ESSAYS ON THE DESIGN OF PROCUREMENT MECHANISMS

Inauguraldissertation zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften
eingereicht an der Wirtschaftswissenschaftlichen Fakultät der
Heinrich Heine Universität Düsseldorf

2019

vorgelegt von
Philippe Gillen
aus Dudelange

Referent: Prof. Dr. Alexander Rasch
Korreferent: Prof. Achim Wambach, Ph.D.

To Jenna.

ACKNOWLEDGEMENTS

Most importantly, I want to thank my lovely wife Jenna for her support, this journey would not have been possible without you. Thanks to my family and friends for keeping me on track.

My deepest gratitude goes to my supervisors Achim Wambach for sparking my joy for economics and Alexander Rasch for patiently helping me during my first steps as an economist. I'd like to thank all of my coauthors Nicolas Fugger, Vitali Gretschko, Alexander Rasch, Tobias Riehm, Achim Wambach, Peter Werner and Christopher Zeppenfeld. Working on our joint projects has been an absolute pleasure. Special thanks to Nicolas and Tobias for keeping me sane during this process.

Without the help and encouragement from my colleagues at the University of Cologne, the ZEW, and the Heinrich Heine University Düsseldorf this thesis would have been impossible. Thanks to all of you, I had a fantastic time.

CONTENTS

CHAPTER 1

INTRODUCTION

The analysis and design of procurement mechanisms and auctions in general has become an important topic in economic research over the last decades. With a volume of 17% of the European GDP in 2007 in public procurement alone, this interest is not purely academic.[1] Especially in private procurement, small improvements relative to the total procurement costs can lead to large increases in the profit margins. This is confirmed by the consulting company Oliver Wyman. They report that suppliers are responsible for roughly 60% of the value added of a car.[2] For public procurement on the other hand, the objective is usually an efficient allocation. But efficiency can also be important in private procurement, since the relations are often long-term.

It is therefore interesting to pursue research that is either motivated by observations of real-life procurement or to look for well-studied behavioral biases in the economic literature and to give advice on how a buyer could exploit these biases with a suitable procurement mechanism. The essays in this dissertation therefore always pursue the goal of including a pertinent management implication besides being academically relevant.

In the first three chapters, we consider such a behavioral bias, namely loss aversion. In Chapter 2 titled *Procurement Design with Loss Averse Bidders*, which is joint work with Nicolas Fugger and Tobias Riehm, we show that it is beneficial for a buyer to conduct a multi-stage mechanism if bidders are loss averse. In such a multi-stage mechanism, bidders participate in a fixed number of stages and submit a bid in each one of them. The rules of the mechanism include how many stages are conducted and which bidders

[1] Internal Market Scoreboard, n^o 19, July 2009

[2] `https://www.oliverwyman.com/our-expertise/industries/automotive/procurement.html`

advance at each stage. Also, a payment rule is specified which determines what participants have to pay after the mechanism concludes.

First, we derive a revenue equivalence principle. For a fixed multi-stage structure, the auctioneer's revenue is not dependent on the payment rule she chooses. This result considerably simplifies the analysis and allows us to concentrate on the structure of multi-stage mechanisms.

The main result and management implication of this paper is the introduction of a simple and easily implementable two-stage mechanism, the tournament. This mechanism always leads to a decrease in procurement costs compared to any (standard) single-stage auction. Bidders are sorted into two separate groups in a first stage, the semifinals. The lowest bidder of each group then advances to the final. The bidder submitting the lowest offer in the final wins the award process.

Finally, we derive the optimal efficient two-stage mechanism. Taking into account the bidders' degree of loss aversion, the optimal mechanism induces a level of risk that optimally exploits the bidders' loss aversion.

Chapter 3 with the title *Auction Experiments with a Real Effort Task*, which is joint work with Nicolas Fugger and Tobias Riehm, aims to develop a novel experimental tool set to increase the external validity of auction experiments. We propose and test a simple experimental design based on money and effort. When investigating auctions in the laboratory, economic researchers usually rely on induced values experiments. While this gives the researcher a lot of control, it abstracts from two well-known phenomena that both can potentially limit the external validity of results from the lab: Two-dimensional outcome evaluation and common values.

There is ample evidence in the economic literature that both these phenomena are present in most real world auctions, and that both are important drivers of bidding behaviour. Therefore, one has to be cautious when giving practitioners advice based on induced values experiments. Our design aims to account for these phenomena. In a first step, bidders familiarize themselves with the real-effort task we chose, the slider task. In an incentivized test round lasting four minutes, bidders solve as many slider tasks as possible and are remunerated per solved unit. We then let subjects bid on a monetary prize. They submit a bid that expresses the number of slider tasks they would maximally solve in case of winning the auction, i.e. how much effort they are willing to spend in order to receive this prize.

We then implemented a between-subjects design with a varying number of bidders between treatments. If subjects were one-dimensional utility maximizers with a purely private valuation, they would determine the level of effort they are maximally willing to spend for the monetary prize, and bid exactly that amount, independent of the number of competitors. On the other hand, agents that act according to two-dimensional prospect theory bid more aggressively when the number of bidders is low, as a high winning probability leads to an increased attachment to the prize. The same applies if there is a common value component in conducting the slider task. When bidding against a high number of bidders, winning is 'bad news' with a higher probability since a higher number of other bidders estimated a lower common value component.

In line with the reference dependent two-dimensional prospect theory and common value predictions, we observe significantly higher bids if the number of competitors is low. We hence argue that our design enables researchers to increase the external validity of auction experiments. Moreover, the slider task allows experimenters to control for participants' abilities, while at the same time having the advantages of real effort tasks.

Chapter 4 with the title *Preferences and Decision Support in Competitive Bidding*[3], which is joint work with Nicolas Fugger, Alexander Rasch and Christopher Zeppenfeld, aims to understand a discrepancy between theory and real life. The paper is motivated by the observation that the theoretical strategic equivalence between the static first-price sealed-bid auction and the dynamic Dutch auction breaks down empirically.

The three prevalent explanations for the empirically robust difference are opportunity costs, preferences, and complexity of bidding. In a laboratory experiment, we investigate the role of (non-standard) preferences

[3]Financial support from the German Research Foundation (DFG) through the research unit Design & Behavior is gratefully acknowledged. We also want to thank the Center for Social and Economic Behavior (C-SEB) at the University of Cologne. An earlier version of this work is published in Zeppenfeld (2015).

and complexity while controlling for opportunity costs. In line with the experimental literature, we find significant differences between both auction formats if decision support is absent. However, the difference between the formats becomes insignificant once we provide decision support regarding the winning probability. If the differences were driven by preferences, they should be independent of this level of decision support. This indicates that the non-equivalence is caused by differing complexity rather than non-standard preferences.

Staying in a procurement context, but moving away from behavioral models, in Chapter 5 with the title *Commitment in First-Price Auctions*[4], I study the role of commitment in a first-price auction environment. I compare a standard first-price auction with commitment to a first-price auction where renegotiation is possible. In the first-price auction with renegotiation, bidders submit an initial offer that the auctioneer can observe. In the second stage, the auctioneer selects a winner and makes a counteroffer. There is no commitment on the auctioneer's side to accept an offer as is or to choose the lowest bidder. I show theoretically that this implies that bidders pool on bids that reveal no information about their costs. This means that, in equilibrium, the auctioneer has to implement the ex-ante optimal take-it-or-leave-it offer. In the standard first-price auction on the other hand, the auctioneer has to choose one of the offers. In that case, a unique, separating equilibrium exists.

I then take both mechanisms into the laboratory. Contrary to theoretical predictions, I observe no significant difference in the offers between the setting with renegotiation and the standard first-price auction. Also, I find evidence that first-stage offers are correlated with the private information of the bidders in both settings. Still, auctioneers are not able to fully exploit the information in theses offers.

In Chapter 6 with the title *Pre-Auction or Post-Auction Qualification?*[5], which is joint work with Vitali Gretschko and Alexander Rasch, we com-

[4]Financial support from the German Research Foundation (DFG) through the research unit Design & Behavior is gratefully acknowledged.

[5]This chapter is published as Gillen et al. (2017).

pare auctions with bidder qualification before or after the bidding process. Bidder qualification plays an important role in real-life auctions and procurement procedures. Verifying the qualification of a bidder is costly for the buyer and the potential sellers. Interestingly, in most procurement processes, the bidders are required to undergo qualification before the actual awarding of the project. Typically, this is explained by the risk of qualification failure. However, there are many situations where this risk is not an issue but where qualification requirements are nevertheless in place.

We address this puzzling observation by constructing a model as simple as possible without the risk of qualification failure. We then analyze whether an auctioneer should demand proof of bidders' qualification before or after the auction. Under pre-auction qualification, there is an exclusion effect. Depending on the cost realization of the bidder, his expected surplus could still be below the qualification costs every bidder has to pay. Under post-auction qualification, only the winner has to undergo costly qualification. Still, the bidder will keep in mind that he has to pay the qualification costs after the auction and therefore increases his offer by exactly this amount.

We show that interestingly, pre-auction qualification is more profitable if the qualification cost is sufficiently high. If qualification costs rise, the cost of participation increases. However, less bidders participate which means that the winning probability increases. This increase in the winning probability dampens the increase of the exclusion effect and the marginal increase goes to zero. With post-auction qualification, bids increase linearly with the increase in qualification costs and the marginal increase is one for all cost levels. Thus, pre-auction qualification yields higher revenues.

Finally, in Chapter 7 with the title *Bid Pooling in Reverse Multi-Unit Dutch Auctions – An Experimental Investigation*[6], which is joint work with Alexander Rasch, Peter Werner and Achim Wambach, we move from single-unit to multi-unit auctions. We experimentally investigate reverse multi-

[6]This chapter is published as Gillen et al. (2016). Financial support of the German Research Foundation (DFG) through the Gottfried Wilhelm Leibniz Prize awarded to Axel Ockenfels and through the Research Unit "Design & Behavior" (FOR 1371) is gratefully acknowledged.

unit Dutch auctions in which bidders compete to sell their single unit to a buyer who wants to purchase several objects. We show that this auction format is prone to higher prices than predicted by standard theory and is characterized by bid pooling. Furthermore, we set up a theoretical framework to show that these experimental results can be organized by boundedly rational bidding strategies. We distinguish between myopic bids consisting of a simple backward-looking heuristic and sophisticated bids where agents anticipate the behavior of others and choose their optimal bids according to their expectations but may make mistakes.

Our study yields three insights. (i) Bids are substantially higher than Nash equilibrium bids predicted by standard economic theory; (ii) these higher-than-predicted prices gradually decline in later periods; and (iii) bid pooling (or simultaneous bidding) is frequently observed – the majority of bidders submit their bids immediately after the first bidder has sold his unit.

CHAPTER 2

PROCUREMENT DESIGN WITH LOSS AVERSE BIDDERS

Abstract

We show that it is beneficial for a buyer to conduct a multi-stage mechanism if bidders are loss averse. In a first step, we derive a revenue equivalence principle. Fixing the multi-stage structure, the revenue is independent of the chosen payment rule. Secondly, we introduce a simple two-stage mechanism which always leads to a decrease in procurement costs compared to any single-stage auction. Finally we derive the optimal efficient two-stage mechanism.

2.1 INTRODUCTION

Procurement plays an important role both in the public and private sector. In Europe public procurement represented around 17% of the GDP in 2007.[1] In many sectors of the industry the role of procurement is even more pronounced. The consulting company Oliver Wyman reports that suppliers are responsible for roughly 60% of the value added of a car.[2] Hence, small savings in average procurement costs translate to a substantial increase in overall profit margins.

In the past few decades reverse auctions have been established as one of the main tools to select suppliers and to determine prices in many industries. Depending on factors like size or complexity of a project, the procurement designer usually commits to a certain auction format. In the academic literature on auctions, it is typically assumed that the auction designer chooses between a first-price or second-price payment rule and decides if

[1] Internal Market Scoreboard, n^o 19, July 2009

[2] `https://www.oliverwyman.com/our-expertise/industries/automotive/procurement.html`

she wants to conduct a static or dynamic auction. In the static formats, each bidder submits a sealed bid and the lowest bidder gets the contract. The dynamic formats typically considered are the Dutch auction and the English auction. In the English auction the price is decreased over time and bidders can drop out. It ends when the second-last bidder drops out. The winner is the last active bidder and he is paid the last displayed price. In the English auction the price increases over time and the first bidder who accepts the current price receives the contract and is paid the accepted price. In addition to the four auction formats described, the auction designer could also determine the number of stages.

In single-stage auctions, suppliers hand in an offer once and the contract is allocated based on these offers. In multi-stage auctions, the first rounds are usually conducted to reduce the set of suppliers that can participate in the final round.[3] Talks with practitioners suggest that especially in strategically important projects, multi-stage auctions are the preferred choice.

Interestingly, economic theory suggests that the use of multi-stage mechanisms cannot increase revenues above those that are achievable by one-stage mechanisms[4] when agents have standard preferences. However, if bidders are loss averse, the auction designer can increase her revenue by conducting multi-stage mechanisms. Proceeding to the next stage affects a bidder's winning probability and he therefore adjusts his reference point. The auction designer can exploit her influence on the bidders' reference points. Following Kőszegi and Rabin (2006), we assume reference points are based on rational expectations.[5]

A supplier who proceeds to the final stage of the multi-stage mechanism updates his winning probability. He knows that winning is now more likely

[3]Note that in these mechanisms, suppliers are typically restricted to hand in (weakly) more attractive offers in subsequent rounds.

[4]We consider settings in which the time between the different stages is rather short and suppliers cannot adjust their product during the auction.

[5]There is an ongoing debate on how the reference point is formed. Some studies suggest that it is mainly driven by expectations, whereas others hold that it is mostly given by the status quo. For a discussion, see Heffetz and List (2014) and references therein.

than before. Loss aversion implies that such a bidder gets more attached to winning and is willing to make a more attractive offer, since losing in the final round would cause a high disutility. These additional gains and losses are anticipated by the agent before the auction and factored into his first-round bid. A straightforward way of implementing such a mechanism is by conducting a two-stage tournament. Suppliers compete in two semifinals and only the best supplier of each semifinal proceeds to the final stage.[6]

In line with von Wangenheim (2019), we assume that bidders evaluate outcomes in two dimensions, a money dimension and a good dimension.[7] Consider a key account manager working for a supplier of a car manufacturer. When competing for a strategically important contract, he thinks in two independent dimensions: In the money domain, all monetary details such as his own costs, negotiated piece prices, investments etc. are captured. Independent of these details, the manager evaluates his chances of winning the contract and therefore getting a high level of recognition within his company. If this is the case, the buyer of the car manufacturer could exploit this behavior when designing her procurement mechanism.

In this paper, we first derive a revenue equivalence principle for bidders that are loss averse in the good domain. For a fixed multi-stage structure, meaning which and how many bidders advance in the individual stages, the auctioneer's revenue is not dependent on the payment rule she chooses. This result considerably simplifies the analyses and allows us to concentrate on the structure of multi-stage mechanisms. Furthermore, as a side result, this entails that all single-stage static auctions lead to the same expected costs.

The main result of this paper is that the symmetric two-stage tournament always leads to a decrease in procurement costs compared to any (standard) single-stage auction. This result is robust, as it does not require

[6]If the number of suppliers is odd, one can conduct semifinals that are symmetric in expectation.

[7]Lange and Ratan (2010) compare how the consideration of a one-dimensional reference point differs from the consideration of a two-dimensional reference point. They show that it can strongly affect predictions and argue that in most real world settings the consideration of a two-dimensional reference point is more reasonable.

knowledge about bidders' loss aversion. Hence, by conducting such a mechanism, the procurement designer's revenue strictly improves compared to all standard auctions if agents are loss averse, and makes no difference if not.

Finally, we derive the optimal efficient two-stage mechanism. When conducting two-stage mechanisms the procurement designer is confronted with a trade-off: On the one hand, she wants to maximize the attachment to winning the contract, and hence induce large winning probabilities to low-cost types. On the other hand, she cannot neglect high-cost types, either. If high-cost types have an already very low chance of winning the project, they might insure themselves from a deviation from their expectation by bidding even lower. Taking into account the bidders' degree of loss aversion, the optimal mechanism thus creates the level of uncertainty that optimally solves this trade-off.

2.2 RELATED LITERATURE

Our paper contributes to the literature on expectations-based loss aversion. The concept of loss aversion has been studied since the seminal paper of Kahneman et al. (1990). In their paper, they introduce the endowment effect and experimentally show that a subjects' valuation for a certain good increases when they are physically endowed with the good. According to this strand of literature subjects have a reference point and a deviation from this reference point in direction of losses has a larger impact on utility than a deviation in direction of gains.

A discussion around the formation of these reference points has risen in the literature. Kőszegi and Rabin (2006) suggest that the reference point is based on rational expectations. In an auction, this means that bidders have a certain probability of winning in mind and feel losses and gains compared to these expectations. As a consequence, a bidder expecting to win a good with a high probability suffers more from not winning than if he gauged his chances of winning as slim.

Our paper is most closely related to von Wangenheim (2019), who compares a sealed-bid second-price auction to an English auction assuming that bidders are loss averse and that their reference point is given by rational expectations. While both formats are strategically equivalent in independent private value settings if bidders have standard preferences, he shows that the second-price sealed-bid auction dominates the English auction if bidders are loss averse. The intuition is as follows: At the beginning of the English auction a bidder has the same chance of winning as in the second-price sealed-bid auction. However, during the course of the English auction the winning probability decreases and the bidder becomes less attached to the good. As a consequence, his willingness to pay decreases and he will drop out before the price is reached that he would have bid in the second-price sealed-bid auction.

Similar to von Wangenheim (2019), Ehrhart and Ott (2014) compare two standard auction formats for bidders with reference-dependent preferences. Comparing the Dutch auction to the English auction they show that the Dutch auction outperforms the English auction. The intuition is closely related to von Wangenheim (2019) and to our paper. For a given valuation a bidder has the same winning probability at the beginning of the Dutch auction and the English auction. However, while the winning probability decreases during the course of the English auction, it increases during the course of the Dutch auction. Hence, the attachment to the good is larger in the Dutch auction and bidders are thus willing to bid more aggressively. Similarly, a bidder who advances a stage in our setting also updates his winning probability and therefore his attachment to the good increases. This, in return, increases the bid he is willing to submit.

Banerji and Gupta (2014) and Rosato and Tymula (2019) provide evidence for the effect of expectations-based loss aversion in auction environments. In a setting in which participants compete in a second-price auction for a real good, they observe that bidders bid less when their winning probability was smaller. This observation stands in contrast to the predictions of standard theory which implies that subjects have a dominant strategy of bidding their true valuation independent of their winning probability. In contrast to that, loss aversion implies that a bidder with a higher chance of winning is more attached to the good and, hence, willing to bid more.

In contrast to the existing paper on auctions with loss averse bidders, we do not concentrate on comparing standard auction formats but investigate the following question: How can an auctions designer exploit bidders' loss aversion to increase her revenue?

Given this research question our work is also related to Maskin and Riley (1984) who also investigate how the auction designer can increase her revenue if bidders have a behavioral bias, in their case risk aversion. Similar to us, they present an optimal mechanism that needs to be fine-tuned to bidders' risk preferences and seems too complex to be implemented in practice. While our management implication is that simple two-stage mechanisms outperform one-stage auctions if bidders are loss averse, they show that first-price auctions outperform second-price auctions if bidders are risk averse.

Another related paper is Engelbrecht-Wiggans and Katok (2007). They analyze how the auction designer can exploit regret aversion of bidders. They show that the right information design, namely revealing the best bid but concealing all other bids, allows the auction designer to increase her revenue.

2.3 MODEL

In this section, we introduce the formal model. We consider $n \geq 2$ ex-ante symmetric bidders competing for one indivisible good. The value v_i of bidder $i \in \{1, \ldots, n\}$ for the good is privately drawn from a distribution F, $v_i \overset{\text{iid.}}{\sim} F[0,1]$. F is assumed to have a differentiable density f which is strictly positive on its support $[0,1]$. Moreover, F is common knowledge. Bids are placed after learning the value for the good.

For loss aversion we follow Kőszegi and Rabin (2006). We assume that bidders are loss averse in the good domain g representing the item the winner of the auction receives.[8] Furthermore, we assume bidders to be narrow-bracketers, following the definition of von Wangenheim (2019). Let x^m be the price a bidder pays if he wins and x^g a binary variable that is equal to one if the bidder wins the good and zero else. For an outcome $x = (x^c, x^g)$, valuation v for the good, and the reference consumption $r^g \in \{0, 1\}$, agent's utility is given by

$$u(x \middle| r^g) = x^c + vx^g + \mu^g(vx^g - vr^g). \tag{1}$$

Following Kőszegi and Rabin (2006), we assume μ^g to be a piecewise linear function with a kink at zero,

$$\mu^g(y) = \begin{cases} \eta^g y & \text{if } y \geq 0 \\ \lambda^g \eta^g y & \text{if } y < 0. \end{cases} \tag{2}$$

Here μ^g denotes the gain-loss utilities in the good dimension, where $\eta^g > 0$ and $\lambda^g > 1$. We assume non-dominance-of-gain-loss-utility, which means for a multi-stage mechanism with k stages $\eta^g(\lambda^g - 1) \leq 1/k$.[9] The importance of the non-dominance-of-gain-loss-utility bounds on η^i and λ^i are laid out in Herweg et al. (2010). To summarize, if $\eta^g(\lambda^g - 1) > 1/k$, a decision maker might choose stochastically dominated choices because he ex-ante expects to experience a net loss. For example, such a decision maker might choose a payment of zero over a lottery with slim chances of winning a strictly positive amount of money to avoid the disappointment, should he lose.

The interpretation of this gain-loss utility is that bidders perceive, in addition to their classical utility, a feeling of gain or loss, depending on the deviation from their reference consumption.

The reference point in our paper is assumed to be determined by rational expectations following Kőszegi and Rabin (2006).

[8] We assume that bidders are not loss averse in the money domain. This assumption is in line with Horowitz and McConnell (2003), who argue that the endowment effect is "highest for non-market goods, next highest for ordinary private goods, and lowest for experiments involving forms of money."

[9] This bound for non-dominance-of-gain-loss-utility is derived in Section 2.4.1.

2.3.1 EQUILIBRIUM CONCEPT

Following von Wangenheim (2019), we adapt Kőszegi and Rabin (2006)'s equilibrium concept under uncertainty, according to which bidders form their strategy after learning their valuation. We apply the concept of unacclimated personal equilibria, which is, as argued by Kőszegi and Rabin (2006), the appropriate concept in auction settings. Fixing the opponents' strategies, let $H(b, v_i)$ denote i's payoff distribution given his draw v_i from a continuous distribution $F(v)$ and his bid b. A bid b^* constitutes an unacclimated personal equilibrium (UPE), if for all b

$$U[H(b^*, v_i)\Big|H(b^*, v_i)] \geq U[H(b, v_i)\Big|H(b^*, v_i)]. \tag{3}$$

That means, given your reference point is determined by the payoff distribution resulting from an (exogenous) bid b^*, it is a best response to bid b^*.

2.3.2 MULTI-STAGE MECHANISMS

In a multi-stage mechanism, bidders participate in k stages and submit a bid in each one of them. The rules of the mechanism include how many stages there are and which bidders advance to the next stage. Bids are binding and cannot be lowered in subsequent stages.

As an example, consider four bidders and a mechanism with two stages. In the first stage, the semi-final, all four bidders submit an offer. The two bidders with the highest offers then advance to the final, where they submit another offer. The highest offer in the final is then the winner of the auction.[10]

In this section, we introduce the formal notation for multi-stage mechanisms. To completely characterize a multi-stage structure, we need to define the number of stages k and for each of the k stages, which of the bidders advance to the next stage. For N bidders, let $\mathfrak{B} = \{{}^jB_1, {}^jB_2, \ldots, {}^jB_N\}$ be the set of bids for each bidder in a stage $j \leq k$. We restrict ourselves to multi-stage mechanisms that are symmetric in expectation. This means that in

[10] We call this mechanism the "play-offs", it is analysed in section section 2.4.1.

some stage j of the mechanism, each bidder has the same expectation of number of opponents he is facing even if there are asymmetric groups.[11] Borrowing from order statistics notation, a multi-stage mechanism is then defined by $\left(\boldsymbol{\mu}, \mathfrak{M}\right)$, with $\boldsymbol{\mu}$ the payment rule and

$$\mathfrak{M} = \left\{ \underbrace{\left\{ o_1, \overbrace{\bigcup_{i=1}^{a_1} \left\{ {}^1B_i^{(o_1)} \right\}}^{\substack{\text{set of bidders}\\ \text{advancing to}\\ \text{stage 2}}} \right\}}_{\text{Stage 1}}, \underbrace{\left\{ o_2, \bigcup_{i=1}^{a_2} \left\{ {}^2B_i^{(o_2)} \right\} \right\}}_{\text{Stage 2}}, \ldots, \underbrace{\left\{ o_k, \overbrace{\left\{ {}^kB_1^{(o_k)} \right\}}^{\substack{\text{highest bidder}\\ \text{wins the good}\\ \text{in final stage}}} \right\}}_{\text{Stage k}} \right\}. \tag{4}$$

Here the o_j are the number of bidders per subgroup in stage j and a_j the number of bidders advancing from stage j to $j+1$.[12] It must hold that $a_i \leq o_i$ and $o_j \leq N$ where N is the total number of bidders.

[11] If there are asymmetric groups, the probability of being matched to a specific group has to be stated.

[12] This implies that only the highest o_{j+1} bidders of each subgroup advance from stage j to $j+1$.

2.4 ANALYSIS

The theory section is structured as follows. In Section 2.4.1 we derive general properties of the equilibrium bidding behavior in one- and multi-stage mechanisms. In Section 2.4.2 we show that fixing the multi-stage structure implies a revenue equivalence principle: the chosen payment rule does not affect the expected revenue of a mechanism. We then present a robust, easily implementable improvement over one-stage mechanisms in Section 2.4.3 and finally derive the optimal efficient two-stage mechanism in Section 2.4.4.

2.4.1 BIDDING BEHAVIOR

One-Stage Mechanisms

Assume that the bidders have standard preferences and let bidders participate in a standard auction A.[13] Further assume that the other bidders bid according to an increasing and absolutely continuous bidding function β^A. The payment rule of the auction is denoted by $\boldsymbol{\mu}^A(b_i, b_{-i})$ and the expected payment by $m^A(b_i)$. Define $G(b) := F_1^{(N-1)} \circ \beta^{A^{-1}}(b)$ the winning probability with a bid b in the auction. Then the expected utility of bidder i having value v_i and bidding b is given by

$$u_i^A(v_i, b) = G(b)v - m^A(b). \tag{5}$$

We now introduce loss aversion with bidders being loss averse only in the good domain. Given a reference bid b^*, the expected utility is then given

[13] Krishna (2009) defines a standard auction as an auction where the person who bids the highest amount is awarded the object.

by

$$\begin{aligned} u_i^A(v_i, b|b^*) &= G(b)v - m^A(b) \\ \boxed{\text{feeling of gain, good domain}} \quad &+ G(b)(1 - G(b^*))\mu^g(v - 0) \qquad (6) \\ \boxed{\text{feeling of loss, good domain}} \quad &+ (1 - G(b))G(b^*)\mu^g(0 - v) \\ &= G(b)v - m^A(b) \\ &+ G(b)(1 - G(b^*))\eta^g v \qquad (7) \\ &+ (1 - G(b))G(b^*)\eta^g\lambda^g(-v) \end{aligned}$$

Bidders optimize u_i^A with respect to b.

Multi-Stage Mechanisms

As a first step, we show that in equilibrium, bidders submit the same bid in every stage of the mechanism if non-dominance-of-gain-loss-utility holds.

Proposition 1. *In a multi-stage mechanism, bidders submit the same bid in every stage.*

Proof. Consider bidder i. Assume the other bidders bid according to an increasing, absolutely continuous bidding function β_j^{MS}, where j denotes the stage. The structure of the multi-stage mechanism, i.e. how many bidders advance in the individual stages and how many opponents they face in each stage, is then encoded in the probabilities to reach the individual stages of the mechanism. Let ϕ_j be defined such that $\phi_j \circ F \circ \beta_j^{MS^{-1}}$ is the probability of reaching stage $j+1$ given the bidder reached stage j. [14] Let $\vec{b} = (b_1, b_2, \ldots, b_k)$ be the vector of bids of bidder 1. This means that the ex-ante probability to win the auction is given by

$$\text{Prob}^{\text{ex-ante}}(\text{win with } b) = \prod_{j=1}^{k} \phi_j(b_j) =: H\left(\vec{b}\right). \qquad (8)$$

Note that to simplify the notation, we define that advancing to stage $k+1$ means winning the auction.

[14]The $\phi_j \circ F$ are expressions of probability and thus inherit the properties of the original distribution functions.

It is useful to define the probability to reach stage l, given that the bidder reached stage i. Let $i < l$. Then Φ_i^l is given by

$$\text{Prob}\left(\text{get to } l \text{ with } \vec{b} \,\middle|\, \text{get to } i \text{ with } \vec{b}\right) \tag{9}$$

$$= \frac{\text{Prob}\left(\text{get to } l \text{ with } \vec{b} \,\&\, \text{get to } i \text{ with } \vec{b}\right)}{\text{Prob}\left(\text{get to } i \text{ with } \vec{b}\right)} \tag{10}$$

$$= \frac{\text{Prob}\left(\text{get to } l \text{ with } \vec{b}\right)}{\text{Prob}\left(\text{get to } i \text{ with } \vec{b}\right)} = \prod_{j=i}^{l-1} \phi_j(b_j) := \Phi_i^l\left(\vec{b}\right) \tag{11}$$

The probability to win the auction given the bidder reached stage i is given by

$$\text{Prob}(\text{win with } \vec{b} \mid \text{reached } i) = \Phi_i^{k+1}\left(\vec{b}\right). \tag{12}$$

For each stage l, given a reference bid b_l^*, the bidder experiences a gain-loss utility in expectation. On one hand, the bidder might win with his bid b_l but has expected to lose with the reference bid b_l^*. He then experiences a gain in the good domain with respect to the reference outcome. On the other hand, the bidder might lose in one of the stages with his bid b_l but has expected to win with the reference bid b_l^*. He then experiences a loss in the good domain. This holds true for every stage.

Consider a standard auction based payment rule, $\boldsymbol{\mu}^{MS}$. The expected payment of the multi-stage mechanism composed by the expected amount a bidder has to pay and the probability of him having to pay it,

$$m^{MS}\left(\vec{b}\right) = \text{Prob}\left(\text{having to pay with } \vec{b}\right) \mathbb{E}\left[\boldsymbol{\mu}^{MS} \,\middle|\, \vec{b}, \vec{b}_{-i}\right] \tag{13}$$

$$=: P^{\text{pay}}\left(\vec{b}\right) \mathbb{E}\left[\boldsymbol{\mu}^{MS} \,\middle|\, \vec{b}, \vec{b}_{-i}\right]. \tag{14}$$

For the first-, second-, ...-price auction, we have $P^{\text{pay}}\left(\vec{b}\right) = H\left(\vec{b}\right)$, while for the all-pay auction we have that $P^{\text{pay}}\left(\vec{b}\right) = 1$. Generally, $P^{\text{pay}}\left(\vec{b}\right)$ either depends linearly on the ϕ_i for $i \in \{1, \dots, k\}$ or is constant.[15] This

[15] The fringe case where $\mathbb{E}\left[\boldsymbol{\mu}^{MS} \,\middle|\, \vec{b}, \vec{b}_{-i}\right]$ consists of a lottery that explicitly depends on a ϕ_i with $i \in \{1, \dots, k-1\}$ is excluded here. The lottery may depend on $\vec{b}$.

means that it holds for all $j < k$,

$$\frac{\partial\, m^{MS}\left(\vec{b}\right)}{\partial\left(\phi_j\left(\vec{b}\right)\right)} \leq \frac{m^{MS}\left(\vec{b}\right)}{\phi_j\left(\vec{b}\right)}. \tag{15}$$

Combining the results from above, we arrive at the following utility function for loss averse bidders in multi-stage mechanisms,

$$\begin{aligned} u^{MS}(v_i, \vec{b} \mid \vec{b}^*) = &\, H\left(\vec{b}\right) v_i - m^{MS}\left(\vec{b}\right) \\ &+ \underbrace{\sum_{i=1}^{k} \Phi_1^{k+1}\left(\vec{b}\right)\left(1 - \Phi_i^{i+1}\left(\vec{b}^*\right)\right)\mu^g(v-0)}_{\text{expecting to win with } \vec{b}, \text{ to lose with } \vec{b}^*} \qquad (16)\\ &+ \underbrace{\sum_{i=1}^{k} \Phi_0^{i}\left(\vec{b}\right)\left(1 - \Phi_i^{i+1}\left(\vec{b}\right)\right)\Phi_i^{k+1}\left(\vec{b}^*\right)\mu^g(0-v)}_{\text{expecting to lose with } \vec{b}, \text{ to win with } \vec{b}^*} \\ = &\, H\left(\vec{b}\right) v_i - m^{MS}\left(\vec{b}\right) \\ &+ \sum_{i=1}^{k} \Phi_1^{k+1}\left(\vec{b}\right)\left(1 - \Phi_i^{i+1}\left(\vec{b}^*\right)\right)\eta^g v \qquad (17)\\ &+ \sum_{i=1}^{k} \left(\Phi_0^{i}\left(\vec{b}\right) - \Phi_0^{i+1}\left(\vec{b}\right)\right)\Phi_i^{k+1}\left(\vec{b}^*\right)\eta^g\lambda^g(-v). \end{aligned}$$

Note that we can bound m^{MS} from above depending on v_i and $\vec{b}^*$ since a bidder will not submit bids that result in a negative expected utility,

$$\begin{aligned} & u^{MS}(v_i, \vec{b} \mid \vec{b}^*) \overset{!}{>} 0 \\ \Rightarrow \quad & m^{MS}\left(\vec{b}\right) \overset{!}{<} H\left(\vec{b}\right) v_i \\ &\quad + \sum_{i=1}^{k} \Phi_1^{k+1}\left(\vec{b}\right)\left(1 - \Phi_i^{i+1}\left(\vec{b}^*\right)\right)\eta^g v \qquad (18)\\ &\quad + \sum_{i=1}^{k} \Phi_0^{i}\left(\vec{b}\right)\left(1 - \Phi_i^{i+1}\left(\vec{b}\right)\right)\Phi_i^{k+1}\left(\vec{b}^*\right)\eta^g\lambda^g(-v). \end{aligned}$$

Also note that the right-hand side does not contain b_j outside of ϕ_j for all $j \in \{1, \ldots, k{-}1\}$. This means that for the first $k{-}1$ stages, a bidder can directly optimize over the probability of advancing to the next stage instead of optimizing over the bids that induce probabilities. Our equilibrium concept is UPE, this implies that the first-order condition for $i \in \{1, \ldots, k-1\}$, is given by

$$\left.\frac{\partial u^{MS}(v_i, \vec{b} \mid \vec{b^*})}{\partial\left(\phi_i(b_i)\right)}\right|_{\vec{b^*}=\vec{b}} = \frac{\prod\limits_{j=1}^{k} \phi_j(b_j)}{\phi_i(b_i)} v_i - \frac{\partial m^{MS}\left(\vec{b}\right)}{\partial\left(\phi_i(b_i)\right)} \tag{19}$$

$$+ \frac{\partial}{\partial\left(\phi_i(b_i)\right)} \sum_{l=1}^{k} \prod_{j=1}^{k} \phi_j(b_j)\Big(1 - \Phi_l^{l+1}\left(\vec{b^*}\right)\Big)\eta^g v_i \Bigg|_{\vec{b^*}=\vec{b}} \tag{20}$$

$$+ \frac{\partial}{\partial\left(\phi_i(b_i)\right)} \sum_{l=1}^{k} \prod_{j=0}^{l-1} \phi_j(b_j)\Phi_l^{k+1}\left(\vec{b^*}\right)\eta^g \lambda^g(-v_i) \Bigg|_{\vec{b^*}=\vec{b}} \tag{21}$$

$$- \frac{\partial}{\partial\left(\phi_i(b_i)\right)} \sum_{l=1}^{k} \prod_{j=0}^{l} \phi_j(b_j)\Phi_l^{k+1}\left(\vec{b^*}\right)\eta^g \lambda^g(-v_i) \Bigg|_{\vec{b^*}=\vec{b}}. \tag{22}$$

We now rearrange the terms. (20) simplifies to

$$\sum_{l=1}^{k} \frac{\prod_{j=1}^{k} \phi_j(b_j)}{\phi_i(b_i)} \Big(1 - \phi_l(b_l)\Big)\eta^g v_i. \tag{23}$$

For (21), we get

$$\sum_{l=i+1}^{k} \frac{\prod_{j=1}^{l-1} \phi_j(b_j)}{\phi_i(b_i)} \prod_{j=l}^{k} \phi_j(b_j)\eta^g \lambda^g(-v_i) = \sum_{l=i+1}^{k} \frac{\prod_{j=1}^{k} \phi_j(b_j)}{\phi_i(b_i)} \eta^g \lambda^g(-v_i) \tag{24}$$

$$= \frac{\prod_{j=1}^{k} \phi_j(b_j)}{\phi_i(b_i)} \eta^g \lambda^g(-v_i)(k-i). \tag{25}$$

For (22), we get

$$-\sum_{l=i}^{k} \frac{\prod_{j=1}^{l} \phi_j(b_j)}{\phi_i(b_i)} \prod_{j=l}^{k} \phi_j(b_j)\eta^g \lambda^g(-v_i) = -\sum_{l=i}^{k} \frac{\prod_{j=1}^{k} \phi_j(b_j)}{\phi_i(b_i)} \phi_l(b_l)\eta^g \lambda^g(-v_i). \tag{26}$$

Define

$$\alpha := \frac{\prod_{j=1}^{k} \phi_j(b_j)}{\phi_i(b_i)}. \tag{27}$$

We arrive at

$$
\begin{aligned}
\left.\frac{\partial u^{MS}(v_i,\vec{b}\mid\vec{b}^*)}{\partial\left(\phi_i(b_i)\right)}\right|_{\vec{b}^*=\vec{b}} = & -\frac{\partial m^{MS}\left(\vec{b}\right)}{\partial\left(\phi_i(b_i)\right)} \\
& + \alpha v_i + \eta^g v_i \alpha \sum_{l=1}^{k}(1-\phi_l(b_l)) \qquad (28) \\
& - \eta^g\lambda^g v_i\alpha(k-i) + \eta^g\lambda^g v_i \sum_{l=i}^{k}\phi_l(b_l) \\
\geq & -\frac{m^{MS}\left(\vec{b}\right)}{\phi_i(b_i)} + \alpha v_i + \eta^g v_i\alpha\sum_{l=1}^{k}(1-\phi_l(b_l)) \qquad (29) \\
& - \eta^g\lambda^g v_i\alpha(k-i) + \eta^g\lambda^g v_i \sum_{l=i}^{k}\phi_l(b_l) \\
\geq & -\left[H\left(\vec{b}\right)v_i + \sum_{i=1}^{k}\Phi_1^{k+1}\left(\vec{b}\right)\left(1-\Phi_i^{i+1}\left(\vec{b}^*\right)\right)\eta^g v_i\right. \\
& \left.+\sum_{i=1}^{k}\Phi_0^{i}\left(\vec{b}\right)\left(1-\Phi_i^{i+1}\left(\vec{b}\right)\right)\Phi_i^{k+1}\left(\vec{b}^*\right)\eta^g\lambda^g(-v_i)\right]_{\vec{b}^*=\vec{b}} \qquad (30) \\
& + \alpha v_i + \eta^g v_i\alpha\sum_{l-1}^{k}(1-\phi_l(b_l)) \\
& - \eta^g\lambda^g v_i\alpha(k-i) + \eta^g\lambda^g v_i \sum_{l=i}^{k}\phi_l(b_l) \\
= & -\alpha v_i + \eta^g(\lambda^g-1)v_i\alpha\sum_{l=1}^{k}(1-\phi_l(b_l)) \\
& + \alpha v_i + \eta^g v_i\alpha\sum_{l-1}^{k}(1-\phi_l(b_l)) \qquad (31) \\
& - \eta^g\lambda^g v_i\alpha(k-i) + \eta^g\lambda^g v_i \sum_{l=i}^{k}\phi_l(b_l) \\
= & \ \eta^g\lambda^g v_i\alpha(i-\sum_{l=1}^{i-1}\phi_l) > 0. \qquad (32)
\end{aligned}
$$

Note that we need to make sure that the expression in the brackets in step (30) is positive for all ϕ_j. This means it needs to hold that

$$\alpha v_i - \eta^g(\lambda^g - 1) v_i \alpha \sum_{l=1}^{k} (1 - \phi_l(b_l)) \overset{!}{\geq} 0 \tag{33}$$

$$\Leftrightarrow \quad -\eta^g(\lambda^g - 1) \overset{!}{\geq} \frac{-1}{\sum_{l=1}^{k}(1 - \phi_l(b_l))} \tag{34}$$

$$\Leftrightarrow \quad \eta^g(\lambda^g - 1) \overset{!}{\leq} \frac{1}{\sum_{l=1}^{k}(1 - \phi_l(b_l))} \tag{35}$$

$$\Leftrightarrow \quad \eta^g(\lambda^g - 1) \overset{!}{\leq} \min_{\phi} \frac{1}{\sum_{l=1}^{k}(1 - \phi_l(b_l))} \tag{36}$$

$$\Rightarrow \quad \eta^g(\lambda^g - 1) \overset{!}{\leq} \frac{1}{k}. \tag{37}$$

For every stage, a bidder experiences gain-loss utility. All-in-all, this means that the non-dominance of gain-loss utility has to hold for every stage, in total $\eta^g(\lambda^g - 1) \overset{!}{\leq} \frac{1}{k}$.

Interpreting ϕ_j as the distribution of bids that a bidder needs to beat in expectation to order to advance to stage $j + 1$, (32) implies that a bidder will always want to induce the highest possible probability to advance to the final stage with his bid $\vec{b}$. This implies that a bidder will cap his bids in stages 1 to $k - 1$ by the bid he submits in the final, pay-off relevant stage. A bidder therefore optimizes

$$\begin{aligned} u^{MS}(v_i, b|b^*) = {} & G(b)v_i - m^{MS}(b) \\ & + \sum_{i=1}^{k} \Phi_1^{k+1}(b)\Big(1 - \Phi_i^{i+1}(b^*)\Big)\eta^g v \\ & + \sum_{i=1}^{k} \Big(\Phi_0^{i}(b) - \Phi_0^{i+1}(b)\Big)\Phi_i^{k+1}(b^*)\eta^g\lambda^g(-v) \end{aligned} \tag{38}$$

over b.

□

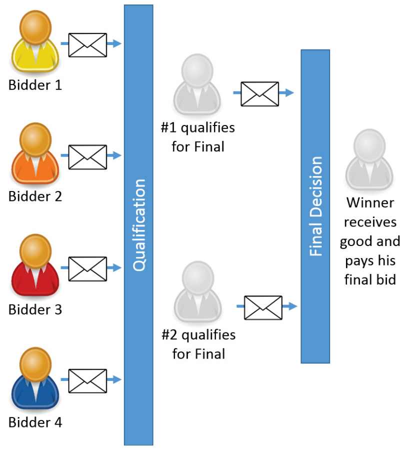

Figure 2.1: The first-price sealed-bid play-offs.

Example: First-Price Sealed-Bid Play-Offs

To get an idea what such a multi-stage mechanism can look like and of how to apply what we have derived so far, let us take a look at the following multi-stage mechanism with four bidders. As can be seen in *Figure 2.1*, the *FPSB play-offs* consists of 2 stages.

1. Out of the four bidders, the two highest bidders are advancing to the second stage.

2. Out of the two remaining bidders, the highest bid wins.

We can write $\mathfrak{M}^{PO}$ as

$$\mathfrak{M}^{PO} = \left\{ \underbrace{\left\{ 4, \left\{ B_1^{(4)}, B_2^{(4)} \right\} \right\}}_{\text{Stage 1}}, \underbrace{\left\{ 2, \left\{ B_1^{(2)} \right\} \right\}}_{\text{Stage 2}} \right\}. \tag{39}$$

The payment rule $\boldsymbol{\mu}$ is given by the first-price auction payment rule. With proposition 1, we can assume bidders to bid the same in every stage. Assume the other bidders bid according to an increasing, absolutely continuous bidding function β^P. In the first stage, bidder i advances if he beats at least the second highest opponent. This yields

$$\phi_1 \circ F = F_2^{(3)} = 3F^2 - 2F^3. \tag{40}$$

Given that the bidder reached stage two, the bidder wins if he beats his strongest opponent,

$$\phi_2 \circ F \circ \beta^{P^{-1}}(b) = \text{Prob}\left(b > \beta^P\left(v_1^{(3)}\right) \,\middle|\, b > \beta^P\left(v_2^{(3)}\right)\right) \tag{41}$$

$$= \frac{F\left(\beta^{P^{-1}}(b)\right)^3}{3F\left(\beta^{P^{-1}}(b)\right)^2 - 2F\left(\beta^{P^{-1}}(b)\right)^3}. \tag{42}$$

The underlying auction format is the first-price auction, the expected payment is given by $m^T(b) = G(b)b$. The utility is then given by

$$\begin{aligned} u^P(v_i, b|b^*) = \; & G(b\left(v_i - b\right) \\ & \underbrace{+F_1^{(3)}\left(\beta^{P^{-1}}(b)\right)\left(1 - F_2^{(3)}\left(\beta^{P^{-1}}(b^*)\right)\right)\eta^g v}_{\text{win but would have lost in stage 1 with } b^*} \\ & \underbrace{+F_1^{(3)}\left(\beta^{P^{-1}}(b)\right)\left(1 - \frac{F_1^{(3)}\left(\beta^{P^{-1}}(b^*)\right)}{F_2^{(3)}\left(\beta^{P^{-1}}(b^*)\right)}\right)\eta^g v}_{\text{win but would have lost in stage 2 with } b^*} \\ & \underbrace{+\left(1 - F_2^{(3)}\left(\beta^{P^{-1}}(b)\right)\right)F_1^{(3)}\left(\beta^{P^{-1}}(b^*)\right)\eta^g \lambda(-v)}_{\text{don't advance to 2nd stage but would have won with } b^*} \\ & \underbrace{+F_2^{(3)}\left(\beta^{P^{-1}}(b)\right)\left(1 - \frac{F_1^{(3)}\left(\beta^{P^{-1}}(b)\right)}{F_2^{(3)}\left(\beta^{P^{-1}}(b)\right)}\right)\frac{F_1^{(3)}\left(\beta^{P^{-1}}(b^*)\right)}{F_2^{(3)}\left(\beta^{P^{-1}}(b^*)\right)}\eta^g \lambda(-v)}_{\text{get to 2nd stage \& lose but would have won with } b^*}. \end{aligned} \tag{43}$$

We are interested in finding the equilibrium bidding function for this multi-stage auction. Our equilibrium concept is UPE, this implies that the first-order condition is given by

$$\left.\left(\frac{\partial u^P(v_i, b|b^*)}{\partial b}\right)\right|_{b^* = \beta^P(v_i)} \stackrel{!}{=} 0. \tag{44}$$

In equilibrium it holds that $b = \beta^P(v_i)$. To simplify notation, let $F_m^{(3)} =: F_m$. The resulting ordinary differential equation admits a closed form solution,

$$\beta^P(v_i) = \frac{1}{F_1(v_i)} \int_0^{v_i} s \Bigg(f_1(s) + \eta^g \lambda^g \Bigg(f_2(s) F_1(s) - \Big(f_2(s) - f_1(s) \Big) \frac{F_1(s)}{F_2(s)} \Bigg) + \eta^g f_1(s) \left(2 - \frac{F_1(s)}{F_2(s)} - F_2(s) \right) \Bigg) ds. \tag{45}$$

2.4.2 REVENUE EQUIVALENCE PRINCIPLE

In this section, we show that once we fix the multi-stage structure of the procurement mechanism, a revenue equivalence principle holds. This means that an auctioneer does not need to worry about the payment rule of her mechanism.[16]

Proposition 2 (Revenue equivalence principle for loss averse bidders)**.** *Suppose that values are independently and identically distributed and that bidders are loss averse in the good domain. Fix the multi-stage structure* $\mathfrak{M}$*. For every standard auction payment rule* $\boldsymbol{\mu}$*, any symmetric and increasing equilibrium such that the expected payment of a bidder with value zero is zero, yields the same expected revenue to the seller.*

Proof. Consider multi-stage mechanism $MS = \Big(\boldsymbol{\mu}, \mathfrak{M}\Big)$, with $\boldsymbol{\mu}$ some standard auction payment rule, and fix a symmetric, strictly increasing equilibrium bidding function β^{MS}. Let $m^{MS}(v_i)$ be the equilibrium expected payment in the mechanism by bidder i with value v_i. Suppose that β^{MS} is such that $m^{MS}(0) = 0$. Define the ex-ante expected gain-loss utility in the good domain Θ^g such that

$$\begin{aligned} \Theta^g(b|\, b^*) := & \sum_{i=1}^{k} \Phi_1^{k+1}(b) \Big(1 - \Phi_i^{i+1}(b^*)\Big) \eta^g v \\ & + \sum_{i=1}^{k} \Big(\Phi_0^i(b) - \Phi_0^{i+1}(b)\Big) \Phi_i^{k+1}(b^*) \eta^g \lambda^g (-v), \end{aligned} \tag{46}$$

[16] We consider payment rules based on standard auctions as defined by Krishna (2009). A standard auction is an auction where the highest bidder wins.

yielding

$$u^{MS}(v_i, b|b^*) = G(b)v_i - m^{MS}(b) + \Theta^g(b|\, b^*). \tag{47}$$

Consider bidder i and suppose other bidders are following the equilibrium strategy β^{MS}. Consider the expected payoff of bidder i with value v_i deviating from the equilibrium bidding strategy. β^{MS} is bijective, meaning that every sensible bid b can be expressed such that $b = \beta^{MS}(z)$. The bidding function β^{MS} constitutes a UPE if and only if the utility function $u_i^{MS}(v_i, b|\beta^{MS}(v_i))$ attains its maximum at $b = \beta^{MS}(v_i)$ for all v_i. Bidder i's expected payoff is given by

$$\begin{aligned} u^{MS}\left(v_i, \beta^{MS}(z)|\beta^{MS}(v_i)\right) = {} & G(\beta^{MS}(z))v_i - m^{MS}(z) \\ & + \Theta\left(\beta^{MS}(z)|\, \beta^{MS}(v_i)\right). \end{aligned} \tag{48}$$

The first-order condition is given by

$$\begin{aligned} \frac{\partial u^{MS}\left(v_i, \beta^{MS}(z)|\beta^{MS}(v_i)\right)}{\partial z} = {} & f_1^{(N-1)}(z)v_i - \frac{\partial}{\partial z}m^{MS}(z) \\ & + \frac{\partial}{\partial z}\Theta\left(\beta^{MS}(z)|\, \beta^{MS}(v_i)\right) \overset{!}{=} 0. \end{aligned} \tag{49}$$

In equilibrium it is optimal to report $z = v_i$ and it holds that $b^* = \beta^{MS}(v_i)$, so we obtain that for all y,

$$\frac{\partial}{\partial y}m^{MS}(z) = f_1^{(N-1)}(y)y + \left(\frac{\partial}{\partial y}\Theta\left(\beta^{MS}(y)|\, \beta^{MS}(z)\right)\right)\Bigg|_{z=y}. \tag{50}$$

This means that

$$\begin{aligned} m^{MS}(v_i) = {} & \cancel{m^{MS}(0)} + \int_0^{v_i} f_1^{(N-1)}(y)y\,dy \\ & + \int_0^{v_i} \left(\frac{\partial}{\partial y}\Theta\left(\beta^{MS}(y)|\, \beta^{MS}(z)\right)\right)\Bigg|_{z=y} dy. \end{aligned} \tag{51}$$

While the right-hand side depends on the multi-stage structure $\mathfrak{M}$, it does not depend on the payment rule $\boldsymbol{\mu}$. □

The result holds for $k \geq 1$ stages, so one-stage mechanisms are included. A first application of the RET for loss averse bidders is to rank the English auction with loss averse bidders.

Proposition 3. *All static standard auction formats yield higher expected revenues with loss averse bidders than the English auction.*

Proof. From von Wangenheim (2019) we know that the English auction performs worse than the second-price auction revenue-wise. We can apply *Proposition 2* to complete the proof. □

2.4.3 A ROBUST IMPROVEMENT OVER ONE-STAGE MECHANISMS

A mechanism that is to be implemented in real-life and that exploits bidders' loss aversion should not depend on the parameters for loss aversion. An auctioneer cannot hope to be able to accurately estimate these parameters in a way that would help her design a mechanism. We will show that for a parameter space that includes the empirically found loss aversion parameters, it is beneficial for the seller to implement a simple two-stage mechanism for every value realization of every distribution function if there are more than two bidders.[17]

For an even number of bidders, $2N$, consider randomly pairing two groups of N bidders and then advance the highest bidder of each pairing to the final. For an odd number of bidders, $2N+1$, consider randomly pairing of one group of N bidders and one group of $N+1$ bidders. Bidders do not know in which group they are selected, they only know the a priori probability of being in the group with N bidders is 0.5. Again, the highest bidder of each pairing advances to the final. We call this multi-stage structure a tournament, it can be seen in figure 2.2. We can write $\mathfrak{M}^T$ as

$$\mathfrak{M}^{T,\text{even}} = \left\{ \underbrace{\left\{ N, \left\{ B_1^{(N)} \right\} \right\}}_{\text{Stage 1}}, \underbrace{\left\{ 2, \left\{ B_1^{(2)} \right\} \right\}}_{\text{Stage 2}} \right\}. \tag{52}$$

[17] See Gächter et al. (2007) for an empirical study on individual-level loss aversion. They present evidence that λ^g lies around 2.

$$\mathfrak{M}^{T,\text{odd}} = \left\{ \underbrace{\left\{ \{N_{P=\frac{1}{2}}, N+1_{P=\frac{1}{2}}\}, \left\{ B_1^{(N_{P=\frac{1}{2}}, N+1_{P=\frac{1}{2}})} \right\} \right\}}_{\text{Stage 1}}, \underbrace{\left\{ 2, \left\{ B_1^{(2)} \right\} \right\}}_{\text{Stage 2}} \right\}. \tag{53}$$

As shown in Proposition 2, the payment rule we choose is not relevant for the revenue. For the proof, we choose the first-price auction payment rule.

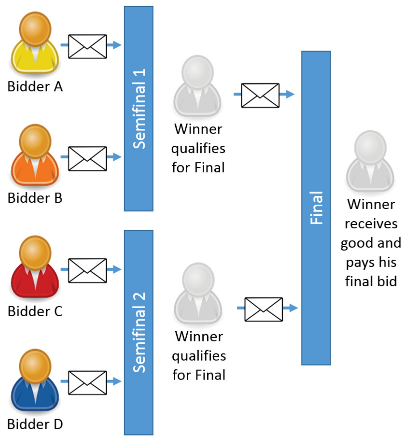

Figure 2.2: The tournament multi-stage structure $\mathfrak{M}^T$ for four bidders.

Proposition 4. *Assume an even number of bidders $2N \geq 4$ that are loss averse in the good domain. Assume that $\lambda \leq \frac{2N-1}{N-1}$. Then for all $\eta \geq 0$ the revenue is higher in the tournament than in any one-stage mechanism.*

Corollary 1. *Assume an even number of bidders $2N \geq 4$ that are loss averse in the good domain. Assume that $\lambda \leq \frac{2N-1}{N-1}$. In the case of the first-price, second-price or all-pay auction as underlying auction format, bids are higher in the tournament than in the corresponding one-stage mechanism for all types.*

Proposition 5. *Assume an odd number of bidders $2N+1 \geq 3$ that are loss averse in the good domain. Assume that $\lambda \leq \frac{4N}{2N-1}$. Then for all $\eta \geq 0$ the revenue is higher in the tournament than in any one-stage mechanism.*

Corollary 2. *Assume an even number of bidders $2N+1 \geq 3$ that are loss averse in the good domain. Assume that $\lambda \leq \frac{4N}{2N-1}$. In the case of the first-price, second-price or all-pay auction as underlying auction format, bids are higher in the tournament than in the corresponding one-stage mechanism for all types.*

Proposition 6. *For $\lambda^g \leq 2$, the tournament yields higher bids than the respective one-stage auction for all types.*

Proof of Proposition 4. Consider the first-price auction payment rule. We start with the one-stage mechanism. Assume the other bidders bid according to an increasing, absolutely continuous bidding function β^{FP} and let $G(b) = F_1^{(N-1)}\left(\beta^{FP^{-1}}(b)\right)$. The expected payment is given by

$$m^{FP}(b) = G(b)b. \tag{54}$$

The utility function is given by

$$\begin{aligned} u_i^{FP}(v_i, b|b^*) = G(b)(v-b) \\ + G(b)(1-G(b^*))\eta^g v \\ + (1-G(b))G(b^*)\eta^g\lambda^g(-v). \end{aligned} \tag{55}$$

The bidding function β^{FP} constitutes a UPE if and only if the utility function $u_i^{FP}(v_i, b|\beta^{FP}(v_i))$ attains its maximum at $b = \beta^{FP}(v_i)$ for all v_i. Differentiating u^{FP} with respect to b and plugging in the equilibrium condition $b = \beta^{FP}(v_i)$ yields the ODE,

$$\beta^{FP'}(v_i)F_1(v_i) + \beta^{FP}(v_i)f_1(v_i) \stackrel{!}{=} v_i f_1(v_i)\Big(1 + (1-F_1(v_i))\,\eta^g + F_1\eta^g\lambda^g\Big) \tag{56}$$

This ODE admits a closed form solution,

$$\begin{aligned} \beta^{FP}(v_i) &= \frac{1}{F_1(v_i)}\int_0^{v_i} s f_1(s)\Big(1+\eta^g+F_1(s)\eta^g(\lambda^g-1)\Big)ds & (57)\\ &= \frac{1}{F_1(v_i)}\int_0^{v_i} s f_1(s)\Big(1+\eta^g\Big(1-F_1(s)\Big)+\eta^g\lambda^g F_1(s)\Big)ds. & (58) \end{aligned}$$

The equilibrium bidding function for the tournament can be derived explicitly, too. With Proposition 1, we can assume bidders bid the same in every stage. Assume that the other bidders bid according to an increasing, absolutely continuous bidding function β^T. In the first stage, bidder i advances if he beats his $N-1$ opponents. This yields

$$\phi_1 \circ F = F_1^{(N-1)}. \tag{59}$$

This implies that advancing to the second stage is not informative in any way about the value of the remaining opponent. The intuition behind this can be understood by considering the mechanism with four bidders. Given the bidder won the first round, he may have beaten his toughest opponent already. But he also might have beaten the second or third highest bidding one,

$$\begin{aligned} &\text{Prob(get to 2nd round with } b) && (60)\\ &= \frac{1}{3}F_1(\beta^{T^{-1}}(b)) + \frac{2}{3}\Big(\frac{1}{2}F_2(\beta^{T^{-1}}(b)) + \frac{1}{2}F_3(\beta^{T^{-1}}(b))\Big) && (61)\\ &= F(\beta^{T^{-1}}(b)). && (62) \end{aligned}$$

Given that the bidder reached stage two, the bidder wins if he beats the winner of the second group given he got there,

$$\begin{aligned} \phi_2 \circ F \circ \beta^{T^{-1}}(b) &= \text{Prob}\left(b > \beta^T\left(v_1^{(N)}\right) \middle| \, b > \beta^T\left(v_1^{(N-1)}\right)\right) && (63)\\ &= \frac{F_1^{(N-1)}\left(\beta^{T^{-1}}(b)\right) F_1^{(N)}\left(\beta^{T^{-1}}(b)\right)}{F_1^{(N-1)}\left(\beta^{T^{-1}}(b)\right)} && (64)\\ &= F_1^{(N)}\left(\beta^{T^{-1}}(b)\right). && (65) \end{aligned}$$

As mentioned before, we have $m^T(b) = F_1^{(2N)}\left(\beta^{T^{-1}}(b)\right) b$. Then the utility is given by

$$\begin{aligned} u^T(v_i, b|b^*) =\ & F_1^{(2N-1)}\left(\beta^{T^{-1}}(b)\right)\left(v - b\right)\\ &+ F_1^{(2N-1)}\left(\beta^{T^{-1}}(b)\right)\left(1 - F_1^{(N-1)}\left(\beta^{T^{-1}}(b^*)\right)\right)\eta^g v_i\\ &+ F_1^{(2N-1)}\left(\beta^{T^{-1}}(b)\right)\left(1 - F_1^{(N)}\left(\beta^{T^{-1}}(b^*)\right)\right)\eta^g v_i\\ &+ \left(1 - F_1^{(N-1)}\left(\beta^{T^{-1}}(b)\right)\right) F_1^{(2N-1)}\left(\beta^{T^{-1}}(b^*)\right)(-\eta^g\lambda^g v_i)\\ &+ \left(F_1^{(N-1)}\left(\beta^{T^{-1}}(b)\right) - F_1^{(2N-1)}\left(\beta^{T^{-1}}(b)\right)\right) F_1^{(N)}\left(\beta^{T^{-1}}(b^*)\right)(-\eta^g\lambda^g v_i). \end{aligned} \tag{66}$$

We are interested in finding the equilibrium bidding function for this multi-stage auction. Our equilibrium concept is UPE, this implies that the first-order condition is given by

$$\left(\frac{\partial u^T(v_i, b|b^*)}{\partial b}\right)\Bigg|_{b^*=\beta^T(v_i)} \stackrel{!}{=} 0. \tag{67}$$

We have

$$\begin{aligned}
\frac{\partial}{\partial b}u^T(v_i, b|b^*)\Big|_{b^*=\beta^T(v_i)} =\; & f_1^{(2N-1)}\left(\beta^{T^{-1}}(b)\right)\left(v-b\right)\frac{1}{\beta^{T\prime}(\beta^{T^{-1}}(b))} \\
& - F_1^{(2N-1)}\left(\beta^{T^{-1}}(b)\right) \\
& + f_1^{(2N-1)}\left(\beta^{T^{-1}}(b)\right)\left(1-F_1^{(N-1)}(v_i)\right)\eta^g v_i \frac{1}{\beta^{T\prime}(\beta^{T^{-1}}(b))} \\
& + f_1^{(2N-1)}\left(\beta^{T^{-1}}(b)\right)\left(1-F_1^{(N)}(v_i)\right)\eta^g v_i \frac{1}{\beta^{T\prime}(\beta^{T^{-1}}(b))} \\
& + f_1^{(N-1)}\left(\beta^{T^{-1}}(b)\right)F_1^{(2N-1)}(v_i)\,\eta^g\lambda^g v_i \frac{1}{\beta^{T\prime}(\beta^{T^{-1}}(b))} \\
& - f_1^{(N-1)}\left(\beta^{T^{-1}}(b)\right)F_1^{(N)}(v_i)\,\eta^g\lambda^g v_i \frac{1}{\beta^{T\prime}(\beta^{T^{-1}}(b))} \\
& + f_1^{(2N-1)}\left(\beta^{T^{-1}}(b)\right)F_1^{(N)}(v_i)\,\eta^g\lambda^g v_i \frac{1}{\beta^{T\prime}(\beta^{T^{-1}}(b))}.
\end{aligned} \tag{68}$$

In equilibrium it holds that $b = \beta^T(v_i)$. The resulting ordinary differential equation for β^T admits a closed-form solution,

$$\begin{aligned}
\beta^T(v_i) = \; & \frac{1}{F_1^{(2N-1)}(v_i)}\int_0^{v_i} s\Bigg[f_1^{(2N-1)}(s) \\
& + \eta^g f_1^{(2N-1)}(s)\left(2-F_1^{(N-1)}(s)-F_1^{(N)}(s)\right) \\
& + \eta^g\lambda^g\Bigg(f_1^{(N-1)}(s)F_1^{(2N-1)}(s) \\
& - f_1^{(N-1)}(s)F_1^{(N)}(s) \\
& + f_1^{(2N-1)}(s)F_1^{(N)}(s)\Bigg)\Bigg]ds.
\end{aligned} \tag{69}$$

For $\beta^T(v_i) \geq \beta^{FP}(v_i)$ to hold for all v_i, a sufficient condition is that we can rank the arguments of the integrals. As a reminder, $\beta^{FP}(v_i)$ is given by

$$\beta^{FP}(v_i) = \frac{1}{F_1(v_i)}\int_0^{v_i} sf_1(s)\Big(1+\eta^g\Big(1-F_1(s)\Big)+\eta^g\lambda^g F_1(s)\Big)ds, \tag{70}$$

with $F_1 = F_1^{(2N-1)}$ and $f_1 = f_1^{(2N-1)}$. Note that the first term stemming from the standard preferences equilibrium bidding function is identical in both bidding functions. What is left are the gain-loss utility terms. This means it has to hold that

$$\begin{aligned} \eta^g f_1^{(2N-1)}(s)\left(2 - F_1^{(N-1)}(s) - F_1^{(N)}(s)\right) \\ -\eta^g f_1^{(2N-1)}(s)\Big(1 - F_1(s)\Big) \\ +\eta^g \lambda^g \Bigg(f_1^{(N-1)}(s) F_1^{(2N-1)}(s) - f_1^{(N-1)}(s) F_1^{(N)}(s) \\ +f_1^{(2N-1)}(s) F_1^{(N)}(s)\Bigg) \\ -\eta^g \lambda^g \eta^g f_1^{(2N-1)}(s) F_1(s) \overset{!}{\geq} 0. \end{aligned} \tag{71}$$

Note that all terms, except the third one, include $f_1^{(2N-1)}(s) = (2N - 1)F^{2N-2}(s)f(s)$. Using the definition of the first-order statistic density for distribution functions, we have

$$f_1^{(N-1)}(s)F_1^{(2N-1)}(s) = (N-1)F^{N-2}(s)f(s)F^{2N-1}(s) \tag{72}$$

$$= \frac{N-1}{2N-1}(2N-1)F^{2N-2}(s)f(s)F^{N-1} \tag{73}$$

$$= \frac{N-1}{2N-1} f_1^{(2N-1)}(s)F_1^{(N-1)}. \tag{74}$$

Similarly, we can write the third term as

$$f_1^{(N-1)}(s)F_1^{(2N-1)}(s) - f_1^{(N-1)}(s)F_1^{(N)}(s) + f_1^{(2N-1)}(s)F_1^{(N)}(s) \tag{75}$$

$$= f_1^{(2N-1)}(s)\left(\frac{N-1}{2N-1}F_1^{(N-1)}(s) - \frac{N-1}{2N-1} + F_1^{(N)}(s)\right). \tag{76}$$

With this, (71) simplifies to

$$(\lambda^g - 1)F_1^{(N)}(s)\cancel{\Big(1 - F_1^{(N-1)}(s)\Big)} - \left(\frac{N-1}{2N-1}\lambda^g - 1\right)\cancel{\Big(1 - F_1^{(N-1)}(s)\Big)} \overset{!}{\geq} 0 \tag{77}$$

$$\Leftrightarrow \quad (\lambda^g - 1)F_1^{(N)}(s) - \frac{N-1}{2N-1}\lambda^g + 1 \overset{!}{\geq} 0 \tag{78}$$

$$\Rightarrow \quad \frac{N-1}{2N-1}\lambda^g - 1 \overset{!}{\leq} 0 \tag{79}$$

$$\Leftrightarrow \quad \lambda^g \overset{!}{\leq} \frac{2N-1}{N-1}. \tag{80}$$

To prove the corollary, we define

$$\gamma^{OS}(s) = s\Big(1 + \eta^g\Big(1 - F_1(s)\Big) + \eta^g\lambda^g F_1(s)\Big). \tag{81}$$

Note that γ^{OS} is given by the argument of the integral of β^{FP}. Similarly, define γ^T as the argument of the integral of β^T. Note that we have shown under which conditions it holds that $\gamma^{OS}(s) \leq \gamma^T(s)$. It is straightforward to compute that in the case of the second-price auction payment rule, bidding functions are given by

$$\beta^{SP}(v) = \gamma^{OS}(v) \tag{82}$$

$$\beta^T(v) = \gamma^T(v). \tag{83}$$

In the case of the all-pay auction, the bidding functions are given by

$$\beta^{SP}(v) = \int_0^v \gamma^{OS}(s)f_1(s)ds \tag{84}$$

$$\beta^T(v) = \int_0^v \gamma^T(s)f_1(s)ds. \tag{85}$$

Combining the results from this section concludes the proof to the corollary.

□

Proof of Proposition 5. Again, we consider the first-price auction payment rule. We derive the equilibrium bidding function for the tournament with an odd number of bidders in a similar way as for an even number of bidders.

With proposition 1, we can assume bidders bid the same in every stage. Assume the other bidders bid according to an increasing, absolutely continuous bidding function β^T. In the group with N bidders, a bidder advances if he beats his $N-1$ paired opponents. This yields

$$\phi_1 \circ F = F^{N-1}. \tag{86}$$

Given that the bidder reached stage two, the bidder wins if he beats the winner of the second group with $N+1$ bidders, given he got there,

$$\phi_2 \circ F \circ \beta^{T^{-1}}(b) = F_1^{(N+1)}\left(\beta^{T^{-1}}(b)\right). \tag{87}$$

In the group with $N+1$ bidders, a bidder advances if he beats his N paired opponents. This yields

$$\phi_1 \circ F = F^{N}. \tag{88}$$

Given that the bidder reached stage two, the bidder wins if he beats the winner of the second group with N bidders, given he got there,

$$\phi_2 \circ F \circ \beta^{T^{-1}}(b) = F_1^{(N)}\left(\beta^{T^{-1}}(b)\right). \tag{89}$$

Again we have $m^T(b) = F_1^{(2N)}\left(\beta^{T^{-1}}(b)\right) b$. Then the utility is given by

$$\begin{aligned}
u^T(v_i, b|b^*) = \; & F_1^{(2N)}\left(\beta^{T^{-1}}(b)\right)\left(v - b\right) \\
+\frac{1}{2}\Bigg[& F_1^{(2N)}\left(\beta^{T^{-1}}(b)\right)\left(1 - F_1^{(N-1)}\left(\beta^{T^{-1}}(b^*)\right)\right)\eta^g v_i \\
& + F_1^{(2N)}\left(\beta^{T^{-1}}(b)\right)\left(1 - F_1^{(N+1)}\left(\beta^{T^{-1}}(b^*)\right)\right)\eta^g v_i \\
& + \left(1 - F_1^{(N-1)}\left(\beta^{T^{-1}}(b)\right)\right) F_1^{(2N)}\left(\beta^{T^{-1}}(b^*)\right)(-\eta^g\lambda^g v_i) \\
& + F_1^{(N-1)}\left(\beta^{T^{-1}}(b)\right) F_1^{(N+1)}\left(\beta^{T^{-1}}(b^*)\right)(-\eta^g\lambda^g v_i) \\
& - F_1^{(2N)}\left(\beta^{T^{-1}}(b)\right) F_1^{(N+1)}\left(\beta^{T^{-1}}(b^*)\right)(-\eta^g\lambda^g v_i)\Bigg] \\
+\frac{1}{2}\Bigg[& F_1^{(2N)}\left(\beta^{T^{-1}}(b)\right)\left(1 - F_1^{(N)}\left(\beta^{T^{-1}}(b^*)\right)\right)\eta^g v_i \\
& + F_1^{(2N)}\left(\beta^{T^{-1}}(b)\right)\left(1 - F_1^{(N)}\left(\beta^{T^{-1}}(b^*)\right)\right)\eta^g v_i \\
& + \left(1 - F_1^{(N)}\left(\beta^{T^{-1}}(b)\right)\right) F_1^{(2N)}\left(\beta^{T^{-1}}(b^*)\right)(-\eta^g\lambda^g v_i) \\
& + F_1^{(N)}\left(\beta^{T^{-1}}(b)\right) F_1^{(N)}\left(\beta^{T^{-1}}(b^*)\right)(-\eta^g\lambda^g v_i) \\
& - F_1^{(2N)}\left(\beta^{T^{-1}}(b)\right) F_1^{(N)}\left(\beta^{T^{-1}}(b^*)\right)(-\eta^g\lambda^g v_i)\Bigg].
\end{aligned} \tag{90}$$

The bracketed expression starting in the second line accounts for the case the bidder is sorted into the N-bidder group, the bracketed expression starting in the seventh line accounts for the case the bidder is sorted into the $N+1$-bidder group. We are interested in finding the equilibrium bidding function for this multi-stage auction. Our equilibrium concept is UPE, this implies that the first-order condition is given by

$$\left(\frac{\partial u^T(v_i, b|b^*)}{\partial b}\right)\Bigg|_{b^*=\beta^T(v_i)} \stackrel{!}{=} 0. \tag{91}$$

Leaving out the arguments of the functions for the sake of readability, we have

$$\begin{aligned}\beta^{T'}(\beta^{T^{-1}}) \cdot \frac{\partial}{\partial b} u^T(v_i, b|b^*)\Big|_{b^*=\beta^T(v_i)} &= f_1^{(2N)}\Big(v-b\Big) - F_1^{(2N-1)}\beta^{T'}(\beta^{T^{-1}}) \\ &+\frac{1}{2}\Bigg[f_1^{(2N)}\left(1-F_1^{(N-1)}\right)\eta^g v_i + f_1^{(2N)}\left(1-F_1^{(N+1)}\right)\eta^g v_i \\ &\quad + f_1^{(N-1)}F_1^{(2N)}\eta^g\lambda^g v_i - \left(f_1^{(N-1)} - f_1^{(2N)}\right)F_1^{(N+1)}\eta^g\lambda^g v_i\Bigg] \\ &+\frac{1}{2}\Bigg[f_1^{(2N)}\left(1-F_1^{(N)}\right)\eta^g v_i + f_1^{(2N)}\left(1-F_1^{(N)}\right)\eta^g v_i \\ &\quad + f_1^{(N)}F_1^{(2N)}\eta^g\lambda^g v_i - \left(f_1^{(N)} - f_1^{(2N)}\right)F_1^{(N)}\eta^g\lambda^g v_i\Bigg]. \end{aligned} \tag{92}$$

In equilibrium it holds that $b = \beta^T(v_i)$. The resulting ordinary differential equation for β^T admits a closed form solution,

$$\begin{aligned}\beta^T(v_i) = \frac{1}{F_1^{(2N)}(v_i)} \int_0^{v_i} s \Bigg[& f_1^{(2N)}(s) \\ &+ \frac{\eta^g}{2} f_1^{(2N)}(s)\left(4 - F_1^{(N-1)}(s) - F_1^{(N+1)}(s) - 2F_1^{(N)}(s)\right) \\ &+ \frac{\eta^g\lambda^g}{2}\Bigg(f_1^{(N-1)}(s)F_1^{(2N)}(s) + f_1^{(N)}(s)F_1^{(2N)}(s) \\ &\quad - \left(f_1^{(N-1)}(s) - f_1^{(2N)}(s)\right)F_1^{(N+1)}(s) \\ &\quad - \left(f_1^{(N)}(s) - f_1^{(2N)}(s)\right)F_1^{(N)}(s)\Bigg)\Bigg] ds. \end{aligned} \tag{93}$$

Again, as a sufficient condition we want to show that we can rank the arguments of the integrals. The equilibrium bidding function of the first-price auction is now given by

$$\beta^{FP}(v_i) = \frac{1}{F_1^{(2N)}(v_i)} \int_0^{v_i} s f_1^{(2N)}(s)\Big(1 + \eta^g\Big(1 - F_1^{(2N)}(s)\Big) + \eta^g\lambda^g F_1^{(2N)}(s)\Big) ds. \tag{94}$$

As before, the first term stemming from the standard preferences equilibrium bidding function is identical in both bidding functions. What is left are the gain-loss utility terms. This means it has to hold that,

$$\begin{aligned}\frac{\eta^g}{2} f_1^{(2N)}(s)\left(4 - F_1^{(N-1)}(s) - F_1^{(N+1)}(s) - 2F_1^{(N)}(s)\right)&\\ -\eta^g f_1^{(2N)}(s)\Big(1 - F_1^{(2N)}(s)\Big)&\\ +\frac{\eta^g\lambda^g}{2}\Bigg[f_1^{(N-1)}(s)F_1^{(2N)}(s) + f_1^{(N)}(s)F_1^{(2N)}(s)&\\ - \left(f_1^{(N-1)}(s) - f_1^{(2N)}(s)\right) F_1^{(N+1)}(s)&\\ - \left(f_1^{(N)}(s) - f_1^{(2N)}(s)\right) F_1^{(N)}(s)\Bigg]&\\ -\eta^g\lambda^g\eta^g f_1^{(2N)}(s)F_1^{(2N)}(s) &\overset{!}{\geq} 0.\end{aligned} \tag{95}$$

Note that all terms, except the third one, include $f_1^{(2N)}(s) = 2NF^{2N}(s)f(s)$. Using the definition of the first-order statistic density for distribution functions and leaving the arguments of the functions out, we can write the third term as

$$\begin{aligned}&f_1^{(N-1)}F_1^{(2N)} + f_1^{(N)}F_1^{(2N)} - \left(f_1^{(N-1)} - f_1^{(2N)}\right) F_1^{(N+1)}\\ &\qquad - \left(f_1^{(N)} - f_1^{(2N)}\right) F_1^{(N)}\\ = \;& f_1^{(2N)}\left(\frac{N-1}{2N}F_1^{(N-1)} + \frac{1}{2}F_1^{(N)} - \frac{N-1}{2N} + F_1^{(N+1)} - \frac{1}{2} + F_1^{(N)}\right).\end{aligned} \tag{96}$$

This inequality can be solved analytically for three bidders and has to be solved numerically for more than three bidders. For three bidders the inequality simplifies to

$$\frac{1}{2} - F + \frac{1}{2}F^2 + \lambda^g\Big(-\frac{1}{4} + \frac{3}{4}F - \frac{1}{2}F^2\Big) \overset{!}{>} 0. \tag{97}$$

Since only F appears, but not its argument, we can solve the inequality without inverting F. The extremum of the left-hand side is attained at $F = \frac{3\lambda-4}{4(\lambda-1)}$, but since the coefficient of the F^2-terms is given by $\frac{1}{2}(1-\lambda)$, this is a maximum. This means that the minimum for valid valued of F is at $F = 0$ or $F = 1$. For $F = 1$, the left-hand side is always equal to zero. For $F = 0$, we have

$$\frac{2-\lambda}{4} \overset{!}{>} 0. \tag{98}$$

N	1	2	3	4	5	6
Total bidders	3	5	7	9	11	13
λ^{crit}	2.0	2.6484	2.3995	2.2856	2.2222	2.1818
$\frac{4N}{2N-1}$	4.0	2.6667	2.4000	2.2857	2.2222	2.1818

Table 2.1: Critical λ^g values for different number of bidders.

This is fulfilled for all $\lambda \leq 2$. For $N > 1$, meaning $5, 7, 9\ldots$ bidders, an analytic solution is not tractable. The inequality can however be solved numerically. The results can be found in Table 2.1, the code to compute the critical lambdas can be found in Appendix 2.6.

From the proof of the case with an even number of bidders, one might expect that the critical λ^g-values are given by the expression for an even number of bidders plus half a bidder per group in expectation,

$$\frac{2(N+\frac{1}{2})-1}{N+\frac{1}{2}-1} = \frac{4N}{2N-1}. \tag{99}$$

While this expression closely approximates the critical λ^gs for more than four bidders, the actual λ^g-values are somewhat smaller than this, as can be seen in Table 2.1. This is due to the fact that the order statistics for the $N+1$- and N-bidder groups depend non-linearly on the number of bidders.

The corollary is proven the exact same way as in the case for an even number of bidders. □

Proof of Proposition 6. The minimal critical λ^g is given by $\lambda^g = 2$. Together with Proposition 4 and Proposition 5, this means that for $N \geq 3$ bidders, an auctioneer is always better off if she conducts a tournament instead of the corresponding one-stage mechanism. □

Note that we derived the critical λ^g-values such that every type bids higher in the tournament than in the corresponding one-stage mechanism. If the auctioneer is solely interested in expected revenue, then the critical

λ^g-values are significantly higher but depend on the distribution function and generally need to be determined numerically.

An exception is the case for $N = 4$ bidders and the uniform distribution. Here, the difference between the expected payment in tournament vs the corresponding one-stage mechanism is given by

$$\mathbb{E}\Big[m^T - m^{FP}\Big] = \frac{1}{840}\eta^g(\lambda^g + 27). \tag{100}$$

This expression is strictly positive for all admissible λ^g and η^g, meaning that the tournament always yields higher revenues than the corresponding one-stage mechanism in this setting. The same result can be derived for a total of three bidders in the case of uniformly distributed values. For $N > 4$, the critical λ^g-values have to be determined numerically, even for the uniform distribution.

2.4.4 OPTIMAL EFFICIENT TWO-STAGE MECHANISM

We have already shown that the tournament poses a strict improvement over one-stage mechanisms if the auctioneer is facing loss averse bidders. Restricting ourselves to two stages, one might ask what the optimal efficient mechanism looks like. In this section we derive and discuss the optimal efficient two-stage mechanism.

Proposition 7. *Assume bidders are loss averse in the good domain and assume a general two-stage mechanism* $\Big(\boldsymbol{\mu}, \mathfrak{M}\Big)$ *that induces* $\varphi_1(s) = \phi_1 \circ F(s)$ *and* $\varphi_2(s) = \phi_2 \circ F(s)$*. Then the expected payment of a bidder with value* v *is given by*

$$\begin{aligned} m^{TS}(v) = \int_0^v s\Big(& f_1(s) + \eta^g f_1(s)\big[2 - \varphi_1(s) - \varphi_2(s)\big] \\ & + \eta^g\lambda^g\big[F_1(s)\varphi_1'(s) + f_1(s)\varphi_2(s) - \varphi_1'(s)\varphi_2(s)\big]\Big)ds. \end{aligned} \tag{101}$$

Proof. We start the proof by choosing the first-price payment rule. We will then use Proposition 2 to show that we can choose any standard payment rule after we have derived the two-stage structure $\mathfrak{M}$. With proposition 1, we can assume bidders bid the same in every stage. Assume the other bidders bid according to an increasing, absolutely continuous bidding function β^{TS}. Note that

$$\varphi_1\left(\beta^{TS^{-1}}(b)\right)\varphi_2\left(\beta^{TS^{-1}}(b)\right) = F_1\left(\beta^{TS^{-1}}(b)\right) \tag{102}$$

Then the utility is given by

$$\begin{aligned} u^{TS}(v_i, b|b^*) = \; & F_1\left(\beta^{TS^{-1}}(b)\right)(v_i - b) \\ & + F_1\left(\beta^{TS^{-1}}(b)\right)\left(1 - \varphi_1\left(\beta^{TS^{-1}}(b^*)\right)\right)\eta^g v \\ & + F_1\left(\beta^{TS^{-1}}(b)\right)\left(1 - \varphi_2\left(\beta^{TS^{-1}}(b^*)\right)\right)\eta^g v \\ & + \left(1 - \varphi_1\left(\beta^{TS^{-1}}(b)\right)\right)F_1\left(\beta^{TS^{-1}}(b^*)\right)\eta^g\lambda(-v) \\ & + \left(\varphi_1\left(\beta^{TS^{-1}}(b)\right) - F_1\left(\beta^{TS^{-1}}(b)\right)\right)\varphi_2\left(\beta^{TS^{-1}}(b^*)\right)\eta^g\lambda(-v). \end{aligned} \tag{103}$$

The bidding function β^{TS} constitutes a UPE if and only if the utility function $u_i^{TS}(v_i, b|\beta^{TS}(v_i))$ attains its maximum at $b = \beta^{TS}(v_i)$ for all v_i. Differentiating u^{TS} with respect to b and plugging in the equilibrium condition $b = \beta^{FP}(v_i)$ yields the ODE

$$\begin{aligned} F_1(s)\,\beta^{TS}(s) + f_1(s)\,\beta^{TS'}(s) = s\Bigg(& f_1(s) + \eta^g f_1(s)\Big[2 - \varphi_1(s) - \varphi_2(s)\Big] \\ & + \eta^g\lambda^g\Big[F_1(s)\varphi_1'(s) + f_1(s)\varphi_2(s) - \varphi_1'(s)\varphi_2(s)\Big]\Bigg). \end{aligned} \tag{104}$$

It follows that

$$\begin{aligned} \beta^{TS}(v) = \frac{1}{F_1(v)}\int_0^v s\Bigg(& f_1(s) + \eta^g f_1(s)\Big[2 - \varphi_1(s) - \varphi_2(s)\Big] \\ & + \eta^g\lambda^g\Big[F_1(s)\varphi_1'(s) + f_1(s)\varphi_2(s) - \varphi_1'(s)\varphi_2(s)\Big]\Bigg)ds \end{aligned} \tag{105}$$

and

$$m^{TS}(v) = \int_0^v s\Bigg(f_1(s) + \eta^g f_1(s)\Big[2 - \varphi_1(s) - \varphi_2(s)\Big] + \eta^g\lambda^g\Big[F_1(s)\varphi_1'(s) + f_1(s)\varphi_2(s) - \varphi_1'(s)\varphi_2(s)\Big]\Bigg)ds. \tag{106}$$

□

Proposition 8. *Assume bidders are loss averse in the good domain and assume a general two-stage mechanism that induces $\varphi_1(s) = \phi_1 \circ F(s)$ and $\varphi_2(s) = \phi_2 \circ F(s)$. Then the expected revenue for the auctioneer is given by*

$$\mathbb{E}[R] = N\int_0^1 s(1 - F(s))\Big(f_1(s) + \eta^g f_1(s)\big[2 - \varphi_1(s) - \varphi_2(s)\big] + \eta^g\lambda^g\big[F_1(s)\varphi_1'(s) + f_1(s)\varphi_2(s) - \varphi_1'(s)\varphi_2(s)\big]\Big)ds. \tag{107}$$

Proof. Again, assume the other bidders bid according to an increasing, absolutely continuous bidding function β^{TS} and use the interim results of Proposition 7. Define

$$\Gamma(s) = s\Big(f_1(s) + \eta^g f_1(s)\big[2 - \varphi_1(s) - \varphi_2(s)\big] + \eta^g\lambda^g\big[F_1(s)\varphi_1'(s) + f_1(s)\varphi_2(s) - \varphi_1'(s)\varphi_2(s)\big]\Big). \tag{108}$$

The expected revenue is given by

$$\mathbb{E}[R] = N\int_0^1\int_0^v \Gamma(s)ds\ f(v)dv. \tag{109}$$

Partial integration yields

$$\int_0^1\int_0^v \Gamma(s)ds\ f(v)dv = \left[\int_0^v \Gamma(s)ds\ F(v)\right]_{v=0}^{v=1} - \int_0^1 \Gamma(s)\ F(s)ds \tag{110}$$

$$= \int_0^1 \Gamma(s)ds - \int_0^1 \Gamma(s)\ F(s)ds \tag{111}$$

$$= \int_0^1 \Big(1 - F(s)\Big)\Gamma(s)ds. \tag{112}$$

□

Proposition 9 (Optimal two-stage structure). *Assume bidders are loss averse in the good domain. Then the optimal two-stage structure is given by*

Stage 1: *With probability $\frac{1}{\lambda}$ bidders get to the second stage with probability 1. With probability $\frac{\lambda-1}{\lambda}$ only the strongest bidder advances to stage 2 and has thus won the auction.*

Stage 2: *If bidders got to stage 2 with probability 1, the strongest bidder wins the auction.*

Bidders are left unaware whether the branch in which everyone advances to the second stage was selected or if the auction took place in the first stage. The only information they receive is whether they have reached stage two or not and after the second stage, whether they have won the auction or not. The interpretation here is that this mechanism induces just the right amount of risk, a bidder in stage 2 does not know whether he beat his opponents already or if the "real" auction is yet to come. This takes care of lower types who do not need to insure themselves against their expectations by bidding even lower, while it encourages strong bidders to bid even higher.

Proof. The proof is structured in two parts. In a first step we optimize the expected revenue functional for general distribution functions and φ_1 and φ_2. In the second step, we show that the optimal φ_i-functions are equivalent to admissible φ_i, meaning that they satisfy the conditions from section 2.3.2. Assume the other bidders bid according to an increasing, absolutely continuous bidding function β^{TS} and use the interim results of Proposition 7.

We have

$$\mathbb{E}[R] = N\int_0^1 \Big(1 - F(s)\Big)\Gamma(s)ds =: N\int_0^1 J(s,\varphi_1,\varphi_1',\varphi_2)ds, \tag{113}$$

with

$$\begin{aligned}\Gamma(s) =& s\Big(f_1(s) + \eta^g f_1(s)\big[2 - \varphi_1(s) - \varphi_2(s)\big] \\ &+ \eta^g\lambda^g\big[F_1(s)\varphi_1'(s) + f_1(s)\varphi_2(s) - \varphi_1'(s)\varphi_2(s)\big]\Big).\end{aligned} \tag{114}$$

We need to find φ_1 and φ_2 that maximize the functional

$$\int_0^1 J(s, \varphi_1, \varphi_1', \varphi_2) ds. \tag{115}$$

A candidate for the optimal φ_i is given by solving the constrained Euler-Lagrange equations for our functional. We will nonetheless begin with the unconstrained Euler-Lagrange equations,

$$\begin{cases} \dfrac{\partial}{\partial \varphi_1} J(s, \varphi_1, \varphi_1', \varphi_2) - \dfrac{d}{ds}\left(\dfrac{\partial}{\partial \varphi_1'} J(s, \varphi_1, \varphi_1', \varphi_2)\right) = 0 \\ \dfrac{\partial}{\partial \varphi_2} J(s, \varphi_1, \varphi_1', \varphi_2) - \dfrac{d}{ds}\left(\dfrac{\partial}{\partial \varphi_2'} J(s, \varphi_1, \varphi_1', \varphi_2)\right) = 0 \\ \varphi_1(1) = 1 \\ \varphi_2(1) = 1. \end{cases} \tag{116}$$

The initial values of the φ_i are the only natural choice: For reasons of efficiency, the highest possible type should always advance with certainty. The probability that two bidders are of the highest possible type is zero. Prescribing values for $\varphi_i(0)$ could lead to distortions since it might be optimal to have an atom on 0. Note that J does not depend on φ_2', so the Euler-Lagrange equations simplify to

$$\begin{cases} \dfrac{\partial}{\partial \varphi_1} J(s, \varphi_1, \varphi_1', \varphi_2) - \dfrac{d}{ds}\left(\dfrac{\partial}{\partial \varphi_1'} J(s, \varphi_1, \varphi_1', \varphi_2)\right) = 0 & (a_1) \\ \dfrac{\partial}{\partial \varphi_2} J(s, \varphi_1, \varphi_1', \varphi_2) = 0 & (b_1) \\ \varphi_1(1) = 1 & (a_2) \\ \varphi_2(1) = 1. & (b_2) \end{cases} \tag{117}$$

This system of ordinary differential equations is closed-form solvable for general distribution functions. We begin with the initial value problem

$(b_1), (a_2), (b_2)$.

$$\begin{cases} s(1-F(s))\Big[-\eta^g f_1(s) + \eta^g\lambda^g\Big(f_1(s)-\varphi_1'(s)\Big)\Big] = 0 & (b_1) \\ \varphi_1(1) = 1 & (b_2) \\ \varphi_2(1) = 1 & (b_2) \end{cases} \tag{118}$$

$$\Leftrightarrow \quad \begin{cases} \varphi_1'(s) = \dfrac{f_1(s)(\lambda^g - 1)}{\lambda} & (b_1) \\ \varphi_1(1) = 1 & (b_2) \end{cases} \tag{119}$$

$$\Rightarrow \quad \varphi_1(s) = \frac{1 + F_1(s)(\lambda^g - 1)}{\lambda^g}. \tag{120}$$

For the second initial value problem $(a_1), (a_2), (b_2)$, we have

$$\begin{cases} s(1-F(s))\Big[-\eta^g f_1(s) - \eta^g\lambda^g\Big(f_1(s)-\varphi_2'(s)\Big)\Big] \\ \quad -\eta^g\lambda^g\Big(1-F(s)-sf(s)\Big)\Big(F_1(s)-\varphi_2(s)\Big) = 0 & (b_1) \\ \varphi_1(1) = 1. & (b_2) \\ \varphi_2(1) = 1. & (b_2) \end{cases} \tag{121}$$

Note that the ODE only depends on φ_2, as was the case with $(b_1), (a_2), (b_2)$ and φ_1. After rearranging and applying the product rule, we arrive at

$$\varphi_2(s) = F_1(s) - \frac{1}{s(1-F(s))}\int_s^1 \frac{y(1-F(y))f_1(y)}{\lambda^g}dy. \tag{122}$$

This means that for the unconstrained optimization problem, the solution is given by

$$\begin{cases} \varphi_1(s) = \dfrac{1 + F_1(s)(\lambda^g - 1)}{\lambda^g} \\ \varphi_2(s) = F_1(s) - \dfrac{1}{s(1-F(s))}\displaystyle\int_s^1 \frac{y(1-F(y))f_1(y)}{\lambda^g}dy. \end{cases} \tag{123}$$

Note that $\varphi_1(s)\varphi_s(s) \neq F_1(s)$, meaning that these do not satisfy the conditions from section 2.3.2. We now show that choosing $\varphi_1(s)$ and $\varphi_2(s)$ according to the solutions of the unconstrained Euler-Lagrange equations is equivalent to choosing $\varphi_2(s) = \frac{F_1(s)}{\varphi_1(s)}$.

Choosing $\varphi_1(s)$ according to (123), the expressions of $\int_0^1 J(s,\varphi_1,\varphi_1',\varphi_2)ds$ that involve $\varphi_2(s)$ are given by

$$\int_0^1 s(1-F(s))\eta^g\varphi_2(s)\Big[-f_1(s)+\lambda^g f_1(s)-\lambda^g\varphi_1'(s)\Big]ds \tag{124}$$
$$= \int_0^1 s(1-F(s))\eta^g\varphi_2(s)\Big[f_1(s)(\lambda^g-1)-\cancel{\lambda^g}\frac{f_1(s)(\lambda^g-1)}{\cancel{\lambda^g}}\Big]ds \tag{125}$$
$$= 0. \tag{126}$$

This implies that once we have chosen $\varphi_1(s)$ as the solution of the unconstrained optimization problem and therefore independent of $\varphi_2(s)$, it does not matter which $\varphi_2(s)$ we choose, as long as it remains measurable. Therefore our final φ_i are given by

$$\begin{cases} \varphi_1(s) = \dfrac{1+F_1(s)(\lambda^g-1)}{\lambda^g} \\ \varphi_2(s) = \dfrac{\lambda^g F_1(s)}{1+F_1(s)(\lambda^g-1)}. \end{cases} \tag{127}$$

□

This two-stage structure optimizes the revenue for the seller. We can even show that bids of *all* types are higher than in the one-stage variants of the mechanism and not just overall revenue.

Proposition 10. *Assume bidders are loss averse in the good domain and consider either the first-price auction, the second-price auction or the all-pay auction. Equilibrium bids in the optimal two-stage structure are higher than in the corresponding one-stage mechanism.*

Proof. First note that replacing the φ_i in Γ by (127) yields

$$\begin{aligned}\Gamma(s) =& s\Big(f_1(s)+\eta^g f_1(s)\big[2-\varphi_1(s)-\varphi_2(s)\big] \\ &+\eta^g\lambda^g\big[F_1(s)\varphi_1'(s)+f_1(s)\varphi_2(s)-\varphi_1'(s)\varphi_2(s)\big]\Big) \\ &= sf_1(s)\left(1+\eta^g\left(2-\frac{1}{\lambda^g}\right)+\eta^g\frac{(\lambda^g-1)^2}{\lambda}F_1(s)\right).\end{aligned} \tag{128}$$

Define

$$\gamma^{OS}(s) = s\Big(1 + \eta^g\big(1 - F_1(s)\big) + \eta^g\lambda^g F_1(s)\Big) \tag{129}$$

$$\gamma^{Opt}(s) = s\left(1 + \eta^g\left(2 - \frac{1}{\lambda^g}\right) + \eta^g \frac{(\lambda^g - 1)^2}{\lambda} F_1(s)\right). \tag{130}$$

We have

$$\gamma^{Opt}(s) \overset{!}{\geq} \gamma^{OS}(s) \tag{131}$$

$$\Leftrightarrow \quad 2 - \frac{1}{\lambda^g} + \frac{(\lambda^g - 1)^2}{\lambda} F_1(s) \overset{!}{\geq} 1 - F_1(s) + \lambda^g F_1(s) \tag{132}$$

$$\Leftrightarrow \quad F_1(s) - 1 \overset{!}{\leq} 0, \tag{133}$$

which is always true. This means that the ranking holds for the first-price auction. One can easily compute that in the case of the second-price auction as underlying mechanism, bidding functions are given by

$$\beta^{SP}(v) = \gamma^{OS}(v) \tag{134}$$

$$\beta^{Opt}(v) = \gamma^{Opt}(v). \tag{135}$$

In the case of the all-pay auction, the bidding functions are given by

$$\beta^{SP}(v) = \int_0^v \gamma^{OS}(s) f_1(s) ds \tag{136}$$

$$\beta^{Opt}(v) = \int_0^v \gamma^{Opt}(s) f_1(s) ds. \tag{137}$$

This concludes the proof. □

2.5 CONCLUSION

In this paper we investigate how a buyer should design her procurement mechanism when bidders are loss averse. Loss aversion implies that the willingness to pay of a bidder depends on the probability he assigns to winning the auction. We show that a simple two-stage mechanism, the tournament, outperforms any one-stage mechanism revenue-wise if bidders are not too loss averse. As a robustness-check, we show that the buyer's revenue is not dependent on the payment rule she implements. Once the structure of the multi-stage mechanism is fixed, a revenue equivalence principle holds. Finally, we derive the optimal, efficient two-stage mechanism. This mechanism is, in contrast to the tournament, dependent on the degree of loss aversion of the bidders and therefore difficult to implement in real-life procurement.

Our analysis opens the door to further research. On the one hand, it might be interesting to investigate whether a buyer could further improve her revenue if she were to implement a three-stage (or even more stages) mechanism. Numerical simulations suggest that the answer is no, but the problem quickly becomes untractable even for a fixed cost distribution like the uniform distribution. On the other hand, one could expand the model to include bidders that are loss averse in the money domain, too. The revenue equivalence principle that we derived fails in that case, as shown by Eisenhuth and Ewers (2012). In their paper, they show that the all-pay auction yields higher revenues than the first-price auction in a setting similar to ours. This implies that the optimal mechanism for two or more stages will depend on the payment rule the buyer implements, making the optimization problem a lot harder.

2.6 APPENDIX

MATHEMATICA CODE

The code takes a starting value an then shoots λ-values until the minimum of the function *Func* crosses 0.

```
Func[N_, l_] :=
 (1 - x^(N - 1) / 2 - x^(N + 1) / 2 - x^(N) + x^(2 N)) + l * (((N - 1) / (4 N)) * x^(N - 1) +
     x^N / 4 - ((N - 1) / (4 N)) - 1 / 4 + x^(N + 1) / 2 + (x^N) / 2 - x^(2 N))

a = 1;
step = 0.0001;
temp = 0;
startvalue = 2.153;
While[a > 0,
 sumsteps = temp;
 a = FindMinimum[{Func[7, startvalue + sumsteps], 0 < x < .3}, x][[1]];
 temp = sumsteps + step;
 If[a < 0,
  Print["lambda=" <> ToString[NumberForm[startvalue + sumsteps - step, 10]]]]
]
```

CHAPTER 3

AUCTION EXPERIMENTS WITH A REAL-EFFORT TASK

Abstract

We propose a novel design for auction experiments based on effort and money. Participants bid a number of sliders in order to win a monetary prize. If successful, a participant has to solve a real-effort task, namely the slider task. The design allows us two capture two-dimensional prospect theory and common value effects. In a second step, we test our design in the laboratory. We find evidence for both loss-aversion and common values.

3.1 MOTIVATION

When investigating auctions in the laboratory, economic researchers usually rely on induced values experiments. This means that each participant is assigned a value v for a (hypothetical) good. A participant's payoff associated with getting the good is given by the difference between his induced value v and the price p he has to pay for the good. Induced values experiments grant the researcher a lot of control, which is an advantage when for example hypotheses about a specific bidding function are tested in the laboratory. However, compared to real world auctions, this design choice abstracts from two well-known phenomena that both can potentially limit the external validity of results from the lab: Two-dimensional outcome evaluation and common values. We propose and test a simple experimental design based on money and effort that can account for both these phenomena.

In the vast majority of economic research, agents are assumed to evaluate their outcomes in one dimension. Indeed, assuming a one-dimensional outcome evaluation is without loss of generality if agents have standard preferences, in the sense that '[they] maximize a global utility function over lifetime consumption $U(x|s)$' (DellaVigna, 2009). However, Lange and Ratan (2010) show that theoretical predictions differ between one-dimensional (induced values auctions) and two-dimensional settings (real good auctions), e.g., if agents are loss averse. Furthermore, Abeler et al. (2011) provide experimental evidence for a multidimensional evaluation in a setting in which participants perform a real effort task and earn money.

Consider the following situation: You discover that a certain good you always wanted to own is offered in an eBay auction. A day before the auction ends you have determined your willingness to pay and bid exactly that amount. If you have standard preferences and a private valuation for the good, your bid should be equal your willingness to pay. After submitting your bid, you learn that you are currently the highest bidder, which stays the case until one minute before the auction ends. Then you learn that another person outbid you.

If agents have standard preferences, nothing else would happen. Bidding above your predefined private valuation cannot be rationalized by any one-dimensional, standard-preferences model, in which your payoff is simply $v - p$. The same applies to induced values experiments, where paying more than the induced valuation would lead to negative payoffs.

However, if you compare outcomes to expected outcomes in multiple dimensions, you might increase your bid. One minute prior to the end of the auction, your expected outcome is "I will receive the good" in the good dimension, and "I will spend the second-highest bid" in the money dimension. Losing the auction in the last second would imply a large deviation in the good dimension. As a result, you'd rather deviate a little in the money dimension and bid above your valuation.

Kahneman et al. (1990) showed experimentally that the valuations for goods are not exogenous. In line with them, we argue that the willingness to pay for a good depends on the selling mechanism. When you believe the probability of winning a good is high, you become more attached to it, which in turn leads to a higher willingness to pay.

In addition to two-dimensional outcome evaluation, in most real-world auctions bidders are confronted with some common value component in the auctioned good, meaning that there is information on my own valuation in the bids of my competitors. Take again procurement as an example: Suppliers usually have some uncertainty about their actual cost. This might stem from future commodity prices, wages, or changes in the specification after the sourcing process. Hence a very low bid of competitors might mean that I overestimated these future costs. Even when consumption goods are auctioned off, some common value component might be present. Other bidders might e.g. have better (or different) information on the availability and prices of the good in other outlets. In addition, there is a large strand of literature showing that common value auctions lead to different predictions than pure private value auctions (for an extensive review, see Kagel and Levin (2002)).

To summarize, induced values experiments do not account for two-dimensional outcome evaluation and common value components. Since both these phenomena are present in most real world auctions, and both are important drivers of bidding behaviour, one has to be very cautious when giving practitioners advice based on induced values experiments.

We propose a novel design to increase external validity of auction experiments, based on effort and money. In a first step, bidders can familiarize themselves with the real-effort task, the slider task, in an incentivized test round lasting four minutes. In that time, bidders solve as many slider tasks as possible and are remunerated per solved unit. We then let subjects bid on a prize of 10 Euros. We asked participants to submit bids that express the number of slider tasks they would maximally solve in case of winning the auction, i.e. how much effort they are willing to spend in order to receive 10 Euros. We implemented a between-subjects design with a varying number of bidders between treatments ($N = 2$ vs. $N = 8$). Moreover, we chose the second-price auction. It has the desirable property that with standard preferences, the dominant strategy is independent of beliefs about the number of bidders, their valuations, or their strategies. In our design, if subjects were one-dimensional utility maximizers with a purely private valuation, they would determine the level of effort they are maximally willing to spend for 10 Euros, and bid exactly that amount. Based on standard theory and experiments with induced values, we would thus expect no difference in behavior between the treatments. However, as we show in Chapter 2, theoretical predictions differ when agents act according to two-dimensional prospect theory. Bids are higher when the number of bidders is low, as a high winning probability leads to an increased attachment to the prize of 10 Euros. The same applies if there is a common value component in conducting the slider task. When bidding against seven bidders, winning is 'bad news' with a higher probability since seven other bidders estimated a lower common value component.

In line with the reference dependent two-dimensional and common value predictions, we observed significantly higher bids for $N = 2$. On average, subjects were willing to solve roughly 30% more slider tasks when they had one instead of seven opponents. This result is robust to regressions where we control for demographic characteristics as well as participants' test scores, i.e. the amount of sliders they were able to solve in an incentivized test round. We hence argue that (in contrast to induced value experiments) our design enables researchers to increase external validity of auction experiments. Moreover, as pointed out by Gill and Prowse (2019), the slider task allows experimenters to control for participants' abilities, while at the same time having the advantages of real effort tasks.

In addition, we conducted three treatments to investigate the main driver behind our results: To isolate two-dimensional loss aversion from common values, we let bidders bid against computerized competitors. Evidence from these treatments is mixed: On the one hand, we did not observe a significant difference in bids depending on the ex-ante winning probability of bidders bidding against computers. On the other hand, we didn't observe a significant difference between bids against computerized and human competitors, either. While the former result is in favor of common values as main driver, the latter opposes this hypothesis.

Our results are in line with Rosato and Tymula (2019), who investigate bidding behavior in second-price auctions. They auction off several real goods sequentially and find that subjects bid more if they face less competition and hence have a higher probability of getting the good. Banerji and Gupta (2014), find that participants bid less aggressively when they faced stronger computerized competitors. They employ a BDM mechanism in which participants bid against a computerized opponent in a second-price auction.

Notably, compared to real good auctions of Rosato and Tymula (2019), our design has three important advantages: Firstly, we do not observe a concentration of bids at very low values, which is often the case when real goods are sold. Students might have the expectation to leave a laboratory experiment with a certain amount of money, not with a good. Secondly, due to the incentivized test round our design enables researchers to control for valuations of participants, i.e. a participant's pace in solving slider tasks. Thirdly, we argue that experiments with the proposed design are less expensive than real good experiments. In real good experiments all participants are usually endowed with a certain amount of money which at the same time serves as upper bound for bids. In order to allow all participants to express their true willingness to pay for certain goods, these upper bounds need to be quite high. Alternatively, experimenters have to use goods with low values, which in turn aggravates the problem of bid concentration around zero. Using our design, one does not have to define and endow all participants with that artificial upper bound.

Finally, due to remarkable analogies to practices in industry, especially in procurement, our design adds additional realism to the existing experimental literature. When bidding on a procurement contract, suppliers usually have a good idea of their true costs, based on internal calculations and estimations on future developments, e.g. in commodity markets. Furthermore, they tend to have some beliefs about their competitors, i.e. a supplier might know whether they are a high- or a low-cost supplier. Yet they do not know their exact costs, as well as the exact distribution that their competitors draw their costs from. The same holds true for participants in our experiment. They know how the task works and how long it would take them to fulfill a certain amount of tasks given that they keep their initial pace. They also might have an idea on how well they perform, or how high their opportunity costs of staying in the lab are compared to other participants. Yet they are faced with similar uncertainties as suppliers in procurement: On the one hand, there are uncertainties with regards to the actual costs of effort (resulting from changes in pace or an unexpected evolvement of marginal pain in each slider task) and on the other hand, there is no common distribution function that all bidders draw their valuation from.

3.2 THEORY

3.2.1 MODEL

In this section, we introduce the formal model. We consider $n \geq 2$ bidders competing for one indivisible good in a second-price auction. The value v_i of bidder $i \in \{1, \ldots, n\}$ for the good is privately drawn from a distribution F, $v_i \overset{\text{iid.}}{\sim} F[0,1]$. F is assumed to have a differentiable density f which is strictly positive on its support $[0,1]$. Moreover, f is common knowledge. Analogous to the standard setting, where the value of the good is measured in monetary units, i.e. in the dimension bidders submit their bids, we assume that the value is measured in slider tasks. Hence bidders draw the amount of slider tasks they are willing to solve in order to receive 10 Euros.

Bids are placed after learning the value for the good. The bidder submitting the highest bid is awarded the good and has to pay the second highest bid.

Bidders are assumed to be loss-averse following Kőszegi and Rabin (2006). We assume two distinct dimensions of loss aversion, a currency domain c in which bidders submit their bids, and a prize domain g representing the item the winner of the auction receives. Furthermore, we assume bidders to be narrow-bracketers, following the definition of von Wangenheim (2019). This means that the bidders' gain-loss utility is evaluated separately for each dimension. Summarizing, for outcome $x = (x^c, x^g)$, valuation v for the good, and reference consumptions r^c and r^m, agent's utility is given by

$$u(x\big|r^g, r^c) = x^c + vx^g + \mu^g(vx^g - vr^g) + \mu^c(x^c - r^c).$$

Following Kőszegi and Rabin (2006), we assume μ^i to be a piecewise linear function with a kink at zero,

$$\mu^g(y) = \begin{cases} \eta^g y & \text{if } y \geq 0 \\ \lambda^g \eta^g y & \text{if } y < 0, \end{cases} \qquad \mu^c(y) = \begin{cases} \eta^c y & \text{if } y \geq 0 \\ \lambda^c \eta^c y & \text{if } y < 0. \end{cases}$$

The μ^i denote the gain-loss utilities in the respective dimension, with $\eta^i > 0$, $\lambda^i > 1$ and $\eta^i(\lambda^i - 1) \leq 1$ for $i \in \{g, c\}$. The interpretation is that bidders perceive, in addition to their classical utility, a feeling of gain or loss, depending on the deviation from their reference consumption.

The reference point in our paper is assumed to be determined by rational expectations following Kőszegi and Rabin (2006).

3.2.2 EQUILIBRIUM CONCEPT

We follow Kőszegi and Rabin (2006)'s and von Wangenheim (2019)'s equilibrium concept under uncertainty, according to which bidders form their strategy after learning their valuation. We apply the concept of unacclimated personal equilibria, which is, as argued by Kőszegi and Rabin (2006), the appropriate concept in auction settings. Fixing the opponents' strategies, let $H(b, v_i)$ denote i's payoff distribution given his draw v_i from a continuous distribution $F(v)$ and his bid b. A bid b^* constitutes an unacclimated personal equilibrium (UPE), if for all b

$$U[H(b^*, v_i)\big|H(b^*, v_i)] \geq U[H(b, v_i)\big|H(b^*, v_i)].$$

That means, given your reference point is determined by the payoff distribution resulting from an (exogenous) bid b^*, it is a best response to bid b^*.

3.2.3 ANALYSIS

It is well-known that if agents are loss averse only in the prize domain, bidders in second-price auctions bid more aggressively when the number of bidders is low. Yet, in our setting it is arguable that loss aversion in the currency domain, i.e. the amount of slider tasks participants have to solve, also plays a role. Still it seems very plausible that students in the lab are more concerned about receiving money than solving slider tasks, and hence face a higher degree of loss aversion in the prize domain. In this section, we hence derive and analyze bidding behavior in second-price auctions, showing that when agents are more loss averse in the price domain the result above still holds true.

Assume all bidders except bidder i bid according to a strictly increasing bidding function β. Let $G(x) := F^{n-1}(x)$. The utility of bidder with value v, who is loss averse in both the good and the currency domain, bids b and has a reference point of b^*, is given by

$$\begin{aligned} u_i(v_i, b_i|b^*) = {} & G\Big(\beta^{-1}(b_i)\Big)v - \int_0^{b_i} s\beta(s)dG\big(\beta^{-1}(s)\big) \\ & + G\Big(\beta^{-1}(b_i)\Big)\Big(1 - G\Big(\beta^{-1}(b^*)\Big)\Big)\mu^g(v-0) \\ & + \Big(1 - G\Big(\beta^{-1}(b_i)\Big)\Big)G\Big(\beta^{-1}(b^*)\Big)\mu^g(0-v) \\ & + \int_0^{b} \Bigg(\int_0^{b^*} \mu^c(t-s)dG\big(\beta^{-1}(t)\big) \\ & \qquad + \int_{b^*}^{\infty} \mu^c(0-s)dG\big(\beta^{-1}(t)\big)\Bigg) dG\big(\beta^{-1}(s)\big) \\ & + \int_b^{\infty} \Bigg(\int_0^{b^*} \mu^c(t-0)dG\big(\beta^{-1}(t)\big) \\ & \qquad + \int_{b^*}^{\infty} \mu^c(0-0)dG\big(\beta^{-1}(t)\big)\Bigg) dG\big(\beta^{-1}(s)\big) \end{aligned} \tag{138}$$

As shown by von Wangenheim (2019), the equilibrium bidding function for n bidders is given by

$$\begin{aligned}\beta_n^{II}(v) =& \frac{1+\eta_g+\eta_g\left(\lambda_g-1\right)F^{n-1}(v)}{1+\eta_c\lambda_c}v \\ &+\int_0^v \left[\frac{\eta_c\left(\lambda_c-1\right)\left(1+\eta_g+\eta_g\left(\lambda_g-1\right)F^{n-1}(s)\right)}{(1+\eta_c\lambda_c)^2}s\right. \\ &\left.\exp\left(\frac{\eta_c\left(\lambda_c-1\right)}{1+\eta_c\lambda_c}\Big(F^{n-1}(v)-F^{n-1}(s)\Big)\right)\right]dF(s).\end{aligned} \tag{139}$$

Theorem 1. *If bidders are loss averse in both the currency (subscript c) and the prize domain (subscript g), and it holds that bidders are more loss averse in the prize domain in the sense that*

$$\lambda_g \geq \lambda_c\frac{\eta_c(1+\eta_g)}{\eta_g(1+\eta_c)}+\frac{\eta_g-\eta_c}{\eta_g(1+\eta_c)}, \tag{140}$$

then $\beta_n^{II}(v) > \beta_m^{II}(v)$ for all v and $n < m$. Sufficient conditions are given by

$$\begin{cases}\Lambda_g \geq \Lambda_c & \text{if } \eta_g \leq \eta_c \\ \lambda_g \geq \lambda_c & \text{if } \eta_g > \eta_c.\end{cases} \tag{141}$$

Proof. We need to show that

$$\Delta(v;n,m) := \beta_n^{II}(v)-\beta_m^{II}(v) > 0 \tag{142}$$

if $m > n$. Define

$$a(x,y;n,m) := \exp\left(\tilde{c}\left(x^n-y^n\right)\right)-\exp\left(\tilde{c}\left(x^m-y^m\right)\right) \tag{143}$$

$$b(x,y;n,m) := y^n\exp\left(\tilde{c}\left(x^n-y^n\right)\right)-y^m\exp\left(\tilde{c}\left(x^m-y^m\right)\right), \tag{144}$$

where

$$\tilde{c} := \frac{\eta_c\left(\lambda_c-1\right)}{1+\eta_c\lambda_c}. \tag{145}$$

With this, we have

$$\begin{aligned}\Delta(v;n,m) =& \frac{\eta_g\left(\lambda_g-1\right)v}{1+\eta_c\lambda_c}\Big(F^{n-1}(v)-F^{m-1}(v)\Big)\\&+\int_0^v \frac{\eta_c\left(\lambda_c-1\right)}{(1+\eta_c\lambda_c)^2}\Bigg[(1+\eta_g)a\left(F(v),F(s);n,m\right)\\&\qquad+\Lambda_g b\left(F(v),F(s);n,m\right)\Bigg]sdF(s)\\ >& \int_0^v \frac{\eta_c\left(\lambda_c-1\right)}{(1+\eta_c\lambda_c)^2}\Bigg[(1+\eta_g)a\left(F(v),F(s);n,m\right)\\&\qquad+\Lambda_g b\left(F(v),F(s);n,m\right)\Bigg]sdF(s).\end{aligned} \tag{146}$$

A sufficient condition for $\Delta(v;n,m)>0$ to hold is that

$$(1+\eta_g)a\left(F(v),F(s);n,m\right)+\Lambda_g b\left(F(v),F(s);n,m\right)>0. \tag{147}$$

Following the definitions of a and b and because F and exp are strictly increasing, we have, for $s\leq v$

$$(1+\eta_g)a\left(F(v),F(s);n,m\right)+\Lambda_g b\left(F(v),F(s);n,m\right)\overset{!}{>}0 \tag{148}$$

$$\Leftrightarrow\ (1+\eta_g)a\left(v,s;n,m\right)+\Lambda_g b\left(v,s;n,m\right)\overset{!}{>}0 \tag{149}$$

$$\Leftrightarrow\ (1+\eta_g)a\left(1,s;n,m\right)+\Lambda_g b\left(1,s;n,m\right)\overset{!}{>}0. \tag{150}$$

Note that $a\left(1,s;n,m\right)\leq 0$ and $b\left(1,s;n,m\right)\geq 0$ for all s. Also note that $b\left(1,s;n,m\right)>-a\left(1,s;n,m\right)$ for all $s\in(0,1)$. This means there exist $\tilde{q}\in(0,\infty)$ such that $\tilde{q}\,a\left(1,s;n,m\right)+b\left(1,s;n,m\right)=0$ for one or multiple $s\in(0,1)$. Let $q=\min\{\tilde{q}\}$. Then

$$q\,a\left(1,s;n,m\right)+b\left(1,s;n,m\right)\geq 0 \tag{151}$$

for all $s\in[0,1]$. Let $\tilde{s}\in(0,1)$ be such that

$$q\,a\left(1,\tilde{s};n,m\right)+b\left(1,\tilde{s};n,m\right)=0. \tag{152}$$

For inequality (147) to hold, it then needs to hold that

$$\frac{1+\eta_g}{\Lambda_g}=\frac{1+\eta_g}{\eta_g(\lambda_g-1)}\overset{!}{<}q. \tag{153}$$

Rearranging yields

$$\lambda_g \overset{!}{>} \lambda_g^*(\eta_g, q) := \frac{1+\eta_g+\eta_g q}{\eta_g q}. \tag{154}$$

We have that

$$\frac{\partial}{\partial q}\lambda_g^*(\eta_g, q) = -\frac{1+\eta_g}{(\eta_g q)^2} < 0. \tag{155}$$

This means if q increases, the inequality for (154) admits smaller λ_g. The "worst case" to check is therefore the smallest q.

Note that

$$\frac{\partial}{\partial s}q = \frac{\partial}{\partial s}\frac{-b\,(1, s; n, m)}{a\,(1, s; n, m)} < 0. \tag{156}$$

Since $\tilde{s} \in (0, 1)$, and q strictly decreasing in s, we need to check the limit case $s \to 0$,

$$\lim_{s\to 1} q = \frac{1}{\tilde{c}} - 1 = \frac{1+\eta_c}{\eta_c(\lambda_c - 1)} =: q^*. \tag{157}$$

We can now plug this q^* into λ_g^* from (154), yielding

$$\lambda_g^*(\eta_g, q^*) = \lambda_c\frac{\eta_c(1+\eta_g)}{\eta_g(1+\eta_c)} + \frac{\eta_g-\eta_c}{\eta_g(1+\eta_c)}. \tag{158}$$

Concerning the sufficient conditions, let us first consider the case $\eta_g \le \eta_c$. Assume it holds that $\Lambda_c < \Lambda_g$, meaning $\eta_c(\lambda_c - 1) < \eta_g(\lambda_g - 1)$. With $\lambda_i > 1$ and $0 < \eta_i < 1$ for $i \in \{g, m\}$, this is equivalent to

$$\lambda_g > \lambda_c\frac{\eta_c}{\eta_g} + \frac{\eta_g-\eta_c}{\eta_g}. \tag{159}$$

For $\eta_g \le \eta_c$, we have that

$$\begin{aligned} &\lambda_c\frac{\eta_c}{\eta_g} + \frac{\eta_g-\eta_c}{\eta_g} - \lambda_c\frac{\eta_c(1+\eta_g)}{\eta_g(1+\eta_c)} - \frac{\eta_g-\eta_c}{\eta_g(1+\eta_c)} \\ &\qquad = \Big(\lambda_c - 1\Big)\frac{\eta_c(\eta_c-\eta_g)}{\eta_g(1+\eta_c)} \ge 0. \end{aligned} \tag{160}$$

Therefore it follows that if $\eta_g \le \eta_c$ and $\Lambda_g \ge \Lambda_c$, then $\lambda_g \ge \lambda_g^*$. For the second case where $\eta_g \ge \eta_c$, we have that

$$\lambda_c > \lambda_c\frac{\eta_c(1+\eta_g)}{\eta_g(1+\eta_c)} + \frac{\eta_g-\eta_c}{\eta_g(1+\eta_c)}, \tag{161}$$

so $\lambda_g > \lambda_c$ is sufficient in the case $\eta_g \ge \eta_c$. □

3.3 EXPERIMENT

In this section, we introduce our experimental design and state our hypotheses for the experiment.

3.3.1 DESIGN

In each experimental treatment all subjects participated in a second-price sealed-bid auction. In this auction bidders competed for a fixed payment of 10 Euros and bid how many slider tasks they were willing to solve. The bidder who placed the highest offer won. The number of sliders the winner had to solve was equal to the second highest bid.

After the auction took place the winner had a total of 90 minutes to solve the slider task. Only if the winner managed to solve the required number of sliders the winner received 10 Euros, otherwise the winner received no payment.[1] Losing bidders left the laboratory before winners started to solve the slider tasks.

The auction stage was preceded by a first stage in which participants familiarized with the slider task. In this stage participants had 4 minutes to solve slider tasks. For each slider solved they received 4 Cents. At this point in time they did not yet receive the instructions for the auction stage.

We conducted a total of 5 different treatments. We had 2 treatments in which all bidders were human, in one of the treatments we conducted an auction with 2 bidders (H_2) and in the other treatment we conducted an auction with 8 bidders (H_8). In our 3 treatments with computerized competitors we had one treatment with one computerized competitor (C_2^{2000}) and one treatment with 7 computerized competitors (C_8^{2000}). In both treatments the bids of the computerized competitors were uniformly distributed between 0 and 2000. In the remaining treatment (C_2^{4000}) participants bid against a single computerized competitor with bids uniformly distributed between 0 and 4000.

Screenshots of the experiment can be found in the appendix.

[1]All winners managed to solve the required numbers of sliders.

3.3.2 ORGANIZATION

The experiments were conducted in the Cologne Laboratory for Economic Research (CLER) at the University of Cologne, Germany. Using the recruiting system ORSEE (Greiner, 2015), we invited a random sample of the CLER's subject pool via email. Our participants were mostly undergraduate students from the University of Cologne, with different beackground with regards to their major. The whole experiment was computerized using the programming environment oTree (Chen et al., 2016). Upon their arrival at the laboratory, participants were randomly assigned to one of two rooms to either the two-bidder or the eight-bidder treatment. Both treatments were conducted simultaneously and are described in section 4.2. Participants were grouped into cohorts of two and eight respectively. Moreover, participants were seated in visually isolated cubicles and read instructions on their screens (see Appendix 5.5.1) describing the rules of the game.

In total, 112 subjects participated in the experiment, with 48 subjects participating in the two-bidder second-price auctions and 64 subjects participating in the eight-bidder second-price auctions. An overview on participants and their demographics can be found in Table 1.

Payoffs were stated in EUR. Participants were paid out in private after the completion of the experiment. All 112 participants were paid their total net earnings. The average payoff for the entire experiment was 9.59 EUR corresponding to approx. 10.84 USD at the time of the payment.

In order to prevent selection effects as much as possible, we conducted the treatments we primarily want to compare in parallel. Participants were invited to the same experimental session and randomly assigned to one of two treatments that ran simultaneously. Table 3.1 displays which treatments were conducted in parallel.

3.3.3 HYPOTHESES

Standard theory predicts that bidders behave the same in both treatments. That is, agents determine their ”valuation”, i.e. the amount of slider tasks they are maximally willing to solve in order to receive 10 EUR, and then bid

Table 3.1: Experimental sessions

Sessions	Treatment 1	Treatment 2
1	H_2	H_8
2	C_2^{2000}	C_8^{2000}
3	C_2^{2000}	C_2^{4000}

Table 3.2: Descriptive statistics and summary

	H_2	H_8	C_2^{2000}	C_8^{2000}	C_2^{4000}
Age	24.75	26.25	24	23	24.5
Share of females	0.33	0.32	0.51	0.47	0.41
Lab experience	$15-20$	$10-15$	$15-20$	$10-15$	$15-20$
Observations	48	64	84	47	41
Test score	51.9	50.4	55.1	54.9	63.8
Bid	736	551	784	707	934

exactly that amount. Bidding one's true valuation is a dominant strategy in the second-price auction with private values, independent of risk-aversion or beliefs about others, and therefore the bids should not depend on the number of bidders present in the auction. This leads to the following hypothesis:

Hypothesis 1. *We observe no difference in the bids between the treatments.*

When agents are loss-averse, a relatively high ex ante winning probability leads to a relatively strong attachement to the prize of 10 Euros. A strong attachement to the prize increases agents' willingness to work and hence lets them bid more aggressively as compared to a situation where the ex ante winning probability is low. This leads to the following alternative hypthesis:

Hypothesis 2. *We observe higher bids in the "2 bidder" treatments than in the "8 bidder" treatments.*

3.3.4 SUMMARY

A summary of our data can be found in Table 3.2. We denote participants experienced if they have participated in more than 10 laboratory experiments. Test score denotes how many sliders the participant solved during the initial, incentivized four-minute test. Participants do not exhibit a significant difference in this score between the two treatments. Bids are distributed between 10 and 4000.

3.3.5 RESULTS

We start the analysis of our experiment by comparing the bidding behavior in the treatments that were conducted in parallel. This is most similar to the analyses conducted by Banerji and Gupta (2014) and Rosato and Tymula (2019). Afterwards, we will also take into consideration data generated in the first part of the experiment, in which participants got used to the slider-task, and demographic information. Since computerized bidders in treatments C_2^{2000} and C_8^{2000} could not bid above 2000, we censored bids at 2000. Six out of 284 bids were larger than 2000.

Result 1. *When two human bidders competed (H_2) they bid more aggressively than in the case in which eight human bidders (H_8) competed (Mann-Whitney-U test, $p = 0.0329$).*

This result is in line with Rosato and Tymula (2019) who find that increasing the number of bidders decreases average bids. Possible explanations are loss-averse bidders or a common-value effect. While the former explanation predicts a similar effect in treatments with computerized competitors, meaning that lower winning probability implies lower bids, the latter explanation implies that no effect should be observable when comparing treatments in which participants bid against computerized competitors.

Result 2. *Bids do not differ between C_2^{2000} and C_8^{2000} as well as between C_2^{2000} and C_2^{4000} (MW, $p = 0.4597$ and $p = 0.3590$)[2].*

[2]Significance does not change if we consider all C_2^{2000} sessions.

Figure 3.1: Cumulative bid distributions

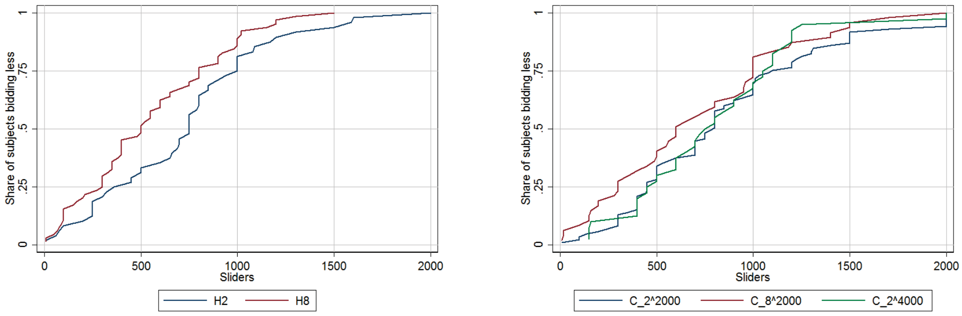

(a) Human competitors (b) Computerized competitors

Looking at the treatments with computerized competitors, we do not find further evidence for loss-aversion. This result is in contrast to Banerji and Gupta (2014) who find that participants bid less aggressively when they faced stronger computerized competitors. Figure 3.1 displays the cumulative bid distributions for the different treatments.

Table 3.3 compares bidding behavior in between treatments with human and computerized competitors. Sessions in which competitors were human serve as a baseline. Computer is a dummy variable that is equal to one if the competitors were computerized and zero otherwise. Similarly, Female is a dummy variable indicating the gender of the participant. Age indicates participants' age and Lab experience how often a subject participated in lab experiments before. The regression shows that the performance in the first part of the experiment is a good predictor of the bid. At the same time we find no evidence that it makes a difference for participants whether they bid against a human or a computerized competitor. In case of a strong common value effect, one would expect a significant difference given that the computer bid is uninformative. Furthermore, demographics have no significant influence on bids.

Table 3.4 compares bidding behavior in treatments with human competitors taking into account the performance in the first part of the experiment and demographics. The treatment $H2$ serves as a baseline and $H8$ is a dummy variable, being equal to one for the $H8$ treatment and zero other-

Table 3.3: Regression comparing bidding against human and computerized competitors

	(I) Bid	(II) Bid
Test score	18.97***	19.28***
	(11.91)	(11.58)
Computer	59.29	44.93
	(1.09)	(0.81)
Female		-9.380
		(-0.17)
Age		6.330
		(1.92)
Lab experience		-7.648
		(-0.89)
Constant	-319.8***	-452.3**
	(-3.50)	(-3.18)
Observations	284	275[3]

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.0001$

wise. The analysis confirms the former result, showing that it is not driven by different abilities or demographic factors.

Table 3.5 compares bidding behavior in treatments with computerized competitors taking into account the performance in the first part of the experiment and demographics. The C_2^{2000} treatment serves as a baseline and C_2^{4000} and C_8^{2000} are dummy variables indicating the treatment. The analysis confirms the former result, showing that the result is not driven by different abilities or demographic factors. However, it suggests that older participants bid more aggressively.

Table 3.4: Regression comparing in treatments with human competitors

	(I) Bid	(II) Bid
Test score	14.77^{***}	14.56^{***}
	(4.05)	(3.49)
H8	-205.1^{*}	-201.9^{*}
	(-2.35)	(-2.19)
Female		-82.95
		(-0.82)
Age		2.941
		(0.47)
Lab experience		-2.218
		(-0.14)
Constant	216.9	218.3
	(0.90)	(0.63)
Observations	112	110[4]

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.0001$

Table 3.5: Regression comparing in treatments with human competitors

	(I) Bid	(II) Bid
Test score	19.44***	19.70***
	(11.15)	(10.98)
C_8^{2000}	80.50	103.0
	(0.99)	(1.25)
C_2^{4000}	-73.27	-50.58
	(-0.96)	(-0.65)
Female		10.23
		(0.15)
Age		8.337*
		(2.15)
Lab experience		-16.55
		(-1.60)
Constant	-286.2**	-461.2**
	(-2.69)	(-2.90)
Observations	172	165[5]

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.0001$

3.4 CONCLUSION

In this paper we propose and test a novel design for auction experiments. In our design, bidders submit bids in terms of slider tasks they are willing to solve in order to receive a certain amount of money. By using different dimensions for bids and good, our design can be exploited to increase external validity of auction experiments. Notably, our design can capture two practically important phenomena that induced values auctions abstract from: Two-dimensional outcome evaluation and common value components. As auction theorists have shown, the existence of either of these two phenomena can lead to qualitatively different predictions as compared to predictions based on induced values experiments (for the former see e.g. Lange and Ratan (2010) and for the latter see e.g. Kagel and Levin (2002)).

Testing our design, we conduct second-price auctions with a varying number of bidders. If agents are either loss-averse and evaluate their outcome in multiple dimensions, or if the auctioned good has a common value component, theory predicts that bids in second-price auctions are decreasing in the number of bidders. This has already been confirmed experimentally Banerji and Gupta (2014) or Rosato and Tymula (2019) in real good experiments. By conducting additional treatments where agents bid against computers, we investigate if our results are mainly driven by two-dimensional loss aversion or common values (which do not play a role when playing against a computer). However, based on these treatments we cannot confirm nor reject common values as a driver behind our results. On the one hand, bids do not differ significantly if the ex-ante probability of winning against a computer is varied, which is in favor of common values as main driver. On the other hand, bids do also not differ significantly between treatments with computerized and human competitors, contradicting the hypothesis of common values as main driver. We however argue that, by manipulating information about the slider task and other bidders, scholars can exploit our design to choose the extent to which common values play a role.

Our contribution is hence that when conducting auction experiments, the design choice should depend on the relevant environment that is investigated. If agents are bidding on objects that have only monetary value to them (e.g. for resale or pure investments), induced values experiments are a natural and appropriate choice. Yet, whenever outcomes are evaluated in multiple dimensions or the auctioned good has a common value component, our design or, if applicable, real good experiments should be preferred.

3.5 APPENDIX

3.5.1 INSTRUCTIONS

H_2 *Treatment*

Übersicht

Herzlichen Dank für Ihre Teilnahme an diesem Experiment. Während des Experiments ist es Ihnen nicht erlaubt, mit anderen Teilnehmern zu kommunizieren, Mobiltelefone zu benutzen, oder andere Programme auf dem Computer zu starten. Sollten Sie gegen diese Regeln verstoßen, müssen wir Sie vom Experiment und all seinen Auszahlungen ausschließen. Für Ihr Erscheinen zu diesem Experiment erhalten Sie 4€.

Das Experiment besteht aus zwei Teilen, wobei Ihr Verhalten im ersten Teil des Experiments keinen Einfluss auf den zweiten Teil des Experiments hat. Im ersten Teil des Experiments können Sie sich mit der Schieberegler-Aufgabe vertraut machen. Die Schieberegler-Aufgabe wird auf der nächsten Seite erklärt.

Ihre Auszahlung ergibt sich als Summe Ihrer Verdienste aus den beiden Teilen des Experiments und den 4€ für Ihr Erscheinen.

Falls Sie Fragen haben, heben Sie bitte Ihre Hand. Ein Experimentleiter wird dann an Ihren Platz kommen, um Ihre Frage leise zu beantworten.

Teil 1: Schieberegler-Aufgabe

Die Schieberegler-Aufgabe ist am rechten Bildrand illustriert. Bei der Schieberegler-Aufgabe müssen Sie den Schieberegler mit der Maus so positionieren, dass er auf der 50 einrastet. Dann gilt die Schieberegler-Aufgabe als gelöst. In diesem Fall zeigt das Feld unter dem Schieberegler eine 50 an.

Im ersten Teil des Experiments haben Sie 4 Minuten Zeit, um sich mit der Schieberegler-Aufgabe vertraut zu machen. Während der Schieberegler-Aufgabe werden Sie auf Ihrem Bildschirm jeweils 48 Schieberegler-Aufgaben sehen, die Sie in beliebiger Reihenfolge bearbeiten können. Sie sehen außerdem einen Knopf „Seite abschließen", die Anzahl der Schieberegler-Aufgaben, die Sie auf den vorherigen Seiten schon gelöst haben und die verbleibende Zeit.

Klicken Sie auf „Seite abschließen", so wird geprüft, wie viele Schieberegler-Aufgaben Sie auf der Seite gelöst haben. Sie gelangen dann zur nächsten Seite, die wieder 48 Schieberegler-Aufgaben beinhaltet.

Am oberen linken Bildrand wird Ihnen angezeigt, wie viele Schieberegler-Aufgaben Sie auf den vorherigen Seiten bereits erfolgreich gelöst haben.

Sie sehen außerdem die verbleibende Zeit. Ist die Zeit abgelaufen, so wird automatisch überprüft, wie viele Schieberegler-Aufgaben Sie auf der aktuellen Seite gelöst haben und die Gesamtzahl der gelösten Schieberegler-Aufgaben ermittelt. Für jede Schieberegler-Aufgabe, die Sie innerhalb der 4 Minuten gelöst haben, erhalten Sie 4 Cent.

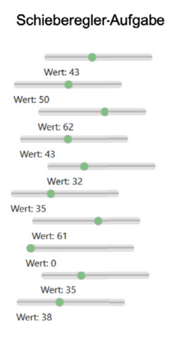

Figure 3.2: Instructions page 1 and 2 for the H_2 treatment

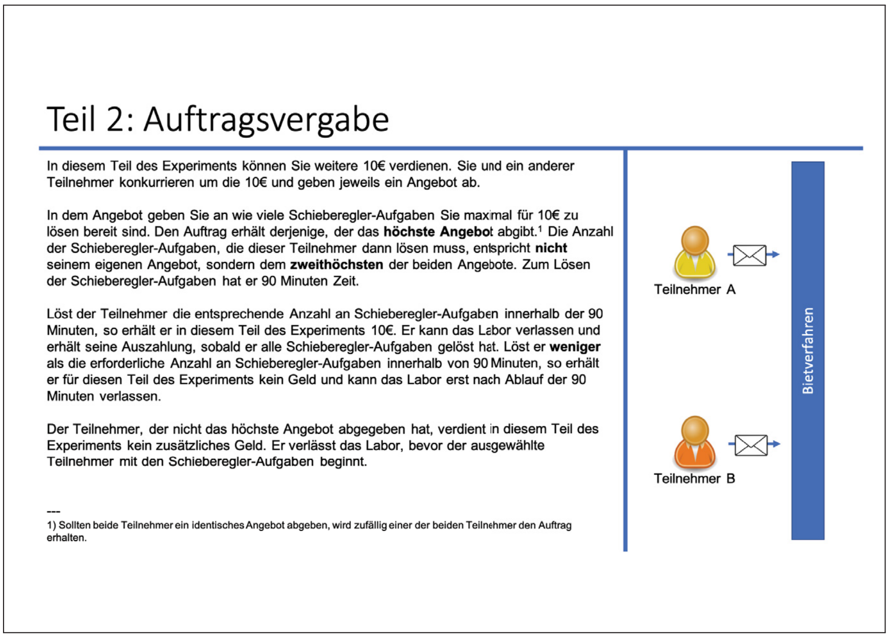

Figure 3.3: Instructions pages 3 and 4 for the H_2 reatment

H_8 *Treatment*

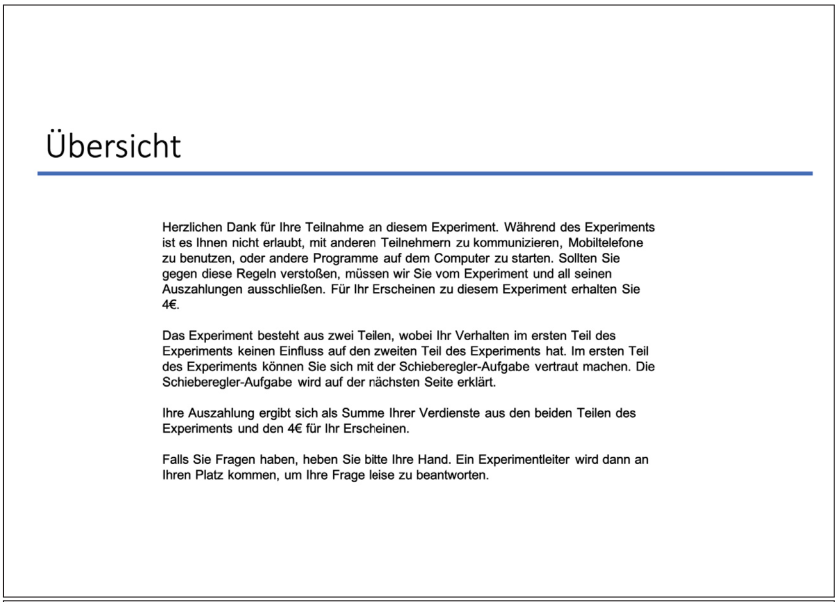

Übersicht

Herzlichen Dank für Ihre Teilnahme an diesem Experiment. Während des Experiments ist es Ihnen nicht erlaubt, mit anderen Teilnehmern zu kommunizieren, Mobiltelefone zu benutzen, oder andere Programme auf dem Computer zu starten. Sollten Sie gegen diese Regeln verstoßen, müssen wir Sie vom Experiment und all seinen Auszahlungen ausschließen. Für Ihr Erscheinen zu diesem Experiment erhalten Sie 4€.

Das Experiment besteht aus zwei Teilen, wobei Ihr Verhalten im ersten Teil des Experiments keinen Einfluss auf den zweiten Teil des Experiments hat. Im ersten Teil des Experiments können Sie sich mit der Schieberegler-Aufgabe vertraut machen. Die Schieberegler-Aufgabe wird auf der nächsten Seite erklärt.

Ihre Auszahlung ergibt sich als Summe Ihrer Verdienste aus den beiden Teilen des Experiments und den 4€ für Ihr Erscheinen.

Falls Sie Fragen haben, heben Sie bitte Ihre Hand. Ein Experimentleiter wird dann an Ihren Platz kommen, um Ihre Frage leise zu beantworten.

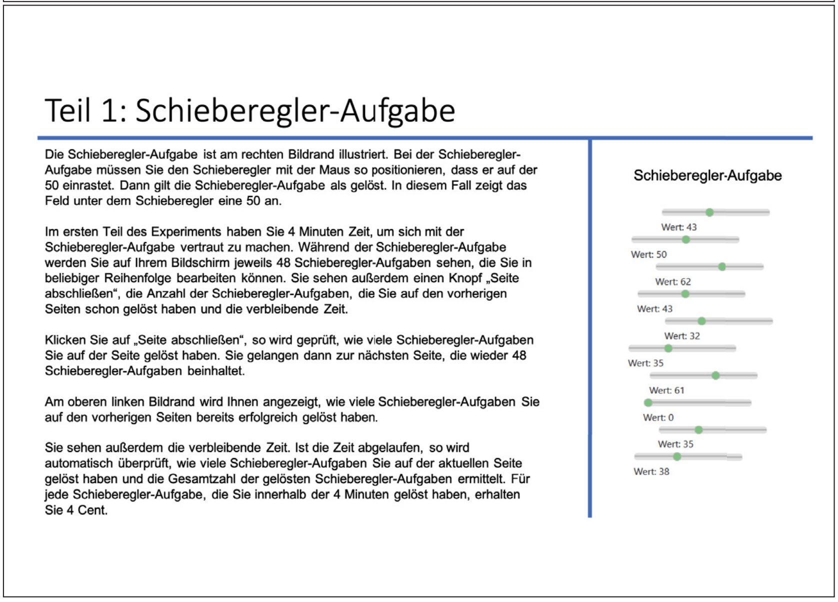

Teil 1: Schieberegler-Aufgabe

Die Schieberegler-Aufgabe ist am rechten Bildrand illustriert. Bei der Schieberegler-Aufgabe müssen Sie den Schieberegler mit der Maus so positionieren, dass er auf der 50 einrastet. Dann gilt die Schieberegler-Aufgabe als gelöst. In diesem Fall zeigt das Feld unter dem Schieberegler eine 50 an.

Im ersten Teil des Experiments haben Sie 4 Minuten Zeit, um sich mit der Schieberegler-Aufgabe vertraut zu machen. Während der Schieberegler-Aufgabe werden Sie auf Ihrem Bildschirm jeweils 48 Schieberegler-Aufgaben sehen, die Sie in beliebiger Reihenfolge bearbeiten können. Sie sehen außerdem einen Knopf „Seite abschließen", die Anzahl der Schieberegler-Aufgaben, die Sie auf den vorherigen Seiten schon gelöst haben und die verbleibende Zeit.

Klicken Sie auf „Seite abschließen", so wird geprüft, wie viele Schieberegler-Aufgaben Sie auf der Seite gelöst haben. Sie gelangen dann zur nächsten Seite, die wieder 48 Schieberegler-Aufgaben beinhaltet.

Am oberen linken Bildrand wird Ihnen angezeigt, wie viele Schieberegler-Aufgaben Sie auf den vorherigen Seiten bereits erfolgreich gelöst haben.

Sie sehen außerdem die verbleibende Zeit. Ist die Zeit abgelaufen, so wird automatisch überprüft, wie viele Schieberegler-Aufgaben Sie auf der aktuellen Seite gelöst haben und die Gesamtzahl der gelösten Schieberegler-Aufgaben ermittelt. Für jede Schieberegler-Aufgabe, die Sie innerhalb der 4 Minuten gelöst haben, erhalten Sie 4 Cent.

Figure 3.4: Instructions page 1 and 2 for the H_2 treatment

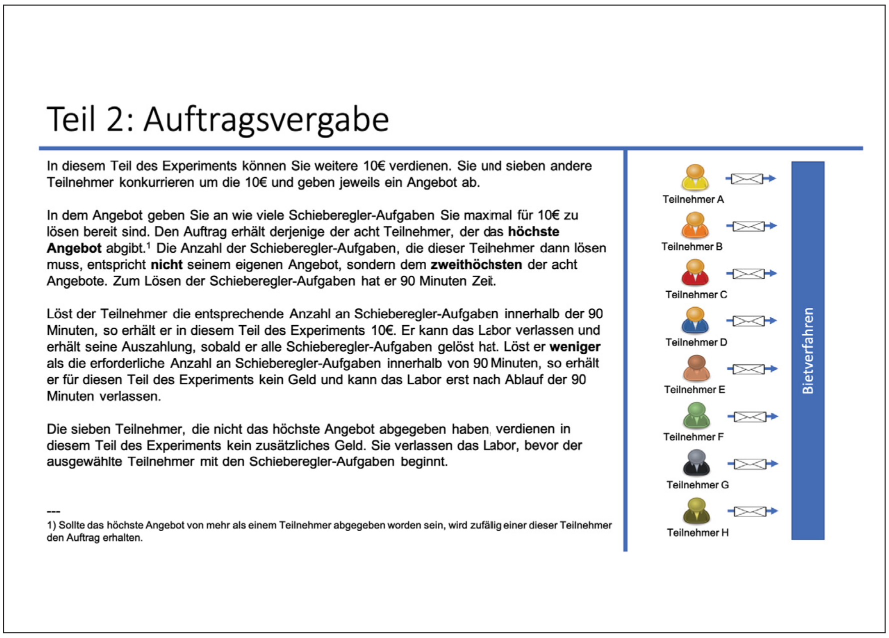

Figure 3.5: Instructions page 3 for the H_8 reatment

C_2^{2000} *Treatment*

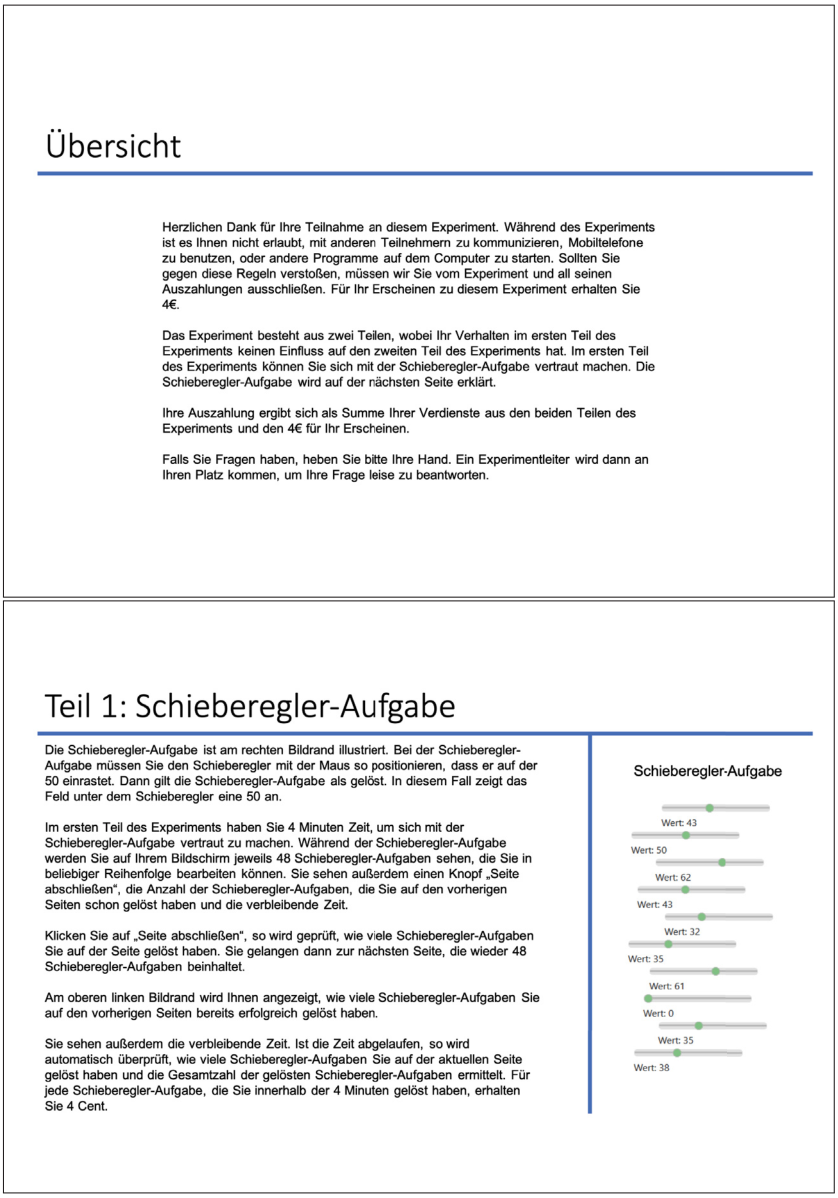

Übersicht

Herzlichen Dank für Ihre Teilnahme an diesem Experiment. Während des Experiments ist es Ihnen nicht erlaubt, mit anderen Teilnehmern zu kommunizieren, Mobiltelefone zu benutzen, oder andere Programme auf dem Computer zu starten. Sollten Sie gegen diese Regeln verstoßen, müssen wir Sie vom Experiment und all seinen Auszahlungen ausschließen. Für Ihr Erscheinen zu diesem Experiment erhalten Sie 4€.

Das Experiment besteht aus zwei Teilen, wobei Ihr Verhalten im ersten Teil des Experiments keinen Einfluss auf den zweiten Teil des Experiments hat. Im ersten Teil des Experiments können Sie sich mit der Schieberegler-Aufgabe vertraut machen. Die Schieberegler-Aufgabe wird auf der nächsten Seite erklärt.

Ihre Auszahlung ergibt sich als Summe Ihrer Verdienste aus den beiden Teilen des Experiments und den 4€ für Ihr Erscheinen.

Falls Sie Fragen haben, heben Sie bitte Ihre Hand. Ein Experimentleiter wird dann an Ihren Platz kommen, um Ihre Frage leise zu beantworten.

Teil 1: Schieberegler-Aufgabe

Die Schieberegler-Aufgabe ist am rechten Bildrand illustriert. Bei der Schieberegler-Aufgabe müssen Sie den Schieberegler mit der Maus so positionieren, dass er auf der 50 einrastet. Dann gilt die Schieberegler-Aufgabe als gelöst. In diesem Fall zeigt das Feld unter dem Schieberegler eine 50 an.

Im ersten Teil des Experiments haben Sie 4 Minuten Zeit, um sich mit der Schieberegler-Aufgabe vertraut zu machen. Während der Schieberegler-Aufgabe werden Sie auf Ihrem Bildschirm jeweils 48 Schieberegler-Aufgaben sehen, die Sie in beliebiger Reihenfolge bearbeiten können. Sie sehen außerdem einen Knopf „Seite abschließen", die Anzahl der Schieberegler-Aufgaben, die Sie auf den vorherigen Seiten schon gelöst haben und die verbleibende Zeit.

Klicken Sie auf „Seite abschließen", so wird geprüft, wie viele Schieberegler-Aufgaben Sie auf der Seite gelöst haben. Sie gelangen dann zur nächsten Seite, die wieder 48 Schieberegler-Aufgaben beinhaltet.

Am oberen linken Bildrand wird Ihnen angezeigt, wie viele Schieberegler-Aufgaben Sie auf den vorherigen Seiten bereits erfolgreich gelöst haben.

Sie sehen außerdem die verbleibende Zeit. Ist die Zeit abgelaufen, so wird automatisch überprüft, wie viele Schieberegler-Aufgaben Sie auf der aktuellen Seite gelöst haben und die Gesamtzahl der gelösten Schieberegler-Aufgaben ermittelt. Für jede Schieberegler-Aufgabe, die Sie innerhalb der 4 Minuten gelöst haben, erhalten Sie 4 Cent.

Figure 3.6: Instructions page 1 and 2 for the C_2^{2000} treatment

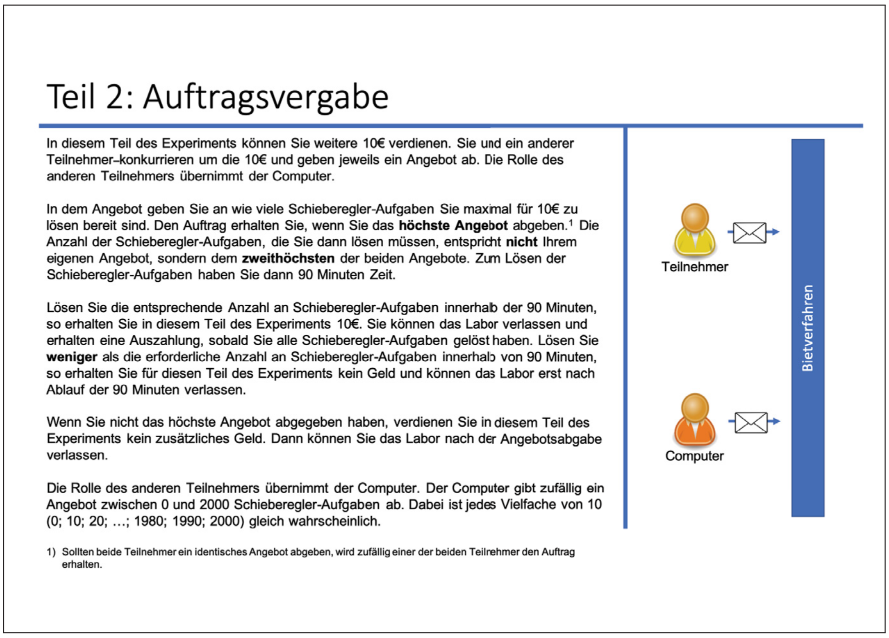

Figure 3.7: Instructions page 3 for the C_2^{2000} reatment

C_8^{2000} *Treatment*

Übersicht

Herzlichen Dank für Ihre Teilnahme an diesem Experiment. Während des Experiments ist es Ihnen nicht erlaubt, mit anderen Teilnehmern zu kommunizieren, Mobiltelefone zu benutzen, oder andere Programme auf dem Computer zu starten. Sollten Sie gegen diese Regeln verstoßen, müssen wir Sie vom Experiment und all seinen Auszahlungen ausschließen. Für Ihr Erscheinen zu diesem Experiment erhalten Sie 4€.

Das Experiment besteht aus zwei Teilen, wobei Ihr Verhalten im ersten Teil des Experiments keinen Einfluss auf den zweiten Teil des Experiments hat. Im ersten Teil des Experiments können Sie sich mit der Schieberegler-Aufgabe vertraut machen. Die Schieberegler-Aufgabe wird auf der nächsten Seite erklärt.

Ihre Auszahlung ergibt sich als Summe Ihrer Verdienste aus den beiden Teilen des Experiments und den 4€ für Ihr Erscheinen.

Falls Sie Fragen haben, heben Sie bitte Ihre Hand. Ein Experimentleiter wird dann an Ihren Platz kommen, um Ihre Frage leise zu beantworten.

Teil 1: Schieberegler-Aufgabe

Die Schieberegler-Aufgabe ist am rechten Bildrand illustriert. Bei der Schieberegler-Aufgabe müssen Sie den Schieberegler mit der Maus so positionieren, dass er auf der 50 einrastet. Dann gilt die Schieberegler-Aufgabe als gelöst. In diesem Fall zeigt das Feld unter dem Schieberegler eine 50 an.

Im ersten Teil des Experiments haben Sie 4 Minuten Zeit, um sich mit der Schieberegler-Aufgabe vertraut zu machen. Während der Schieberegler-Aufgabe werden Sie auf Ihrem Bildschirm jeweils 48 Schieberegler-Aufgaben sehen, die Sie in beliebiger Reihenfolge bearbeiten können. Sie sehen außerdem einen Knopf „Seite abschließen“, die Anzahl der Schieberegler-Aufgaben, die Sie auf den vorherigen Seiten schon gelöst haben und die verbleibende Zeit.

Klicken Sie auf „Seite abschließen“, so wird geprüft, wie viele Schieberegler-Aufgaben Sie auf der Seite gelöst haben. Sie gelangen dann zur nächsten Seite, die wieder 48 Schieberegler-Aufgaben beinhaltet.

Am oberen linken Bildrand wird Ihnen angezeigt, wie viele Schieberegler-Aufgaben Sie auf den vorherigen Seiten bereits erfolgreich gelöst haben.

Sie sehen außerdem die verbleibende Zeit. Ist die Zeit abgelaufen, so wird automatisch überprüft, wie viele Schieberegler-Aufgaben Sie auf der aktuellen Seite gelöst haben und die Gesamtzahl der gelösten Schieberegler-Aufgaben ermittelt. Für jede Schieberegler-Aufgabe, die Sie innerhalb der 4 Minuten gelöst haben, erhalten Sie 4 Cent.

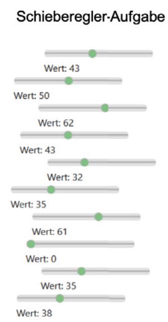

Figure 3.8: Instructions page 1 and 2 for the C_8^{2000} treatment

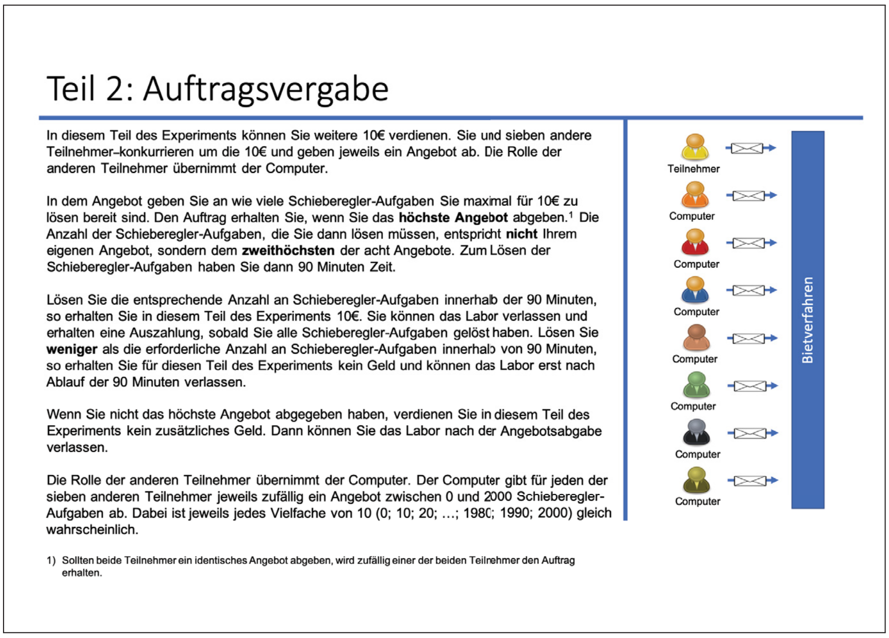

Figure 3.9: Instructions page 3 for the C_8^{2000} reatment

C_2^{4000} *Treatment*

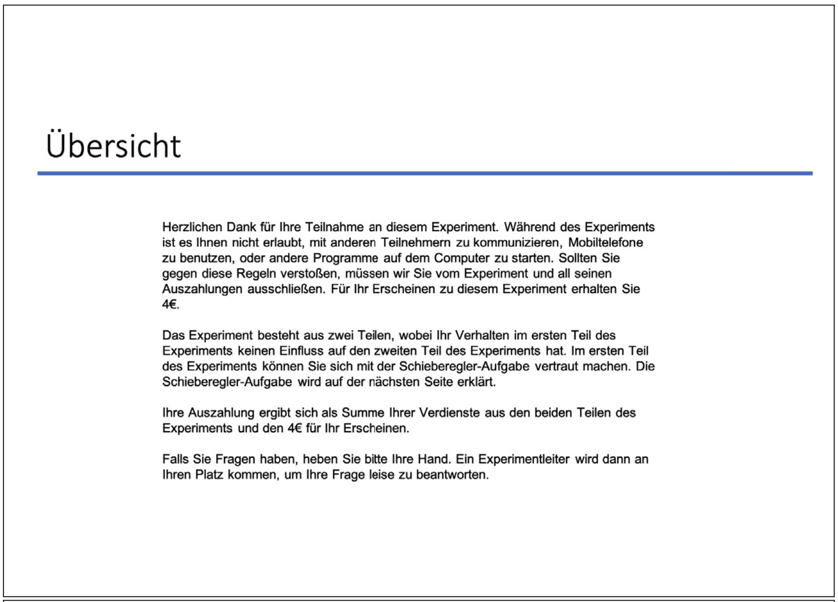

Übersicht

Herzlichen Dank für Ihre Teilnahme an diesem Experiment. Während des Experiments ist es Ihnen nicht erlaubt, mit anderen Teilnehmern zu kommunizieren, Mobiltelefone zu benutzen, oder andere Programme auf dem Computer zu starten. Sollten Sie gegen diese Regeln verstoßen, müssen wir Sie vom Experiment und all seinen Auszahlungen ausschließen. Für Ihr Erscheinen zu diesem Experiment erhalten Sie 4€.

Das Experiment besteht aus zwei Teilen, wobei Ihr Verhalten im ersten Teil des Experiments keinen Einfluss auf den zweiten Teil des Experiments hat. Im ersten Teil des Experiments können Sie sich mit der Schieberegler-Aufgabe vertraut machen. Die Schieberegler-Aufgabe wird auf der nächsten Seite erklärt.

Ihre Auszahlung ergibt sich als Summe Ihrer Verdienste aus den beiden Teilen des Experiments und den 4€ für Ihr Erscheinen.

Falls Sie Fragen haben, heben Sie bitte Ihre Hand. Ein Experimentleiter wird dann an Ihren Platz kommen, um Ihre Frage leise zu beantworten.

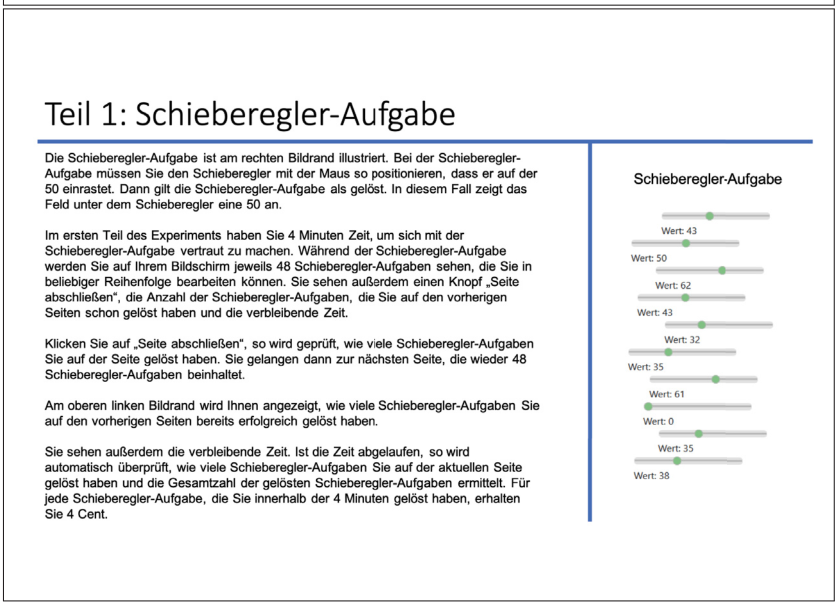

Teil 1: Schieberegler-Aufgabe

Die Schieberegler-Aufgabe ist am rechten Bildrand illustriert. Bei der Schieberegler-Aufgabe müssen Sie den Schieberegler mit der Maus so positionieren, dass er auf der 50 einrastet. Dann gilt die Schieberegler-Aufgabe als gelöst. In diesem Fall zeigt das Feld unter dem Schieberegler eine 50 an.

Im ersten Teil des Experiments haben Sie 4 Minuten Zeit, um sich mit der Schieberegler-Aufgabe vertraut zu machen. Während der Schieberegler-Aufgabe werden Sie auf Ihrem Bildschirm jeweils 48 Schieberegler-Aufgaben sehen, die Sie in beliebiger Reihenfolge bearbeiten können. Sie sehen außerdem einen Knopf „Seite abschließen“, die Anzahl der Schieberegler-Aufgaben, die Sie auf den vorherigen Seiten schon gelöst haben und die verbleibende Zeit.

Klicken Sie auf „Seite abschließen“, so wird geprüft, wie viele Schieberegler-Aufgaben Sie auf der Seite gelöst haben. Sie gelangen dann zur nächsten Seite, die wieder 48 Schieberegler-Aufgaben beinhaltet.

Am oberen linken Bildrand wird Ihnen angezeigt, wie viele Schieberegler-Aufgaben Sie auf den vorherigen Seiten bereits erfolgreich gelöst haben.

Sie sehen außerdem die verbleibende Zeit. Ist die Zeit abgelaufen, so wird automatisch überprüft, wie viele Schieberegler-Aufgaben Sie auf der aktuellen Seite gelöst haben und die Gesamtzahl der gelösten Schieberegler-Aufgaben ermittelt. Für jede Schieberegler-Aufgabe, die Sie innerhalb der 4 Minuten gelöst haben, erhalten Sie 4 Cent.

Figure 3.10: Instructions page 1 and 2 for the C_2^{4000} treatment

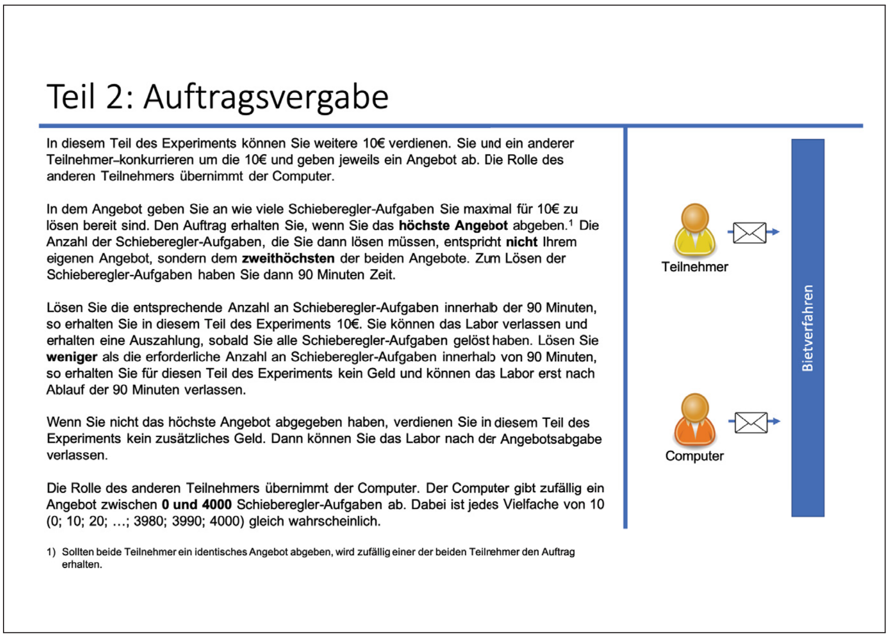

Figure 3.11: Instructions page 3 for the C_2^{4000} reatment

CHAPTER 4

PREFERENCES AND DECISION SUPPORT IN COMPETITIVE BIDDING

Abstract

We examine bidding behavior in first-price sealed-bid and Dutch auctions, which are strategically equivalent under standard preferences. We investigate whether the empirical breakdown of this equivalence is due to (non-standard) preferences or due to the different complexity of the two formats (i.e., a different level of mathematical/individual sophistication needed to derive the optimal bidding strategy). We first elicit measures of individual preferences and then manipulate the degree of complexity by offering various levels of decision support. Our results show that the equivalence of the two auction formats only breaks down in the absence of decision support. This indicates that the empirical breakdown is caused by differing complexity between the two formats rather than non-standard preferences.

4.1 INTRODUCTION

The first-price sealed-bid auction (FSPBA) and the Dutch auction (DA) are two of the most frequently used auction formats. In an FPSBA, bidders simultaneously submit "sealed" bids to the seller and the highest bidder receives the object and pays his bid. In a DA, the seller starts at a high initial ask price and gradually decreases the ask price until the first bidder stops the auction, receives the item, and pays the stop price. With slight variations, both the FPSBA and the DA generate billions of dollars in revenue each year. Governments and private firms frequently use the FPSBA for procurement in construction and to subcontract with suppliers. Variants of the FPSBA are also used to organize online labor markets for freelancers

(Hong et al., 2016). The DA is traditionally used to sell flowers in the Netherlands and, for example, the annual sales of the seven Dutch flower auctions exceeded 2.9 billion dollars in 1996 (Kambil and Van Heck, 1998). Federal banks and firms use variants of the DA to sell securities and refinance credit.[1] Furthermore, the DA can be found on fish and fresh-produce markets (e.g., Cassady, 1967).

With regard to the actual implementation of auctions, offline auctions as a mechanism to buy and sell goods is not a new phenomenon, but the specific use of online auctions has experienced tremendous growth in the new media era (Hennig-Thurau et al., 2010).[2] According to Ariely and Simonson (2003), the popularity of online auctions is due to the following three particular features: First, online auctions overcome geographical limitations, such that people from all over the world have the opportunity to submit their bids in any auction. Second, electronic auctions on the Internet allow for more flexibility among sellers and bidders, because the duration of an auction can be several days (or even weeks), and there is the possibility of asynchronous bidding. Third, auctions can be organized at substantially lower costs, which translates into lower commission fees and hence higher participation rates among sellers and buyers.

The increase in the use of Internet-based auctions has led to a rise in the demand for expert services. Indeed, there is an increasing number of consulting firms specializing in auctions (e.g., Market Design Inc.) and major economic consulting companies offer services regarding auctions and bidding (e.g., The Brattle Group, NERA). These services typically include all aspects relevant for setting up and participating in auctions (e.g., bid tracking, bidding strategy, auction rules and design, training, provision of input to regulators). Moreover, the design of decision support systems (DSS) has also attracted considerable interest (e.g., Hass et al., 2013; Park et al., 2010; Kayande et al., 2009; Todd and Benbasat, 1991). For example, several

[1]Note, however, that these examples typically auction off multiple units and that the auctions are then modified such that they usually do not discriminate between different bidders but apply a uniform-pricing rule.

[2]See also Haruvy and Popkowski Leszczyc (2009), who provide an overview of the implications of economically relevant aspects that are characteristic of Internet auctions.

patents have been filed for (automated) bid-advising systems that account for the auction structure and risk attitudes of rival bidders based on historical data among other things (see, e.g., Guler et al., 2002, 2003, 2009; Zhang and Guler, 2013). At the same time, technological advances and the use of Internet auctions means that relevant information can be provided more easily and faster in the course of an online auction. We take these observations as a starting point to address the implications of decision support systems in different formats of (online) auctions.[3]

Theory suggests that the FPSBA and the DA yield the same revenue as both formats are strategically equivalent. However, this strong theoretical result breaks down empirically. Previous research suggests three possible explanations: opportunity costs (Carare and Rothkopf, 2005; Katok and Kwasnica, 2007), preferences (Weber, 1982; Nakajima, 2011; Lange and Ratan, 2010; Belica and Ehrhart, 2013; Ehrhart and Ott, 2014), and complexity of the decision (Cox et al., 1983). We analyze the role of preferences and complexity while controlling for opportunity costs. Our results indicate that the non-equivalence is driven by the difference in complexity of competitive bidding in the two auction formats rather than by individual (non-standard) preferences.

The empirical breakdown of this equivalence is a robust observation in experimental settings both in the laboratory and in the field. However, the direction of the deviation is non-conclusive. On the one hand, Coppinger et al. (1980) and Cox et al. (1982) find that the FPSBA yields higher revenue than the DA in a controlled laboratory setting. On the other hand, in a field experiment on an Internet auction platform, Lucking-Reiley (1999) finds that the DA generates higher revenue than the FPSBA.

Differences in opportunity costs can explain these differences. In a DA, bidders have an incentive to accept a high price and stop the auction early, because they have to frequently monitor the price clock or even have to physically return to the auction site to check for updates in prices as long

[3] Adomavicius et al. (2013) and Bichler et al. (2017) who analyze the role of decision support systems in combinatorial auctions.

as the auction is running. Such costs do not occur in the (static) FPSBA which ends immediately after the (simultaneous) submission of bids.

Carare and Rothkopf (2005) show theoretically that such increased opportunity costs increase the optimal bid. In a DA, Cox et al. (1983) and Katok and Kwasnica (2007) analyze the trade-off between opportunity costs and additional utility from suspense, i.e., from a joy of gambling. Both articles provide evidence that increasing opportunity costs by increasing payoffs or by decreasing the clock speed, respectively, increases bids in a DA.[4] In contrast to their approach, our goal is to assess the predictive power of different preference-based theories for observed bidding and to analyze the effect of complexity. Hence, we eliminate confounding differences in opportunity costs by holding the time per auction format and thus the opportunity costs from participation constant. In addition, we hold the action set, i.e., the set of feasible bids, constant across the two formats which allows a direct comparison of the two auctions.

In the absence of opportunity costs, the strategic-equivalence result rests on the assumption that bidders have standard preferences, i.e., they derive utility only from realized personal payoffs. In addition, the utility function is global in the sense that the effect of wealth changes does not depend on whether such changes occur in the gain or loss domain or whether they are certain or generated by a lottery. With regard to the departures from standard preferences, we study expectations-based reference-dependent and Allais-type preferences. We focus on these two specifications, because they are frequently used to explain decision making under uncertainty.[5]

Under reference dependence, the bidder compares gains and losses in wealth relative to a reference point (Kahneman and Tversky, 1979). In this comparison, the bidder is assumed to be loss averse and puts more weight

[4] In contrast to the observation by Katok and Kwasnica (2007), there is anecdotal evidence from the Dutch flower auctions that faster clock speeds result in higher prices (Kambil and Van Heck, 1998).

[5] Reference dependence as proposed by Kahneman and Tversky (1979) is the most cited theory on risky decision making (Kim et al., 2006). Allais-type preferences are an early critique of expected utility theory (EUT) (Allais, 1953) and are empirically very robust in explaining deviations from predictions under standard preferences (Kahneman and Tversky, 1979; Camerer, 1989; Weber, 2007).

on negative deviations from this reference point (losses) than on equivalent positive deviations (gains). Loss aversion contradicts the global-utility assumption of standard preferences because the bidder considers changes in wealth with respect to a local reference point. The specification of the reference point is subject to debate. Kőszegi and Rabin (2006, KR) propose expectations-based reference dependence, i.e., the reference point is stochastic and given by the rational expectations that the individual holds over the outcomes of a risky decision. In the following, we will denote expectations-based reference-dependent preferences as *KR preferences*.

Individuals with *Allais-type (AT) preferences* prefer outcomes that are generated with certainty to the same outcomes that are generated by a risky lottery (e.g., Andreoni and Sprenger, 2010). This difference is most prevalent in the Allais paradox (Allais, 1953). Here, subjects prefer a degenerate lottery over a risky one with a higher expected value but reverse their choice if both lotteries are monotonically transformed and become both risky (the so-called common-ratio effect, CRE). This reversal is inconsistent with standard preferences as it violates the crucial independence axiom of EUT (Savage, 1954; Anscombe and Aumann, 1963). According to this axiom, decisions between lotteries should not depend on consequences that do not differ between the lotteries.

We make use of data from a two-stage experiment in which we first elicit the preferences of all subjects that participate in our experiment. In this first stage, we utilize the procedure of Abdellaoui et al. (2007) and elicit individual preferences in a fully non-parametric procedure, i.e., without imposing any assumption on the functional form of utility. Furthermore we measure to what extent participants exhibit Allais-type preferences by utilizing a metric version of the CRE (e.g., Beattie and Loomes, 1997; Dean and Ortoleva, 2014; Schmidt and Seidl, 2014).

Preference theories assume Bayesian rationality in the sense that bidders derive and process probabilities correctly. However, bidding in auctions can be a demanding problem. In deriving the optimal bid, the bidder faces a trade-off between increasing his winning probability by submitting a higher bid and increasing his winning profit by submitting a lower bid. Individual preferences determine the optimal bid that balances these diametric effects. However, this optimization requires a certain level of mathematical sophistication. It is thus possible that the observed differences between bidding behavior is due to different levels of complexity of the two auction formats. In other words, bidders can make mistakes, e.g., in deriving the winning probability associated with their bid, and these mistakes might differ between the two formats.

We design a DSS to reduce the complexity and assist bidders in deriving the optimal bid that corresponds to their individual preferences. We vary the auction format within-subjects and the level of decision support between-subjects. Subjects either have no decision support (No DSS treatment) or they have medium (Medium DSS treatment) or full support (Full DSS treatment) to assist bidding. The decision support system is a computerized overlay displaying additional information. Medium DSS shows the winning probability whereas Full DSS additionally provides expected profits. Although this information is redundant for fully rational decision makers, it is non-trivial to derive and providing such information greatly reduces the complexity of optimal bidding.[6, 7]

Our results highlight the role of decision support systems. In line with the literature, we find significant differences between auction formats when bidders do not receive decision support. However, differences vanish between participants once we provide decision support. This indicates that the observed differences in bidding behavior between the FPSBA and the DA are due to different levels of complexity rather than non-standard pref-

[6] As such, this is different, for instance, from setups in which bidders get information which is not readily available and must be acquired at a cost (see, e.g., Gretschko and Wambach, 2014, for a model with heterogenous prior information and Gretschko and Rajko, 2015, for an experimental treatment), or in which bidding by rival bidders in multi-object auctions conveys important information (see, e.g., Gretschko et al., 2014).

[7] Our implementation of decision support is primarily a mean to analyze the role of complexity in competitive bidding, the design of such DSS is also of interest in itself. Several patents have been filed for (automated) bid-advising systems that take into account, for example, the auction structure and risk attitudes of rival bidders based on historical data (see, e.g., Guler et al., 2002, 2003, 2009; Zhang and Guler, 2013). Our DSS implementation resembles such automated bidding advice that estimates competitors' bidding behavior in a given auction format.

erences. In addition, our tests show that bidding behavior strongly depends on participants' risk aversion. The influence of individual loss aversion and Allais-type preferences is not significant and cannot explain differences in bidding behavior. Our results thus highlight that from a consulting perspective, it seems to be more important to support decision makers in the derivation of optimal bidding strategies than to focus on the choice of the auction format.

The paper proceeds as follows. The next section introduces the model environment and theoretically analyzes the effect of different preference specifications on optimal bidding in the FPSBA and the DA. Section 4.3 presents our experimental design and our implementation of decision support. We report our results in Section 4.4. Section 4.5 concludes.

4.2 THEORY

In this section, we first describe the two auction mechanisms. We then characterize the equilibria in both auction formats for standard preferences (SP), Kőszegi-Rabin (KR) preferences, and Allais-type (AT) preferences. We analyze the optimal bidding behavior of one bidder given a bidding strategy of the competitor.

In both auction formats, two bidders compete for one indivisible item and the highest bidder wins. Let $P = \{p_1, \ldots, p_n\}$ be a discrete price grid. In the FPSBA, each bidder places a *bid* $b \in P$ at which he is willing to buy the item. In the DA, each bidder decides for every *ask* $a \in P$ whether to accept it or not. In the FPSBA, the *price* corresponds to the highest bid, whereas in the DA, it corresponds to the highest accepted ask. The winning bidder receives the item and pays the price. If the bidder does not win the auction, he does not receive the item and does not pay anything.[8] In both auction formats bidders face a trade-off between improving their probability of winning and increasing their profit in case of winning.

To derive the equilibrium bidding strategy in the discrete FPSBA, we follow Cai et al. (2010). For the dynamic course of the DA, we adopt the

[8]Ties are broken at random with equal probability to receive the item.

modeling approach of Bose and Daripa (2009). In the DA, the seller starts the auction with the highest ask p_n. She then approaches each bidder sequentially asking whether or not the bidder accepts that ask. Which bidder is asked first is randomly determined at the beginning of each offer. Each bidder has the same chance to be asked first. In case that the bidder who is asked first rejects the offer, the seller offers the same ask to the other bidder.

4.2.1 STANDARD PREFERENCES

The term standard preferences covers all preferences that are purely outcome-based and only consider the own payoff. This means an individual has standard preferences if the utility function is global and only depends on one's own payoff DellaVigna (2009) .

Proposition 1 (Standard Preferences)**.** *The FPSBA and the DA are strategically equivalent, which implies that they yield the same revenue (Vickrey, 1961).*

The crucial observation to this result is that the information revealed during the descending of the price clock in the DA does not change the trade-off between a bidder's winning probability and his profit in case of winning. Suppose a bidder bids $b = p_k$ in an FPSBA. This bidder enters a DA with the plan to accept the ask $a = p_k$, because the ex-ante problem is identical for the two formats. As the price clock is approaching p_k, two things may happen. First, the competitor accepts an ask greater than p_k. In this case, the auction ends and the bidder cannot react to this information. Second, the price continues to fall which increases the probability to win. However, the marginal trade-off stays the same. This is due to the fact that a bidder derives his optimal bidding strategy under the assumption that he has the highest valuation. Hence, the bidder sticks to his plan and waits for the ask p_k.

4.2.2 EXPECTATIONS-BASED REFERENCE POINTS

In contrast to individuals with standard preferences, an individual with reference-dependent preferences does not only care about his absolute payoff, but also compares the outcome to a reference point. Therefore, the utility function of such a bidder consists of two parts. First, the term $u(x)$ corresponds to utility derived from payoff x as under standard preferences. Second, the term $n(x, r)$ corresponds to gain-loss utility that evaluates the outcome x against a reference level r (Kahneman and Tversky, 1979). Following the approach of Kőszegi and Rabin (2006) the gain-loss utility is defined piece-wise as

$$n(x, r) = \mu\left(u(x) - u(r)\right),$$

where

$$\mu(z) := \begin{cases} \eta z & \text{if } z > 0 \\ \eta\lambda z & \text{if } z \leq 0. \end{cases}$$

Here $\eta > 0$ determines how important the relative component is compared to the absolute payoff. Furthermore, λ represents the level of loss aversion which weighs negative deviation from the reference point (losses) relative to positive deviations (gains). If $\lambda > 1$, the bidder is loss averse, i.e., losses hurt him more than equally sized gains please him. If $\lambda = 1$, the agent is loss-neutral, and if $\lambda < 1$, the agent is gain-seeking. Total utility is the sum of both parts and given by $u^{\mathrm{KR}}(x, r) = u(x) + n(x, r)$. We follow the literature and focus on the effect of loss aversion by assuming that utility of payoff $u(x)$ is linear. Hence, gain-loss utility $n(x, r)$ is a two-piece linear function.

Kőszegi and Rabin (2006) assume that the reference point is stochastic and formed by the rational expectations of the bidder. They introduce the concept of a *personal equilibrium* which requires that the bidder has rational expectations about his own behavior and behaves consistently with his plans. Specifically, they propose that the bidder evaluates each possible outcome x under the winning probability $\Pr(x|b)$ against all other possible outcomes under this distribution. This modification has recently been successful in describing various empirical observations from laboratory endowment effects to labor supply in the field (e.g., Sprenger, 2010; Ericson and Fuster, 2011; Crawford and Meng, 2011).

Proposition 2 (Expectations-based reference point). *A revenue ranking of the FPSBA and the Dutch auction is not possible.*

In the FPSBA, loss aversion implies that bidders want to reduce the difference between the payoff in case of winning and in case of losing the auction. As a consequence, subjects with a higher degree of loss aversion place higher bids than less loss-averse subjects. In the FPSBA, there exists an almost everywhere unique optimal bidding strategy (Eisenhuth and Ewers, 2012).

In contrast to the FPSBA, there might be several consistent bidding strategies in the DA. For example, it may be optimal for a subject to accept a high offer p if it planned to do so, whereas it is optimal for the same subject to wait for a smaller offer p' if her initial plan was to accept only a small offer p'. Different plans induce different reference points and thereby different optimal bidding strategies. Since several reasonable reference points can exist in the DA, we do not get a unique bidding prediction but a set of optimal bidding strategies. Applying a refinement and identifying the bidding strategy with the highest expected utility might not be possible as the optimality of a bidding strategy can change during the dynamic course of the auction (Ehrhart and Ott, 2014).

As shown in the Appendix 4.6.2, it may well be the case that for a given valuation the lowest optimal bid in the DA is lower than the optimal bid in the FPA, whereas the highest optimal bid in the DA is higher than the optimal bid in the FPA. As a consequence, a revenue ranking is not possible in general.

4.2.3 ALLAIS-TYPE PREFERENCES

Allais-type preferences violate the independence (or substitution) axiom, which is essential for EUT (Allais, 1953; Savage, 1954; Anscombe and Aumann, 1963). The independence axiom states that an individual who is indifferent between two lotteries should also be indifferent between these lotteries if the probabilities of both lotteries are multiplied by $\rho \in (0, 1]$. That is, if one scales the probabilities of both lotteries by a common ratio, the preference ordering is not affected under EUT. Grimm and Schmidt (2000) show that this independence requirement is a necessary and sufficient condition for strategic equivalence between the FPSBA and the DA.

Kahneman and Tversky (1979) report that subjects have a preference for certainty, i.e., outcomes in a degenerate lottery. In their experiment, a majority of individuals reveals that they prefer a degenerate lottery over a risky one but reverse this choice if both lotteries are scaled by ρ such that both now become risky. Thus, participants violate the independence requirement. This so-called "Allais paradox" (Allais, 1953) is empirically very robust, although reverse Allais-type preferences (i.e., a preference for risky outcomes if a certain outcome is available) have also been observed experimentally (Camerer, 1989; Weber, 2007).

Proposition 3 (Allais-type preferences)**.** *The DA yields higher revenue than the FPSBA if bidders have Allais-type preferences. The FPSBA generates higher revenue if bidders have reverse Allais-type preferences (Weber, 1982; Nakajima, 2011).*

The intuition is that the current price in the DA is augmented by a psychological premium for certainty for individuals with Allais-type preferences. This premium makes it more attractive to accept a high price in the DA than in the FPSBA in which all bids imply uncertainty. In other words, the DA offers a certain payoff in the given round against a risky lottery (prices in future rounds), whereas the FPSBA only offers a risky lottery.[9]

[9]We note that this overbidding only works given our organization of the DA, because we resolve the order in which the seller approaches the two bidders at the beginning of each period. If we had broken ties at random after each round, which is frequently done in DA implementations, the current price would actually be risky as well and Allais-type preferences would coincide with standard preferences.

4.3 EXPERIMENT

In this section, we first introduce our experimental design and then review previous research that examines the equivalence of the first-price sealed-bid auction and the Dutch auction experimentally.

4.3.1 DESIGN

Each subject participated in 18 FPSBA and 18 DA. Each auction consists of one participant and one bidding robot as bidders. The valuations of the participant are drawn from the set $\{6, 10, 14, 18, 22, 26, 30, 34, 38\}$ EUR. In each format, every participant is assigned each valuation twice in order to make participants' bidding behavior as comparable as possible. The bidding robot draws one price from $P = \{0, 1, \ldots, 21\}$ EUR according to a uniform distribution. This realization is the robot's bid in the FPSBA and its stopping price in the DA. We use a bidding robot as the competitor for three reasons. First, we do not want our results to be confounded by other-regarding preferences that are not considered in any of the models presented in Section 4.2. Second, we effectively reduce the strategic problem to a decision problem by fixing the strategy of the competitor. This makes it easier for subjects to focus on their optimal strategy by breaking the dynamics of higher-order beliefs.[10] Third, we are able to precisely calculate the winning probability and the expected profit. The provision of this information depends on the DSS treatment status.

Auction formats

In our experiment, we analyze the following two auction formats:

- **FPSBA** In the FPSBA, the computer screen informs the participants about their valuation and features a testing area. In this area, participants can explore the consequences of a particular bid on their

[10]Note that most work that analyzes strategic interaction in auctions assumes that subjects' preferences are common knowledge and that only valuations are private information. However, one cannot ensure common knowledge in reality.

profit and, depending on the DSS treatment, on the winning probability and the expected profit (see below). Participants are further informed about the remaining time of the current round. Finally, they enter their actual bid and submit this bid by pressing a button. After submitting their bid, participants are immediately informed whether they have won the auction and about the remaining time the current auction lasts. When the round has timed out, a feedback screen informs the subjects about their valuations, the winning bid, whether or not they received the item, and their profit for the this round.

- **DA** In the DA, the computer screen informs participants about their valuation and displays the current price, the time until the next price, and the next price. As in the FPSBA, participants are informed about their profit given both the current and the next price. Depending on the DSS treatment, participants are also informed about the probability to be offered the current price and the next price as well as the associated expected profits (see below). Finally, participants can accept the current price by pressing a button. After either the participant or the computer bidder has accepted the current price, participants are immediately informed whether they have won the auction and about the remaining time the current auction lasts. When the round has timed out, participants receive the same feedback as in the FPSBA.

Decision support system

The theoretical analysis on the role of preferences in Section 4.2 highlights the fact that deriving the optimal bid depends on the following aspects: (i) the profit from winning with the chosen bid, $v^i - b^i$, (ii) the probability to win with the chosen bid, $\Pr(\text{win}|b^i)$, and (iii) the expected utility derived from the combination of the former two. The latter depends on the individual preferences whereas the former two are identical across all theories. Hence, we design a DSS that assists the bidder by providing (i) the profit from winning, (ii) the winning probability, and (iii) the expected profit which is the product of (i) and (ii).

Any deviation from bidding predictions can result from two sources: an omitted preference specification or problems in deriving the optimal bid. Our DSS allows us to disentangle the role of preferences from the impact of a lack of mathematical sophistication (complexity). This is because in the experiment, we fix the bidding strategy of the competitor and hence reduce the strategic problem of finding mutual best responses to the problem of finding a one-sided best response (i.e., an optimization or decision problem). We can thus objectively state expected profits and winning probabilities that should help participants derive the bid that maximizes the expected utility based on their actual preference specification. In other words, we implement the DSS to analyze whether observed bids are due to the underlying preferences or the complexity of the auction.

Specifically, the DSS varies between participants regarding the information a bidder receives during an auction. There are three nested levels of DSS: No, Medium, and Full DSS. In the FPSBA, the information is given for the current test bid. In the DA, the information is given for both the current and the next price. We vary the information content of the DSS between participants. The information content in each condition is as follows:

- **No DSS** In the FPSBA, subjects see the *profit if bid is successful* which is the profit their test bid would generate given that they won the auction. In the DA, subjects see the *profit at given price* which is the profit they would make if they accept the current price or if they now decide to accept the next price.

- **Medium DSS** Subjects have the same information as in No DSS. In addition, in the FPSBA, they also see the *winning probability* of their test bid which is the probability of having a higher bid than the competitor plus the probability of having the same bid and being selected as winner by the tie-breaking rule. In the DA, subjects receive the *probability to be offered the given price* for both the current and the next price. The probability to receive the current price p_k is trivially given by 1. However, the probability to be offered the next ask, H^i_k, is highly non-trivial to derive (see Section 4.6.2 for details).

- **Full DSS** Subjects have the same information as in Medium DSS. In addition, in the FPSBA, they also see the expected profit of their test bid. In the DA, subjects see the expected profit of the next price. In the FPSBA, the expected profit is the product of the winning probability and the profit if the bid is successful. In the DA, the expected profit is the product of the probability to be offered the given price and the profit at the given price.

We are not aware of any other work that incorporates decision support in auctions. Armantier and Treich (2009) elicit both subjective probabilities and risk preferences in an attempt to find an explanation for overbidding in experimental first-price auctions. The authors report that participants underestimate their winning probability and overbid. Furthermore, they investigate the effect of a feedback system regarding winning probabilities. The feedback is implemented as follows. Participants are asked to predict their winning probability and they are given feedback regarding the precision of their prediction at the end of each round. As such, their feedback system is designed to induce learning whereas learning is not necessary in our setup as participants are given support before (FPSBA) or during (DA) the auction. They show that overbidding is reduced if their feedback system is in place.

Subjects

Table 4.1 provides an overview of participants characteristics in the different treatments.

Risk aversion is measured as the are under the curve on the gain domain, i.e. the integral of the estimated utility function on the gain domain. We normalize the domain of utility to [0,1] by dividing each elicited gain by the maximum gain. We interpolate linearly between the elicited points and use a geometric approach to calculate the area. In case of risk aversion the measure is smaller 0.5. A risk seeking individual has a measure larger than 0.5 and a risk neutral subject has a measure equal to 0.5.

Loss aversion relates the slope of utility in the gain domain to its slope in the loss domain. Kahneman and Tversky (1979) define loss aversion by $-u(-x) > u(x)$ for every $x > 0$. We measure the coefficient of loss aversion as the mean of $-u(-x)/u(x)$ for all elicited values x.

Allais-type preferences are measured by metric measure of the common-ratio effect (CRE) to assess the preference reversal due to violations of the independence axiom. Participants exhibiting the common-ratio effect show a preference reversal such that, they have a preference for certain outcomes. Participants with a CRE of 0 are consistent with expected utility theory, a CRE larger zero indicates Allais-type preferences and subjects with a CRE smaller zero have reverse Allais-type preferences.

Subjects' numeracy is rated according to a combination of the Schwartz et al. (1997) and the *Berlin Numercy Test* that assess the understanding of fundamental concepts of probability. Subjects have to answer seven questions and the variable numeracy reflects how many of these questions were answered correctly.

Table 4.1: Summary statistics by treatment

Treatment	No DSS		Medium DSS		Full DSS		
First format	FPSBA	DA	FPSBA	DA	FPSBA	DA	p-value
Risk aversion	0.461	0.466	0.499	0.528	0.441	0.526	0.52
	(0.167)	(0.105)	(0.106)	(0.140)	(0.143)	(0.103)	
Loss aversion	1.842	1.396	1.673	1.352	2.088	1.407	0.08
	(0.860)	(0.474)	(0.713)	(0.506)	(0.842)	(0.450)	
Allais-type	2	2.714	3.857	3.667	4.333	2.267	0.90
	(13.90)	(2.301)	(4.605)	(6.199)	(9.566)	(18.25)	
Numeracy	4.333	4.714	3.929	4.833	4.167	4.867	0.24
	(1.291)	(2.199)	(1.141)	(1.267)	(1.403)	(0.990)	
Participants	15	14	14	12	12	15	

Notes: Reported are means of each variable with standard deviation in parentheses. The last column presents the results of a Kruskal-Wallis tests for the equality of populations.

Organization

The auctions were the second stage of the experiment. In the first stage, which was conducted one week before the second, participants' preferences were elicited. Detailed results are reported in Zeppenfeld (2015).[11] Both stages of the experiment were conducted in the Cologne Laboratory for Economic Research (CLER) at the University of Cologne, Germany.[12] Using the recruiting system *ORSEE* (Greiner, 2015) , we invited a random sample of the CLER's subject pool via email. The whole experiment was computerized using the programming environment *z-tree* (Fischbacher, 2007).

In both stages, payoffs were stated in Euros (EUR). Participants were paid out in private for the entire course of experimentation after the completion of the second stage. In the second stage, one auction of each auction format was randomly chosen to be payoff-relevant. All 82 participants were paid their total net earnings, i.e., their earnings from the auctions and their earnings from first stage of the experiment. The average payoff for the entire experiment was 36.63 EUR corresponding to approx. 45.54 USD at the time of the payment.[13]

4.3.2 OPPORTUNITY COSTS AND ACTION SETS

Previous research argues that differences between the two mechanisms come from the heterogeneous organization of the two auctions. The FPSBA is faster, as it only requires to place simultaneous bids and the winner can be announced immediately after all bids are collected. The DA, on the other hand, requires a certain time interval for the clock to reach the desired price level of an individual bidder. Hence, a bidder in a DA faces substantial waiting costs. Carare and Rothkopf (2005) analyze the effect of transaction

[11] The first stage of the experiment was the same for all participants and participants only learned their earnings of the first part until the very end of the entire experiment, i.e., after they completed the second stage.

[12] See `www.lab.uni-koeln.de`.

[13] The first stage elicited preference parameters across gains and losses. Total net payoffs across the entire experiment range from −3.00 EUR (−3.73 USD) to 98.45 EUR (122.41 USD). The one subject who accumulated negative payoffs paid in cash at the end of experiment.

costs that accrue from the necessity to return to the auction site to check whether the desired price level has been reached. Not surprisingly, facing these additional costs, a bidder is willing to stop the auction at a higher price to avoid the need to return to the auction site.

Cox et al. (1983) and Katok and Kwasnica (2007) analyze the following trade-off experimentally. Despite the fact that bidders face transaction and/or opportunity costs from slow DA's, they also enjoy the "waiting game", as it implies a certain level of suspense. Cox et al. (1983) do not find that tripling payoffs, and therewith increasing the opportunity costs of playing the waiting game, significantly increases bids in a DA. Hence, they reject the hypothesis of "suspense utility". Katok and Kwasnica (2007) find that increasing the clock time, i.e., the time between consecutive price ticks, significantly increases bids in a DA. Slow clocks increase opportunity costs which have to be paid no matter if the bidder wins the auction or not. Katok and Kwasnica (2007) note that in the laboratory, these opportunity costs correspond most likely to participants' value of leaving the laboratory earlier. Hence, a bidder is willing to accept a higher ask to reduce the time to complete the experiment and save opportunity costs.

We account for opportunity costs in two ways. First, we hold opportunity costs constant across treatments. We follow Turocy et al. (2007) and keep the time per mechanism constant. This means that we fix the absolute time per mechanism irrespective of how fast participants decide (FPSBA) or how early they stop (DA). One round of bidding in the FPSBA always lasts 60 seconds.[14] One round of bidding in the DA always lasts 220 seconds, i.e., ten seconds per price tick (see below for a motivation). If a participant accepts a current ask, he wins the auction, but the next round does not start before the 220 seconds are over.[15] Second, all subjects play both the FPSBA and the DA.

[14]If participants do not enter a valid bid by the end of this time limit, they do not participate in the auction in that round.

[15]In both mechanisms, after the auction has ended, participants see a screen showing the remaining time until the round is completed and whether or not they have won the auction.

Katok and Kwasnica (2007) show that the clock speed has great impact on the bids in a DA due to the implied differences in opportunity costs. Because we hold opportunity costs constant, this is not an argument in our experiment. Participants in the FPSBA have 60 seconds to arrive at a bid that balances the trade-off between the winning probability and the profit in case of winning. We determine the clock speed in the DA based on two considerations. On the one hand, the trade-off between two consecutive price ticks in a DA is easier to compute and participants should need less time. On the other hand, we have to provide some time for the reference point to form. We therefore decide on a clock speed of ten seconds. This is the same clock speed as in the middle treatment in Katok and Kwasnica (2007). However, in contrast to their experiment subjects cannot reduce the duration of the DA in our experiment, as each DA lasts for 220 seconds.

In addition to controlling opportunity costs, we also hold action sets constant across the two mechanisms. In Cox et al. (1983), participants' bids are rounded to the next feasible bid in the DA. Participants can then either confirm or alter this rounded bid. In Katok and Kwasnica (2007), participants can bid integers in the FPSBA, whereas price decrements in the DA were five tokens. In contrast, in our design, participants in the FPSBA face the same set of possible prices as in the DA. This is a direct transfer of our model environment to the laboratory and ensures strict comparability between the two mechanisms.

4.4 RESULTS

In this section, we report the results of the second stage of our laboratory experiment and focus on the comparison of the FPSBA and the DA. We only consider winning bids, because we only observe a participant's bid in the DA if a participant stopped the auction and won. In order to derive a one-dimensional measure of individual bidding behavior, we first conduct OLS regressions without constants for each participant. Regressing without a constant corresponds to the assumption that a bidder with a valuation of zero behaves rational and places a bid of zero. This gives us the average slope of a subjects bidding function. The steeper the slope the more aggressive is the subject's bidding behavior. Each participant represents one independent observation, because there was no interaction between participants. We report results of non-parametric Wilcoxon signed rank (SR), Mann-Whitney-Wilcoxon (MWW), or Kruskal-Wallis (KW) tests.

In line with the observations by Coppinger et al. (1980) and Cox et al. (1982), we find that individuals place higher bids in the FPSBA than in the DA (MWW: $p = 0.0183$). However, a closer look reveals that bidders only place higher bids in the FPSBA than in the DA if they get no decision support (MWW: $p = 0.0046$). The No DSS treatment is comparable to standard experimental auction designs. If bidders get (some) decision support, the differences vanish (MWW: Medium DSS $p = 0.1498$ and Full DSS $p = 0.6256$). Table 4.4 complements these tests controlling for bidder characteristics. It confirms the observation that bids in the DA are substantially lower than in the FPSBA in absence of decision support ($p < 0.001$) and that this differences vanish once support is provided (Medium DSS $p = 0.1628$, Full DSS $p = 0.8044$).

In the FPSBA, the provision of decision support changes the bidding behavior significantly (KW: $p = 0.0704$). Bidders who receive decision support (Medium DSS, Full DSS) place lower bids than bidders without decision support (No DSS; MWW: $p = 0.0214$). In contrast, the influence of decision support is overall not significant in the DA (KW: $p = 0.1224$). However, we find some evidence that the effect of decision support works in the opposite direction compared to the FPSBA, i.e., bidders who only receive limited decision support (No DSS, Medium DSS) place smaller bids than those bidders who get full decision support (Full DSS; MWW: $p = 0.0424$).

Figure 4.1 illustrates the bidding behavior and Table 4.2 presents the results of Tobit panel regressions analyzing the influence of elicited preferences and of decision support in the FPSBA and the DA. Controlling for individual characteristics, the regressions support the results of our non-parametric tests. The provision of decision support (Medium DSS, Full DSS) decreases bids in the FPSBA. In contrast to that, in the DA the provision of Medium DSS does not influence bidding behavior ($p = 0.679$) and the influence of Full DSS is also not significant ($p = 0.106$).

The regressions further show that risk-averse bidders place higher bids, which is in line with other experimental studies (See for example Bichler et al., 2015). Our measures of individual loss aversion and Allais-type preferences have no or only marginal influence on bidding behavior. Theories based on Allais-type preferences predict higher bids in the DA than in the FPSBA, something we do not observe. In the DA we find some indication that subjects with a higher numeracy score place lower bids. However, the significance vanishes if we do not control for risk aversion.

Table 4.3 complements Table 4.2 and examines if the elicited preferences (risk aversion, loss aversion, Allais-tpye preferences) and characteristics (numeracy) have different effects on bidding behavior in the two auction formats. We only find weak evidence that a higher numeracy score leads ceteris paribus to lower bids in the DA than in the FPSBA, but no indication that any of the elicited preferences can explain differences in bidding behavior. Cox et al. (1983) argue that differences between the two mechanisms result from violations of Bayes' rule and indirectly test this conjecture by tripling individual payoffs which increases opportunity costs from miscalculations. In contrast, our design is a direct test of the impact of cognitive limitations and we find additional evidence for this conjecture.

Similar to the other experimental papers that compare bidding behavior in the FPSBA to bidding behavior in the DA (Cox et al., 1983; Katok and Kwasnica, 2007), participants in our experiment first played 18 rounds in the DA and then another 18 rounds in the FPSBA.[16] In contrast to the findings of Cox et al. (1983); Katok and Kwasnica (2007), we find that neither subjects who first participate in the FPSBA nor subjects who start in the DA change their bidding behavior when the auction format changes (SR: FPSBA $\rightarrow$ DA, $p = 0.3888$; DA $\rightarrow$ FPSBA, $p = 0.1973$). This within-participant consistency is in contrast to the literature and we relate this finding to the strict comparability of the two formats in our experiment.

[16]In order to control for order effects, about half of the participants played in reverse order.

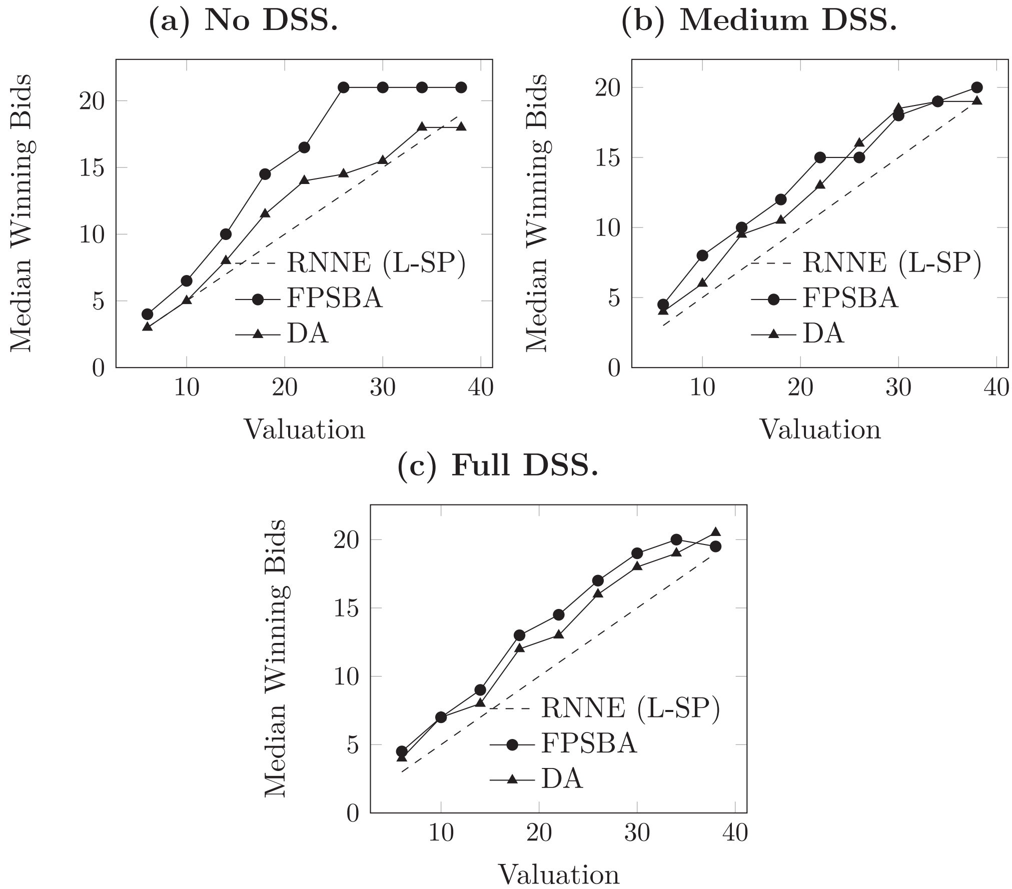

Notes: Depicted are medians of the winning bids for each valuation and format separated by decision support. The reference line is the risk-neutral Nash equilibrium (RNNE) given by Linear SP (L-SP). Participants in No DSS do not receive additional information. In treatment Medium DSS, participants receive information about the winning probability (FPSBA) or the probability to receive the next price (DA). In treatment Full DSS, participants receive the same information as in Medium DSS and, in addition, the expected profit associated with their bid.

Figure 4.1: Median winning bids across decision support.

Hence, our bidding data indicates that a constant action set and fixed opportunity costs are necessary for consistency between the two formats.[17] The other cited experiments that also vary the order of the two formats do not find a similar consistency in bidding even in absence of decision support. We think that the consistency in our data stems from the direct compa-

[17]Opportunity costs include, e.g., monitoring costs (Carare and Rothkopf, 2005) or costs from participating in the experiment (Katok and Kwasnica, 2007).

Table 4.2: Tobit panel regressions of the influence of preferences on winning bids in periods 1 to 18.

	Winning bid			
	(1)	(2)	(3)	(4)
	FPSBA	DA	FPSBA	DA
Valuation	0.523***	0.479***	0.524***	0.479***
	(0.0134)	(0.0154)	(0.0134)	(0.0154)
Allais-type	-0.0222	0.00507	-0.0156	0.00747
	(0.0374)	(0.0289)	(0.0347)	(0.0275)
Risk aversion	6.101**	10.02***	6.202**	9.533***
	(2.974)	(2.983)	(2.754)	(2.928)
Loss aversion	0.488	-0.468	0.441	-0.536
	(0.508)	(0.745)	(0.475)	(0.709)
Numeracy	0.200	-0.469**	0.115	-0.472**
	(0.300)	(0.233)	(0.280)	(0.222)
midDSS			-2.007**	-0.329
			(0.786)	(0.795)
fullDSS			-1.670**	1.218
			(0.812)	(0.753)
Period	0.0865***	0.0336	0.0859***	0.0338
	(0.0237)	(0.0267)	(0.0237)	(0.0267)
Constant	-2.553	-0.00274	-1.005	0.00709
	(2.652)	(2.199)	(2.520)	(2.096)
Observations	443	448	443	448
Participants	41	41	41	41

Standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Notes: Reported are results of tobit panel regressions with an upper limit at the highest possible bid of 21.

rability of the two formats in our design by using the same price grid and holding opportunity costs constant. Only bidders in the No DSS treatment who start bidding in the FPSBA change their bidding behavior and place lower bids when the auction format changes to a DA (SR: $p = 0.0995$). This observation might indicate that, in absence of decision support, the FPSBA is more complex than the DA.

4.5 CONCLUSION

We examine the role of decision support and preferences in first-price sealed-bid and Dutch auctions. In a laboratory experiment, we elicit participants' preferences and vary the degree of decision support to account for the complexity in deriving the optimal bid. We confirm the frequently observed non-equivalence of the first-price and Dutch auction under the absence of decision support. In addition, we observe that any differences in bidding behavior between the two mechanisms vanish once we provide decision support, which indicates that differences in bidding behavior are due to different levels of complexity. Differences between the two auction formats based on preferences should be independent of the level of decision support. We use the elicited individual preferences of all participants to explain bidding behavior. We find no indication that non-standard preferences explain the empirical differences. Our results thus indicate that the empirical breakdown of equivalence is primarily caused by the complexity of the bidding decision rather than by bidders' preferences. This observation should be taken into account in real-world business interactions involving auctions.

In the experiment, the implemented DSS is perfect in the sense that we can precisely calculate the respective probabilities and expected values due to the fixed bidding strategy of a bidding robot. Obviously, this is not directly implementable in real auctions. However, the availability of historical bid data promotes the design of decision support systems similar to our implementation. Thus, our findings on the differences in auction formats indicate that the higher revenue in the FPSBA is less relevant in real auctions in which bidders are likely to have such support.

4.6 APPENDIX

4.6.1 TABLES

Table 4.3: Tobit panel regression of the influence of preferences and numeracy on differences between winning bids in the FPBSA and the DA in periods 1 to 18.

	Winning Bid
Valuation	0.501***
	(0.0103)
Period	0.0577***
	(0.0181)
Constant	-1.730
	(2.548)
Risk aversion	6.053**
	(2.873)
Loss aversion	0.490
	(0.490)
Allais-type	-0.0225
	(0.0360)
Numeracy	0.216
	(0.290)
DA	0.843
	(3.358)
DA × Risk aversion	3.914
	(4.217)
DA × Loss aversion	-0.937
	(0.913)
DA × Allais	0.0280
	(0.0469)
DA × Numeracy	-0.672*
	(0.377)
Observations	891
Participants	82

Standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Notes: The upper limit in the Tobit regression is the maximum bid of 21. It was placed in 174 out of 891 observations. DA is a dummy variable that is zero if the auction format is a FPSBA and is one in case of a DA.

Table 4.4: Tobit panel regression of the influence of decision support in the FPSBA and the DA in periods 1 to 18.

	Winning bid
Valuation	0.502***
	(0.0103)
Period	0.0581***
	(0.0181)
Constant	1.417
	(1.682)
Allais-type	-0.00370
	(0.0220)
Risk aversion	6.498***
	(1.973)
Loss aversion	0.169
	(0.388)
Numeracy	-0.159
	(0.172)
midDSS	-2.245***
	(0.786)
fullDSS	-1.729**
	(0.815)
DA	-3.251***
	(0.794)
DA × midDSS	2.088*
	(1.131)
DA × fullDSS	3.041***
	(1.130)
Observations	891
Participants	82

Standard errors in parentheses

* $p < .10$, ** $p < .05$, *** $p < .01$

Notes: The upper limit in the Tobit regression is the maximum bid of 21. It was placed in 174 out of 891 observations. DA is a dummy variable that is zero if the auction format is a FPSBA and is one in case of a DA.

Table 4.5: Average winning bids for periods 1 to 18.

Valuation	No DSS			Medium DSS			Full DSS			KW test	
	FPSBA	DA	p-value	FPSBA	DA	p-value	FPSBA	DA	p-value	p-value FPSBA	p-value DA
6	4.25	7.42	0.8710	4.25	4.00	0.5541	4.25	3.67	0.4450	0.9905	0.9191
10	6.67	6.00	0.3417	7.31	5.93	0.1234	7.38	6.57	0.6310	0.5114	0.6148
14	10.20	8.35	0.0397	10.38	10.50	0.9575	8.67	8.55	0.7575	0.1396	0.1450
18	14.39	11.04	0.0042	12.57	10.91	0.1105	11.25	11.80	0.1498	0.2347	0.7103
22	15.29	12.05	0.0740	14.54	13.83	0.5108	14.67	13.42	0.2008	0.7910	0.6215
26	18.88	14.50	0.0022	15.18	15.71	0.6428	17.09	17.15	0.8142	0.0150	0.0878
30	19.71	16.14	0.0019	17.96	15.73	0.3084	18.00	18.20	0.786	0.0189	0.1328
34	20.20	17.65	0.0062	18.68	17.04	0.1268	18.83	19.17	0.6750	0.0219	0.1285
38	20.20	17.86	0.0190	19.35	18.42	0.4404	18.50	19.77	0.1287	0.0265	0.1372
Average	15.87	13.33	-	14.57	13.86	-	14.93	15.01	-	-	-

Notes: Reported are the average winning bids for periods 1 to 18 and the probability that bids in the different formats are drawn from the same distribution based on the Wilcoxon-Mann-Whitney U-test. The Kruskal-Wallis (KW) test reports whether there is any significant difference across decision support systems for a given auction format.

4.6.2 THEORY

We consider a situation in which the bidder faces one (non-strategic) competitor either in a FPSBA or in a DA. Let $P = \{p_1, p_2, \cdots, p_n\}$ be the common price grid, i.e. the set of possible bids in the FPSBA and the set of possible offers in the DA. Let p_k denote the kth- smallest possible price in this price grid. Let the price grid be uniformly spaced, with $p_k - p_{k-1} = \delta$ for all k.

The probability that the competitor places a bid smaller or equal p_k in the FPSBA is given by $F(p_k)$. $F(p_k)$ also denotes the probability that the highest price offer the competitor is going to accept in a DA is smaller or equal p_k.

For large η and λ the utility of a bidder is mainly driven by the relative outcomes, i.e. by his gain loss utility, and not by absolute outcomes. Consequently, it may be the case that a bidder who has a strictly positive chance of making strictly positive profits and faces no risk of a loss prefers not to participate in the auction. In the following we assume that bidder's expected utility is increasing in his valuation, which rules out such implausible predictions and guarantees monotone bidding functions. This assumption is referred to as *no dominance of gain-loss utility* in Herweg et al. (2010).

First-Price Sealed-Bid Auction

In the FPSBA both participants place a bid $b_i \in P$ and the participant who places the higher bid wins. In case of a tie both participants have a winning probability of one half. The expected profit of a bidder with valuation v bidding b_k is given by

$$\Pi(b_k, v) = \left[F(b_{k-1}) + \frac{F(b_k) - F(b_{k-1})}{2}\right] \cdot (v - b_k) \tag{162}$$

$$= \frac{F(b_k) + F(b_{k-1})}{2} \cdot (v - b_k) \tag{163}$$

$$=: P_\omega^k \cdot (v - b_k). \tag{164}$$

When relative outcomes are evaluated as

$$\mu(x) := \begin{cases} \eta x & x \geq 0 \\ \eta\lambda x & x < 0, \end{cases} \tag{165}$$

the expected utility of a bidder with KR preferences bidding b_k is given by

$$\begin{aligned} U(b_k, v) \quad = \quad & P_\omega^k \cdot (v - b_k) \\ & + P_\omega^k \cdot (1 - P_\omega^k) \cdot \mu(v - b_k) \\ & + P_\omega^k \cdot (1 - P_\omega^k) \cdot \mu(b_k - v) \end{aligned} \tag{166}$$

and optimal bids are given by

$$b_{FP}^*(v) = \arg\max_{b \in P} \{U(b, v)\}. \tag{167}$$

As the price grid starts at 0, bidders can always place bids smaller their valuation. For this reason the relevant part of the piece-wise defined utility function is given by

$$U(b_k, v) = P_\omega^k \cdot (v - b_k) - P_\omega^k \cdot (1 - P_\omega^k) \cdot (v - b_k) \cdot \eta(\lambda - 1). \tag{168}$$

Let v_k be the valuation for which a bidder is indifferent between bidding p_k and p_{k+1}. Given that these v_k are increasing in k the optimal bidding strategy $\beta_{FP}(v)$ is monotone and it is optimal for bidders to bid p_k for all bidders with a valuation between v_{k-1} and v_k. These indifference values are given by

$$U(b_k, v_k) \stackrel{!}{=} U(b_{k+1}, v_k) \tag{169}$$

$$\Leftrightarrow v_k = b_k + \delta \frac{\overbrace{P_\omega^{k+1} - P_\omega^{k+1}(1 - P_\omega^{k+1})\eta(\lambda - 1)}^{:=\Omega_{k+1}}}{\underbrace{P_\omega^{k+1} - P_\omega^k - \eta(\lambda - 1)\left(P_\omega^{k+1}(1 - P_\omega^{k+1}) - P_\omega^k(1 - P_\omega^k)\right)}_{:=\Lambda_k = \Omega_{k+1} - \Omega_k}} \tag{170}$$

The *no dominance of gain-loss utility* assumption implies a restriction on values for η and λ:

$$\begin{aligned} & \frac{\partial U(b,v)}{\partial v} = P_\omega^k - P_\omega^k \cdot (1-P_\omega^k)\eta(\lambda-1) \overset{!}{\geq} 0 \\ \Leftrightarrow \quad & \eta(\lambda-1) \overset{!}{\leq} \min_{k\in\{1,\dots,n\}} \left\{ \frac{1}{1-P_\omega^k} \right\} \\ \Leftrightarrow \quad & \eta(\lambda-1) \overset{!}{\leq} \frac{1}{1-P_\omega^1}, \end{aligned} \tag{171}$$

(171) implies that $\Omega_k \geq 0$ and $\Lambda_k \geq 0$ for all k and we get

$$\begin{aligned} v_k - v_{k-1} &= \overbrace{b_k - b_{k-1}}^{=\delta} + \delta \left[\frac{\Omega_{k+1}}{\Lambda_k} - \frac{\Omega_k}{\Lambda_{k-1}} \right] \\ &= \frac{\delta}{\Lambda_k \Lambda_{k-1}} \left[\Lambda_k \Lambda_{k-1} + \Omega_k \Lambda_k - \Omega_{k+1} \Lambda_{k-1} \right] \\ &= \frac{\delta}{\Lambda_k \Lambda_{k-1}} \left[\Lambda_{k-1}(\lambda_k - \Omega_{k+1}) + \Omega_k \Lambda_k \right] \\ &= \frac{\delta \Omega_k}{\Lambda_k \Lambda_{k-1}} \left[\Lambda_k - \Lambda_{k-1} \right] > 0. \end{aligned}$$

The bidding strategy is then given by

$$\beta_{FP}(v) = \begin{cases} 0 & \text{if } v \in [0, v_1] \\ b_k & \text{if } v \in (v_k, v_{k+1}], \end{cases} \tag{172}$$

with $v_{n+1} = 1$ if $v_k \leq 1$. Else if $v_k > 1$ for any k, β_{FP} is adjusted accordingly.

Dutch Auction

In the DA participants sequentially receive decreasing offers $a_j \in P$ starting with p_n. A participant who receives an offer can either accept or reject it. In case of acceptance the auction ends immediately. If the participant who receives the offer p_k first rejects, the other participant will also receive the offer p_k. If the other participant rejects p_k, too, the new offer will be p_{k-1}. Which participant receives the offer p_{k-1} first is randomly determined. This modeling approach is also used by Bose and Daripa (2009).

Every time the bidder receives an offer he has the choice between accepting or waiting for a lower offer. Let H_k be the probability that the bidder will receive an offer p_{k-1} given that he rejects offer p_k. The probability H_k can be split in two parts. First, ρ_k denotes the probability that the price step p_{k-1} is reached, i.e. the probability that the good is not sold at p_k. Second, ϕ_k denotes the probability that the bidder receives an offer p_{k-1} given that the price step p_{k-1} is reached. Consequently, $H_k = \rho_k \cdot \phi_k$.

COMPUTATION OF ρ_K In order to derive the probability ρ_k of reaching the next price step p_{k-1} we first determine how likely it is that the bidder receives the first offer at p_k given that he receives an offer p_k. First, denote by $\#_k^i \in \{1, 2\}$ the position of the bidder in period k. Second, denote by A_k the event that the bidder receives the offer p_k.

$$\Pr\{\#_k = 1|A_k\} = \frac{\Pr\{\#_k = 1\} \cdot \Pr\{A_k|\#_k = 1\}}{\Pr\{\#_k = 1\} \cdot \Pr\{A_k|\#_k = 1\} + \Pr\{\#_k = 2\} \cdot \Pr\{A_k|\#_k = 2\}} \tag{173}$$

$$= \frac{\frac{1}{2}}{\frac{1}{2} + \frac{1}{2} \cdot \frac{F(p_k)}{F(p_{k+1})}} \tag{174}$$

$$= \frac{F(p_{k+1})}{F(p_{k+1}) + F(p_k)}. \tag{175}$$

Consequently, the probability that the bidder is asked second at p_k given that he is asked at p_k is given by

$$\Pr\{\#_k = 2|A_k\} = 1 - \Pr\{\#_k = 1|A_k\} \tag{176}$$

$$= \frac{F(p_k)}{F(p_{k+1}) + F(p_k)}. \tag{177}$$

Given that the bidder is asked second, $\#_k = 2|A_k$, his rejection of the offer p_k directly implies that the price step p_{k-1} is reached. However, if the bidder is asked first, $\#_k = 1|A_k$, his rejection only implies that the price step p_{k-1} is reached if the competitor also rejects p_k given that she already rejected p_{k+1}, which happens with probability $F(p_k)/F(p_{k+1})$. Hence, the probability ρ_k that price step p_{k-1} will be reached given that the bidder rejects the offer p_k is given by

$$\rho_k = \Pr\{\#_k = 2|A_k\} \cdot 1 + \Pr\{\#_k = 1|A_k\} \cdot \frac{F(p_k)}{F(p_{k+1})} \tag{178}$$
$$= \frac{2 \cdot F(p_k)}{F(p_{k+1}) + F(p_k)}. \tag{179}$$

COMPUTATION OF ϕ_K Given that the price step p_{k-1} is reached the probability of being asked first is one half. In this case the bidder receives an offer with certainty. If the opponent is asked first, which also happens with a probability of one half, the bidder receives the item only if the competitor refuses the offer p_{k-1}. The probability that the competitor refuses the offer p_{k-1} given that she refused p_k is given by $F(p_{k-1})/F(p_k)$. Hence, the probability of receiving an offer p_{k-1} given that price step p_{k-1} is reached is given by

$$\phi_k = \frac{1}{2} + \frac{1}{2} \cdot \frac{F(p_{k-1})}{F(p_k)}. \tag{180}$$

COMPUTATION OF H_K Combining the probability ρ_k of reaching the next price step p_{k-1} with the probability ϕ_k of receiving an offer given that the price step p_{k-1} is reached, gives us the probability H_k of receiving another offer when rejecting p_k.

$$H_k = \rho_k \cdot \phi_k \tag{181}$$
$$= \frac{F(p_k) + F(p_{k-1})}{F(p_k) + F(p_{k+1})}. \tag{182}$$

BIDDING Let $\mathrm{R}\left(p_j|p_k\right)$ denote the probability that the bidder will be receive (or has received) an offer p_j given that he is currently offered p_k,

$$\mathrm{R}\left(p_j|p_k\right) := \begin{cases} \frac{F(p_j)+F(p_{j+1})}{F(p_k)+F(p_{k+1})} & j \leq k \\ 1 & j > k. \end{cases} \tag{183}$$

Note that for some $a < b < c$,

$$\mathrm{R}\left(a|b\right)\mathrm{R}\left(b|c\right) = \mathrm{R}\left(a|c\right).$$

The expected profit of a bidder with valuation v planning to accept offer p_j who is currently offered $p_k \geq p_j$ is given by

$$\Pi(p_j, v|p_k) = \mathrm{R}\,(p_j|p_k) \cdot (v - p_j). \tag{184}$$

A bidder with KR preferences conceives a plan at the beginning of the auction, namely accepting the offer $r \in \{p_1, ..., p_m\}$ and evaluates his profit compared to a reference outcome determined by his plan. The utility of such a bidder with valuation v who planned to accept offer r from accepting the current offer p_k is given by

$$u_k = v - p_k + (1 - \mathrm{R}\,(r|p_k)) \cdot \mu\,(v - p_k) + \mathrm{R}\,(r|p_k) \cdot \mu\,(r - p_k)\,. \tag{185}$$

Defining

$$u(x, r|y) = v - x + (1 - \mathrm{R}\,(r|y)) \cdot \mu\,(v - x) + \mathrm{R}\,(r|y) \cdot \mu\,(r - x)\,, \tag{186}$$

We now analyze two cases:

1. $p_j < r < p_k$***:***

Then, the expected utility from waiting for an offer p_j is given by,

$$U(p_j, v, r|p_k) = R(r|p_k)\Big[(1 - \mathrm{R}(p_j|r))[\mu(r - v)] + R(p_j|r)\,[v - p_j + \mu(r - p_j)]\Big]. \tag{187}$$

2. $r < p_j < p_k$***:***

Then, the expected utility from waiting for an offer p_j is given by,

$$U(p_j, v, r|p_k) = R(p_j|p_k)\Big[(1 - \mathrm{R}(r|p_j))[v - p_j + \mu(v - p_j)] + R(r|p_j)\,[v - p_j + \mu(r - p_j)]\Big]. \tag{188}$$

The bidder prefers to accept now over waiting if and only if

$$u_{k,r} \geq \max_{p_j < p_k} \{U(p_j, v, r|p_k)\} \tag{189}$$

Determining the indifference values $v_{k,r}$ gives us the bidding function,

$$\beta_r(v) = \begin{cases} 0 & \text{if} \quad v \in [0, v_1] \\ p_k & \text{if} \quad v \in (v_k, v_{k+1}], \end{cases} \tag{190}$$

with $v_{m+1} = 1$.

These strategies define best responses to the distribution of competitor's bids $F(x)$. It is easy to see that bidding strategies depend on the reference point r, i.e. the bidders plan when to accept an offer. As a consequence multiple personal equilibria are possible.

First-Price Sealed-Bid Auction vs. Dutch Auction

For subjects with KR preferences it is not possible to make a general statement about the revenue ranking of the FPSBA and the Dutch auction. In the following we provide examples that prove this statement.

Figure 4.2: Equilibrium bids in Dutch auctions and FPSBA

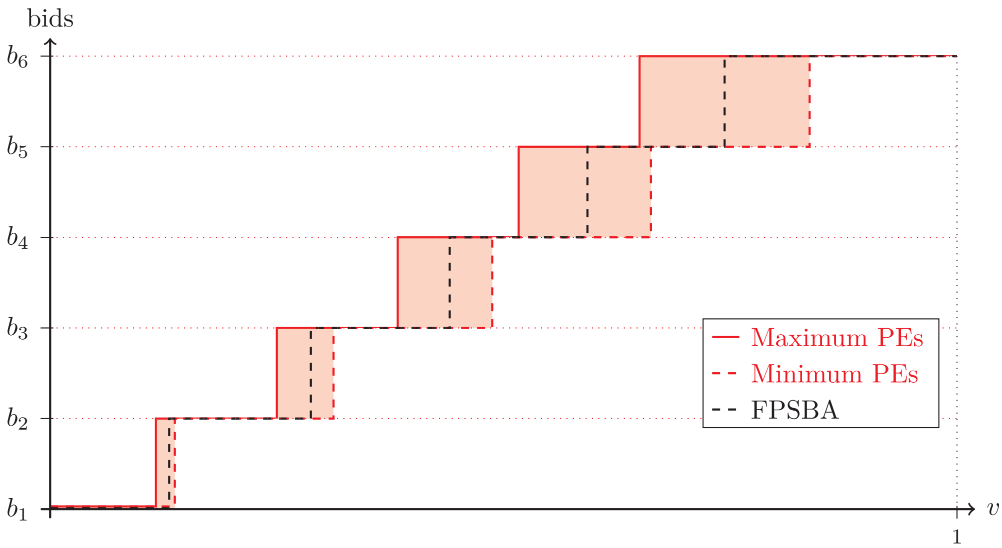

Notes: This figure shows the lowest and the highest personal equilibrium bids in the DA and the unique equilibrium bidding strategy in the FPSBA for $\lambda = 2.5$ and $\eta = 0.5$. The revenue ranking of the two auction format depends on the equilibrium selection in the DA.

4.6.3 INSTRUCTIONS

This section provides the instruction in German (original) and English (translated) separated by parts 1 and 2. Each part consists of part A and part B. Part B was always distributed after part A had been conducted. Experiment 1 was identical for each participant. Experiment 2 was counterbalanced, i.e., half of the participants received the first-price sealed-bid auction in part A followed by the Dutch auction in part B. The other half faced the reversed order. We present the instructions for the full-DSS treatment where subjects had full information. The instructions for the other treatments are the same and only exclude parts of the decision support which is reported in parentheses within the instructions.

Übersicht

Dieser Teil des Experiments besteht aus 18 Runden, die jeweils die gleiche Abfolge an Entscheidungen haben. Am Ende wird eine der 18 Runden zufällig durch den Computer ausgewählt und ausgezahlt. Alle Runden haben dabei die gleiche Wahrscheinlichkeit ausgewählt zu werden.

Erstpreisauktion

Sie nehmen an einer Erstpreisauktion teil, in der Sie ein Produkt erwerben können. Zu Beginn jeder Runde erfahren Sie, welchen Wert das Produkt für Sie hat. Dieser Wert wird aus der Menge

{ 6 €, 10 €, 14 €, 18 €, 22 €, 26 €, 30 €, 34 €, 38 € }

gezogen. Jeder Wert kommt genau zweimal vor. Die Reihenfolge ist jedoch zufällig bestimmt.

Sie befinden sich in einer Gruppe mit einem anderen Bieter. Der andere Bieter ist ein Bietroboter.

In der Auktion kann ein ganzzahliges Gebot zwischen 0 € und 21 € abgegeben werden. Der andere Bieter wählt sein Gebot zufällig zwischen 0 € und 21 €. Jedes Gebot ist dabei gleich wahrscheinlich.

Der Bieter, der das höchste Gebot abgegeben hat, gewinnt die Auktion und erhält das Produkt. Der Preis des Produkts entspricht diesem höchsten Gebot. Falls Sie und der andere Bieter das gleiche Gebot abgeben, erhalten Sie das Produkt mit 50% Wahrscheinlichkeit.

Falls Sie die Auktion gewinnen, ist Ihr Gewinn gegeben durch:

Gewinn = Wert – Gebot.

Falls Sie die Auktion nicht gewinnen, beträgt Ihr Gewinn 0.

Entscheidungshilfe

Bevor Sie Ihr echtes Gebot eingeben, können Sie verschiedene Gebote testen, wofür Ihnen ein Testbereich zur Verfügung steht.

Im Testbereich sehen Sie:

[Treatments: No DSS, Medium DSS, Full DSS]

- **Gewinn, falls Gebot erfolgreich**
  Der Gewinn, falls das aktuelle Testgebot erfolgreich wäre. Dieser wird wie folgt berechnet:
  Gewinn = Wert – Gebot.

[Treatments: Medium DSS, Full DSS]

- **Gewinnwahrscheinlichkeit**
  Die Wahrscheinlichkeit, dass Sie mit einem Gebot in Höhe des Testgebots die Auktion gewinnen.

[Treatments: Full DSS]

- **Erwarteter Gewinn**
  Durchschnittlicher Gewinn, den Sie mit dem Gebot erwarten können. Dieser wird wie folgt berechnet:
  Erwarteter Gewinn = (Gewinnwahrscheinlichkeit) x (Gewinn, falls Gebot erfolgreich).

Gebotsabgabe

- Um Ihr finales Gebot abzugeben, tippen Sie eine Zahl aus der erlaubten Menge der Gebote in das vorgesehene Feld ein. Anschließend klicken Sie auf „Gebot abgeben".
- Sie haben in jeder Runde 60 Sekunden Zeit, Ihr finales Gebot abzugeben. Sollten Sie kein Gebot in den 60 Sekunden abgeben haben, nehmen Sie in dieser Runde nicht an der Auktion teil.

Hinweis

Eine Runde dauert immer 60 Sekunden, unabhängig davon zu welchem Zeitpunkt Sie Ihr Gebot abgegeben haben. Nachdem Sie und der andere Bieter ein finales Gebot abgegeben haben, ist die Auktion zwar beendet, aber die Runde endet erst, wenn die 60 Sekunden abgelaufen sind.

Ergebnis

Übersicht

Dieser Teil des Experiments besteht aus 18 Runden, die jeweils die gleiche Abfolge an Entscheidungen haben. Am Ende wird eine der 18 Runden zufällig durch den Computer ausgewählt und ausgezahlt. Alle Runden haben dabei die gleiche Wahrscheinlichkeit ausgewählt zu werden.

Tickerauktion

Sie nehmen an einer Tickerauktion teil, in der Sie ein Produkt erwerben können. Zu Beginn jeder Runde erfahren Sie, welchen Wert das Produkt für Sie hat. Dieser Wert wird aus der Menge

{ 6 €, 10 €, 14 €, 18 €, 22 €, 26 €, 30 €, 34 €, 38 € }

gezogen. Jeder Wert kommt genau zweimal vor. Die Reihenfolge ist jedoch zufällig bestimmt.

Sie befinden sich in einer Gruppe mit einem anderen Bieter. Der andere Bieter ist ein Bietroboter.

In der Auktion startet der Preis bei 21 € und wird alle 10 Sekunden um 1 € gesenkt. Bei jedem neuen Preis wird zufällig einer der Bieter zuerst gefragt, ob er diesen Preis annehmen möchte. Nimmt der gefragte Bieter den Preis an, so endet damit die Auktion. Lehnt der gefragte Bieter ab, so wird der gleiche Preis dem verbleibenden Bieter angeboten. Beide Bieter haben die gleiche Wahrscheinlichkeit zuerst gefragt zu werden.

Der andere Bieter wählt zufällig einen Preis zwischen 0 € und 21 € aus, zu dem er annehmen würde. Jeder mögliche Preis hat dabei die gleiche Wahrscheinlichkeit ausgewählt zu werden.

Sie gewinnen die Auktion und erhalten das Produkt, falls Sie vor dem anderen Bieter einen Preis annehmen.

Falls Sie die Auktion gewinnen, ist Ihr Gewinn gegeben durch:

Gewinn = Wert – Preis.

Falls Sie die Auktion nicht gewinnen, beträgt Ihr Gewinn 0.

Entscheidungshilfe

Sie sehen auf dem Bildschirm den aktuellen Preis, den nächsten Preis sowie die Zeit bis zum nächsten Preis.

Zusätzlich sehen Sie:

[Treatments: No DSS, Medium DSS, Full DSS]

- **Gewinn bei gegebenem Preis**
  Der Gewinn, falls Sie den Preis annehmen würden. Dieser wird wie folgt berechnet:
  Gewinn bei gegebenem Preis = Wert – Preis.

[Treatments: Medium DSS, Full DSS]

- **Wahrscheinlichkeit, Preis angeboten zu bekommen**
  Die Wahrscheinlichkeit, dass Sie den jeweiligen Preis annehmen können.

[Treatments: Full DSS]

- **Erwarteter Gewinn**
  Durchschnittlicher Gewinn, den Sie erwarten können, wenn Sie sich jetzt entscheiden den jeweiligen Preis anzunehmen. Dieser wird wie folgt berechnet:
  Erwarteter Gewinn = (Wahrscheinlichkeit, Preis angeboten zu bekommen) x (Gewinn, bei gegebenem Preis).

Hinweis

Eine Runde dauert immer 220 Sekunden, unabhängig davon welchen Preis Sie annehmen. Nachdem Sie oder der andere Bieter einen Preis angenommen haben, ist die Auktion zwar beendet, aber die Runde endet erst, wenn die 220 Sekunden abgelaufen sind.

Ergebnis

Nach jeder Runde sehen Sie das Ergebnis der Runde. Hier erfahren Sie den Preis, ob Sie das Produkt erhalten haben und wie hoch Ihr Gewinn ist.

Overview

This part of the experiment consists of 18 rounds which have the same course of decisions. At the end, one of the 18 rounds will be randomly selected by the computer and paid out. All rounds have the same probability to be selected.

First-Price Auction

You will participate in a first-price auction in which you can acquire a product. At the beginning of each round, you will learn which value this product has for you. The value will be drawn from the set

{6 €, 10 €, 14 €, 18 €, 22 €, 26 €, 30 €, 34 €, 38 €}.

Each value occurs exactly twice. The order, however, is random.

You are in a group with one other bidder. This other bidder is a bidding robot.

In the auction, you can enter an integer bid between 0 € and 21 €. The other bidder will choose his bid randomly between 0 € and 21 €. Every bid is equally likely.

The bidder with the highest bid wins the auction and receives the product. The price of the product is given by this highest bid. If you and the other bidder submit the same bid, you have a 50% chance to receive the product.

If you win the auction, your profit is given by:

Profit = Value – Bid.

If you do not win the auction, your profit is 0.

Decision Support

Before you enter your actual bid, you can test different bids for which a testing area is provided for you.

In the testing area, you will see:

[Treatments: No DSS, Medium DSS, Full DSS]

- **Profit if bid was successful**
  The profit if the actual profit was successful. It is calculated as follows:
  Profit = Value – Bid.

[Treatments: Medium DSS, Full DSS]

- **Winning Probability**
  The probability that you win the auction with a bid equal to the test bid.

[Treatments: Full DSS]

- **Expected Profit**
  Average profit that you can expect with the bid. It is calculated as follows:
  Expected Profit = (Winning Probability) x (Profit if bid is successful).

Bid Submission

- To submit your final bid, type in a number out of the feasible set of bids into the respective field. Then, click on "submit bid".
- In each round, you have 60 seconds to submit your final bid. If you do not submit a bid within these 60 seconds, you will not participate in the auction in this round.

Note

One round always lasts for 60 seconds, independently of when you submit your bid. After you and the other bidder submitted a final bid, the auction end but the round will only end after the 60 seconds have elapsed.

Result

After each round, you will see the result of that round. Here you learn the price, whether or not you received the product, and how large your profit is.

Overview

This part of the experiment consists of 18 rounds which have the same course of decisions. At the end, one of the 18 rounds will be randomly selected by the computer and paid out. All rounds have the same probability to be selected.

Ticker Auction

You will participate in a ticker auction in which you can acquire a product. At the beginning of each round, you will learn which value this product has for you. The value will be drawn from the set

{6 €, 10 €, 14 €, 18 €, 22 €, 26 €, 30 €, 34 €, 38 €}.

Each value occurs exactly twice. The order, however, is random.

You are in a group with one other bidder. This other bidder is a bidding robot.

In the auction, the price starts at 21 € and will decrease by 1 € every 10 seconds. At every new price, one of the bidders is randomly asked whether or not he wants to accept the price. If the bidder accepts the price, the auction ends. If the bidder rejects the price, the same price is offered to the remaining bidder. Both bidders have the same probability to be asked first.

The other bidder will randomly choose a price a price between 0 € and 21 € which he would accept. Each feasible price has the same probability to be chosen.

You will win the auction and receive the product if you accept a price before the other bidder does.

If you win the auction, your profit is given by:

Profit = Value – Bid.

If you do not win the auction, your profit is 0.

Decision Support

On your screen, you see the current price, the next price, and the time until the next price is shown.

In addition, you will see:

[Treatments: No DSS, Medium DSS, Full DSS]

- **Profit at given price**
  The profit if you accepted the current price. It is calculated as follows:
  Profit at given price = Value – price.

[Treatments: Medium DSS, Full DSS]

- **Probability to be offered the given price**
  The probability that you can accept the respective price.

[Treatments: Full DSS]

- **Expected Profit**
  Average profit that you can expect if you decide now to accept the respective price. It is calculated as follows:
  Expected Profit = (Probability to be offered this price) x (Profit at given price).

Note

One round always lasts 220 seconds, independently of which price you accept. After you or the other bidder accepted a price, the auction ends but the round will only end after the 220 seconds have elapsed.

Result

After each round, you will see the result of that round. Here you learn the price, whether or not you received the product, and how large your profit is.

4.6.4 SCREENS IN THE LAB EXPERIMENT

Notes: Depicted is the computer interface used in the first-price sealed-bid auction. The individual valuation is depicted at the very top. Participants have a test button *Test-Gebot* (*Test bid*) that allows to enter a bid. Depending on the decision support, the following information is calculated from the test bid: *Profit falls Test-Gebot erfolgreich* (*Profit if bid was successful*) (No, Medium, and Full DSS), *Gewinnwahrscheinlichkeit* (*Winning probability*) (Medium and Full DSS), and *Erwarteter Profit* (*Expected profit*) (Full DSS). A timer displays the remaining time to submit a real bid that can be entered in the text field in the lower right corner and submitted by pressing the button *Gebot abgeben* (*Submit bid*).

Figure 4.3: Computer Interface: FPSBA.

Notes: Depicted is the computer interface used in the Dutch auction. The individual valuation is depicted at the very top. The screen shows the current price, the time until the next price, and the next price. Depending on the decision support, the following information is calculated automatically: *Gewinn bei gegebenem Preis* (*Profit at given price*) (No, Medium, and Full DSS), *Wahrscheinlichkeit, Preis angeboten zu bekommen* (*Probability to be offered the given price*) (Medium and Full DSS), and *Erwarteter Gewinn* (*Expected profit*) (Full DSS). The current price can be accepted by pressing the button *Preis annehmen* (*Accept price*).

Figure 4.4: Computer Interface: DA.

CHAPTER 5

COMMITMENT IN FIRST-PRICE AUCTIONS

Abstract

We study the role of commitment in a first-price auction environment. We devise a simple two-stage model in which bidders first submit an initial offer that the auctioneer can observe and then make a counteroffer. There is no commitment on the auctioneer's side to accept an offer as is or even to choose the lowest bidder. We compare this setting to a standard first-price auction both theoretically and experimentally. While theory suggests that the offers and the auctioneer's revenue should be higher in a standard first-price auction compared to the first-price auction with renegotiation, we cannot confirm these hypotheses in the experiment.

5.1 INTRODUCTION

The question whether to commit to clear rules when selecting the winner plays a large role in most real-life procurement processes. The multi-attribute nature of the goods or services to be procured makes a binding price-only auction a suboptimal choice. In this type of auction, the buyer cannot account for factors that she deems relevant for her awarding decision in the auction itself. From her perspective, a non-binding negotiation format where she chooses the winner after having seen all the offers might seem attractive. This non-commitment to rules on how a winner is chosen allows for flexibility when taking other, non-price attributes, into account. To support this, Jap (2002) points out that many auctions in procurement are carried out in a non-binding fashion.

This paper investigates the role of commitment in a concise setting and examines whether participants react to commitment, or a lack thereof, in a first-price auction. We compare a standard first-price auction with commitment to a first-price auction where renegotiation is possible, while varying as little as possible between the two settings. In our simple two-stage mechanism, bidders first submit an offer that the auctioneer can observe. In the second stage, the auctioneer then selects a winner and can make a counteroffer. There is no commitment on the auctioneer's side to accept an offer as is or to choose the lowest bidder.[1] In theory and considering that the auctioneer makes a counteroffer, this means that bidders pool on bids that reveal no information about their costs. This means, in equilibrium, bids are uninformative and the auctioneer implements the ex-ante optimal take-it-or-leave-it offer.

We then take these mechanisms into the laboratory where we benchmark the theoretical model of step two against a standard first-price auction. Contrary to theoretical predictions, we observe no significant difference in the offers between the setting with renegotiation and the standard first-price auction. Also, we find evidence that first-stage offers are correlated to the private information of the bidders in both settings.

There is evidence in the literature that having a binding auction, or an auction with commitment, is an important factor when designing the procurement process. The most related study was conducted by Fugger et al. (2016). They show that conducting auctions without commitment can lead to non-competitive prices. In their study, a quality component is introduced that is unknown to the auctioneer before the auction. The authors then compare two settings of a dynamic reverse auction: with and without commitment. The auctioneer conducts either a price-only auction, where the lowest bid wins, or a buyer-determined auction. In the latter, she chooses the winner after having seen all the offers and qualities. Since bidders do not know their quality ranking, they cannot be sure that a reduction in price leads to a higher winning probability. Therefore, bidders lack an incentive to submit competitive offers and collusion on high prices prevails. They show theoretically that these non-competitive offers become

[1] Even if the auctioneer did commit to choosing the lowest offer, the offers would still be uninformative.

profitable once the auctioneer does not commit to clear rules on how the winner is chosen. This theoretical finding is then confirmed via a laboratory experiment. Our study is focussed on keeping the mechanism as simple as possible to isolate the role of commitment. There exists one type of equilibria in both our settings with clear predictions, collusion is not profitable. While offers in the standard first-price auction are competitive, theory predicts that bidders pool on offers that reveal no information about their type in the first-price auction with renegotiation.

Commitment has been studied mainly in the multi-attribute literature and the optimal mechanism-design without commitment literature. Che (1993) analyzes the role of commitment in multi-attribute auctions. If the auctioneer is able to commit to a scoring rule, then the optimal scoring rule undervalues quality with respect to the auctioneer's utility. If not, the only scoring rule she can implement is given by her utility. In contrast to this paper, their perspective is to derive optimal buyer behavior in the presence and absence of commitment power. They also theoretically show the importance and benefits of commitment. We, on the other hand, focus on bidder behavior in settings with and without commitment.

In the optimal mechanism-design literature Vartiainen (2013) shows that if a sequentially rational auctioneer cannot commit to the mechanism rules, the only mechanism she can implement is a variant of the English auction. Mechanisms in which offers directly depend on a bidder's type are generally not possible. The English auction has the property that the winner of the auction does not reveal his offer (and type), while in our first-price auction, this is not the case. Also, in our paper, the auctioneer cannot choose the procurement mechanism and is bound, depending on the setting, to either a standard first-price auction or a first-price auction with renegotiation. McAfee and Vincent (1997) assume more structure. The auctioneer sets a reserve price but cannot commit to not reauction the good if the reserve-price is not met. They show that in this case, the revenue of the auction drops to the static auction without reserve price. This is related to our setting, where in the first-price auction with renegotiation, the buyer can enforce a reserve-price via take-it-or-leave-it offer if the offers do not meet her expectations. This is possible because the buyer still wields some commitment power, namely that reauctioning is not possible. Once the chosen bidder has declined the counteroffer, no deal is made.

Related to commitment in auctions is Tan (1996). The author studies a procurement setting where a buyer is privately informed about her own demand. If the buyer is able to commit a reserve price, it is always in her interest to do so. This means that she reveals her private information. In our setting, the auctioneer does not possess private information. Also, communication is only possible from the suppliers to the auctioneer in the form of offers.

The paper is organized as follows. In section 5.2, we develop the model and analyze it. In section 5.3, we describe our experimental setting and present the results.

5.2 MODEL

In this section, we introduce the formal model. We consider an auctioneer and n bidders that compete for one indivisible good in a two-stage mechanism.[2] We assume that both bidders and the auctioneer are risk-neutral profit maximizers.

Bidders' values for the good are independently and identically distributed according to a cumulative distribution function F over the set $V = \{\underline{v}, \ldots, \bar{v}\}$, $\underline{v} \geq 0$ and $V \subset \mathbb{N}^0$.[3] The auctioneer assigns zero value to the good.

In the first stage, bidders send an offer to the auctioneer. Offers are binding, the auctioneer may acquire the good for any offer that was submitted. The set of possible offers is given by $B = \mathbb{N}^0$.

In the second stage, the two settings we compare differ. In the first-price auction with renegotiation, there is no commitment on the auctioneer's side. She observes the offers and can choose the winner arbitrarily. She then

[2] We write our model as a selling rather than a procurement mechanism, since our experiment is framed as a selling auction, too. This has the advantage that we have consistency in notation throughout the paper. This is, of course, without loss of generality.

[3] The exact spacing between types and bids is not important, as long as the spacing in the bid and type spaces stays constant.

makes a counteroffer to the chosen bidder or accepts the offer as-is. The bidder can accept or decline the counteroffer. In the standard first-price auction, the auctioneer observes the offers and chooses one of them.

Utilities are identical in both settings. For the auctioneer, her utility is given by the price paid by the winner of the auction if the trade takes place. The utility of the chosen bidder with value v winning with a price of p is given by

$$u_b(v; p) = v - p. \tag{191}$$

The bidder that was not chosen has a utility of zero. If no trade takes place, the utility of everyone is zero.

5.2.1 ANALYSIS

We show that there exists a continuum of equilibria in the no-commitment setting. Each equilibrium is characterized by bidders mixing over a subset $U_o \subset B$ such that $\max\{U_o\} \leq \underline{v}$. Note that U_o may contain only one element, p_o. In that case, bidders pool on p_o. For $\underline{v} = 0$, the equilibrium is unique.

Proposition 1. *The equilibria are characterized by*

1. *Bidders: randomize over a subset $U_o \subset B$ such that $\max\{U_o\} \leq \underline{v}$ in the first stage*
2. *Auctioneer:*
   - *i) observes that all offers are $\in U_o$: she chooses a bidder at random and makes a counteroffer. The counteroffer p_{co} is equal to the ex-ante optimal take-it-or-leave-it offer, $p_{co} = \arg\max_{p \in \{\underline{v}, \ldots, \bar{v}\}} (1 - F(p))p$.*
   - *ii) observes one or multiple offers $\notin U_o$: she chooses a deviating offer and makes a counteroffer that is equal to $\bar{v}$.*

Proof. We start by showing that the proposed behavior indeed forms an equilibrium. Bids are binding, so every bidder submitting offers above his value has an incentive to deviate to a lower offer. This means for any value larger than $\underline{v}$, there is a non-zero possibility that the bidder cannot make that offer. This means bidders cannot pool on any value larger than $\underline{v}$ and it follows that bidders pool by mixing over a subset $U_o \subset B$ such that $\max\{U_o\} \leq \underline{v}$. Off-equilibrium beliefs of the auctioneer are given by $\mu(v_i = \bar{v}|o_i \notin U_o) = 1$, meaning that if she observes any signal $\notin U_o$ in the first stage, she assumes that the bidder is of the highest type. Hence, deviating always yields a revenue of zero for the bidder, he receives a counteroffer of $\bar{v}$. The intuition behind these off-equilibrium beliefs comes from how an auctioneer would eliminate possible types.

Suppose there are $n+1$ types, $V = \{0, 1, \ldots, \bar{v}\}$ and let $0 < p_{co} < \bar{v}$. If the auctioneer observes an offer of 1, this bidder must be of type $v \in \{1, \ldots, \bar{v}\}$. Bidders of type $v = 1$ can send offers of 0 or 1. But the smallest counteroffer an auctioneer could commit to after observing 1 would be 1. This means bidders of value $v = 1$ are indifferent and submit only offers of 0. Bidders of type $v = 2$ can send offers of 0, 1 or 2. Since bidders of type $v = 1$ do not send offers of 1, the smallest counteroffer an auctioneer could commit to after observing 1 or 2, would be 2. Therefore, the expected profit of a bidder of type $v = 2$ who submits an offer of 1 or 2 is 0 and $\frac{1}{2}\max\{2 - p_{co}, 0\}$ if he submits an offer of 0. It follows that, like the type-$v = 1$ bidders, bidders of type $v = 2$ will only submit offers of 0. This argument can be chained n times until only the bidder of type $\bar{v}$ is left. The best response of the auctioneer facing these uninformative offers is setting the optimal reserve price, $p_{co} = \arg \max_{p \in \{\underline{v}, \ldots, \bar{v}\}} (1 - F(p))p$.

We will now show that this is the unique type of equilibrium. Bidders are assumed to be profit maximizers. This means they will accept any counteroffer that is larger or equal than their value. Assume the bidders submit offers according to a separating equilibrium bidding function.[4] A separating equilibrium bidding function implies that the auctioneer can infer the value of each bidder from the offer. She would then make a counteroffer to the bidder with the highest and extract full surplus in the second stage equal to

[4]This equilibrium bidding function does not need to be monotone.

his value for the good. The bidder would then accept this counteroffer and make a profit of zero. This means bidders would always prefer to imitate a lower type. This rules out the existence of any separating equilibrium. The same logic can be applied over any subset of V.

It is left to show that there exist no partial pooling equilibria where bidders pool on multiple offers in $V\backslash\{\underline{v}\}$. Consider a setting with m pooling offers $s_i \in V\backslash\{\underline{v}\}$ with $i \in \{1, 2, ..., m\}$. W.l.o.g., let $v_i \in V$ be the lowest type that sends s_i. Let p_i, $i \in \{1, 2, ..., m\}$, be the respective prices an auctioneer sets after observing that the highest bid is s_i. W.l.o.g., let $p_1 < p_2 < \cdots < p_m$. The auctioneer cannot commit to any $p_i \leq v_i$ since she knows that the lowest type sending the signal s_i has a value of v_i. On the other hand, a bidder having a value of v_i will never send a signal that results in a price $p \geq v_i$ since deviating to a lower signal would earn him a strictly positive expected payoff. This is a contradiction to the assumption that v_i is the lowest type sending signal s_i, meaning that no pooling equilibrium with one or multiple offers in $V\backslash\{\underline{v}\}$ can exist. □

In the standard first-price auction, bidders send an offer to the auctioneer in the first stage. The set of possible offers is the same as before, $B = \mathbb{N}_0$. The auctioneer then observes these offers and has to choose one of the offers. She cannot make a counteroffer. This setting is equivalent to a first-price auction: A profit maximizing auctioneer will always select the highest offer. First-price auctions are well-studied in the literature, see for example Krishna (2009), for the discrete case see Chwe (1989) and Cai et al. (2010). The equilibrium bidding function for a bidder with value v bidding against $n-1$ other bidders is approximated well by the continuous equilibrium bidding function if there are a sufficient number of bid steps,

$$\beta^I(v) = \frac{1}{F_1^{(n-1)}(v)} \int_0^v y f_1^{(n-1)}(y) dy. \tag{192}$$

Proposition 2. *Bids in the standard first-price auction are higher or equal than in the first-price auction with renegotiation.*

Proposition 3. *The standard first-price auction is more efficient than the first-price auction with renegotiation.*

Proof. From Chwe (1989) and Cai et al. (2010), we know that for any type, bids are higher or equal than $\underline{v}$. This is in contrast to the first-price auction with renegotiation where every submitted offer is smaller or equal to $\underline{v}$, proving Proposition 2. In the standard first-price auction, the good is always sold in equilibrium. In the first-price auction with renegotiation, the bidder might reject the counteroffer, making this format inefficient. This proves Proposition 3. □

5.2.2 QUANTAL RESPONSE EQUILIBRIUM

The predictions for the first-price auction with renegotiation are extreme in the sense that for any deviation from the equilibrium, the auctioneer's counteroffer jumps from the optimal take-it-or-leave-it offer to $\bar{v}$. Bidders are assumed to perfectly understand that the auctioneer can infer their type in any separating bidding strategy and that their bid should not contain any information about their type. In comparison to the standard first-price auction, where small errors only lead to small changes in winning probability and expected payment, the first-price auction with renegotiation leaves no room for errors. Still, in real-life situations, bidders and the auctioneer might err due to, for example, cognitive limitations. In this section, we are interested in what happens when we relax the assumption that players' choices are always optimal and allow them to make mistakes.

One equilibrium concept choice to account for these type of deviations is the quantal response equilibrium (QRE) (McKelvey and Palfrey, 1995). In this section we model both the first-price auction with renegotiation and the standard first-price auction settings and derive the corresponding response functions. We begin with the first-price auction with renegotiation. Consider $n = 2$ bidders. Let $T = \{0, ..., 10\}$ be the set of possible types.[5] The action space of the bidders is given by $A^B = T = \{t_i\}_{i\in\{0,...,10\}}$. The

[5] We consider a reduced version with eleven types of the experiment that has 101 types. This is due to computational limitations when numerically solving the QRE.

auctioneer's action space is given by

$$A^A = \{(t_0, b^1), (t_1, b^1), \dots, (t_{10}, b^1), (t_0, b^2), \dots, (t_{10}, b^2)\} = \{a_i^A\}_{i \in \{0,\dots,21\}}, \tag{193}$$

where the first eleven entries denote a counteroffer of t_i to bidder one while the other entries denote the counteroffers to bidder two. Note that both offers and counteroffers are capped by the highest possible type. In QRE, every action of every player is chosen with a positive probability depending on the expected utility of said action and on a precision parameter, $\lambda \in [0, \infty)$. We use the logit QRE concept described in Goeree et al. (2016) in chapter 3.3.

Consider bidder 1. Let σ_{ij}^B be the probability that a bidder of type t_i submits an offer of t_j. Let σ_{ijk}^A be the probability that, given the bids of bidder one and two, $b^1 = t_i$ and $b^2 = t_j$, the auctioneer chooses the action a_k^A. The weighting function depends on the expected utilities. Then the expected utility of bidder 1 being of type t_i and submitting an offer of t_j is given by

$$U_1^B(t_i, t_j, \sigma^B, \sigma^A) = \underbrace{\sum_{k=0}^{10} \sum_{l=0}^{k} \sigma_{kl}^B}_{\text{bidder 2: } b^2} \underbrace{\sum_{m=0}^{21} \sigma_{jlm}^A}_{\text{auction.: } a_m^A} \begin{cases} t_i - t_m & a_m^A \in (\cdot, b^1) \ \& \ m \leq i \\ 0 & a_m^A \notin (\cdot, b^1) \\ 0 & a_m^A \in (\cdot, b^1) \ \& \ m > i \end{cases} \tag{194}$$

$$= \sum_{k=0}^{10} \sum_{l=0}^{k} \sigma_{kl}^B \sum_{m=0}^{i} \sigma_{jlm}^A (t_i - t_m) \tag{195}$$

$$:= U_{1,ij}^B(\sigma^B, \sigma^A). \tag{196}$$

Analogously, the expected utility of the second bidder being of type t_i and submitting an offer of t_j is given by

$$U_2^B(t_i, \sigma^B, \sigma^A) = \sum_{k=0}^{10} \sum_{l=0}^{k} \sigma_{kl}^B \sum_{m=0}^{21} \sigma_{jlm}^A \begin{cases} t_i - t_m & a_m^A \in (\cdot, b^2) \ \& \ m \leq i + 11 \\ 0 & a_m^A \notin (\cdot, b^2) \\ 0 & a_m^A \in (\cdot, b^2) \ \& \ m > i + 11 \end{cases} \tag{197}$$

$$= \sum_{k=0}^{10} \sum_{l=0}^{k} \sigma_{kl}^B \sum_{m=11}^{i+11} \sigma_{jlm}^A (t_i - t_{m-11}) \tag{198}$$

$$:= U_{2,ijk}^B(\sigma^B, \sigma^A). \tag{199}$$

The expected utility of the auctioneer having received the offers $b^1 = t_i$ and $b^2 = t_j$ and taking action a_k^A is given by

$$U_{ijk}^A(\sigma^B, \sigma^A) = \begin{cases} \sum\limits_{m=0}^{10} \sigma_{mi}^B \begin{cases} t_k & m \geq k \\ 0 & m < k \end{cases} & a_k \in (\cdot, b^1) \\ \sum\limits_{m=0}^{10} \sigma_{mj}^B \begin{cases} t_k & m + 11 \geq k \\ 0 & m + 11 < k \end{cases} & a_k \in (\cdot, b^2) \end{cases} \tag{200}$$

$$= \begin{cases} \sum\limits_{m=k}^{10} \sigma_{mi}^B t_k & a_k \in (\cdot, b^1) \\ \sum\limits_{m=k-11}^{10} \sigma_{mj}^B t_k & a_k \in (\cdot, b^2). \end{cases} \tag{201}$$

The logit QRE response function to determine the σ's in the quantal response equilibrium is generally of the form

$$\sigma_i = \frac{e^{\lambda U(\sigma_i)}}{\sum_{\sigma_j} e^{\lambda U(\sigma_j)}}. \tag{202}$$

In our case, we have the following system of equations,

$$\begin{aligned} \sigma_{ij}^B &= \frac{e^{\lambda U_{1,ij}^B(\sigma^B, \sigma^A)}}{\sum_{k=0}^{i} e^{\lambda U_{1,ik}^B(\sigma^B, \sigma^A)}} \quad \forall\, t_i, t_j \in T \\ \sigma_{ij}^B &= \frac{e^{\lambda U_{2,ij}^B(\sigma^B, \sigma^A)}}{\sum_{k=0}^{i} e^{\lambda U_{2,ik}^B(\sigma^B, \sigma^A)}} \quad \forall\, t_i, t_j \in T \\ \sigma_{ijk}^A &= \frac{e^{\lambda U_{ijk}^A(\sigma^B, \sigma^A)}}{\sum_{m=0}^{21} e^{\lambda U_{ijm}^A(\sigma^B, \sigma^A)}} \quad \forall\, t_i, t_j \in T \text{ and } \forall\, a_k^A \in A^A. \end{aligned} \tag{203}$$

We make some assumptions on the behavior of both the bidders and the auctioneer. The bidders cannot submit bids strictly higher than their type, so $\sigma^B_{ij} = 0$ for all $j > i$. The auctioneer takes this into account and forgoes strictly dominated choices when submitting the counteroffer. Therefore, she does not make counteroffers lower than the highest of offers she has received. This means $\sigma^A_{ijk} = 0$ for all $k < \max\{i, j\}$ and $11 < k < \max\{i, j\} + 11$. Additionally, we assume that the auctioneer chooses the highest of the two bidders for the counteroffer, $\sigma^A_{ijk} = 0$ for $k < 12$ if $i < j$ and $\sigma^A_{ijk} = 0$ for $k > 11$ if $i > j$. While this assumption increases the pressure on prices, it does not change the results qualitatively and makes the presentation of the results easier. This is due to the fact that the probability of an action a^A_k then depends only on the highest bid, which yields a probability matrix that is easier to interpret. The standard first-price auction is modeled analogously, see Appendix 5.5.2.

In Goeree et al. (2016) it is shown that for $\lambda \to \infty$, the QRE converges to the unique Bayes-Nash-equilibrium derived in the last section. This means QRE gives us three predictions for the behavior of bidders and auctioneer:

Proposition 4. *For the limit case $\lambda \to \infty$, bids are lower in the first-price auction with renegotiation than in the standard first-price auction.*

Proof. As shown in Goeree et al. (2016) section 3, the logit QRE converges to the unique Bayes-Nash-equilibrium derived in the last section. Then the results derived in that section apply. □

Proposition 5. *For the limit case $\lambda \to 0$, bids are identical in the first-price auction with renegotiation and the standard first-price auction. Bidders are unresponsive to expected payoffs and submit all valid offers with equal probability.*

Proof. For $\lambda \to 0$, the system (203) simplifies to

$$\begin{aligned} \sigma^B_{ij} &= \frac{1}{i+1} \quad \forall\, t_i, t_j \in T \\ \sigma^B_{ij} &= \frac{1}{i+1} \quad \forall\, t_i, t_j \in T \\ \sigma^A_{ijk} &= \frac{1}{22} \quad \forall\, t_i, t_j \in T \text{ and } \forall\, a^A_k \in A^A. \end{aligned} \tag{204}$$

□

Proposition 6. *In contrast to the unique Bayes-Nash-equilibrium derived in the last section, for any $\lambda > 0$, there is correlation between the offers submitted by the bidders and the counteroffer submitted by the auctioneer.*

Proof. For $\lambda > 0$, bidders submit all offers larger than zero and smaller or equal than their type with strictly positive probability. The auctioneer then conditions her counteroffer on the bids she received and forgoes strictly dominated actions, namely those counteroffers smaller than the highest of offers. This means that there exists a correlation between the offers and the counteroffer. □

We can numerically compute the equilibrium probability weights as described in Goeree et al. (2016) for different λ values. The results can be found in Figure 5.1 – Figure 5.3.

The QRE of the standard first-price auction can be found in Figure 5.1. As expected, for the higher λ-value, the offers are less "washed out" around the standard equilibrium bidding strategy of around $v/2$.

The QRE of the first-price auction with renegotiation can be found in Figure 5.2 and Figure 5.3. For the bidders, one can still see some pressure to pool offers in the $\lambda = 15$ case, while in the more error-prone $\lambda = 3$ case, offers start resembling those of the first-price auction. For the counteroffers, the auctioneer makes use of the information she gets from the bidders and mixes her response.

In conclusion, we might observe offers in the first-price auction with renegotiation that are closer to the offers in the standard first-price auction than standard theory would predict.

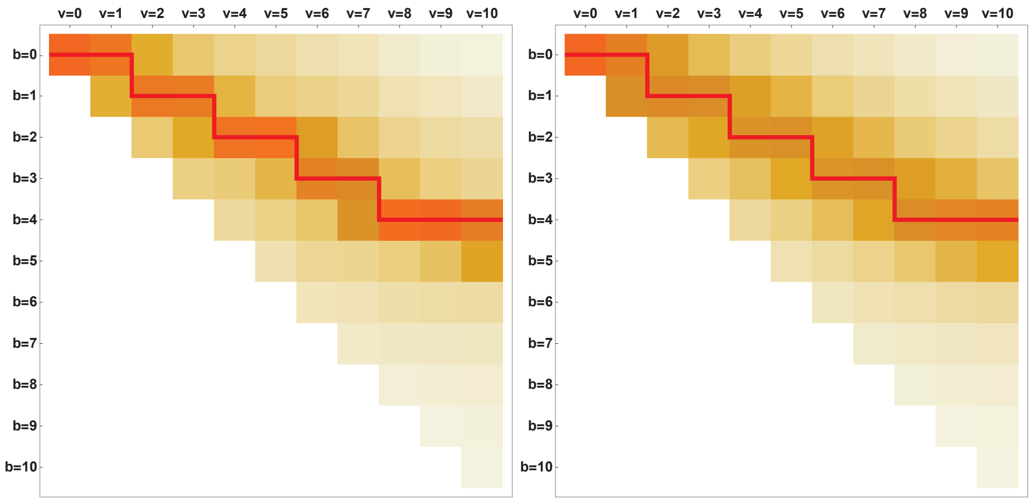

(a) $\lambda = 15$ (b) $\lambda = 3$

Figure 5.1: Numerical QRE of the standard first-price auction for two different values of λ. The rows represent the probability a certain offer is submitted for each of the types (rows). A darker shade represents a higher probability. The red line marks the offer with the highest probability for each type.

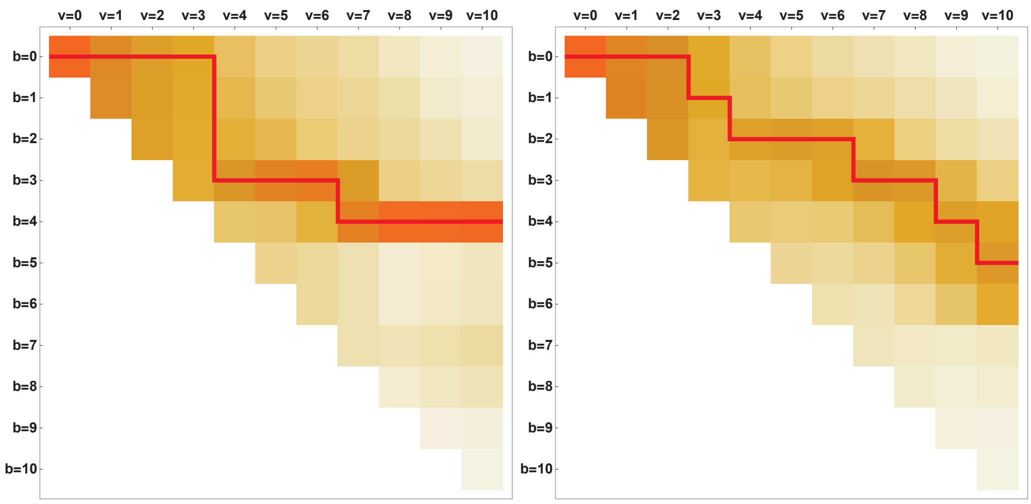

(a) $\lambda = 15$ (b) $\lambda = 3$

Figure 5.2: Numerical QRE of the first-price auction with renegotiation **for the bidders** for two different values of λ. The rows represent the probability a certain offer is submitted for each of the types (rows). A darker shade represents a higher probability. The red line marks the offer with the highest probability for each type.

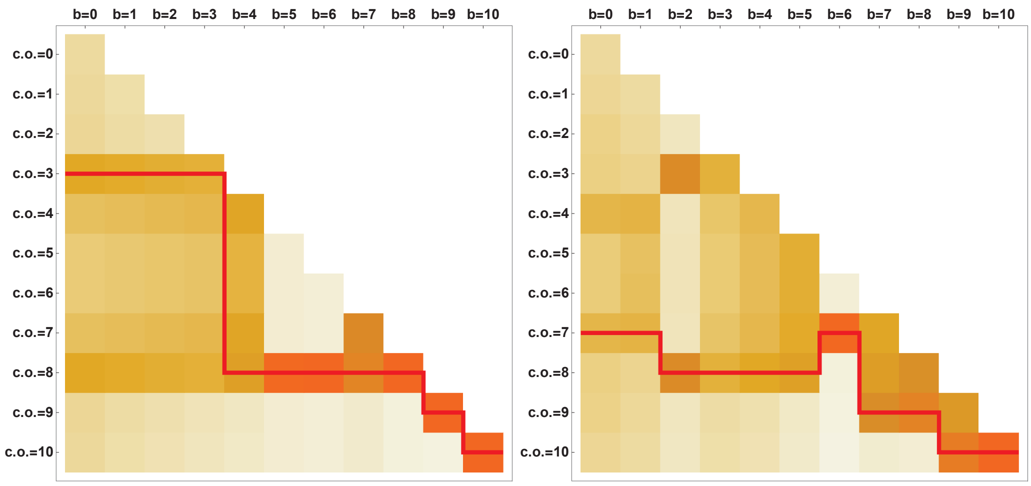

(a) $\lambda = 15$ (b) $\lambda = 3$

Figure 5.3: Numerical QRE of the first-price auction with renegotiation **for the auctioneer** for two different values of λ. The rows represent the probability a certain counteroffer is submitted after a certain highest offer (columns) was observed. A darker shade represents a higher probability. The red line marks the counteroffer with the highest probability for each highest offer.

5.3 EXPERIMENT

In this section, we introduce our experimental design and state our hypotheses for the experiment.

5.3.1 EXPERIMENTAL DESIGN

We conducted three different treatments: the standard first-price auction (FPA), the first-price auction with renegotiation (FPR) and the first-price auction with renegotiation and feedback (FPRF). In all settings, the valuations of the bidders are drawn from the set $\{0, 1, ..., 100\}$ ECU, all valuations are equally likely. In the FPA treatment, both bidders can submit offers $\in \{0, 1, ..., 100\}$ in a first stage. The auctioneer than observes these offers and chooses one of them at will. In the FPR treatment, the auctioneer can additionally make a counteroffer. The counteroffer is automatically accepted if it is below or equal to the value of the chosen bidders, and is rejected if it is higher than his value. This is done to reduce noise from an additional decision of the participants.

The FPRF treatment includes additional feedback for the auctioneer: After each round finishes, the offers and values of the two bidders are revealed to her. With standard preferences, this does not have any implications on the equilibrium bidding strategies derived in 5.2.1.

5.3.2 ORGANIZATION

The experiments were conducted in the Cologne Laboratory for Economic Research (CLER) at the University of Cologne, Germany. Using the recruiting system ORSEE (Greiner, 2015), we invited a random sample of the CLER's subject pool via email with cash as the only incentive offered. Our participants were mostly students at the University of Cologne, mostly undergraduates, from a variety of majors, and they therefore represent the larger university community. The whole experiment was computerized using the programming environment oTree (Chen et al., 2016). Upon their arrival at the laboratory, participants were seated in visually isolated cubicles and read instructions on their screens (see Appendix 5.5.1) describing the rules of the game. Following this, they were handed control questions which they had to answer correctly to proceed.

In total, 138 subjects participated in the experiment, with 36 subjects participating in the FPA treatment, 48 subjects participating in the FPR treatment and 54 subjects participating in the FPRF treatment.

Payoffs were stated in ECU, the conversation rate used was 1ECU = 0.01 EUR. Participants were paid out in private after the completion of the experiment. All 138 participants were paid their total net earnings. The average payoff for the entire experiment was 16.17 EUR corresponding to approx. 18.95 USD at the time of the payment.

Participants were randomly assigned to one of two rooms where two different treatments were conducted simultaneously. We randomly assigned one of the two roles, bidder and auctioneer, to every participant. Participants kept their assigned role for the whole experiment. Participants were grouped into cohorts of six where two auctioneers and four bidders were matched randomly in each of the 50 rounds within a cohort.

5.3.3 HYPOTHESIS

Our theory predicts that offers in the FPR and FPRF treatments are not correlated with the value of the respective bidder, they submit offers of zero in equilibrium. With this, we can state the following hypotheses:

Hypothesis 1. *There is no correlation between the value and the offers in the FPR and FPRF treatments.*

Hypothesis 2. *Offers are lower in the FPR and FPRF treatments than in the FPA treatment.*

When we compare offers between the FPR and the FPRF treatment, the difference in feedback could improve learning in the FPRF treatment. The equilibrium bidding strategy in this setting requires a certain depth of reasoning, a bidder needs to understand that any separating bidding strategy leads to full surplus extraction. The additional feedback allows the auctioneer to see how much money she “left on the table” in each round. This, in turn, should lead to higher counteroffers which should lead bidders to adjust their offers downwards. Thus, we expect that the additional feedback pushes bidders closer to the equilibrium bidding strategy. This is also related to our QRE results from Section 5.2.2: The additional feedback could lead to less errors, or a higher λ value.

Hypothesis 3. *Offers in the FPRF treatment are lower than in the FPR treatment.*

Theory predicts that counteroffers of the auctioneer do not depend on the offers received in the first stage.

Hypothesis 4. *There is no correlation between the offer of the chosen bidder and the counteroffer of the auctioneer in the FPR and FPRF treatments.*

The next hypothesis concerns the revenue of the auctioneer. In the FPA treatment, the competition between bidders helps the auctioneer while in the FPR and FPRF treatments, she can only propose the ex-ante optimal take-it-or-leave-it offer to one of the bidders. A numerical simulation confirms this intuition. We can also approximate the revenues with continuous types, since our bid grid is fine enough. For uniformly distributed values, $F = \mathcal{U}[0, 100]$, the bidding strategy simplifies to

$$\beta^I(v) = \frac{v}{2}. \tag{205}$$

The expected revenue for the first-price auction for the uniform distribution over the interval $[0, 100]$ is given by

$$\mathbb{E}\Big[R\Big] = \frac{100}{3}. \tag{206}$$

The optimal take-it-or-leave-it offer in the same setting is given by 50. Offering this to one of the bidders at random results in an expected revenue of $\frac{100}{4} = 25$.

Hypothesis 5. *The auctioneer's revenue is strictly higher in the FPA treatment than in the FPR and FPRF treatments.*

If the additional feedback of the FPRF really leads to less errors and with that to a QRE that is closer to the unique Bayes-Nash equilibrium, than revenue should be lower in the FPRF treatment.

Hypothesis 6. *The auctioneer's revenue is lower in the FPRF treatment than in the FPR treatment.*

Related to Proposition 3, the FPA should be efficient, while theory predicts that the FPR and the FPRF are not.

Hypothesis 7. *The FPA is more efficient, in the sense that the bidder with the highest value is more often the winner, than in the FPR and FPRF treatments.*

For the FPA treatment, we should observe that bidders bid according to the equilibrium bidding strategy (205).

	Mean	Std. Dev.	Min	Max
FPA				
Participants	36	–	–	–
Values	50.56	29.41	0	100
Offers	34.56	21.46	0	95
FPR				
Participants	48	–	–	–
Values	49.58	28.97	0	100
Offers	36.81	23.19	0	98
Counteroffers	53.08	18.24	1	100
FPRF				
Participants	54	–	–	–
Values	50.52	28.99	0	100
Offers	34.25	21.73	0	99
Counteroffers	53.49	18.66	1	99

Table 5.1: Summary statistics for the treatments.

Hypothesis 8. *Bidders submit offers according to the equilibrium bidding function of $\beta^I(v) = v/2$ in the FPA treatment.*

5.3.4 RESULTS

We begin with the hypotheses concerning the bidding strategy of the bidders in FPR and FPRF treatments and the comparison with the FPA treatment, Hypothesis 1 and Hypothesis 2.

As a reminder, the equilibrium offers are given by zero in these two settings. However, we observe only four out of 68 bidders who submit an offer of zero when their value is larger than five and of these, only three do so more than once. Also as can be seen in table 5.2, value has a significant influence on the offers in the FPR and FPRF treatments. Thus, we must reject Hypothesis 1.

In the FPR treatment, the average offer is given by 36.82, while in the FPRF treatment, it is given by 34.25, see table 5.1. In the FPA treatment, the average offer is given by 34.56. While the treatment dummy for the FPR treatment has a significant effect on the offers (see Table 5.3), it is

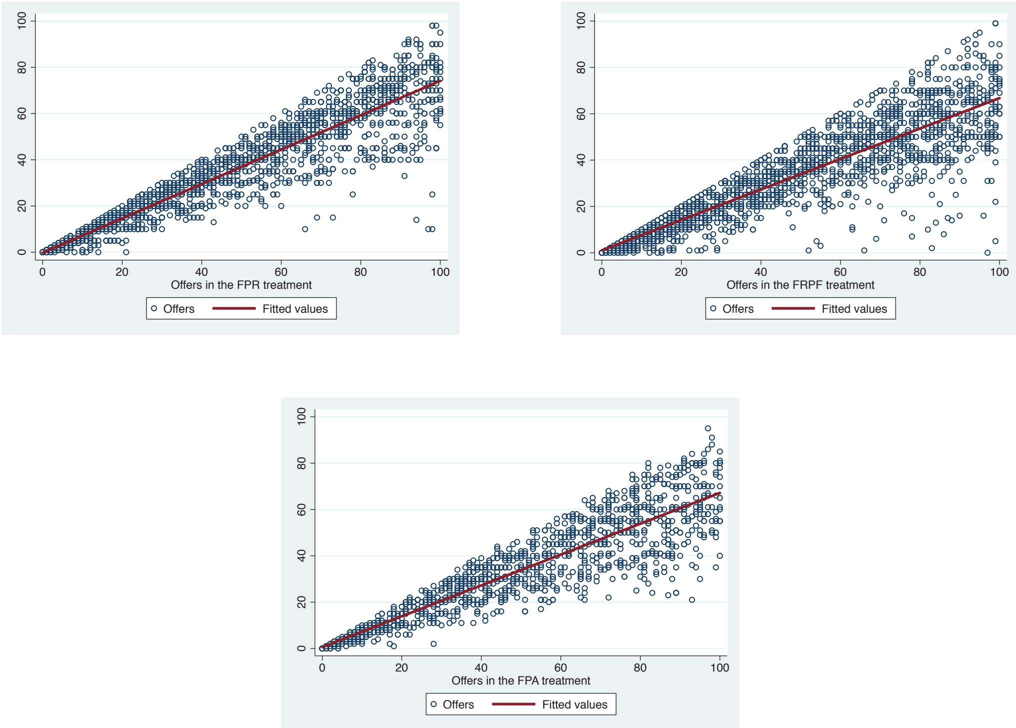

Figure 5.4: Offers and the corresponding linear regressions in the FPR (left), the FPRF (right) and the FPA treatment (below).

positive, contrary to Hypothesis 2. We find no significant difference in the offers between the FPRF and the FPA treatments. Thus, we must reject Hypothesis 2 as well.

Result to Hypothesis 1 Offers are correlated with the respective values in the FPR and FPRF treatments (p=0.000, linear regression).

Result to Hypothesis 2 There is no significant difference between the offers in the FPA and the FPRF treatment (linear regression, $p = 0.6897$). Offers are significantly higher in the FPR treatment than in the FPA treatment (linear regression, $p = 0.097$).

However, offers in the FPR are significantly higher than in the FPRF treatment (students' t-Test $p = 0.06$). This can also be seen in Table 5.3.

Result to Hypothesis 3 Offers in the FPRF treatment are significantly lower than in the FPR treatment (students' t-Test $p = 0.0587$).

The average counteroffer is given by 53.08 in the FPR treatment and 53.49 in the FPRF treatment, which are both slightly higher than the ex-ante optimal take-it-or-leave-it offer of 50 in the unique equilibrium. A regression of the counteroffer on the offer of the chosen bidder suggests a high correlation between the two in both treatments (see Table 5.4). Therefore me must reject Hypothesis 4, as predicted by our analysis of the QRE (Proposition Proposition 6) in the FPR setting.

Result to Hypothesis 4 Counteroffers in the FPR and FPRF treatments are correlated with the offer of the chosen bidder. (linear regression p=0.000).

The revenues for the auctioneer are very similar in all three treatments (means: FPA: 46.47 FPR: 47.96 FPRF: 45.37). All three average revenues are higher than expected from theory but with a prediction of around 33 ECU in the first-price auction setting and around 25 ECU in the FPR and FPRF treatments (numerical simulations), we can conclude that the auctioneers were able to exploit some of the private information shared by the bidders.

Result to Hypothesis 5 There is no significant difference between the revenues in the FPR and FPRF treatments with respect to the FPA treatment (students' t-Test: FPA-FPR: $p = 0.6883$; FPA-FPRF: $p = 0.3652$).

The difference between the FPR and the FPRF treatment is indeed significant.

Result to Hypothesis 6 The revenue in the FPRF treatment is significantly lower than in the FPR treatment (students' t-Test: $p = 0.0757$).

Regarding the efficiency, we observe no significant differences between the treatments.

Result to Hypothesis 7 There is no significant difference concerning the efficiency between the FPA and the FPR, and the FPA and the FPRF (students' t-Test, FPA-FPR: $p = 0.6310$; FPA-FPRF: $p = 0.2929$).

Summary statistics for the FPA treatment can be found in table 5.1. We observe overbidding in line with the experimental literature, the average offer is given by 34.56, the median offer is 33. From table 5.2, we must reject *Hypothesis 5*. The slope is significantly different from 0.5.

Result to Hypothesis 8 Bidders bid significantly higher than predicted in the FPA treatment (students' t-Test $p = 0.000$).

Dependent variable: Offer						
Treatment:	**FPA**		**FPR**		**FPRF**	
Value	0.66***	0.66***	0.74***	0.74***	0.66***	0.66***
	(22.35)	(22.33)	(31.16)	(31.14)	(28.78)	(28.88)
β_0	0.85*	1.20	-0.24	0.30	0.88	1.92**
	(1.66)	(1.16)	(-0.44)	(0.26)	(1.47)	(2.00)
Period		-0.01		-0.02		-0.04
		(-0.36)		(-0.56)		(-1.46)
Observations	1200		1600		1800	

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.001$

Table 5.2: Panel regression estimates for the offers in the FPR, FPRF and FPA treatments

	Offer
Value	0.689***
	(46.24)
Period	-0.0146
	(-0.38)
FPR	3.249*
	(1.66)
FPRF	0.751
	(0.40)
FPR × Period	-0.00612
	(-0.11)
FPRF × Period	-0.0282
	(-0.58)
Constant	-0.232
	(-0.17)
Observations	4600

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5.3: Panel regression estimates for the effect of the treatment variables on the offers of the bidders

Dependent variable: Counteroffer		
	FPR	FPRF
Offer of chosen bidder	0.798***	0.919***
	(53.39)	(72.56)
Period	0.00582	0.0576***
	(0.29)	(3.58)
Constant	12.75***	9.089***
	(9.02)	(8.45)
Observations	800	900

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5.4: Panel regression estimates for the effect of the treatment variables on the counteroffers of the auctioneers

5.4 CONCLUSION

In this paper, we investigate how bidders react to commitment in first-price auctions in a simple and concise setting. While theory clearly predicts that the offers of the bidders should be higher in the standard first-price auction than in the first-price auction with renegotiation, we cannot verify this experimentally. The same holds true for the revenue of the auctioneer and the efficiency of the mechanisms, however, we find evidence for neither hypothesis. Offers are informative of the bidders' type in the first-price auction with renegotiation but the auctioneers are not able to lever this information into profit. This could reduce the pressure to pool of the bidders, as the quantal response equilibrium analysis insinuates. For real-life procurement, this would mean that a buyer does not need to focus on the commitment of her mechanism and can expect competitive offers, even when the rules on how a winner is selected are not clear. On the other hand, there have been studies that show a strong reaction to a lack of commitment by laboratory participants. This opens the door for further research. For example, it would be interesting to understand how a mechanism can convey commitment in a way that bidders understand and react to varying amounts of it.

5.5 APPENDIX

5.5.1 INSTRUCTIONS

FPA Treatment

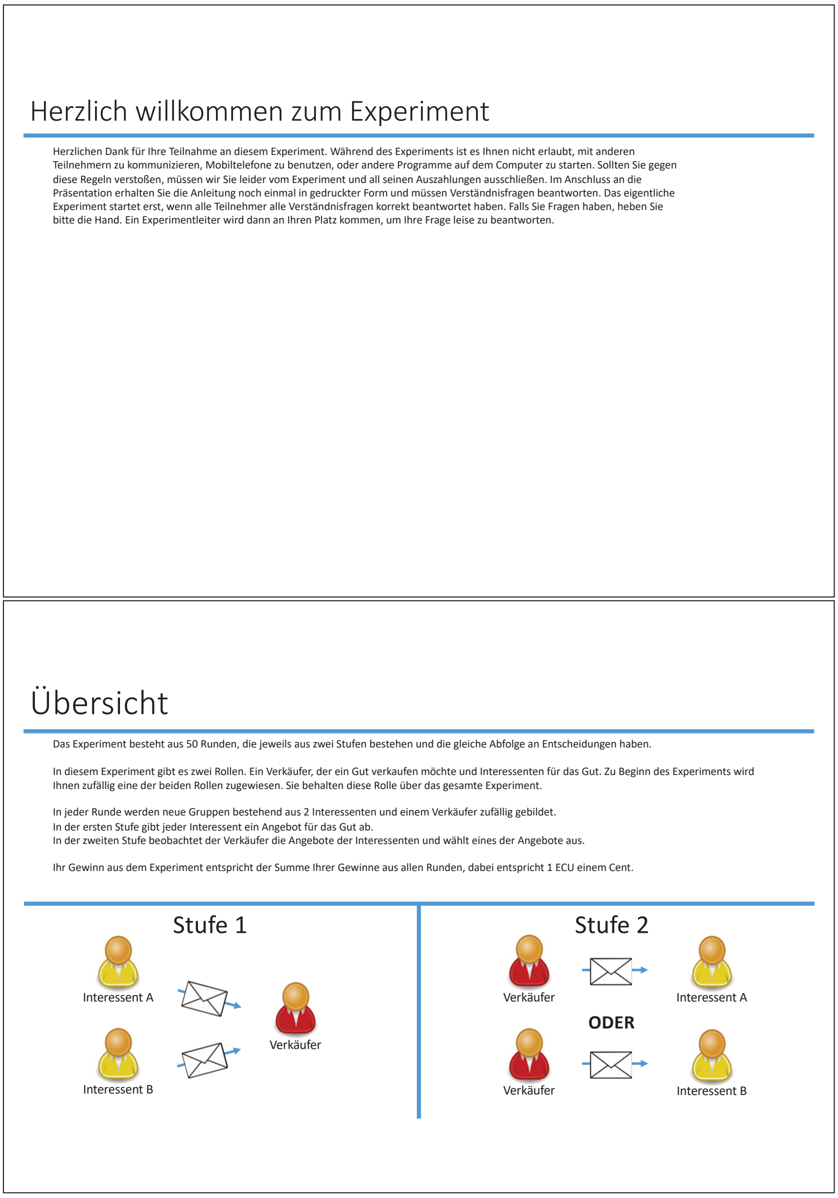

Figure 5.5: Instructions page 1 and 2 for the FPA treatment

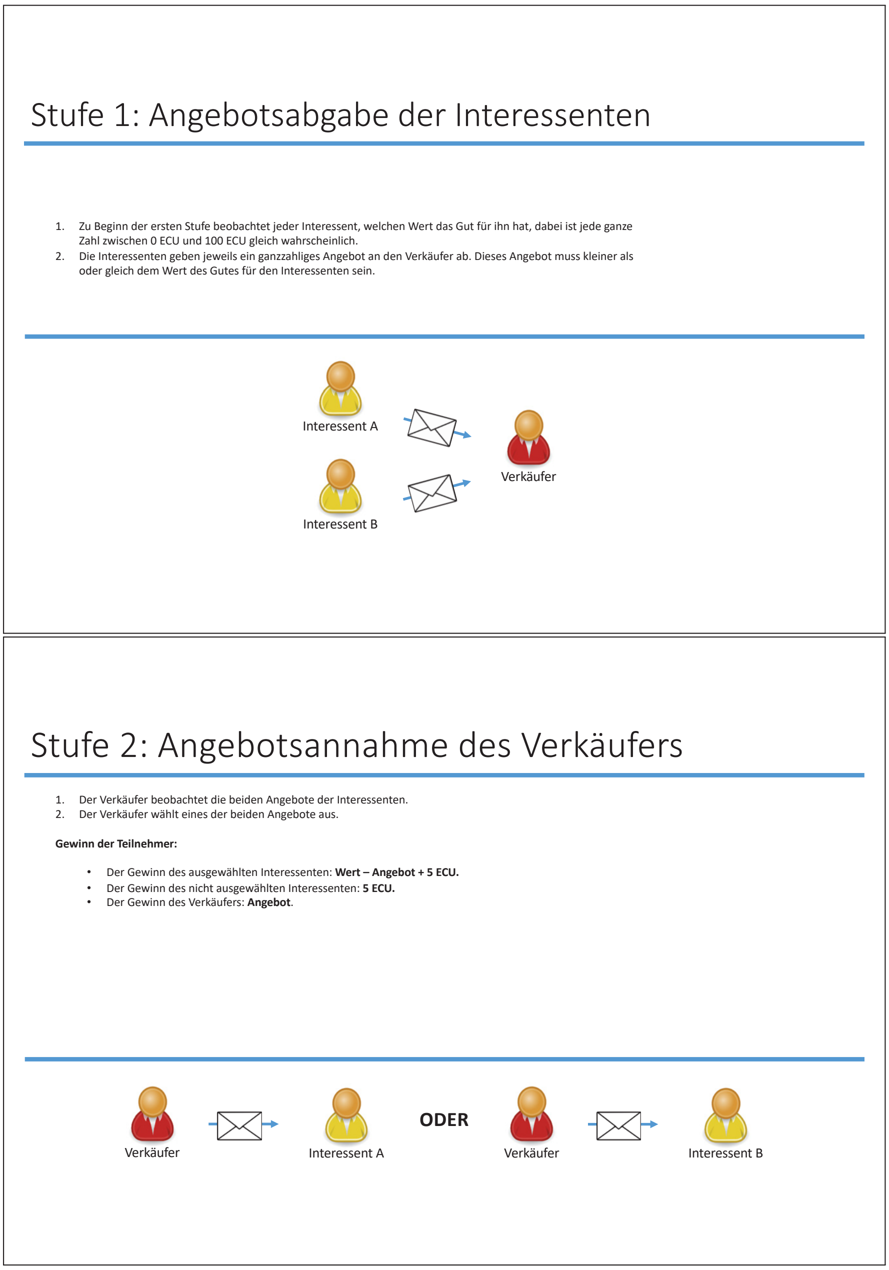

Figure 5.6: Instructions pages 3 and 4 for the FPA reatment

FPR Treatment

Herzlich willkommen zum Experiment

Herzlichen Dank für Ihre Teilnahme an diesem Experiment. Während des Experiments ist es Ihnen nicht erlaubt, mit anderen Teilnehmern zu kommunizieren, Mobiltelefone zu benutzen, oder andere Programme auf dem Computer zu starten. Sollten Sie gegen diese Regeln verstoßen, müssen wir Sie leider vom Experiment und all seinen Auszahlungen ausschließen. Im Anschluss an die Präsentation erhalten Sie die Anleitung noch einmal in gedruckter Form und müssen Verständnisfragen beantworten. Das eigentliche Experiment startet erst, wenn alle Teilnehmer alle Verständnisfragen korrekt beantwortet haben. Falls Sie Fragen haben, heben Sie bitte die Hand. Ein Experimentleiter wird dann an Ihren Platz kommen, um Ihre Frage leise zu beantworten.

Übersicht

Das Experiment besteht aus 50 Runden, die jeweils aus zwei Stufen bestehen und die gleiche Abfolge an Entscheidungen haben.

In diesem Experiment gibt es zwei Rollen. Ein Verkäufer, der ein Gut verkaufen möchte und Interessenten für das Gut. Zu Beginn des Experiments wird Ihnen zufällig eine der beiden Rollen zugewiesen. Sie behalten diese Rolle über das gesamte Experiment.

In jeder Runde werden neue Gruppen bestehend aus 2 Interessenten und einem Verkäufer zufällig gebildet.
In der ersten Stufe gibt jeder Interessent ein Angebot für das Gut ab.
In der zweiten Stufe beobachtet der Verkäufer die Angebote der Interessenten und wählt einen der Interessenten aus, um diesem ein Gegenangebot zu machen.

Ihr Gewinn aus dem Experiment entspricht der Summe Ihrer Gewinne aus allen Runden, dabei entspricht 1 ECU einem Cent.

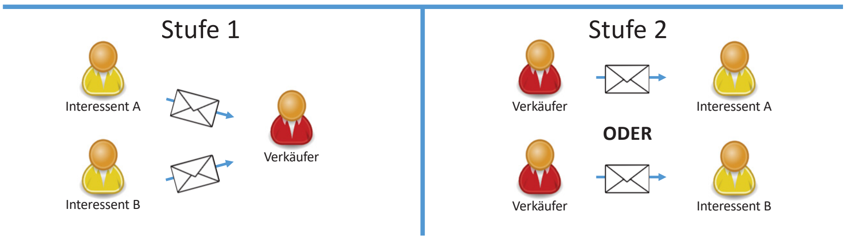

Figure 5.7: Instructions page 1 and 2 for the FPR treatment

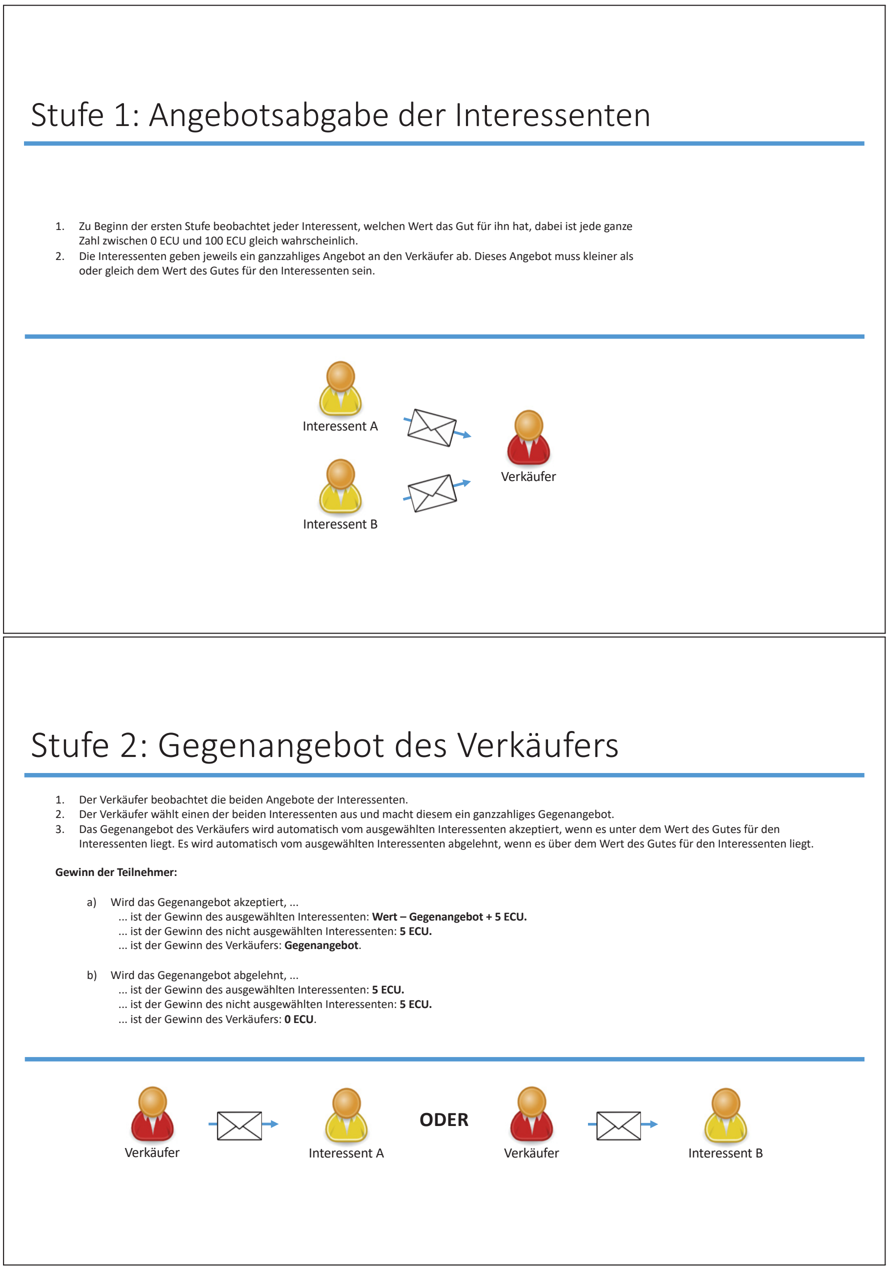

Figure 5.8: Instructions pages 3 and 4 for the FPR treatment

FPRF Treatment

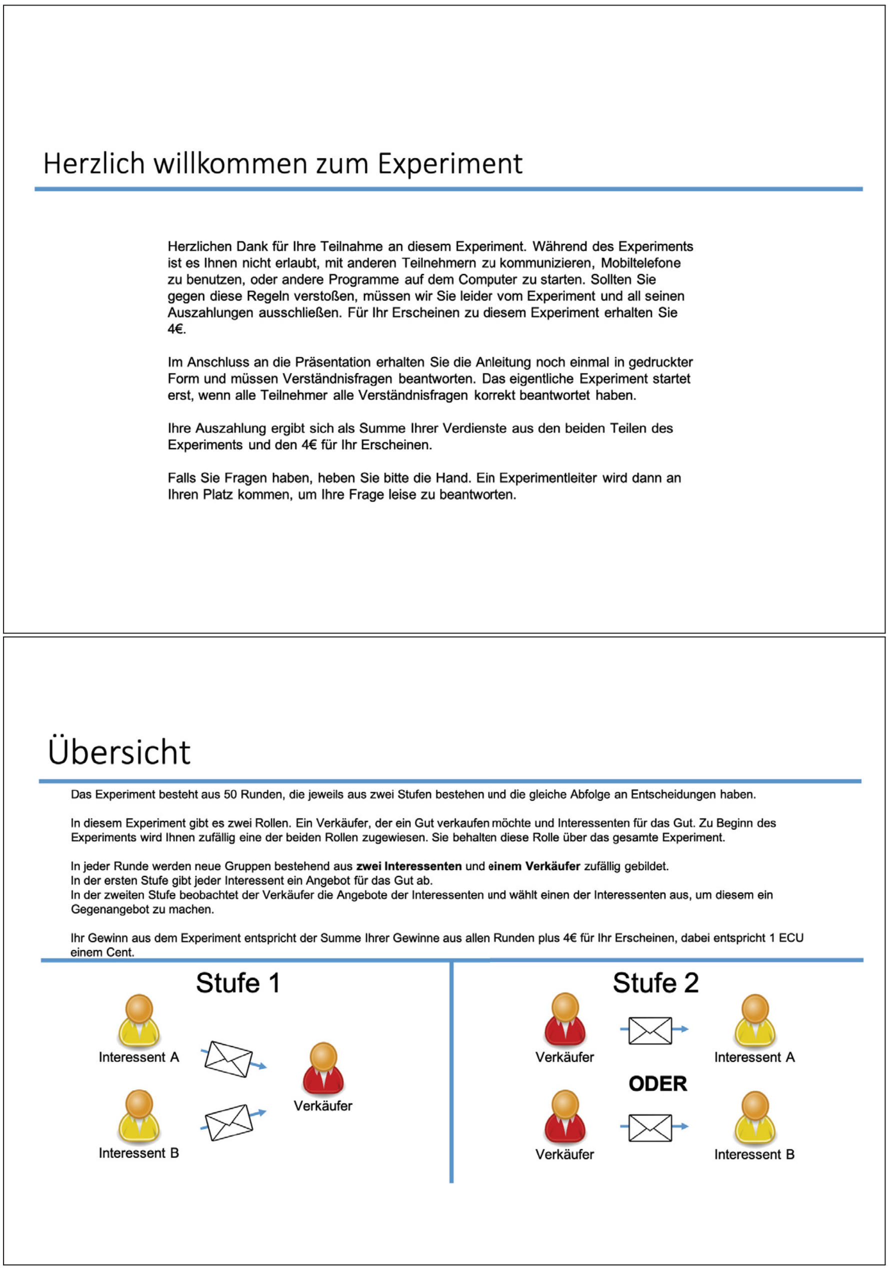

Figure 5.9: Instructions page 1 and 2 for the FPRF treatment

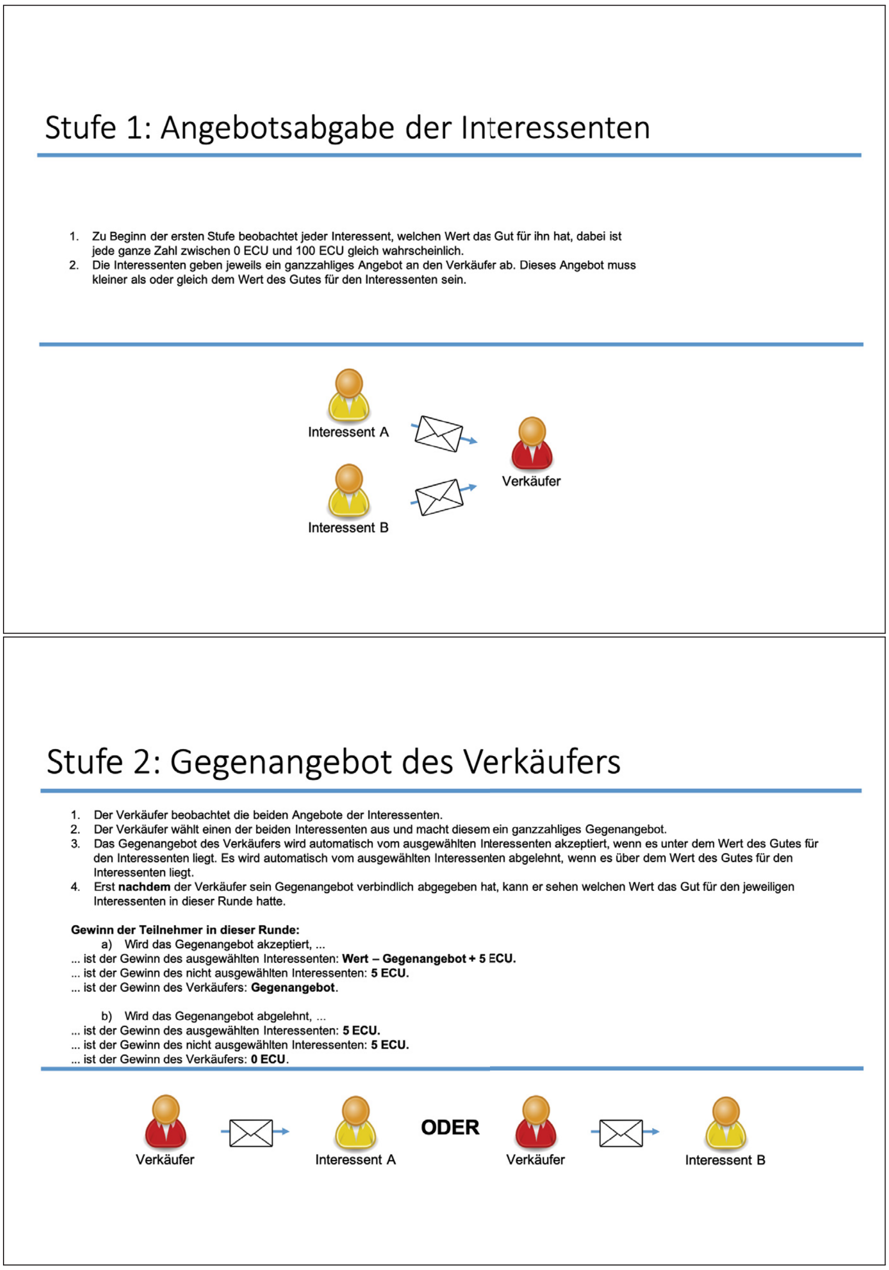

Figure 5.10: Instructions pages 3 and 4 for the FPRF treatment

5.5.2 LOGIT QRE FOR THE FPA

The action space of the bidders is given by $A^B = T = \{t_i\}_{i\in\{0,\dots,10\}}$. The auctioneer's action space is given by $A^A = \{b^1, b^2\} = \{a_i^A\}_{i\in\{1,\dots,2\}}$, either she chooses the offer of bidder one or the offer of bidder two. The expected utility of bidder 1 being of type t_i and submitting an offer of t_j is given by

$$U_1^B(t_i, t_j, \sigma^B, \sigma^A) = \underbrace{\sum_{k=0}^{10}\sum_{l=0}^{k}\sigma_{kl}^B}_{\text{bidder 2: } b^2} \underbrace{\sum_{m=1}^{2}\sigma_{jlm}^A}_{\text{auction.: } a_m^A} \begin{cases} t_i - t_j & a_m^A = b^1 \\ 0 & a_m^A \neq b^1 \end{cases} \tag{207}$$

$$= \sum_{k=0}^{10}\sum_{l=0}^{k}\sigma_{kl}^B\sigma_{jl1}^A(t_i - t_j) \tag{208}$$

$$:= U_{1,ij}^B(\sigma^B, \sigma^A). \tag{209}$$

Analogously, the expected utility of bidder 2 being of type t_i and submitting an offer of t_j is given by

$$U_2^B(t_i, t_j, \sigma^B, \sigma^A) = \sum_{k=0}^{10}\sum_{l=0}^{k}\sigma_{kl}^B \sum_{m=1}^{2}\sigma_{jlm}^A \begin{cases} t_i - t_j & a_m^A = b^2 \\ 0 & a_m^A \neq b^2 \end{cases} \tag{210}$$

$$= \sum_{k=0}^{10}\sum_{l=0}^{k}\sigma_{kl}^B\sigma_{jl2}^A(t_i - t_j) \tag{211}$$

$$:= U_{2,ij}^B(\sigma^B, \sigma^A). \tag{212}$$

The expected utility of the auctioneer having received the offers $b^1 = t_i$ and $b^2 = t_j$ and taking action a_k^A is given by

$$U_{ijk}^A(\sigma^B, \sigma^A) = \begin{cases} t_i & a_k = b^1 \\ t_j & a_k = b^2 \end{cases} \tag{213}$$

This yields the following system of equations,

$$\begin{aligned} \sigma_{ij}^B &= \frac{e^{\lambda U_{1,ij}^B(\sigma^B,\sigma^A)}}{\sum_{k=0}^{i} e^{\lambda U_{1,ik}^B(\sigma^B,\sigma^A)}} \quad \forall\, t_i, t_j \in T \\ \sigma_{ij}^B &= \frac{e^{\lambda U_{2,ij}^B(\sigma^B,\sigma^A)}}{\sum_{k=0}^{i} e^{\lambda U_{2,ik}^B(\sigma^B,\sigma^A)}} \quad \forall\, t_i, t_j \in T \\ \sigma_{ijk}^A &= \frac{e^{\lambda U_{ijk}^A(\sigma^B,\sigma^A)}}{\sum_{m=1}^{2} e^{\lambda U_{ijm}^A(\sigma^B,\sigma^A)}} \quad \forall\, t_i, t_j \in T \text{ and } \forall\, a_k^A \in A^A. \end{aligned} \tag{214}$$

CHAPTER 6

PRE-AUCTION OR POST-AUCTION QUALIFICATION?

Abstract

We compare auctions with bidder qualification before or after the bidding process. We show that although post-auction qualification is more efficient, the auctioneer prefers pre-auction qualification when bidders' qualification costs are high.

6.1 INTRODUCTION

Bidder qualification plays an important role in real-life auctions and procurement procedures; however, verifying the qualification of a bidder is costly for both the buyer and potential sellers. For example, it may involve product testing, the inspection of production facilities, or the acquisition of industry-norm qualifications.[1] Interestingly, in most procurement processes, bidders are required to undergo qualification before the actual awarding of a project, which results in significant costs even for bidders who fail to win the project. Typically, this is explained by the risk of qualification failure. That is, if the bidder who wins the auction does not qualify, the procurement procedure will have to be repeated.

However, there are many situations in which the risk of qualification failure is not an issue, but qualification requirements are nevertheless in place. First, in real-life auctions, buyers often deal with a limited set of known suppliers, and qualification failure is relatively rare.[2] Second, the auctioneer may feel obligated or be required to document the fact that

[1]See Wan and Beil (2009) for details.

[2]For example, in the automotive industry, buyers and engineers are highly reluctant to include unknown suppliers in procurement processes. As a result, usually only a few well-known suppliers participate, but each supplier's product is still extensively tested before the procurement process starts.

bidders meet certain standards (e.g., a DIN certification requirement by the government, environmental obligations) for the benefit of less-informed third parties (e.g., authorities, courts, superiors). Third, if the buyer's requirements for the qualification are known to potential bidders, they may decide to invest in the assets that are needed to achieve qualification. In this case, qualification costs in our model can be reinterpreted as necessary investment costs. Our question then reduces to whether the auctioneer should require bidders to invest prior to the auction or after the winner is known. Fourth, our model also applies to a setting in which the sellers know for certain (e.g., from previous experience or for technological reasons) whether they will achieve qualification or not. In this situation, it does not make sense for a bidder who would eventually fail qualification to participate in the auction. As a consequence, every bidder who participates in the auction will be qualified. Note that qualification is still necessary to deter unqualified suppliers from participating. Absent the risk of qualification failure, it is puzzling that the cost advantages of qualifying only the winner of the procurement process are rarely realized.

We address this observation by constructing a simplistic model without the risk of qualification failure. We then analyze when an auctioneer should demand proof of bidders' qualification: before or after the auction.[3] Under pre-auction qualification, all bidders must become qualified before they can participate in the auction. Thus, there is an exclusion effect. Given that the cost of qualification must be paid prior to participation and qualified bidders may fail to win the auction, the expected surplus of bidders with valuations above but close to the qualification cost is negative. These bidders therefore refrain from participation even if their valuation is higher than the qualification cost. Under post-auction qualification, only the winner must undergo costly qualification, resulting in a bid-shading effect. During the bidding procedure, bidders will keep in mind that upon winning they will have to pay the cost of qualification. Consequently, bidders will shade their bids by an amount that is equal to the qualification cost. The ex-

[3]Even though typically both the buyer and the sellers bear the cost of qualification, it is convenient to assume that bidders incur all such costs.

clusion effect reduces revenues from pre-auction qualification, whereas the bid-shading effect reduces revenues from post-auction qualification.

We show that, interestingly, pre-auction qualification is more profitable if the qualification costs are sufficiently high. With pre-auction qualification, each bidder decides whether to enter the auction based on the incurred cost and his or her winning probability. If qualification costs rise, the cost of participation increases. At the same time, however, fewer bidders participate, which means that the winning probability increases. This increase in the winning probability dampens the increase in the exclusion effect, and the marginal increase eventually goes to zero for very high qualification costs. With post-auction qualification, the bid-shading effect increases linearly with the increase in qualification costs, and the marginal increase is one for all cost levels. Thus, as qualification costs increase, the bid-shading effect becomes more important and pre-auction qualification yields higher revenues.

This result is somewhat surprising, as previous studies have focused on the risk of project failure (due to bankruptcy, lack of expertise, etc.) in order to explain the prevalence of pre-auction qualification in real-life procurement settings (see below). Without the risk of project failure, as in our model, the superiority of post-auction qualification might appear self-evident. However, our analysis highlights that pre-auction qualification may be beneficial for the auctioneer even in the absence of the risk of project failure.

Wan and Beil (2009) show that when the buyer can delay part or all of qualification until after the auction, his or her expected payments can increase due to the fact that bids from unqualified bidders will be discarded. However, in their setting, delaying some or all qualification saves costs. Thus, the authors conclude, the standard practice of pre-auction qualification screenings can be improved upon by the judicious use of post-auction qualification. In their model, when the risk of qualification failure disappears, post-auction qualification becomes unambiguously dominant. This is not true in our case, as even without the risk of qualification failure, we show that pre-auction qualification may be desirable. In their setting, the auctioneer draws suppliers from an infinite pool, and pre-qualification implies that the auctioneer will pre-qualify suppliers until he or she finds n suppliers that pass the pre-qualification threshold. Thus, if the risk of qualification failure disappears, auctioneers still bear the costs of pre-qualifying n potential bidders if they insist on using full or partial pre-qualification. In our setting, only bidders that have sufficiently high valuations enter the auction and pay the qualification costs. Moreover, in their setting, the risk of qualification failure has different effects on the bidding behavior in the two qualification scenarios. When there is no risk of qualification failure and the auctioneer picks the same number of bidders in both scenarios, this difference disappears and bidders submit the same bids in the two qualification regimes. In contrast, in our setting, bidders never submit identical bids in the two scenarios due to the diverging exclusion and bid-shading effects. Overall, pre-qualification may become more attractive compared to the setting found in Wan and Beil (2009).

Wan et al. (2012) consider a setting in which a potentially unqualified entrant competes in an open descending reverse auction with a qualified incumbent. They find that, in the case of post-auction qualification, the incumbent will hold back on bidding in the auction. This is due to the fact that when the entrant fails qualification, the incumbent will be asked to deliver the project even if he or she lost the auction in the first place. Again, when the risk of qualification failure disappears, post-auction qualification becomes unambiguously dominant.[4]

6.2 MODEL

There are N risk-neutral bidders competing in a second-price sealed-bid auction for one indivisible object (with $N \geq 2$). Before the auction starts, each bidder privately observes his or her valuation $v_i \in [0, 1]$, where $i \in$

[4]Both Wan and Beil (2009) and Wan et al. (2012) assume that the auctioneer bears all the costs of qualification. However, they also point out that who bears the qualification costs does not change the results. Thus, the only difference compared to the assumptions in the present setup is the fact that suppliers may fail qualification.

$\{1, \ldots, N\}$.[5] Valuations are independently distributed according to a common distribution function F that is assumed to be absolutely continuous and to have a density f with $\operatorname{supp} f = [0, 1]$ and $f(1) < \infty$. The seller's valuation is assumed to be equal to zero. In order to be selected as the winner of the auction, the bidders must be qualified. Qualification comes at a cost c that must be borne by bidders and that is common knowledge among all bidders (with $c \in [0, 1]$). We explicitly abstract from the possibility that any of the bidders may fail to become qualified.[6]

We compare the following two scenarios. Under *pre-auction qualification*, all bidders decide simultaneously whether to invest in qualification after they have learned their valuation but before the auction starts.[7] Under *post-auction qualification*, a bidder must invest in qualification only if he or she wins the auction.

6.3 EQUILIBRIUM BIDDING

Finding a symmetric equilibrium bidding strategy under pre-auction qualification is equivalent to finding a symmetric equilibrium bidding strategy in an auction with participation or bidding costs.[8] For the case of post-auction qualification, the decision problem of the bidder with valuation v is equivalent to the decision problem of a bidder in an auction without qualification costs and an ex-post valuation of $v - c$. Thus, equilibrium bidding in the second-price auction can be characterized as follows.

Proposition 1. *In the pre-auction scenario, there exists a symmetric equilibrium in which only bidders with* $v \geq \underline{v}(c)$ *are qualified, where* $\underline{v}(c)$ *is*

[5] We frame the setting as a selling rather than a procurement auction. This is without loss of generality and has the advantage that most readers are more familiar with the notation.

[6] This is due to the fact that we are interested in determining whether the requirement of pre-auction qualification has positive effects other than minimizing the risk that a bidder may fail qualification.

[7] We assume that bidders who enter the auction do not observe the number of competitors. However, as the second-price auction is dominance solvable, relaxing this assumption does not change the results.

[8] See, e.g., Menezes and Monteiro (2000), Kaplan and Sela (2006), Tan and Yilankaya (2006), and Celik and Yilankaya (2009).

implicitly defined by

$$\underline{v}(c)F(\underline{v}(c))^{N-1} = c.$$

The increasing equilibrium bidding function $\beta_{\text{pre}}(v) : [\underline{v}(c), 1] \to [0, 1]$ *is then given by*

$$\beta_{\text{pre}}(v) = v.$$

In the post-auction scenario, only bidders with $v \geq c$ *undergo qualification. The symmetric and increasing equilibrium bidding function* $\beta_{\text{post}}(v) : [c, 1] \to [0, 1]$ *is given by*

$$\beta_{\text{post}}(v) = v - c.$$

Bidding under pre-auction qualification is driven by what we call the exclusion effect. As $\underline{v}(c) \geq c$, there are bidders who have valuations that are greater than the qualification costs but who nevertheless do not participate in the auction.[9] With post-auction qualification, bidders take into account that upon winning they will have to pay the qualification cost. Consequently, bidders shade their bids, resulting in what we call the bid-shading effect. The trade-off between this bid-shading effect and the exclusion effect is what will drive our results concerning revenue.[10]

6.4 EFFICIENCY AND REVENUE

Conditional on the object being sold, the allocation in both qualification regimes is efficient.[11] Moreover, with post-auction qualification, only winners pay for qualification, and they do so if and only if their valuation is

[9] Note that as $F(\underline{v}(c))^{N-1} \leq 1 \; \forall c \in [0, 1]$, it follows immediately that $\underline{v}(c) \geq c$.

[10] For the comparison of the two qualification regimes, we only consider symmetric equilibria. However, with post-auction qualification, the symmetric equilibrium is the unique equilibrium in undominated strategies. With pre-auction qualification, the symmetric equilibrium is also unique if $F(c)$ is concave (see Tan and Yilankaya, 2006, for details). If $F(c)$ is not concave, asymmetric equilibria (in the sense that bidders use different cut-offs when they decide whether to participate in the auction) may exist.

[11] In fact, in all standard auctions, the allocation of the object is efficient among those bidders who entered the auction when the bidders use symmetric cut-offs. Moreover, in a symmetric equilibrium, the decision of whether to enter the auction is the same for all auction types, as the borderline entrant only wins against those bidders who do not enter. Overall, revenue equivalence is preserved in our setting. Thus, the results from this section apply to all standard auction formats.

greater than or equal to c. Thus, from an ex-ante point of view, post-auction qualification is always more efficient than pre-auction qualification, not only in terms of allocative efficiency, but also in terms of overall surplus taking into account the qualification costs. We summarize this finding in the following proposition, which also comments on the optimal mechanism:

Proposition 2. *In the equilibrium of the second-price auction with post-auction qualification described in Proposition 1, ex-ante and ex-post efficiency is maximized for all mechanisms in which the winning bidder needs to be qualified.*[12]

Proof. See *Appendix A.* □

Interestingly, despite the higher efficiency, revenues can be lower in the scenario with post-auction qualification, as the following proposition shows:

Proposition 3. *Post-auction qualification yields higher revenues than pre-auction qualification whenever*

$$c < \Pr\left(Y_2^{(N)} \leq \underline{v}(c) \middle| Y_2^{(N)} \geq c\right) \mathbb{E}\left[Y_2^{(N)} \middle| c \leq Y_2^{(N)} \leq \underline{v}(c)\right], \tag{215}$$

where $Y_2^{(N)}$ denotes the second order statistic of N draws from F.

Proof. See *Appendix B.* □

Revenues in both scenarios depend on the second-highest value of the bidders. The bid-shading effect in the post-auction qualification regime lowers the second-highest value by c. The exclusion effect implies that if the second-highest value is between c and $\underline{v}(c)$, the bidder with this value chooses not to participate, and thus the revenue is zero in the pre-auction qualification scenario. The same bidder would participate in the auction

[12] In particular, the efficiency of the setting with post-auction qualification is higher than that of any symmetric or asymmetric equilibrium of the auction with pre-auction qualification. Note that Blume and Heidhues (2004) show that in a second-price auction with more than two bidders an effective reserve price implies the uniqueness of the equilibrium. Thus, if $N > 2$, the second-price auction with post-auction qualification has a unique equilibrium in which the ex-ante and ex-post efficiency is maximized. If $N = 2$, other asymmetric equilibria may exist. However, in this case, the described equilibrium is the unique symmetric equilibrium.

with post-auction qualification, yielding a positive revenue. Thus, the left-hand side of inequality (215) corresponds to the revenue loss due to the bid-shading effect, whereas the right-hand side corresponds to the revenue loss due to the exclusion effect.[13]

Unfortunately, the right-hand side of inequality (215) is not monotone either in the number of bidders N or in the qualification cost c. This makes comparative statics somewhat difficult. However, some general results can still be derived:

Proposition 4. *For each N, there exists a cutoff value c' such that revenues from pre-auction qualification are strictly higher than revenues from post-auction qualification for all $c \geq c'$.*

Proof. See *Appendix C.* □

This surprising result indicates that when qualification costs are high, the seller benefits from asking each bidder to undergo qualification before entering the auction. This is due to the fact that with pre-auction qualification, each bidder decides whether to enter the auction based on the incurred cost and his or her winning probability. As qualification costs rise, the cost of participation increases. However, fewer bidders participate, and thus the winning probability also increases. This increase in winning probability mitigates the increase in the exclusion effect, and the marginal increase goes to zero. Under post-auction qualification, the bid-shading effect increases linearly with the increase in qualification costs, and the marginal increase is one for all cost levels. Thus, the bid-shading effect becomes more important, which eventually results in higher revenues under pre-auction qualification.

For lower levels of the qualification cost, the revenue ranking depends on the distribution of the valuations, as highlighted in the following proposition:

[13]Our comparison of revenue is valid only for symmetric equilibria in the scenario with pre-auction qualification. However, as noted in footnote 10, asymmetric equilibria may exist under pre-auction qualification. Such asymmetric equilibria might generate higher revenues than the symmetric equilibrium considered in the present paper (see Celik and Yilankaya, 2009, for details).

Proposition 5. *Suppose the limit*

$$\lim_{v \to 0} \frac{F(v)}{vf(v)} \tag{216}$$

exists. Then, if

$$\lim_{v \to 0} \frac{F(v)}{vf(v)} < N^2 - 2N + 1 =: k(N), \tag{217}$$

there exists a cutoff value c'' such that revenues from post-auction qualification are higher than revenues from pre-auction qualification for all $c \leq c''$.

If

$$\lim_{v \to 0} \frac{F(v)}{vf(v)} > k(N),$$

then there exists a cutoff value c''' such that revenues from pre-auction qualification are higher than revenues from post-auction qualification for all $c \leq c'''$.

Proof. See *Appendix D.* □

Condition (217) can be easily verified for a large class of distribution functions. In particular, condition (217) holds true for all convex distribution functions. Thus, if F is convex or if N is relatively high, post-auction qualification yields higher revenues for low qualification costs. This is due to the fact that convex distributions assign a relatively small winning probability to low valuations. In such a case, a bidder deciding whether to invest in qualification faces a relatively unfavorable probability of winning the object and thus does not enter even if the cost of qualification is low. This is also true when the number of bidders is relatively high. In such cases, therefore, the exclusion effect outweighs the bid-shading effect when the qualification costs are small.[14]

To illustrate our results, we provide three examples.

[14] When the auctioneer has more degrees of freedom in designing the auction than only choosing the timing of the qualification, post-auction qualification dominates pre-auction qualification in terms of revenue for any initial value of the qualification costs. This is due to the fact that the optimal auction with pre-auction qualification involves subsidies to all losing bidders. This can be replicated with post-auction certification without cost. See Menezes and Monteiro (2000) for details.

Uniform distribution

For the uniform distribution, we have $\lim_{v\to 0} F(v)/vf(v) = \lim_{v\to 0} v/v = 1$. Since $N^2 - 2N + 1 > 1$ for all $N > 2$, post-auction qualification yields higher revenues for low qualification costs and $N > 2$. Note that the case in which $N = 2$ is not covered by *Proposition 5*, as $N^2 - 2N + 1 = 1$, but it is easily verified that post-auction qualification will yield higher revenues for all qualification costs in this situation. The results are illustrated in *Figure 6.1a*. *Figure 6.2a* identifies when an auctioneer prefers to employ pre-auction or post-auction qualification, depending on the qualification costs and the number of bidders.

$$F(v) = v^2$$

In this case, $\lim_{v\to 0} F(v)/vf(v) = \lim_{v\to 0} v^2/2v^2 = 1/2$. Since $N^2-2N+1 > 1/2$ for all $N > 1.7$, post-auction qualification will yield higher revenues for any number of bidders and low qualification costs. The difference in revenue between pre-auction and post-auction qualification for $N = 2$ in this case is illustrated in *Figure 6.1b*.

$$F(v) = \sqrt{v}$$

In this case, we have $\lim_{v\to 0} F(v)/vf(v) = \lim_{v\to 0} 2\sqrt{v}\sqrt{v}/v = 2$. Because $N^2 - 2N + 1 > 2$ for all $N > 2$, post-auction qualification yields higher revenues for low qualification costs and $N > 2$. *Figure 6.2b* identifies when an auctioneer would prefer to employ pre-auction or post-auction qualification, depending on the qualification costs and the number of bidders.

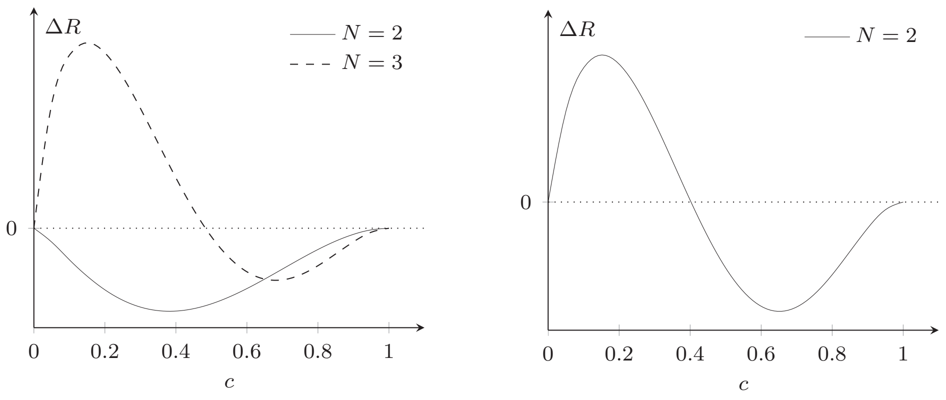

(a) Uniform distribution. (b) $F(v) = v^2$.

Figure 6.1: Difference in expected revenue ΔR (where $\Delta R = \mathbb{E}[R_{\text{post}}] - \mathbb{E}[R_{\text{pre}}]$).

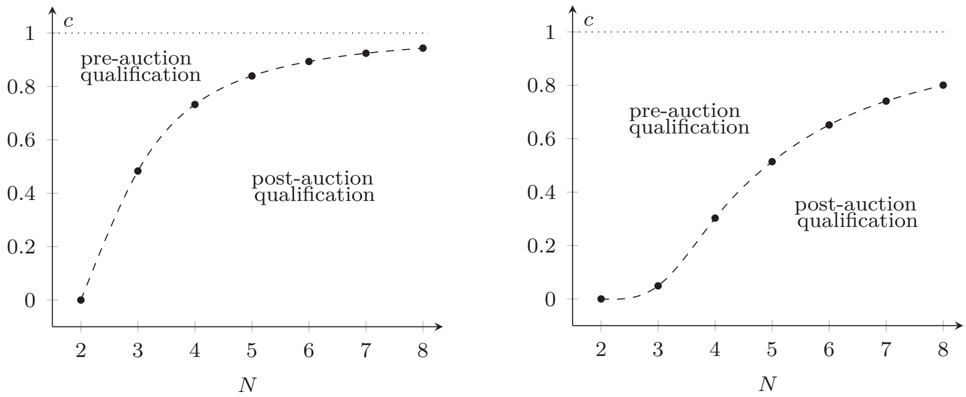

(a) Uniform distribution. (b) $F(v) = \sqrt{v}$.

Figure 6.2: Optimality of qualification regimes from the auctioneer's point of view (where the black circles indicate identical revenues in both qualification settings).

6.5 APPENDIX

6.5.1 PROOF OF PROPOSITION 2

Proof. An efficient mechanism has the following three properties: (i) it allocates the object if and only if $v \geq c$; (ii) if the object is allocated, it goes to the bidder with the highest valuation; and (iii) only the winner is qualified. The monotonicity of the equilibrium described in Proposition 1 implies that the object is allocated to the bidder with the highest valuation, i.e., property (ii) is satisfied. The fact that bidders only become qualified if $v \geq c$ ensures that the object is allocated if and only if the winner's valuation is greater than the qualification cost. Hence, property (i) is fulfilled. Finally, post-auction qualification implies that in equilibrium only the winner undergoes qualification, which means that property (iii) is satisfied. □

6.5.2 *PROOF OF* PROPOSITION 3

Proof. An auctioneer can expect to generate revenue (denoted by R)

$$\mathbb{E}\left[R_{\text{pre}}\right] = \Pr\left(Y_2^{(N)} > \underline{v}(c)\right)\mathbb{E}\left[Y_2^{(N)}\middle|Y_2^{(N)} > \underline{v}(c)\right]$$

from pre-auction qualification and revenue

$$\mathbb{E}\left[R_{\text{post}}\right] = \Pr\left(Y_2^{(N)} > c\right)\left(\mathbb{E}\left[Y_2^{(N)}\middle|Y_2^{(N)} > c\right] - c\right)$$

from post-auction qualification. Note that it holds that

$$\begin{aligned}\Pr\left(Y_2^{(N)} > c\right)\mathbb{E}\left[Y_2^{(N)}\middle|Y_2^{(N)} > c\right] =& \Pr\left(Y_2^{(N)} > \underline{v}(c)\right)\mathbb{E}\left[Y_2^{(N)}\middle|Y_2^{(N)} \geq \underline{v}(c)\right] \\ &+ \Pr\left(\underline{v}(c) \geq Y_2^{(N)} \geq c\right)\mathbb{E}\left[Y_2^{(N)}\middle|\underline{v}(c) \geq Y_2^{(N)} \geq c\right].\end{aligned}$$

Using this, we now consider the difference in the expected revenue between the two qualification regimes,

$$\begin{aligned}&\Delta R := \mathbb{E}\left[R_{\text{post}}\right] - \mathbb{E}\left[R_{\text{pre}}\right] \geq 0 \\ \Leftrightarrow\ &\Pr\left(Y_2^{(N)} > c\right)\left(\mathbb{E}\left[Y_2^{(N)}\middle|Y_2^{(N)} > c\right] - c\right) - \Pr\left(Y_2^{(N)} > \underline{v}(c)\right)\mathbb{E}\left[Y_2^{(N)}\middle|Y_2^{(N)} > \underline{v}(c)\right] \geq 0 \\ \Leftrightarrow\ &\Pr\left(Y_2^{(N)} > \underline{v}(c)\right)\mathbb{E}\left[Y_2^{(N)}\middle|Y_2^{(N)} \geq \underline{v}(c)\right] \\ &+ \Pr\left(\underline{v}(c) \geq Y_2^{(N)} \geq c\right)\mathbb{E}\left[Y_2^{(N)}\middle|\underline{v}(c) \geq Y_2^{(N)} \geq c\right] - c\Pr\left(Y_2^{(N)} > c\right) \\ &- \Pr\left(Y_2^{(N)} > \underline{v}(c)\right)\mathbb{E}\left[Y_2^{(N)}\middle|Y_2^{(N)} > \underline{v}(c)\right] \geq 0 \\ \Leftrightarrow\ &\frac{\Pr\left(\underline{v}(c) \geq Y_2^{(N)} \geq c\right)}{\Pr\left(Y_2^{(N)} > c\right)}\mathbb{E}\left[Y_{(2)}^{(N)}\middle|\underline{v}(c) \geq Y_2^{(N)} \geq c\right] - c \geq 0 \\ \Leftrightarrow\ &\Pr\left(\underline{v}(c) \geq Y_2^{(N)}\middle|Y_2^{(N)} > c\right)\mathbb{E}\left[Y_2^{(N)}\middle|\underline{v}(c) \geq Y_2^{(N)} \geq c\right] \geq c.\end{aligned}$$

Because $0 \leq \Pr\left(\underline{v}(c) \geq Y_2^{(N)}\middle|Y_2^{(N)} > c\right)\mathbb{E}\left[Y_2^{(N)}\middle|\underline{v}(c) \geq Y_2^{(N)} \geq c\right] \leq 1$, this inequality is well defined and the proposition holds. □

6.5.3 PROOF OF PROPOSITION 4

Proof. As the revenue is equal to zero for $c = 1$ in both scenarios, we will show that the difference in revenue ΔR has a local maximum at $c = 1$. Consider the following derivatives:[15]

$$\frac{\partial}{\partial c}\Delta R = \underline{v} f_2^{(N)}\left(\underline{v}\right)\underline{v}_c + F_2^{(N)}(c) - 1$$

and

$$\frac{\partial^2}{\partial c^2}\Delta R = \underline{v}_c^2\left(f_2^{(N)}(\underline{v}) + \underline{v} f_2^{(N)\prime}(\underline{v})\right) + \underline{v} f_2^{(N)}(\underline{v})\frac{\partial^2 \underline{v}}{\partial c^2} + f_2^{(N)}(c),$$

with

$$\underline{v}_c := \frac{\partial \underline{v}}{\partial c} = \frac{1}{F_1^{(N-1)}(\underline{v}) + \underline{v}(n-1)F_1^{(N-2)}(\underline{v})f(\underline{v})}.$$

Note that for $\lim_{c\to 1}$, both $\underline{v}_c$ and $\frac{\partial \underline{v}_c}{\partial c}$ stay bounded. Also note that for all probability densities f, $f_2^{(N)}(1) = 0$ and $\lim_{v\to 1} f_2^{(N)\prime}(v) < 0$.

Then,

$$\lim_{c\to 1}\frac{\partial}{\partial c}\Delta R = 0 \tag{218}$$

and

$$\lim_{c\to 1}\frac{\partial^2}{\partial c^2}\Delta R = \lim_{c\to 1}\underline{v}_c^2 \underline{v} f_2^{(N)\prime}(v) < 0. \tag{219}$$

The second limit is a strict inequality because we assumed $f(1) < \infty$ which means that $\lim_{c\to 1}\underline{v}_c > 0$. Finally note that expression (218) provides the possible extremum in $c = 1$ whereas expression (219) confirms that it is a local maximum. Thus, in a neighborhood $U = (c', 1)$, $\Delta R < 0$ for all $c \in U$. □

6.5.4 PROOF OF PROPOSITION 5

Proof. As the difference in revenue ΔR is equal to zero for $c = 0$ (revenue equivalence), we will show that the first derivative of ΔR is positive or

[15] We write $F_k^{(N)}$ to denote the distribution of the k-th order statistic of N draws and $f_k^{(N)}$ to denote its density.

negative, respectively. Observe that

$$\begin{aligned}\frac{\partial}{\partial c}\Delta R &= \frac{N(N-1)(1-F(\underline{v}))f(\underline{v})}{\frac{F(\underline{v})}{\underline{v}} + (N-1)f(\underline{v})} - 1 + F_2^{(N)}(c)\\ &= \frac{N(N-1)(1-F(\underline{v}))}{\frac{F(\underline{v})}{\underline{v}f(\underline{v})} + N - 1} - 1 + F_2^{(N)}(c).\end{aligned}$$

Now there are two possible scenarios:

1. $\lim_{c\to 0} \frac{\partial}{\partial c}\Delta R \to \infty$, and
2. $\lim_{c\to 0} \frac{\partial}{\partial c}\Delta R < \infty$ for all N.

If the first possibility holds, we find that

$$\lim_{c\to 0} \frac{N(N-1)(1-F(\underline{v}))}{\frac{F(\underline{v})}{\underline{v}f(\underline{v})} + N - 1} \to \infty.$$

However, this means that $F(\underline{v})/\underline{v}f(\underline{v}) + N - 1 \to 0$, which implies that $F(\underline{v})/vf(\underline{v}) \to 1 - N < 0$. This contradicts

$$\begin{cases} F \text{ probability distribution}\\ f \text{ density}\\ \underline{v} \in [0,1], \end{cases}$$

which means that the second case is the relevant one. Because we are dealing with a rational function, and since the limits on both sides exist, we can write

$$\lim_{v\to 0} \frac{N(N-1)(1-F(\underline{v}))}{\frac{F(\underline{v})}{\underline{v}f(\underline{v})} + N - 1} = \frac{\lim_{v\to 0} N(N-1)(1-F(\underline{v}))}{\lim_{v\to 0} \frac{F(\underline{v})}{\underline{v}f(\underline{v})} + N - 1}.$$

We can therefore identify three possible scenarios:

1. $\lim_{c\to 0} F(\underline{v})/\underline{v}f(\underline{v}) = 0$,
2. $\lim_{c\to 0} F(\underline{v})/\underline{v}f(\underline{v}) = m > 0$, and
3. $\lim_{c\to 0} F(\underline{v})/\underline{v}f(\underline{v}) = \infty$.

For the first two scenarios,

$$\begin{aligned}
\lim_{c\to 0}\frac{\partial}{\partial c}\Delta R &= \lim_{c\to 0}\frac{N(N-1)(1-F(\underline{v}))}{\frac{F(\underline{v})}{\underline{v}f(\underline{v})}+(N-1)} - 1 + F_2^{(N)}(c) \\
&= \frac{\lim_{c\to 0}(N(N-1)(1-F(\underline{v})))}{\lim_{c\to 0}\left(\frac{F(\underline{v})}{\underline{v}f(\underline{v})}+(N-1)\right)} - 1 \\
&= \frac{N(N-1)}{m+(N-1)} - 1.
\end{aligned}$$

For $m = 0$, this results in $\lim_{c\to 0}\frac{\partial}{\partial c}\Delta R = N - 1$. For $m > 0$, this leaves us with the condition

$$m < N^2 - 2N + 1 = k(N)$$

for $\lim_{c\to 0}\frac{\partial}{\partial c}\Delta R > 0$.

In the third scenario, $\frac{\partial}{\partial c}\Delta R = -1 < 0$. □

CHAPTER 7

BID POOLING IN REVERSE MULTI-UNIT DUTCH AUCTIONS – AN EXPERIMENTAL INVESTIGATION

Abstract

In this paper, we experimentally investigate reverse multi-unit Dutch auctions in which bidders compete to sell their single unit to a buyer who wants to purchase several objects. Our study yields three insights. (i) Bids are substantially higher than Nash equilibrium bids predicted by standard economic theory; (ii) these higher-than-predicted prices gradually decline in later periods; and (iii) bid pooling (or simultaneous bidding) is frequently observed—the majority of bidders submit their bids immediately after the first bidder has sold his unit. A model that distinguishes between myopic and sophisticated bidding strategies helps to organize these patterns both on the aggregate and on the individual level.

7.1 INTRODUCTION

Multi-unit Dutch auctions and their procurement counterparts are implemented in a variety of real-world markets. In these settings, bid pooling (also known as bidding frenzy), i.e., many bidders submit bids at the same time/clock price, and crashes, i.e., situations where bidders withhold bids, are frequently observed phenomena (see Bulow and Klemperer, 1994). One example for the use of multi-unit Dutch auctions in practice is the sale of new securities by US underwriters. There, an initial price is maintained or supported as long as either an issue is sold out or demand turns out so low such that a significant price decrease is necessary. Moreover, multi-unit Dutch auctions are also used in commodity markets such as fish markets or markets for fresh produce (see Cassady, 1967; Romeu, 2000). Furthermore, tickets for concerts, shows, etc. are typically sold on a first-come-first-serve basis which can be interpreted as a multi-unit descending auction.[1]

Despite the practical importance of multi-unit Dutch procurement auctions in reality and the (potentially) detrimental effects of bid pooling for buyers resulting from an allocative inefficiency as products are not necessarily supplied by the most efficient sellers, related empirical and experimental evidence on the basic multi-unit Dutch auction is scarce.[2] Our study contributes to the literature by showing that this auction format is prone to higher prices than predicted by standard theory and is characterized by bid pooling. Furthermore, we set up a theoretical framework to show that these experimental results can be organized by boundedly rational bidding strategies.[3]

In the auction that we analyze, each subject can sell at most one unit and faces the same commonly known costs to produce it. In many environments where inputs are procured through reverse auctions, transparent information about costs seems to be a plausible assumption. For example, cost structures are transparent for industry sectors such as for raw material as well as standardized and upstream products. These products are typically characterized by a relatively small value added or sunk R&D costs. More generally, as Haruvy and Katok (2013) point out, bidder-specific attributes may be well known in markets where a relative small number of bidders repeatedly interact with each other.

[1]In these auction-like settings, all buyers pay the same price but they may incur different opportunity costs depending on the point in time they decide to purchase their tickets. Another application is the problem of how to deal with stranded passengers when flights are overbooked. The use of auctions as a solution to the airline-overbooking problem was suggested by Simon (1968) (see also Simon, 1994).

[2]There are several contributions on Dutch multi-unit auctions (see McCabe et al., 1990; Katok and Roth, 2004; Goeree et al., 2006; Kwasnica and Sherstyuk, 2007) which all analyze different environments (heterogeneous units, externalities between units, etc.) compared to the present study.

[3]Sherstyuk (1999, 2002) analyzes multi-unit English auctions and shows that collusion (without bid pooling) occurs if bidders can match their offers. This is in accordance with standard theoretical predictions as she allows for bid matching which means that competitors may immediately match any deviating bid rendering deviation unprofitable. Contrary to that, our observations in a multi-unit Dutch auction cannot be explained by standard theory.

In every round of our game, a buyer starts the price clock at some low price and the selling price is increased continuously. At any price, four bidders can decide to sell their product or remain in the auction and wait for a higher price. As soon as the buyer has obtained the desired number of objects, the auction ends. In this setup, standard economic theory predicts that all bidders accept to sell the good either at a price equal to costs or at the start price if the start price is above costs. Therefore, our design shares important features of Bertrand-style competition. At the same time, bidders prefer higher bids to increase revenues.

Indeed, we find that bids in our experiment are substantially above the Nash equilibrium price with rational bidders and only gradually approach it over the periods. Bid pooling is a predominant pattern—the majority of bids within an auction occur immediately after the first supplier submitted a bid. Also, bidders seem to focus on reference prices equal to the highest successful bids in the previous auction when they decide about accepting the clock price.

We propose a framework that integrates bounded rationality into the derivation of bidding functions by assuming that bids are heterogeneous with respect to their strategic sophistication. In this framework, we distinguish between *myopic* bids consisting of a simple backward-looking heuristic and *sophisticated* bids where agents anticipate the behavior of others and choose their optimal bids according to their expectations but may make mistakes. This approach can organize our experimental observations on the aggregate, suggesting that on the individual level, about half of the bids are sophisticated whereas the other half are myopic.

Bidding behavior in descending multi-unit auctions has been theoretically analyzed in Bulow and Klemperer (1994), Martínez-Pardina and Romeu (2011), as well as Gretschko et al. (2014) . These articles show that any symmetric equilibrium in this auction format is inefficient as bid pooling occurs under standard assumptions with bidder heterogeneity. In the present setting, however, bidders are homogeneous with respect to their costs so that simultaneous bidding should only be observed at the start price. Yet, under the assumption that subjects are boundedly rational, bid pooling is predicted at prices substantially above the start price. Moreover, in our framework, sophisticated bidders maximize their profit if they just preempt bid pooling, i.e., accept the price clock just before all others do. As these subjects will not enter the auction immediately if the highest price in the previous period is sufficiently far away from the minimum price but aim at undercutting it gradually, prices will decline over time and eventually converge to the starting price.

Our experimental design can be linked to other classes of experiments. With its equilibrium of placing a bid equal to the lowest possible price, our design is related to investigations of Bertrand competition with homogenous products in which participants have to decide about the price they charge for the goods. Experimental studies in this area have found that the realized prices typically range above the equilibrium prices of rational profit-maximizing players—at least when the number of competitors is sufficiently low (see, for example, Dufwenberg and Gneezy, 2000, Muren and Pyddoke, 2006, Hinloopen and Soetevent, 2008, and Fonseca and Normann, 2012).

At the same time, there are important design aspects that distinguish our setting from Bertrand competition. First and foremost, the sequence of the price increases in our setting makes bidding a dynamic decision problem from the subjects' perspective. Due to the continuous price increase in the course of each auction, bidders in our experiment can react to the bids of others and bids are potentially placed in sequences rather than simultaneously as it is the case in Bertrand markets. Related to this, competitors in Bertrand markets face a constant demand whereas in our case, the demand for the objects may shrink during an auction and bidders may adapt their strategies to changes in the demand. Second, competitors in Bertrand experiments typically pick one price from an interval of possible prices whereas bidders in our setting face a repeated binary decision. This is due to the fact that for each price step, they have to choose whether to enter the auction or to wait. Third, aiming at higher prices in Bertrand markets is rather risky for the firms. If one firm is underbid by a competitor, it earns zero profits as the competitor serves the whole market. On the contrary, if one supplier in our setting has not yet entered the auction at the time a competitor has already sold the object, there is still the chance to make positive profits when the object is sold later in the auction.

In addition, our design shares important features with centipede games due to its complete information structure and the sequential decision-making of subjects (McKelvey and Palfrey, 1992). In the standard version of the centipede game, two players repeatedly choose whether to exit the game or to pass the decision to the other player. Whereas total payoffs increase with the number of times the decision is passed on, each player has the incentive to exit at every stage. Under standard assumptions, rational players immediately exit. Yet, most experimental variations of this game find strong deviations from the Nash equilibrium predictions, with a substantial probability that players pass the decision to others even in later stages of the game (see, for example, Rapoport et al., 2003, Murphy et al., 2006, Palacios-Huerta and Volij, 2009, Levitt et al., 2011, and the references cited therein).[4] Our setup differs from "classic" centipede games in several important aspects. First, in the present setting, more than two players interact with each other who decide simultaneously at each stage whether or not to sell their items. Second, as our game does not end once a single player has moved, the relation of bidders and goods in our design does not produce one winner and $n-1$ losers (as in centipede games), but $n-1$ winners and one loser. Third, depending on the bids placed in the auctions, winners in our auction may obtain substantially different payoffs.

Finally, our experiment is related to clock games as introduced by Brunnermeier and Morgan (2010). In a clock game, several players have to decide when to sell an asset whose value increases exponentially over time. If a player does not sell the asset, he receives an "end of game" payoff that is stochastically determined and relevant for all players. At some point in time, each player receives a private signal that the value of the asset exceeds the "end of game" payoff. If the decisions to sell the asset are observable,

[4]Moreover, in a recent article, Cox and James (2012) compare Dutch auctions and centipede games with private information about payoffs and highlight that the mode of presentation (clock or tree structure) has a decisive impact on behavior.

the model predicts that players initially wait but that there is "herding" after the first asset has been sold. This pattern is confirmed in experimental tests. Similar to the theoretical papers on descending multi-unit auctions, the pooling of bids is explained by strictly rational behavior.

The paper proceeds as follows. First, we describe the experimental design and derive the theoretical prediction given standard assumptions (*Section 2*). We then proceed with reporting our experimental results (*Section 3*) and suggest a behavioral model of the interaction between sophisticated and myopic bidding strategies to organize the observed patterns in our data (*Section 4*). The last section discusses our findings and concludes.

7.2 EXPERIMENTAL DESIGN AND THEORETICAL PREDICTIONS

Our experimental auction was conducted as follows. $K = 4$ subjects in the role of sellers were endowed with one item they could sell in a reverse Dutch auction. A maximum of $G = 3$ goods could be sold in a single auction (the items were sold to the experimenter; there was no human buyer involved). The selling price started at 20 Experimental Currency Units (ECU) and was increased by 5 ECU every five seconds as long as less than three sellers had sold their good at any of the previous price steps. A participant who wanted to sell at the current price could do so by clicking a 'Sell' button on the screen. If the number of bids exceeded the number of requested items, sellers were randomly chosen.[5] During the auction, active bidders were informed at each price step how many items were still to be sold. The auction automatically stopped at a price cap of 100 ECU. After the auction ended, all bidders were given full information about the prices paid for each of the three objects and about unsuccessful bids (if any). The experiment included 20 repetitions of this reverse Dutch auction. We implemented a

[5] This feature is similar to some treatments of the oral auction studies by Sherstyuk (1999) and (2002) where, after observing the bid of a competitor at a given price, a bidder could decide to match it. The possibility of matching bids strongly facilitated collusion. In our design, a subject can realize only after a given bid price that a competitor has entered the auction, thereby running the risk that all items have already been sold.

partner matching where four subjects were randomly assigned to each other prior to the experiment and interacted with the same participants throughout the 20 rounds. The partner matching was chosen for two reasons. First, the real-world motivation for our study—procurement auctions with transparent cost structures of the competitors—applies in particular for markets where a relatively small number of bidders repeatedly interact with each other. Second, as we assume that bounded rationality may have an impact on bidding patterns and realized prices, a partners matching should be a challenging environment to test this conjecture as the repeated interaction with the same competitors should provide better opportunities to gain experience and to learn optimal behavior than in a strangers matching. This is true in particular because in our setting, bidders receive full information about all bids in every round.[6] Finally, to reduce complexity for the participants, costs for the item were normalized to zero which was public knowledge.

The experimental sessions were run in the Cologne Laboratory for Economic Research. Subjects were recruited with Greiner's online recruitment system ORSEE (Greiner, 2004). The experimental software was programmed with z-Tree (Fischbacher, 2007). We conducted three sessions with a total of 88 subjects yielding 22 statistically independent bidder groups. The majority of the participants were students with a major in economics, business administration, or related fields. Subjects arrived at the laboratory, were randomly assigned to a cubicle, and received written instructions.[7] After the experiment, subjects were privately paid out and left the laboratory. Experimental payoffs were converted at a rate of 100 ECU = 1€. The average payoff was 14.63€ (including a show-up fee of 7.50€[8])

[6] At the same time, we acknowledge that repeated game effects may arise in our setting. Therefore, an interesting extension of our experiment would be to investigate the bidding patterns and market prices that emerge if subjects are matched with new competitors in every period.

[7] Instructions translated from German can be found in *Appendix A*.

[8] Given that the expected payoff resulting from equilibrium play accounted for only $3/4 \times 0.2€ \times 20 = 3.00€$, the high show-up fee was chosen to ensure that the average payoff in our experiment would not be significantly below the typical level.

with a standard deviation of 2.44€. Each session lasted approximately one and a half hours.

In the above setup where the costs of all sellers are fixed to zero and known to all bidders, standard economic theory predicts Bertrand-style competition among bidders who try to undercut each other to win the contract. As a result, we should observe that under standard behavioral assumptions, every bidder immediately presses the 'Sell' button at the start price of 20 ECU.[9]

7.3 EXPERIMENTAL RESULTS

In a first step, we present the aggregate outcomes of the experimental auctions. As will become clear, bidding behavior significantly departs from standard predictions.

Calculated over all auctions, the average bid accounts for 48.0 ECU and is thus substantially higher than the Nash equilibrium bid under standard assumptions. *Figure 7.1* displays the time trend for average bids and the corresponding standard deviations. Initially, the average bid equals 65.6 ECU but bids decline over time. This negative time trend is confirmed if we calculate Spearman rank correlation coefficients between average bids and the number of rounds for each bidder group and perform a Sign-test for the 22 statistically independent correlation coefficients ($p < 0.001$, two-sided Sign-test). In the last rounds, average bids approach the equilibrium bid of 20 ECU. Yet, bids are still higher than the equilibrium bids for a substantial share of bidder groups. Even in the final round, the lowest bid is above 20 ECU in 7 out of 22 bidder groups.

We can thus state the first result:

Result 1. *Average bids are substantially higher than the start price and gradually approach it over time.*

Table 7.1 confirms that prices decrease only slowly throughout our experiment. Here we compare the dynamics of the highest prices achieved

[9]We choose a start price of 20 ECU in order to get a unique equilibrium under standard assumptions.

Figure 7.1: Average bid prices over all auctions (in ecu)

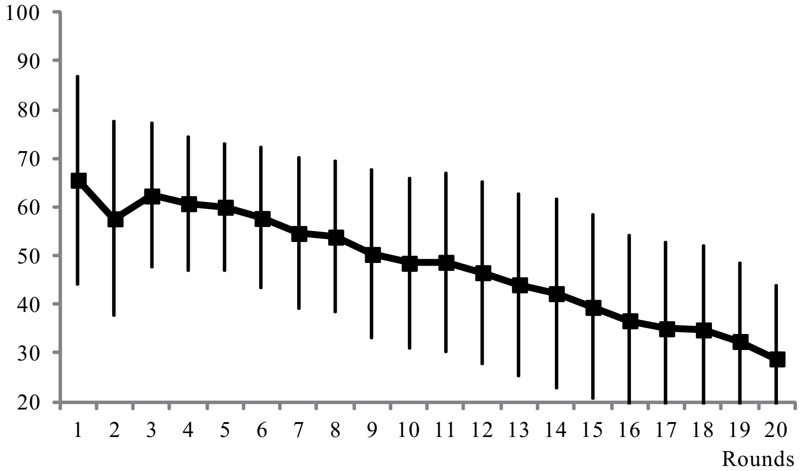

The figure displays the time trend of the mean bid price at which sellers entered an auction and the corresponding standard deviations.

within a bidder group. Realized prizes are rather stable as calculated over all 20 rounds, the highest price in an auction is identical to the previous period in more than 40% of the cases. Also, if the highest selling price declines from one round to the next, it only decreases by 5 ECU (= 1 price step) in the large majority of the cases. Moreover, auctions where prices drop by more than two price steps are rare. In some 11% of the cases, the highest selling price even increases in the present round. Finally, the patterns concerning price dynamics do not seem to be vary strongly across time intervals.[10]

If we consider bidding dynamics within one particular auction by calculating the bid spread (defined as the difference in price steps between the highest and the lowest realized selling price), we find substantial evidence for bid pooling. From rounds 3–15, the median bid spread equals exactly one price step (i.e., 5 ECU). Calculated over all periods, more than 85% of all “subsequent” bids—i.e., bids that follow the initial bid in a particular

[10]The exception is the final time interval (periods 16–20) where the share of auctions where the highest prices stay constant rises by about 20 percentage points which is due to the fact that many auctions have already reached the minimum price of 20 by then.

auction—are exactly one price step above the price of the first bidder. From round 16 on, the median price spread equals zero as the majority of auctions reach or approach the minimum price of 20 ECU. *Table 7.2* separately lists the percentage shares of auctions with bid spreads of a given number of price steps for each five-period interval. Indeed, if we calculate Sign-tests separately for each time interval, we find that median bid spreads across bidder groups are always significantly smaller than 10 ECU or two bid steps (all p-values are smaller than $p < 0.001$). These results further illustrate the pattern that in the large majority of cases, the last successful bidder in an auction enters at most one price step after the first bidder.

We therefore arrive at our second result.

Result 2. *Once the first object is sold, the majority of the bids are exactly one price step above the initial bid (bid pooling).*

In the next step, we aim at gaining more insights into the path dependency of bidding behavior. Our conjecture is that bidders focus on the outcomes from the past auction when they decide at which price they enter a given auction. The highest price achieved in the preceding auction is the obvious candidate for the reference price as it reflects the bidders' common

Table 7.1: Distribution of changes in the highest price achieved in two consecutive auctions

Price steps	< 0	0	1	2	> 2
All rounds	10.8%	44.5%	35.6%	4.5%	4.5%
2–5	19.3%	36.4%	28.4%	6.8%	9.1%
6–10	9.1%	40.9%	42.7%	5.5%	1.8%
11–15	11.8%	40.0%	41.8%	3.6%	2.7%
16–20	4.5%	59.1%	28.2%	2.7%	5.5%

The table lists the percentage shares of auctions in which the highest achieved price in round $t+1$ is either smaller than (by 1, 2 or >2 price steps), equal to (0 price steps), or larger than ($<$ 0 price steps) the highest achieved price in round t. One price step equals 5 ECU.

Table 7.2: Distribution of bid spreads over all auctions

Price steps	0	1	2	> 2
All rounds	30.5%	60.7%	4.1%	4.8%
1–5	10.9%	64.5%	10.0%	14.5%
6–10	24.5%	73.6%	1.8%	0.0%
11–15	27.3%	66.4%	4.5%	1.8%
16–20	59.1%	38.2%	0.0%	2.7%

The table lists the percentage shares of auctions with bid spreads of 0, 1, 2, and > 2. The bid spread is defined as the difference of price steps between the lowest and the highest successful bid; one price step equals 5 ECU.

goal to maximize revenues.[11] To test this conjecture, we calculate simple Tobit models with the individual bid as the dependent variable to account for the fact that bids in our setting cannot be lower than 20 ECU, clustering standard errors on the level of the experimental bidder group.

The results are listed in *Table 7.3*. Model 1 includes the number of periods and the highest winning bid in the previous period. Here, the variable for periods enters with a negative and significant sign which is in line with the general downward trend in prices. Importantly, the coefficient of *Highest price (t-1)* is positive and highly significant, suggesting that it is indeed an important focal point for bidders. This effect is robust also in alternative specifications. In Model 2, we additionally include the second-highest and the third-highest bid from the previous period, which, however, are both insignificant. In particular, the first winning bid of the previous period (or the lowest price one or more bidders accepted) which might be interpreted as a signal for the competitiveness of the bidders in the auction (variable *Third highest price (t-1)*) does not have an additional impact on bidding decisions. In Model 3, we add a dummy variable equal to one if a particular bidder was not able to sell his good in period $t-1$ and therefore ended up with zero profits for this auction (*Did not sell (t-1)*). The coefficient of this dummy variable is negative and significant,

[11] A somewhat related argument is made by Suetens and Potters (2007) to explain dynamic behavioral patterns in Bertrand settings. When information about previous behavior is provided, subjects might imitate the best performer in the last round.

Table 7.3: Determinants of individual bids

Model No.	1	2	3
Dependent variable	*Price bid (t)*	*Price bid (t)*	*Price bid (t)*
Period	-0.257***	-0.231***	-0.230***
	(0.063)	(0.073)	(0.073)
Highest price (t-1)	0.980***	1.218***	1.218***
	(0.032)	(0.264)	(0.264)
Second highest price (t-1)		-0.214	-0.213
		(0.281)	(0.281)
Third highest price (t-1)		-0.025	-0.026
		(0.097)	(0.098)
Did not sell (t-1)			-0.853***
			(0.326)
Constant	-1.095	-1.764	-1.561
	(2.089)	(2.094)	(2.102)
Observations	1,532	1,532	1,532
Log-Pseudolikelihood	-4,645	-4,641	-4,640

Tobit models are calculated to account for the fact that bids cannot be smaller than 20 ECU. Robust standard errors clustered on experimental matching groups are listed in parentheses. *** indicates a significance level of $p < 0.01$. The variable *Period* denotes the number of the particular period.

indicating more aggressive subsequent bidding behavior after a bidder was not successful. At the same time, the effect of *Highest price (t-1)* remains significant.

Hence, our third result is the following.

Result 3. *The highest winning bid in the previous auction is significantly correlated with individual bids in a given auction.*

We can thus conclude from our experimental results that (i) bids deviate substantially from standard predictions, (ii) bid pooling is pervasive, and (iii) the highest winning bid from the preceding period appears to be an important focal point for bidding behavior in the present period. In the next section, we model the impact of players' strategic sophistication in our setting.

7.4 STRATEGIC SOPHISTICATION

In this section, we argue that our experimental results can be organized by the notion that bids differ in the degree of sophistication and can largely be categorized into two classes of bidding strategies.

Myopic and sophisticated bidding

In the following, we develop a model of bidding behavior to organize the main findings from our experiment. In particular, bidding may be characterized by (i) a myopic bidding strategy (denoted by subscript m) or (ii) a sophisticated strategy (denoted by subscript s). Whereas the myopic bidding strategy is a simple heuristic which requires less intellectual effort, the sophisticated bidding strategy anticipates the existence of myopic and other sophisticated bids and aims at playing a best response to a composition of different bids.

In the first auction, bidders have no anchor to which they can adjust their bids. Under the assumption that bidders who follow a myopic bidding strategy randomize over the $\{20, 25, \ldots, 100\}$ interval and given that bidders using a sophisticated strategy play best responses, prices above the minimum price can be explained.[12]

From the second round on, we assume that bidders take the auction outcomes realized in the previous round into account when deciding about their current bid. In particular, we hypothesize that even a myopic bidding strategy does not completely ignore the history of the game but takes the accepted bids as reference points. As shown in the last section, the highest price achieved in the preceding auction is positively related to bids in the present auction so that this price seems to serve as a reference point for bidding behavior.

Applying this idea to our setting, we assume that the myopic strategy consists of bidding exactly the highest realized price for which the last object was sold in the last auction. This is naïve in the sense that the strategy does not take into account that other bidders have an incentive to enter

[12]We formalize this idea in *Appendix 7.6.2*.

the current auction earlier. By focussing on the last object sold, a bidder pursuing a myopic bidding strategy would furthermore, once an object is sold in the current auction, react to the resulting reduction in demand by accepting the next possible price to still realize a sale. As a result, we can summarize the myopic bidding strategy $b_m\left(g, \bar{b}_{t-1}\right)$ from the second auction onward as follows

$$b_m\left(g, \bar{b}_{t-1}\right)\Big|_{t\geq 2} = \begin{cases} \bar{b}_{t-1} & \text{if } g = 3 \\ 20 & \text{if } g < 3. \end{cases} \tag{220}$$

Here $g \in \{1, 2, 3\}$ denotes the number of objects the buyer still wants to purchase and $\bar{b}_{t-1}$ denotes the highest price achieved in the previous auction.

The sophisticated bidding strategy depends on the share of myopic and sophisticated bids. Let $x \in [0, 1]$ denote the share of sophisticated bids. Remember that bidders have full information about winning bids at the end of the auction. A bidder following the sophisticated bidding strategy assumes that the myopic bids are equal to $\bar{b}_{t-1}$. Thus, in order to ensure that he can sell his good, this bidder should enter the auction at a price lower than $\bar{b}_{t-1}$. The best answer for a share x of sophisticated bids is to bid one price step below $\bar{b}_{t-1}$ if

$$\begin{aligned} \bar{b}_{t-1} - 10 &\overset{!}{\leq} \left(1 - x^3 + \frac{3}{4}x^3\right)\left(\bar{b}_{t-1} - 5\right) \\ \Leftrightarrow \qquad 5 &\overset{!}{\geq} \frac{1}{4}x^3\left(\bar{b}_{t-1} - 5\right). \end{aligned} \tag{221}$$

The left-hand side of inequality (221) is the additional reduction in profits if the bidder were to bid two price steps ahead of $\bar{b}_{t-1}$ instead of one. The right-hand side is the potential loss when bidding $\bar{b}_{t-1}-5$. A loss occurs if all other bids in the group are sophisticated (which happens with probability x^3) and the bidder does not win the resulting lottery (which happens with probability 1/4).[13] Finally, if one or more units have already been sold, the bidding strategy consists of accepting the clock price immediately as the

[13] For example, if the share of sophisticated bids is 50%, then it is optimal for a bidder following the sophisticated bidding strategy to enter exactly one price step below $\bar{b}_{t-1}$ for any $\bar{b}_{t-1} \in \{20, ..., 100\}$.

bidders with the myopic strategy do so as well. If the share of sophisticated bids is not too large, i.e., $x \lesssim 0.6$, which holds true in the auctions we observe,[14] the sophisticated bidding strategy is given by

$$b_s\left(g, \bar{b}_{t-1}, x\right)|_{t\geq 2} = \begin{cases} \max\left\{20, \bar{b}_{t-1} - 5\right\} & \text{if } g = 3 \\ 20 & \text{if } g < 3. \end{cases} \tag{222}$$

In the following, we extend the analysis and allow for mistakes. Here, we assume that errors occur only in the sophisticated bidding strategy.[15] This means that, although a bidder would prefer to follow the sophisticated bidding strategy, he might err and place a different bid. Due to fact that in our model myopic bids will be placed at $\bar{b}_{t-1}$, mistakes by sophisticated bidders can only occur downwards, i.e., before the price reaches $\bar{b}_{t-1} - 5$. Let ϵ be the probability of making an error at all. Then, with probability $1 - \epsilon$ the sophisticated bidding function is given by

$$b_s\left(g, \bar{b}_{t-1}, x, \epsilon\right)|_{t\geq 2} = \begin{cases} \max\left\{20, \bar{b}_{t-1} - 5\right\} & \text{if } g = 3 \\ 20 & \text{if } g < 3. \end{cases} \tag{223}$$

With probability ϵ, the bidder makes a mistake and the bidding function is given by

$$b_s\left(g, \bar{b}_{t-1}, x, \epsilon\right)|_{t\geq 2} = \begin{cases} \max\left\{20, \bar{b}_{t-1} - 5F\left(a\right)a\right\} & \text{if } g = 3 \\ 20 & \text{if } g < 3. \end{cases} \tag{224}$$

Here F denotes the probability to tremble by a steps where $a \in \mathbb{N}_{\geq 2}$.[16]

The bidding behavior characterized by expressions (220) and (223)–(224) implies that in presence of sophisticated bids, realized prices will decline over time as sophisticated players underbid the prices achieved in previous auctions.[17]

[14] For higher values of x, there may be unraveling to the minimum price.

[15] This assumption seems plausible, because bidders need to perform elaborate computations to follow the sophisticated strategy. On the contrary, following the myopic strategy, a bidder simply copies a price in the next period.

[16] Assuming that the mistakes are distributed according to the Poisson distribution, one can even derive a trembling-hand perfect equilibrium (Selten, 1975). This can be found in *Appendix 7.6.3*.

[17] The term "realized price" here refers to the initial price at which one or more bidders enter the auction. Note that whether the prices at which subsequent bids are submitted

Observation 1. *Prices decrease from the second period on if the share of sophisticated bids is larger than zero.*

Moreover, as long as the bidders following a sophisticated bidding strategy do not err by to many price steps, we expect that initial bids do not decrease strongly in subsequent rounds.

Observation 2. *Prices decline slowly and gradually converge toward the minimum price over time.*

These two predictions of our model match the qualitative patterns found in the data.

The idea that the interaction between sophisticated and naïve strategies influences outcomes has been applied in related settings. For example, in the context of Bertrand games, Dufwenberg and Gneezy (2000) show how the presence of "noise bidders" might lead to prices above the Nash equilibrium under standard assumptions. Also, the basic intuition that bids can be attributed to one of the two strategies in our auction setting is similar to the idea behind level-k models of bounded rationality (see Crawford et al., 2013, for an overview of the literature) and related approaches (see, e.g., the cognitive hierarchy model by Camerer et al., 2004). These models assume that players maximize payoffs on the basis of simplified beliefs over other players' actions. Optimal choices are derived by iterated best response, with more sophisticated players applying a higher number of steps of iterative reasoning. Moreover, similar to our model, optimal actions are determined anticipating the behavior of less sophisticated players.[18]

---

are equal to or lower than the highest price achieved in the previous round depends on the initial bid in the present round.

[18]Many studies (see, for example, Duffy and Nagel, 1997, Ho et al., 1998, Bosch-Domènech et al., 2002, Kocher and Sutter, 2005, Costa-Gomes and Crawford, 2006 and the references cited therein) have provided empirical support for level-k models. A robust phenomenon is that the majority of choices is in line with one or two steps of iterated reasoning; the share of more sophisticated players is typically small. For recent discussions on the general validity of the level-k approach, see, e.g., Penczynski (2011) and Georganas et al. (2015). Investigations on the empirical relevance of level-k thinking in static auction settings yielded mixed results (see Crawford and Iriberri, 2007, Ivanov et al., 2010, Georganas, 2011, and Kirchkamp and Reiss, 2011). More generally, other studies showed the importance of boundedly rational bidding strategies in auctions (see,

Empirical fit of the model

In the next step, we have to check whether our proposed model of heterogeneous bidding behavior makes reasonable predictions for the distribution of bids and errors across auctions. To do so, we first consider every *initial* bid from auctions $t \geq 2$. For these auctions and bids, we can unambiguously determine whether the particular bid is sophisticated or myopic. The sophisticated, mistake-afflicted, bidding strategy consists of entering one or more steps ahead of the myopic bids (the highest price achieved in the preceding auction). Contrary to that, follow-up bids in auctions $t \geq 2$ where the initial bid was more than one price step below $\bar{b}_{t-1}$ cannot be attributed to the sophisticated or the myopic strategy.[19] For instance, in an auction where the initial bid is placed two price steps below the highest bid in the previous auction, the bids that immediately follow the first bid can be either sophisticated or myopic. We therefore choose an indirect approach and implicitly derive the underlying distribution of bid types.

To do so, we can derive the ex-ante probabilities of observing between one and four bids (as four bidders interact) at a given number of steps below the highest previous bid in each auction. As stated above, first bids below the highest bid in the previous round are attributed to the sophisticated bidding strategy. *Table 7.5* in the appendix lists all corresponding ex-ante probabilities that n players submit either myopic or sophisticated bids and also gives the number of observations per case from period $t \geq 2$ for the 325 out of the 418 auctions (77.8%) which can be used for the following analysis.[20]

for example, Cooper and Fang, 2008, Shachat and Wei, 2012, and Kirchkamp and Reiß, 2014).

[19]The exception are auctions where the initial bid is placed one price step below the highest price achieved in the previous period. Bidders who do not enter at this price are sure to follow the myopic strategy according to our definition.

[20]We cannot use data from auctions in the first round because it is not possible here to separate sophisticated and myopic bids from each other. Due to the assumption that without a bidding history, the myopic strategy consists of randomizing over the strategy space, every observed bid in the first auction could in principle be sophisticated or myopic. Moreover, initial bids exceeding the highest previous level in any auction in periods $t \geq 2$ or initial bids at $\bar{b}_{t-1}$ by less than four bidders are not captured by our approach and therefore excluded (altogether 43 auctions). Finally, we have to exclude 50

We use the input from *Table 7.5* to specify a system of equations that define the distance between the ex-ante and the observed probabilities of n initial bids of myopic and sophisticated type in an auction. We capture the errors made by sophisticated types by assuming that these are distributed according to a discrete probability mass function. We then assume that the error is Poisson-distributed, $P(y;\lambda) = \lambda^y e^{-\lambda}/y!$, where λ specifies where the probability mass function has its peak, i.e., which erroneous step is most likely. By comparing the ex-ante probabilities with the observed number of auctions for a given number of initial bids, we can calculate the share of sophisticated bids that best fits this system and λ.[21]

Our results are summarized in *Table 7.4*. First of all, we observe that the shares of myopic and sophisticated players are roughly equal. The large share of myopic bids (47.1% calculated over auctions 2–20) follows from the fact that the majority of initial bids are submitted at one price step below the highest previously achieved price. This finding illustrates that a large share of bids follow the heuristic of remaining in the auction until the highest price of the previous period is reached. If we take the dynamics of the game into account by breaking the calculation down to five-period intervals, we observe that the shares of sophisticated and myopic bids constantly range at around one half throughout the rounds. Only in the last five rounds of the game, the share of sophisticated bids increases and accounts for some 65%. This suggests a small shift towards more sophisticated strategies at the end of the game.

Concerning bidders' proneness to making errors, our estimation suggests that the bidders are rather accurate in their bidding behavior. Given that a player bids less than $\bar{b}_{t-1}$, the overall probability to bid exactly one price step below $\bar{b}_{t-1}$ accounts for 78%. Moreover, if bidders make mistakes, it is very likely that the mistakes are only small.[22] Conditional on underbidding

auctions from the analysis where goods are sold for prices near or at the Nash equilibrium under standard assumptions. In these cases, a distinction between sophisticated and myopic bids and equilibrium play is no longer possible.

[21] The exact approach, functions and derivation can be found in *Appendix D*.

[22] We acknowledge that with our estimation strategy, we may understate the propensity of making mistakes. As prices successively approach the Nash equilibrium bid for

Table 7.4: Estimated distribution of bid types and errors

Periods	# of auctions	Myopic	Sophisticated	λ
2–20	325	47.1	52.9	0.26
2–5	73	50.5	49.5	0.36
6–10	100	47.3	52.7	0.19
11–15	92	49.9	50.1	0.26
16–20	60	34.9	65.1	0.21

The table lists the estimated distribution of bid types and the parameter λ of the Poisson distribution for auctions 2–20 and for four separate time intervals.

by more than one price step, the probability of erring by exactly one bid step (i.e. bidding $\bar{b}_{t-1} - 10$) is 88%. If we consider time dynamics, we find that bidders tend to make mistakes more often in the beginning of the experiment (periods 2–5 where the corresponding probability is around 30%) than in later rounds of the game—in the last time interval (rounds 16–20), the probability decreases to some 19%.

All in all, the estimated distribution of sophisticated and myopic bids corroborates our observation from the descriptive statistics that the convergence to minimum bids is driven by gradual responses to the bidding history and not by some kind of "eureka" learning such that some bidders switch to Nash equilibrium play from some period on.

7.5 DISCUSSION AND CONCLUSION

We have conducted a procurement experiment in which subjects acted as sellers in repeated reverse multi-unit Dutch auctions. Empirical and experimental evidence on this auction format is scarce despite its practical importance and despite the (potentially) detrimental effects of bid pooling for buyers. In our setting, we observe that bids are substantially above the

rational players, bidders have less and less possibility to underbid by may price steps. Yet our results do not change substantially if we restrict our analysis to auctions in which the highest previous bid was 40 ECU or higher. For these auctions, sophisticated bidders still can erroneously underbid by three steps. Repeating our estimation for the restricted sample, 52.1% (47.9%) of the bids are classified as sophisticated (myopic); of the sophisticated bids, 75.7% are estimated to be exactly one price step below the highest bid from the previous round.

Nash equilibrium prediction under standard assumptions and that bid pooling is a predominant pattern. Moreover, bidding behavior reacts strongly to the highest prices achieved in the previous auction. We explain these results with a bidding model that distinguishes between sophisticated and myopic bidding strategies.

Our model can capture the qualitative patterns found in the data. Average prices achieved in the auctions are substantially above the Nash equilibrium price in the beginning of the game. In nearly all auctions, the variance of bids is only small. The difference between most bids is only one price step and prices only slowly converge to the equilibrium level for rational, profit-maximizing bidders. Moreover, based on our estimates roughly half of the bids seem to result from myopic bidding whereas the other half applies a more sophisticated strategy.

Hence, our experiment indicates the importance of boundedly rational bidding strategies. At the same time, learning processes among the bidders in our setting seem to be limited. As our estimation of the underlying bidder behavior reveals, the shares of myopic and sophisticated strategies are roughly constant over the rounds. Only at the end of the game, we observe a moderate shift toward more sophisticated strategies. Yet, recent studies on level-k thinking suggest the importance of learning for strategic sophistication and, in particular, that subjects are heterogenous concerning these learning processes. Gill and Prowse (ming) find a systematic connection between the cognitive ability of subjects and the levels of reasoning exhibited in a repeated beauty-contest game. Moreover, subjects with a high cognitive ability are also able to learn from the strategic interaction with other subjects as their observed level of reasoning responds to the level of reasoning of the subjects they are matched with. Agranov et al. (2012) observe in one-shot beauty contest games that the degree of sophistication some players show crucially depends on exogenously manipulated beliefs on how other subjects play. In the light of these results, learning processes in general and especially the way how sophisticated players respond to the (perceived or observed) sophistication of others would be an interesting next step to investigate in auction settings.[23]

An important related question is how this insight can be utilized from a market-design perspective. In the procurement setting we explore, bids are substantially above minimum bids and converge only slowly over time as a substantial part of the bidders adhere to the simplest possible heuristic—bidding the highest observed price in the previous auction. In our case, the auctioneer would substantially benefit from higher levels of sophistication on the aggregate as this would countervail the observed inflation of prices. Any institutional change fostering the strategic sophistication of bidders that prevents simple bidding heuristics would be highly desirable. For example, the results by Haruvy and Katok (2013) suggest that the transparency of different procurement auction formats may influence the auctioneer's payoffs. A deeper understanding of the interaction between bidder sophistication and the market environment is a promising avenue for further research.

[23]In classic level-k models, subjects are assumed to stick to the strategy associated with their types. In a recent approach to modeling level-k-thinking, Ho and Su (2013) assume that players may change their strategy after observing unexpected behavior of others. In their model, players are in principal capable of playing according to any level-k-thinking type but choose their behavior as to maximize payoffs given their beliefs about the distribution of the thinking types of the other players.

Table 7.5: Ex-ante probabilities for specific bids and actual observations over auctions 2–20

Bid type	# of bidders	Probabilities	# of observations $N_k^{\{n\}}$
Myopic	4	$(1-x)^4$	2
Sophisticated	1	$4P(0;\lambda)x(1-x)^3$	73
	2	$6\left(P(0;\lambda)x\right)^2(1-x)^2$	55
	3	$4\left(P(0;\lambda)x\right)^3(1-x)$	51
	4	$\left(P(0;\lambda)x\right)^4$	18
Sophisticated with errors	1	$4\left(1-P(0;\lambda)x\right)(1-x)^3$ $+12\left(1-P(0;\lambda)x\right)(1-x)^2P(0;\lambda)x$ $+12\left(1-P(0;\lambda)x\right)(1-x)\left(P(0;\lambda)x\right)^2$ $+4\left(1-P(0;\lambda)x\right)\left(P(0;\lambda)x\right)^3$	113
	2	$6\left(1-P(0;\lambda)x\right)^2(1-x)^2$ $+12\left(1-P(0;\lambda)x\right)^2(1-x)P(0;\lambda)x$ $+6\left(1-P(0;\lambda)x\right)^2\left(P(0;\lambda)x\right)^2$	9
	3	$4\left(1-P(0;\lambda)x\right)^3(1-x)$ $+4\left(1-P(0;\lambda)x\right)^3P(0;\lambda)x$	4
	4	$\left(1-P(0;\lambda)x\right)^4$	0

In the third column, the table lists the ex-ante probabilities that $n \in \{1,2,3,4\}$ bidders have submitted a myopic, sophisticated or erroneous sophisticated bid in an auction. The last column lists the corresponding, observed number of experimental auctions for these bids. We cannot capture bidding behavior by subjects in auctions in which 1, 2 or 3 first bids were placed at $\bar{b}_{t-1}$ (12, seven, and five cases) because the bidders who have not entered by then are not classified by our approach.

7.6 APPENDIX

7.6.1 INSTRUCTIONS

Below we include a translation from German of the instructions that we used in the experiment.

Welcome to our experiment! In this experiment you can earn money. How much you will get depends on your own decisions and on the decisions of the other participants. **From now on please do not communicate with other participants. If there are any questions, please raise your hand!** We will come to you and answer your question. If you break this rule, we will have to exclude you from this experiment and all payments.

In this experiment ECU is used as the currency. **At the end of the experiment, your ECU-payoffs will be converted to Euro and paid out in cash.** The conversion rate is **100 ECU = 1 Euro**.

Procedures

The experiment consists of **20 auctions in total**. In each auction you have the opportunity **to sell a fictitious good**.

At the beginning of the experiment **four sellers** are randomly chosen and randomly assigned to each other. These sellers **will interact in each of the 20 auctions**.

At most **three goods** are sold in each auction. Thus, not every seller will be able to sell her good. An auction proceeds as follows:

The price starts at 20 ECU. The sellers then have **5 seconds to decide whether they want to sell their goods at this price** by clicking on the **button "Sell the good at this price"**. The remaining time is displayed in the upper right corner of the screen.

Sellers who do not want to sell their good at this price do not have to do anything. After 5 seconds, the experiment proceeds automatically.

If all sellers have made their decisions **and not all three goods were sold, the price will be raised by 5 ECU to 25 ECU** and all remaining sellers in the auction have 5 seconds to decide whether they want to sell their goods at the new price.

If not all three goods are sold at this price, the price will again be raised by 5 ECU to 30 ECU and all remaining sellers in the auction decide again.

The price will be raised **by 5 ECU-steps** until **either the three goods are all sold or the price reaches the upper limit of 100 ECU**, without three sellers having sold at this price.

If **more sellers want to sell their goods** at a certain price **than goods are demanded in the auction**, it will be **randomly determined** which seller is allowed to sell her good.

The sellers have no costs. **This means that sellers who sell their goods in the auction receive a payoff equally to their selling prices.** Sellers who did not sell their good in an auction do not receive a payoff from this auction.

After each auction, the sellers will be informed about all prices that were realized in this auction.

Concluding remarks

At the end of the experiment, the **sum of payoffs from all 20 auctions** will be converted into Euro and paid to you. **In addition**, you will receive an amount of **7.50 Euro for your participation** irrespective of the decisions in the experiment.

7.6.2 A MORE ELABORATE APPROACH TO MODELING FIRST-ROUND BIDDING

We consider the simplest case where a bidder following the sophisticated bidding strategy expects the three other bidders to pursue the myopic strategy. In the first auction, as there is no information on bidding behavior from earlier periods available, a sophisticated bidder optimizes his bidding behavior under the assumption that myopic bids follow a discrete uniform distribution on the interval $\{20, 25, \dots, 100\}$.[24] In our case, this assumption reflects that in the very first auction, there is no anchor to which players can adjust their bids. To determine the optimal sophisticated bidding behavior at the current price b, we need to distinguish between three cases depending on the number of products that have already been sold during the auction.

Suppose none of the competing bidders has submitted a bid at price $b-5$. Now let $p := (b-(b-5))/(100-(b-5)) = 5/(105-b)$ be the probability that a competing bidder bids at the current price $b > 20$. Then, the expected payoff of a player with the sophisticated strategy from submitting a bid at b is given by

$$\mathbb{E}\left[\pi_s^{\text{bid}}\right] = b\left(1 - p^3 + \frac{3}{4}p^3\right) = b\left(1 - \frac{p^3}{4}\right). \tag{225}$$

Note that $1-p^3$ is the probability that at most two other competitors bid at the same time which means that the bidder sells his product with certainty. Similarly, p^3 represents the probability that all other bidders simultaneously submit bids in which case the bidder has a winning probability of only $3/4$.

Let $\tilde{p} := (b + 5 - b)/(100 - b) = 5/(100 - b)$ be the probability that the myopic bid is $b+5$. Similar to the case where the bidder submits a bid, the expected payoff from waiting at the current price b then amounts to

[24] Note that this is also a typical assumption in the literature on auction settings that explain bidding behavior with level-k models of bounded rationality (see Crawford and Iriberri, 2007, as well as Kirchkamp and Reiss, 2011).

$$\mathbb{E}\left[\pi_s^{\text{wait}}\right] = (b{+}5)\left(\left(1-\tilde{p}+\frac{1}{2}\tilde{p}\right)3p^2\left(1-p\right)+\left(1-\tilde{p}^2+\frac{2}{3}\tilde{p}^2\right)3p\left(1-p\right)^2 + \left(1-\tilde{p}^3+\frac{3}{4}\tilde{p}^3\right)\left(1-p^3-3p^2\left(1-p\right)-3p\left(1-p\right)^2\right)\right). \quad (226)$$

Solving $\mathbb{E}[\pi_s^{\text{bid}}] = \mathbb{E}[\pi_s^{\text{wait}}]$ gives $b \approx 85.07$ as the (relevant) solution. Hence, if no product has been sold, the sophisticated strategy consists of accepting a clock price of 90 (where $\mathbb{E}[\pi_s^{\text{bid}}] > \mathbb{E}[\pi_s^{\text{wait}}]$).

Suppose next that one of the competing bidders has already sold his product at a price lower than $b-5$. Then, the expected payoff for a bidder following the sophisticated strategy and submitting a bid at b is given by

$$\mathbb{E}\left[\pi_s^{\text{bid}}\right] = b\left(1-p^2+\frac{2}{3}p^2\right) = b\left(1-\frac{p^2}{3}\right).$$

Analogously, waiting for another tick of the price clock yields an expected payoff of

$$\mathbb{E}\left[\pi_s^{\text{wait}}\right] = (b+5) \times\left(\left(1-\tilde{p}+\frac{1}{2}\tilde{p}\right)2p\left(1-p\right)+\left(1-\tilde{p}^2+\frac{2}{3}\tilde{p}^2\right)\left(1-p^2-2p\left(1-p\right)\right)\right).$$

Again, solving $\mathbb{E}[\pi_s^{\text{bid}}] = \mathbb{E}[\pi_s^{\text{wait}}]$ gives $b \approx 76.33$ as the (relevant) solution. Hence, if one product has been sold, the sophisticated bid is 80.

Last, consider the case where only one more product can be sold to the buyer. Then, the expected payoffs amount to

$$\mathbb{E}\left[\pi_s^{\text{bid}}\right] = b\left(1-p+\frac{1}{2}p\right) = b\left(1-\frac{p}{2}\right).$$

Analogously, waiting for another tick of the price clock results in an expected payoff of

$$\mathbb{E}\left[\pi_s^{\text{wait}}\right] = (b+5)\left(1-\tilde{p}+\frac{1}{2}\tilde{p}\right)\left(1-p\right).$$

From $\mathbb{E}[\pi_s^{\text{bid}}] = \mathbb{E}[\pi_s^{\text{bid}}] \Leftrightarrow b = 48.75$, it follows that the sophisticated bid is 50.

The bidding behavior by the two strategies in the first round is the given by

$$b_m(g, \bar{b}_{t-1}, x)\Big|_{t=1} \sim U[20, 100]$$

and

$$b_s(g, \bar{b}_{t-1}, x, \epsilon)\Big|_{t=1} = \begin{cases} 90 & \text{if } g = 3 \\ 80 & \text{if } g = 2 \\ 20 & \text{if } g = 1. \end{cases}$$

This bidding behavior implies that actual bids in the first period of the reverse multi-unit Dutch auction are significantly higher than the one predicted by standard economic theory which equals 20 ECU.

Note that although we only covered the case where one player applies a sophisticated strategy, the above argument also holds for the case where two to four players follow the sophisticated bidding strategy. This is due to the fact that if the bidder applying the sophisticated bidding strategy expects three other players following a sophisticated strategy as well, he is going to enter the auction at a price of 20 for $g \in \{1, 2, 3\}$. As the bidding strategy continuously depends on the distribution of bid types and given the intermediate value theorem, sophisticated bids higher than 20 can be supported for certain ranges of shares x.

7.6.3 TREMBLING-HAND EQUILIBRIUM

Denote the highest winning bid in the previous round by $\bar{b}_{t-1}$. Assume a share of x bidders follow the sophisticated bidding strategy of bidding $\bar{b}_{t-1} - 5$ but might err in doing so. Assume the errors are distributed according to the Poisson distribution $P(\lambda, k)$ where k is the number of steps of deviation and λ the variance/expected value of the distribution. The share $1 - x$ follow the myopic strategy. A best response to this setup might be to bid $\bar{b}_{t-1} - 5$ depending on x and λ. This is the case if the following holds:

$$\begin{aligned}\bar{b}_{t-1} - 10 \leq &(1-(P(\lambda,0)x)^3)(\bar{b}_{t-1}-5) + \frac{3}{4}(P(\lambda,0)x)^3(\bar{b}_{t-1}-5)\\ &+(1-(1-x)^3)\sum_{i=1}^{\frac{\bar{b}_{t-1}}{5}-5} P(\lambda,i)\frac{2}{3}(\bar{b}_{t-1}-5i)\\ &+(1-(1-x)^2)\sum_{i=1}^{\frac{\bar{b}_{t-1}}{5}-5} P(\lambda,i)^2\frac{1}{2}(\bar{b}_{t-1}-5i)\end{aligned}$$

The left-hand side is the reduction in profits if one were to bid $\bar{b}_{t-1} - 10$ instead of $\bar{b}_{t-1} - 5$. The first line of the right-hand side is the probability that there is at least one bidder following the myopic bidding heuristic plus the probability that all other bids are sophisticated and non-erring which leads to a winning probability of 3/4. The second line is the probability that there is one trembling bidder who follows the sophisticated bidding strategy and submits an initial bid strictly smaller than one step below $\bar{b}_{t-1}$ times the resulting winning probability and profit when everyone enters at the next step. The third line follows the same logic given that there are two sophisticated trembling bids at the same bid step strictly smaller than one step below $\bar{b}_{t-1}$. This can be rearranged to

$$\begin{aligned}5 \geq &\frac{1}{4}(P(\lambda,0)x)^3(\bar{b}_{t-1}-5)\\ &-(1-(1-x)^3)\sum_{i=1}^{\frac{\bar{b}_{t-1}}{5}-5} P(\lambda,i)\frac{2}{3}(\bar{b}_{t-1}-5i)\\ &-(1-(1-x)^2)\sum_{i=1}^{\frac{\bar{b}_{t-1}}{5}-5} P(\lambda,i)^2\frac{1}{2}(\bar{b}_{t-1}-5i).\end{aligned}$$

This is satisfied for large ranges of λ and x. In particular, it is fulfilled for all combinations of λ, x, and $\bar{b}_{t-1}$ that we find empirically in Section 7.4. The only exception here are the periods 16–20, where x and λ are such that bidding $\bar{b}_{t-1} - 5$ is not trembling-hand perfect for the whole support of $\bar{b}_{t-1}$. However, in these five periods, the observed values for $\bar{b}_{t-1}$ are low enough that together with $x \approx 0.65$ and $\lambda \approx 0.208$, they form a trembling-hand perfect equilibrium again.

7.6.4 ESTIMATION PROCEDURE

In what follows, we first characterize the estimation procedure (*Subsection 7.6.4*) and then show that it is a valid approach (*Subsections 7.6.4–7.6.4*).

Procedure and Intuition

We can derive the ex-ante probabilities of observing between one and four bids (as four bidders interact) at a given number of steps below the highest previous bid in each auction. *Table 7.5* lists all corresponding ex-ante probabilities that n bids are of either myopic or sophisticated type. By comparing the ex-ante probabilities with the observed number of auctions with a given number of initial bids, we can estimate the distribution of sophisticated and myopics strategies that fits our data best.

We capture the errors in the sophisticated bidding strategy by assuming that these are distributed according to a discrete probability mass function. High price undercuts are rare and our model shows that underbidding by exactly one price step is a trembling hand perfect equilibrium for large parameter spaces. Transferring this idea to our formulation of the error term, we need a distribution function that allows for relatively high probabilities for small errors and vice versa. For this reason, we assume that the errors are Poisson-distributed, $P(y;\lambda) = \lambda^y \mathrm{e}^{-\lambda}/y!$, where y is the number of steps under $\bar{b}_{t-1}$, and estimate λ, the parameter that measures the expected probability of placing a bid based on erroneous beliefs in our case.

Let x denote the share of sophisticated bids. The functions in *Table 7.5* minus the observed shares define the distance between the ex-ante and the observed probabilities of n initial bids of either type. Let us denote these functions $f_j^{\{n\}}$ where n is the number initial bids of either type and $j \in \{m, s, se\}$ for myopic, sophisticated or sophisticated with error. Take the example where exactly $n = 2$ initial bids is classified as type $j = s$ in 55 of all 325 auctions. This means that according to the above definition, we have a function

$$f_2^{\{s\}} = 6\left(P(0;\lambda)x\right)^2 (1-x)^2 - \frac{55}{325}.$$

From the table it becomes clear that we are looking at an over-determined system of equations. To estimate the probabilities given the number of observations per case (i.e., the values for x and λ), we transform the system into a minimization problem by defining a function $f : \mathbb{R}^2 \to \mathbb{R}^9$ with equations $f_j^{\{n\}}$ as components.[25] For a classical solution $\tilde{\mathbf{x}} = \left(\tilde{x}, \tilde{\lambda}\right)$, it holds that

$$\|f(\tilde{\mathbf{x}})\|_2 = 0.$$

In combination with the positive homogeneity of a norm, one can minimize the norm of f under the constraints

$$0 \overset{!}{\leq} x \overset{!}{\leq} 1.$$

With $\mathbf{x} = (x, \lambda)$, this problem can be written as:

$$\min_{\mathbf{x} \in [0,1] \times \mathbb{R}_+} \{\|f(\mathbf{x})\|_2\}$$

Thus, we look for those ex-post probabilities $\mathbf{x}$ in the model functions such that the distance between x and the observed share is minimized.

We use our approach—rather than a maximum likelihood estimation conducted with individual data—because based on our model assumption concerning bidding behavior, each initial bid in an auction below the highest price in the previous auction is unambiguously assigned to the sophisticated bidding strategy so that the observed distribution of initial bids is deterministic. Hence, our method finds the best estimate for the underlying distribution of sophisticated and myopic individual bids. In addition, our approach provides us with an estimate of the distribution of errors among the share of sophisticated bids.

Choice of Norm

In the minimization problem, we use the Euclidean (or 2-)norm. Since one could also minimize the L^1 or even an L^p norm, this choice may not be clear but actually follows naturally from the problem. One can only minimize a

[25] We provide a formal proof concerning the validity of this procedure below.

function $f : \mathbb{K} \mapsto \mathbb{R}$ with $\mathbb{K} \subset \mathbb{R}^n$ but in our case, the function maps to $\mathbb{R}^9$ so one would minimize

$$\langle f^\intercal, f \rangle$$

where $\langle \cdot, \cdot \rangle$ is the standard inner product in $\mathbb{R}$. The norm induced by $\langle \cdot, \cdot \rangle$ is the 2-norm

$$\|\mathbf{x}\|_2 = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$$

and therefore, the choice of the norm follows directly from the problem.

Existence and Uniqueness of the Minimum

The existence of a minimum (and not just an infimum) is guaranteed because $\|f(\mathbf{x})\|_2$ continuous. The constraints require x to be in the compact space $[0, 1]$. λ represents the expected value of the Poisson distribution and is therefore bounded by $\lambda_{\text{max}} = 15$ because that is the maximum number of steps a bidder can erroneously deviate (given $\bar{b}_{t-1} = 100$). Continuous functions attain a minimum on compact spaces (Theorem of Weyerstraß).

Every norm is a convex function because by the triangle inequality and the positive homogeneity, it holds that

$$\forall \Theta \in (0, 1) \;\; \forall x, y \in \mathbb{K}$$

$$\|\Theta x + (1 - \Theta) y\| \leq \Theta \|x\| + (1 - \Theta) \|y\| .$$

Therefore, the worst-case scenario is that $\|f(\mathbf{x})\|$ is constant for a small space $\mathbb{S} \subset [0, 1]$ around a critical point. However, it can be easily check that for every critical point $\mathbf{x}$ found, the Hessian matrix is strictly positive definite. It follows that the minimum is unique.

BIBLIOGRAPHY

Abdellaoui, M., H. Bleichrodt, and C. Paraschiv (2007). Loss aversion under prospect theory: A parameter-free measurement. *Management Science 53*(10), 1659–1674.

Abeler, J., A. Falk, L. Goette, and D. Huffman (2011). Reference points and effort provision. *American Economic Review 101*(2), 470–92.

Adomavicius, G., S. P. Curley, and A. Gupta (2013). Impact of information feedback in continuous combinatorial auctions: An experimental study of economic performance. *MIS Quarterly 37*(1), 55–76.

Agranov, M., E. Potamites, A. Schotter, and C. Tergiman (2012). Beliefs and endogenous cognitive levels: An experimental study. *Games and Economic Behavior 75*, 449–463.

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école Américaine. *Econometrica 21*(4), 503–546.

Andreoni, J. and C. Sprenger (2010). Certain and uncertain utility: The Allais paradox and five decision theory phenomena. *Working Paper*.

Anscombe, F. J. and R. J. Aumann (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, 199–205.

Ariely, D. and I. Simonson (2003). Buying, bidding, playing, or competing? Value assessment and decision dynamics in online auctions. *Journal of Consumer Psychology 13*(1&2), 113–123.

Armantier, O. and N. Treich (2009). Star-shaped probability weighting functions and overbidding in first-price auctions. *Economics Letters 104*(2), 83–85.

Banerji, A. and N. Gupta (2014). Detection, identification, and estimation of loss aversion: Evidence from an auction experiment. *American Economic Journal: Microeconomics 6*(1), 91–133.

Beattie, J. and G. Loomes (1997). The impact of incentives upon risky choice experiments. *Journal of Risk and Uncertainty 14*, 155–168.

Belica, M. and K.-M. Ehrhart (2013). Reference-dependent Bayesian games. *Working Paper*.

Bichler, M., K. Guler, and S. Mayer (2015). Split-award procurement auctions—Can Bayesian equilibrium strategies predict human bidding behavior in multi-object auctions? *Production and Operations Management 24*(6), 1012–1027.

Bichler, M., Z. Hao, and G. Adomavicius (2017). Coalition-based pricing in ascending combinatorial auctions. *Infomation Systems Research.*

Blume, A. and P. Heidhues (2004). All equilibria of the vickrey auction. *Journal of Economic Theory 114*(1), 170–177.

Bosch-Domènech, A., J. G. Montalvo, R. Nagel, and A. Satorra (2002). One, two, (three), infinity, ...: Newspaper and lab beauty-contest experiments. *American Economic Review 92*, 1687–1701.

Bose, S. and A. Daripa (2009). A dynamic mechanism and surplus extraction under ambiguity. *Journal of Economic Theory 144*(5), 2084–2114.

Brunnermeier, M. K. and J. Morgan (2010). Clock games: Theory and experiments. *Games and Economic Behavior 68*(2), 532–550.

Bulow, J. and P. Klemperer (1994). Rational frenzies and crashes. *Journal of Political Economy 102*, 1–23.

Cai, G. G., P. R. Wurman, and X. Gong (2010, 2014/08/28). A note on discrete bid first-price auction with general value distribution. *International Game Theory Review 12*(01), 75–81.

Camerer, C. F. (1989). An experimental test of several generalized utility theories. *Journal of Risk and Uncertainty 2*, 61–104.

Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics 119*, 861–898.

Carare, O. and M. Rothkopf (2005). Slow dutch auctions. *Management Science 51*(3), 365–373.

Cassady, R. (1967). *Auctions and Auctioneering*. University of California Press.

Celik, G. and O. Yilankaya (2009). Optimal auctions with simultaneous and costly participation. *B.E. Journal of Theoretical Economics (Advances) 9*(1), Article 24.

Che, Y.-K. (1993). Design competition through multidimensional auctions. *The RAND Journal of Economics*, 668–680.

Chen, D. L., M. Schonger, and C. Wickens (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance 9*, 88 – 97.

Chwe, M. S.-Y. (1989). The discrete bid first auction. *Economics Letters 31*(4), 303–306.

Cooper, D. J. and H. Fang (2008). Understanding overbidding in second price auctions: An experimental study*. *The Economic Journal 118*(532), 1572–1595.

Coppinger, V. M., V. L. Smith, and J. A. Titus (1980). Icentives and behavior in english, dutch and sealed bid auctions. *Economic Inquiry 18*(1), 1–22.

Costa-Gomes, M. A. and V. P. Crawford (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review 96*, 1737–1768.

Cox, J. C. and D. James (2012). Clocks and trees: Isomorphic Dutch auctions and centipede games. *Econometrica 80*, 883–903.

Cox, J. C., B. Roberson, and V. L. Smith (1982). Theory and behavior of single object auctions. *Research in Experimental Economics 2*, 1–43.

Cox, J. C., V. L. Smith, and J. M. Walker (1983). A test that discriminates between two models of the dutch-first auction non-isomorphism. *Journal of Economic Behavior & Organization 4*(2–3), 205–219.

Crawford, V. P., M. A. Costa-Gomes, and N. Iriberri (2013). Structural models of non-equilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature 51*, 5–62.

Crawford, V. P. and N. Iriberri (2007). Level-*k* auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica 75*, 1721–1770.

Crawford, V. P. and J. Meng (2011). New york city cab drivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. *American Economic Review 101*(5), 1912–32.

Dean, M. and P. Ortoleva (2014). Is it all connected? A testing ground for unified theories of behavioral economics phenomena. *Working Paper*.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic literature 47*(2), 315–72.

Duffy, J. and R. Nagel (1997). On the robustness of behaviour in experimental 'beauty contest' games. *Economic Journal 107*, 1684–1700.

Dufwenberg, M. and U. Gneezy (2000). Price competition and market concentration: an experimental study. *International Journal of Industrial Organization 18*, 7–22.

Ehrhart, K. and M. Ott (2014). Reference-dependent bidding in dynamic auctions. Technical report, working paper.

Eisenhuth, R. and M. Ewers (2012). Auctions with loss averse bidders. Technical report, Citeseer.

Engelbrecht-Wiggans, R. and E. Katok (2007, October). Regret in auctions: theory and evidence. *Economic Theory 33*(1), 81–101.

Ericson, K. M. M. and A. Fuster (2011). Expectations as endowments: Evidence on reference-dependent preferences from exchange and valuation experiments. *Quarterly Journal of Economics 126*(4), 1879–1907.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*, 171–178.

Fonseca, M. A. and H.-T. Normann (2012). Explicit vs. tacit collusion—the impact of communication in oligopoly experiments. *European Economic Review 56*(8), 1759–1772.

Fugger, N., E. Katok, and A. Wambach (2016). Collusion in dynamic buyer-determined reverse auctions. *Management Science 62*(2), 518–533.

Gächter, S., E. J. Johnson, and A. Herrmann (2007). Individual-level loss aversion in riskless and risky choices.

Georganas, S. (2011). English auctions with resale: An experimental study. *Games and Economic Behavior 73*(1), 147–166.

Georganas, S., P. J. Healy, and R. A. Weber (2015). On the persistence of strategic sophistication. *Journal of Economic Theory 159*(PA), 369–400.

Gill, D. and V. Prowse (2019). Measuring costly effort using the slider task. *Journal of Behavioral and Experimental Finance 21*, 1–9.

Gill, D. and V. Prowse (forthcoming). Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. *Journal of Political Economy*.

Gillen, P., V. Gretschko, and A. Rasch (2017). Pre-auction or post-auction qualification? *Economic Theory Bulletin 5*(2), 139–150.

Gillen, P., A. Rasch, A. Wambach, and P. Werner (2016). Bid pooling in reverse multi-unit dutch auctions: an experimental investigation. *Theory and Decision 81*(4), 511–534.

Goeree, J. K., C. A. Holt, and T. R. Palfrey (2016). *Quantal Response Equilibrium - A Stochastic Theory of Games*. Princeton University Press.

Goeree, J. K., T. Offerman, and A. Schram (2006). Using first-price auctions to sell heterogeneous licenses. *International Journal of Industrial Organization 24*, 555–581.

Greiner, B. (2004). An online recruitment system for economic experiments. In K. Kremer and V. Macho (Eds.), *Forschung und wissenschaftliches Rechnen 2003, GWDG Bericht 63*, pp. 79–93. Göttingen, Germany: Gesellschaft für Wissenschaftliche Datenverarbeitung.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association 1*(1), 114–125.

Gretschko, V. and A. Rajko (2015). Excess information acquisition in auctions. *Experimental Economics 18*, 335–355.

Gretschko, V., A. Rasch, and A. Wambach (2014). On the strictly descending multi-unit auction. *Journal of Mathematical Economics 50*, 731–751.

Gretschko, V. and A. Wambach (2014). Information acquisition during a descending auction. *Economic Theory 55*(3), 731–751.

Grimm, V. and U. Schmidt (2000). Equilibrium bidding without the independence axiom: A graphical analysis. *Theory and Decision 49*(4), 361–374.

Guler, K., L. Fine, K. Chen, A. Karp, T. Liu, H. Tang, F. Safai, R. Wu, and A. Zhang (2002, November 21). Automated decision support system for designing auctions. US Patent App. 09/858,251.

Guler, K., T. Liu, and H. Tang (2003). Joint estimation of bidders' risk attitudes and private information. US Patent App. 09/904,311.

Guler, K., T. Liu, and H. K. Tang (2009, 05). Method and system for automated bid advice for auctions.

Haruvy, E. and E. Katok (2013). Increasing revenue by decreasing information in procurement auctions. *Production and Operations Management 22*, 19–35.

Haruvy, E. and P. T. L. Popkowski Leszczyc (2009). Internet auctions. *Foundations and Trends in Marketing 4*(1), 1–75.

Hass, C., M. Bichler, and K. Guler (2013). Optimization-based decision support for scenario analysis in electronic sourcing markets with volume discounts. *Electronic Commerce Research and Applications 12*(3), 152–165.

Heffetz, O. and J. A. List (2014). Is the endowment effect an expectations effect? *Journal of the European Economic Association 12*(5), 1396–1422.

Hennig-Thurau, T., E. C. Malthouse, C. Friege, S. Gensler, L. Lobschat, A. Rangaswamy, and B. Skiera (2010). The impact of new media on customer relationships. *Journal of Service Research 13*(3), 311–330.

Herweg, F., D. Müller, and P. Weinschenk (2010). Binary payment schemes: Moral hazard and loss aversion. *American Economic Review 100*(5), 2451–77.

Hinloopen, J. and A. R. Soetevent (2008). Laboratory evidence on the effectiveness of corporate leniency programs. *RAND Journal of Economics 39*(2), 607–616.

Ho, T.-H., C. Camerer, and K. Weigelt (1998). Iterated dominance and iterated best response in experimental "p-beauty contests". *American Economic Review 88*, 947–969.

Ho, T.-H. and X. Su (2013). A dynamic level-k model in sequential games. *Management Science 59*, 452–469.

Hong, Y., C. Wang, and P. A. Pavlou (2016). Comparing open and sealed bid auctions: Evidence from online labor markets. *Information Systems Research 27*(1), 49–69.

Horowitz, J. K. and K. E. McConnell (2003). Willingness to accept, willingness to pay and the income effect. *Journal of Economic Behavior & Organization 51*(4), 537–545.

Ivanov, A., D. Levin, and M. Niederle (2010). Can relaxation of beliefs rationalize the winner's curse? An experimental study. *Econometrica 78*, 1435–1452.

Jap, S. D. (2002). Online reverse auctions: Issues, themes, and prospects for the future. *Journal of the Academy of Marketing Science 30*(4), 506–525.

Kagel, J. H. and D. Levin (2002). *Common value auctions and the winner's curse*. Princeton University Press.

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1990). Experimental tests of the endowment effect and the coase theorem. *Journal of political Economy 98*(6), 1325–1348.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–292.

Kambil, A. and E. Van Heck (1998). Reengineering the dutch flower auctions: A framework for analyzing exchange organizations. *Information Systems Research 9*(1), 1–19.

Kaplan, T. R. and A. Sela (2006). Second price auctions with private entry costs. *Available at SSRN 442521*.

Katok, E. and A. M. Kwasnica (2007). Time is money: The effect of clock speed on seller's revenue in dutch auctions. *Experimental Economics 11*(4), 344–357.

Katok, E. and A. E. Roth (2004). Auctions of homogeneous goods with increasing returns: Experimental comparison of alternative "Dutch" auctions. *Management Science 50*, 1044–1063.

Kayande, U., A. De Bruyn, G. L. Lilien, A. Rangaswamy, and G. H. Van Bruggen (2009). How incorporating feedback mechanisms in a dss affects dss evaluations. *Information Systems Research 20*(4), 527–546.

Kőszegi, B. and M. Rabin (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics 121*(4), 1133–1165.

Kim, E. H., A. Morse, and L. Zingales (2006). What has mattered to economics since 1970. *Journal of Economic Perspectives 20*(4), 189–202.

Kirchkamp, O. and J. P. Reiss (2011). Out-of-equilibrium bids in first-price auctions: Wrong expectations or wrong bids. *Economic Journal 121*, 1361–1397.

Kirchkamp, O. and J. P. Reiß (2014). Heterogeneous bids in auctions with rational and boundedly rational bidders—theory and experiment. Technical report, Working paper.

Kocher, M. G. and M. Sutter (2005). The decision maker matters: Individual versus group behaviour in experimental beauty-contest games. *Economic Journal 115*, 200–223.

Kőszegi, B. and M. Rabin (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics 121*(4), 1133–1165.

Krishna, V. (2009). *Auction theory*. Academic press.

Kwasnica, A. M. and K. Sherstyuk (2007). Collusion and equilibrium selection in auctions. *Economic Journal 117*, 120–145.

Lange, A. and A. Ratan (2010). Multi-dimensional reference-dependent preferences in sealed-bid auctions–how (most) laboratory experiments differ from the field. *Games and Economic Behavior 68*(2), 634–645.

Levitt, S. D., J. A. List, and S. E. Sadoff (2011). Checkmate: Exploring Backward Induction among Chess Players. *American Economic Review 101*(2), 975–90.

Lucking-Reiley, D. (1999). Using field experiments to test equivalence between auction formats: Magic on the internet. *American Economic Review 89*(5), 1063–1080.

Martínez-Pardina, I. and A. Romeu (2011). The case for multi-unit single-run descending-price auctions. *Economics Letters 113*, 310–313.

Maskin, E. and J. Riley (1984). Optimal auctions with risk averse buyers. *Econometrica 52*(6), 1473–1518.

McAfee, R. and D. Vincent (1997). Sequentially optimal auctions. *Games and Economic Behavior 18*(2), 246 – 276.

McCabe, K. A., S. J. Rassenti, and V. L. Smith (1990). Auction institutional design: Theory and behavior of simultaneous multiple-unit generalizations of the Dutch and English auctions. *American Economic Review 80*, 1276–1283.

McKelvey, R. D. and T. R. Palfrey (1992). An experimental study of the centipede game. *Econometrica 60*, 803–36.

McKelvey, R. D. and T. R. Palfrey (1995). Quantal response equilibria for normal form games. *Games and economic behavior 10*(1), 6–38.

Menezes, F. M. and P. K. Monteiro (2000). Auctions with endogenous participation. *Review of Economic Design 5*, 71–89.

Muren, A. and R. Pyddoke (2006). Collusion without communication. *Information Economics and Policy 18*(1), 43–54.

Murphy, R., A. Rapoport, and J. Parco (2006). The breakdown of cooperation in iterative real-time trust dilemmas. *Experimental Economics 9*, 147–166.

Nakajima, D. (2011). First-price auctions, dutch auctions, and buy-it-now prices with allais paradox bidders. *Theoretical Economics 6*(3), 473–498.

Palacios-Huerta, I. and O. Volij (2009). Field centipedes. *American Economic Review 99*(4), 1619–35.

Park, S., G. E. Bolton, L. Rothrock, and J. Brosig (2010). Towards an interdisciplinary perspective of training intervention for negotiations: Developing strategic negotiation support contents. *Decision Support Systems 49*(2), 213–221.

Penczynski, S. P. (2011). Strategic thinking: The influence of the game. Working Paper.

Rapoport, A., W. E. Stein, J. E. Parco, and T. E. Nicholas (2003). Equilibrium play and adaptive learning in a three-person centipede game. *Games and Economic Behavior 43*(2), 239–265.

Romeu, A. (2000). Some applications of simulation-based and semi-nonparametric estimation methods in microeconometrics. PhD Dissertation, Universitat Autónoma de Barcelona.

Rosato, A. and A. A. Tymula (2019). Loss aversion and competition in vickrey auctions: Money ain't no good. *Games and Economic Behavior 115*, 188–208.

Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons, Ltd.

Schmidt, U. and C. Seidl (2014). Reconsidering the common ratio effect: The roles of compound idependence, reduction, and coalescing. *IfW Kiel Working Paper 1930*.

Schwartz, L. M., S. Woloshin, W. C. Black, and G. Welch (1997). The role of numeracy in understanding the benefit of screening mammography. *Annals of Internal Medicine 127*(11), 966–972.

Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory 4*(1), 25–55.

Shachat, J. and L. Wei (2012). Procuring Commodities: First-Price Sealed-Bid or English Auctions? *Marketing Science 31*(2), 317–333.

Sherstyuk, K. (1999). Collusion without conspiracy: An experimental study of one-sided auctions. *Experimental Economics 2*, 59–75.

Sherstyuk, K. (2002). Collusion in private value ascending price auctions. *Journal of Economic Behavior & Organization 48*, 177–195.

Simon, J. L. (1968). An almost practical solution to airline overbooking. *Journal of Transport Economics and Policy 2*, 201–202.

Simon, J. L. (1994). Origins of the airline oversales auction system. *Regulation 17*, 48–52.

Sprenger, C. (2010). An endowment effect for risk: Experimental tests of stochastic reference points. *Working Paper*.

Suetens, S. and J. Potters (2007). Bertrand colludes more than Cournot. *Experimental Economics 10*, 71–77.

Tan, G. (1996). Optimal procurement mechanisms for an informed buyer. *The Canadian Journal of Economics / Revue canadienne d'Economique 29*(3), 699–716.

Tan, G. and O. Yilankaya (2006). Equilibria in second price auctions with participation costs. *Journal of Economic Theory 130*, 205–219.

Todd, P. and I. Benbasat (1991). An experimental investigation of the impact of computer based decision aids on decision making strategies. *Information Systems Research 2*(2), 87–115.

Turocy, T. L., E. Watson, and R. C. Battalio (2007). Framing the first-price auction. *Experimental Economics 10*(1), 37–51.

Vartiainen, H. (2013). Auction design without commitment. *Journal of the European Economic Association 11*(2), 316–342.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance 16*(1), 8–37.

von Wangenheim, J. (2019). English versus vickrey auctions with loss averse bidders. Technical report, CRC TRR 190 Rationality and Competition.

Wan, Z. and D. R. Beil (2009). Rfq auctions with supplier qualification screening. *Operations Research 57*(4), 934–949.

Wan, Z., D. R. Beil, and E. Katok (2012). When does it pay to delay supplier qualification? theory and experiments. *Management Science 58*(11), 2057–2075.

Weber, B. J. (2007). The effects of losses and event splitting on the allais paradox. *Judgment and Decision Making 2*(2), 115–125.

Weber, R. J. (1982). The allais paradox, dutch auctions, and alpha-utility theory. *Northwestern University Discussion Paper 536*.

Zeppenfeld, C. (2015). *Decisions under Uncertainty: Preferences, Institutions, and Social Interaction*. Ph. D. thesis, University of Cologne.

Zhang, B. and K. Guler (2013). Determination of a bid value associated with a selected bidder. US Patent 8,548,882.

CURRICULUM VITAE

PERSONAL DETAILS

Name	Philippe Gillen
Date of birth	January 16, 1989
Place of birth	Dudelange

EDUCATION

2015	Master of Science in Mathematics, Universität zu Köln
2013	Bachelor of Science in Mathematics, Universität zu Köln
2013	Bachelor of Science in Physics, Universität zu Köln
2009	Graduation from secondary education, Lycée de Garçons Esch

Köln, December 3, 2019