Hypothetical thinking and the winner's curse: An experimental investigation*

Johannes Moser†

March 28, 2017

Abstract

There is evidence that bidders fall prey to the winner's curse because of mistakes in hypothetical thinking. I provide a lab experiment with two stages to investigate this relationship. In stage I the subjects participate in a non-standard common value auction called the wallet game in which a naïve bidding strategy can lead to both: winner's curse and loser's curse. In stage II the subjects in the treatment group learn whether their initial bid was the winning bid or not with the possibility to change this bid. In this sense the bidders face the same decision problems as in stage I again but the need for hypothetical thinking is reduced in stage II. The overall pattern of the data suggests that the problem of winner's and loser's curse can be weakened by giving the subjects ex ante feedback about their bid, when both are regarded separately.

This is a preliminary draft. Please do not cite or distribute without permission of the author.

JEL-Classification: D03, D44, D82, D83, C91

Keywords: Hypothetical thinking, winner's curse, common value auctions

*This paper is partly based on my master's thesis circulated under the title "The Winner's Curse: Investigations of an anomaly".

†University of Regensburg, Department of Economics, johannes.moser@ur.de

1 Introduction

Imagine the following scenario: Several bidders compete for a particular oil field with an unknown but common value in a second-price sealed bid auction. After all bidders submitted their bids the auctioneer looks through the bids and searches for the winning bid. Before the auctioneer announces the result officially she tells the winning bidder in private that he is the one who submitted the highest bid. She gives him a chance to submit a new bid if he would like to. When the winning bidder receives this information he might update his valuation for the oil field and come to the conclusion that the true value is lower than his estimate because he is the one who submitted the highest bid. Obviously he should have taken this into account beforehand. The same behavior can be found in the famous Groucho Marx quote: "Please accept my resignation. I don't want to belong to any club that will accept me as a member", which indicates a failure in hypothetical thinking resulting in the winner's curse.

I show in an experimental setup that bidders in a common value auction are more likely to avoid the winner's curse (irrational overbidding) and the loser's curse (irrational underbidding) when they are informed whether their bid is the winning bid or not - an information bidders receive usually only at the very end of an auction. The overall results suggest that a substantial part of irrational bidding behavior in common value auctions can be explained by mistakes in hypothetical thinking. At the same time there is also a downside of receiving information about one's bid since the subjects differ only imperfectly between situations in which decreasing (increasing) a bid is rational and those in which it is not. Additionally my findings cast doubts on the accuracy of belief-based models like *cursed equilibrium* (Eyster and Rabin, 2005) in explaining the winner's curse.

The winner's curse is a phenomena whose existence is empirically well documented in field and laboratory studies. It describes a situation, in an auction context, in which individuals systematically tend to overbid, relative to the true value of an object. Thus the winner of an auction might be the actual loser, because he payed a price that exceeds the value of the auctioned good. Empirically this phenomena was first described by Capen et al. (1971). They showed that many oil companies in the 1960's and 1970's had to report a drop in profit rates, which they ascribed to irrational bidding behavior like systematic overbidding in auctions for drilling rights. Later also experimental evidence for this phenomena was found in a large number of lab studies (see for example Bazerman and Samuelson, 1983; Charness and Levin, 2009; Ivanov et al., 2010; Brocas et al., 2017). Most of these studies conclude that the origins of the winner's curse can be seen in bounded rational behavior of the agents. In this regard there are two main classes of theories: Explanations concerning errors in conditional reasoning on future events and explanations concerning the belief formation of individuals. Furthermore there are also some psychological arguments which see the reason for the winner's curse in emotional (Astor et al., 2013) or social (Van den Bos et al., 2008) aspects of winning. In the economic literature there is an ongoing debate whether belief-based models like *cursed equilibrium* and level-k model (Eyster and Rabin, 2005; Crawford and Iriberri, 2007)[1] or approaches concerning contingent reasoning on hypothetical future events (Charness and Levin, 2009; Ivanov et al., 2010; Esponda and Vespa, 2014) are more accurate to explain the winner's curse.

My experiment takes a closer look on the issue of contingent reasoning by providing an auction game in which the bidders learn whether a bid which was considered as optimal ex ante is the winning bid or not (but without learning their payoff yet) and with the possibility to change this bid. This treatment intervention is similar to the *sequential* treatment in Esponda and Vespa (2014) where participants in an election learned whether their vote was pivotal or not before they had to cast a vote.

Contingent reasoning refers to the ability of thinking through hypothetical scenarios. There is evidence that people have difficulties to engage in this cognitive task. Whereas this is well documented in the psychological literature (see for example Evans et al., 2007; Nickerson, 2015; Singmann et al., 2016), economists devoted little attention to this issue for a long time. However, in the more recent economic literature this topic appears more and more frequently (see for example Charness and Levin, 2009; Louis, 2013; Esponda and Vespa, 2014; Koch and Penczynski, 2014; Li, 2015; Ngangoue and Weizsacker, 2015; Levin et al., 2016; Esponda and Vespa, 2016).

The concept of contingent reasoning has still received very little formal treatment in the economic literature. The first one who provided a formal definition of this mental process is Li (2015) by introducing the concept of obviously strategy-proof (OSP) mechanisms which can be seen as an extension of strategy-proof mechanisms. By definition of his paper a mechanism is OSP if and only if an optimal strategy can be found without the necessity of performing contingent reasoning. An OSP mechanism requires an equilibrium in obviously dominant strategies. A strategy is obviously dominant if such a cognitively limited agent can recognize a strategy as weakly dominant. More formally:

> A strategy S_i is *obviously dominant* if, for any deviating strategy S_i', starting from any earliest information set where S_i and S_i' diverge, the best possible outcome from S_i' is no better than the worst possible outcome from S_i. (Li, 2015)

Li (2015) showed that participants in *private value* ascending bid auctions perform better than in theoretically equivalent second-price auctions what he ascribed to the fact that the first ones are OSP while the second ones are only strategy-proof. However, in non-trivial *common value* auctions there is normally no weakly dominant strategy hence

[1] Both models fall into the category of belief-based models, since the cause manifesting in the winner's curse is seen in the belief formation of individuals. The general assumptions of the Bayesian Nash Equilibrium are still fulfilled in the sense that subjects best-respond to beliefs but the assumption about the consistency of beliefs is relaxed. In the cursed equilibrium the degree of *cursedness* is given by $\chi \in [0, 1]$, i.e. the belief that with some probability χ the actions of the opponents do not depend on their types. A value of 0 is equivalent to the usual Bayesian Nash Equilibrium, whereas a value of 1 corresponds to a setting in which the players do not assume any correlation between the actions of a player and his type which is also denoted as the *fully* cursed equilibrium.

they are not even strategy-proof no matter whether they are conducted as ascending bid or second-price auction. A crucial point of the winner's curse is that it appears typically in common value auctions, because of the inherent adverse selection issue, but not in private value auctions and the concept of OSP mechanisms cannot explain why subjects behave differently in these two kinds of auctions. In this sense the concept of OSP is only partly useful for explaining why bidders fall prey to the winner's curse.[2] However, it can be a powerful tool to explain why people make less mistakes in ascending bid auctions compared to sealed bid auctions when the valuations are private.

The experiment presented in this paper is a second-price sealed bid auction, similar to the *wallet game* proposed by Klemperer (1998) and the model used in Avery and Kagel (1997) which is a non-standard common value auction. The basic idea of the game is the following: Two players, indexed by $i = 1, 2$, receive a private iid signal x_i drawn from some commonly known distribution. In a second-price sealed bid auction they bid for an object worth $v = x_1 + x_2$. A more detailed and formal description of this game will be provided in Section 3. The reasons for choosing this special form are the following: First, it is simple enough, such that it can easily be understood by inexperienced subjects while it still represents a common value auction context and second, the same cognitive mistakes can lead to both over- and underbidding, relative to the symmetric equilibrium strategy (i.e. there can be both a winner's and a loser's curse (see also Holt and Sherman, 1994)). The latter point is suitable to control for psychological explanations, like emotional factors of winning. Henceforth a *winner's curse* can be understood as winning an auction but with a negative payoff and a *loser's curse* as losing an auction but the bidder could have won with a positive payoff.

My experiment consists of two stages. In stage I the subjects participate in the wallet game against a random opponent. In stage II the subjects play the same auctions again against the decisions of the former opponent but this time the subjects in the treatment group are informed whether their initial bid was the winning bid or not. This is a *dynamic* treatment intervention since the bidders receive information about some realized event before they have to come up with a bid.[3] Apart from knowing whether a certain bid is the winning or losing bid the subjects face exactly the same decision problem as in stage I if we abstract from social preferences.[4] It is important to stress that the treatment intervention does not lead to a complete elimination of contingent reasoning in the sense of OSP mechanisms, since even with knowing whether a certain bid is the winning (or losing) bid there are still many more contingencies left and the game is still not solvable in weakly dominant strategies, hence it is not even strategy-proof. In this sense my *information* treatment reveals in fact only some contingency for the bidder but a crucial one for a cognitively limited agent.

[2] In general Li's definition of contingent reasoning is very strong and according to his definition even a simple prisoner's dilemma requires contingent reasoning although it is solvable in dominant strategies.

[3] See also Esponda and Vespa (2016) for a distinction between static and dynamic choice situations.

[4] Since the subjects do not directly interact with their former opponents in stage II but only play against their decisions, the payoff of the respective opponent is not affected anymore by the own decisions.

The aim of this paper is not to come up with a generalized formal theory of contingent reasoning but to look at one crucial aspect of it in an experimental setup. My contribution to the current literature is to provide an experiment in which the requirements for contingent reasoning concerning winning or losing are manipulated in a real auction context. To the best of my knowledge, none of the recent approaches provide a framework with a *dynamic* treatment intervention in a sealed bid auction.[5]

Additionally the effects of my treatment intervention can be clearly attributed to mistakes in hypothetical thinking and cannot be explained by cursed equilibrium since a crucial assumption of these belief-based models is that by definition no (or only a partly) correlation between the other players' actions and types is assumed. Hence for a "cursed" bidder the information whether his bid is higher or lower than the bid of the opponent does not provide him further information which would be relevant for updating his bid.

Finally I am making use of a game that includes both, the risk of a winner's curse and a loser's curse, to control for psychological explanations and to check whether bid shading after the treatment intervention is due to proper Bayesian updating or just a rule of thumb. An advantage of my design is that bid shading is only useful for low signals but not for high signals. In this sense the bidders have to differentiate between those two kinds of signals.

This paper is organized as follows. Section 2 will provide an overview about the current literature concerning the winner's curse and hypothetical thinking. Section 3 will present the underlying theoretical model for the experiment. Section 4 presents the experimental design. Section 5 presents the results of the experiment. Section 6 discusses the results and concludes the paper.

2 Literature review

There is a large body of experimental literature related to the winner's curse. Some of the most influential experimental publications concerning the winner's curse and hypothetical thinking are presented in the following part. At the same time I want to point out how my experiment can complement the existing literature in a meaningful way.

Esponda and Vespa (2014) created a common value voting experiment where a subject interacts together with two computers. The main task for the subject was to submit a vote for a ball which was either red or blue. Since the vote of the subject was only relevant when the ball was blue, the optimal choice for the subject was always to vote for blue. To distinguish between hypothetical thinking and information extraction problems Esponda and Vespa (2014) conducted a *simultaneous* voting treatment and a *sequential*

[5] So for example the experiment by Charness and Levin (2009) is rather a single decision maker problem in an adverse selection environment, the experiment conducted by Esponda and Vespa (2014) eliminates the requirements for hypothetical thinking, but in a voting game and in the experiments by Koch and Penczynski (2014) and Levin et al. (2016) the treatment interventions are still static ones in the sense that the bidders do not receive information about a realized event which was hypothetical ex ante.

voting treatment with the difference that the player in the sequential form was able to see the decisions of the computers before he had to make his choice. Esponda and Vespa (2014) showed that subjects in the sequential form made significantly less errors compared to the simultaneous form. They concluded that most of the subjects are in fact able to correctly extract information from the computers' actions, but that they have problems to think in hypothetical situations (here: being the pivotal voter). Similarly Ngangoue and Weizsacker (2015) found in a lab experiment that traders in a financial market setting perform better in a sequential trading mechanism where no Bayesian updating on hypothetical events was required.

In a second study Esponda and Vespa (2016) analyzed further classical economic anomalies and games like Ellsberg and Allais paradoxes, second-price auctions and elections focusing on the "Sure-Thing Principle" which was first pointed out by Savage (1972). The authors showed that even nudging the subjects with the issue of hypothetical thinking had a significant effect on their behavior.

Charness and Levin (2009) conducted an experiment constructed as a simple individual choice problem similar to an *acquiring a company game* which is based on a lemon market (Akerlof, 1970) to investigate the driving mechanisms behind the winner's curse. The results of their paper revealed that most of the subjects systematically overbid even though the strategy of the computerized seller was known and hence there was no need for the subjects to form beliefs about his behavior. This pattern was reduced in a setting, where the bidding task was transformed into a set of simple lotteries with no requirement of thinking in hypothetical situations. However, transforming the initial game into a simple lottery task changes the whole structure of the game and so it remains difficult to extract a causal effect of hypothetical thinking.

Ivanov et al. (2010) used a similar approach as Charness and Levin (2009) but they conducted their experiment in a real auction context using the *maximal game*.[6] The aim of their experiment was not to look on the effect of contingent reasoning but rather to disprove that the winner's curse is driven by inconsistent beliefs. Ivanov et al. (2010) observed significant overbidding, which was not reduced in a modified setting of the maximal game where belief-based models, like cursed equilibrium, had few explanatory power. However, Costa-Gomes and Shimoji (2015) criticized that Ivanov et al. (2010) misused some of the game theoretical concepts arguing that some of their findings can indeed be explained by belief-based models. Similarly, Camerer et al. (2016) argued that the observed behavior of the subjects in Ivanov et al. (2010) can be explained by belief-based models if they are combined with a *quantal* response model. Under this extension the assumption of perfect best-reply behavior in cursed equilibrium and level-k model is

[6] Ivanov et al. (2010) criticized that the acquiring a company game used in Charness and Levin (2009) basically represents a lemon market (Akerlof, 1970) and not a common value auction. So it seems to be problematic to extend the findings from their experiment to common value auctions in general. They also claimed that it can make a difference whether a subject plays against other people or against a computer. In fact, Ivanov et al. (2010) also used a computer treatment, but in their case the computer mimicked the subject's own past strategy.

relaxed and stochastic choices are allowed. They showed that this extended model fits very well to the data of Ivanov et al. (2010). Contrary to Costa-Gomes and Shimoji (2015), Camerer et al. (2016) also believe that, if the *perfect* best-reply assumption of cursed equilibrium and level-k model is maintained, they are very bad in predicting the behavior in maximal value games.

Koch and Penczynski (2014) conducted an auction game similar to the one in Kagel and Levin (1986) to investigate how explanations concerning contingent reasoning and belief-based models interact.[7] They used a transformed version of this game with no private but only common information to identify the effect of contingent reasoning on the bidding behavior of the subjects.[8] The authors claimed that there is no requirement to think in hypothetical situations in the transformed game. In the sense of OSP mechanisms this is not true and even in a more psychological definition of contingent reasoning this remains questionable since in their transformed model the value of the object is not deterministic but it depends stochastically on the bidding behavior of the subjects. The requirement for hypothetical thinking is only affected in a sense that both players know that they have the same information and thus face the same decision problem. The justification for their setup is that the strategic nature of the original game, concerning best-respond functions and equilibria, is still maintained, but this comes at the cost of a higher complexity what makes it difficult to compare both games. However, the work of Koch and Penczynski (2014) is a very important contribution to the current research, since it provides the first joint analysis of both main explanations for the winner's curse.

Levin et al. (2016) used an experiment to investigate the impact of Bayesian updating and non-probabilistic reasoning (referred to as *contingent reasoning* in this paper) on avoiding the winner's curse. They used common-value Dutch and common-value first-price auctions based on the model in Kagel et al. (1987) and compared both versions to quantify the effect of non-probabilistic reasoning. Additionally they measured the skills of the participants in Bayesian updating and non-probabilistic reasoning through a questionnaire which the subjects had to answer before they participated in the auctions. Levin et al. (2016) showed that both cognitive skills had an significant effect on avoiding the winner's curse resulting in higher earnings for the subjects. The authors used three different forms of auctions as treatments: An active clock Dutch auction (AD) where the auction ends when the first subjects stops the clock, an silent clock Dutch auction (SD) where the auction ends, when the last subjects stops the clock (and no subject receives any feedback whether he is the highest bidder) and a first-price sealed bid auction (FP). Although all of these auction are strategically equivalent, Levin et al. (2016) claimed that the requirements for hypothetical thinking in the AD treatment are reduced since in the moment of stopping the clock the subject obviously knows that he is the highest

[7] So far Koch and Penczynski (2014) had been the only ones who combined contingent reasoning and belief-based models in their experiment.

[8] In contrast to my design Koch and Penczynski (2014) focused on a related but different aspect of contingent reasoning.

bidder. The results of Levin et al. (2016) showed that the winning bidders lost money in all treatments but substantially less in the AD treatment, and the authors suggested that the subjects can better handle the winner's curse if the requirements for hypothetical thinking are reduced. Note that the AD treatment is still no dynamic treatment intervention because strictly speaking only *after* stopping the clock the subject learns that he is the highest bidder.[9] The advantage of my design is that the bidders in the treatment group receive this information *before* they have to come up with a bid.

To conclude this section: Many attempts have been used to investigate the impact of hypothetical thinking in different common value environments. The novelty of my design is that I use a dynamic treatment intervention (i.e. the subjects learn about a realized event before they have to come up with a decision) in a real common value auction. At the same time I leave all other parameters unchanged. This is for example not given in the works of Charness and Levin (2009) and Koch and Penczynski (2014). In this sense I am able to extract a clean measure of this cognitive ability and to identify a causal effect on the behavior in common value auctions.

3 The model

In the following section a formal description of the *wallet game* will be presented. First the general model and finally a simplified model of this game which is used in the experiment. The theoretical derivations of the respective bidding functions are primarily based on the works of Eyster and Rabin (2005) and Crawford and Iriberri (2007).

General model

The general wallet game with N bidders can formally be described as follows:

There is a set of N players with $N \geq 2$. Each player i with $i \in \{1, 2, \ldots, N\}$ receives a signal x_i from a uniform distribution with range $[x_{min}, x_{max}]$ which is common knowledge. The bidders compete for an object worth $v = \sum_{i=1}^{N} x_i$ in an auction. The utility function is assumed as symmetric across i such that $u_i(\mathbf{x}) = v = \sum_{i=1}^{N} x_i$ with $\mathbf{x} = [x_1, x_2, \ldots, x_N]^T$ (see also Crawford and Iriberri, 2007).

Given that bidder i sees a signal of x_i the expected value of the object is given by

$$r(x_i) := E[V|X_i = x_i] = x_i + (N-1) \cdot E[X_j] = x_i + (N-1) \cdot \frac{x_{min} + x_{max}}{2} \tag{1}$$

The expected value of the object given that bidder i sees a signal of x_i and the highest signal of all other bidders is y, is given by

$$v(x_i, y) := E[V|X_i = x_i, Y_i = y] \tag{2}$$

[9]A further problem of Dutch auctions is that the bidders have to make a decision pressed for time. This can lead to various effects on an emotional level (see for example Adam et al., 2012, 2015).

where Y_i is defined as $max_{j \neq i}\, x_j$. In this sense $E[V|X_i = x_i, Y_i \leq x_i]$ is the expected value of the object conditional on winning the auction.[10] It can be shown that bidding

$$\begin{aligned} v(x_i, x_i) := E[V|X_i = x_i, Y_i = x_i] &= 2 \cdot x_i + (N-2) \cdot E[X_j|X_j \leq x_i] \\ &= 2 \cdot x_i + \frac{N-2}{2} \cdot (x_i - x_{min}) \end{aligned} \tag{3}$$

is an equilibrium strategy of the wallet game in a second-price auction. This is the unique symmetric equilibrium of this game (Eyster and Rabin, 2005; Crawford and Iriberri, 2007).

Henceforth using the bidding function $b_i(x_i) = r(x_i)$ refers to *naïve bidding* and using the bidding function $b_i(x_i) = v(x_i, x_i)$ refers to *sophisticated bidding.* In this context a bidder who plays the naïve strategy $b_i(x_i) = r(x_i)$ overbids for any $x_i < \frac{(N-1) \cdot x_{min} + x_{max}}{N}$ and underbids for any $x_i > \frac{(N-1) \cdot x_{min} + x_{max}}{N}$, relative to the symmetric equilibrium strategy (see also Crawford and Iriberri, 2007).

Model with two players

For only two players we have

$$r(x_i) = x_i + \frac{x_{min} + x_{max}}{2} \tag{4}$$

and

$$v(x_i, x_i) = 2 \cdot x_i \tag{5}$$

It is easy to see that equation (4) describes bidding the own signal plus the expectation of the other's signal, whereas equation (5) describes bidding twice the own signal.

Additionally this game has also a continuum of asymmetric equilibria. For two players $i \in \{1, 2\}$ and for all $\alpha > 1$ the following strategy pairs are Bayesian Nash Equilibria (see *Proofs* in Appendix B):

$$b_1(x_1) = \alpha \cdot x_1, \qquad b_2(x_2) = \frac{\alpha}{\alpha - 1} \cdot x_2 \tag{6}$$

Additionally there can be even more equilibrium strategies like bidding $b_1(x_1) = x_1$ and $b_2(x_2) = 2 \cdot x_{max}$. In this sense almost any bid can be rationalized as potential equilibrium strategy. However, all of those asymmetric equilibria involve weakly dominant bids at least for some x_i (for one player) since bidding above $x_i + x_{max}$ is weakly dominated when this game is conducted as a second-price auction (see *Proofs* in Appendix B). This makes asymmetric equilibria less plausible.[11]

[10] Under the assumption that all bidders have the same bidding function $b_i(x_i)$ which is monotonically increasing in x_i.

[11] So for example choosing the maximal possible bid in a second-price sealed bid auction with *private values* can also be part of an equilibrium strategy when all the other bidders always bid 0. However, this is merely rational when a bidder does not know the bidding strategy of the other bidders.

Simplified model with two players

I now present a slightly modified version of the wallet game with only *low* and *high* signals possible and with a discrete signal space. This is the version used in the experiment.

In this simplified version two players, indexed by $i = 1, 2$, receive a private signal x_i drawn from the set $X = \{0, 1, \ldots, 9, 10, 50, 51, \ldots, 59, 60\}$, with each value equally likely and with replacement. So in total 22 different signals are possible (11 low and 11 high signals). The players compete for an object worth $v = x_1 + x_2$ in a second-price sealed bid auction.

In case of a tie the players with the higher signal wins. If the signals are also equal, both players receive a payoff of 0.[12] The players are allowed to choose a bid b_i in the range of $[0, 120]$, with only integer values possible.[13]
The payoff of player 1 (analogously for player 2) is thus given by:

$$\pi_1 = \begin{cases} x_1 + x_2 - b_2 & \text{if } b_1 > b_2 \\ x_1 + x_2 - b_2 & \text{if } b_1 = b_2 \wedge x_1 > x_2 \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

Since the symmetric equilibrium in the two player wallet game does not depend on the distribution of the signals x_1 and x_2 the sophisticated bidding function remains the same in this simplified version of the game (Klemperer, 1998). Also the naïve bidding function remains unchanged since the expected value of the other's signal is the same as in a continuous setting (only the variance increases because of the gap).

The bidding functions for naïve and sophisticated bidding are:

$$r(x_i) = x_i + 30 \tag{8}$$

and

$$v(x_i, x_i) = 2 \cdot x_i \tag{9}$$

A useful property of this simplified model is that in this setup the sophisticated bidding strategy is also a best-response for the naïve bidding strategy. So even if a sophisticated player assumes that his opponent bids naïvely he still best-responds by using the sophisticated bidding function. The intuition is simple: If the other player uses the naïve bidding strategy the best response is to win always for high signals and to lose always for low signals. This is given when following the sophisticated bidding rule.

As a concluding remark it is important to note that bidding $b_i(x_i) = r(x_i) = x_i + 30$ can be explained by both: mistakes in contingent reasoning and cursed equilibrium or

[12] So there are no chance elements and winning is always deterministic. In the symmetric equilibrium both bidders would also receive a payoff of 0 in case of a tie.

[13] In a first trial session of the experiment I chose a maximal bid of 150. However, some of the participants found this confusing and so I decided to restrict the bidding range. A maximal bid of 120 seems natural since the maximal possible value of the good is 120.

level-k model (Eyster and Rabin, 2005; Crawford and Iriberri, 2007).[14] This elucidates a general dilemma in behavioral and experimental economics: even though a model fits well to the data it is not clear whether it has indeed explanatory power. Hence it remains important to disentangle competing theories.

4 Experimental design

Implementation

The experiment was conducted in the Regensburg Economics Science Lab (RESL). For the technical implementation the software zTree was used (Fischbacher, 2007) and for the recruitment of participants the online recruitment system ORSEE was used (Greiner et al., 2004). In total 5 sessions were conducted with overall 72 participants (33 males and 39 females). For each session I had between 10 and 18 participants (always an even number). Table 1 shows an overview of all sessions.

The currency used in the experiment were *Taler*. All the signals and bids in the experiment were expressed in terms of Taler. The exchange rate was 1 Euro = 10 Taler. The participants were payed out in Euro at the end of the experiment. The average payment was 16.02 Euro.

At the beginning of the experiment each participant was endowed with 50 Taler. At the end of the experiment, the participants received their initial endowment plus (minus) their generated earnings (losses) in both stages of the experiment (in total 6 rounds were payoff-relevant). Additionally a show-up fee of 4 Euro was payed to each subject which was guaranteed no matter what decisions the subject made during the experiment. So each subject earned at least 4 Euro. If the losses exceeded 50 Taler the participants only received their show-up fee. 4 out of 72 participants suffered from higher losses.

Session	Participants	Males/Females	Treat. A/Treat. B	Average earnings
1	18	7/11	9/9	16.89 EUR
2	16	7/9	6/10	17.21 EUR
3	18	8/10	12/6	13.25 EUR
4	10	4/6	4/6	15.58 EUR
5	10	7/3	8/2	17.99 EUR
Total	72	33/39	39/33	16.02 EUR

Table 1: Summary of all sessions

[14] A bidder who does not condition his bid on winning ignores the adverse selection issue inherent in this kind of auction and only considers his private signal and a (fully) "cursed" bidder implicitly assumes that the opponent will bid independently of his signal what makes $b_i(x_i) = x_i + 30$ a best response.

Basic setup

In the experiment the simplified model of the wallet game with only two players is used. The experiment is divided into *two stages*. The subjects are informed that there is a second stage but they receive the details only after finishing stage I. In stage I the players participate in 15 rounds of the wallet game (i.e. each subject receives successively 15 random signals drawn from the set $X = \{0, 1, \ldots, 9, 10, 50, 51, \ldots, 59, 60\}$ *with replacement*).[15] Each subject gets randomly matched with another subject of the group (e.g. subject k and subject l). This constellation is stable for all 15 auctions. In this sense the first signal of subject k is matched with the first signal of subject l, the second signal of k is matched with the second signal of l and so on. The subjects receive no immediate feedback after submitting their bids but only learn their payoff at the very end of the experiment (i.e. after finishing stage II). So there should be neither endowment effects nor learning effects. Three randomly selected rounds are payoff-relevant. A typical decision screen of stage I is shown by Figure 2 in Appendix A.

Before starting with the actual task, all participants are asked to answer eight control questions and to participate in five testing rounds of the wallet game without monetary payoff but with immediate feedback about their hypothetical payoff after each bid to give them a practical understanding of the game. However, the subjects receive no feedback about the bid or the signal of the opponent. The opponent in the testing rounds is undertaken by a computer who uses the bidding strategy $b_j(x_j) = x_j + 30$. The participants are not explicitly informed about the strategy of the computer and they only learn that the bidding function of the computer is monotonically increasing in his signal.[16]

Stage II is basically a repetition of stage I.[17] All subjects receive the same 15 signals as in stage I in the same order. In stage II the subjects play against an computerized opponent who mimics the behavior of their former opponent from stage I. So subject k plays against the decisions of subject l in stage I and vice versa. So each subject faces exactly the same decision problems as in stage I if we abstract from social preferences.[18] As in stage I the bids and the signals of the opponent are not observable. The rules for bidding and winning are the same as in stage I and the same three randomly selected rounds are again payoff-relevant. In stage II the subjects are randomly assigned to either treatment A or B.

---

[15] The actual explanation in the experiment is that each player receives an envelope with a random amount of money inside (see *Instructions* in Appendix D).

[16] The reason for using the naïve bidding function as a strategy for the computer in the testing rounds is that deviating from the sophisticated bidding function $b_i(x_i) = 2 \cdot x_i$ is much more harmful, when the other player uses the naïve bidding function. If the computer would have used the sophisticated bidding function, the subjects are less likely to realize that deviating from this strategy is a bad idea. In the presented setup $b_i(x_i) = 2 \cdot x_i$ is a best-response for both: the naïve and the sophisticated bidding function.

[17] The subjects received the instructions for stage II only after all subjects completed stage I.

[18] In contrast to stage I the decisions do not affect the payoff of the opponent anymore. So if a subject has preferences concerning the other player's payoff the decision problem might be different for him.

Treatment A - *Information* (Treatment group)

The subjects who receive treatment A (*Information*) are able to see for each signal whether their initial bid from stage I was HIGHER or LOWER than the respective bid of the opponent.[19] E.g. if a subjects sees that his initial bid of $b_i = \bar{z}$ was HIGHER than the bid of his opponent he knows that submitting a bid of $b_i' = z \geq \bar{z}$ results in winning the auction for sure. Conversely, if a subjects sees that his bid $b_i = \underline{z}$ was LOWER than the bid of his opponent, he knows that submitting a bid of $b_i' = z < \underline{z}$ results in losing the auction for sure. In this sense there is no requirement anymore to condition on the hypothetical event of winning (or losing) for a certain range of bids especially for the bid which was considered as optimal in stage I. An example of an typical screen is given by Figure 3 in Appendix A.

Treatment B - *No information* (Control group)

The subjects in treatment B (*No information*) face exactly the same situation as those in Treatment A except for the point that they do not get any information about the bid of their opponent. Instead of HIGHER or LOWER they only see *???* on their screen. However, all subjects are informed about both treatments (i.e. the subjects in treatment B know how treatment A looks like and vice versa). An example of an typical screen is given by Figure 4 in Appendix A. The general structure of the treatments is illustrated by Figure 1.

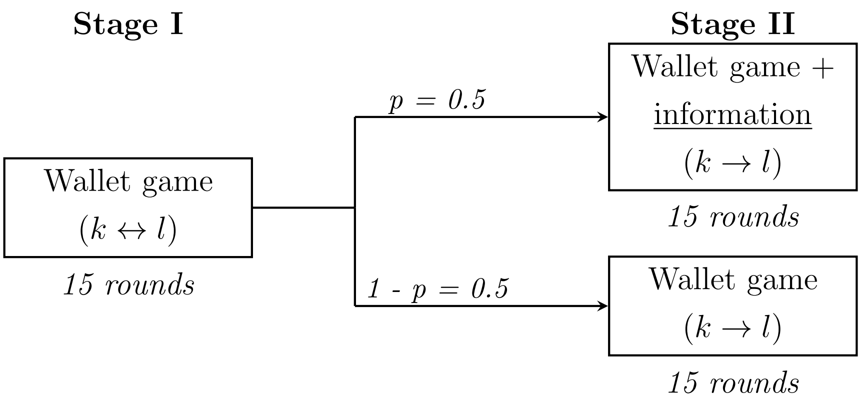

Figure 1: Illustration of the treatments

[19] If the bids are equal the subjects also get the message "LOWER", so LOWER means lower or equal.

5 Results

In this section the results of the experiment will be presented. For each of the 72 subjects I have 15 observations which means that I collected $N = 1080$ observations in total. One observation corresponds to one auction from the perspective of a particular subject and includes the bids for both stages. The overall analysis is on an aggregate level and when looking at the treatment effects I compare different subgroups, which differ in terms of treatment (A, B), signal (*low*, *high*) and information (*LOWER*, *HIGHER*), with a Diff-in-Diff approach.[20]

First I will give a descriptive overview about the overall bidding pattern in stage I when there is no difference for the subjects and then I will present the effects of the treatment intervention in terms of bidding behavior and the resulting profits.[21] Finally I will provide evidence that the observed behavior of the treatment group is to a large extent driven by an actual updating of the opponent's signal when receiving information which is not hypothetical anymore and not because the subjects followed a simple decision rule like "always decrease when you see HIGHER and always increase when you see LOWER".

Result 1. *Overall bidding pattern.*
The median bids for low signals are *above* the Nash prediction and the median bids for high signals are *below* the Nash prediction in stage I (see Figures 5 and 6 in Appendix A). Especially for high signals the median bids are fitted very well by the naïve bidding function $b_i(x_i) = x_i + 30$. The results of a Wilcoxon sign rank test show that for *high signals* the hypothesis that the actual bids are equal to bids resulting from the naïve bidding function cannot be rejected ($p = 0.624$).

Overall we can observe an increased appearance of the winner's curse for low signals (16.67 % in stage I) and of the loser's curse for high signals (17.02 % in stage I). *Conditional on winning* the rate for the winner's curse increases to 59.31 % for low signals and *conditional on losing* the rate for the loser's curse increases to 56.47 % for high signals. These are very high rates especially when considering that the auctions were conducted as *second-price* auctions. This shows clearly that the problem of irrational bidding behavior is a considerable one.

[20] Since the subjects in treatment B do not receive information about their bid in stage II, this feedback can be regarded as *hypothetical* and corresponds to the information they *would have* received if they had been in treatment A.

[21] For the respective comparisons of the differences between treatment and control group I used a Clustered Wilcoxon rank sum test based on the Rosner-Glynn-Lee method, since the observations are not independent (Rosner et al., 2003). As an additional robustness check I also repeated these calculations on an aggregate level by using the individual means of the respective outcome variables (see Appendix C).

	Low (LOST)	**Low** (WON)	**Low**	**High** (LOST)	**High** (WON)	**High**
No curse	338 91.11 %	59 40.69 %	397 76.94 %	74 43.53 %	378 95.94 %	452 80.14 %
Winner's curse	0 0.00 %	86 59.31 %	86 16.67 %	0 0.00 %	16 4.06 %	16 2.84 %
Loser's curse	33 8.89 %	0 0.00 %	33 6.40 %	96 56.47 %	0 0.00 %	96 17.02 %
Total	371 100.00 %	145 100.00 %	516 100.00 %	170 100.00 %	394 100.00 %	564 100.00 %

Table 2: Winner's curse and loser's curse in Stage I for different constellations of signals. Winner's curse: Won but with a negative payoff. Loser's curse: Lost but could have won the auction with a positive payoff.

Result 2. *Deviation from Nash bid.*
For the constellations "low signal and information HIGHER" and "high signal and information LOWER" the absolute deviation from the Nash bid is lower in stage II compared to stage I for the subjects in the treatment group on average. Both effects are significantly larger than in the respective control group. For the other constellations the effect is reverse. Overall, for both low and high signals, there is a very small effect which is on the edge of significance, indicating that on average the subjects in the treatment group deviate even further away from the Nash prediction in stage II than the subjects in the control group. Table 3 reports the changes in the average absolute deviation from the Nash bid from stage I to stage II for different constellations of signal and information and compares the respective values between treatment and control group.

Constellation	Means (Std. deviation)			Observations A / B	Wilcoxon p-value
Change in absolute deviation from Nash bid					
	Information (A)	*No information (B)*	*Difference*		
Low signals (HIGHER)	-16.48 (15.21)	-11.41 (28.87)	-5.07	61 / 79	0.011
Low signals (LOWER)	14.16 (15.93)	5.51 (21.80)	8.65	212 / 164	0.000
High signals (HIGHER)	6.18 (12.19)	1.13 (10.49)	5.05	215 / 176	0.000
High signals (LOWER)	-18.82 (16.30)	-3.74 (13.43)	-15.09	97 / 76	0.000
All signals (HIGHER)	1.17 (15.96)	-2.76 (19.12)	3.93	276 / 255	0.110
All signals (LOWER)	3.81 (22.17)	2.58 (19.99)	1.23	309 / 240	0.078

Notes: The last column reports the p-values of a Clustered Wilcoxon rank sum test using the Rosner-Glynn-Lee method (Rosner et al., 2003). Clusters are on subject level.

Table 3: Summary Table - Deviation from Nash bid

Result 3. *Profits.*
For low signals paired with the information HIGHER the change in average profits from stage I to stage II is positive and significantly higher in the treatment group than in the control group. The same is true for high signals paired with the information LOWER. Conversely, for low signals paired with the information LOWER and for high signals paired with the information HIGHER the change in average profits of the subjects in the control group is significantly less negative than for those in the treatment group.[22]

It can also be seen that the positive effect of the information HIGHER for *low signals* is greater in absolute terms than the negative effect of the information HIGHER for *high signals*. However, the second constellation occurs more often and so the overall effect in the treatment group is not significantly different from the one in the control group. The same pattern can be seen for the information LOWER, where the effect is reverse (i.e. the positive effect of the information LOWER for *high signals* is greater in absolute terms than the negative effect of the information LOWER for *low signals*).

Figures 11 and 12 in Appendix A show the average profits in stage I and II for different constellations of signal and information and compares them to the average profits when the bidders would have unilaterally followed the symmetric Nash bidding rule. Table 4 reports the changes in average profits from stage I to stage II for different constellations of signal and information and compares the respective values between the treatment and control group.

As a robustness check I also used an alternative measure for profitability in which any positive payoff is transformed into 1 and any negative payoff is transformed into -1. In this sense there is only a distinction between whether an auction was won profitable, unprofitable or lost. In contrast to the actual profit, the magnitude of this transformed profit is not affected by the opponent's bid what makes it a cleaner measure for sophisticated bidding. The overall results do not change when using this transformed measure for profits (see Appendix C). As a further robustness check I repeated the calculations of Tables 3 and 4 on a *subject level* by using the individual means for the respective values. Here again, the overall results do not change substantially (see Appendix C).

[22]For low signals paired with the information LOWER the difference between treatment and control group is only significant at the 10%-level.

Constellation	Means (Std. deviation)			Observations A / B	Wilcoxon p-value
	Change in average profits				
	Information (A)	*No information (B)*	*Difference*		
Low signals (HIGHER)	5.84	2.06	3.77	61 / 79	
	(9.76)	(18.95)			0.045
Low signals (LOWER)	-3.76	-1.79	-1.97	212 / 164	
	(10.83)	(8.00)			0.071
High signals (HIGHER)	-1.62	-0.20	-1.41	215 / 176	
	(7.39)	(2.71)			0.024
High signals (LOWER)	5.68	0.57	5.11	97 / 76	
	(13.82)	(9.37)			0.047
All signals (HIGHER)	0.03	0.50	-0.47	276 / 255	
	(8.54)	(10.79)			0.534
All signals (LOWER)	-0.80	-1.05	0.25	309 / 240	
	(12.61)	(8.51)			0.789

Notes: The last column reports the p-values of a Clustered Wilcoxon rank sum test using the Rosner-Glynn-Lee method (Rosner et al., 2003). Clusters are on subject level.

Table 4: Summary Table - Change in average profits

Result 4. *Differentiation between low and high signals.*
By construction of the game and the observed bidding behavior in most of the cases it is profitable for the bidders to increase the initial bid for high signals and to decrease the initial bid for low signals. Thus, when receiving either the information HIGHER or LOWER it is important to distinguish between those two kinds of signals instead of following the simple decision rule "decrease when you receive HIGHER and increase when you receive LOWER". The results show that when receiving the information HIGHER the subjects in the treatment group strongly differentiate between low and high signals. The decrease rate for low signals is 78.7 % and only 42.3 % for high signals. For the information LOWER the subjects in the treatment group do not differentiate between low and high signals. The increase rate for low signals is 83.0 % and 89.7 % for high signals. However, the difference is not statistically significant. After receiving the information LOWER it seems to be very tempting to increase the initial bid. This can be an indicator of an actual joy of winning (or rather disappointment of losing) but this is also in line with theories about hypothetical thinking. If the initial bid was higher it was actually the relevant bid. If the initial bid was lower, this is not the case and the subjects have to engage in hypothetical thinking again before coming up with a new bid.

The subjects in the control group who *would have* received the information HIGHER or LOWER differentiate between low and high signals for both kinds of (hypothetical) information. However, the fraction of decreasing and increasing a bid in stage II is much lower compared to the treatment group. Figures 7 and 8 in Appendix A provide an graphical overview of the bid changing behavior in stage II. Table 5 reports the fractions of decreasing and increasing the initial bids in stage II for different constellations of treatment and information.

Constellation	Fractions		Observations	Fisher's exact
	Fraction of decreased bids			
	Low signals	*High signals*		
HIGHER (Information)	0.787	0.423	61 / 215	0.000
HIGHER (No information)	0.456	0.165	79 / 176	0.000
	Fraction of increased bids			
	Low signals	*High signals*		
LOWER (Information)	0.830	0.897	212 / 97	0.168
LOWER (No information)	0.415	0.618	164 / 76	0.004

Notes: The last column reports the p-values of Fisher's exact test.

Table 5: Summary Table: Differentiation between low and high signals

Result 5. *Changing of the bids in stage II.*
Figures 9 and 10 in Appendix A show histograms of the bids in stage II for the subjects in the treatment group for the constellations "low signal and information HIGHER" and "high signal and information LOWER" *conditional* on increasing or decreasing, respectively. These are the two constellations where the subjects in the treatment group actually profited from the information they receive. Additionally these are the constellations in which the respective *information* provides a relatively unambiguous hint about the opponent's signal.[23] As already shown in Table 3 the average deviation from the Nash prediction decreases from stage I to stage II resulting in higher average profits for the subjects in the treatment group. However, it remains unclear whether the rational increasing or decreasing of the bids is because the subjects realized that the events of winning and losing provide information about the signal of the opponent or is just a rule of thumb.

In Figure 9 it can be seen that there is an accumulation of bids below 20 in stage II for low signals when receiving the information HIGHER (conditional on decreasing the initial bid). This is actually an indicator that the subjects indeed realized that if they win with a low signal, the other player has most likely also a low signal and hence they bid exactly for this case (if both players have low signals, the *maximal* value of the good

[23]For example when receiving the information LOWER for a low signal it is not so clear whether the opponent has a low or a high signal. The same is true for the information HIGHER when having a high signal.

is 20). Equivalently for high signals paired with the information LOWER there should be an accumulation of bids above 100, when the subjects realize that the other player has most likely a high signal, when losing with a high signal (if both players have high signals, the *minimal* value of the good is 100). For this constellation the accumulation is less distinct but still noticeable (see Figure 10 in Appendix A).

The results of Table 5 already showed that at least for the information HIGHER there is a distinction between low and high signals. Combined with the results of Figures 9 and 10 it seems to be very plausible to assume that the changing behavior in stage II of the subjects in the treatment group is to a large extent driven by an actual updating of the opponent's signal. This is also in line with the answers of a questionnaire, which was provided after the actual experiment. 79.49 % (31 out of 39) of the participants in the treatment group answered that the information they received helped them indeed to get a better estimate of the opponent's signal.

6 Conclusion and discussion

This paper investigates whether subjects in a common value auction perform better when they already learn ex ante, before the final payoffs are known, whether their bid is the winning bid or not - an information bidders in a sealed bid auction usually receive only at the very end of the auction. The results show that there is a significant effect of the *information* treatment on the bidding behavior of the participants. For *low signals* the information that the bid was HIGHER helps the bidders to correct their bids downwards resulting in much lower rates of the winner's curse. On the other hand for *high signals* the information that the bid was LOWER helps the bidders to increase their bids resulting in bids closer to the Nash prediction and lower rates of the loser's curse. However, for *high signals* paired with the information HIGHER and for *low signals* paired with the information LOWER there is a negative effect and the bidders in the treatment group depart even further from the Nash prediction. The negative effects of the latter two constellations are much smaller than the positive effects of the first ones in absolute terms but the constellations "low signal and information LOWER" and "high signal and information HIGHER" appear by construction more often than their respective counterparts what results in a neutralization of the positive effect of the *information* treatment on the average profit. So overall the bidders in the *information* treatment are not better off than those who receive *no information* and we can see that additional information can be even negative for the bidders.[24]

However, if we regard the two phenomena of winner's and loser's curse separately, the

[24] In related studies like Charness and Levin (2009) and Koch and Penczynski (2014) the optimal behavior was mainly given by choosing a bid as low as possible. In my design the bidders have to differentiate between low and high signals and decreasing a bid is not always optimal. This additional hurdle shows that additional information can be negative for the bidders when they differentiate only imperfectly between situations in which decreasing (increasing) a bid is rational and those in which it is not.

information treatment is indeed helpful for the bidders. The information HIGHER helps to significantly reduce the winner's curse and the information LOWER helps to significantly reduce the loser's curse if we look at those constellations in which the respective information provides a relatively unambiguous hint about the opponent's signal. For the information HIGHER the bidders in the treatment group strongly differentiate between low and high signals (though imperfectly), resulting in a much higher decreasing rate for low signals, what indicates that the changing of the bids is not just a rule of thumb but rather due to Bayesian updating. This claim is also supported by the pattern that *conditional on decreasing* a bid for low signals after receiving the information HIGHER, most of these newly chosen bids are below 20. This suggests that the bidders indeed realized that the other bidder has most likely also a low signal. For the information LOWER there is no significant differentiation between low and high signals but in this case there is again a requirement of conditioning on winning for the bidders. So for example a naïve Bayesian updater might correctly assume that the value of the good is higher than expected when receiving the information LOWER and increase his initial bid. The problem is that the new bid is not chosen conditional on winning what results in an even more frequent appearance of the winner's curse for low signals.

A further conclusion is that the overall behavior of the subjects in the treatment group in stage II cannot be explained by cursed equilibrium since a "cursed" bidder assumes by definition that the bids and signals of the opponent are not (or only partly) correlated. Hence such a bidder ought not react on the information of winning or losing because he implicitly assumes that the bid of the opponent provides no valid information about the true value of the good. By definition the bid of a "cursed" player is already evaluated conditional on winning in stage I and hence there would be no need to change it in stage II. This casts doubts whether the initial bidding behavior in stage I can be explained by cursed equilibrium unless one assumes that a bidder can suffer from both: a cursed system of beliefs and the inability of thinking in hypothetical situations. However, it is not clear how or whether the effects of both cognitive mistakes add up. So far Koch and Penczynski (2014) had been the only ones who looked at both combined in a lab setting but more research is needed especially concerning the interaction of both cognitive mistakes. In general my findings support the claim of Ivanov et al. (2010) who stated that bidders in common value auctions might act "as if" they have cursed beliefs.[25]

As a concluding remark there is to say that mistakes in hypothetical thinking seem to explain a substantial part of irrational bidding behavior in common value auctions. However, even without the necessity of conditioning on winning there still exists a significant deviation from optimal behavior which remains unexplained.

[25]Ivanov et al. (2010) were one of the first authors who claimed that bidders in common value auctions might just act "as if" they have cursed beliefs since they observed seemingly cursed behavior in a context where belief-based models had few explanatory power.

References

Adam, M. T., Krämer, J., and Müller, M. B. (2015). Auction fever! How time pressure and social competition affect bidders arousal and bids in retail auctions. *Journal of Retailing*, 91(3):468–485.

Adam, M. T., Krämer, J., and Weinhardt, C. (2012). Excitement up! Price down! Measuring emotions in dutch auctions. *International Journal of Electronic Commerce*, 17(2):7–40.

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500.

Astor, P. J., Adam, M. T., Jähnig, C., and Seifert, S. (2013). The joy of winning and the frustration of losing: A psychophysiological analysis of emotions in first-price sealed-bid auctions. *Journal of Neuroscience, Psychology, and Economics*, 6(1):14–30.

Avery, C. and Kagel, J. H. (1997). Second-price auctions with asymmetric payoffs: An experimental investigation. *Journal of Economics & Management Strategy*, 6(3):573–603.

Bazerman, M. H. and Samuelson, W. F. (1983). I won the auction but don't want the prize. *Journal of Conflict Resolution*, 27(4):618–634.

Brocas, I., Carrillo, J. D., and Castro, M. (2017). Second-price common value auctions with uncertainty, private and public information: experimental evidence. *Journal of Behavioral and Experimental Economics.*

Camerer, C., Nunnari, S., and Palfrey, T. R. (2016). Quantal response and nonequilibrium beliefs explain overbidding in maximum-value auctions. *Games and Economic Behavior*, 98:243–263.

Capen, E. C., Clapp, R. V., Campbell, W. M., et al. (1971). Competitive bidding in high-risk situations. *Journal of Petroleum Technology*, 23(6):641–653.

Charness, G. and Levin, D. (2009). The origin of the winner's curse: A laboratory study. *American Economic Journal: Microeconomics*, 1(1):207–236.

Costa-Gomes, M. A. and Shimoji, M. (2015). A comment on "Can relaxation of beliefs rationalize the winner's curse?: An experimental study". *Econometrica*, 83(1):375–383.

Crawford, V. P. and Iriberri, N. (2007). Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica*, 75(6):1721–1770.

Esponda, I. and Vespa, E. (2014). Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6(4):180–202.

Esponda, I. and Vespa, E. (2016). Contingent preferences and the sure-thing principle: Revisiting classic anomalies in the laboratory. *working paper*.

Evans, J. S. B. et al. (2007). *Hypothetical thinking: Dual processes in reasoning and judgement*, volume 3. Psychology Press.

Eyster, E. and Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73(5):1623–1672.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, 10(2):171–178.

Greiner, B. et al. (2004). The online recruitment system orsee-a guide for the organization of experiments in economics. Technical report, Max Planck Institute of Economics, Strategic Interaction Group.

Holt, C. A. and Sherman, R. (1994). The loser's curse. *The American Economic Review*, 84(3):642–652.

Ivanov, A., Levin, D., and Niederle, M. (2010). Can relaxation of beliefs rationalize the winner's curse?: An experimental study. *Econometrica*, 78(4):1435–1452.

Kagel, J. H., Harstad, R. M., and Levin, D. (1987). Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, 55(6):1275–1304.

Kagel, J. H. and Levin, D. (1986). The winner's curse and public information in common value auctions. *The American Economic Review*, 76(5):894–920.

Klemperer, P. (1998). Auctions with almost common values: The wallet game and its applications. *European Economic Review*, 42(3):757–769.

Koch, C. and Penczynski, S. P. (2014). The winner's curse: Conditional reasoning & belief formation. *working paper*.

Levin, D., Peck, J., and Ivanov, A. (2016). Separating bayesian updating from non-probabilistic reasoning: An experimental investigation. *American Economic Journal: Microeconomics*, 8(2):39–60.

Li, S. (2015). Obviously strategy-proof mechanisms. *Available at SSRN 2560028*.

Louis, P. (2013). The barrel of apples game: Contingent thinking, learning from observed actions, and strategic heterogeneity. *working paper*.

Ngangoue, K. and Weizsacker, G. (2015). Learning from unrealized versus realized prices. *working paper*.

Nickerson, R. (2015). *Conditional Reasoning: The Unruly Syntactics, Semantics, Thematics, and Pragmatics of "if"*. Oxford University Press.

Rosner, B., Glynn, R. J., and Ting Lee, M.-L. (2003). Incorporation of clustering effects for the wilcoxon rank sum test: a large-sample approach. *Biometrics*, 59(4):1089–1098.

Savage, L. J. (1972). *The foundations of statistics*. Courier Corporation.

Singmann, H., Klauer, K. C., and Beller, S. (2016). Probabilistic conditional reasoning: Disentangling form and content with the dual-source model. *Cognitive Psychology*, 88:61–87.

Van den Bos, W., Li, J., Lau, T., Maskin, E., Cohen, J. D., Montague, P. R., and McClure, S. M. (2008). The value of victory: Social origins of the winner's curse in common value auctions. *Judgment and Decision Making*, 3(7):483–492.

A Figures

Decision screens

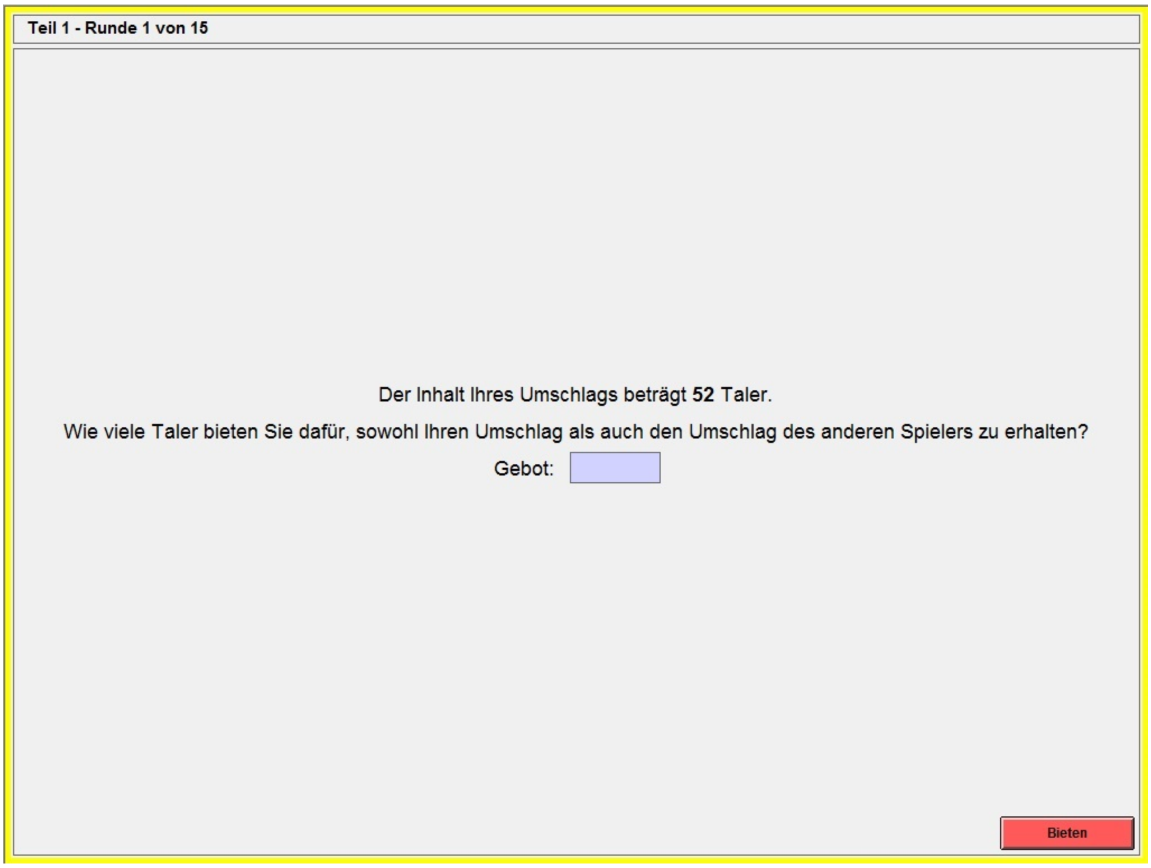

Figure 2: Typical decision screen in stage I. *The value of your envelope is 52 Taler. How many Taler do you want to bid to receive your envelope and the envelope of the other player?*

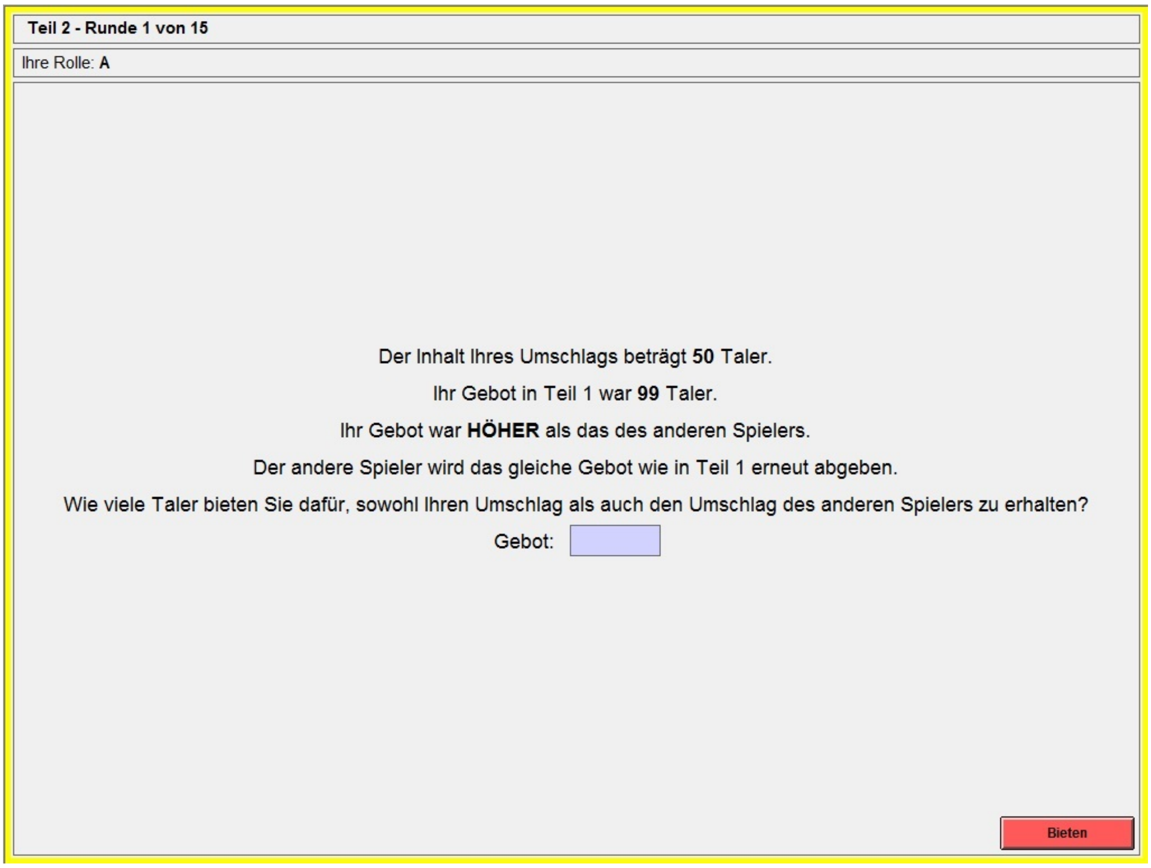

Figure 3: Typical decision screen in stage II (treatment A). *The value of your envelope is 50 Taler. Your bid in part 1 was 99 Taler. Your bid was HIGHER than the bid of the other player. The other player will choose the same bid as in part I again. How many Taler do you want to bid to receive your envelope and the envelope of the other player?*

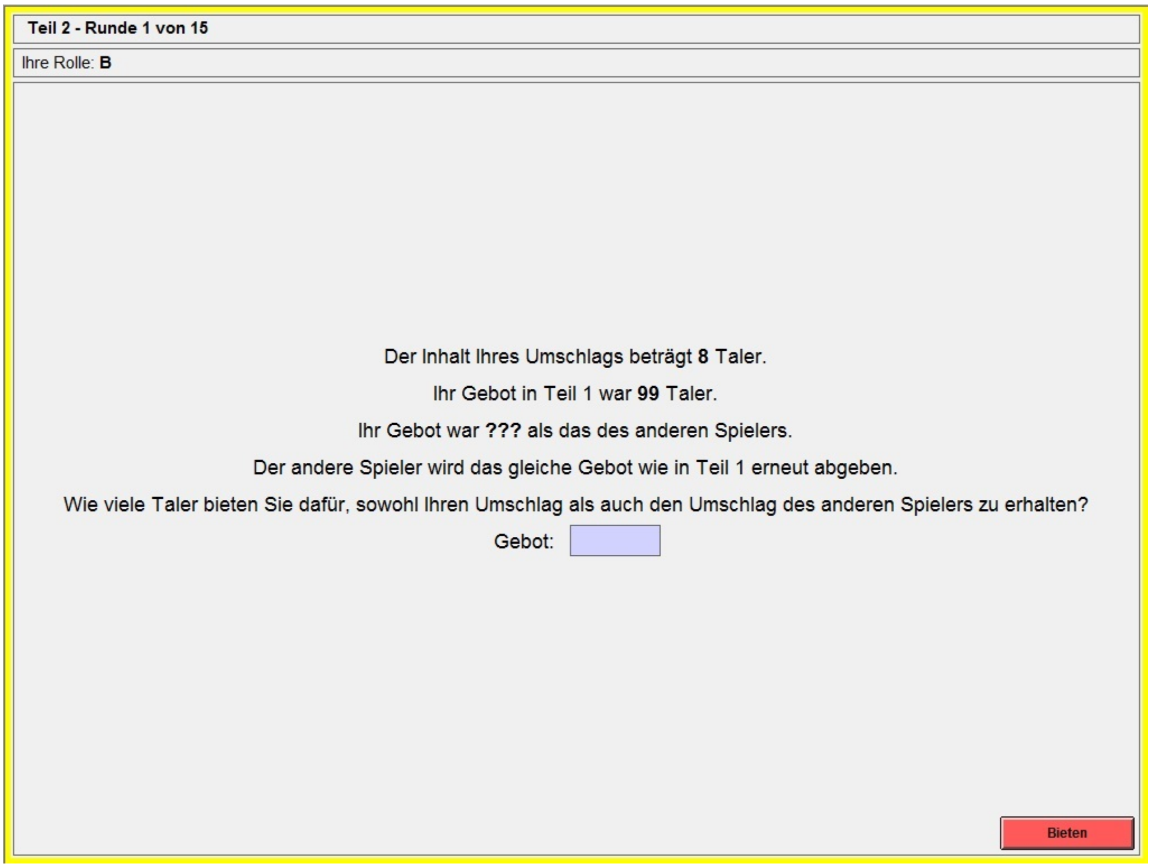

Figure 4: Typical decision screen in stage II (treatment B). *The value of your envelope is 8 Taler. Your bid in part 1 was 99 Taler. Your bid was ??? than the bid of the other player. The other player will choose the same bid as in part I again. How many Taler do you want to bid to receive your envelope and the envelope of the other player?*

Bidding behavior

Median bids

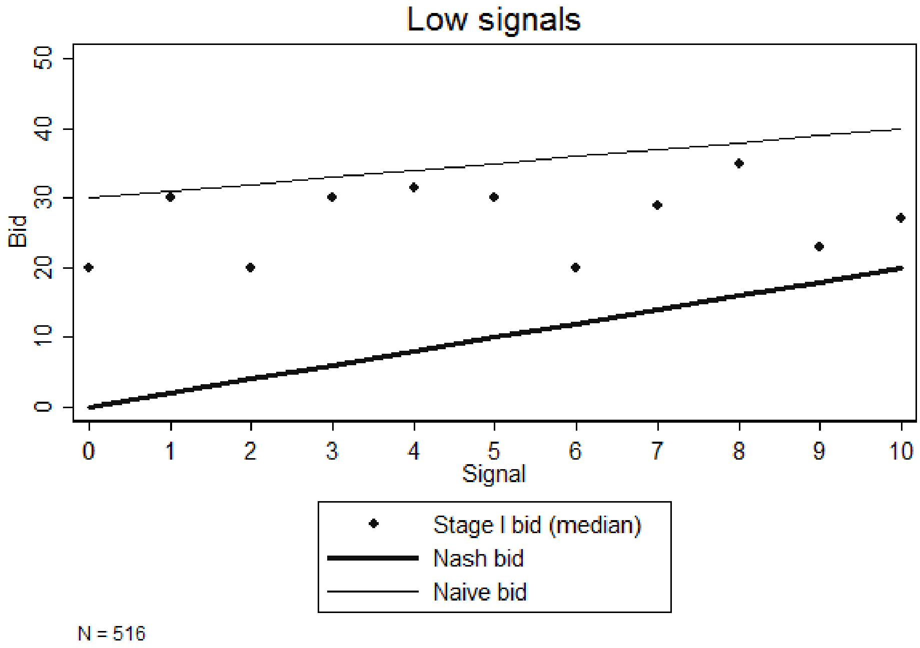

Figure 5: Median bids - Low signals

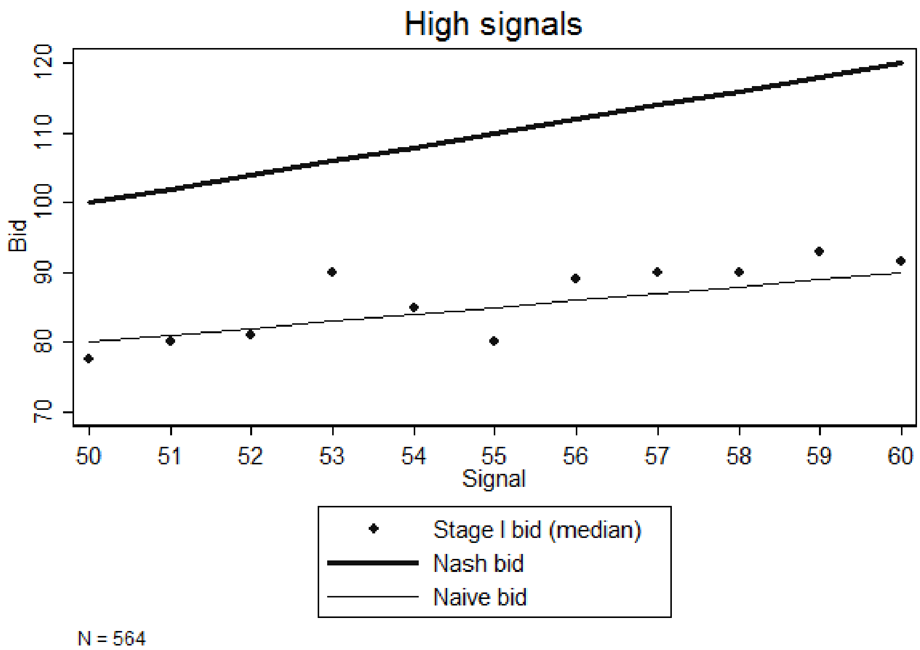

Figure 6: Median bids - High signals

Changing of the bids in stage II

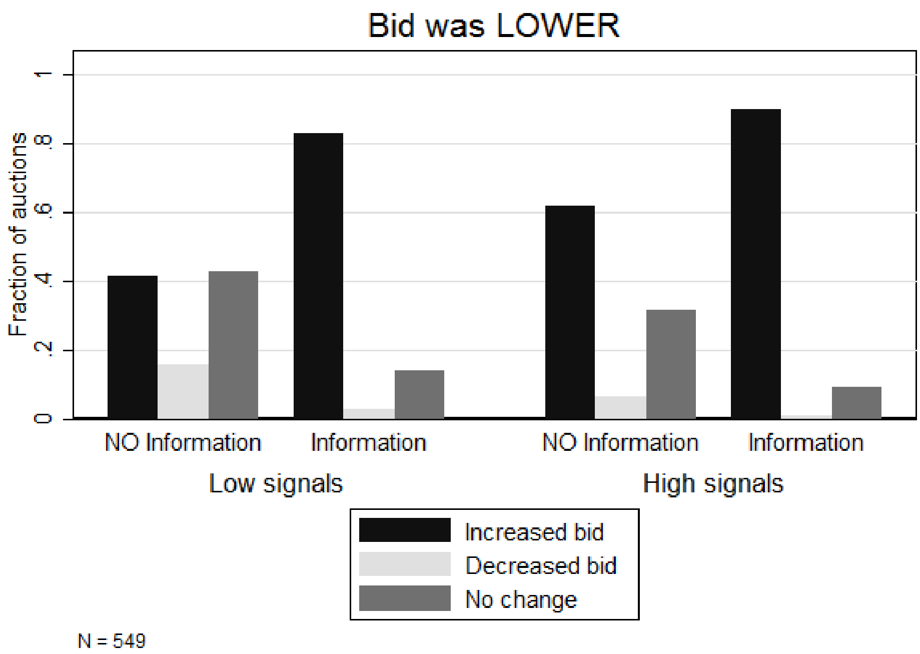

Figure 7: Changing of the bids when bid was LOWER

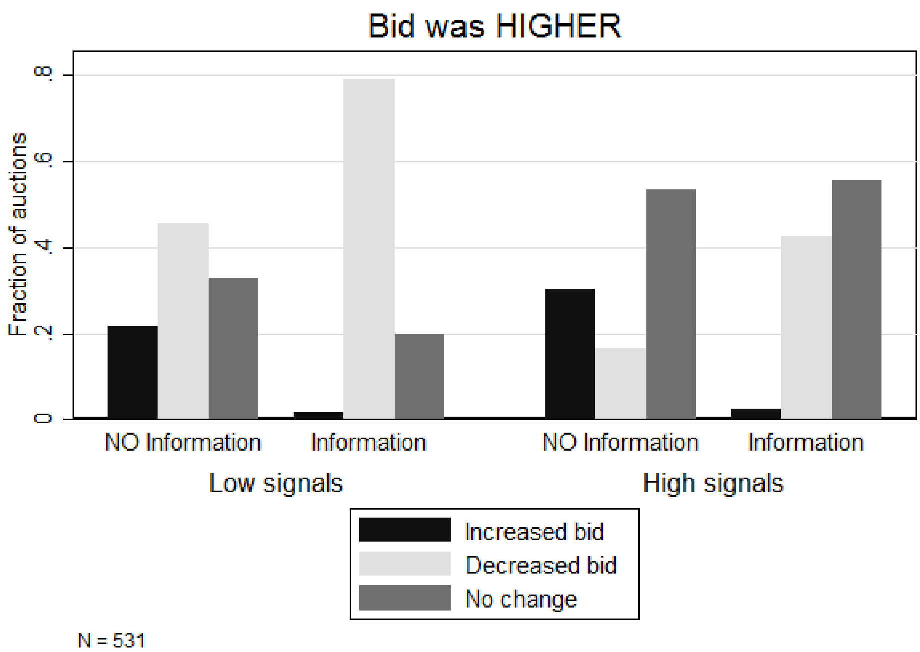

Figure 8: Changing of the bids when bid was HIGHER

Changing of the bids in stage II

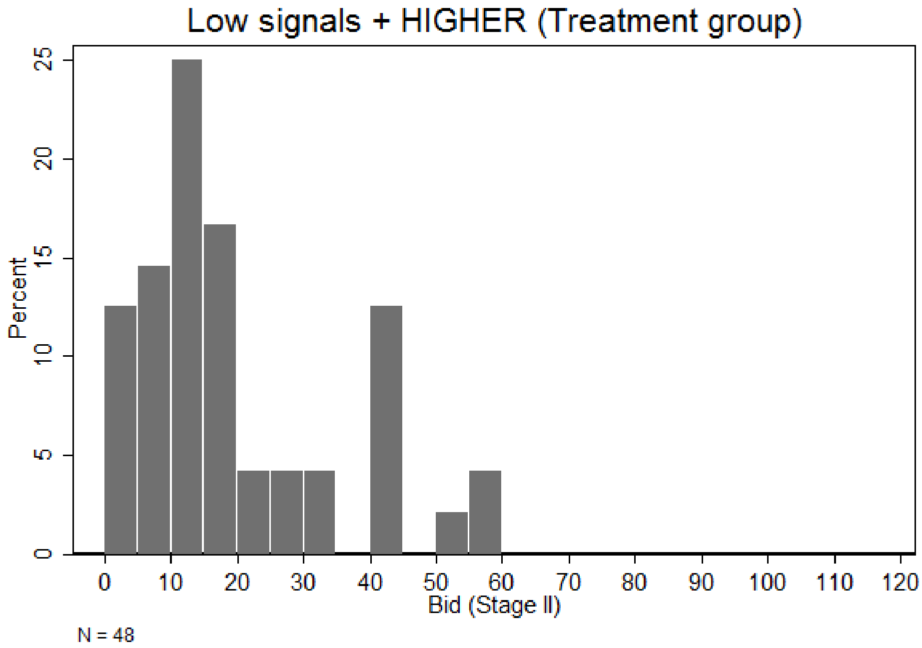

Figure 9: Stage II bid (conditional on *decreasing*)

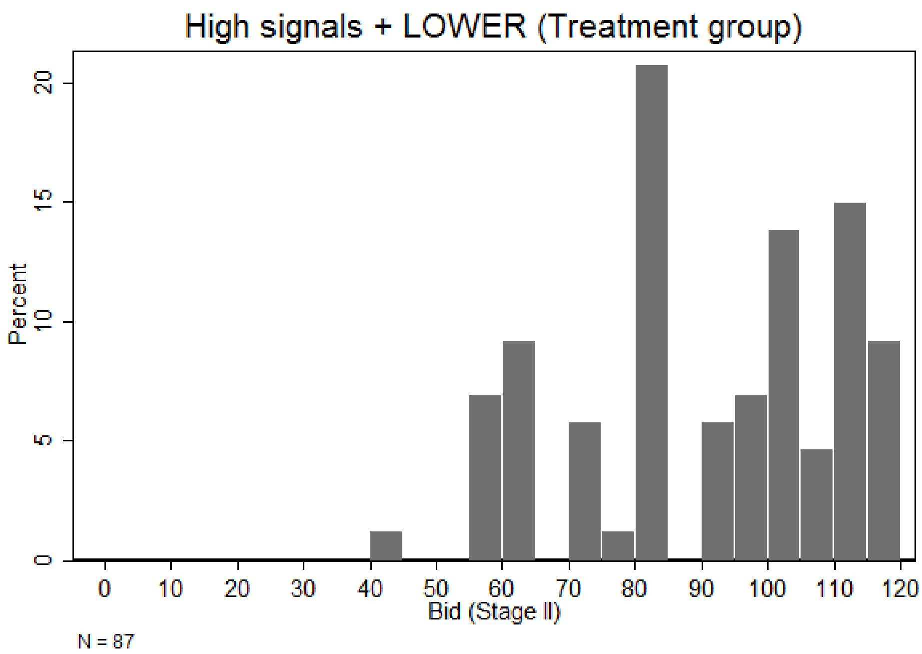

Figure 10: Stage II bid (conditional on *increasing*)

Profits

Average profits

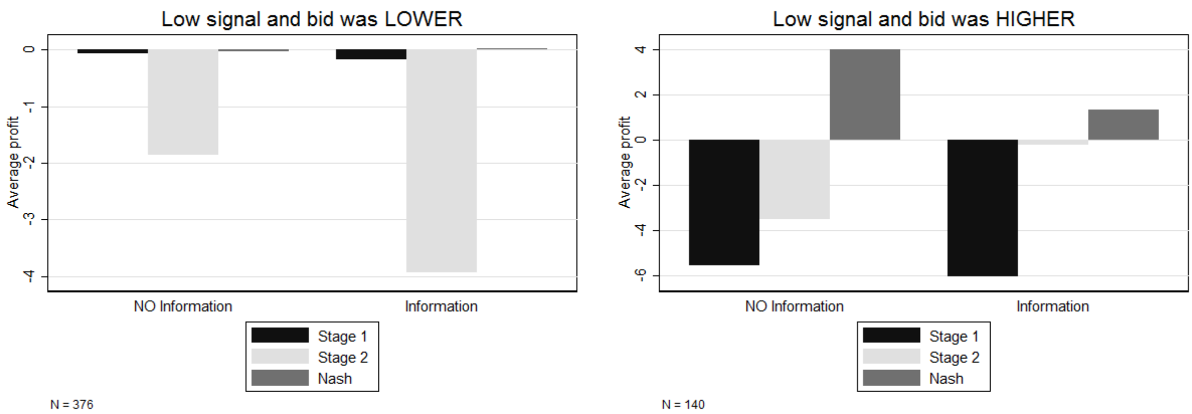

Figure 11: Average profits for low signals

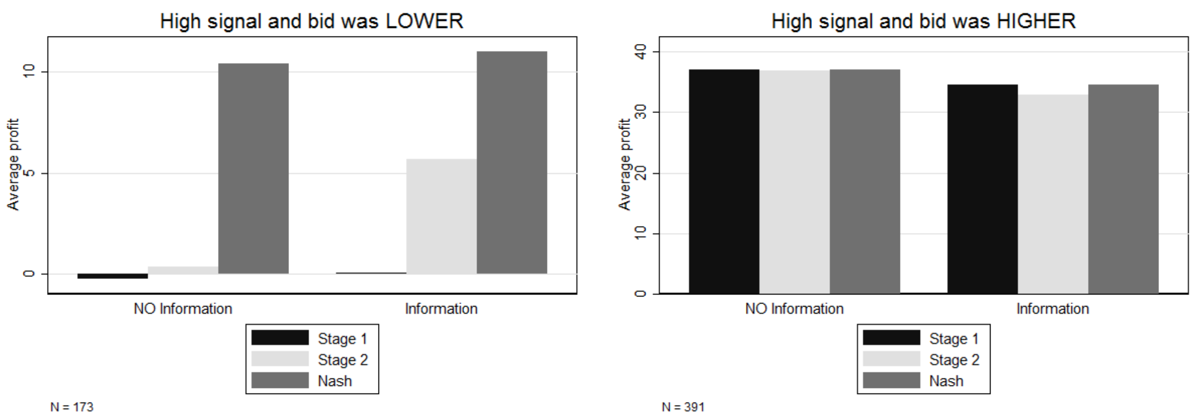

Figure 12: Average profits for high signals

Average profits (transformed)

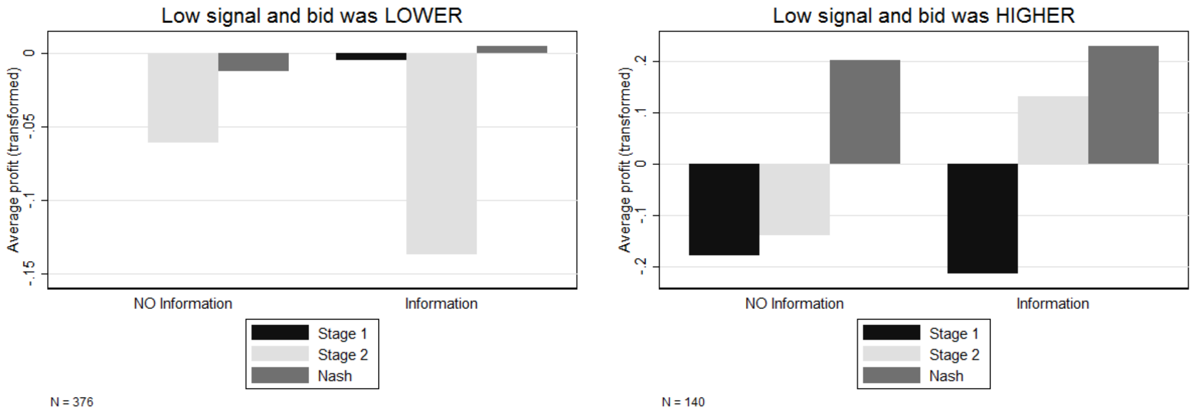

Figure 13: Average profits (transformed) for low signals

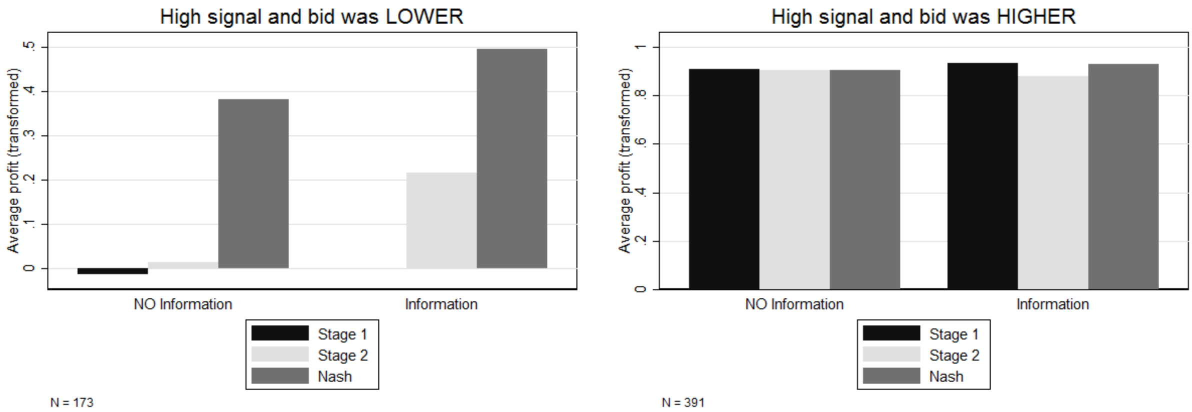

Figure 14: Average profits (transformed) for high signals

B Proofs

Proof. In the two player wallet game, conducted as second-price sealed bid auction, any bid outside the interval $[x_i, x_i + x_{max}]$ is weakly dominated. Consider two players, indexed by $i = 1, 2$.

(i) Bidding below x_i is weakly dominated by bidding x_i. If player 1 would have lost the auction with $b_1(x_1) = x_1$, deviating to a bid below x_1 ($b_1'(x_1) < b_1(x_1)$) would not change the result because $b_1'(x_1) < b_1(x_1) \leq b_2(x_2)$. If player 1 would have won the auction with $b_1(x_1) = x_1$ he receives a payoff of at least 0, because the value of the object is at least x_1 and $b_2(x_2) \leq b_1(x_1) = x_1$. Deviating to a bid below x_1 can lead to losing an auction that generated a positive payoff. This is the case if player 1 bids $x_1 - \delta$ and player 2 bids $x_1 - \epsilon$ with $\delta > \epsilon \geq 0$. Thus player 1 gives up an auction that guarantees a payoff of at least $x_1 + x_2 - x_1 + \epsilon = x_2 + \epsilon \geq 0$. If $\epsilon > \delta > 0$ player 1 still wins when he deviates, but the payoff does not change, because he still receives $x_1 + x_2 - x_1 + \epsilon = x_2 + \epsilon$ as before.

(ii) Bidding above $x_i + x_{max}$ is weakly dominated by bidding $x_i + x_{max}$. If player 1 would have won the auction with $b_1(x_1) = x_1 + x_{max}$ he receives either a positive, negative or zero payoff. Deviating to a bid above $x_1 + x_{max}$ (($b_1'(x_1) > b_1(x_1)$) would not change the result because $b_1'(x_1) > b_1(x_1) \geq b_2(x_2)$. If player 1 would have lost the auction with $b_1(x_1) = x_1 + x_{max}$ he receives a payoff of 0. Deviating to a bid above $x_1 + x_{max}$ can lead to winning the auction, but the payoff is at best 0, because the value of the object is at most equal to $x_1 + x_{max}$ and $b_2(x_2) \geq b_1(x_1) = x_1 + x_{max}$. So if player 1 deviates to $x_1 + x_{max} + \delta$ and player 2 bids $x_1 + x_{max} + \epsilon$ with $\delta > \epsilon \geq 0$, player 1 faces a loss of at least $\epsilon \geq 0$.

□

Proof. In the two player wallet game, conducted as second-price sealed bid auction, bidding $b_1(x_1) = \alpha \cdot x_1$ and $b_2(x_2) = \frac{\alpha}{\alpha-1} \cdot x_2$ are equilibrium strategies for any $\alpha > 1$.

Consider two players, indexed by $i = 1, 2$. Suppose, without loss of generality, player 1 follows the bidding rule $b_1(x_1) = \alpha \cdot x_1$. Player 2 knows that the price he has to pay in the winning case is equal to $p = \alpha \cdot x_1$. So winning is beneficial for him as long as $x_1 + x_2 \geq p \Leftrightarrow x_2 + \frac{p}{\alpha} \geq p \Leftrightarrow x_2 \geq p \cdot \frac{\alpha-1}{\alpha} \Leftrightarrow p \leq \frac{\alpha}{\alpha-1} \cdot x_2$. Since we have a second-price auction player 2 will bid exactly $b_2(x_2) = \frac{\alpha}{\alpha-1} \cdot x_2$.

Now suppose player 2 follows the bidding rule $b_2(x_2) = \frac{\alpha}{\alpha-1} \cdot x_2$. Player 1 knows that the price he has to pay in the winning case is equal to $p = \frac{\alpha}{\alpha-1} \cdot x_2$. So winning is beneficial for him as long as $x_1 + x_2 \geq p \Leftrightarrow x_1 + p \cdot \frac{\alpha-1}{\alpha} \geq p \Leftrightarrow x_1 \geq \frac{p}{\alpha} \Leftrightarrow p \leq \alpha \cdot x_1$. Since we have a second-price auction player 1 will bid exactly $b_1(x_1) = \alpha \cdot x_1$.

Hence bidding $b_1(x_1) = \alpha \cdot x_1$ and $b_2(x_2) = \frac{\alpha}{\alpha-1} \cdot x_2$ are equilibrium strategies for any $\alpha > 1$. Note that this argumentation does not relay on the distribution of the signals x_1 and x_2. □

C Robustness

Transformed profit

As an alternative measure for profitability I used a more robust measure of the profits in which any positive payoff is transformed into 1 and any negative payoff is transformed into -1. So there is only a distinction between whether an auction was won profitable, unprofitable or lost. Unlike the actual profit, the magnitude of this transformed profit is independent of the opponent's bid.

Figures 13 and 14 in Appendix A show the average profits (transformed) in stage I and II for different constellations of signal and information and compares them to the average profits (transformed) when the bidders would have unilaterally followed the symmetric Nash bidding rule. Table 6 reports the changes in average profits (transformed) from stage I to stage II for different constellations of signal and information and compares the respective values between the treatment and control group.

Constellation	Means (Std. deviation)			Observations A / B	Wilcoxon p-value
	Change in average profits (transformed)				
	Information (A)	*No information (B)*	*Difference*		
Low signals (HIGHER)	0.34	0.04	0.31	61 / 79	
	(0.60)	(0.49)			0.030
Low signals (LOWER)	-0.13	-0.06	-0.07	212 / 164	
	(0.35)	(0.29)			0.071
High signals (HIGHER)	-0.06	-0.01	-0.05	215 / 176	
	(0.23)	(0.08)			0.024
High signals (LOWER)	0.22	0.03	0.19	97 / 76	
	(0.56)	(0.33)			0.041
All signals (HIGHER)	0.03	0.01	0.02	276 / 255	
	(0.38)	(0.28)			0.498
All signals (LOWER)	-0.02	-0.03	0.01	309 / 240	
	(0.46)	(0.30)			0.796

Notes: The last column reports the p-values of a Clustered Wilcoxon rank sum test using the Rosner-Glynn-Lee method (Rosner et al., 2003). Clusters are on subject level.

Table 6: Summary Table - Change in average profits (transformed)

Individual means

Tables 7 and 8 are equivalents to Tables 3 and 4, but with the difference that the observations are aggregated on the subject level. For each subject and constellation the individual mean of the respective outcome variable was used.

Constellation	Means (Std. deviation)			Observations A / B	Wilcoxon p-value
	Change in absolute deviation from Nash bid				
	Information (A)	*No information (B)*	*Difference*		
Low signals (HIGHER)	-13.12	-8.31	-4.81	24 / 25	
	(12.70)	(21.29)			0.027
Low signals (LOWER)	13.65	3.94	9.70	39 / 31	
	(11.23)	(15.33)			0.000
High signals (HIGHER)	6.80	1.22	5.58	38 / 32	
	(11.08)	(10.62)			0.000
High signals (LOWER)	-18.79	-3.22	-15.57	26 / 24	
	(15.78)	(11.70)			0.000
All signals (HIGHER)	1.85	-1.40	3.25	38 / 32	
	(8.88)	(10.32)			0.059
All signals (LOWER)	4.76	2.07	2.68	39 / 31	
	(12.41)	(14.46)			0.055

Notes: The last column reports the p-values of a Wilcoxon rank sum test.

Table 7: Summary Table - Deviation from Nash bid (individual means)

Constellation	Means (Std. deviation)			Observations A / B	Wilcoxon p-value
	Change in average profits				
	Information (A)	*No information (B)*	*Difference*		
Low signals (HIGHER)	6.17	3.48	2.69	24 / 25	
	(8.67)	(13.09)			0.031
Low signals (LOWER)	-3.86	-1.45	-2.42	39 / 31	
	(6.37)	(4.00)			0.035
High signals (HIGHER)	-1.83	-0.38	-1.46	38 / 32	
	(4.72)	(2.12)			0.085
High signals (LOWER)	7.78	2.49	5.29	26 / 24	
	(13.67)	(7.15)			0.029
All signals (HIGHER)	-0.07	0.35	-0.41	38 / 32	
	(3.18)	(4.46)			0.398
All signals (LOWER)	-1.06	-0.82	-0.24	39 / 31	
	(7.46)	(4.13)			0.279

Notes: The last column reports the p-values of a Wilcoxon rank sum test.

Table 8: Summary Table - Change in average profits (individual means)

D Instructions

We would like to welcome you to this economic experiment! During the experiment you have the possibility to conduct a task that is explained in detail in the following instructions. In the experiment you can win a non-negligible amount of money. The amount of your payoff depends on your decisions, on the other participants' decisions and on chance. During the experiment it is forbidden to communicate with the other participants. Please read through the instructions at hand thoroughly. Should you have questions before or during the experiment, please raise your hand and an experimenter will come to your seat.

General Structure

The experiment consists of **two** parts. Now, **part 1** will be explained. After part 1 ends you will receive separate instructions for part 2. Your decisions in part 1 do not influence your payoff in part 2. During the whole experiment you can earn *Taler*. These will be converted into *Euros* after the experiment. The conversion rate is

10 Taler = 1 EURO

At the begin of the experiment you are endowed with **50 Taler**. Your experimental credit at the end of the experiment consists of these 50 Taler plus your profits and minus your losses in part 1 and part 2. If you lose more than 50 Taler in the course of the experiment, your experimental credit drops down to 0 Taler. At the end of the experiment you receive your experimental credit in **EUR**. Anyway, independent of your decisions in the course of the experiment, you will receive **4 EUR** show-up fee at the end of the experiment. Your final payout will be calculated as follows:

Final Payout = 4 EUR + experimental credit from part 1 and part 2 (in EUR)

Important remark

All numerical examples that are used in the instructions and later on in the control questions for exemplification consist of arbitrary values and do not give a hint for optimal behavior in this experiment!

Part 1

Basic idea

This experiment's underlying task is the following:

- Together with **one other player** you participate in an auction
- **You and the other player** receive each a randomly selected sealed **envelope that contains money**

- There are **red** and **blue** envelopes
- A **red** envelope contains a random integer amount between **0 and 10 Taler** (all values are equally likely)
- A **blue** envelope contains a random integer amount between **50 and 60 Taler** (all values are equally likely)
- Both colors are **equally likely**
- Thus, all together there are 22 different amounts an envelope can contain and every amount is **equally likely**
- The colors and the amounts in both envelopes are **independent of one another** and it is also possible that both players receive the same amount

The following combinations are possible:

Player 1 *Envelope (0-10 Taler)*	**Player 2** *Envelope (0-10 Taler)*
Player 1 *Envelope (0-10 Taler)*	**Player 2** *Envelope (50-60 Taler)*
Player 1 *Envelope (50-60 Taler)*	**Player 2** *Envelope (0-10 Taler)*
Player 1 *Envelope (50-60 Taler)*	**Player 2** *Envelope (50-60 Taler)*

Both players are allowed to open their own envelope. This implies that every player gets to know his own amount but not the other player's amount and color.
Then, both players participate in an auction, in which the highest bidder can win **both envelopes** and the money the envelopes contain.
Both players can submit a bid once. The highest bidder wins **both envelopes** and pays the bid of the inferior bidder. The inferior bidder does not receive an envelope and does not have to pay anything - thus, he does neither make profit nor losses.

The rules in detail

Bidding
You and the other player can submit once an *integer* bid between **0 and 120 Taler**. These bids are made *secretly*, i.e. the other player does not see what bid you have made and vice versa.

Winning and Losing
The winner and the payoff are determined as follows:
You win, if:

1. Your bid is *higher* than the other player's bid
2. Your bid equals the other player's bid and your envelope contains *more* money

If your bids as well as the contents of the envelopes are equal, both players receive a payoff of 0. If your bid is lower than the other players bid, you lose and receive a payoff of 0.

Payoff
If you have won the auction your payoff is calculated as follows:
You receive the money of **both envelopes** and pay for this the **other player's bid**. Thus, you do *not pay your own bid*, but the bid of the inferior bidder. This implies that in the winning case you must pay *at most* your own bid.
Is the amount of both envelopes higher than the bid you have to pay, you make profit. Is the amount of both envelopes lower than the bid you have to pay, you make a loss. If you have lost the auction, you receive a payoff of 0 - thus, you do neither make profit nor losses.

Example
Assume, you have 50 Taler in your envelope and the other player would have 10 Taler in his envelope (every player only knows his own amount). You bid 90 Taler and the other player bids 45 Taler (every player only knows his own bid). You win the auction, because you have submitted the higher bid. So, you win both envelopes and pay the other player's bid. Your profit in this round would be $60-45 = 15$ Taler. The other player's profit would be 0 Taler.

The course of the experiment

Trial phase
At first you will bid in 5 *trial rounds* for the envelopes. During these 5 rounds you do not play against another human player, but against a computer. These rounds are **not relevant for your payoff** and their only purpose is to gain an understanding of the game and its general course. In every round you and the computer will receive a randomly selected amount between 0 and 10 or 50 and 60. This implies for you that you see in every round a randomly selected integer number between 0 and 10 or 50 and 60 on your screen. **This number symbolizes the content of an envelope** (you can find an example screenshot at the bottom of this page). You only get to know your own amount, but not your computer opponent's amount and vice versa. Now, you can submit once (per round) any integer bid between 0 and 120 Taler. Overall you participate in 5 auctions. After every round (i.e. after every bid) you receive an immediate feedback on your hypothetical profit or loss. However, you do not get to know the opponent's bid or the amount the computer received. The computer is programmed to choose a bid that depends on his amount, i.e. the higher the computer's amount the higher the bid the

computer chooses.

[SCREEN 1]

Main phase
After the trial phase you bid in 15 rounds for the envelopes and now you can receive an actual monetary payoff. In the main phase you do not compete with a computer, but with a human player. Your opponent will be randomly drawn from this room. You do not know who your opponent is nor does your opponent. The general course will be similar as in the trial phase. This means, in every round you and the other player will see a respective amount on your screen, which is a randomly selected number between 0 and 10 or between 50 and 60 that **symbolizes the content of an envelope**. In every round you and the other player can submit any integer bid between 0 and 120 Taler. But now, you do not receive an immediate feedback after each bid and you only get to know at the very end of the experiment (**i.e. only when part 2 is finished**) your final payoff. Out of the 15 rounds **3 randomly selected rounds are relevant for your payoff**. The other 12 rounds do not influence your payoff. Your profit or loss from this three randomly selected rounds is added to or subtracted from your initial endowment of 50 Taler, respectively.

Example
Assume, round 1,2 and 3 are randomly selected for the payoff. In round 1 your payoff is 30 Taler, in round 2 your payoff is -5 Taler and in round 3 your payoff is 0 Taler. Your profit in part 1 would be 25 Taler. Your current experimental credit would be $50+25 = 75$ Taler.

[SCREEN 2]

Part 2

Now, you will bid for the envelopes once again. For this purpose you and your previous opponent will receive the **same amounts** as in part 1 once again in the **same order**. In contrast to the previous part you do *not bid directly* against your opponent, but against **his decisions** that he made in part 1. This implies that your opponent does *not make new decisions* in part 2 and he will bid in every round exactly as in part 1. As you only compete with your opponent indirectly, your decisions in part 2 do not influence the payoff of your opponent in part 2.
In **part 2** you are randomly assigned to a role: either **A** or **B**. Both roles are equally likely. Your role is determined before the beginning of the 15 rounds by a virtual coin toss. Your respective role is constant for all 15 rounds and is displayed at the top of the screen.

If your role is **A**, you see for every amount additionally on the screen, whether your respective bid in part 1 was **HIGHER** or **LOWER** than the other player's bid. If the bids are equal, it will also be displayed "**LOWER**" (thus, "HIGHER" means *strictly higher* and "LOWER" means *lower or equal*). If your role is **B**, you receive no further information in part 2 and instead of HIGHER or LOWER it is only displayed "???".
Independent of your role you are once again in every round allowed to submit a bid, which can be as in part 1 between 0 and 120 Taler. The rules for winning and losing are exactly as in part 1 and also your payoff is computed equally. As already in part 1 the same three randomly selected rounds determine your payoff (i.e. if round 1, 2 and 3 were selected in part 1, round 1, 2 and 3 also determine your payoff in part 2). Your profit or loss from part 2 is added to or subtracted from your current experimental credit. Also in part 2 you do not receive immediate feedback after every bid, but you get to know your final payoff only at the very end of the experiment.

Summarized

- Part 2 is a repetition of the main phase of part 1
- In part 2 you have the same opponent as already in part 1
- Now you do not compete with your opponent directly, but with his decisions he made in part 1
- Your opponent will bid in part 2 exactly as in part 1
- You are randomly assigned to either role A or B
- If your role is A, you additionally see, whether your bid in part 1 was higher or lower than the other player's bid
- If your role is B, you do not receive additional information

[SCREEN 3]

[SCREEN 4]

Example (role - A)
Part 1 - 1st round: You have 50 Taler in your envelope and the other player has 10 Taler in his envelope (every player only knows his own amount). You bid 90 Taler and the other player bids 45 Taler (every player only knows his own bid).
Part 2 - 1st round: Now you receive once again an envelope that contains 50 Taler and the other player again receives an envelope that contains 10 Taler (every player only knows his own amount). Now, you see on your screen that your bid in part 1 was **HIGHER** than the other player's bid (as 90 is larger than 45 - anyway, also in part 2 you do not

know the other player's exact bid). Now, you can submit any bid between 0 and 120 Taler once again. The other player bids 45 Taler as in part 1.

Example (role - B)
Part 1 - 1st round: You have 50 Taler in your envelope and the other player has 10 Taler in his envelope (every player only knows his own amount). You bid 90 Taler and the other player bids 45 Taler (every player only knows his own bid).
Part 2 - 1st round: Now you receive once again an envelope that contains 50 Taler and the other player again receives an envelope that contains 10 Taler (every player only knows his own amount). You do not receive further information about the other player's bid on your screen. Now, you can submit any bid between 0 and 120 Taler once again. The other player bids 45 Taler as in part 1.