (8.46)$$

Completely analogously I obtain for the expected seller payoff

$$EPO_1(\alpha_0, \beta_0, \alpha_1, \sigma_1) = \Theta(\alpha_0 + \beta_0 - \alpha_1) \int_{\frac{\alpha_1 - \alpha_0}{\beta_0}}^1 dv_0 \int_0^{\min(1, E_1)} dv_1 (p - v_1) \quad (8.47)$$

The value of the expected payoff for both, bidder and seller, depends critically on the value of E_1 because the upper border of the integral over E_1 is given by $\min(1, E_1)$. Therefore, I have to distinguish three regions of the four-dimensional phase space $(\alpha_0, \beta_0, \alpha_1, \sigma_1)$, depending on the value of $E_1 = \alpha_0 \sigma_1 + \beta_0 v_0 \sigma_1 - \alpha_1 \sigma_1$:

$$\begin{aligned} \text{Case 1: } &E_1 \leq 1 \quad \forall v_0 \in (0, 1) \\ \text{Case 2: } &E_1 \geq 1 \quad \forall v_0 \in (0, 1) \\ \text{Case 3: } &E_1 < 1 \quad \text{for some } v_0 \in (0, 1) \text{ and } E_1 > 1 \text{ for some } v_0 \in (0, 1) \end{aligned} \quad (8.48)$$

In the following I treat the three cases one by one.

8.2.2 The NE Region

Starting with case 1, a sufficient and necessary condition for $E_1 = \alpha_0\sigma_1 + \beta_0v_0\sigma_1 - \alpha_1\sigma_1 \leq 1$ $\forall v_0 \in (0, 1)$ is that the condition is satisfied for $v_0 = 1$ since $v_0 < 1$ gives an even lower value of E_1 . Therefore the condition $E_1 \leq 1$ reduces to

$$\alpha_0\sigma_1 + \beta_0\sigma_1 - \alpha_1\sigma_1 \leq 1 . \quad (8.49)$$

For later reference I denote this value-independent expression on the left side of the equation as E_2 :

$$E_2 \stackrel{!}{=} \alpha_0\sigma_1 + \beta_0\sigma_1 - \alpha_1\sigma_1 \quad (8.50)$$

Note that the NE derived by Chatterjee and Samuelson (1983) lies in the region $E_2 \leq 1$ because

$$\frac{1}{12} \frac{3}{2} + \frac{2}{3} \frac{3}{2} - \frac{1}{4} \frac{3}{2} = \frac{3}{4} < 1 \quad (8.51)$$

In this region $\min(E_1, 1) = E_1 \forall v_0$ and therefore the expected bidder payoff is given as

$$\begin{aligned} EPO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1) &= \Theta(\alpha_0 + \beta_0 - \alpha_1) \int_{\frac{\alpha_1 - \alpha_0}{\beta_0}}^1 dv_0 \int_0^{E_1} dv_1 (v_0 - p) = \\ &= -\Theta(\alpha_0 + \beta_0 - \alpha_1) \frac{\sigma_1}{12\beta_0^2} \{(\alpha_0 - \alpha_1 + \beta_0)^2(\beta_0(3\beta_0 - 4) + \alpha_1(3\beta_0 - 2) + \alpha_0(3\beta_0 + 2))\} . \end{aligned} \quad (8.52)$$

Note that the term $\Theta(\alpha_0 + \beta_0 - \alpha_1)$ simply says that the expected payoff is 0 if the parameters are such that a trade can never happen. If $\alpha_1 > \alpha_0 + \beta_0$, then, even for the extreme value combination $v_0 = 1, v_1 = 0$, no trade would happen because the amount that the seller asks is more than the bidder is willing to give. Surely then, also for all other value combinations, no trade would happen since they would lead the bidder to bid even less and the seller to ask more.

In completely the same way, the seller payoff in this phase space regions is obtained as

$$\begin{aligned} EPO_1(\alpha_0, \beta_0, \alpha_1, \sigma_1) &= \Theta(\alpha_0 + \beta_0 - \alpha_1) \int_{\frac{\alpha_1 - \alpha_0}{\beta_0}}^1 dv_0 \int_0^{E_1} dv_1 (p - v_1) = \\ &= -\Theta(\alpha_0 + \beta_0 - \alpha_1) \frac{\sigma_1}{12\beta_0} \{(\alpha_0 + \beta_0 - \alpha_1)^2(\beta_0(2\sigma_1 - 3) + \alpha_0(2\sigma_1 - 3) - \alpha_1(2\sigma_1 + 3))\} . \end{aligned} \quad (8.53)$$

Of course, the same closed-form solutions for the expected payoff of the bidder and the seller could have also been obtained by integrating first over v_0 and in the second step over v_1 . However, the choice of the appropriate integration borders has to be adapted. I have calculated the expected payoff with reversed order of integration in a Mathematica-Notebook that can be obtained from me.

The best response parameter settings are – like in the case of first-price auctions – obtained by solving the FOCs for the bidder and the seller. The bidder chooses her parameters (α_0, β_0) such as to maximize EPO_0 whereas the seller chooses (α_1, σ_1) to maximize EPO_1 . For the

linear parameters α_0, α_1 , I get two solutions each, for σ_1 I get a unique solution:

$$\begin{aligned}
\frac{\partial}{\partial \alpha_0} EPO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1) &\stackrel{!}{=} 0 \Rightarrow \\
\alpha_0^{(1)}(\beta_0, \alpha_1, \sigma_1) &= \alpha_1 - \beta_0 \\
\alpha_0^{(2)}(\beta_0, \alpha_1, \sigma_1) &= \frac{\beta_0(2 - 3\beta_0) + \alpha_1(2 - \beta_0)}{2 + 3\beta_0}, \\
\frac{\partial}{\partial \alpha_1} EPO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1) &\stackrel{!}{=} 0 \Rightarrow \\
\alpha_1^{(1)}(\alpha_0, \beta_0, \sigma_1) &= \alpha_0 + \beta_0 \\
\alpha_1^{(2)}(\alpha_0, \beta_0, \sigma_1) &= \frac{(\alpha_0 + \beta_0)(2\sigma_1 - 1)}{3 + 2\sigma_1}, \\
\frac{\partial}{\partial \sigma_1} EPO_1(\alpha_0, \beta_0, \sigma_1) &\stackrel{!}{=} 0 \Rightarrow \\
\sigma_1(\alpha_0, \beta_0, \alpha_1) &= \frac{3(\alpha_0 + \alpha_1 + \beta_0)}{4(\alpha_0 - \alpha_1 + \beta_0)}.
\end{aligned} \tag{8.54}$$

For $\beta_0(\alpha_0, \alpha_1, \sigma_1)$, Mathematica gives four solutions. One solution, $\beta_0^{(1)} = \alpha_1 - \alpha_0$, is the mirror image of the first set of solutions for $\alpha_0^{(1)}$ and $\alpha_1^{(1)}$. The three other solutions are in general given by complex numbers. However, one of the solutions has the nice feature that in the NE and in a region around the NE, the imaginary part of that solution is equal to 0. Since the parameters have to be real-valued numbers, this is the apt solution to approximate the behavior of bidders around the NE.² The other two complex solutions have a non-vanishing imaginary part in this region.

I refrain here from explicitly writing the solution because this would need about one page and doesn't provide any interesting insights. The Mathematica Notebook with the explicit solution for β_0 can of course be obtained from the author.

So, there exist two sets of solutions that describe the parameter evolution of bidder and seller around the NE. The first set of solutions consists of the equations

$$\begin{aligned}
\alpha_0(\alpha_1, \beta_0, \sigma_1) &= \alpha_1 - \beta_0 \\
\beta_0(\alpha_0, \alpha_1, \sigma_1) &= \alpha_1 - \alpha_0 \\
\alpha_1(\alpha_0, \beta_0, \sigma_1) &= \alpha_0 + \beta_0 \\
\sigma_1(\alpha_0, \beta_0, \alpha_1) &= \frac{3(\alpha_0 + \alpha_1 + \beta_0)}{4(\alpha_0 - \alpha_1 + \beta_0)}.
\end{aligned} \tag{8.55}$$

Parameter substitution of α_1 in the last equation yields

$$\sigma_1(\alpha_0, \beta_0, \alpha_1) = \frac{3(\alpha_0 + \alpha_0 + \beta_0 + \beta_0)}{4(\alpha_0 - \alpha_0 - \beta_0 + \beta_0)} = \infty. \tag{8.56}$$

So the seller asks in this scenario a value-independent reserve price of $\alpha_1 = \alpha_0 + \beta_0$ ³. Since the bid is always smaller than the ask,

$$b_0(v_0) = \alpha_0 + \beta_0 v_0 < \alpha_0 + \beta_0 = \alpha_1 = s_1(v_1) \quad \forall v_0 \in (0, 1), \tag{8.57}$$

this is not a NE. A unilateral increase of α_0 or β_0 by the bidder or a decrease of α_1 by the seller would increase the revenue of both participants. Therefore, this set of solution minimizes the expected payoff of both participants as opposed to the NE strategy combination that would locally maximize it.

Note that the formula for the EPO_0 and EPO_1 correctly predicts zero expected payoff in this phase space region because of the multiplicative term $\Theta(\alpha_0 + \beta_0 - \alpha_1)$.

²I have to assess in the future whether the imaginary part really vanishes in all of the phase-space region with $E_1 \leq 1$. But for all the parameter combinations that I tried out, it did vanish.

³Since the term $\frac{v_1}{\sigma_1}$ vanishes for $\sigma_1 \rightarrow \infty$.

The more interesting case is the second set of solutions which entails the Chatterjee and Samuelson (1983)- NE.

$$\begin{aligned}
\alpha_0(\alpha_1, \beta_0, \sigma_1) &= \frac{\beta_0(2 - 3\beta_0) + \alpha_1(2 - \beta_0)}{2 + 3\beta_0} \\
\beta_0(\alpha_0, \alpha_1, \sigma_1) &= [\text{A very nasty expression (see Mathematica)}] \\
\alpha_1(\alpha_0, \beta_0, \sigma_1) &= \frac{(\alpha_0 + \beta_0)(2\sigma_1 - 1)}{3 + 2\sigma_1} \\
\sigma_1(\alpha_0, \beta_0, \alpha_1) &= \frac{3(\alpha_0 + \alpha_1 + \beta_0)}{4(\alpha_0 - \alpha_1 + \beta_0)}.
\end{aligned} \tag{8.58}$$

The solution to this system of equations is the NE derived in Chatterjee and Samuelson (1983),

$$\{\alpha_0^*, \beta_0^*, \alpha_1^*, \sigma_1^*\} = \left\{ \frac{1}{12}, \frac{2}{3}, \frac{1}{4}, \frac{3}{2} \right\}, \tag{8.59}$$

since

$$\begin{aligned}
\alpha_0(\beta_0^*, \alpha_1^*, \sigma_1^*) &= \frac{\frac{2}{3}(2 - 2) + \frac{1}{4}(2 - \frac{2}{3})}{2 + 2} = \frac{\frac{1}{4} \cdot \frac{4}{3}}{4} = \frac{1}{12} \\
\beta_0(\alpha_0^*, \alpha_1^*, \sigma_1^*) &= [\text{See Mathematica Notebook}] = \frac{2}{3} \\
\alpha_1(\alpha_0^*, \beta_0^*, \sigma_1^*) &= \frac{(\frac{1}{12} + \frac{2}{3})(2 \cdot \frac{3}{2} - 1)}{3 + 3} = \frac{\frac{3}{4} \cdot 2}{6} = \frac{1}{4} \\
\sigma_1(\alpha_0^*, \beta_0^*, \alpha_1^*) &= \frac{3(\frac{1}{12} + \frac{1}{4} + \frac{2}{3})}{4(\frac{1}{12} - \frac{1}{4} + \frac{2}{3})} = \frac{3 \cdot 1}{4 \cdot \frac{1}{2}} = \frac{3}{2}.
\end{aligned} \tag{8.60}$$

These is exactly the parameter combination that is derived in Chatterjee and Samuelson (1983) for the NE. In the following section I will investigate the phase space dynamics around the NE.

8.2.3 Dynamics around the NE

In the following I use the set of equations (8.58) to investigate the phase space dynamics around the Nash Equilibrium. In particular I am interested wether the NE is attracting, repelling or wether it shows saddle point behavior. For this I consider what happens if bidders repeatedly adapt their parameters to the parameters of their opponents.

From an interpretational point of view I am describing here the long term dynamics of repeated double auctions or repeated bargaining situations under truncated fictitious play: I describe bidders and sellers who stick to their strategies for sufficiently many rounds that the opponent can accurately assess them.⁴ After this time, both opponents simultaneously update their strategies according to the set of equations (8.58). In particular, I am interested wether this dynamics leads bidders from near the NE into the NE or away from it, i.e., I am interested whether the NE is attracting or repelling.

As a first assessment of the system's behavior, I computed the table of comparative statics around the NE: Table 8.1 shows the direction of change of the row-parameter if the column-parameter is increased and the other parameters are set to their NE values. A plus means that an increase in the column-parameter induces an increase in the row-parameter, a minus means that an increase in the column-parameter induces a decrease in the row-parameter.

The comparative static results are quite intuitive:

Looking at the first column, an increase in the reservation price α_0 allows the bidder to decrease her linear component β_0 . However, it allows the seller to ask more by increasing α_1 or reducing σ_1 .⁵

The analogous reasoning holds for β_0 . Note that the mutual negative feedback between α_0 and β_0 induces a vicious cycle between α_0 and β_0 : Increasing β_0 decreases α_0 which in turn increases β_0 again. Therefore the NE is unstable.

⁴This is the same situation as for truncated fictitious play in the case of 1PAs.

⁵Remember that the seller strategy was given by $\alpha_1 + \frac{\beta_1}{\sigma_1}$, so an decrease in σ means that the seller asks more.

| | α_0 | β_0 | α_1 | σ_1 |
|------------|------------|-----------|------------|------------|
| α_0 | + | - | + | 0 |
| β_0 | - | + | + | 0 |
| α_1 | + | + | + | + |
| σ_1 | - | - | + | + |

Table 8.1: Comparative Statics around the NE Region in a Double Auction

An increase in the sellers minimum ask α_1 forces the bidder to bid more on average and hence to increase α_0 and/or β_0 . On the other hand, it allows the seller to reduce her average ask by increasing σ_1 .

Intuitively it is also clear that an increase in σ_1 leads to an increase in α_1 . However, it is not obvious to me, why the bidder parameters α_0 and β_0 do not depend directly on the linear strategy parameter σ_1 of the seller. α_0 and β_0 depend on σ_1 only indirectly via α_1 . This means that a change in σ_1 has an impact on α_0 and β_0 only two rounds later. To find a plausible interpretation of this remains a task for future research.

To get a more detailed understanding of the dynamics I wrote a small program in Mathematica. The program can be obtained from the author. The program flow is as follows:

- Start with an initial configuration $(\alpha_0^{Start}, \alpha_1^{Start}, \beta_0^{Start}, \sigma_1^{Start})$ near the NE
- Calculate the next round's strategy parameter values as the perfect memory best response to the old round's strategies. The according formulas are given in (8.58)
- Repeat this procedure; record and plot the values of the parameters over time
- In each time step also check that the system stays in the appropriate parameter range.

First, I have to control that the imaginary part of β_0 is 0, $\Im(\beta_0) \stackrel{!}{=} 0$. This guarantees that the strategies don't get complex. Second, take care that the system doesn't leave the appropriate region of the parameter space that I used for deriving the expected payoff:

$$E_2 = \alpha_0\sigma_1 + \beta_0\sigma_1 - \alpha_1\sigma_1 \stackrel{!}{\leq} 1 \quad 6$$

I started with an initial configuration $\{\alpha_0 = \alpha_0^* + 0.0001, \beta_0^*, \alpha_1^*, \sigma_1^*\}$ near the NE. Slight deviations of other parameter values yield qualitatively the same results. Figure 8.1 shows that the NE is not stable. A small initial perturbation is of little effect in the first few rounds and the system initially stays near the NE. However, approximately around round 10, the parameters begin to oscillate visibly. The oscillations increase ever more until in round 22, the system leaves the phase space region with $E_2 \leq 1$.

Of course, the mathematically exact way to proof that the NE is not stable would be to analyze the Eigenvalues of the Jacobian of the set of equations (8.58). This is a topic for further research. For me the results that I obtained by the simulation with Mathematica seem sufficient for the following

Conjecture 8.1 *The NE that Chatterjee and Samuelson (1983) derived for double auctions is repelling. Even under perfect memory best response dynamics the players will never move to the NE. Instead, once they have left the NE they will be pushed even further away.*

The phase space dynamics drives the system from the parameter region that entails the NE into the other parameter regions that I denoted by Case 2 and Case 3. In the following section I shortly investigate these phase space regions.

8.2.4 Phase Space Regions off the NE

The NE region was described by the condition $E_1 = \alpha_0\sigma_1 + \beta_0v_0\sigma_1 - \alpha_1\sigma_1 \leq 1$.

In region 2 I consider the part of the phase space where $E_1 \geq 1$. The condition can be written out as

$$E_1 = \alpha_0\sigma_1 + \beta_0v_0\sigma_1 - \alpha_1\sigma_1 \geq 1 \quad \forall v_0 \in (0, 1). \quad (8.61)$$

⁶In general, in all initial parameter constellations that I investigated, the first time that $\Im(\beta_0) \neq 0$ is after the first time that $E_2 \geq 1$. This suggests, that we can ignore the control for $\Im(\beta_0)$ because this is always satisfied in the phase space region $E_2 \leq 1$. However, a proof of this hypothesis is a topic for further research.

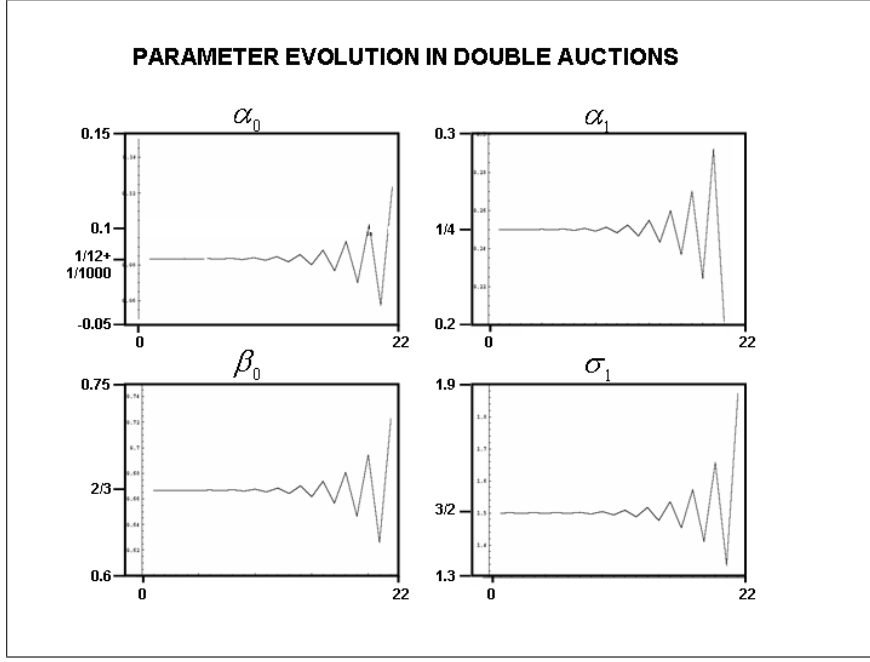


Figure 8.1: Parameter Evolution in Double Auctions Near the NE

Since $\beta_0 v_0 \sigma_1 > 0$, a sufficient condition for this is

$$\alpha_0 \sigma_1 - \alpha_1 \sigma_1 \geq 1. \quad (8.62)$$

Then, $\min(1, E_1) = 1$ and equations (8.52) and (8.53) for the expected payoffs reduce to

$$\begin{aligned} EPO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1) &= \Theta(\alpha_0 + \beta_0 - \alpha_1) \int_{\frac{\alpha_1 - \alpha_0}{\beta_0}}^1 dv_0 \int_0^1 dv_1 (v_0 - p) = \\ &= -\frac{1}{2\beta_0^2 \sigma_1} \{(\alpha_0 + \beta_0 - \alpha_1)(2(\alpha_0 - \alpha_1)\sigma_1 + \beta_0^2 + \beta_0(1 + (\alpha_0 + 3\alpha_1 - 2)\sigma_1))\} \end{aligned} \quad (8.63)$$

and

$$\begin{aligned} EPO_1(\alpha_0, \beta_0, \alpha_1, \sigma_1) &= \Theta(\alpha_0 + \beta_0 - \alpha_1) \int_{\frac{\alpha_1 - \alpha_0}{\beta_0}}^1 dv_0 \int_0^1 dv_1 (p - v_1) = \\ &= \frac{1}{4\beta_0 \sigma_1} \{(\alpha_0 + \beta_0 - \alpha_1)(1 + (\alpha_0 + \beta_0 + 3\alpha_1 - 2)\sigma_1)\}. \end{aligned} \quad (8.64)$$

Note that the Θ -function is always one since

$$E_1 \geq 1 \Leftrightarrow \alpha_0 + \beta_0 - \alpha_1 \geq \frac{1}{\sigma_1} \geq 0. \quad (8.65)$$

Solving the FOCs yields the equations

$$\begin{aligned} \alpha_0 &= -\frac{\beta_0 + \beta_0^2 - 4\alpha_1 \sigma_1 + 2\alpha_1 \beta_0 \sigma_1 + \beta_0^2 \sigma_1}{2(2 + \beta_0)\sigma_1} \\ \alpha_1 &= -\frac{1 - 2\sigma_1 - 2\alpha_0 \sigma_1 - 2\beta_0 \sigma_1}{6\sigma_1} \\ \sigma_1 &= -\frac{2}{\alpha_0 + 3\alpha_1 - \beta_0 - 2} \end{aligned} \quad (8.66)$$

Again, for β_0 there are three solutions, where only one of them is real.

Region 3 has the condition $E_1 = 1$ for some $\tilde{v}_0 \in (0, 1)$. This condition can be rewritten as a condition on the parameters:

$$\begin{aligned} \alpha_0\sigma_1 + \beta_0\tilde{v}_0\sigma_1 - \alpha_1\sigma_1 &= 1 \Leftrightarrow \\ \Leftrightarrow \tilde{v}_0 &= \frac{1}{\beta_0\sigma_1} + \frac{\alpha_1}{\beta_0} - \frac{\alpha_0}{\beta_0} \in (0, 1) \end{aligned} \quad (8.67)$$

Therefore, $E_1 < 1$ iff $v_0 < \frac{1}{\beta_0\sigma_1} + \frac{\alpha_1}{\beta_0} - \frac{\alpha_0}{\beta_0}$ and $E_1 > 1$ iff $v_0 > \frac{1}{\beta_0\sigma_1} + \frac{\alpha_1}{\beta_0} - \frac{\alpha_0}{\beta_0}$. Accordingly, the expected payoffs are given as

$$\begin{aligned} EPO_0(\alpha_0, \beta_0, \alpha_1, \sigma_1) &= \\ = \Theta(\alpha_0 + \beta_0 - \alpha_1) \{ &\int_{\frac{\alpha_1 - \alpha_0}{\beta_0}}^{\frac{\alpha_1 - \alpha_0}{\beta_0} + \frac{1}{\beta_0\sigma_1}} dv_0 \int_0^{E_1} dv_1 (v_0 - p) + \int_{\frac{\alpha_1 - \alpha_0}{\beta_0} + \frac{1}{\beta_0\sigma_1}}^1 dv_0 \int_0^1 dv_1 (v_0 - p) \} \end{aligned} \quad (8.68)$$

and

$$\begin{aligned} EPO_1(\alpha_0, \beta_0, \alpha_1, \sigma_1) &= \\ = \Theta(\alpha_0 + \beta_0 - \alpha_1) \{ &\int_{\frac{\alpha_1 - \alpha_0}{\beta_0}}^{\frac{\alpha_1 - \alpha_0}{\beta_0} + \frac{1}{\beta_0\sigma_1}} dv_0 \int_0^{E_1} dv_1 (p - v_1) + \int_{\frac{\alpha_1 - \alpha_0}{\beta_0} + \frac{1}{\beta_0\sigma_1}}^1 dv_0 \int_0^1 dv_1 (p - v_1) \} \end{aligned} \quad (8.69)$$

The FOCs yield

$$\begin{aligned} \alpha_0 &= \frac{2 + 4\alpha_1\sigma_1 + \beta_0(-1 - 2(\alpha_1 + \beta_0)\sigma_1)}{2(2 + \beta_0)\sigma_1} \\ \alpha_1 &= \frac{-3 + 2(1 + \alpha_0 + \beta_0)\sigma_1}{6\sigma_1} \\ \sigma_1 &= \frac{6}{4 + 3\alpha_0 - 9\alpha_1 + 3\beta_0} \end{aligned} \quad (8.70)$$

Like in the other two phase space regions, there again exist three solutions for β_0 that can be in general complex but where one of them seems to be real in the whole region. The dynamic assessment of this phase space region via simulations suggests that this phase space region is left within one round after it is entered.

In conclusion this section suggest that the mathematical formalism that I developed for the analysis of perfect memory best response in 1PAs is a promising route for further research on more general auction setups. Of course, the analysis in this section is not comprehensive: I derived some times only sufficient but not necessary conditions for the determination of the different phase-space regions. This means that there are some parts of the phase-space that I didn't analyze yet. Moreover, also for the regions that I investigated much remains to be done – like a rigorous proof that always the same solution of β_0 is real within a given region or an intuitive explanation why the solutions for β_0 are so much more complicated than the respective equations for the other parameters.

The switching of the phase space dynamics between the different regions and thereby a comprehensive description of the – probably chaotic – behavior of mutually adapting players in double auctions under truncated fictitious play with sufficient long time between the updatings remains a topic for future research. In this chapter I just wanted to line out some of the possibilities that the methodology of perfect memory best response provides and to show that the NE that was derived in Chatterjee and Samuelson (1983) is repelling.

Chapter 9

The Auction Simulator

9.1 Motivation for Simulation

As noted in the introduction, the idea of simulating auctions is not new, see e.g., Andreoni and Miller (1995), Dawid (1999) or Byde (2002). Still, there are various reasons why I developed my own computational auction simulation, the Auction Simulator(AS), in parallel to the mathematical models about bounded rational behavior in auctions:

- First of all, the development of the simulation provided guidance in identifying the driving forces of bounded rational bidding in IPOBAs. It forced me to think through the various decisions that bounded rational bidders face. Therefore, the model gave significant help in structuring the analysis and the proofs in the mathematical part of the thesis.
- Even if several people have worked through mathematical proofs, this is still not a guarantee that they have no flaws. Therefore it is of great help to be able to double-check the quantitative predictions from the pertinent theorems 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 most predictions of the mathematical analysis. Therefore, it is a good starting place to investigate learning in more general auction setups in the same conceptual framework that are mathematically nasty or simply too complicated for a closed-form solution.

Because of all these reasons I programmed the Auction Simulator. At the current stage, the program can simulate single-unit open-bid double-auctions under best response and payoff-exponential quantal response dynamics with myopic, incomplete and perfect memory. The strategy set is at the current stage restricted to linear ones. By setting appropriate parameters, pure bidder or seller first-, second- or higher-price auctions can be simulated. The AS can deal with an arbitrary number of participants and with a wide range of random value distributions. In most simulations I investigated single sided auctions.¹ The simulation of double auctions with the AS, is at the current stage not very insightful since, as discussed above, the NE in double auctions is not given by linear bidding strategies.² Still, the program is designed to be easily extendable to the simulation of double auctions in the future.

9.2 Program Architecture

The AS is programmed in SWARM. SWARM is an object oriented programming language that was designed at the Santa Fe Institute. In principle it is a library of Objective C and

¹In fact, the AS always simulates double auctions. However, by choosing fixed values and /or strategies for one side of the market this reduces effectively to the simulation of single sided auctions. I will say more on this in the next chapter.

²Compare the chapter about Double Auctions.

Java routines. It is published under the open GNU license and is freely downloadable from <http://wiki.swarm.org>. It runs under Linux, Windows, Apple and Unix operating systems and guarantees the transferability of code among different platforms.

SWARM has an active user community that is organized in several newsgroups. Questions on programming technicalities are usually answered within few hours. A variety of tutorials are available which provide also complete programming newbies with the skills to write their first applications.

The power of SWARM is that it provides a number of routines for facilitating common tasks in agent based modeling (ABM) while still allowing for all the flexibility of the underlying programming languages Objective C and Java. The term agent based models defines computer models that are built "bottom-up". The programmer puts considerable effort into defining how one single agent reacts to her environment. After, a multitude of agents are put together and the model calculates the emergent properties of the system: Each agent's behavior influences the other agents' reactions which in turn influence the behavior of the agent again.

A dummy model of such a system is an ant colony. Even if the individual ant shows very little intelligence and mainly reacts to the pheromone trails of other ants, the ant colony at a collective level shows intelligence in identifying the shortest ways to food resources and coordinating the ants on these paths.

Another example would be financial market models – as discussed in the third chapter – where traders with heterogeneous technical trading strategies interact in a common market. The trading strategies influence supply and demand of the investors. This influences the price which in turn influences again the trading strategies of the investors.

Several famous models have been (re)built in SWARM. As a prominent example, the Santa Fe Artificial Stock Market (SFASM) - though originally developed in plain Objective C - has been ported to SWARM and extended by several types of investors, including investors that are modeled via neural networks. The extension of the SFASM by new features is an ongoing effort. The latest SWARM implementation of the SFASM is downloadable, e.g., from <http://sourceforge.net/projects/artstkmkt>. Luna and Stefansson (2000) collect articles on economic simulations that were implemented in SWARM. For a recension of that book and of the discussed models see <http://jasss.soc.surrey.ac.uk/4/2/reviews/johnson.html>. For a book that focuses more on the building of simulations than on the economic intuitions behind them see Luna and Perrone (2002). A comprehensive overview of books and articles that deal with SWARM simulations can be found at <http://www.swarm.org/pubs.html>.

Figure 9.1 schematically shows the program architecture of the AS. The parts have the following

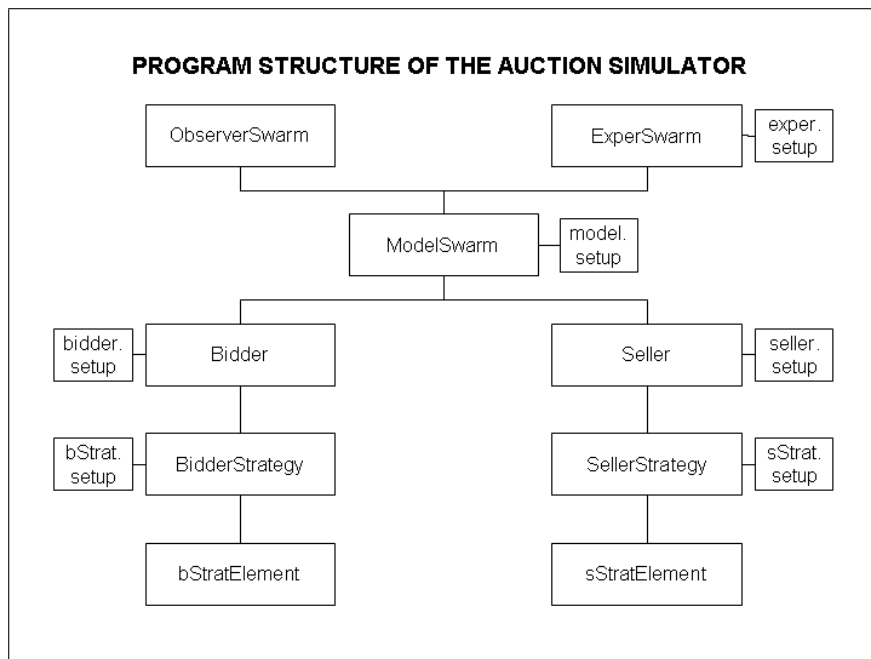


Figure 9.1: Program Structure of the Auction Simulator

purposes:

ObserverSwarm: This object allows the real-time observation of simulation runs. The `ObserverSwarm` manages graphical output in windows that show online the time series of bidder and seller strategies, time series of prices and payoffs as well as the histograms of strategies.

ExperSwarm: This object governs the writing of simulation results to files for further evaluation with statistical programs.

exper.setup: This file determines simulation technicalities like ranges and step sizes of parameter values or random seeds.

ModelSwarm: This object governs the simulation flow as described in the next section. The `ModelSwarm` holds lists with competing sellers and bidders and manages the assignment of payoffs to strategies under various auction formats. The allocation of payoffs will be explained below.

model.setup: This file determines the auction format used in the simulation, the number of buyers and sellers and the random seed for real-time observation.

Bidder resp. **Seller:** These objects represent the individuals that participate in the auction. In particular, each bidder possesses a list of strategies that are updated according to best response or quantal response via Fixed Strategies (FSs) or Genetic Algorithms (GAs). For details on the updating mechanism see below.

bidder.setup resp. **seller.setup:** This file determines the details of the simulation of learning like the number of strategies, whether FSs or GAs are used, the `memoryStrength`, etc.

BidderStrategy resp. **SellerStrategy:** These objects hold the constant, linear and square parameters for each strategy.³

bStrat.setup resp. **sStrat.setup:** These files determine whether constant, linear and quadratic terms are included in the strategies.

bStratElement resp. **sStratElement:** These objects hold the values of the strategy parameters.

9.3 Main Aspects of the Program

9.3.1 Simulation Flow

The program simulates an arbitrary number of bidders, `numBidders` and sellers, `numSellers`. Each participant i has access to a private set of `numStrategies` strategies. Strategies are real numbers β_i for bidders resp. σ_i for sellers. Each player chooses an active strategy that determines the bid $\beta_i v_i$ resp. the ask $\frac{v_i}{\sigma_i}$ that she 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` rounds. In the FS case, simply a new strategy is chosen as active while the strategy set remains fixed. In the GA case, at these times, the strategy population is additionally replaced by a new generation.

A model run consists of the following steps:

1. in the beginning, participants 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. players' values are randomly chosen from the underlying value distribution
 - ii. players submit their bids and asks according to their active strategy
 - iii. the auction module determines the winners
 - iv. players update the current payoffs of their active and passive strategies
 - (c) players evaluate the payoff, their strategies (would) have generated during the `numRounds` auctions and choose an active one for the next `numRounds` auctions according to best response or quantal response dynamics

³The feature for constant and square parameters is not fully implemented yet, currently I can simulate only linear strategies. However, the future extension to nonlinear value dependent strategies is already considered in the program architecture.

- (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. players 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.

9.3.2 Determination of Payoffs

The core of the auction simulator is the determination of payoffs for active and inactive strategies. At the moment, the auction simulator can simulate arbitrary k^{th} price auctions with an arbitrary number of participants. Figure 9.2 shows exemplarily the basic principle for an auction with three bidders and three sellers: The left most part of the Figure shows for each

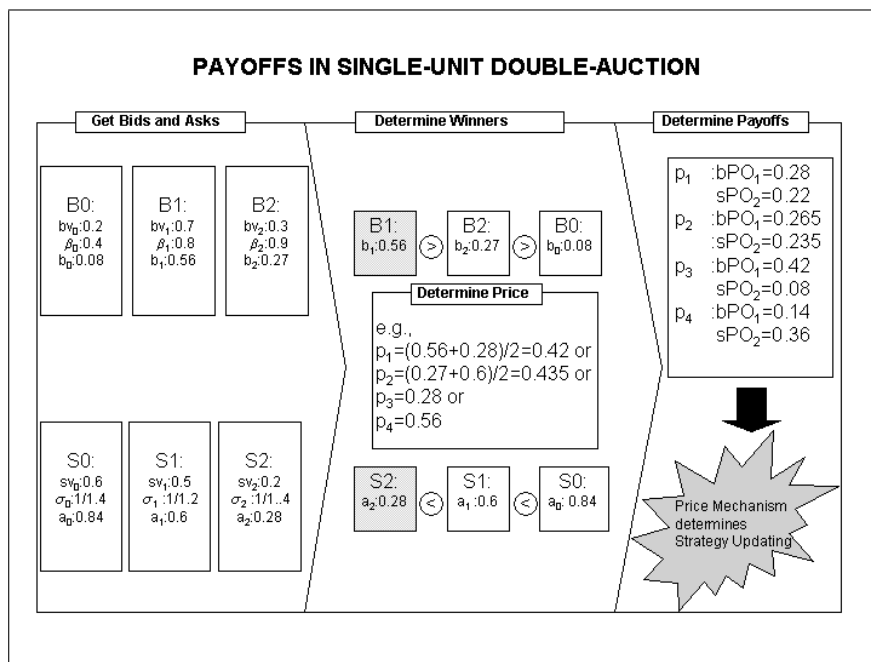


Figure 9.2: Payoffs in single-unit double-auctions

bidder and each seller her value, her strategy and her bid resp. ask. The winning players the highest bidder B1 with a bid of 0.56 and the lowest seller S2 with an ask of 0.28. The middle part of the figure shows, how the prices can be calculated in various ways to reflect different auction formats:

P1: (settings: `auctionType=1`, `priceDet=0.5`). This simulates a first-price double-auction where the price is set in the middle between the highest bid and the lowest ask.

P2: (settings: `auctionType=2`, `priceDet=0.5`). This simulates a second-price double-auction where the price is set in the middle between the second highest bid and the second lowest ask.

P3: (settings: `auctionType=1`, `priceDet=1`). This simulates a first-price auction where the winning seller gets exactly her ask but no more. In connection with setting a fixed very high bid for each bidder⁴, this setup allows the simulation of pure seller auctions.

P4: (settings: `auctionType=1`, `priceDet=0`). This simulates a first-price auction where the winning bidder pays exactly her bid but no less. In connection with setting a fixed value of 0 for each seller, this setup allows the simulation of pure bidder auctions.

The rightmost part of Figure 9.2 shows the bidder and seller payoffs for the different price

⁴to ensure that each winning ask is payed for

determination mechanisms. The Figure illustrates only a fraction of all the possible parameter combinations. In addition, the AS is also able to simulate 3rd and higher price auctions (`auctionType=3, 4, ...`), provided that sufficiently many bidders and sellers participate in the auction.

9.4 Learning in the Auction Simulator

9.4.1 Real GAs

In the Auction Simulator I use real GAs instead of ones with bit string populations. Real GAs are faster than bit string GAs, note however that the theoretical properties are less well understood than for bit string GAs (see, e.g., Holland (1992) and Lux and Schornstein (2002)). In particular, it is not obvious what the equivalent of Holland's Schema Theorem should be. Still, also with this lack of theoretical foundation, real GAs seem to work, which is why they are used in practice. To check for the differences between real and bit string GAs, I also implemented in an earlier version bit string GAs with the AS. The simulation results are basically the same as in the case of real GAs, however, the simulations take a much longer time. Therefore I report in this thesis only the results of the simulations with real GAs.

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.⁵ 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 and select the strategies that determine future play:

The best strategy is chosen as the active one for the next generation. Therefore, GAs at the current stage are apt for the simulation of best response but not of quantal response play. 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 the `breedingList`. Two strategies β_i and β_j are considered as different if $|\beta_i - \beta_j| > \text{strategyDistance}$.⁶ 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 applying 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 selection for crossover is better suited than fitness proportional selection for the simulation of best response dynamics. 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 very small. Having selected the two strategies, denote the lower by β^{\min} and the larger by β^{\max} . A new strategy is constructed by selecting a number between $\beta^{\min}(1 - \text{crossoverPar}(\beta^{\max} - \beta^{\min}))$ and $\beta^{\max}(1 + \text{crossoverPar}(\beta^{\max} - \beta^{\min}))$ with uniform probability. The extension of the crossover interval offsets the tendency of crossover to equalize all strategies⁷.

The third step is to subject the offspring to the mutation operator:

Strategy β is with probability `mutationProb` 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.

⁵For the inactive strategies this is the payoff, they *would* have generated if they had been the active one.

⁶This prevents the GA from getting stuck in too homogenous populations

⁷Usually the literature suggests a normal distribution for the crossover. However, in my experience a uniform distribution yields better results.

Together with the `numElite` unchanged strategies, the newly constructed strategies form the next generation strategy set.

9.4.2 Fixed Strategies

As an alternative simulation tool I employ fixed strategies:

In the beginning I partition the strategy space $(0, 1)$ into a grid of equidistant strategies, e.g. `numStrategies`=100 ranging from 0.00 to 0.99. In the first round, 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 $PO_0^{\text{curr};(i)}(\beta_0) = v_0(1 - \beta_0)\Theta(\beta_0v_0 - \beta_1v_1)$.

After the first `numRounds` auctions, each strategy's payoffs are added up and yield the payoff, the strategy generated in the first `numRounds` auctions:

$$PO_0^{\text{numRounds};(1)}(\beta_0) = \sum_{i=1}^{\text{numRounds}} PO_0^{\text{curr};(i)}(\beta_0) \quad (9.1)$$

The strategy that generated the highest payoff in the first `numRounds` rounds is then selected as the active one for the next `numRounds` rounds. If no strategy has positive payoff⁸, the currently active strategy remains active also in the next `numRounds`. The case that no strategy has positive payoff happens if $v_0 < \beta_1v_1$ for all auctions within the last generation. This is the equivalent of the rule $\beta_0^{\text{new}} = \beta_0^{\text{old}}$ in the mathematical part if no feasible strategy could have won the auction. Denote

$$PO_0^{\text{cum};(1)}(\beta_0) = PO_0^{\text{numRounds};(1)}(\beta_0) \quad (9.2)$$

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:

$$PO_0^{\text{cum};(2)}(\beta_0) = \text{memoryStrength}PO_0^{\text{cum};(1)}(\beta_0) + PO_0^{\text{numRounds};(2)}(\beta_0) \quad (9.3)$$

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

9.4.3 Interpretation of GA and FS Results

It is a well known fact that GAs are very efficient in finding extreme values of functions - even if these are quite rough with many local minima and maxima. However, GAs perform quite poorly in optimizing functions that fluctuate wildly over time. Why? In each generation the population tries to concentrate on the current maximum and therefore it is pure luck whether the population has in the next generation enough individuals near next generation's maximum to identify it accurately. In this sense, the GA interpolates between the optima at consecutive times by leaving behind 'streaks' of old solutions like when you clean a window with a bad cloth. The GA 'remembers' past results and intermingles them with current results in a highly nonlinear fashion.

Having said this it should be clear that in GA simulations we might expect the analytically derived results only if we let players aggregate payoff over sufficiently many rounds before they update the GA. So, the GA should be able to reproduce the predictions of my analysis for perfect memory play if we let it aggregate payoffs over as many rounds as possible⁹.

If, however, the GA is updated after every round, the computational results will not correspond to the analytical results for myopic best response.

The other simulation technique I use are FSs. They are a very mechanical approach to model decision making: In the beginning, a fixed set of strategies is defined and in each step, all

⁸This happens especially likely if `numRounds` is low.

⁹To be precise, with GAs I can only simulate truncated fictitious play not fictitious play which would be the equivalent of perfect memory play. But as the number of rounds increases after which the GA is updated, truncated fictitious play converges to fictitious play.

strategies are evaluated. However, by this simple approach I am able to reproduce the analytical results on all time scales. Moreover, only with FSs, I can use the parameter `memoryStrength` to simulate the effects of continuously varying memory strength on the experimental outcome. I will highlight the differences between FS and GA simulations in the first section of the next chapter.

Chapter 10

Double-check of Analytic Results

This chapter aims at ensuring that the AS functions correctly. To this end I check whether the AS correctly reproduces the analytic results of the previous chapters. The correct functioning and the identification and analysis of possible error sources will ensure in the next chapter that I can more reliably interpret the behavior of bounded rational bidders for cases where I have no closed form solutions.

For the interpretation of the parameter settings in the simulations please consult appendix B

10.1 Bidder Response against Fixed Bidding Strategies

Hypothesis:

The AS confirms the predictions of Theorems 4.1, 4.5, 6.1 and 6.3: Means of best response and quantal response in a two bidder first-price auction with values $v_{0,1} \sim U(0, 1)$ and one bidder playing a fixed strategy β_1 follow the equations that are stated in the pertinent theorems.

Parameter Settings:

Table 10.1 shows the basic parameter setup. The parameters in `exper.setup` determine the number and the length of simulation runs. The parameters in `model.setup` determine the auction setup like the chosen format and the number of participants on each side of the market. `bidder.setup` and `seller.setup` specify the individual behavior of sellers respective bidders.

The parameters in this simulation have the following interpretation:

Unless noted otherwise, each of the experiments in this chapter consists of 4 runs with random

| exper.setup | | | | | |
|----------------------|-------------------|--------------------|--------------|-------------------|-----------|
| numGenerations | 10.000 | recordStart | 4000 | | |
| lowerBoundSeed | 10 | upperBoundSeed | 14 | | |
| model.setup | | | | | |
| numRounds | 1(mBR),1.000(pBR) | auctionType | 1 | priceDet | 0 |
| RWmean | 0 | transactionFee | 0 | numBidders | 2 |
| numSellers | 2 | | | | |
| bidder.setup | | | | | |
| bLearningType | 1(FS), 2(GA) | bSelectionType | 1(QR), 2(BR) | bMemoryStrength | 0 |
| bFixedStrategyMarker | 1 | bFixedBeta | 0.1 ... 0.9 | bFixedValueMarker | 0 |
| bFixedValue | 0 | bOfixedValueMarker | 0 | bOfixedValue | 0 |
| bValDistShape | 1 | bValDet1 | 0 | bValDet2 | 1 |
| bNumStrategies | 999(FS), 100(GA) | bMinStrategy | 0 | bMaxStrategy | 1 |
| bStrategyDistance | 0.0002 | bNumParents | 50 | bNumElite | 20 |
| bCrossOverPar | 0.1 | bMutationProb | 10 | bMutationType | 0 |
| bMutationPar | 0.01 | bXi | 1 | | |
| seller.setup | | | | | |
| sFixedStrategyMarker | 3 | sFixedStrategy | 1 | sFixedValueMarker | 3 |
| sFixedValue | 0 | sNumStrategies | 2 | all others | arbitrary |

Table 10.1: Standard Parameter Settings for the Auction Simulator

seeds 10, 11, 12 and 13 respectively (`lowerBoundSeed=10`, `upperBoundSeed=14`). Each run consists of 6.000 generations after a phase-in of 4.000 generations (`numGenerations=10.000`, `recordStart=4.000`). The simulation records active strategies, values, payoffs and the price for two bidders and sellers each.¹

For the basic simulation setup of myopic best response in 1POBAs with FSs I use the following parameter settings: I consider two bidders (`numBidders=2`) in a 1POBA (`auctionType=1`) where strategies are updated in each round (`numRounds=1`). Each bidder has to pay exactly her bid (`priceDet=0`, `transactionFee=0`). In each round, values are drawn anew (`RWmean=0`) from $U(0,1)$ (`bValDistShape=1`, `bValDet1=0`, `bValDet2=1`) for both bidders (`bFixedValueMarker=0`, `b0FixedValueMarker=0`).

Bidder 0 uses a best response rule (`bSelectionType=2`) for strategy updating. This is simulated via FSs (`bLearningType=1`). Only the last `numRounds` payoffs are considered for strategy updating (`bMemoryStrength=0`), so by setting `bNumRounds=1` I simulate a myopically adapting bidder 0. In contrast, bidder 1 uses a fixed strategy (`bFixedStrategyMarker=1`) with different values for β_1 (`bFixedBeta=0.1, \dots, 0.9`). Each bidder has 999 equidistant strategies (`bNumStrategies=999`) between 0 and 1 (`bMinStrategy=0`, `bMaxStrategy=1`).

On the seller side I instruct both sellers (`numSellers=2`) not to ask anything for the asset (`sFixedValueMarker=3`, `sFixedValue=0`). Since the strategies are unequal to 0 (`sFixedStrategyMarker=3`, `sFixedStrategy=1`), this guarantees an ask of 0 by both sellers.² Therefore, this "double auction" mimics a pure two bidder auction where the seller(s) posts no reserve price. Each bidder is only influenced by the other bidder's bid. To speed up the simulation I equip the sellers only with 2 strategies (`sNumStrategies=2`).

For the simulation of myopic best response play with GAs (`bLearningType=2`) I use the same settings, with a few changes:

I regard it as sufficient to simulate the evolution of a population of 100 strategies (`bNumStrategies=100`). To ensure that successful strategies are not lost, the best 20 strategies always stay in the population (`bNumElite=20`). The best 50 strategies (`bNumParents=50`) – the `numElite` best and the 30 next best – are used as parents to breed the remaining 80 strategies by using the crossover- and the mutation-operator with rank-proportional selection. Strategies are only acknowledged as different if they differ by at least 0.0002 (`bStrategyDistance=0.0002`). This prevents the convergence of the population to a suboptimal state. Crossover is simulated by extending the interval between two strategies by the factor 0.1 to each side (`bCrossOverParameter=0.1`) and choosing a new strategy value with uniform probability from this interval. With a probability of 1% (`bMutationProb=10`), strategies are mutated by up to $\pm 1\%$ (`bMutationPar=0.01`) with uniform probability (`bMutationType=0`).

The implementation of linear myopic quantal response is available only with FSs. The simulation has the same settings as myopic best response with FSs, only that `bSelectionType=2`. The choice of `bXi=1` guarantees that we model linear and not payoff-exponential quantal response play.

For perfect memory play I use the same settings as for the corresponding myopic play. The difference is that now I set `numRounds=1000` and `bMemoryStrength=1`. (Parameter settings for `memoryStrength` only have an effect on simulations with FS). So, the first recorded strategies in generation 4.000 of each run are already based on the information that was obtained in 4 million auctions.

Results:

All results presented in Tables 10.2 - 10.5 are structured in the same format:

The first block shows the fixed strategy of bidder 1 and the theoretical prediction for bidder 0s response. The remaining blocks show the simulation results for FSs and - in the case of best response - for GAs. The first column in each block reports the mean of the 24.000 experimental data points that were recorded during the simulation. The second column reports the deviation between mean and theoretical prediction and the third column the variance of the 4 individual means of the simulation runs.

For second-price auctions I performed the analogous experiments with GAs and FSs. As theoretically predicted, strategies in 2PAs under best response converge to 1, however I have to use perfect memory strength to obtain this result, otherwise the simulation finds the NE strategy

¹For technical reasons, the simulation needs at least two participants on each side of the market.

²If strategies were 0, the seller asks $\frac{0}{0}$ which is indeterminate.

| $\bar{\beta}_0^{\text{mBR}}(\beta_1)$ via FSs and GAs | | | | | | | |
|---|--------------------------------|--------------------------------|--|---|--------------------------------|--|---|
| β_1 | $\beta_0^{\text{th}}(\beta_1)$ | $\beta_0^{\text{FS}}(\beta_1)$ | $\text{dev}(\beta_0^{\text{th}}, \beta_0^{\text{FS}})$ | $\sigma_{\text{mean}}(\beta_0^{\text{FS}})$ | $\beta_0^{\text{GA}}(\beta_1)$ | $\text{dev}(\beta_0^{\text{th}}, \beta_0^{\text{GA}})$ | $\sigma_{\text{mean}}(\beta_0^{\text{GA}})$ |
| 0.1 | 0.14750 | 0.14767 | 0.11% | 0.95% | 0.17254 | 16.97% | 1.56% |
| 0.2 | 0.23438 | 0.23442 | 0.02% | 1.79% | 0.25585 | 9.16% | 0.98% |
| 0.3 | 0.30070 | 0.30334 | 0.88% | 1.36% | 0.31690 | 5.39% | 0.76% |
| 0.4 | 0.35407 | 0.35375 | -0.09% | 1.27% | 0.35994 | 1.66% | 1.07% |
| 0.5 | 0.39772 | 0.39692 | -0.20% | 1.31% | 0.39608 | -0.41% | 0.60% |
| 0.6 | 0.43321 | 0.43041 | -0.65% | 0.97% | 0.41821 | -3.46% | 0.42% |
| 0.7 | 0.46129 | 0.45880 | -0.54% | 0.76% | 0.43185 | -6.38% | 1.56% |
| 0.8 | 0.48210 | 0.47834 | -0.78% | 1.27% | 0.43996 | -8.74% | 1.29% |
| 0.9 | 0.49530 | 0.49320 | -0.42% | 0.97% | 0.44469 | -10.22% | 1.80% |

Table 10.2: Simulation of Myopic Best Response

| $\beta_0^{\text{pBR}}(\beta_1)$ via FSs and GAs | | | | | | | |
|---|--------------------------------|--------------------------------|--|---|--------------------------------|--|---|
| β_1 | $\beta_0^{\text{th}}(\beta_1)$ | $\beta_0^{\text{FS}}(\beta_1)$ | $\text{dev}(\beta_0^{\text{th}}, \beta_0^{\text{FS}})$ | $\sigma_{\text{mean}}(\beta_0^{\text{FS}})$ | $\beta_0^{\text{GA}}(\beta_1)$ | $\text{dev}(\beta_0^{\text{th}}, \beta_0^{\text{GA}})$ | $\sigma_{\text{mean}}(\beta_0^{\text{GA}})$ |
| 0.1 | 0.18231 | 0.18194 | -0.20% | 0.27% | 0.18162 | -0.38% | 0.11% |
| 0.2 | 0.28390 | 0.28389 | 0.00% | 0.21% | 0.28266 | -0.37% | 0.07% |
| 0.3 | 0.36598 | 0.36610 | 0.03% | 0.14% | 0.36271 | -0.35% | 0.08% |
| 0.4 | 0.43685 | 0.43711 | 0.06% | 0.07% | 0.43522 | -0.37% | 0.01% |
| 0.5 | 0.5 | 0.49920 | -0.16% | 0.14% | 0.48905 | -2.19% | 0.08% |
| 0.6 | 0.5 | 0.50081 | 0.16% | 0.20% | 0.49601 | -0.8% | 0.04% |
| 0.7 | 0.5 | 0.49931 | -0.14% | 0.28% | 0.49629 | -0.74% | 0.04% |
| 0.8 | 0.5 | 0.50014 | 0.03% | 0.06% | 0.49552 | -0.90% | 0.08% |
| 0.9 | 0.5 | 0.49946 | -0.11% | 0.44% | 0.49520 | -0.96% | 0.02% |

Table 10.3: Simulation of Perfect Memory Best Response

| $\bar{\beta}_0^{\text{mQR}}(\beta_1)$ via FSs | | | | |
|---|--------------------------------|--------------------------------|--|---|
| β_1 | $\beta_0^{\text{th}}(\beta_1)$ | $\beta_0^{\text{FS}}(\beta_1)$ | $\text{dev}(\beta_0^{\text{th}}, \beta_0^{\text{FS}})$ | $\sigma_{\text{mean}}(\beta_0^{\text{FS}})$ |
| 0.1 | 0.43167 | 0.43243 | 0.18% | 0.79% |
| 0.2 | 0.48959 | 0.48990 | 0.06% | 1.04% |
| 0.3 | 0.53380 | 0.53591 | 0.40% | 0.75% |
| 0.4 | 0.56938 | 0.56938 | 0.07% | 0.77% |
| 0.5 | 0.59848 | 0.59844 | -0.01% | 0.82% |
| 0.6 | 0.62214 | 0.62090 | -0.20% | 0.63% |
| 0.7 | 0.64086 | 0.64018 | -0.11% | 0.83% |
| 0.8 | 0.65473 | 0.65348 | -0.19% | 1.16% |
| 0.9 | 0.66353 | 0.66294 | -0.09% | 0.99% |

Table 10.4: Simulation of Myopic Quantal Response

| $\beta_0^{\text{pQR}}(\beta_1)$ via FSs | | | | |
|---|--------------------------------|--------------------------------|--|---|
| β_1 | $\beta_0^{\text{th}}(\beta_1)$ | $\beta_0^{\text{FS}}(\beta_1)$ | $\text{dev}(\beta_0^{\text{th}}, \beta_0^{\text{FS}})$ | $\sigma_{\text{mean}}(\beta_0^{\text{FS}})$ |
| 0.1 | 0.38518 | 0.38563 | 0.12% | 1.12% |
| 0.2 | 0.42236 | 0.42276 | 0.09% | 0.99% |
| 0.3 | 0.45015 | 0.45049 | 0.08% | 0.93% |
| 0.4 | 0.47039 | 0.47073 | 0.07% | 0.91% |
| 0.5 | 0.48428 | 0.48454 | 0.05% | 0.91% |
| 0.6 | 0.49293 | 0.49318 | 0.05% | 0.89% |
| 0.7 | 0.49756 | 0.49780 | 0.05% | 1.25% |
| 0.8 | 0.49948 | 0.49972 | 0.05% | 0.90% |
| 0.9 | 0.49997 | 0.50020 | 0.05% | 0.90% |

Table 10.5: Simulation of Perfect Memory Quantal Response

of 1 but loses it again.³ For quantal response play in 2PAs, as theoretically predicted, the result depends on the cutoff-level of the strategy set.

Discussion:

For all setups, the AS replicates the analytical predictions of bounded rational play in 1PAs very accurately with FSs. The means of all runs are without exception less than 1% away from the analytic predictions as can be seen from Tables 10.2-10.5. Naturally, perfect memory results are more accurate than the results of myopic play since I used 1.000 times more auctions to assess them.

Since "a picture says more than 1000 words resp. numbers" I graphed the simulation results of Tables 10.2 and 10.3. FSs in comparison to the theoretical predictions for myopic and perfect memory best response. The solid line in Figure 10.1 gives the theoretical prediction of myopic resp. perfect memory best response against a fixed strategy whereas the crosses give the simulation results. As one can easily see, there are virtually no deviations between the line and the crosses. Similarly, also a plot of the results for quantal response play would show virtually no deviations between theoretical predictions and simulation results. With GAs, a look at Table




Figure 10.1: Theory vs Simulation under Best Response with FSs

10.3 shows that the Auction Simulator does a good job for the simulation of perfect memory play. However, its results are persistently a little too low. The most significant underbidding of over 2% occurs if bidder 1 plays the NE strategy $\beta_1 = 0.5$. This can be explained by the concave kink of the best response function $\beta_0^{pBR}(\beta_1)$ that we see in Figure 4.7 respectively the lower part of Figure 10.1.

To understand the underbidding intuitively assume that bidder 1 sets her strategy fixed to the NE value $\beta_1 = \frac{1}{2}$. Bidder 0 does not know this but she selects her strategy as the one that performed best in the last 1.000 rounds. This is about the same as if she would estimate her opponent's bid from the last 1.000 rounds as

$$\bar{\beta}_1 = \sum_{i=1}^{1.000} \frac{b_i}{0.5} \tag{10.1}$$

³In theory, for FS, a positive but arbitrarily small memoryStrength is needed for convergence. In the simulations I had to use perfect memory strength. The reason is the finite working precision of computers: The computer uses internally 8 digits behind the dot. In 2PAs, the strategies 0.999 and 1 give the same result in 999 of 1000 auctions. If I use however for instance a memoryStrength of 0.5, then 27 generations after an auction where the strategy 1 performed better than the strategy 0.999, this difference is forgotten by the computer again (since $\frac{1}{2}^{27} < 10^{-8} < \frac{1}{2}^{26}$). Therefore I set the variable memoryStrength to 1, so that the computer can not forget anything. In the simulation of 1PAs the finite working precision is no problem since in each round each strategy yields different payoff and therefore the effect of past returns is not so important.

and play her perfect memory best response against this estimate according to Figure 4.7. Her estimate $\bar{\beta}_1$ is after each 1.000 rounds approximately normally distributed around the true value of $\frac{1}{2}$. In half of the 1.000 auction-rounds bidder 0 therefore estimates her opponent's strategy too low. In these cases her best response is below the NE. In the cases where she overestimates her opponent's strategy, her best reply is just the NE strategy $\frac{1}{2}$ ⁴. Therefore, bidder 0 underbids on average significantly if her opponent plays the NE strategy $\beta_1 = \frac{1}{2}$. She underbids the less, the narrower the distribution of her estimates is. In the simulation, this can be achieved by increasing `numRounds`.

For lower values of β_1 between 0 and $\frac{1}{2}$, a similar reasoning holds but the underbidding of bidder 0 is less pronounced there since the slope of the best response function is more symmetric to the left and the right of β_1 than at $\beta_1 = \frac{1}{2}$.

For FSs, there occurs no significant underbidding for perfect memory best response. The reason is that for FS simulation I set `memoryStrength=1` which reduces the stochasticity of value realizations: In the end of each of the 4 simulation runs, bidder 0 bases her response on the accumulated information she gets from 10 million independent value realizations.

For the simulation of myopic strategy updating with GAs, Table 10.2 shows systematic deviations between theoretical predictions and simulation results. For low values of β_1 , the GA significantly overbids whereas for high β_1 it significantly underbids. As argued above, this simply reflects the fact that GAs perform poorly in optimizing functions that fluctuate over time.

In conclusion, FS are the better tool to guarantee simulation results that correspond to the mathematical models of bounded rationality. First of all, FS – in contrast to GAs - don't show intertemporal smoothing under myopic play. Second, FS simulations can be influenced by the parameter `memoryStrength`. On the one hand, this means that they don't suffer from underbidding under perfect memory best response because `memoryStrength` can be set to 1. On the other hand, varying the parameter `memoryStrength` provides a possibility to smoothly interpolate between myopic and perfect memory play. The GA lacks this feature because it is not clear, how payoff-information should be transferred between generations of strategies.

The claim that the AS with FS modeling also correctly simulates intermediate memory strengths requires a little explanation: I have shown that myopic play (`memoryStrength= 0, numRounds= 1`) is correctly simulated by the AS. On the other hand, also perfect memory play (`memoryStrength= 1, numRounds= 1`) is simulated correctly⁵. So, since the AS replicates the correct solutions for the extreme settings `memoryStrength= 0` resp. `= 1`, I can be sure that it also delivers the correct results for intermediate memory capabilities like, e.g., `memoryStrength= 0.3`.

For all these reasons, I will stick to FS for most of the remaining thesis. However, this doesn't mean that the implementation of the GA is in vain. As I will discuss at the end of the next chapter, the GA is better suited to locate NEs than FSs.

10.2 Seller Best Response against Fixed Asking Strategies

Hypothesis:

The AS correctly reproduces the predictions of Theorems 7.1 and 7.2 for the behavior of two bounded rational sellers under myopic and perfect memory best response where one seller plays a fixed strategy.

Parameter Settings:

The settings are symmetric to the two bidder case. The differences are shown in table 10.2: The setting of `priceDet= 1` ensures that the winning seller gets the full payoff $PO_i = v_i(\frac{1}{\sigma_i} - 1)$. Each bidder bids a fixed amount of 1.000 (`bFixedStrategyMarker= 3, bFixedBeta= 1.000, bFixedValueMarker= 3, bFixedValue= 1`). Because of `sMinStrategy= 0.001` and $v \sim U(0, 1)$, the sellers' asks are always between 0 and $1/0.001 = 1.000$. Therefore, each ask is absorbed by

⁴Note that in the simulations, responding bidders choose sometimes even strategies slightly above the NE. The reason is that they base their strategies only on finitely many value realizations and not on the mathematical theory. However, the mathematical theory still shows, why underbidding is much more likely than overbidding.

⁵Though I have made the simulations with `numRounds= 1.000` to increase the accuracy of the simulation, analogous results for perfect memory play could also have been obtained by setting `numRounds= 1` if, e.g., I would have simulated 1.000 times more generations.

| | | | | | |
|----------------------|-----|----------------|-------------|-------------------|---|
| model.setup | | | | | |
| priceDet | | 1 | | | |
| bidder.setup | | | | | |
| bFixedStrategyMarker | 3 | bFixedBeta | 1000 | bFixedValueMarker | 3 |
| bFixedValue | 1 | bNumStrategies | 2 | | |
| seller.setup | | | | | |
| sFixedStrategyMarker | 1 | sFixedSigma | 0.1,...,0.9 | | |
| sNumStrategies | 999 | sMinStrategy | 0.001 | | |

Table 10.6: Specific Parameter Settings for Pure Seller Auctions

| $\bar{\sigma}_0^{\text{mBR}}(\sigma_1)$ via FSs | | | | |
|---|----------------------------------|----------------------------------|--|--|
| σ_1 | $\sigma_0^{\text{th}}(\sigma_1)$ | $\sigma_0^{\text{FS}}(\sigma_1)$ | $\text{dev}(\sigma_0^{\text{th}}, \sigma_0^{\text{FS}})$ | $\sigma_{\text{mean}}^2(\sigma_0^{\text{FS}})$ |
| 0.1 | 0.14750 | 0.14744 | -0.04% | 1.56% |
| 0.2 | 0.23438 | 0.23345 | -0.40% | 0.60% |
| 0.3 | 0.30070 | 0.29917 | -0.51% | 1.56% |
| 0.4 | 0.35407 | 0.35197 | -0.59% | 1.50% |
| 0.5 | 0.39772 | 0.39714 | -0.14% | 1.21% |
| 0.6 | 0.43321 | 0.43207 | -0.26% | 1.02% |
| 0.7 | 0.46129 | 0.45818 | -0.67% | 0.82% |
| 0.8 | 0.48210 | 0.47839 | -0.77% | 0.71% |
| 0.9 | 0.49530 | 0.49358 | -0.35% | 0.40% |

Table 10.7: Simulation of Myopic Best Response for Seller

the bidder-side of the market, the only relevant parameter for the updating seller is therefore her opponent's strategy. Again, I keep the strategy of one seller fixed (`sFixedStrategyMarker=1`) at $\sigma_1 = 0.1, \dots, 0.9$ (`sFixedSigma=0.1 \dots 0.9`) while the other seller updates against this strategy. For a sufficiently accurate assessment of the seller's strategy I use 999 possible strategies (`sNumStrategies=999`) while I use only 2 strategies for each bidder in order to save simulation time.

Results:

The structure of Tables 10.7 and 10.8 is the same as in the previous section. I did not make experiments with GAs since, as noted in the previous section, FSs are much better apt for reproducing the analytic results for myopic as well as for perfect memory play.

Discussion:

Again, all the results for FS are well within the 1% range of the theoretical predictions and it can not be rejected that they are equal. The correct simulation of myopic and perfect memory best response for pure seller auctions shows that the AS simulates the seller side correctly. In this section I checked only for best response play since I derived no analytic predictions for quantal response of sellers. But this isn't even necessary to still guarantee the correctness of the

| $\sigma_0^{\text{pBR}}(\sigma_1)$ via FSs | | | | |
|---|----------------------------------|----------------------------------|--|--|
| σ_1 | $\sigma_0^{\text{th}}(\sigma_1)$ | $\sigma_0^{\text{FS}}(\sigma_1)$ | $\text{dev}(\sigma_0^{\text{th}}, \sigma_0^{\text{FS}})$ | $\sigma_{\text{mean}}(\sigma_0^{\text{FS}})$ |
| 0.1 | 0.1250 | 0.1245 | -0.40% | 0.48% |
| 0.2 | 0.2353 | 0.2354 | 0.04% | 0.21% |
| 0.3 | 0.3333 | 0.3330 | -0.10% | 0.48% |
| 0.4 | 0.4211 | 0.4212 | 0.03% | 0.33% |
| 0.5 | 0.5 | 0.4998 | -0.04% | 0.26% |
| 0.6 | 0.5 | 0.5001 | 0.02% | 0.40% |
| 0.7 | 0.5 | 0.5000 | 0.00% | 0.88% |
| 0.8 | 0.5 | 0.4978 | -0.44% | 0.40% |
| 0.9 | 0.5 | 0.5003 | 0.06% | 0.72% |

Table 10.8: Simulation of Perfect Memory Best Response for Seller

simulation results: From the last section we know that quantal response is simulated correctly for the bidder side, so the AS translates correctly payoffs into selection probabilities. Since the same part of the program is responsible for the seller side, the selection probabilities are also determined correctly for the seller side. On the other hand, this section shows that also the payoff-updating routines are correct for the seller side. Payoff-updating and the translation of payoffs into selection probabilities are the only two ingredients for the simulation of quantal response. Since both routines work correctly, quantal response must be simulated correctly for the seller side. An analogous argument as in the previous section shows that the AS correctly simulates imperfect memory seller play because it gives the correct predictions for myopic best response and for perfect memory best response.

10.3 Linear Strategy Best Response in Double Auctions

Hypothesis:

The AS correctly reproduces the predictions of Theorems 8.1, 8.2, 8.3 and 8.4 for double auctions with one bidder and one seller where agents are restricted to linear strategies and one participant plays a fixed strategy. As already noted in the chapter about double auctions, the simulation of linear strategies in double auctions provides additional data points to gauge the correctness of the simulation results in the linear case. However – even under perfect memory play – it is not a realistic description of behavior since the participants’ strategy space does not entail the NE of double auctions.

The more realistic simulation of affine bidding strategies is only possible when the Auction Simulator is extended to allow for nonlinear strategies. This is a project for future research. I will say more on this in the last chapter.

Parameter Settings: The non-obvious parameter settings for double auctions with one bidder who adapts against one seller with fixed strategy σ_1 are given as follows:

| model.setup | | | | | | | |
|-------------|---|------------|---|-------------|---|----------|-----|
| numBidders | 2 | numSellers | 2 | auctionType | 1 | priceDet | 0.5 |

| bidder.setup | | | |
|----------------------|---|------------|---|
| bFixedStrategyMarker | 1 | bFixedBeta | 0 |
| bFixedValueMarker | 0 | | |

| seller.setup | | | |
|----------------------|---|-------------|-------------|
| sFixedStrategyMarker | 3 | sFixedSigma | 0.1, ., 0.9 |
| sFixedValueMarker | 1 | sFixedValue | 1000 |

The setting of `priceDet= 0.5` guarantees that the price is set in the middle between the highest bid and the lowest ask, if a trade occurs. As already noted before, due to technical reasons, the number of bidders and seller in the auction simulation has to be at least two for each side of the market (`numBidders= 2`, `numSellers= 2`). Therefore, I have to keep the second participant on each side of the market "artificially" out of the game. On the bidder side, setting `bFixedStrategyMarker= 1` and `bFixedBeta= 0` lets one bidder bid 0 while the other bidder updates her strategy. On the seller side, both sellers play a fixed strategy (`sFixedStrategyMarker= 3`) between 0.1 and 0.9. However, one seller has a fixed value (`sFixedValueMarker= 1`) of 1.000 (`sFixedValue= 1.000`), so even with a low strategy of 0.1, she asks 100 – more than the bidder could afford because she has a maximal bid of 1. Therefore, this seller gets involved in no trade and effectively only one seller with a fixed strategy and values drawn from $U(0, 1)$ participates in the auction.

In the case of one adapting seller, the pertinent settings for bidder and seller are interchanged.

Results:

The results are given in the well-known format in Tables 10.9-10.12.

Discussion:

The Auction Simulator gives the correct simulation results for double auctions if we confine the strategy space to linear strategies and one participant plays a fixed strategy. This shows

| $\bar{\beta}_0^{\text{mBR}}(\sigma_1)$ via FSs | | | | |
|--|---------------------------------|---------------------------------|--|---|
| σ_1 | $\beta_0^{\text{th}}(\sigma_1)$ | $\beta_0^{\text{FS}}(\sigma_1)$ | $\text{dev}(\beta_0^{\text{th}}, \beta_0^{\text{FS}})$ | $\sigma_{\text{mean}}(\beta_0^{\text{FS}})$ |
| 0.1 | 0.5 | 0.5127 | 2.54% | 1.14% |
| 0.2 | 0.5 | 0.5103 | 2.06% | 2.34% |
| 0.3 | 0.5 | 0.506 | 1.20% | 3.78% |
| 0.4 | 0.5 | 0.4999 | -0.02% | 1.92% |
| 0.5 | 0.5 | 0.4940 | -0.12% | 1.28% |
| 0.6 | 0.5 | 0.5018 | 0.36% | 2.16% |
| 0.7 | 0.5 | 0.4986 | -0.28% | 2.18% |
| 0.8 | 0.5 | 0.4958 | -0.84% | 1.54% |
| 0.9 | 0.5 | 0.4943 | -1.14% | 1.58% |

Table 10.9: Simulation of Myopic Best Response Bidder Adaptation in Double Auction

| $\beta_0^{\text{pBR}}(\sigma_1)$ via FSs | | | | |
|--|---------------------------------|---------------------------------|--|---|
| σ_1 | $\beta_0^{\text{th}}(\sigma_1)$ | $\beta_0^{\text{FS}}(\sigma_1)$ | $\text{dev}(\beta_0^{\text{th}}, \beta_0^{\text{FS}})$ | $\sigma_{\text{mean}}(\beta_0^{\text{FS}})$ |
| 0.1 | 0.6667 | 0.6670 | 0.05% | 0.44% |
| 0.2 | 0.6667 | 0.6677 | 0.16% | 0.48% |
| 0.3 | 0.6667 | 0.6648 | -0.28% | 0.51% |
| 0.4 | 0.6667 | 0.6659 | -0.11% | 0.42% |
| 0.5 | 0.6667 | 0.6681 | 0.22% | 0.41% |
| 0.6 | 0.6667 | 0.6664 | -0.04% | 0.17% |
| 0.7 | 0.6667 | 0.6656 | -0.16% | 0.14% |
| 0.8 | 0.6667 | 0.6665 | -0.02% | 0.21% |
| 0.9 | 0.6667 | 0.6660 | -0.10% | 0.29% |

Table 10.10: Simulation of Perfect Memory Best Response Bidder Adaptation in Double Auction

| $\bar{\sigma}_0^{\text{mBR}}(\beta_1)$ via FSs | | | | |
|--|---------------------------------|---------------------------------|--|--|
| β_1 | $\sigma_0^{\text{th}}(\beta_1)$ | $\sigma_0^{\text{FS}}(\beta_1)$ | $\text{dev}(\sigma_0^{\text{th}}, \sigma_0^{\text{FS}})$ | $\sigma_{\text{mean}}(\sigma_0^{\text{FS}})$ |
| 0.1 | 0.5 | 0.5127 | 2.54% | 1.14% |
| 0.2 | 0.5 | 0.5103 | 2.06% | 2.34% |
| 0.3 | 0.5 | 0.506 | 1.20% | 3.78% |
| 0.4 | 0.5 | 0.4999 | -0.02% | 1.92% |
| 0.5 | 0.5 | 0.4940 | -0.12% | 1.28% |
| 0.6 | 0.5 | 0.5018 | 0.36% | 2.16% |
| 0.7 | 0.5 | 0.4986 | -0.28% | 2.18% |
| 0.8 | 0.5 | 0.4958 | -0.84% | 1.54% |
| 0.9 | 0.5 | 0.4943 | -1.14% | 1.58% |

Table 10.11: Simulation of Myopic Best Response Seller Adaptation in Double Auction

| $\sigma_0^{\text{pBR}}(\beta_1)$ via FSs | | | | |
|--|---------------------------------|---------------------------------|--|--|
| β_1 | $\sigma_0^{\text{th}}(\beta_1)$ | $\sigma_0^{\text{FS}}(\beta_1)$ | $\text{dev}(\sigma_0^{\text{th}}, \sigma_0^{\text{FS}})$ | $\sigma_{\text{mean}}(\sigma_0^{\text{FS}})$ |
| 0.1 | 0.75 | 0.7491 | -0.12% | 0.44% |
| 0.2 | 0.75 | 0.7494 | -0.08% | 0.17% |
| 0.3 | 0.75 | 0.7477 | -0.31% | 0.32% |
| 0.4 | 0.75 | 0.7495 | -0.07% | 0.13% |
| 0.5 | 0.75 | 0.7494 | -0.08% | 0.17% |
| 0.6 | 0.75 | 0.7481 | -0.25% | 0.32% |
| 0.7 | 0.75 | 0.7501 | 0.01% | 0.15% |
| 0.8 | 0.75 | 0.7492 | -0.11% | 0.28% |
| 0.9 | 0.75 | 0.7519 | 0.25% | 0.15% |

Table 10.12: Simulation of Perfect Memory Best Response Adaptation in Double Auction

| $\bar{\beta}_0^{\text{mBR}}(\beta_1, \rho)$ via FSs | | | | | | | | | |
|---|-------------------------------------|-----------------------|--------|-------------------------------------|-----------------------|--------|-------------------------------------|-----------------------|-------|
| β_1 | $\beta_0^{\text{th}}(\beta_1, 0.1)$ | β_0^{FS} | dev | $\beta_0^{\text{th}}(\beta_1, 0.5)$ | β_0^{FS} | dev | $\beta_0^{\text{th}}(\beta_1, 0.9)$ | β_0^{FS} | dev |
| 0.2 | 0.29258 | 0.29265 | 0.08% | 0.69315 | 0.69215 | -0.21% | 0.94825 | 0.94890 | 0.07% |
| 0.5 | 0.42745 | 0.42691 | -0.30% | 0.69315 | 0.69215 | -0.21% | 0.94825 | 0.94890 | 0.07% |
| 0.8 | 0.50615 | 0.50245 | -1.44% | 0.71986 | 0.71850 | -0.26% | 0.94825 | 0.94890 | 0.07% |

Table 10.13: Simulation of Myopic Best Response with Seller Reserve Price

that the AS correctly simulates the interaction between the bidder- and the seller side of the market. Together with the previous sections, we see that the AS would also correctly simulate, e.g., imperfect memory quantal response in double auctions with linear strategies. Note however that the simulation doesn't give the correct results under purely myopic mutual adaptation: Though the myopic best response is on average $\frac{1}{2}$ against each strategy of the opponent, under mutual adaptation the strategies don't converge to $\frac{1}{2}$. The reason is – like in the case of pure bidder auctions – the finiteness of the underlying strategy space. I will comment on this effect in more detail below in section 10.6.

10.4 Myopic Best Response with Reserve Price

Hypothesis:

The AS correctly reproduces the predictions of Theorem 4.2 for myopic best response of one bidder to another bidder who bids according to a fixed strategy β_1 and a seller who posts a fixed reserve price ρ as given by Theorem 4.2.

Parameter Settings:

The only difference to the pertinent settings without reserve price is that instead of setting `sFixedValue= 0`, I set `sFixedValue=0.1, 0.5, 0.9`.

Results:

The first column in Table 10.13 indicates the fixed strategy of bidder 1. The remainder of the table is segmented into three identical blocks for the reserve prices 0.2, 0.5 and 0.8. In each block the first column reports the theoretical value, the second column the simulation result with FS and the third column the deviation between the two.

Discussion:

The results show again – like in the case of double auctions – that the interaction between bidders and sellers is correctly simulated by the AS. Therefore, the AS is a solid basis for an extension to investigate the mutual interaction of several bidders and sellers in the future. Note that the results are exactly the same for a given value of ρ as long as $\beta_1 \leq \rho$. As explained in section 4.2.3 this results from the fact that for $\beta_1 \leq \rho$ bidder 1s bid is always below the seller reservation price and therefore irrelevant for bidder 0s considerations.

10.5 Strategy Distribution under Myopic Best Response

Hypothesis:

The AS correctly reproduces the predictions of Theorem 4.3 for the strategy distribution of a myopic best response bidder 0 against a fixed strategy of bidder 1.

Parameter Settings:

The parameter settings are the same as in the first simulation that are given in Table 10.1.

Results:

Figure 10.2 shows the frequency histogram of bidder 0's strategies if she plays her myopic best response against a fixed strategy $\beta_1 = 0.3$ of bidder 1.

Discussion:

Though I did not prove it by statistical methods, it is obvious that under myopic best response,

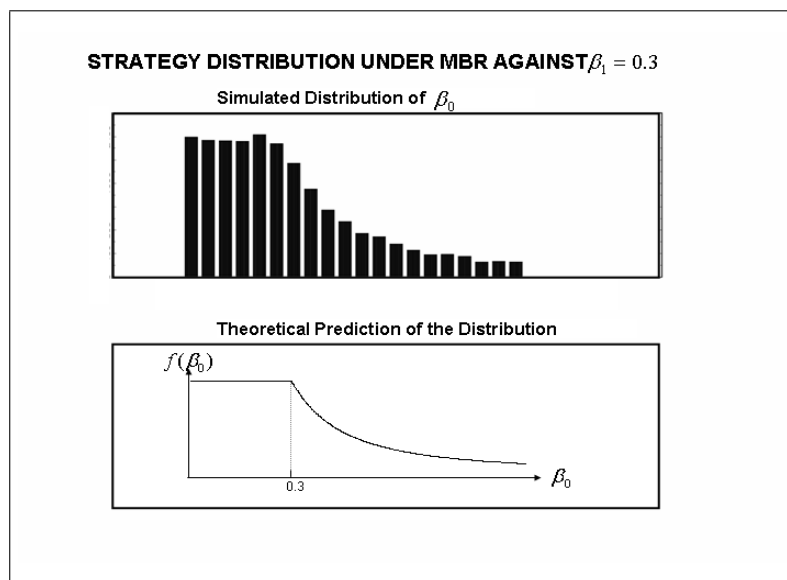


Figure 10.2: Strategy Distribution under mBR against $\beta_1 = 0.3$

the simulation results for the strategy histogram $f(\beta_0)$ against a fixed strategy of bidder 1 fits the theoretical prediction. This result gives additional proof for the reliability of the simulation results on the micro-level. While the previous sections showed that means are simulated correctly, this section shows that also the distribution of strategies around these means is simulated correctly.

10.6 Convergence under Mutual Best Response Adaptation

Hypothesis:

The AS correctly reproduces the predictions of Theorems 4.4 and 4.6 for two mutually adapting bidders under myopic and perfect memory best response.

Parameter Settings:

The settings are analogous to adaptation against a fixed strategy only that

```

bidder.setup
bFixedStrategyMarker 0

```

This guarantees that both bidders update their strategies.

Results:

For myopic best response, Table 10.14 reports the simulation outcome for different settings of `bNumStrategies`. The first column gives `numStrategies`, the second column the simulation result. The third column gives the theoretical prediction of equation (4.49) for the average strategy of the bidders when they have a lowest strategy of $\epsilon = 1/\text{bNumStrategies}$. The fourth column gives the deviation between the theoretical and the simulated value. To visualize the time series properties of mutually adapting myopic best response bidders, Figure 10.3 shows the strategies of both bidders over 100 time steps in the middle of the simulation. The solid line is bidder 0's strategy while the dashed line shows bidder 1's strategy.

For perfect memory best response, the upper part of Figure 10.4 shows the development of β_0 during the first 100.000 rounds if `bNumRounds=1` and `bMemoryStrength=1`. I set here `numRounds=1` – instead of `numRounds=1.000` as in the previous sections – to show the pure effect of perfect memory best response without using fictitious play truncation.

| $\bar{\beta}_0^{\text{mBR}}$ via FS | | | |
|-------------------------------------|-----------------------|-----------------------|--|
| bNumStrat | β_0^{FS} | β_0^{th} | $\text{dev}(\beta_0^{\text{FS}}, \beta_0^{\text{th}})$ |
| 9 | 0.3009 | 0.2996 | 3% |
| 99 | 0.1194 | 0.112 | 6% |
| 999 | 0.0608 | 0.0573 | 6% |
| 9999 | 0.0387 | 0.0348 | 10% |
| 49999 | 0.0282 | 0.0262 | 7% |

Table 10.14: Simulation of Myopic Best Response under Mutual Adaptation

The lower part of Figure 10.4 shows both bidders' strategies if they follow not perfect memory play put truncated Fictitious play. I show here 100 strategy updatings in the middle of a simulation if updating occurs every 1.000 auctions (`numRounds`= 1.000 and `memoryStrength`= 0). Again, the solid line shows bidder 0's strategy and the dashed line bidder 1's strategy. To get a more accurate assessment of the convergence of perfect memory play if more information is aggregated, I left `numRounds`=1.000 but set `memoryStrength`= 1. Then I simulated 24.000 generations where recording started in generation 1. Table 10.15 reports the strategy averages of the first to the fourth quarter of the simulation.

Discussion:

Table 10.14 shows that in simulations of myopic best response, strategies in general do *not* converge to 0 under mutual adaptation. This is in contradiction to the prediction of Theorem 4.4 that predicts the convergence to a steady state where bidders do not bid anything at all.

However, the means of simulations with increasing numbers of strategies – given in the second column of Table 10.14 – show that this is a finite size effect. As the number of strategies goes to infinity, the strategies converge to 0 under mutual adaptation. The underlying reason for this is that in FS simulations with finitely many strategies there exists a lowest strategy β^{\min} that is unequal to 0. In the case of 1.000 strategies this would be $\beta^{\min} = 1/1.000$. Under myopic best response, updating bidders change their strategy according to

$$\beta^{(t+1)} = nx\left(\frac{\beta_{-i}^{(t)} v_{-i}^{(t)}}{v_i^{(t)}}\right) \quad (10.2)$$

where $nx(arg)$ denotes the strategy in the strategy set of bidder i that is next highest to the argument⁶. Therefore, a bidder who would be led by myopic best response to play a strategy between the lowest strategy and 0 overbids significantly in the simulation. For higher strategies of the bidder, the effect is still there, however it gets ever less significant, as the percentage of overbidding shrinks with increasing strategy. At the lower end of the strategy set, however, this effect constitutes a natural barrier against the convergence to 0:

Even if bidder 1 plays her lowest strategy, bidder 0 has a fixed positive probability that her next strategy is arbitrarily high. The probability to exceed a given strategy is given by

$$P(\beta_i^{(t+1)} > \tilde{\beta}) = P\left(\frac{\beta_{\min} v_{-i}^{(t)}}{v_0} > \tilde{\beta}\right) = \frac{\beta_{\min}}{2\tilde{\beta}} \quad (10.3)$$

and does not decrease over time. So, even if both bidders play the lowest strategies in their strategy set, there is a positive nonvanishing probability that at least one bidder breaks away and plays a much higher strategy in the next time step. Once, the strategy of one bidder has jumped up by chance, the other bidder will tend to use a high strategy after the next updating. The bidder who plays the high strategy at time t will in general jump down to a low level again since she uses the old low strategy of her opponent in the calculation of her new strategy. So bidders' strategies fall in mutual oscillations that eventually die out again – as derived in chapter 4 – because the long-term dynamics pushes strategies towards 0. Figure 10.3 shows these intermittent bursts of strategy oscillations for bidder 0 – solid line – and bidder 1 – dashed line – in the middle of a simulation run. Clearly, this gives rise also to clustered volatility in

⁶If no strategy could have won the auction, bidders keep using their old strategy. However, this rule is not of much interest here, since I want to assess why bidders loose their strategy combination of bidding their lowest strategy. If the opponent uses a very low strategy, she bids low and therefore the opponent has only very seldom a value below the opponent's bid. Therefore, the rule is applied very seldom.

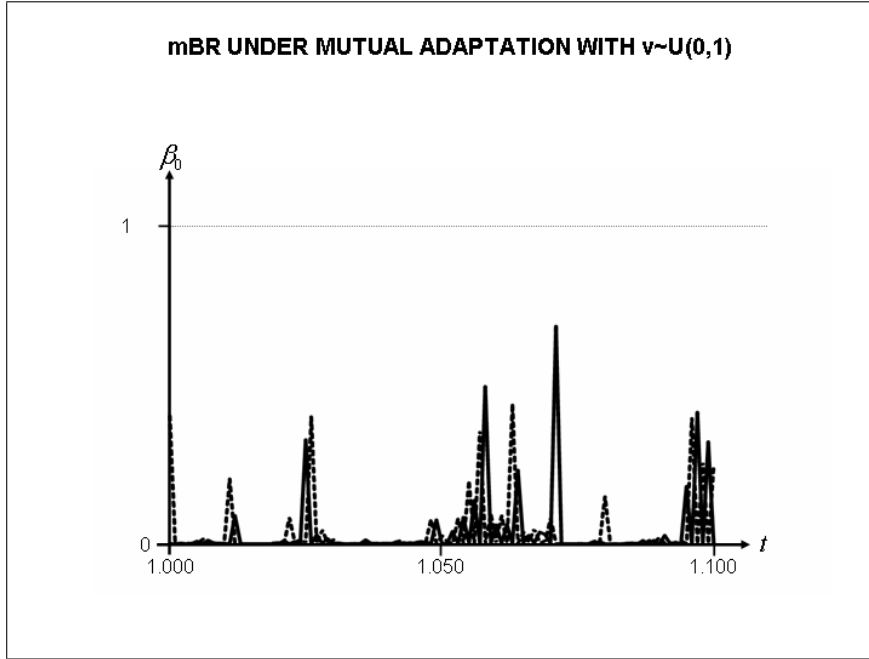


Figure 10.3: Myopic Best Response under Mutual Adaptation with $v_i \sim U(0, 1)$

seller revenues, i.e., the asset prices and returns. As `numStrategies` increases, the bursts occur less and less often and the strategies are most of the time at low levels.

The convergence rate of the average strategy towards 0 with increasing `numStrategies` in Table 10.14 is quantitatively in accordance with the predictions of equation (4.49). The third and the fourth column of Table 10.14 show the theoretical prediction of equation (4.49) for a minimum strategy of $\epsilon = 1/\text{numStrategies}$ and its deviation to the simulation results. The mean strategies are for all investigated cases within 10% of the theoretical prediction of equation 4.49. The difference between the simulation means and the prediction of equation (4.49) arises since the theory assumes a continuum of feasible strategies between the lowest strategy and 1 whereas the simulation also quantizes strategies in this interval. Therefore, also for higher strategies there occurs a certain overbidding which shifts the mean strategies in the simulation above the theoretical prediction.

Figure 10.4 shows the average strategy of bidder 0 under perfect memory best response (fictitious play) and truncated fictitious play.

The upper part shows that under perfect memory best response (`numRounds`= 1 and `memoryStrength`= 1), strategies converge towards the NE. However, in accordance with Theorem 4.6 bidders need infinitely long to arrive in the NE.

As discussed in the mathematical part of the thesis, the underlying reason is that with a setting of `memoryStrength`= 1, bidders never forget anything. A bidder who wants to estimate her opponent's strategy in round t would calculate

$$\bar{\beta}_{-i} = \sum_{j=1}^t \frac{\beta_{-i}^{(j)} v_{-i}^{(j)}}{0.5}. \quad (10.4)$$

Because of the convergence of myopic play towards 0, strategies in the beginning of the simulation are relatively low. The longer the simulation runs, the better the strategy updating of each bidder is described by perfect memory best response against the estimate of the opponent's strategy. This updating is quantified by Theorem 4.5. In particular, when bidders estimate their opponent's strategy as $\bar{\beta}_{-i} < 0.5$, their best response satisfies $\bar{\beta}_{-i} < \beta_{-i}^{\text{pBR}}(\bar{\beta}_{-i}) < 0.5$. Therefore, strategies increase over time. However, the convergence towards the NE is extremely slow since strategies below the NE also induce lower strategies in later rounds.

To assess the long term behavior of perfect memory best response, I set `numRounds`= 1.000 – instead of `numRounds`= 1 – while still keeping `memoryStrength`= 1. Then I simulated 24.000

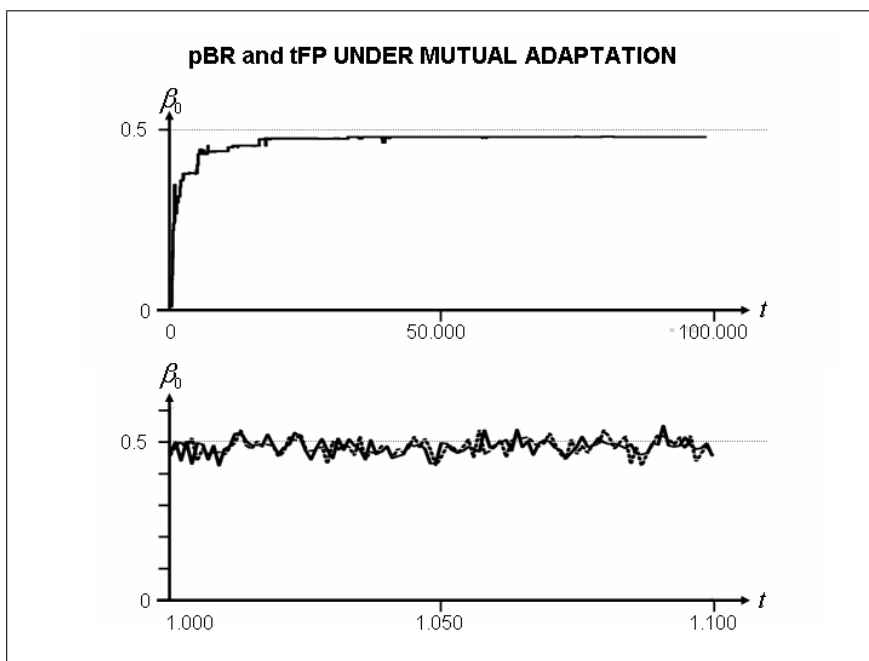


Figure 10.4: Perfect Memory Best Response and truncated Fictitious Play under Mutual Adaptation

generations where recording started in round 1. Table 10.15 shows that the average strategy converges towards the NE. However, even in the fourth quarter of the simulation where bidders based their strategy updating on the results of 18 to 24 *million* auctions, strategies are still some 3% below the NE. This shows that under any realistic assumptions on the frequency of auctions and on the memory capabilities of participants we have to assume significant deviations from NE play in first-price auctions.

In many games, perfect memory best response doesn't converge towards the NE because players can never forget the information from the beginning of the game. An alternative model that often still converges towards the NE is truncated fictitious play. There, players forget the information from previous stages of the game and can thereby more accurately assess the strategies that their opponents currently use. Theorem 4.6 predicts that in first-price auctions, also truncated fictitious play doesn't converge but instead leads to permanent strategy fluctuations below the NE.

The underlying reason here is that a finite number of rounds is not sufficient to assess the opponent's strategy arbitrarily accurate. Even if the opponent would play the NE strategy, the best response would be below the NE half of the times. However, under mutual adaptation, if a bidder plays a strategy below the NE, this increases the chance that her opponent updates to a strategy below the NE even above one half. This mutual pushing to lower strategies eventually comes to a halt because there also is an increasing trend according to Theorem 4.5. In the long term, these two trends interact and lead to a stable strategy distribution with a mean below the NE. The lower part of Figure 10.4 shows the simulation result for `numRounds=1.000` and `memoryStrength=0`. Player 0's strategy is given by the solid line whereas bidder 1's strategy is given by the dashed line. Both strategies show persistent fluctuations around a mean that is well below the NE of 0.5. Moreover we see that the strategies are correlated. They involve in simultaneous up- and down movements. This shows that the main determinant of each bidder's strategy is the opponent's strategy and not the stochasticity of the value realizations. I will say more on truncated fictitious play in the next chapter.

$$\beta_0^{\text{pBR}} \text{ via FS}$$

| generation | β_0^{mean} |
|------------|-------------------------|
| 0-6k | 0.4574 |
| 6-12k | 0.4771 |
| 12k-18k | 0.4818 |
| 18k-24k | 0.4842 |

Table 10.15: Simulation of Perfect Memory Best Response under Mutual Adaptation

$$\bar{\beta}_0^{\text{mQR}}(\beta_1, \xi)$$

| β_1 | $\beta_0^{\text{th}}(\beta_1, 1)$ | $\beta_0^{\text{FS}}(\beta_1, 1)$ | $\beta_0^{\text{th}}(\beta_1, 10)$ | $\beta_0^{\text{FS}}(\beta_1, 10)$ | $\beta_0^{\text{th}}(\beta_1, 100)$ | $\beta_0^{\text{FS}}(\beta_1, 100)$ | $\beta_0^{\text{mBR}}(\beta_1)$ |
|-----------|-----------------------------------|-----------------------------------|------------------------------------|------------------------------------|-------------------------------------|-------------------------------------|---------------------------------|
| 0.2 | 0.48959 | 0.48990 | 0.29818 | 0.29816 | 0.24189 | 0.24113 | 0.23438 |
| 0.5 | 0.59848 | 0.59844 | 0.44791 | 0.44710 | 0.40362 | 0.40189 | 0.39772 |
| 0.8 | 0.65473 | 0.65349 | 0.52525 | 0.52188 | 0.48717 | 0.48273 | 0.48210 |

Table 10.16: Simulation of Myopic Quantal Response with Variable ξ

10.7 Payoff-Exponential Quantal Response

Hypothesis:

The AS correctly replicates the predictions of Theorem 6.2 for myopic payoff-exponential quantal response against a fixed strategy for all ξ . For perfect memory quantal response, the AS correctly interpolates between the predictions of Theorem 6.3 for $\xi = 1$ and Theorem 4.6 for $\xi \rightarrow \infty$.

Parameter Settings:

The settings are the same as for myopic and perfect memory quantal response, only that bXi varies between 1 and 100.

Results:

Table 10.16 shows the simulation results for myopic payoff-exponential quantal response. The rows differ by bidder 1's strategy which is given in the first column. The three blocks in the middle show the theoretical prediction and the simulation result for different values of ξ . The rightmost column gives the theoretical best response value that should be approached for $\xi \rightarrow \infty$. In table 10.17 I show exemplarily the results for the simulation of perfect memory payoff-exponential quantal response with a fixed strategy of $\beta_1 = 0.2$ and varying levels of ξ . The second column shows the theoretical result for $\xi = 1$ and the rightmost column shows the best response result that should be obtained for $\xi \rightarrow \infty$.

Discussion:

We see that for varying ξ the AS gives results that are consistent with the theory. For myopic quantal response, the results are in line with the theoretical predictions for all ξ . For perfect memory play, where no general closed-form solutions were obtainable for intermediate levels of ξ , the simulation replicates the theoretical prediction of $\xi = 1$. With increasing ξ it decreases from there monotonically and converges towards $\beta_0^{\text{pBR}}(\beta_1)$, the theoretical value for $\xi \rightarrow \infty$. Therefore, the results of this section show that payoff-exponential quantal response is simulated correctly by the AS.

$$\beta_0^{\text{pQR}}(0.2, \xi)$$

| β_1 | $\beta_0^{\text{th}}(\beta_1, 1)$ | $\beta_0^{\text{FS}}(\beta_1, 1)$ | $\beta_0^{\text{FS}}(\beta_1, 5)$ | $\beta_0^{\text{FS}}(\beta_1, 10)$ | $\beta_0^{\text{FS}}(\beta_1, 50)$ | $\beta_0^{\text{pBR}}(\beta_1)$ |
|-----------|-----------------------------------|-----------------------------------|-----------------------------------|------------------------------------|------------------------------------|---------------------------------|
| 0.2 | 0.42236 | 0.42276 | 0.34081 | 0.31758 | 0.29143 | 0.28390 |

Table 10.17: Simulation of Perfect Memory Quantal Response with Variable ξ

10.8 Comments on Chapter 5

In the thesis, I do not replicate any results from chapter 5. I want to shortly argue, why: I can not double-check the results from section 5.1 since the AS can only simulate linear strategies. The prediction about the shape for the bidding function can therefore currently not be checked. I plan to do this, when I have extended the AS for the simulation of non-linear bidding functions. I will say more on this in the last chapter.

Section 5.2 dealt with a very specific model that bidders used to estimate the opponent's strategy from her bid. It works in the specified form if there are only two bidders and both are perfectly sure that values are uniformly distributed. These assumption are in contrast to the usual assumption from agent based modeling that bidders react to their environment without knowing the determinants of this environment. Moreover, the results would have to be calculated anew for each different number of bidders and value distribution. Therefore I refrained from implementing this estimation procedure in the AS.

Section 5.3 considered the differences between open and sealed bid auctions. Since I can not decide on a canonical model for the winner's strategy-lowering rule, I refrain at the current level from implementing such a specific rule in the AS. However, it might be an interesting route for further research to experimentally determine how winning bidders update their strategies in sealed bid auctions. Then, I would aim to analyze the effects of this mechanism mathematically and to implement it into the AS.

10.9 Conclusions from this Chapter

In this chapter I have shown that the AS correctly reproduces the mathematical results that I derived in the first part of the thesis: The first two sections show that both, the bidder and the seller side are simulated correctly. The correct results for myopic and perfect memory play strongly suggest that also for intermediate memory strengths, the setups are simulated correctly. The third and the fourth section indicate that also the interaction between bidders and sellers is simulated correctly. Sections 5 and 6 show that not only against fixed strategies but also under mutual adaptation, the results are in general correct. And section seven shows that the AS correctly simulates QR with arbitrary ξ .

I draw two main conclusions from this:

First of all, the mathematical analysis is correct. Wherever mathematical predictions were not matched by simulation results, I was able to explain this by the particularities of the simulation technique. As a particularly interesting route for further mathematical research, the Θ -function-based formalism that I developed for the analysis of perfect memory best response seems to be a reliable tool to assess the phase space dynamics around the NE. Therefore I want in future research to extend the methodology to investigate in more general auction setups, whether NEs are attracting, repelling or show saddle path behavior.

Second, the AS seems to be free of bugs. Especially with FSs, it is a reliable tool to replicate analytical results. The setups that I investigated can be characterized along the following dimensions:

- **Memory Strength:** Myopic vs perfect memory play (interpolation via λ or `numRounds`)
- **Learning Model:** Best response vs quantal response (interpolation via ξ)
- **Participants:** Bidder and Seller Side of a market; with or without interaction (setup via `numBidders` resp. `numSellers`, `priceDet`, respective fixed strategies and/or values)

One can think of this chapter as building a hypercube in the space of auction setups that are reliably assessable with the AS. I rigorously check the vertices of the hypercube and then claim that, if the results in the vertices are correct, also the whole space of auction setups within the hypercube that is spanned by the vertices is simulated correctly. An example for this rather abstract claim:

Myopic and perfect memory best response are simulated correctly, therefore I claim that also

imperfect memory play and truncated fictitious play are simulated correctly as they are interpolations between these two extremes. On the other hand, the parameter ξ interpolates correctly between best and linear quantal response for myopic and perfect memory play. Therefore I assume that the AS also correctly simulates all possible combinations of λ and ξ – any kind of imperfect memory payoff-exponential quantal response. In this sense I can guarantee that – within its functionality – the AS gives for very general auction setups the results that would also be obtained by the according rigorous mathematical analysis - though the analysis is too tedious.

In the following chapter I will follow this route by changing several assumptions and looking at characteristics of the generated time series.

Where do I have to be careful in interpreting the results of the AS?

First, the AS has problems in simulating 2PAs correctly. However, since I know that the error only comes from the finite computation accuracy, I can overcome this problem by setting the memory strength in these simulations equal to 1. Still, the same problem could distort the results in the simulation of 3rd and higher price auctions under imperfect memory best response.⁷

Second, the simulation of myopic bidders in 1PAs under mutual adaptation, is biased towards higher strategies by the finiteness of the strategy set. This severely distorts the dynamics of the auction. As a counter-measure, I should use in such simulations as many strategies as possible. Note however that the problem of overbidding is especially severe for the particular setup of 2 myopic best response bidders with uniform value distributions, since strategies in this setup tend towards 0. By allowing for positive memory strength, truncated fictitious play, normal value distributions or more than two bidders, the lowest strategies are played less often or not at all. Therefore, the degree of overbidding is much lower in such setups and I can expect simulation results that – largely – correspond to the mathematical predictions.

⁷I think, this error in simulating 3rd and higher price auctions is not too big, given that the error in the simulation of 2PAs is also quite small.

Chapter 11

Computational Analysis of Mutually Adapting Bidders

In the previous chapter I discussed in detail that the AS correctly reproduces the mathematical predictions of the first part of the thesis. This serves as a quality check for the AS. However, the auction setups that I considered there were very restricted: They focused only on first-price auctions with two bidders and values from a uniform distribution. In this chapter I want to use the AS to investigate auction setups where I have no mathematical predictions.

11.1 Auction Efficiency

In the previous chapter, I considered two-bidder 1PAs with uniform value distributions and myopic or perfect memory best response. Under these assumptions, the simulation replicated correctly the analytic predictions. In particular, the simulation showed that even for the case of perfect memory best response, it takes bidders infinitely long to converge to the NE. Under truncated fictitious play, strategies fluctuate below the NE. The mutual adaptation leads to an inefficient allocation since strategies β_i are in general different among bidders and therefore the good is not necessarily allocated to the bidder who values it the most. However, we saw also correlation between the strategies, so, while fluctuating, the strategies stay quite close together and therefore we might suspect that the good is at least in most of the cases obtained by the bidder who values it the most. In this section I want to investigate the efficiency of different auction setups. In particular, I am interested in the dependence of auction efficiency on the memory strength, on the number of bidders and on the auction format.

Varying Memory Strength:

Table 11.1 shows the effect of varying the memory strength λ of two mutually adapting bidders in repeated first-price open bid auctions. I chose λ and not `numRounds` as the interpolation variable because I consider imperfect memory best response as a more realistic model for finite memory strength than truncated fictitious play. However, the dynamics arising from these two interpolation methods should be similar as bidders use information from ever more rounds to update their strategies. The average auction efficiency in the fifth row is determined by calculating the ratio of the values of the winning bidder and the winner with the highest valuation, $\sum v_{\text{win}} / \sum v_{\text{max}}$ where the summation goes over all auctions. The deviation of the average auction efficiency from 1 shows, how much potential surplus is wasted by the allocation mechanism in comparison to the pareto optimal allocation. The last row reports the percentage of auctions in which the bidder with the highest valuation got assigned the asset.

The mathematical analysis predicts convergence of the strategies to 0 for myopic best response ($\lambda = 0$) and to the NE for perfect memory best response ($\lambda = 1$). Intuitively, imperfect memory with $\lambda \in (0, 1)$ should therefore smoothly interpolate between these two extremes and lead to underbidding that is the more severe, the lower λ is. This intuition is verified by the

| | | | | | |
|----------------------------------|------|------|------|------|------|
| Memory Strength λ | 0.1 | 0.3 | 0.5 | 0.7 | 0.9 |
| NE Strategy | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 |
| Average Strategy | 0.11 | 0.16 | 0.21 | 0.27 | 0.36 |
| Deviation from NE Strategy | -78% | -68% | -59% | -47% | -29% |
| Average Auction Efficiency | 87% | 91% | 94% | 96% | 99% |
| Percentage of Efficient Auctions | 67% | 73% | 78% | 84% | 91% |

Table 11.1: Effect of Memory Strength on 1POBAs

| | | | | | |
|----------------------------------|------|------|------|------|------|
| Participating Bidders | 2 | 3 | 5 | 10 | 20 |
| NE Strategy | 0.5 | 0.67 | 0.8 | 0.9 | 0.95 |
| Average Strategy | 0.16 | 0.40 | 0.66 | 0.84 | 0.93 |
| Deviation from NE Strategy | -68% | -41% | -18% | -6% | -2% |
| ρ | 68% | 81% | 71% | 58% | 48% |
| Average Auction Efficiency | 91% | 91% | 95% | 97% | 99% |
| Percentage of Efficient Auctions | 73% | 66% | 65% | 66% | 66% |

Table 11.2: Effect of Number of Bidders on 1POBAs

simulation: All memory strengths $\lambda \in (0, 1)$ lead to average strategies above 0 but below the NE. The average strategy increases monotonically in λ , however, even a unrealistically high memory strength of $\lambda = 0.9$ ¹ still leads to an underbidding of 29% and hence to 29% lower expected revenue for the seller when compared to a second-price auction with bidders who bid the true value.² A plot of the average strategy over the memory strength λ would be increasing in λ where the first derivative is the higher, the higher λ is – i.e., with positive second derivative. For the efficiency we see that both measures, average auction efficiency and percentage of efficient auctions are below 100% for all memory strengths. With increasing memory strength the efficiency increases. This means that the correlation between the two bidders' strategies increases, strategies move closer together. However, the average auction efficiency stays for all imperfect memory strengths below the benchmark of perfect efficiency that would be reached by a 2PA with imperfect memory strength.³ Furthermore we see that efficiency losses occur mainly when bidders' valuations are close to each other since the average auction efficiency is significantly above the percentage of efficient auctions. This tendency gets more significant for higher memory strengths: While the average social loss per inefficient auction is 13%/33% \sim 40% for $\lambda = 0.1$, it falls down to 1%/9% \sim 10% for $\lambda = 0.9$.

More than 2 Bidders:

Table 11.2 reports the impact of increasing the number of participants n in first-price auctions. Memory strength was fixed at 0.3.⁴ The first row reports the number n of participating bidders. The second row shows the NE strategy $\beta_i = (n - 1)/n$. The third row gives the deviation between the average strategy and the NE strategy, $1 - \beta^{\text{NE}}/\bar{\beta}$. The fourth row shows the quotient $\rho = (\beta^{\text{NE}} - \bar{\beta})/(1 - \beta^{\text{NE}})$. The remainder of the table is structured as the previous one.

The results show that first-price auctions fail to converge to the NE also with an increasing number of bidders. However, as the number of bidders increases, the deviations from the NE get rapidly smaller. At first sight, one could therefore conclude that for an increasing number of participants, bidders converge into the NE. This convergence could however be an artefact: Two monotonically increasing series that both converge towards 1 with increasing number of bidders – namely the series of NEs and the series of average strategies – naturally converge in absolute terms to each other. However, it is not clear whether they also move closer to each other

¹From an information theoretic point of view, $\lambda = 0.9$ should roughly correspond to truncated fictitious play with $1 + 0.9 + 0.9^2 + 0.9^3 + \dots = 1 + \frac{1}{1-0.9} = 11$ rounds.

²For low memory strengths, the finiteness of the strategy set could – like under myopic play – lead to overbidding when compared to the analytical predictions. However, also then, analytically the qualitative behavior is the same as in the simulation: Theoretically, average strategies for small values would be even smaller, therefore the increase of the average strategy in the memory strength is even bigger.

³As lined out above, the 2PA simulation fails to capture this convergence due to the finite working precision of computers. Still, mathematically it is clear the also 2PAs with any positive memory strength must in Theory converge to the NE and therefore 2PAs are perfectly efficient. Note that this does not affect the results of 1PAs.

⁴The relatively low value of λ is motivated by experimental observations that participants in repeated games mainly react to the outcome of the last rounds.

| Auction Format | 1 | 2 | 3 | 4 | 5 |
|----------------------------------|------|------|-------|------|------|
| NE Strategy (10 bidders) | 0.9 | 1 | 1.125 | 1.29 | 1.5 |
| Average Strategy | 0.84 | 1 | 1.13 | 1.34 | 1.56 |
| Deviation from NE Strategy | -6% | 0% | +1% | +4% | +4% |
| Average Auction Efficiency | 97% | 100% | 97% | 96% | 97% |
| Percentage of Efficient Auctions | 66% | 100% | 62% | 57% | 60% |

Table 11.3: Effect of Auction Format

in relative terms. The natural quantity to assess this as `numBidders` increases is the quotient $\rho = (\beta^{\text{NE}} - \bar{\beta}) / (1 - \beta^{\text{NE}})$. It compares the distances of the NE and the average strategy to 1 and tells us, by how much further away the average strategy is from 1 than the NE is away from 1. With respect to this quantity, there is no clear convergence of the average strategy towards the NE. Even for 20 bidders, the average strategy has about 50% more distance from 1 than the NE. However, the results might still indicate a very slow convergence towards the NE as ρ falls from $\sim 80\%$ for 3 bidders to $\sim 50\%$ for 20 bidders. Simulations with many more bidders could clarify this point in the future.

From the fifth row we see that the amount of wasted social surplus decreases with increasing number of bidders. This means that bidders's strategies move closer together as the number of bidders increases. This is a simple consequence of the fact that increasing the number of bidders leads each bidder to use a higher strategy since competition gets fiercer. Consequently, the strategies are confined in a smaller interval and therefore closer together.

For more than 2 bidders the probability that the bidder with the highest valuation gets assigned the asset stays essentially at $2/3$. Interestingly, this probability seems also largely independent of the memory strength and the auction format - as long as it is not a 2PA (see also Table 11.3). In the future, it would be interesting to identify the underlying reasons for this constant.

Auction Format:

Table 11.3 reports the outcome of simulations to assess the impact of choosing different auction formats. Again I choose $\lambda = 0.3$ and simulate a fixed number of 10 participating bidders. The first row reports the auction format, ranging from 1st- to 5th-price auctions. The rest of the rows are as in the previous simulations where the NE strategy in a k th-price auction with n participants and uniform value distributions is given by

$$b(v) = \left(1 + \frac{k-2}{n-k+1}\right)v, \quad (11.1)$$

see e.g., Wolfstetter (1996).

While first-price open bid auctions lead to underbidding for finite memory strengths, all k th-price auctions with $k \geq 3$ lead strategies above the NE where the overbidding increases with increasing k . So, a revenue-maximizing seller might prefer k th price auctions with high k . However, the increase in expected seller revenue is accompanied by fluctuating bidding strategies and therefore by riskier seller revenue when compared to the seller revenue in a 2PA. Therefore, the auction format that a risk averse seller prefers depends on her specific risk profile.

In conclusion I arrive at the following conjecture:

Conjecture 11.1 *Among all k th-price pure bidder open or sealed bid auctions, the second-price auction is the unique format that yields an optimal allocation of goods under best response dynamics with imperfect memory.*

11.2 Myopic Best Response with Normal Value Distributions

11.2.1 Influence of the Variance

For uniform value distributions I have derived a variety of analytical results and proved them in the last section. But what can be said about the case of normal value distributions?

| $\beta_0^{\text{mBR}}(\sigma_v^2)$ | | | | | | | |
|------------------------------------|------------------------------|-----------------------------|----------------------------|--------------------------|---------------------------|----------------------------|--------------------------|
| $\beta_0^{\text{FS}}(0)$ | $\beta_0^{\text{FS}}(0.001)$ | $\beta_0^{\text{FS}}(0.01)$ | $\beta_0^{\text{FS}}(0.1)$ | $\beta_0^{\text{FS}}(1)$ | $\beta_0^{\text{FS}}(10)$ | $\beta_0^{\text{FS}}(100)$ | $\beta_0^{\text{FS}}(U)$ |
| 1 | 0.9955 | 0.9565 | 0.7437 | 0.1772 | 0.0799 | 0.0621 | 0.0608 |

Table 11.4: Effect of Value Distribution on Strategy

Simple Gedankenexperimente show that the outcome of IPOBAs must depend critically on the variance σ_v^2 of the value distribution for bidders. For very small standard deviation resp. variance of the normal value distribution, $\bar{\beta}_0^{\text{br}}(\beta_1)$ must for all memory strengths be approximately equal to 1: In the extreme case, if both bidders have the same value each, each bidder does optimal if she bids more than her opponent. Naturally, since feasible strategies are bounded from above by one, the process of mutual best response adaptation therefore increases both bidders bids and leads them in the long term to bid exactly their true value.

For very large standard deviations, on the other hand, results must converge to the results for an underlying uniform distribution: In the simulation, the normal value distribution $N(\mu_v, \sigma_v^2)$ gets cut off at $v = 0$ and $v = 2\mu_v$ to avoid negative values but still guarantee symmetry. The higher σ_v , the flatter the normal distribution in the interval $(0, 2\mu_v)$. If a value is outside $(0, 2\mu_v)$, the value gets repeatedly drawn anew until it lies inside $(0, 2\mu_v)$. For $\sigma \rightarrow \infty$, each of the values on $(0, 2\mu)$ gets drawn with equal probability and therefore the distribution converges for $\sigma_v \rightarrow \infty$ to a uniform distribution $U(0, 2\mu_v)$.

Parameter Settings:

The recording in this simulation started in generation 5.000 to guarantee sufficient time for swinging in and lasted for 24.000 generations thereafter. Parameter settings are as for myopic best response under mutual adaptation in section 10.6 except that

```

bidder.setup
valDistShape 0  valDet1  0.001,...,100  valDet2  100

```

Setting `valDistShape= 0` instructs the AS to draw values from a normal distribution. The mean of this normal distribution is given by `valDet2= 100`. The variance is given by `valDet1` and varies in the different runs between 0.001 and 100. Note that in my notation the variance is given in % of the mean. In this particular simulation this also is the absolute value of the variance since $\mu_v = 100$. Additionally I did one run with the parameter settings (`valDistShape= 1`, `valDet1= 0`, `valDet2= 100`) to simulate a uniform distribution on $(0, 100)$.

Note furthermore, that for normal value distributions the average strategy depends crucially on the density of strategies used in the simulation. If neighboring strategies in the simulation are far apart when compared to the typical difference between two value realizations, the model will fail to simulate the dynamics of mutual strategy adaptation correctly. Therefore, I adapted the value of `bMinStrat` in response to the peakedness of the normal value distribution: The lower σ^2 , the closer the strategies are to 1 and therefore I chose `bMinStrat` near 1 in order to maximize the coverage of the used part of the strategy space. I determined the appropriate value of `bMinStrat` in pre-runs to guarantee that the strategy space was still large enough so that strategies would never leave it.

Results:

Table 11.4 shows the mean strategy of bidder 0 for different variances of the underlying normal distributions (given in % of the mean). The leftmost column denotes the extreme case $\sigma_v = 0$ where both bidders always get the same value. The uniform value distribution in the rightmost column is the other extreme case with $\sigma_v \rightarrow \infty$.

Figure 11.1 shows the strategy histograms of bidder 0 for four different underlying value distributions. The figure should be read row-wise starting from the upper-left corner.

Discussion:

Table 11.4 shows that the Auction Simulator smoothly interpolates between the predictions for mean strategies for a fixed value and uniform value distributions. For the case of an underlying uniform value distribution, as I showed in the last chapter, the theoretical correct value of convergence to 0 is not replicated by the AS because of the finite strategy set. As noted there, this is a problem that influences the results particularly strong if strategies are very low.

A look at Figure 11.1 shows, however, that the results for normal value distributions should be not influenced by the finiteness of the strategy set as long as their variance is sufficiently small. Looking at the upper-left plot, we see that for $v \sim N(100, 0.01)$, strategies never drop below 0.84. Therefore, the overbidding of a player with 1000 strategies is at most $\frac{1/1000}{0.84} \sim 0.12\%$. This shows that for myopic best response with sufficiently narrow normal distributions – as opposed to uniform distributions – the finiteness of the strategy set does not influence the results significantly.

From Figure 11.1 we see that strategy distributions are skewed. For low variances of the

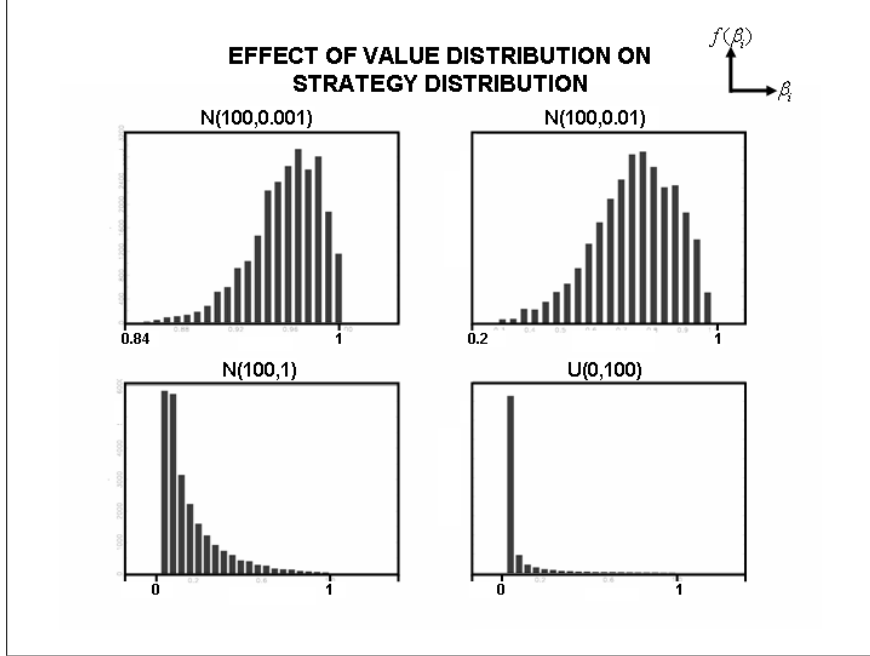


Figure 11.1: Effect of Value Distribution on Strategy Distribution

value distribution, the strategy distribution is left-skewed. As the variance of the normal distribution increases, the strategy distribution changes into an exponential distribution with negative exponent. For uniform distributions and normal distributions that are sufficiently flat so that strategies come arbitrarily close to 0, as I have argued above, the strategy distribution is not simulated correctly, however, the lower right plot shows that the lowest strategy is played most of the time. In the future, it would be interesting to find a one-parameter family of distribution functions that smoothly interpolates between the different strategy distributions by taking the variance σ_v of the underlying value distribution as parameter.

Besides the strategy distribution, another interesting characteristics of the data is shown by a log-log-plot of the absolute seller returns. The returns are calculated in each round as

$$r_t = \frac{PO_t^{\text{sell}} - PO_{t-1}^{\text{sell}}}{PO_{t-1}^{\text{sell}}} \quad (11.2)$$

The plot is made in the following way:

In step 1, all absolute returns are ranked by size: $r^1 > r^2 > \dots$.

In step 2, the logs of the absolute returns are plotted over the logs of their ranks: $f(\ln(\text{rank})) = \ln(r^{\text{rank}})$. This gives a monotonically decreasing plot. If a region within this plot is approximately given by a straight line, this means that within this region, returns are characterized by a power-law.

The log-log-plots in Figure 11.2 show for normal value distributions two different regions.⁵ For a uniform value distribution, such a distinction into two regions does not occur. This indicates, that in the case of normal value distributions, there are two different mechanisms responsible for price changes. One governs the extreme returns that occur in the left part of the plot.

⁵Note that the two regions of the return characteristics only arise if the strategy space is sufficiently dense in the simulation. Otherwise, the simulation fails to capture the two regions.

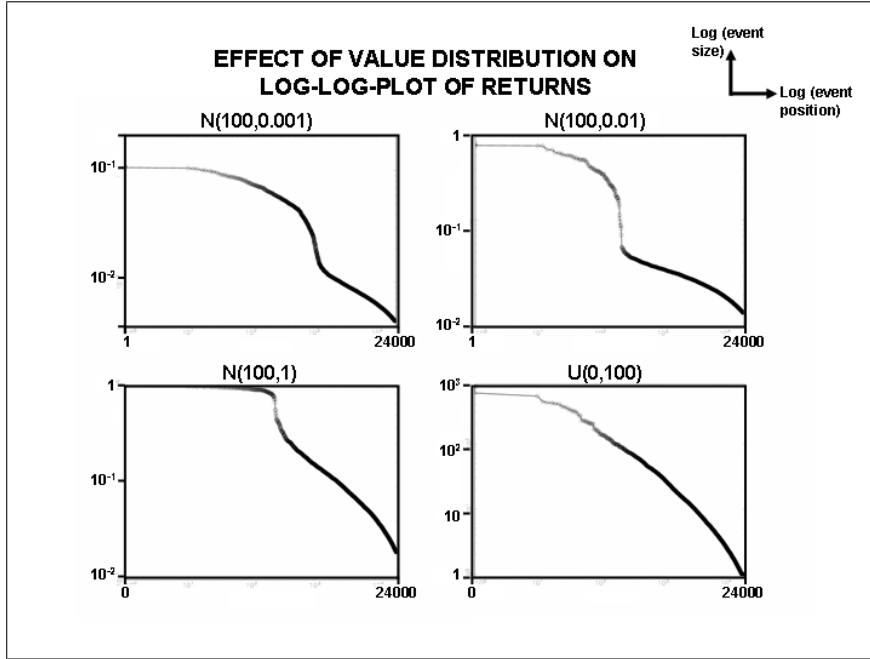


Figure 11.2: Effect of Value Distribution on Log-Log-Plot of Returns

Figure 11.2 shows that these extreme returns are only a fraction of the total sample of returns since the x -axis is logarithmic. The other mechanism governs the more moderate changes that make up the main part of the observations. In the following I investigate the emergence of these two regions in more detail.

11.2.2 Myopic Best Response

Figure 11.3 shows the begin of a simulation with two mutually adapting bidders under best response. In contrast to most other simulations I recorded here the parameters without a phase in, starting at time 0. Each bidder starts with the strategy $\beta_0 = \beta_1 = 0$. Each bidder has 999 strategies between 0 and 1. For the underlying value distribution I chose $\mu_v = 1$ and $\sigma_v^2 = 0.001\%$. The narrowness of this distribution allows me to single out the effects that stem from the mutual interaction of bidders.

We see in the upper two plots of the figure the bidders strategies in the beginning of the simulation. Bidders increase their strategies for the first approximately 2000 generations from 0 to slightly below 1. This increase is accompanied by oscillations with time-varying amplitude – as can be seen by the ‘thickness’ of the strategy plots that changes over time. The price is given by the highest bid. Since the values in each round are very close to 1 – and therefore also close to each other –, in general the bidder with the higher strategy wins the auction. Consequently, the price – shown in the lower left plot – closely follows the highest strategy over time and increases steadily without any jumps or sudden losses.

How can we understand this strategy increase?

Assume – for a simulation with 1.000 strategies between 0 and 1 – that in the beginning both bidders play their lowest strategy $\beta_0^{(0)} = \beta_1^{(0)} = 0.001$. Assume further that $v_0^{(0)} = 0.998$ and $v_1^{(0)} = 1.001$. Then,

$$\beta_0^{(1)} = nx \left(\frac{0.001 \cdot 1.001}{0.998} \right) = nx(0.001003) = 0.002 \quad (11.3)$$

and

$$\beta_1^{(1)} = nx \left(\frac{0.001 \cdot 0.998}{1.001} \right) = nx(0.000997) = 0.001 \quad (11.4)$$

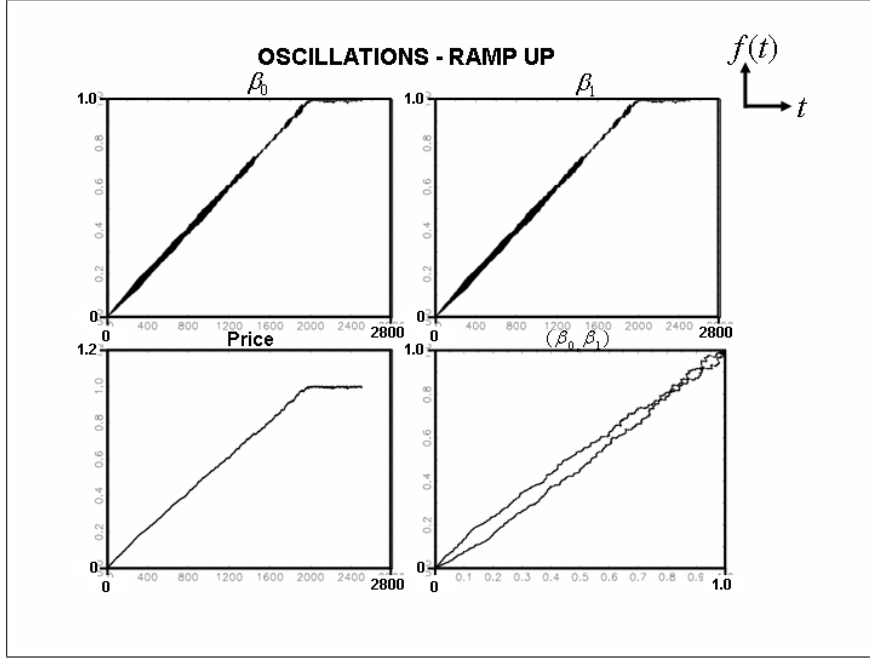


Figure 11.3: Oscillations - Ramp Up

By $nx(\alpha)$ I denote the operator that assigns to any number $\alpha \in (0, 1)$ the next highest strategy in the strategy set. Assume now values for the next round as $v_0^{(1)} = 0.994$ and $v_1^{(1)} = 0.998$. Then,

$$\beta_0^{(2)} = nx\left(\frac{0.001 \cdot 0.998}{0.994}\right) = 0.002, \quad \beta_1^{(2)} = nx\left(\frac{0.002 \cdot 0.994}{0.998}\right) = 0.002 \quad (11.5)$$

If values had been the other way around, $v_0^{(1)} = 0.998$ and $v_1^{(1)} = 0.994$, then

$$\beta_0^{(2)} = nx\left(\frac{0.001 \cdot 0.994}{0.998}\right) = 0.001 \quad \beta_1^{(2)} = nx\left(\frac{0.002 \cdot 0.998}{0.994}\right) = 0.002 \quad (11.6)$$

So, we see the following pattern for the strategies:

$$(0.001, 0.001) \rightarrow (0.002, 0.001) \rightarrow (0.002, 0.002) \text{ resp. } (0.001, 0.002) \quad (11.7)$$

So, on average, strategies increase in the beginning of the simulation run. You might also check this by trying other value realizations. The underlying reason is the finiteness of the strategy grid that leads to overbidding in each time step.

Next I investigate a strategy combination in the middle of the ramp-up:

Assume that $\beta_0^{(25)} = 0.011$ and $\beta_1^{(25)} = 0.008$. Assume for this round value realizations of $v_0^{(25)} = 0.998$, $v_1^{(25)} = 1.002$. Then,

$$\beta_0^{(26)} = nx\left(\frac{0.008 \cdot 1.002}{0.998}\right) = 0.009 \quad \beta_1^{(26)} = nx\left(\frac{0.011 \cdot 0.998}{1.002}\right) = 0.011 \quad (11.8)$$

Conversely, if $v_0^{(25)} = 1.002$ and $v_1^{(25)} = 0.998$ we would obtain

$$\beta_0^{(26)} = nx\left(\frac{0.008 \cdot 0.998}{1.002}\right) = 0.008 \quad \beta_1^{(26)} = nx\left(\frac{0.011 \cdot 1.002}{0.998}\right) = 0.012 \quad (11.9)$$

So we see that bidders' strategies start to oscillate. While bidder 0s strategy was at time 25 above bidder 1s strategy, in time step 26 it is the other way around. The oscillation is largely independent of the value realizations in the respective round, the value realizations mainly influence whether the gap between the high and the low strategy increases or decreases in the next round. Moreover we see that strategies on average continue to increase.

Figure 11.4 shows the first 50 observations of the ramp-up. The first two plots show the strategies of bidder 0 and bidder 1. We see, how in the very first steps the strategies either

increase or stay constant. At around time 5, permanent oscillations set in and their mean increases over time. Comparison between β_0 and β_1 for a specific time step shows that the strategies alternately jump to high and low levels: If the strategy of one bidder is high, the strategy of the other bidder is low and vice versa. The third and the fourth row show the bidders' values for control purposes. As we see, they are purely random and show no signs of oscillations. This highlights graphically that the oscillations are internally generated by the updating rules of the bidders.

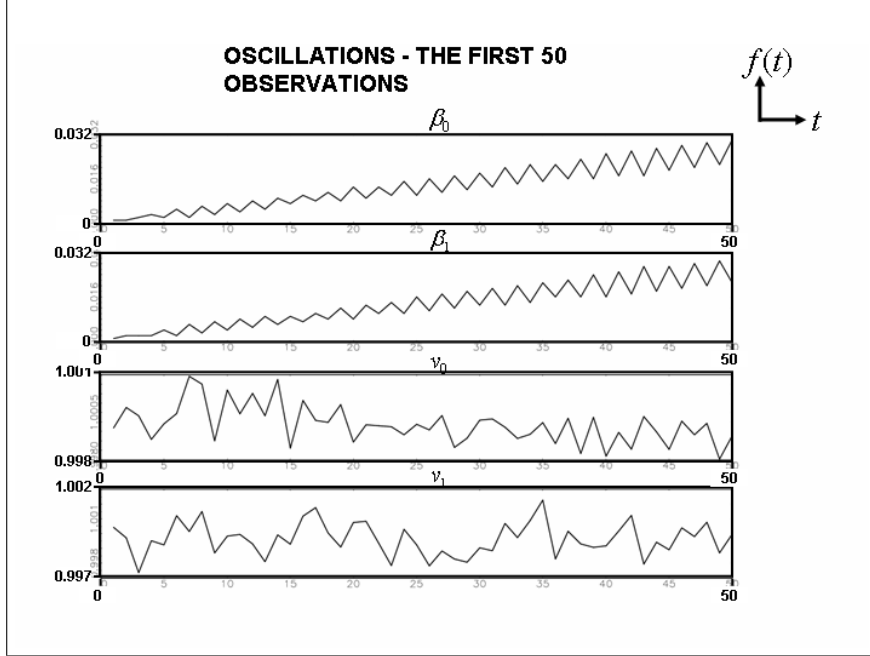


Figure 11.4: Oscillations - The first 50 observations

Fig11.3 shows that at time ~ 2000 the system stops the increase because feasible strategies have to stay below one. Once the value combinations would be such that

$$\text{urenx}\left(\frac{\beta_{-i}^{(t)} \cdot v_{-i}^{(t)}}{v_i^{(t)}}\right) > 1 \quad (11.10)$$

the rule $\beta_i^{(t+1)} = \beta_i^{(t)}$ kicks in and so bidder i keeps her old strategy. Note that the slope of the ramp-up and therefore also the time step at which the kink appears depends on the number of strategies I use in the simulation. For 100 strategies the ramp-up would reach the plateau already after 200 time steps.

The application of the rule $\beta_i^{\text{new}} = \beta_i^{\text{old}}$ changes the behavior of the system dramatically. In Figure 11.5 I show 150 observations between time 3900 and 4050.

Note that scales in this exhibit are much smaller than in Figure 11.4: The fluctuations of strategies and prices in this time window are at the 3% level. I observe only a small window of 150 time steps. At this zoom level we see that both strategies oscillate regularly. Usually, the mean of the oscillations increases over time. However, when strategies have increased by too much and the higher strategy approaches 1, a crash occurs that drops both strategies back to a lower level. From there they start to increase again. Correspondingly we see sharp drops in the time series of prices because the price is mainly determined by the higher of the two strategies. After a crash the price gradually rises again.

To understand the crash behavior better, imagine the following typical sequence taken from the middle of the simulation: In each round, bidders draw random value realizations and calculate

$$\beta_i^{(t+1)} = \text{nx}\left(\frac{\beta_{-i}^{(t)} \cdot v_{-i}^{(t)}}{v_i^{(t)}}\right) \quad (11.11)$$

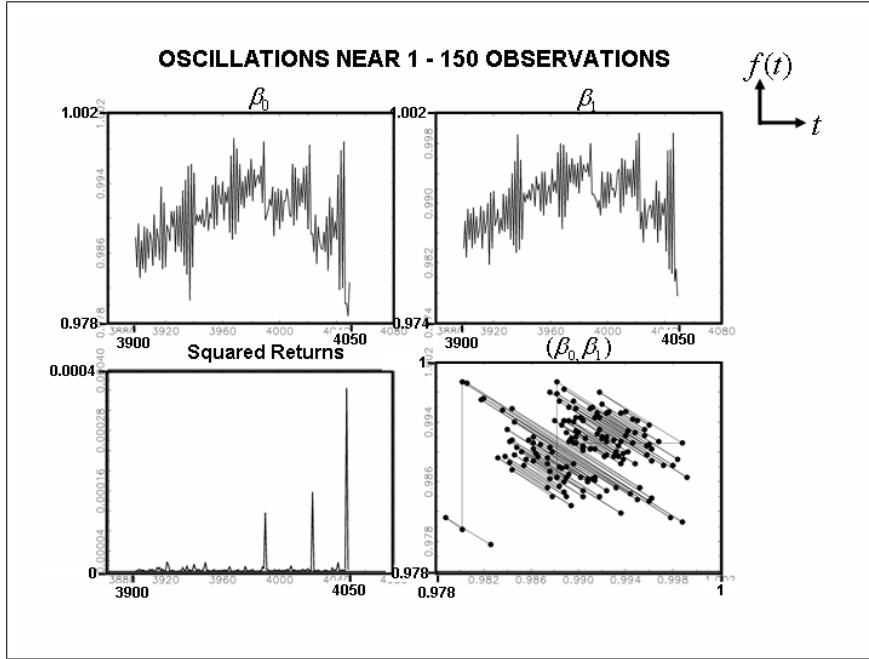


Figure 11.5: Oscillations near 1 - 150 observations

| t | $\beta_0^{(t)}$ | $\beta_1^{(t)}$ | $v_0^{(t)}$ | $v_1^{(t)}$ | $b_0^{(t)}$ | $b_1(t)$ | p | return |
|-----|-----------------|-----------------|-------------|-------------|-------------|----------|--------|---------------|
| 1 | 0.993 | 0.975 | 1.002 | 0.998 | 0.9949 | 0.9370 | 0.9949 | |
| 2 | 0.972 | 0.997 | 1.005 | 1.004 | 0.9786 | 1.0009 | 1.0009 | 0.6% |
| 3 | 0.997 | 0.973 | 0.999 | 1.002 | 0.9960 | 0.9749 | 0.9960 | -0.5% |
| 4 | 0.976 | 0.995 | 0.996 | 1.004 | 0.9720 | 0.9989 | 0.9989 | 0.3% |
| 5 | 0.976 | 0.969 | 0.995 | 1.001 | 0.9711 | 0.9699 | 0.9711 | -2.79% |

Table 11.5: Typical Sequence of mutual Strategy Adaptation under mBR

as next period's strategies. I assume that the bidders have 1000 strategies each, so the strategy space for each bidder is $(0, 0.001, \dots, 0.999)$. Then, a typical set of consecutive data could look as in Table 11.5.

Between time $t = 1$ and $t = 4$ strategies are in the oscillation phase. Next period's strategies are given by $nx(\frac{\beta_{-i}v_{-i}}{v_i}) \sim nx(\beta_{-i})$ since $v_i \sim v_{-i}$ due to the narrowness of the distribution. The price is close to 1. At time $t = 5$, however, bidder 0's strategy would be given by $nx(\frac{0.9989}{0.996}) = nx(1.002992) = 1.003$. Since this strategy is infeasible, player 0 sticks to her old strategy 0.976. Bidder 1, however shows the regular oscillation pattern and drops to 0.969. Consequently, both strategies are now at a low level. The lower right plot in Figure 11.5 shows the scatter plot of strategies where strategy combinations that follow each other in time are connected by lines. From the symmetry around the first median we see that strategies basically exchange their values in each time step. Since strategies increase on average, the points move to the upper right of the plot. Crashes drop them back to a lower level which is visualized by horizontal respective vertical lines.

Since the price follows closely the time series of higher strategies, this behavior induces a sharp price drop and therefore a high negative return: The price is given by the highest bid. In periods of oscillation, the time series of highest bids - basically given by the bids of the bidder who is currently in the top-position - is smoothly increasing. However, if one strategy can not oscillate to a higher level since it would be > 1 , it stays at its previous low level while the other strategy also oscillates to the lower level. So the time series of highest strategies drops suddenly and the consequence is a price crash. The lower left plot in Figure 11.5 shows the corresponding squared returns.

In conclusion, the behavior of mutually adapting myopic best response bidders with normal value distributions depends crucially on the updating rule that bidders use if no feasible bid

could have obtained the item. The special rule that I used in that case was $\beta_0^{\text{new}} = \beta_0^{\text{old}}$. Alternatively, one could imagine several other plausible updating rules. The updating bidder could start bidding higher, $\beta_0^{\text{new}} \sim U(\beta_0^{\text{old}}, 1)$ or even $\beta_0^{\text{new}} = 1$. This would reduce or even eliminate the extreme crashes. Since I see no arguments in favor of any particular updating rule⁶, I refrain from analyzing the alternative updating rules.

11.3 Imperfect Memory Best Response with Normal Value Distributions

In this section I assess the influence of the parameter `memoryStrength` on the outcome of first-price auctions with 2 bidders. For the value distribution I choose $N(1, \sigma_v^2 = 0.1\%)$ ⁷. Figure 11.6 shows the impact of the memory strength on the strategies.

The upper-left plot shows that for myopic play, as discussed in detail above, strategies oscillate,

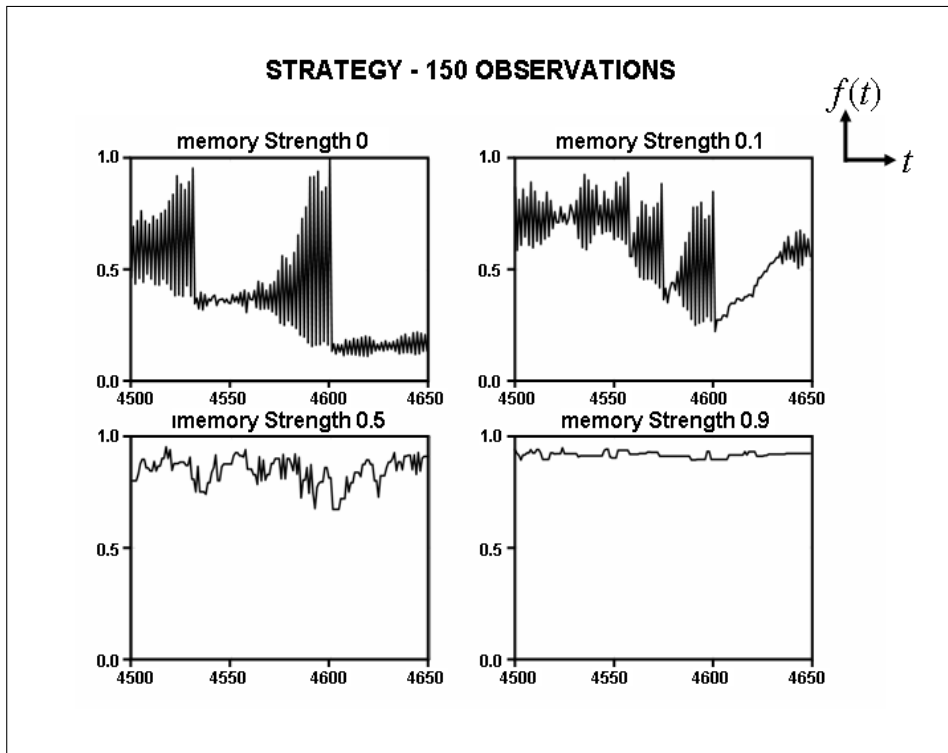


Figure 11.6: Strategy - 150 Observations

while crashing frequently. As the memory strength increases, we see several interrelated effects. For increasing memory strength, also the average strategy increases. This is in principle the same effect that we have already seen in section 11.1 for uniform value distributions: There, increasing the memory strength interpolates smoothly between the underbidding of myopic best response and the (near to) NE-play of perfect memory best response. A similar result holds in the case of a normal value distribution and simulation with linear strategies: Myopic play leads on average to underbidding whereas the linear strategy converges under perfect memory best response, as I will show in section 11.5, to a strategy that is about double as far from 1 as the NE strategy.

With increasing `memoryStrength`, the lock-ins of strategies into oscillations are increasingly often interrupted by phases where they seem to move irregularly. The variance of the strategies' movements over time decreases drastically as the memory strength increases. This is also intuitive since in the extreme case of perfect memory best response (`memoryStrength=1`) strategies

⁶Hints for realistic updating rules could probably be obtained from human experiments

⁷I.e., the standard deviation σ_v of the underlying value distribution is ten times as high as in the detailed analysis of the last section.

would not fluctuate but increase monotonically over time towards 1. For low `memoryStrength` up to ~ 0.6 where strategies still involve at least sometimes in mutual oscillations, the decreased variance of strategies for increasing `memoryStrength` corresponds to a decreased crash size. On the other hand, strategy crashes occur more and more often as `memoryStrength` increases. For sufficiently high `memoryStrength`, like in the lower right plot with `memoryStrength`= 0.9, there are no oscillations any more. Strategies stay constant for relatively long periods. Note however that this is not an artefact of the rule $\beta_i^{\text{new}} = \beta_i^{\text{old}}$ if no feasible strategy could have won: Since both bidders play strategies significantly below 1 and values are very close to 1, bidders' values satisfy in all auctions $v_i^{(t)} > b_{-i}^{(t)}$ and therefore the rule is not applied. The reason for the consecutive time steps of strategies is that as information from more and more auctions is used in the determination of the best response, the influence of one individual auction on the outcome gets less significant. Therefore, even if some feasible strategies could have won the auction, this need not necessarily change the active strategy since the old strategy is still the one that would have maximized the cumulated payoff.

Note that the simulation of imperfect memory play with sufficiently high memory strength, e.g., $\lambda = 0.9$, is robust: The overbidding due to the finite strategy set is of very little consequence since the lowest strategy is ~ 0.9 and therefore simulation with 1.000 strategies leads to very low overbidding. On the other hand, since the bids are significantly below one, also the rule $\beta^{\text{new}} = \beta^{\text{old}}$ is not applied. Therefore, the simulation results in such a setting do not depend on the peculiarities of the simulation but correspond to a large extend to the mathematical results that would be obtained by the corresponding mathematical analysis. Note however that as `numBidders` increases and the auction becomes more competitive, also the strategies increase. Therefore, it might happen that for sufficiently many bidders, the highest bids are often above one and therefore the rule $\beta_0^{\text{new}} = \beta_0^{\text{old}}$ has to be applied. Therefore, simulations of imperfect memory play with many bidders will in general become inrobust.

The phases of constant strategies that we see especially for $\lambda = 0.9$ correspond to an AR-process for the strategies and therefore also for the price: The price shows much longer phases of consecutive time steps above and below the mean than we could expect if it was a random walk. The corresponding autocorrelation function declines very slowly. However currently I did not assess this process in more detail, because the Nash Equilibrium strategies for normal value distributions are non-linear. I will consider the AR-process for normal value distributions in more detail when I have extended the AS for nonlinear strategies. In contrast, in the next section I will consider truncated fictitious play with uniform value distributions. There, the estimation of the AR model has more predictive power because play near the NE can at least be approximated by linear strategies.

The crash characteristics is summarized in a compact way by the return distributions in Figure 11.7.

For low `memoryStrength`, the return distribution is clearly separated into two different parts. The main part of the distribution simply reflects the stochasticity that is introduced into the simulation by the random value realizations. The left cluster reflects the crashes that arise from the updating rule $\beta_0^{\text{new}} = \beta_0^{\text{old}}$ if no feasible strategy would have won. As `memoryStrength` increases, the left tail of the return distribution increases in mass: As `memoryStrength` increases from 0 to 0.1, the left cluster extends over more bins⁸ and also to higher levels within each bin. This shows that the number of crash events increases with the memory strength. On the other hand the crash size decreases in the memory strength so that the left tail moves towards the central part of the return distribution. For sufficiently high memory strengths, the cluster of extreme returns merges into the main part. This is, when the simulation becomes robust and the rule $\beta^{\text{new}} = \beta^{\text{old}}$ is not applied anymore.

11.4 Truncated Fictitious Play with Uniform Value Distributions

From the simulations so far we have seen that the interpretation of simulation results can be problematic. If strategies are close to 0, then the finiteness of the strategy set leads to

⁸A bin in a histogram is the interval on the x -axis for which events are summed up.

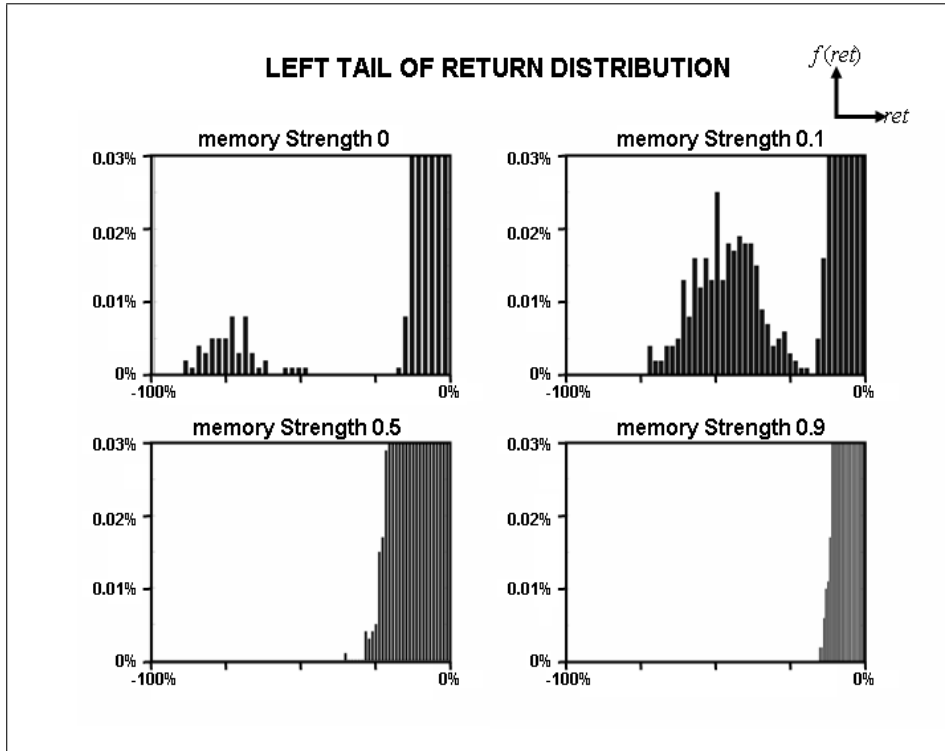


Figure 11.7: Left Tail of Return Distribution

overbidding that fundamentally changes the dynamics of the bidding process when compared to the mathematical predictions. If, on the other hand, strategies are in a round such that the current value couldn't have won the auction, then the results depend on the updating rule that is chosen if the opponent's bid exceeds a bidder's value. This happens especially often when bidders use high strategies.

The simulation results get independent of the number of strategies and also independent of the updating rule, when I investigate bidders under truncated fictitious play with sufficiently high `numRounds`. For the simulations in this section I switch back to a uniform value distribution. Then the NE is given by a linear bidding strategy and the true non-linear strategies under truncated fictitious play get at least approximated near the NE. I consider 10 bidders in a first-price auction and assume that they update every `numRounds`= 100 rounds.

For such a high value of `numRounds`, play is already quite close to the NE which is given in this setting by $\beta_i = 0.9$. Therefore – like in the case of sufficiently high `memoryStrength` – strategies are bounded away from 0. In the particular simulation that I consider here, strategies fluctuate around 0.85. I simulate with 1.000 strategies, so the overbidding is in the range of $0.001/0.85 * 100 = 0.1\%$. Even for the lowest strategy $\beta_i \sim 0.4$ that occurred in the simulation, the overbidding is at most $\sim 0.25\%$. Therefore the finiteness of the strategy grid doesn't influence the simulation results significantly and therefore the results don't depend on the number of strategies chosen as long as the strategy grid is fine enough to accurately simulate the dynamics⁹.

On the other hand, truncated fictitious play also guarantees that the rule $\beta^{(t+1)} = \beta^{(t)}$ does not influence the results: The rule is only applied when extreme value combinations (e.g., high v_1 and low v_0) leave no feasible strategy for the updating bidder. As `numRounds` increases, bidders base their updating decisions on more and more value realizations. Therefore the chance that for a certain bidder, in no round any of her strategies could have won the auction decreases exponentially with `numRounds`. As an example consider a two bidder auction:

I calculate for given β_1 and v_1 the probability that in a specific auction no feasible strategy of

⁹The necessary accuracy of the strategy grid depends only on `numRounds`. The higher `numRounds`, the smaller the strategy changes and the more `numStrategies` we need for an accurate simulation – or at least we need to set `bMinStrategy` and `bMaxStrategy` such that the strategy grid is fine enough in the relevant simulation region around the NE.




Figure 11.8: Truncated Fictitious Play with Uniform Value Distribution

bidder 0 could have won. It is given by $P(v_0 < \beta_1 v_1) = \beta_1 v_1$. Since v_1 comes from a uniform value distribution, the ex ante probability without knowing v_1 is given by

$$P(v_0 < \beta_1 v_1) = \int_0^1 dv_1 \beta_1 v_1 = \frac{\beta_1}{2}. \quad (11.12)$$

Consequently, if updating occurs every n rounds, then the probability that no strategy could ever have won the auction is given by $(\beta_1/2)^n$. So even if the opponent plays her highest feasible strategy, $\beta_1^{(t)} = 1$ and strategies are updated every 10 rounds, the probability that the rule $\beta_0^{(t+1)} = \beta_0^{(t)}$ gets applied in a given time step is only $0.5^{10} \sim 0.1\%$. For `numRounds`= 20, the rule would be applied only in every $1/(0.5)^{20} \sim 1$ millionth time step. If we have more than two bidders participating in the auction, qualitatively the same results hold, though it might be necessary to increase `numRounds` even more to guarantee that the rule $\beta_i^{(t+1)} = \beta_i^{(t)}$ is never applied within a simulation run.

In conclusion, the simulation results of truncated fictitious play with sufficiently high `numRounds` correspond to those that we would also get by a mathematical analysis. The simulation is not influenced by the finiteness of the strategy grid and also not by the specific updating rule if no feasible strategy could have won the auctions.

Figure 11.8 shows typical time series over 200 time steps between 18.500 and 18.700.

The upper left plot shows the strategy of bidder 0. Of course, under mutual adaptation, the strategies of the other bidders look qualitatively the same. As predicted by the mathematical analysis, strategies fluctuate permanently. Moreover we see from the plot that the fluctuations are asymmetric: There are occasional breakouts to extremely low strategies whereas corresponding breakouts to high strategies are not observed. In consequence, the histogram of bidder 0's strategies (not shown) has an asymmetric shape, it is skewed to the left. Table 11.6 quantifies the skewedness: the mean of β_0 is below the median of β_0 – i.e., the strategy value that is obtained by ranking all strategies according to their size and taking exactly the strategy in the middle of this series. This shows that there are less strategies below the mean than above the mean. Therefore, the low strategies show more significant deviations from the mean than those above the mean – the distribution is skewed to the left. The asymmetry of played strategies comes from the shape of the best response function: As `numRounds` $\rightarrow \infty$, bidder 0's strategy would approach her best response to the other bidders' strategies. I conjecture that then – in analogy to the 2-bidder case that I plotted in Figure 4.5, but now for 10 bidders – her best

| | β_0 | β^{best} | β^{max} | β^{av} |
|--------|-----------|-----------------------|----------------------|---------------------|
| mean | 0.858 | 0.849 | 0.907 | 0.858 |
| median | 0.864 | 0.854 | 0.907 | 0.859 |

Table 11.6: Strategy Means and Medians under Truncated Fictitious Play

response against strategies above the NE is the NE whereas her best response against strategies below the NE is also below the NE. Since `numRounds` is finite, the players can not accurately assess their opponents' strategies and therefore play not exactly according to the best response function. However, their reaction on the opponents' play is influenced to a large extent by the shape of the best response function, thus leading them to react to high strategies of their opponents in a much more moderate way than to low strategies.

The upper right plot shows for each time step the strategy of the bidder who has obtained the highest payoff in the `numRounds` rounds. The plot supports the claim that the best response against an environment of mutually adapting bidders is in general not the NE. Again we see visually from the plot or from Table 11.6 that the distribution is skewed to the left. Note that the fluctuations of the best strategies to low strategies are much more modest than the individual strategies. This is intuitively clear: A player with high values throughout the `numRounds` auctions of a certain time step, will play a low strategy in the next round. Such a low strategy occurs then in the time series of the bidder's strategies. However, the chance that she will be the most successful among all the other players with this low strategy in these next `numRounds` auctions – which would put this strategy into the time series of best strategies – is very low, since in these next auctions new value realizations are drawn.

The lower two plots in Figure 11.8 shows the maximum strategy and the average of the strategies used by all bidders in the corresponding time steps. Visually and by comparison with Table 11.6 we see that these strategies are distributed symmetrically around their mean.¹⁰ Note from table 11.6 that the mean of the average strategy is equal to the mean of the individual strategy in the leftmost column. This is no coincidence since bidders are symmetric and in the one case we average over the strategies of one representative player while in the other case we average over the strategies of all players.

The maximum strategy is on average *above* the NE. So, though mathematically the best response function doesn't get higher than the NE¹¹, the finiteness of `numRounds` lets strategies also increase beyond the NE. The average of all strategies played by the bidders, however, is usually below the NE. We know that average seller revenue in a repeated auction is lower than in the NE. But, it is a priori not clear, how important the maximum strategy is in determining the seller revenue when compared to the average strategy: The highest strategy is on average the most likely to win each auction, therefore the maximum strategy is the most important single strategy for determining seller revenue. On the other hand, due to the random value realizations, the highest strategy will not win all of the `numRounds` auctions but all players will win some of the auctions. Therefore, the average strategy is also an important factor in determining the expected seller revenue. The question is, which of the two is more important in determining the seller revenue¹².

To get a feeling for which of the factors is more important for the determination of seller revenue, I computed the effective strategy. This is the strategy that would give the same seller revenue as truncated fictitious play if all bidders would use it constantly as their active strategy. To calculate it I use the seller revenue: Average seller revenue in a first-price auction with 10 bidders where bidders bid according to the NE of $\beta_i = 0.9$ is in the simulation obtained as 0.81838. Average seller revenue in the 1PA with truncated fictitious play is 0.78660. Therefore, the effective strategy of the bidders is given as

$$\beta_i^{\text{eff}} = 0.9 \cdot \frac{0.78660}{0.81838} = 0.86505. \quad (11.13)$$

Comparison with the means in Table 11.6 show that this is much closer to the mean of the average strategy than to the mean maximum strategy. Therefore, the average strategy is the

¹⁰ β_{av} might be slightly skewed to the left. However, this remains to be assessed in longer simulation runs.

¹¹ I assume that this is the case in analogy to the two-bidder case

¹² If, for instance, the average strategy was 5% below the NE, the maximum strategy was 1% above the NE and seller revenue was one 1% below the NE level, then the maximum strategy would be more important in determining the seller revenue.

more important factor for seller revenue than the maximum strategy. The intuition behind this is the following: Bidders' strategies in a given time step typically are between 0.6 and 0.9 whereas values come from $U(0, 1)$. Therefore, the more important determinant of the single bid $b_i(v_i) = \beta_i v_i$ is the bidder's value. So, the bidder with the maximum strategy will not make that many more successful bids than an average bidder and therefore she is not the main driver of seller revenue. An open question is, what would happen if values come from a sufficiently narrow normal distribution. Values would be closer together but strategies as well. It remains a topic for further research to assess whether the effective strategy moves closer to the maximum strategy or whether its position between average and maximum strategy is a constant that is independent of the underlying value distribution. I want to assess this in the future when I have extended the AS for the simulation of nonlinear strategies.

The simulation result for a 2PA with 10 NE bidders is an average seller revenue – cumulated over 100 rounds – of 81.835. This is virtually the same as the 81.838 in the first-price auction in the NE. The equality highlights the validity of the Revenue Equivalence Theorem in this special setting. The numerical value of expected seller revenue is in line with the theoretical prediction, as can be seen by calculating the expected cumulated seller revenue for 100 auctions.: The density of the highest of 10 values from a random distribution is given as

$$dF^n(v) = nF(v)^{n-1}f(v) . \quad (11.14)$$

In the case of $v \sim U(0, 1)$, the CDF is given by $F(v) = v$ and the PDF is given by $f(v) = 1$. Therefore, the expected highest value v^{high} is therefore given as

$$\int_0^1 dv v 10v^9 = \int_0^1 dv 10v^{10} = 10 \frac{1}{11} = \frac{10}{11} \quad (11.15)$$

and the average cumulated seller revenue in 100 auctions where the bidders play according to their NE strategies is given as

$$EPO_{100\text{auctions}}^{\text{sell}} = 100 \cdot 0.9 \cdot \frac{10}{11} = 81.818 \quad (11.16)$$

This is in perfect agreement with the simulation result. The standard deviation of the seller revenue – cumulated over 100 rounds – in the 1PA where bidders use their NE strategies is 0.74 whereas in the second-price auction it is 1.112. This is simply a consequence of the fact that the variance of the second highest of 10 values is higher than the variance of the highest of 10 values. In the literature this observation has been used to claim that risk averse sellers should prefer first-price auctions over second-price auctions.

Under truncated fictitious play, this result doesn't hold anymore. In 1PAs, in addition to the random value realizations, the seller revenue is influenced by the fluctuating strategies. The standard deviation of cumulated seller revenue in the 1PA under truncated fictitious play with `numBidders`= 10 and `numRounds`= 100 is 1.417 which is about 25% above the standard deviation in the second-price auction. In consequence, I would argue that only a sufficiently risk-loving seller should employ first-price auctions to allocate goods. Risk neutral and risk averse sellers would prefer second price auctions since they give more expected revenue and less risk for that revenue.

To assess the dynamics of truncated fictitious play in more detail, I estimate an $AR(p)$ -model¹³ for the average strategy β_{av} . The first step in estimating such a model is to decide on the number of lags of the observable. Allowing for too few parameters might neglect useful information, allowing for too many parameters dilutes the focus on the core drivers of the time evolution. For the model specification I chose p that minimizes the Bayesian Information Criterion (BIC). The BIC is defined as

$$BIC_p = \ln \sigma_p^2 + \frac{\ln T}{T} p \quad (11.17)$$

Here, T is the number of time steps in the sample and p the number of considered lags. The first part in BIC_p gives lower values, the better the goodness of fit is. The second part gets

¹³I do not consider general ARIMA-models since the time series is obviously stationary. Estimation of an $ARMA(1, 1)$ -model gives the same BIC-criterion as an $AR(2)$ -model, so alternatively we could also investigate this model.

the bigger, the more parameters the estimation uses.

Estimation with different settings for p in GAUSS shows that the BIC is minimized for $p = 2$ with a value of -136420 . Therefore I choose the AR-Model

$$\beta_{\text{av}}^{(t)} = \alpha_0 + \alpha_1(\beta_{\text{av}}^{(t-1)} - \alpha_0) + \alpha_2(\beta_{\text{av}}^{(t-2)} - \alpha_0) + \epsilon^{(t)} \quad (11.18)$$

The parameter estimates are given by GAUSS as

$$\alpha_0 = 0.858 \pm 0 \quad \alpha_1 = 0.344 \pm 0.007 (0.330, 0.357) \quad \alpha_2 = -0.030 \pm 0.009 (-0.047, -0.014) ,$$

where the numbers in the bracket give the upper and lower confidence interval at the 2σ -level. All coefficients are significantly different from 0 at the 2σ - level and therefore the $AR(2)$ -equation for the average strategy is given as

$$\beta_{\text{av}}^{(t)} = 0.858 + 0.344(\beta_{\text{av}}^{(t-1)} - 0.858) - 0.03(\beta_{\text{av}}^{(t-2)} - 0.858) + \epsilon^{(t)} \quad (11.19)$$

The equation has the following interpretation: If $\beta_{\text{av}}^{(t-1)}$ was above the long-term average of 0.858, this raises the expected value of the average strategy also at time (t) to about a third of the distance between $\beta_{\text{av}}^{(t-1)}$ and the mean. The distance of $\beta_{\text{av}}^{(t-2)}$ to the mean would via $\beta_{\text{av}}^{(t-1)}$ also influence $\beta_{\text{av}}^{(t)}$'s distance to the mean with a factor of $0.34^2 \sim 0.1$. This influence is reduced to about two thirds by the coefficient of the second lag. That high average strategies promote high average strategies in the next round is intuitively clear: If each bidder faces at time t an environment of high bidding competitors, she will in the next round also use a high strategy. Conversely, by low bidding competitors she is led to use a low strategy in the next round. Compared to a random walk, this $AR(2)$ process leads to longer consecutive periods of the average strategy above the mean respectively below it. So, in repeated auctions we might expect periods of fiercer competition among bidders that interchange with periods of lower competition. This process is purely internally generated and stems only from the mutual adaptation of bidding strategies.

Analogously I estimated an $AR(2)$ -model of the cumulated seller revenue over the time steps of 100 rounds and obtained the parameter estimates

$$\alpha_0 = 78.6 \pm 0 \quad \alpha_1 = 0.285 \pm 0.006 (0.273; 0.298) \quad \alpha_2 = -0.009 \pm 0.004 (-0.017; -0.001) .$$

The process is qualitatively the same as for the average strategy: Periods of relatively high seller revenue change with periods of relatively low seller revenue. However, the parameter α_1 is lower than in the $AR(2)$ model for β_{av} . The reason is that the random value realizations dilute the effects of the time evolution of the strategies. I did not estimate AR-models for the individual strategy β_0 and the minimum strategy β_{min} since they are not symmetrically distributed around their mean and therefore an AR model with a normal noise specification is not apt for their investigation.

By setting `numRounds= 1.000`, I obtain the following $AR(2)$ -models for β_{av}

$$\beta_{\text{av}}^{(t)} = 0.858 + 0.374(\beta_{\text{av}}^{(t-1)} - 0.858) - 0.042(\beta_{\text{av}}^{(t-2)} - 0.858) + \epsilon^{(t)} \quad (11.20)$$

and for the cumulated seller revenue

$$PO_{\text{sell}}^t = 807 + 0.296(PO_{\text{sell}}^{(t-1)} - 807) - 0.005(PO_{\text{sell}}^{(t-2)} - 807) + \epsilon^{(t)} . \quad (11.21)$$

Comparing the coefficients with the coefficient estimates of the simulation with `numRounds= 100` shows that the parameter `numRounds` has virtually no effect on the time evolution of $\beta_{\text{av}}^{(t)}$ and $PO_{\text{sell}}^{(t)}$: The long term dynamics is independent of the number of rounds after which the bidders update their strategies.

The coefficients are quantities that depend only on the number of bidders in the auction¹⁴. Table 11.7 shows the coefficients of the $AR(2)$ estimation for different numbers of bidders. The first-order autocorrelation decreases slowly when we increase the number of bidders. It remains a topic for further research to investigate whether the autocorrelation vanishes for `numBidders` $\rightarrow \infty$. I want to investigate this when I have access to faster computer facilities.

¹⁴The simulations for table 11.7 were made with `numRounds=1.000` for 2 and for 10 bidders and with `numRounds=300` for 5 bidders and 20 bidders. As shown above, the different numbers of rounds have virtually no influence on the results.

| numBidders | 2 | 5 | 10 | 20 |
|------------------------|--------|--------|--------|--------|
| $\alpha_0(\beta_{av})$ | 0.423 | 0.395 | 0.374 | 0.34 |
| $\alpha_1(\beta_{av})$ | -0.009 | -0.032 | -0.042 | -0.041 |
| $\alpha_0(PO_{sell})$ | 0.396 | 0.328 | 0.295 | 0.274 |
| $\alpha_1(PO_{sell})$ | 0.003 | -0.006 | -0.005 | -0.012 |

Table 11.7: Coefficient Estimates for Different Number of Bidders

11.5 Computational Assessment of Nash Equilibria

It is a well known theoretical results that the NE in n^{th} price auctions with underlying uniform value distributions are given by linear bidding strategies. Therefore it is natural to ask if we can use the AS to assess these NEs.

The restriction to linear bidding strategies prevents the AS from finding exact solutions to general auction setups like, e.g., first-price auctions with normal value distributions. Therefore I want to assess as a second question in this section, how big the simulation error for non-uniform value distributions is from restricting the bidding functions to linear ones. I can not expect the AS to replicate the NE correctly. However, I am interested whether I can use the AS at least to replicate comparative static predictions. A route for further research is to extend the AS to allow for the simulation of nonlinear bidding functions. Then, I can use the AS to determine NEs of more general auction setups such as normal or asymmetric value distributions.¹⁵ I will say more on this in the last chapter.

As an example for an auction with linear NE strategies I investigate a 4th price auction with 8 participants. I assume that initially I don't know anything about the NE. For FS I would have to determine a feasible strategy range in the beginning of the simulation. Therefore I use the GA to find the approximate position since the results are then not confined to a strategy interval. I put numRounds=49.999 to limit the effects of fictitious play truncation and thereby come as close to perfect memory best response as possible. I assess the last 50 of 150 generations. The result indicates that the NE is at

$$\bar{\beta}^{\text{GA}} = 1.3998 \pm 0.0133 \quad (11.22)$$

Wolfstetter (1996) analytically derives the NE for a k^{th} price auction with n participants and uniform value distributions as

$$b(v) = \left(1 + \frac{k-2}{n-k+1}\right)v \quad (11.23)$$

In the particular of 8 bidders in a 4th price auction, this amounts to a linear NE strategy of $\beta^{\text{NE}} = 1 + \frac{4-2}{8-4+1} = 1.4$ which is in perfect agreement with the simulation result.

More general, for all k^{th} price auctions with n bidders I was able to get the analytically predicted result for the NE. This shows that the linear NE in all k^{th} price auctions is attracting. Therefore, the AS is apt to quantify correctly the predictions of Wolfstetter (1996) for k^{th} price auctions with uniform value distributions. In particular:

- Equilibrium Bids *exceed* bidders' valuations for $k \geq 3$
- Equilibrium bids increase in k
- Equilibrium bids diminish for $k \geq 3$ if the number of bidders is increased

As a second auction setup I investigate the ability of the AS to correctly predict average strategies for normal distributions. As stated above, the AS can not assess the non-linear bidding functions. Therefore I focus on average strategies since they determine expected seller revenue which is one of the most important characteristics of auctions. Table 11.5 compares the NE strategies for normal value distributions with the bids that AS simulations predict. The normal value distributions have a mean of $\mu = 100$ and differ in their variance σ_v^2 . The NE bidding function $b(v)$ can be calculated by equation 2.4. Note that the NE has no closed-form solution. A Mathematica- Notebook with the numerical calculations can be obtained from the

¹⁵Of course, I can only assess these NEs if they are attracting and not repelling like, e.g., in double auctions.

| σ^2 | 0.001 | 0.01 | 0.1 | 1 | 10 |
|--|--------------|-------------|------------|----------|-----------|
| β^{NE} | 0.999714 | 0.999092 | 0.997144 | 0.990968 | 0.971438 |
| β^{GA} | 0.999439 | 0.998226 | 0.994419 | 0.982510 | 0.945885 |
| $\text{dev}(\beta^{\text{NE}}, \beta^{\text{GA}})$ | -0.03% | -0.09% | -0.27% | -0.85% | -2.63% |

Table 11.8: Comparison of NE with simulation results for $v \sim N(100, \sigma^2)$

author.

The first line in the table gives the average best strategy, calculated as

$$\bar{\beta}^{\text{NE}} = \frac{\text{average bid in the NE}}{\text{average value}} = \frac{\int_0^{2\mu} b(v)f(v)dv}{\mu}. \quad (11.24)$$

The second line gives the simulation results for the linear bidding strategy simulation, where I used a GA run with 500 generations of 49.999 auctions each from which I recorded the last 150. The third line gives the deviation between NE-strategy and simulation results.

We see from the results that the tendency of the AS to underbid when compared to the NE prevails also in the case of normal value distributions. The AS-result is persistently about double as far away from the 1 as the NE. Therefore, the deviation gets ever more pronounced as the variance of the value distribution increases. This is intuitively clear: Remember that in the degenerate case of zero variance, strategies must converge to 1 which is a linear strategy. For bigger σ_v , higher-order terms become increasingly more important and therefore the linear simulation technique of the AS fails to give correct results.

So, if values come from a normal distribution, the AS shows qualitatively correctly the decrease of the mean strategy with the variance of the value distribution. But it fails to replicate the quantitative predictions because it fails to simulate non-linear strategy.

11.6 Conclusion from this Chapter

Mutual adaptation between strategies induces non-trivial dynamics in repeated auctions. Under myopic best response with 2 bidders and uniformly distributed values we saw already in the last chapter that the results depend strongly on the number of strategies that are used in the simulation, since they determine the amount of overbidding near 0. If strategies come close to 1, I showed in this chapter that the specification of the rule for strategy updating if no feasible strategy could have won the auction has large influence on the results. Both of these influencing factors vanish, however, if we investigate truncated fictitious play by increasing the number of rounds after which strategies are updated.

Under truncated fictitious play with 10 bidders, the expected seller revenue is lower than in a second price auction and the seller revenue is riskier. Bidders' strategies fluctuate permanently. On the individual level, this fluctuation is skewed: strategies significantly below the mean are played more often than corresponding high strategies. On the aggregate level, the average strategy shows symmetric fluctuations around a mean below the NE. These fluctuations show autocorrelation, so periods of consecutive fiercer bidding and periods of lower bidding last longer than they would under a random walk. Consequently also for the seller there are periods of higher revenue that interchange with periods of lower revenue.

The intuitive reason is that if a bidder notices that her opponents pursue high strategies, then she is led to use a high strategy after the next updating. Conversely, if she sees that her opponents bid defensively, she does better if she also reduces her strategy in order to maximize her payoff. This shows that repeated auctions can generate emergent properties on the macro-level that stem from the interaction of heterogeneous agents.

Moreover I showed that the simulation of pBR allows us to accurately assess NEs if they are given by linear bidding strategies like, e.g., in k^{th} price auctions. However, the AS can not simulate correctly setups that give rise to non-linear NE-functions.

A realistic simulation of repeated auctions requires a possibility to account for value-dependent bidding strategies. Therefore, I want to equip the Auction Simulator in the future with a possibility to simulate these. I will say more on this in the next, last chapter.

Chapter 12

Conclusion and Areas for Future Research

In this thesis I have developed mathematical and computational methods to investigate evolutionary learning in repeated auctions. 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. Either, she manages to do so (best response) or she makes errors but still plays profitable strategies more often (quantal response). Additionally, bidders are discerned according to their memory strength (myopic, perfect memory or in between).

The results of mathematical and computational analysis show that 1PAs under best response dynamics are characterized by excess volatility and therefore lead to a suboptimal allocation of goods. In contrast, players in second-price auctions find the NE easily - either upfront by simple reasoning or in the course of time by following the best response strategy updating.

For double auctions, the non-convergence results was even more disturbing than in the pure bidder case: I showed that increased sophistication leads in first-price auctions to convergence into the NE. But players in double auctions end up further and further away from the NE, the more sophisticated they are.

The underlying reason for 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, double auctions etc. - depends on the strategies of all other players. This induces a permanent mutual feedback of strategies, thus raising the volatility. Analogously *any value shading auction design must show excess volatility* when compared to an auction design where bidding the true value is dominant because bidders' optimal strategies depend on each other.

For single unit pure bidder auctions the 2PA allocates goods optimally: Since under best response all bidders end up in the NE, this auction format allocates the good to the bidder who values it the most. For multi-unit auctions, corresponding dominant strategy auction formats would be the Ausubel auction (see Ausubel (2004)) or the generalized Vickrey-auction. For single-unit double auctions, McAfee (1992) has proposed a dominant strategy mechanism. This mechanism is not maximally efficient since it potentially prevents the least valuable trade. It remains a subject for further research to compare this inefficiency with the inefficiency that arises in other double auction formats without dominant strategies since bidders don't find the NE. For single-unit single-sided with interdependent values, Perry and Reny (2002) proposes a dominant strategy auction that basically consists of two consecutive series of second-price auctions.

My long-term research goal is the reduction of volatility in Financial Markets. The interaction between brokers on the floor or on the computer screens are repeated double-sided multi-unit auctions with affiliated values. Current order books are designed as two-sided uniform price auctions and therefore give incentives for bid shading. The design and implementation of a dominant-strategy auction format for double-sided multi-unit auctions with affiliated values could therefore have the potential to reduce excess volatility in financial markets.

Other markets where redesigned order books could be implemented are electricity markets and the currently emerging global market for greenhouse emission certificates. These markets are essentially structured in the same way as financial markets. Especially the latter market would be the ideal place for a first implementation of revised order books: It is a new market without strong incumbents and the declared goal is - according to the Kyoto protocol - to use the certificates as efficiently as possible.

Another field where dominant strategy auction formats could increase efficiency is with respect to supply chains and -networks. Supply chains can be viewed as chains of interconnected double auctions (see, e.g., Babaioff and Nisam (2001)). Consider for instance the market for rubber tires for trucks and the two related markets for rubber and for trucks. Then, the stylized supply chain looks as follows: Rubber manufacturers place asks for rubber in the rubber market; tire manufacturers place bids on the rubber market and asks for the tires; truck manufacturers place bids for the tires and asks for the trucks; and finally, customers place bids for trucks.

The usual format for double auctions is the multi-unit-equivalent of the format suggested in Chatterjee and Samuelson (1983). In chapter 8 I showed that the NE of this auction format is repelling so that participants' strategies involve in ever increasing oscillations over time, even if they are initially arbitrarily close to the NE. Moreover, I showed that sophistication of the participants drives their bidding strategies even further away from the NE. Consequently, even a very small volatility of production cost could be magnified by the various double auctions in highly exaggerated price volatility on the consumer side.¹ As a countermeasure, the implementation of dominant strategy formats like the one suggested in McAfee (1992) could significantly reduce this volatility.

A further possible application of auction redesign is given in the context of land conservation auctions. Hailu and Schilizzi (2003) consider land conservation auctions in Australia where conservation contracts for the preservation of wildlife are allocated to private landowners. By simulations they compare the cost that arises to the government if they use fixed payment schemes or auctions in this repeated setting. Specifically, the auctions that the Australian government uses are identified as repeated multi-unit pure-seller discriminatory sealed-ask auctions - the multi-unit equivalent of first-price auctions. In the simulation, landowners use a simple rule where they randomly either stay with their previous ask for the next round or use a simple percentage markup/markdown rule if they have lost/won in the previous auction. The simulation shows that participants involve in permanent mutual adaptation. The authors compare the situation to a basket full of crabs that climb one upon the other to be on top and then drop back to the bottom. On average, the landowners ask more than they would in the NE. Note that this is qualitatively the same result as I derived in this thesis for repeated first-price auctions under best response - the strategies are volatile and the auctioneer gets less utility than in the NE. However, the updating rule is different here.

The deviation from the NE is so big in the simulation that even a fixed price scheme outperforms the auction in terms of efficiency (the landowners with the best ratio of cost/environmental benefit are selected for the contracts) and the total provision of environmental benefits. Therefore, the authors suggest that the government should switch to a fixed price scheme instead of the auctions that are currently conducted.

The results of this thesis suggest, however, that a better way to increase the efficiency of land conservation auctions would be to switch to an Ausubel auction. This auction format has bidding the true value as a weakly dominant strategy so that participants are led to bid in the NE, either by upfront reasoning or by best response learning. Since the Ausubel auction in the NE is maximally efficient, a switch to Ausubel auctions would allocate land conservation contracts more efficient than a fixed price scheme which necessarily leaves some information rent to the landowners.

There are many more economic environments where auctions occur in a repeated setting: IPOs, the first placement of stocks on capital markets, are multi-unit auctions. Since the same big institutional investors participate in virtually all IPOs this is a setting where they might involve in permanent mutual interaction. Treasury bill auctions are conducted by the federal reserve banks in regular frequencies. Again, big institutional investors compete repeatedly against each other. Further examples of repeated auctions are procurement auctions (where the same suppliers compete repeatedly against each other), agricultural or fish auctions and many more. In all these economic systems, a switch to dominant strategy formats *could* significantly in-

¹In a little different context, think about the beer game as a related example of how volatility gets blown up in a supply chain if strategies are interdependent.

crease efficiency and lower risk. However, before implementing dominant strategy formats, one must carefully consider additional constraints. It is for instance well known that second-price auctions allow relatively easily for collusion: The participants can determine their values in a pre-auction. Then, in the main auction, the highest bidder bids her value while all others bid 0. In such a case, the seller would obtain nothing for the asset. A seller should therefore only employ second-price auctions if bidders don't know each other or if she can deter bidders from collusion by sufficiently severe punishments upon detection. Still, the thesis suggest that the gains from a dominant strategy format could in many cases outweigh the additional effort one has to make for the details of the implementation.

In parallel to the development of the mathematical methodology I programmed an agent based computer model, the Auction Simulator (AS). The simulation can capture the simple setups for which mathematical results were obtained but it provides for a much larger space of possible auctions setups. The AS at the current stage simulates repeated k^{th} -price pure-bidder and double auctions with an arbitrary number of participants who pursue linear bidding strategies. Bidders can update their bidding strategies according to best response or quantal response with payoff-linear or payoff-exponential strategy selection probabilities. Memory capabilities of bidders can be smoothly interpolated between myopic and perfect memory. Alternatively, the AS can also interpolate by truncated fictitious play with arbitrary updating intervals. Private independent values can be drawn from uniform or normal distributions. Strategy updating can be simulated by real Genetic Algorithms or by a fixed strategy set for each player. For control purposes, the AS allows for the fixing of the strategy of one or more bidders to investigate the adaptation of the remaining ones.

The AS aims at providing a comprehensive tool for analyzing auctions. The parallel development of the mathematical formalism allows for the rigorous quality-check of the simulation. For simple setups, the AS reproduces the quantitative predictions. Therefore I can use the simulation to reliably assess more complicated auction setups which are mathematical nasty or simply not calculable. In particular, I used the AS in the thesis to assess the dynamics and aggregate statistics of repeated auctions. A second application is the investigation of equilibrium play. By setting the memory strength to 1, the AS converges under best response to the NE in all k^{th} price auctions. By simulating quantal response, it converges analogously to the quantal response equilibrium that predicts behavior in auction experiments better than the NE. So, the AS can also be used as a supplementary tool for the assessment of equilibrium bidding. It is not guaranteed to converge to the equilibrium since the equilibrium needn't be attracting. However, if it converges it provides much more flexibility than the mathematical techniques that are currently used in auction theory.

The current mathematical and computational results provide a first step towards a more realistic description of bidding in auctions than would be possible by game theoretic methods that restrict themselves to closed-form solutions of equilibrium play. However both, mathematics and simulation still need further developments for supplying realistic results.

Nonlinear bidding functions:

As highlighted at numerous places in the thesis, the most important extension of all concerns the strategies: Bidding strategies in auctions are in general nonlinear, they depend on the value of the bidder. While I was already able to use the mathematical methodology to investigate the nonlinear best response to a linear strategy, the simulation is still restricted to linear strategies for all bidders. Therefore I want to enhance the AS by possibilities to simulate nonlinear bidding functions. There are several ways to implement this feature in the program, each with complementary strengths and weaknesses:

GA-implementation: One of the simulation techniques that are already implemented in the simulation are Genetic Algorithms (GAs). They are an especially apt tool to assess equilibrium play by letting bidders accumulate information from many auction rounds before they update their strategies. To implement nonlinear strategies with GAs, it is necessary to implement each strategy as a real-valued vector. One particular way to interpret the vector components is as a Taylor-series expansion of the bidding strategy with the value. By making each strategy consist of e.g. three parameters,

$$\beta_i = (\alpha_i^{(1)}, \alpha_i^{(2)}, \alpha_i^{(3)})$$

$$b_i(v_i) = \alpha_i^{(1)} + \alpha_i^{(2)} v_i + \alpha_i^{(3)} v_i^2,$$

I can model nonlinear bidding strategies with an affine, a linear and a quadratic term. The use

of GAs for the determination of Taylor series coefficients of equilibrium play allows for a rigorous double-check of mathematical predictions. In particular, I want to use GAs in this setting to double-check results on near NE-play in double auctions and on nonlinear best response in first-price auctions where I have already derived closed form mathematical predictions.

A second, more general, extension of the GA functionality would be a non-parametric simulation of value-dependent bidding strategies. The term non-parametric refers to the fact that we merely get a plot of the bidding function but not an equation for it. To understand this methodology assume that I divide the value space into 10 segments of equal size. The bidding strategy is affine² with the same parameters for all values from the same segment, however in general the parameters differ among segment. So, in the case of 10 value segments, a strategy of the GA consists of 10 pairs of real numbers $\beta_i = ((\alpha_i^{(1)}, \gamma_i^{(1)}) ; \dots ; (\alpha_i^{(10)}, \gamma_i^{(10)}))$ that determine the strategy if the value comes from the respective interval. The bid of bidder i for a value in segment j is then given as $b_i(v_i) = \alpha_i^{(j)} + \gamma_i^{(j)} v_i$. If the simulation converges, the steady states of the parameters $(\alpha_i^{(j)}, \gamma_i^{(j)})$ can be used to plot the value-dependent bidding function. Of course, instead of assigning a strategy to each value-segment, one could also assign a bid to each value-segment. This would allow more directly to determine the impact of the size of bid increments on the outcome of auctions, however, the value segments in such simulations would have to be chosen much smaller than if strategies are used and therefore would be more time consuming.

The power of GAs lies in their ability to discover global optima of multi-dimensional problems. The crossover operator leads the strategy population to local optima. The mutation operator, on the other hand, prevents the solution from getting stuck in local optima. However, GAs work only if the underlying function that they are trying to optimize does not fluctuate over time. In the context of auctions this means that they are a suitable tool to assess Nash and quantal response equilibrium bidding functions in auctions. However, if bidders with finite memory capabilities forget information from longer ago, then the fitness function of the strategies fluctuates to frequently too allow their usage.

FS-implementation: A complementary approach to GA-simulation is the modeling of nonlinear bidding strategies via Fixed Strategies (FS). Their advantage over GAs is that they allow for the accurate simulation of different degrees of memory capabilities. However, their simulation is much more time consuming than the simulation with GAs. I want to implement nonlinear bidding strategies with FSs only in a non-parametric way: Again, values are segmented in a number of intervals, say, 10 intervals. For each value segment j , bidder i has a fixed set of affine strategies $\beta_i^{(j):(k,l)} = (\alpha_i^{(j):k}, \gamma_i^{(j):l})$ that assign a bid to each value like in the GA case. In each auction each bidder calculates for each strategy in her strategy set the payoff it would have generated in the auction. If we allow for instance 100 possible values of $\alpha_i^{(j)}$ and $\gamma_i^{(j)}$ between 0 and 1, each bidder has a set of $10 \cdot 100 \cdot 100 = 100.000$ strategies that she has to keep track of in each auction. Since many bidders may participate in the auction and simulation runs typically consist of thousands of auctions, this shows that this simulation methodology is only feasible by the use of massive computer power.

The explicit calculation of payoffs for each possible strategy, allows for a lot of flexibility in the setups: If each bidder chooses the strategy with the highest payoff in each value segment, this simulates best response. If she chooses strategies with a probability proportional to an exponential of the payoff, this simulates quantal response. By multiplying past payoffs with a discount factor, the effects of memory capabilities on the outcome can be determined.

The power of FS lies in their correct functioning on all scales of memory strength. On the other hand, GAs converge more quickly and more reliably to global solutions. By using FS and GA simulations complementarily, I can use the Auction Simulator to investigate dynamic behavior in repeated auctions or equilibrium strategies.

Auction Formats:

I want to enable the AS for the simulation of more general auction setups than k^{th} price auctions. In particular, I want to implement all-pay auctions and multi-unit auction formats like discriminatory, uniform price and Ausubel auctions. In the program multiple units correspond

²The use of affine strategies is the minimum requirement to guarantee the smoothness of the value-dependent strategy function. Still, this approach allows a rich variety of concave or convex bidding functions. Alternatively, one could also think about a GA implementation that assigns only one parameter to each value segment and is interpreted as the slope of the bidding strategy in the respective value segment. However, this approach could suffer the disadvantage that the optimal parameter settings for higher value segments depend on the parameter choice for the lower segments.

to a list of values for each bidder where the strategies for each value in the list can be different. Values in multi-unit auctions should be implementable as homogeneous or heterogeneous with different interrelation structures (supplements or complements) so that the value of the object for sale depends on whether the bidder has already obtained one or several items before. This would allow the simulation of many more auction setups such as they occur in treasury bill auctions or IPOs.

Values:

I want to allow for the possibility that values are drawn from more general distributions than currently possible. On the one hand, the AS should allow for different value distributions among bidders. This would serve to computationally investigate the effects of low balling in auctions with one strong and one weak bidder. Second, it should be possible to simulate affiliated or common values where a bidder's value depends on the other bidders' values. This simulation of non-private values would enhance the realism of the simulation of many auctions. Third, it should be possible to generate more general value distributions than uniform and normal distributions.

Risk Aversion:

Players shall be modeled with different degrees of risk aversion. In particular, the utility function shall be specifiable as a CARA or CRRA function with parameter values that can be freely chosen. The comparison of simulation data with experimental data can later be used to determine the effects of payoffs and experiment framing on bidders' risk aversion in auctions. In combination with the simulation of quantal response, this would allow to investigate in more detail the quantal response equilibrium that is suggested as an explanation for overbidding in auctions in Goeree, Holt, and Palfrey (1999). In particular, I hypothesize that the results therein can be further refined by considering the finite memory capabilities of the bidders.

Updating rules:

To guarantee a realistic assessment of sealed bid auctions and of auctions under myopic play, the AS shall be enhanced by different updating rules for winning bidders in sealed bid auctions and for updating rules for losing bidders, if none of their strategies could have won the auction. These rules shall be chosen based on experimental observations. *Updating Times:*

Currently, the bidders in the AS update their strategies simultaneously at fixed time intervals. The simulation results could change significantly if we allow for different updating schemes such as Poisson distributed updating times or an updating probability that is influenced by the success of the currently active strategy in comparison to the long-term average.

For all these extensions of the simulation, the mathematical methodology should be appropriately extended and applied to simple setups that entail the specific extension. This allows the ongoing quality-check of the simulation. In the future, the mathematical methodology could also be extended to analyze behavior in other economic fields that deal with discontinuous payoff structures – like, e.g., the analysis of oligopolies.

When I have developed an extended version of the Auction Simulator, there are several ways to make use of it. For my own future research I have two particular applications in mind:

Experimental Auctions:

The first application is to compare simulation results with data sets from experimental and maybe also empirical auctions. Thereby it shall be checked, to which extent bounded rational play is a good model of play in real auctions. In particular, it is a possibility to quantify in a rigorous way parameters that determine the behavior of participants like the memory strength λ or the payoff-weight ξ for quantal response: The simulation parameters should be fine-tuned until they replicate optimally the statistics of the experiment. Alternatively, one could also think about the development of a Meta-GA or a Neural Network that searches for the optimal parameter combinations. The results could quantify more precisely the effects of payoff-structures and experiment framing on the outcome of repeated auctions. Note that it is not necessary to conduct new experiments for this but it would be sufficient to obtain old data sets and to subject them to the analysis. Complementarily, the simulation could also be used to analyze the time series properties of empirical repeated auctions. Again, by manual- or GA supported parameter tuning, one could try to obtain more insights into the forces that govern real auctions.

Simulation of Financial Markets:

This aims at extending the simulation by parts that make it possible to simulate financial markets. The goal of this line of research is to identify possible order book designs that reduce

the excess volatility in financial markets. The first step in this line of research is to integrate the underlying concepts of weakly dominant strategy order books for double auctions (see McAfee (1992)), multi-unit auctions (see Ausubel (2004)) and auctions with affiliated values (see Perry and Reny (2002)) into a weakly dominant strategy format for multi-unit double auctions with affiliated values. This auction format shall be implemented in the simulation. Second, the bidder and the seller shall be merged into one agent, the trader. The trader bases her bid- and ask-orders on the wealth, she has already accumulated and her expectation of future prices. The value expectations of the trader consist of a fundamentalist and a chartist component, i.e., of a private and of an affiliated component. Then, I will compare the simulation results of the order book designs with weakly dominant strategies with simulation results of current order book designs. Thereby the effectiveness of order book restructuring in reducing price volatility and the occurrence of bubbles and crashes shall be investigated. In addition, also the effects of speculation and transaction taxes (Tobin tax) on the price stability could be investigated in the simulation.

All in all, I think that a variety of interesting work streams could follow from this thesis. The increasing importance of auctions in the economy makes it vital that we understand their properties in depth. This thesis is a step into this direction and therefore towards a more stable and efficient economic system.

Bibliography

- ANDREONI, J., AND J. MILLER (1995): "Auctions with Artificial Adaptive Agents," *Games and Economic Behavior*, 10, 39–64.
- ARTHUR, W. B., S. N. DURLAUF, AND D. A. LANE (eds.) (1997): *The Economy as an Evolving Complex System*pp. 15–38. Perseus Books, Reading, Massachusetts.
- AUSUBEL, L. M. (2004): "An efficient Ascending-Bid Auction for Multiple Objects," *forthcoming in American Economic Review, Working Paper Version at www.ausubel.com/auction-papers/efficient-ascending-auction-r.pdf*.
- BABAIOFF, M., AND N. NISAM (2001): "Concurrent Auctions Across the Supply Chain," *Working Paper, www.cs.huji.ac.il/~noam/chain.ps*.
- BATTIGALI, P., AND M. SINISCALCHI (2003): "Rationalizable Bidding in First-Price Auctions," *Working paper, <http://faculty.northwestern.edu/faculty/siniscalchi/bs02auct.pdf>*.
- BYDE, A. (2002): "Applying Evolutionary Game Theory to Auction Mechanism Design," *Hewlett-Packard Workingpaper, <http://www.hpl.hp.com/techreports/2002/HPL-2002-321.pdf>*.
- CASTI, J. L. (1997): *Would-Be Worlds*. John Wiley & Sons, inc., first edn.
- CHATTERJEE, K., AND W. SAMUELSON (1983): "Bargaining under Incomplete Information," *Operations Research*, September-October, 835–851.
- CHIARELLA, C., AND G. IORI (2002): "A simulation analysis of the microstructure of double auction markets," *Quantitative Finance*, 2, 1–8.
- CONLISK, J. (1996): "Why Bounded Rationality?," *Journal of Economic Literature*, XXXIV, 669–700.
- COX, J., B. ROBERSON, AND V. L. SMITH (1982): *Theory and Behavior of Single Object Auctions*pp. 1–43. JAI Press, Greenwich.
- COX, J., V. SMITH, AND J. WALKER (1984): "Theory and Behavior of Multiple Unit Discriminative Auctions," *The Journal of Finance*, XXXIX(4), 983–1010.
- COX, J., V. L. SMITH, AND J. M. WALKER (1992): "Theory and Misbehavior of First-Price Auctions: Comment," *American Economic Review*, 82(5), 1392–1412.
- DANIELS, M. G., J. D. FARMER, L. GILLEMOT, G. IORI, AND E. SMITH (2003): "Quantitative Model of Price Diffusion and Market Friction Based on Trading as a Mechanistic Random Process," 90(10).
- DAWID, H. (1999): "On the Convergence of Genetic Learning in a Double Auction Market," *Journal of Economic Dynamics & Control*, 23, 1545–1567.
- DELTAS, G., AND R. ENGELBRECHT-WIGGANS (2001): "Auctions with an Inexpert Bidder," *Economic Letters*, 73, 35–42.
- DELTAS, G., AND T. ENGELBRECHT-WIGGANS (2000): "Naive Bidding," *Working Paper of Illinois at Urbana-Champaign, Office of Research Working Paper Number 00-0118*.

- ENGELMANN, D., AND V. GRIMM (2003): “Bidding Behavior in Mult-Unit Auctions - An Experimental Investigation and some Theoretical Insights,” *Working Paper*, <http://home.cerge-ei.cz/engelmann/WorkingPapers/DerExIERSubmission.pdf>.
- FIBICH, G., AND A. GAVIOUS (2003): “Asymmetric First-Price Auctions- A Perturbation Approach,” *Working Paper*, www.math.tau.ac.il/%7Efibich/Manuscripts/first_rev_final.pdf.
- FRIEDMAN, D. (1992): “Theory and Misbehavior of First-Price Auctions: Comment,” *American Economic Review*, 82(5), 1374–1378.
- FRIEDMAN, L. (1955): “A competitive-bidding strategy,” *Operations Research*, 4, 104–112.
- FUDENBERG, D., AND D. LEVINE (1999): *The Theory of Learning in Games*. MIT Press, Cambridge, London, second edn.
- FUDENBERG, D., AND D. K. LEVINE (1996): “Learning and Evolution in Games,” *Working Paper*; <http://levine.sscnet.ucla.edu/papers/essay/essay7.htm>.
- GARRAT, R., AND J. WOODERS (2004): “Efficiency in Second-Price Auctions: A New Look at Old Data,” *Working Paper*; <http://www.econ.ucsb.edu/papers/wp2-04.pdf>.
- GOEREE, J. K., AND C. A. HOLT (1999): “Stochastic Game Theory: For Playing Games, Not Just Doing Theory,” *Proc. Natl. Acad. Sci. USA*, 96, 10564–10567.
- (2001): “Ten Little Treasures of Game Theory and Ten Intuitive Contradictions,” *American Economic Review*, 91(5).
- GOEREE, J. K., C. A. HOLT, AND T. R. PALFREY (1999): “Quantal Response Equilibrium and Overbidding in Private-Value Auctions,” *working paper*; www.people.virginia.edu/~cah2k/auction.pdf.
- GOLDBERG, D. E. (1989): *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison Wesley Longman, Inc., twentieth edn.
- HAILU, A., AND S. SCHILIZZI (2003): “Learning in a ”Basket of Crabs“: An agent-based computational model of repeated conversation auctions,” *Working Paper*, www.bwl.uni-kiel.de/vwlinstitute/gwrp/wehia/papers/hailu.pdf.
- HARRISON, G. W. (1989): “Theory and Misbehavior of First-Price Auctions,” *American Economic Review*, 79, 749–762.
- (1992): “Theory and Misbehavior of First-Price Auctions: Comment,” *American Economic Review*, 82(5), 1426–1443.
- HAUSCH, D. B. (1986): “Multi-Object Auctions: Sequential vs Simultaneous Sales,” *Management Science*, 32(12).
- HOLLAND, J. (1992): *Adaptation in Natural and Artificial Systems*. MIT Press, Massachusetts, second edn.
- HON-SNIR, S., AND A. SELA (1998): “A Learning Approach to Auctions,” *Journal of Economic Theory*, 82, 65–88.
- JEHLE, G., AND P. RENY (2001): *Advanced Microeconomic Theory*. Addison Wesley.
- JEITSCHKO, T. (1998): “Learning in Sequential Auctions,” *Southern Economic Journal*, 65(1).
- JEITSCHKO, T., AND E. WOLFSTETTER (2002): “Scale Economies and the Dynamics of Recurring Auctions,” *Economic Inquiry*, 20, 403–414.
- JOSHI, S., J. PARKER, AND M. BEDAU (2002): “Financial Markets can be at Sub-optimal Equilibria,” *Computational Economics*, 19(1), 5–23.
- KAGEL, J. H., AND A. ROTH (1992): “Theory and Misbehavior of First-Price Auctions: Comment,” *American Economic Review*, 82(5), 1379–1391.

- (eds.) (1995): *The Handbook of Experimental Economics* pp. 501–585. Princeton University Press, Princeton.
- KLEMPERER, P. (1999): “Auction Theory: A Guide to the Literature,” *Journal of Economic Surveys*, 13(3), 227–286.
- KRISHNA, V. (2002): *Auction Theory*. Academic Press.
- LEBRUN, B. (1999): “First Price Auctions in the Asymmetric N Bidder Case,” *International Economic Review*, 40(1), 125–142.
- LUNA, F., AND A. PERRONE (2002): *Agent-Based Methods in Economics and Finance: Simulations in Swarm*. Kluwer Academic Publishers.
- LUNA, F., AND B. STEFANSSON (2000): *Economic Simulations in SWARM: Agent Based Modelling and Object Oriented Programming*. Kluwer Academic Publishers.
- LUX, T., E. SAMAINDOU, D. STAUFFER, AND E. ZSCHISCHANG (2004): *Microscopic Models of Financial Markets* Springer, in press, Berlin.
- LUX, T., AND S. SCHORNSTEIN (2002): “Genetic Learning as an Explanation of Stylized Facts of Foreign Exchange Markets,” *forthcoming in Journal of Mathematical Economics; in press; Working Paper Verson: www.bwl.uni-kiel.de/vwlinstitut/gwrp/publications/lux_gen_learning.pdf*.
- MASKIN, E. S., AND J. G. RILEY (1999): “Asymmetric Auctions,” *Review of Economic Studies*, 67, 413–438.
- MCAFEE, P. (1992): “A Dominant Strategy Double Auction,” *Journal of Economic Theory*, 6, 434–450.
- MCAFEE, P., AND D. VINCENT (1997): “Sequentially Optimal Auctions,” *Games and Economic Behavior*, 18, 246–276.
- MERLO, A., AND A. SCHOTTER (1992): “Theory and Misbehavior of First-Price Auctions: Comment,” *American Economic Review*, 82(5), 1413–1425.
- MICHALEWICZ, Z. (1999): *Genetic Algorithms + Data Structures = Evolution Programs*. Springer, third edn.
- MITCHELL, M. (1999): *An Introduction to Genetic Algorithms*. MIT Press, Massachusetts, sixth edn.
- MYERSON, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- PERRY, M., AND P. J. RENY (2002): “An Efficient Auction,” *Econometrica*, 70(3), 1199–1212.
- PEZANIS-CHRISTOU, P., AND A. SADRIEH (2004): “Elicited Bid Functions in (a)symmetric First-Price Auctions,” *Working Paper*, <http://www.iae-csic.uab.es/ppc/pub/AsaPap42.pdf>.
- RILEY, J. G., AND W. F. SAMUELSON (1981): “Optimal Auctions,” *American Economic Review*, 71, 381–392.
- RUBINSTEIN, A. (1991): “Comments on the Interpretation of Game Theory,” *Econometrica*, 59(4), 909–924.
- SHILLER, R. J. (1999): *Market Volatility*. MIT Press, Cambridge, Massachusetts, sixth edn.
- SIGMUND, K., AND H. P. YOUNG (1995): “Introduction,” *Games and Economic Behavior*, 11, 103–110.
- VICKREY, W. (1961): “Counterspeculation, Auctions and Competitive Selaed Tenders,” *Journal of Finance*, 16(1), 8–37.
- WOLFSTETTER, E. (1996): “Third- and Higher-Price Auctions,” *Working Paper*; www.wiwi.hu-berlin.de/wt1/papers/1996/third_and_higher_price_auctions.pdf.
- (1999): *Topics in Microeconomics: Industrial Organization, Auctions and Incentives*. Cambridge University Press.

Appendix A

Θ - and δ -Function

The upper part of Figure A.1 shows the Θ -function (Unit step function, Heavyside function).




Figure A.1: Theta- and Delta Function

It assigns 0 to every negative value and 1 to every non-negative value. Formally, it is defined as

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

Note that in the case of continuously distributed values it is irrelevant which specific function value is assigned to $x = 0$ since this event occurs only with zero probability. Other conventions that lead to the same results in the mathematical part of this thesis would be, e.g., $\Theta(0) = 1/2$ or $\Theta(0) = 0$.

The Θ -function (with the convention $\Theta(0) = 1/2$) can be written as the limes of a function series, e.g.,

$$\Theta(x) = \lim_{n \rightarrow \infty} \frac{1 + \tanh(nx)}{2} \quad (\text{A.2})$$

where the hyperbolic tangent of x is defined as

$$\tanh(x) = \frac{\sinh(x)}{\cosh(x)} = \frac{e^x - e^{-x}}{e^x + e^{-x}}. \quad (\text{A.3})$$

By means of the Θ -function, the minimum of two functions $f(x)$ and $g(x)$ can be written as

$$\min(f(x), g(x))(x) = f(x) - \Theta(f(x) - g(x))(f(x) - g(x)) . \quad (\text{A.4})$$

If x is such that $f(x) < g(x)$ then the Θ -term vanishes and the result is $f(x)$. If, however, $f(x) \geq g(x)$ at x , then the Θ -function takes on the value of 1 and therefore the result is $g(x)$. Similarly, the maximum of two functions can be written as

$$\max(f(x), g(x))(x) = f(x) + \Theta(g(x) - f(x))(g(x) - f(x)) . \quad (\text{A.5})$$

The lower part of Figure A.1 shows a function series that converges towards the δ -function. The δ function is the derivative of the Θ -function,

$$\delta(x) = \frac{\partial}{\partial x} \Theta(x) . \quad (\text{A.6})$$

Intuitively we can think about a function that satisfies

$$\delta(x) = \begin{cases} 0 & \text{for } x \neq 0 \\ \infty & \text{for } x = 0 \end{cases} \quad (\text{A.7})$$

and

$$\int_{-\infty}^{\infty} dx \delta(x) = 1 . \quad (\text{A.8})$$

The second condition implicitly quantifies the "magnitude of infinity" that the δ -function takes on at $x = 0$. This is an intuitive way to think about the δ -function but not a rigid mathematical definition. More formally, the δ -function is defined by means of the δ -functional. The δ -functional maps the function space \mathcal{F} of all smooth real-valued functions with compact support into the real numbers by $\delta(f(x)) \stackrel{\text{!}}{=} f(0)$. This functional can be written in integral form as

$$\int_{-\infty}^{\infty} dx \delta(x) f(x) = f(0) \quad (\text{A.9})$$

The δ function, $\delta(x)$ is defined as the kernel of that integral.

The δ -function has the fundamental properties

$$\int_{-\infty}^{\infty} dx f(x) \delta(x - a) = \int_{a-\epsilon}^{a+\epsilon} dx f(x) \delta(x - a) = f(a) \quad (\text{A.10})$$

and

$$\delta(x - a) = 0 \quad \text{for } x \neq a . \quad (\text{A.11})$$

The δ function can be written in a variety of ways as the limes of a series of functions, e.g.,

$$\begin{aligned} \delta(x) &= \frac{1}{\pi} \lim_{\epsilon \rightarrow 0} \frac{\epsilon}{x^2 + \epsilon^2} = \lim_{\epsilon \rightarrow 0} \epsilon |x|^{\epsilon-1} = \\ &= \lim_{\epsilon \rightarrow 0} \frac{1}{2\sqrt{\pi\epsilon}} e^{-\frac{x^2}{4\epsilon}} = \lim_{\epsilon \rightarrow 0} \frac{1}{\pi x} \sin\left(\frac{x}{\epsilon}\right) \end{aligned} \quad (\text{A.12})$$

The value of $0 \cdot \delta(0)$ is not defined.

Appendix B

Simulation Parameters

Parameters in model.setup

| parameter | explanation |
|---------------------------|--|
| numRounds | number of auctions played between strategy updatings |
| auctionType | if ==1: First-Price Open Bid auction if ==2: Second-Price Open Bid auction if ==3: Third-Price Open Bid auction etc. |
| priceDet | if == 0.5: price is set in the middle of auctionType highest bid and auctionType lowest ask if == 1: price = (auctionType) highest Bid if == 0: price = (auctionType) lowest Ask |
| RWMean | if == 0: Mean of normal value distribution given by valDet2 if ==1: Mean of normal value distribution given by last rounds price |
| transactionFee | transaction fee payed in each round by winning bidder and seller |
| randomSeed | sets the initial state of the random generator |
| numBidders | number of bidders participating in the auction |
| numSellers | number of sellers participating in the auction |
| startOfAverageCalculation | number of initial strategy updatings after which the calculation of averages for graphical output starts |
| reportingBidderID | ID of bidder that reports her strategy list |
| reportingSellerID | ID of seller that reports her strategy list |

Parameters in bidder.setup for GAs and Fixed Strategies

| parameter | explanation |
|----------------------|--|
| bLearningType | if==1, FS are used; if==2, GA is used |
| bSelectionType | if==1, fitness proportional selection (Quantal Response); if==2, rank proportional selection (best response) |
| bFixedStrategyMarker | 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 |
| bFixedBeta | this parameter is only important if fixedStrategyMarker \neq 0 if $\in (0, 1)$: $\beta = \text{fixedBeta}$; if $\in (-1, 0)$: $\beta \sim U(-\text{fixedBeta}, 1)$ |
| bFixedValueMarker | if==0: each player has random values if==1: player 1 always has fixed value; all other values are random if==2: player 0 has random value; all other players have fixed value |
| bFixedValue | determines respective values if fixedValueMarker \neq 0 |
| b0FixedValueMarker | if ==0: random v_0 ; if ==1: fixed v_0 |
| b0FixedValue | if player0FixedValueMarker==1, then $v_0 = \text{player0FixedValue}$ |
| bNumStrategies | number of strategies in each players strategy set |
| bValDistShape | if==0: values from normal Dist; if==1: values from uniform Dist |
| bValDet0 | lower bound for uniform Dist, resp. variance for normal Dist |
| bValDet1 | upper bound for uniform Dist, resp. mean for normal Dist |
| bXi | exponential parameter for Quantal Response |

Parameters in bidder.setup for Fixed Strategies only

| parameter | explanation |
|-----------------|---|
| bMemoryStrength | multiplicative weighing factor for previous rounds' payoffs |
| bMinStrategy | value of the minimal strategy |
| bMaxStrategy | value of the maximal strategy |

Parameters in bidder.setup for GAs only

| parameter | explanation |
|--------------------------|--|
| bStrategyDistance | minimum distance between strategies in the breedingList |
| bNumElite | number of fittest strategies that stay unchanged in the population |
| bNumParents | number of fittest strategies that are used for creating offspring by crossover |
| bCrossoverPar | fraction of the interval between two parent strategies by which offspring is allowed to lie outside the interval |
| bMutationType | if==0: uniform Dist between $\beta(1 \pm \text{mutationPar})$ if==1: normal Dist with variance $\beta \cdot \text{mutationPar}$ |
| bMutationProb | probability of mutation of a strategy in promille |
| bMutationPar | determines extent of mutation; see mutationType |

Parameters in seller.setup Symmetric to the parameters in bidder.setup

Appendix C

Abbreviations

| abbreviation | explanation |
|--------------|--|
| 1PA, 2PA | First-Price Auction resp. Second-Price Auction |
| 1POBA, 1PSBA | First-Price Open resp. Sealed Bid Auction |
| 2POBA, 2PSBA | Second-Price Open resp. Sealed Bid Auction |
| ABM | Agent Based Modeling |
| AS | Auction Simulator |
| BR | Best Response |
| CDF | Cumulated Distribution Function |
| DEQ | Differential Equation |
| ESS | Evolutionary Stable Strategy |
| FP | Fictitious Play |
| FS | Fixed Strategy |
| GA | Genetic Algorithm |
| iBR | Imperfect Memory Best Response |
| iQR | Imperfect Memory Quantal Response |
| mBR | Myopic Best Response |
| mQR | Myopic Quantal Response |
| NE | Nash Equilibrium |
| pBR | Perfect Memory Best Response |
| PDF | Probability Distribution Function |
| pmBR | Probability-Weighted Myopic Best Response |
| pQR | Perfect Memory Quantal Response |
| QR | Quantal Response |
| RET | Revenue Equivalence Theorem |
| RD | Replicator Dynamics |
| SFASM | Santa Fe Artificial Stock Market |
| SIPV | Symmetric Independent Private Values Model |

Appendix D

PhD-related Activities besides Writing the Thesis

During my PhD, I have in addition to writing the thesis and making the required exams in the first year of the program, done the following:

- Contributions at Conferences and Seminars
 - WEHIA - Workshop on Heterogeneous Interacting Agents in Kyoto (Talk)
 - World Meeting of the Game Theory Society in Marseille (Poster)
 - Ecomod in Paris (Talk)
 - SMYE - Spring Meeting of Young Economists in Warsaw (Talk)
 - Summer Research Workshop in Bonn (Talk)
 - Brownbag Seminar in Kiel (Talk)
 - Workshop on Computational Finance in Kiel (Talk)
- Papers
 - Working Paper "Revenue Equivalence Revisited"
 - Contribution to the book "Economics and Heterogeneous Interacting Agents" (submitted)
- Additional Classes that I participated in¹
 - Learning and Bounded Rationality (University of Vienna)
 - Auction Theory (University of Vienna)

¹Besides the ones required by the program in the First Year

Appendix E

Lebenslauf

Name Konrad Richter
Geburtsort: Wien
Geburtsdatum: 17. August 1972
Adresse: Gardegasse 3/9, 1070 Wien
e-mail: konrad_richter@mckinsey.com

Ausbildung:

Nov 2004 Abgabe Dissertation in Kiel
Okt 2002 - Aug 2003 Vorlesungen und Klausuren im Rahmen des DAAD-Programms "Quantitative Economics"
Feb 2000 Teilnahme am einmonatigen BWL-Training "Mini-MBA" von McKinsey&Co
Sep 1998 Teilnahme am einmonatigen VWL-Prgramm "School on the Mathematics of Economics: A Primer in Economics for Physicists and Mathematicians" am ITCP in Triest
Apr 1998 Sponson zum Mag. rer. nat. in Physik an der Universiät Wien;
Diplomarbeit: "Gravitational Anomalies and the Atiyah-Singer Family-Index Theorem"
Jun 1994 Abschluß 1. Studienabschnitt Physik mit Auszeichnung
Jun 1993 Abschluß 1. Studienabschnitt Mathematik mit Auszeichnung
Apr 1992 Abschlußprüfung Klavier mit Erwerb der elementaren Lehrberechtigung für Musikschulen
Jun 1990 Abschluss humanistisches Gymnasium mit Auszeichnung

Berufserfahrung:

März 2000 - Apr 2002 Associate bei McKinsey&Co; Projekte in der Automobil-, Bahn- und Finanzdienstleistungsindustrie; Fokus auf strategische Projekte und Risikomanagement; Projektorte: Genf, München, Frankfurt, Detroit, Tokyo u. a.
Jan 1999-Jan 2000 Zivildienst; Betreuung geistig behinderter Menschen
Okt 1998 - Dez 1998 Praktikum bei McKinsey&Co in Frankfurt; Strategieprojekt in der Bankenindustrie
Mai 1998 - Jul 1998 Praktikum bei Procter&Gamble im Bereich Marketing

| | |
|----------------------|---|
| Jun 1997 - Sep 1997 | Praktikum bei McKinsey&Co in Frankfurt; Strategie- und Reorganisationsprojekt im Bankenbereich |
| Okt 1995 - Apr 1996 | Komposition und Einspielung des Soudtracks zum Kinofilm "Vincent" |
| Sep 1995 - März 1998 | Tutor für Mathematik an der Hochschule für Bodenkultur in Wien |

Wissenschaftliche Präsentationen

| | |
|-----------|---|
| Nov 2004 | Workshop on Computational Finance; Vortrag; Kiel |
| Juli 2004 | World Congress of the Game Theory Society; Poster; Marseille |
| Juni 2004 | Ecomod 2004; Vortrag; Paris |
| Mai 2004 | 9 th Workshop on Economic Heterogeneous Interacting Agents; Vortrag; Kyoto |
| Apr 2004 | Spring Meeting of Young Economists; Vortrag; Warschau |

Gedruckt mit der Genehmigung der
Wirtschafts- und Sozialwissenschaftlichen Fakultät der
Christian-Albrechts-Universität zu Kiel

Dekan: Professor Dr. Drexl
Referent: Professor Dr. Lux
Korreferent: Professor Dr. Raff
Tag des Abschlusses der Arbeit: 14. 10. 2004
Tag der mündlichen Prüfung: 9. Juni 2005