

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

Johannes Moser
University of Regensburg

December 19, 2017

Abstract

There is evidence that bidders fall prey to the winner's curse because they fail to extract information from hypothetical events - like winning an auction. This paper investigates experimentally whether bidders in a common value auction perform better when the requirements for this cognitive issue - also denoted by contingent reasoning - are relaxed, leaving all other parameters unchanged. The overall pattern of the data suggests that the problem of irrational over- and underbidding can be weakened by giving the subjects ex ante feedback about their bid, but unlike related studies I also find negative effects of additional information.

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

Keywords: Hypothetical thinking, cursed equilibrium, winner's curse

*Johannes Moser: University of Regensburg, Department of Economics, Universitätsstraße 31, 93040 Regensburg, Germany; johannes.moser@ur.de. I would like to thank Florian Greindl, Julia Löhner, and Matthias Scherr for great research assistance. I would further like to thank the participants of the *Young Economists' Meeting 2017* in Brno and the participants of the *International Conference of the French Association of Experimental Economics 2017 (ASFEE)* in Rennes for useful comments and discussions. Johannes Moser acknowledges funding through the International Doctoral Program "Evidence-Based Economics" of the Elite Network of Bavaria.

1 Introduction

The *winner's curse* is a well-known empirical phenomenon in common value auctions, which was first described by [Capen, Clapp, and Campbell \(1971\)](#). They showed that many oil companies in the 1960's and 1970's had to report a drop in profit rates because of systematic overbidding in auctions for drilling rights. Later also experimental evidence for the winner's curse was found in a large number of lab studies (see, for example, [Bazerman and Samuelson, 1983](#); [Thaler, 1988](#); [Charness and Levin, 2009](#); [Ivanov, Levin, and Niederle, 2010](#)).

I show in an experimental setup that bidders are more likely to avoid the winner's curse when they are informed, already before submitting a bid, whether their bid is the winning bid or not. By giving the subjects this information, I weaken the requirement for them to condition their bid on winning the auction: winning or losing are now not hypothetical anymore, and the adverse selection issue of winning an auction with a common value for all bidders becomes more salient. The findings of my paper suggest that mistakes in hypothetical thinking seem to explain a substantial part of the winner's curse. Thus, approaches that focus on this mental process might be more suitable to explain this phenomenon than belief-based models like cursed equilibrium ([Eyster and Rabin, 2005](#)), which state that the winner's curse is mainly driven by inconsistent beliefs. Herewith, this paper attempts to shed light on the ongoing debate whether models, focusing on the erroneous belief formation of individuals, like cursed equilibrium and level- k ([Eyster and Rabin, 2005](#); [Crawford and Iriberri, 2007](#))¹ or approaches concerning contingent reasoning on hypothetical future events ([Charness and Levin, 2009](#); [Ivanov, Levin, and Niederle, 2010](#); [Esponda and Vespa, 2014](#)) are more suitable to explain the winner's curse.

Contingent reasoning refers to the ability of thinking through hypothetical scenarios and to perform state-by-state reasoning. 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, 2007](#); [Nickerson, 2015](#); [Singmann, Klauer, and Beller, 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](#); [Ngangoue and Weizsäcker, 2015](#); [Levin, Peck, and Ivanov, 2016](#);

¹Both 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 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 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 *fully* cursed equilibrium.

Esponda and Vespa, 2016; Li, 2017; Koch and Penczynski, 2017).

In contrast to belief-based models, like cursed equilibrium, the concept of contingent reasoning has, however, still received very little formal treatment. Li (2017) represents the first attempt to capture this mental process formally by introducing the concept of *obviously strategy-proof* (OSP) mechanisms.² A major contribution of Li’s paper is to explain why subjects perform better in ascending bid auctions compared to sealed bid auctions. However, it cannot explain why common value auctions might be more challenging for the bidders than private value auctions. For this reason, the concept of OSP mechanisms is not sufficient to fully capture the most relevant aspects of the winner’s curse.

Can a broader definition of contingent reasoning explain why bidders fall prey to the winner’s curse? To answer this question I will go back one step from OSP mechanisms and focus only on the events of winning or losing an auction - which provide information about the true value of a good in common but not in private value auctions. Bidders in common value auctions now face two cognitive hurdles. First, they have to be able to recognize and identify several hypothetical scenarios which might possibly occur (e.g., winning or losing an auction) and second, the bidders have to be able to infer information from such hypothetical events. The question is which of these cognitive tasks is more challenging for subjects and to what extent they affect the likelihood of the winner’s curse to occur. There are three possibilities: (i) bidders are perfectly able to perform state-by-state reasoning, but they neglect the informational content of the other bidders’ actions and, hence, the informational content of winning an auction as proposed by cursed equilibrium; (ii) bidders take into account that the bids of the other players are correlated with their signals, but they are not able to identify the relevant states to condition on in the first place; (iii) bidders are neither able to perform state-by-state reasoning, nor do they take into account the informational content of winning an auction.

This paper tests experimentally whether subjects in an auction are able to infer information from the events of winning or losing if these are not hypothetical anymore. For this purpose, I constructed a second-price auction in which the bidders learn whether a bid, which was considered optimal *ex ante*, is the winning bid or not (but without learning their payoff yet) and they receive 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

²According to Li (2017), a mechanism is OSP if and only if an optimal strategy can be found without the necessity of performing contingent reasoning. However, Li’s definition of contingent reasoning is very strict in a game-theoretic sense as it does not only involve conditioning on different broader states of the world or future events (such as conditioning on being the winner in an auction or conditioning on being the pivotal voter in an election), but also conditioning on each possible decision node in a given information set. Additionally, there is only a differentiation between whether a game is OSP or not, but no distinction between “more” or “less” contingent reasoning.

pivotal or not before they had to cast a vote.

The auction model I used in the experiment is based on 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 $v = x_1 + x_2$. This game is played in 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. In this sense, the bidders receive information about some realized event before they have to come up with a bid.³ Apart from knowing whether their bid is the winning or losing bid, the subjects face exactly the same decision problem as in stage I.

My design allows to clearly distinguish between mistakes in contingent reasoning and cursed equilibrium since two crucial assumptions of the latter are that (i) no (or only a partly) correlation between the other players' actions and types is assumed and (ii) bidders still best-respond given their beliefs. 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. For fully cursed bidders this argument is straightforward, but it also holds for partly cursed bidders because of the best-response assumption. By definition, the bid of a cursed bidder is already evaluated conditional on winning in stage I, albeit a partly cursed bidder implicitly assumes that winning is less informative than in equilibrium. This means, the feedback of winning in stage II provides no further relevant information, since it is already included in the decision of a partly cursed bidder in stage I.

The reason for choosing the wallet game, instead of a more standard model for common value auctions, is that in this game a naïve bidding strategy can lead to both over- and underbidding, relative to the symmetric equilibrium strategy, depending on whether the private signal is low or high. In this sense, there can be both a winner's and a loser's curse (see also [Holt and Sherman, 1994](#)).⁴ This property is useful for two reasons. First, I am able to control for psychological explanations, stating that the winner's curse is mainly driven by emotional factors of winning (see, for example, [Van den Bos, Li, Lau, Maskin, Cohen, Montague, and McClure, 2008](#); [Astor, Adam, Jähnig, and Seifert, 2013](#)). Second, and more importantly, I am able

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

⁴A loser's curse can be understood as losing an auction, but the bidder could have won with a positive payoff.

to check whether bid shading in stage II is due to proper Bayesian updating or just a rule of thumb when learning that a certain bid was the winning bid. In my setup, bid shading is only advisable for low signals, but not for high signals. Thus, the subjects have to differentiate between these two kinds of signals, instead of following the simple decision rule “*decrease your bid, when you learn that your bid was the winning bid*”.

The findings of my paper reveal two important observations:

(i) Bidders are more likely to avoid the winner’s curse and the loser’s curse in stage II when they are informed whether their bid is the winning bid or not, given that the respective information has a sufficiently high predictive power concerning the opponent’s signal. This suggests that the crucial cognitive hurdle for bidders in a common value auction is not forming beliefs about the opponents’ behavior, but identifying the relevant states to condition on in the first place.

(ii) Information can also be negative for the bidders, depending on the context. The subjects in my experiment differentiated only imperfectly between situations in which decreasing (increasing) a bid is rational and those in which it is not and they often used simple heuristics instead of making strategic changes. This behavior might be partly explained by an actual *joy of winning* or, to be more precise, a disappointment of losing. When subjects learned that they lost an auction in stage I, most of them increased their bid in stage II, no matter whether they had a low or high private signal. Conversely, when subjects learned that they won an auction in stage I, they acted more strategically and bids for low signals were decreased at a much higher rate compared to bids for high signals (78.7% vs. 42.3%).

This paper is organized as follows. Section 2 will provide an overview about the current literature closely related to my research topic. Section 3 will present the underlying theoretical model for the experiment. Section 4 describes the experimental design. Section 5 presents and discusses the results of the experiment. Section 6 concludes the paper.

2 Literature review

This paper is the first to investigate the direct effect of the information of winning (or losing) in the context of a common value auction. The novelty of my design is that I use a dynamic choice setting à la [Esponda and Vespa \(2014\)](#) in a sealed bid auction. Similar as them, I want to differentiate between mistakes in hypothetical thinking and problems with extracting information from the opponents’ actions. With this setup I am able to test whether potential errors in common value auctions rather occur due to an inability of performing contingent reasoning or due to inconsistent

beliefs as proposed by cursed equilibrium. To the best of my knowledge, none of the recent approaches provided a similar framework. For example, the experiment by [Charness and Levin \(2009\)](#) is rather a single decision maker problem in an adverse selection environment. Additionally they measured the effect of contingent reasoning only indirectly by transforming their initial game into a set of simple lotteries. [Ivanov, Levin, and Niederle \(2010\)](#) showed that belief-based models, like cursed equilibrium, might be not that powerful in explaining the winner's curse, but they provided no alternative explanation. The paper of [Koch and Penczynski \(2017\)](#) is also closely related to my work. So far they are the only ones who combined contingent reasoning and belief-based models in their experiment. However, they focused on a different aspect of contingent reasoning, not directly related to the event of winning an auction. Finally, [Levin, Peck, and Ivanov \(2016\)](#) used Dutch auctions and showed that in this format conditioning on winning is more salient compared to a strategically equivalent first-price auction. All of these papers will be discussed in more detail in the following part.

As explained above, the paper most closely related to my own work is the one by [Esponda and Vespa \(2014\)](#). They created a common value voting experiment where a subject interacted with two computers. The main task for the subject was to submit a vote for a ball which was either red or blue. Due to the commonly known voting algorithm of the computers, the vote of the subject was only relevant when the ball was blue and, hence, the optimal choice for the subject was always to vote for blue. [Esponda and Vespa \(2014\)](#) observed that subjects made significantly less errors in a sequential election, where the voters knew whether they were pivotal or not, compared to a simultaneous election, where the voters were not informed about their pivotality. Similarly, in their lab experiment [Ngangoue and Weizsäcker \(2015\)](#) found that traders in a financial market setting performed better in a sequential trading mechanism where no Bayesian updating on hypothetical events was required.

[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, Levin, and Niederle \(2010\)](#) used a similar approach as [Charness and](#)

Levin (2009), but they conducted their experiment in an actual auction context using the *maximal game*.⁵ The aim of their experiment was not to look at the effect of contingent reasoning, but rather to disprove that the winner’s curse is driven by inconsistent beliefs. Ivanov, Levin, and Niederle (2010) observed significant over-bidding which was not reduced in a modified setting of the maximal game where the beliefs of the participants were explicitly formed and, thus, belief-based models, like cursed equilibrium, had little explanatory power. However, Costa-Gomes and Shimoji (2015) criticized Ivanov, Levin, and Niederle (2010) for the misuse of some of the game theoretical concepts, arguing that some of their findings can indeed be explained by belief-based models. Similarly, Camerer, Nunnari, and Palfrey (2016) argued that the observed behavior of the subjects in Ivanov, Levin, and Niederle (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 relaxed and stochastic choices are allowed. They showed that this extended model fits very well to the data of Ivanov, Levin, and Niederle (2010).⁶

My approach is different from the one in Ivanov, Levin, and Niederle (2010). Rather than showing that bidders do not best-respond, even though their beliefs are explicitly formed, I show that bidders in my setup react on information which is, by definition, not relevant for the updating process of a “cursed” bidder.⁷

Koch and Penczynski (2017) 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. The authors found that the relaxation of both cognitive requirements had a significant effect on avoiding the winner’s curse. To identify the effect of contingent reasoning on the bidding behavior of the subjects, the authors used a transformed version of the original game, where both bidders received the *same* information about an object with a stochastic value, instead of *different* private signals. In this sense, there was no need for the bidders to condition on whether their own signal was relatively low or high. This is different from my approach, where I focus on the direct effect of the information of winning

⁵Ivanov, Levin, and Niederle (2010) criticized that the acquiring a company game, used in Charness and Levin (2009), represents a lemon market and not a common value auction. Thus, it seems 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, Levin, and Niederle (2010) also used a computer treatment, but in their case the computer mimicked the subject’s own past strategy.

⁶However, Camerer, Nunnari, and Palfrey (2016) also assume that, if the *perfect* best-reply assumption of cursed equilibrium and level- k model is maintained, these models are very bad in predicting the behavior in maximal value games.

⁷In the setup of Ivanov, Levin, and Niederle (2010) one could still argue that the belief formation of a cursed bidder takes place in an isolated “black box” and is not affected by information about the opponent’s behavior which is given to such a bidder. This is not an issue in my design.

an auction. Additionally I only changed one parameter in stage II, whereas [Koch and Penczynski \(2017\)](#) used two different games, but with the same best-response functions and equilibria.

[Levin, Peck, and Ivanov \(2016\)](#) 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, Harstad, and Levin \(1987\)](#) 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, and showed that both cognitive skills had a significant effect on avoiding the winner’s curse, resulting in higher earnings for the subjects. The authors also showed that subjects performed better in a Dutch auction, where the auction ended when the first subject stopped the clock, compared to a Dutch auction, where the auction ended when the last subject stopped the clock and no subject received any feedback about whether he is the highest bidder. In the first setting, a bidder knows, in the moment of stopping the clock, that he is the highest bidder and, hence, it is more salient for him to condition his bid on winning compared to the second setting with a “silent” clock. The authors concluded that bidders can better handle the winner’s curse if the requirements for this form of hypothetical thinking are reduced.⁸ Note that this form of a Dutch auction is not comparable to the sequential approach of [Esponda and Vespa \(2014\)](#) 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 stage II receive this information *before* they have to come up with a bid.

3 The model

In the following section I present a formal description of the auction game used in the experiment. This game is based on the *wallet game* proposed by [Klemperer \(1998\)](#) and the model used in [Avery and Kagel \(1997\)](#). Henceforth, this game will be denoted as the *wallet game*, although it is slightly different from the original one.¹⁰

There are two players, indexed by $i = 1, 2$. Each player i receives a signal x_i from

⁸As a further robustness check, the authors also used private-value auctions, where they did not find a significant difference between these two kinds of Dutch auctions.

⁹A 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, Krämer, and Weinhardt, 2012](#); [Adam, Krämer, and Müller, 2015](#)).

¹⁰A detailed analysis of the general wallet game with N bidders can be found in [Eyster and Rabin \(2005\)](#) and [Crawford and Iriberry \(2007\)](#).

the set $X = \{0, 1, \dots, 9, 10, 50, 51, \dots, 59, 60\}$ ($|X| = 22$), with each value equally likely and with replacement (so there are 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. The players are allowed to choose a bid b_i in the range of $[0, 120]$, with only integer values possible. In case of a tie, the player with the higher signal wins. If the signals are also equal, both players receive a payoff of 0. The payoff of player 1 (analogously for player 2) is thus given by:

$$\pi_1 = \begin{cases} x_1 + x_2 - b_2 & \text{if } b_1 > b_2 \\ x_1 + x_2 - b_2 & \text{if } b_1 = b_2 \wedge x_1 > x_2 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

The utility function of both players is assumed to be symmetric across i such that $u_i(\mathbf{x}) = v = x_1 + x_2$ with $\mathbf{x} = [x_1, x_2]^T$ (see also Crawford and Iriberry, 2007).

Proposition 1. *Given that player i sees a signal x_i , the expected value of the object is given by $E[V|X_i = x_i] = x_i + 30$.*

The proof of Proposition 1 is straightforward and, hence, omitted. Henceforth, bidding according to $b_i(x_i) = x_i + 30$ will be denoted as *naïve* bidding, since the bid is not evaluated conditional on winning. It is easy to show that this bidding function cannot be part of a symmetric equilibrium. For example a naïve bidder with a low signal ($x_i \in \{0, \dots, 10\}$) only wins when the other naïve player also has a low signal and the price to pay is always above (or equal to) 30, which results in a negative payoff for the winning bidder.

Proposition 2. *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 $\alpha > 1$. For $\alpha = 2$ we have the unique symmetric equilibrium with $b_i(x_i) = 2 \cdot x_i$ for all $i \in \{1, 2\}$.*

The proof of Proposition 2 can be found in Appendix B.¹¹ The intuitive explanation for the symmetric equilibrium is the following: under the assumption that all bidders have the same bidding function $b_i(x_i)$, which is monotonically increasing in x_i , a rational player anticipates that he only wins if his signal is at least as high as the signal of the opponent, thus, barely if $x_1 = x_2$ holds. Hence, the optimal bid in (the symmetric) equilibrium is given by $b_i(x_i) = 2 \cdot x_i$. Henceforth, bidding according to $b_i(x_i) = 2 \cdot x_i$ will be denoted as *sophisticated* bidding. Note that these equilibrium strategies are neither affected by risk preferences nor the distribution of the signals x_1 and x_2 (see also Klemperer, 1998).

¹¹The proof of uniqueness of the symmetric equilibrium is not presented in this paper, but it can be found in Eyster and Rabin (2005) and Crawford and Iriberry (2007) for the general wallet game.

Proposition 3. *Any bid $b_i(x_i)$ outside the interval $[x_i, x_i + 60]$ is weakly dominated.*

The proof of Proposition 3 can be found in Appendix B. The intuitive explanation is the following: bidding below the private signal can never be optimal in a second-price auction because the value of the good is at least x_i . Bidding above $x_i + 60$ is, likewise, never optimal, since the maximal value of the good is at most $x_i + 60$.

Proposition 4. *All equilibria from Proposition 2, except the symmetric one with $\alpha = 2$, involve weakly dominated bids for one player for at least some signals.*

The proof of Proposition 4 can be found in Appendix B. Because of this issue, asymmetric equilibria are less plausible, since they involve weakly dominant bids for at least some x_i (for one player). Additionally, some bids resulting from a bidding function $b_i(x_i) = \alpha \cdot x_i$ with $\alpha > 2$ are not even feasible because the bidding range is restricted on the interval $[0, 120]$. Consequently, throughout the the analysis of this paper, I will only focus on the symmetric equilibrium, with $b_i(x_i) = 2 \cdot x_i$, as a benchmark for sophisticated bidding. Hence, we have the following benchmark bidding functions for naïve (NVE) and sophisticated (BNE) bidding:

$$b_i^{NVE}(x_i) = x_i + 30 \tag{2}$$

and

$$b_i^{BNE}(x_i) = 2 \cdot x_i. \tag{3}$$

At this point it seems important to explain the reasons for choosing a modified model of the wallet game, with discrete signal space and a gap in the middle, instead of a continuous range between 0 and 60. First, in the original wallet game, the naïve and sophisticated bidding functions get closer, the more one reaches the middle of the signal space. For the value in the middle, which would be 30 in a continuous range, both bidding functions are even identical. Thus, in the original wallet game, it gets more difficult to distinguish between naïve and sophisticated bidding for these intermediate values. Second, in my 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 always win for high signals and to always lose for low signals. This is given when following the sophisticated bidding rule. This makes $b_i^{BNE}(x_i) = 2 \cdot x_i$ a stronger and more credible benchmark for sophisticated bidding, although it is still not a weakly dominant strategy.

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

or level- k model (Eyster and Rabin, 2005; Crawford and Iriberri, 2007).¹² 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

4.1 Implementation

The experiment was conducted in the Regensburg Economics Science Lab (RESL) in February 2017. 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, 2004). In total, 5 sessions were conducted with overall 72 participants (mostly undergraduate students from various fields). For each session I had between 10 and 18 participants (always an even number).

The experimental currency unit (ECU) used in the experiment were *Taler*. All 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. The sessions lasted between 60 and 75 minutes.

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.

4.2 Basic setup

The experiment is divided into *two stages*, both of which are payoff-relevant. 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

¹²A 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 “cursed” bidder implicitly assumes that the opponent will bid independently of his signal which makes $b_i(x_i) = x_i + 30$ a best-response.

set $X = \{0, 1, \dots, 9, 10, 50, 51, \dots, 59, 60\}$ *with replacement*).¹³ Each subject gets randomly matched (fixed stranger matching) with another subject of the group (e.g., subject k and subject l). 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 9 in Appendix A.

Before starting with the actual task, all participants are asked to answer eight control questions (see Appendix D) 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 represented 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.¹⁴

Stage II is, from a theoretical point of view, a repetition of stage I. All subjects receive the same 15 signals as in stage I, in the same order. In stage II, the subjects play against a 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.¹⁵ 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 (and with equal likelihood) assigned to either treatment INF (*information*) or NOINF (*no information*).

¹³The actual explanation in the experiment is that each player receives an envelope with a random amount of money inside (see *Instructions* in Appendix C).

¹⁴The 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 $b_i^{BNE}(x_i) = 2 \cdot x_i$ is much more harmful when the other player uses the naïve bidding function. If the computer had used the sophisticated bidding function, the subjects would have been less likely to realize that deviating from this strategy is a bad idea. In the presented setup $b_i^{BNE}(x_i) = 2 \cdot x_i$ is a best-response for both the naïve and the sophisticated bidding function.

¹⁵In 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.

4.3 Treatments

INF (*Information*)

The subjects who receive treatment INF 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.¹⁶ For example if a subjects sees that his initial bid of $b_i^I = \bar{z}$ was HIGHER than the bid of his opponent, he knows that submitting a bid of $b_i^{II} = z \geq \bar{z}$ results in winning the auction for sure. Conversely, if a subjects sees that his bid $b_i^I = \underline{z}$ was LOWER than the bid of his opponent, he knows that submitting a bid of $b_i^{II} = z < \underline{z}$ 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 a typical screen is given by Figure 10 in Appendix A.

NOINF (*No information*)

The subjects in treatment NOINF face exactly the same situation as those in treatment *INF*, except 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 NOINF know how treatment INF looks like and vice versa). An example of a typical screen is given by Figure 11 in Appendix A. The general structure of the treatments is illustrated by Figure 1.

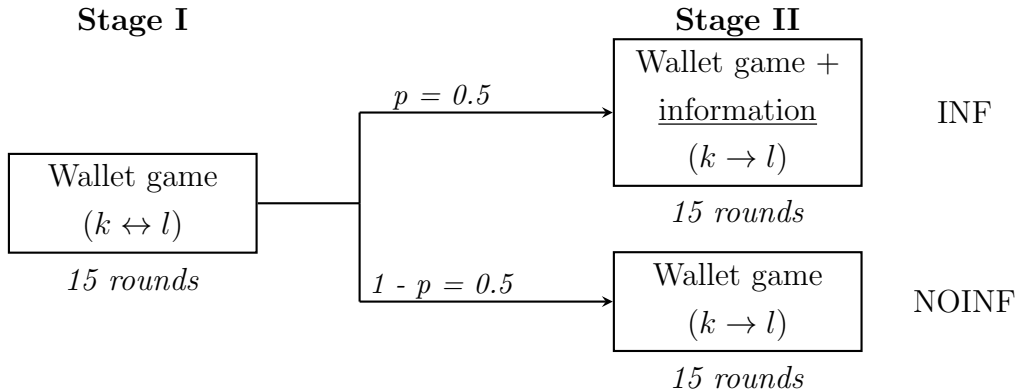


Figure 1: Illustration of the treatments

¹⁶If the bids are equal, the subjects also get the message “LOWER”, so LOWER means lower or equal.

4.4 Behavioral predictions

The following behavioral predictions are based on the assumption of a boundedly rational agent who is cognitively limited in the sense that he is not able to infer information from hypothetical events, like winning or losing an auction. Such an agent will be denoted as *naïve*.¹⁷ A naïve bidder will behave as if the auction is a private value auction since he ignores the adverse selection in the event of winning (for low signals) and the positive selection in the event of losing (for high signals) and will choose a bid that is based on the expected value of the good if he is risk-neutral. Hence, he will bid $b_i(x_i) = x_i + 30$ for any given signal.

If all bidders behave like this, there will always be a winner's curse for the bidders who win with a low signal ($x_i \in \{0, \dots, 10\}$), since they only win when the other player also has a low signal and the price to pay is always above (or equal to) 30. Conversely, all bidders who lose with a high signal ($x_i \in \{50, \dots, 60\}$) fall prey to a loser's curse, since they could have won the auction at a profitable price.

How will such a cognitively limited agent react, when he receives the information HIGHER or LOWER - i.e., when he learns whether his ex ante, as optimally considered bid, is the winning bid or not? If such an agent is able to correctly extract information from *observed events*, which implies that he has consistent beliefs about the behavior of the opponent, he should update his bids at least for some constellations of signal and information. Consider therefore Table 1 and two players, 1 and 2. Player 1 is a naïve bidder who bids according to $b_1(x_1) = x_1 + 30$ in stage I. Player 2 is the opponent of player 1. As a minimum requirement for rationality, player 2 does not use any weakly dominated bids.¹⁸ The first columns show the four different constellations of signal and information that can occur in the experiment from the perspective of player 1. The last column shows the prediction of player 2's signal when player 1 receives the respective information and given that he believes that his opponent does not use any weakly dominated bids.

Proposition 5. *If a player bids according to $b_i(x_i) = x_i + 30$ in stage I and the other player does not use any weakly dominated bids, he can predict in stage II whether the signal of the other player is low or high when (i) he wins with a low signal or (ii) he loses with a high signal.*

¹⁷Note that this paper will not provide a generalizable formal model of naïve behavior as it can be found, for example, in Li (2017).

¹⁸This means for each x_2 he may choose any bid in a range of $[x_2, x_2 + 60]$. So, for low signals the bid can be in a range of $[0, 70]$ and for high signals in a range of $[50, 120]$.

<i>Constellation</i>	Signal	Information	Prediction of other’s signal
ls^H	Low signal	HIGHER	<i>Low signal</i>
ls^L	Low signal	LOWER	<i>Low or high signal</i>
hs^H	High signal	HIGHER	<i>Low or high signal</i>
hs^L	High signal	LOWER	<i>High signal</i>

Table 1: Prediction of other player’s signal for different constellations

The proof of Proposition 5 is provided in Appendix B. We can see that for constellations ls^H and hs^L the respective information provides unambiguous hints about the opponent’s signal for player 1. While for these constellations the weak assumption of an opponent who does not use any weakly dominated bids is sufficient to correctly predict his signal (low or high)¹⁹, this is not the case for constellations ls^L and hs^H where the probabilities for a low or high signal depend on the specific bidding function of the opponent. So for the latter two constellations a sophisticated computation of probabilities is required to predict the expected value of the opponent’s signal correctly. Based on this, clear predictions can be made for constellations ls^H and hs^L , but not for constellations ls^L and hs^H .

- In constellations ls^H and hs^L , bidders will update their bids in stage II, resulting in bids closer to the Nash prediction and higher earnings for the subjects.
- For constellations ls^L and hs^H , the behavior of the bidders can be ambiguous since the expected value of the opponent’s signal depends on the specific beliefs of the bidders and the updating process requires non-trivial computations.

As a concluding remark, it is important to note that cursed equilibrium cannot explain an updating of bids in stage II, since a “cursed” bidder implicitly assumes that the opponent chooses a bid which is independent of his signal, hence the respective information would not be useful for such a bidder.²⁰

¹⁹To be more precise: for constellations ls^H and hs^L it is sufficient to assume that player 2 does not bid more than $x_2 + 60$ for low signals (no weakly dominated overbidding) and not less than x_2 for high signals (no weakly dominated underbidding). Since 91.94% of all bids in stage I fall into this category, this is actually a plausible belief.

²⁰Eyster and Rabin (2005) argued that cursed equilibrium is basically not defined for sequential games and that players in sequential games might be less “cursed” than in simultaneous games. However, in my setup the players only observe whether their bid was higher or lower than the bid of the opponent, but they do not observe the specific action of the other player, as, for example, in an ascending bid auction. Additionally, it can be argued that the concept of “cursedness” would be a rather weak one if “cursedness” would suddenly vanish in the moment of revealing the other player’s action.

5 Results

In this section, the results of the experiment will be presented. First, I will give a descriptive overview about the overall bidding pattern in stage I, before the subjects are assigned to a treatment group. Second, I will present the effects of the *information* treatment (INF) in terms of bidding behavior and the resulting profits.²¹ Finally, I will provide evidence that the observed behavior is to a not negligible 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”.

All bids and profits in the following part are expressed in terms of *experimental currency units* (ECU) with an exchange rate of 1 Euro = 10 ECU.

5.1 Overall bidding pattern in stage I

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 Figure 2). Especially for high signals the median bids are fitted very well by the naïve bidding function $b_i^{NVE}(x_i) = x_i + 30$. 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 ($p = 0.624$).




Figure 2: Median bids in stage I

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).

²¹When I compare the effects between INF and NOINF, for different constellations of signal and information, I look at those subjects in treatment NOINF who *would have* received the respective information if they had been in treatment INF, in order to get an appropriate control group.

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. The rate of the winner’s curse is very high, 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. A further conclusion is that emotional factors of winning seem to play a minor role in stage I, since the rates for winner’s and loser’s curse are very similar.

	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.

5.2 Bidding behavior in stage II

In this subsection I provide histograms and kernel density plots of the bids in stage I and II for different constellations of signal and information for both treatment INF and NOINF. As explained in the Section 3, the symmetric Nash bidding function $b_i^{BNE}(x_i) = 2 \cdot x_i$ serves as a benchmark for sophisticated bidding. The vertical line in the graphs marks the upper bound of Nash bids for low signals ($b_i = 20$) and the lower bound of Nash bids for high signals ($b_i = 100$), respectively.

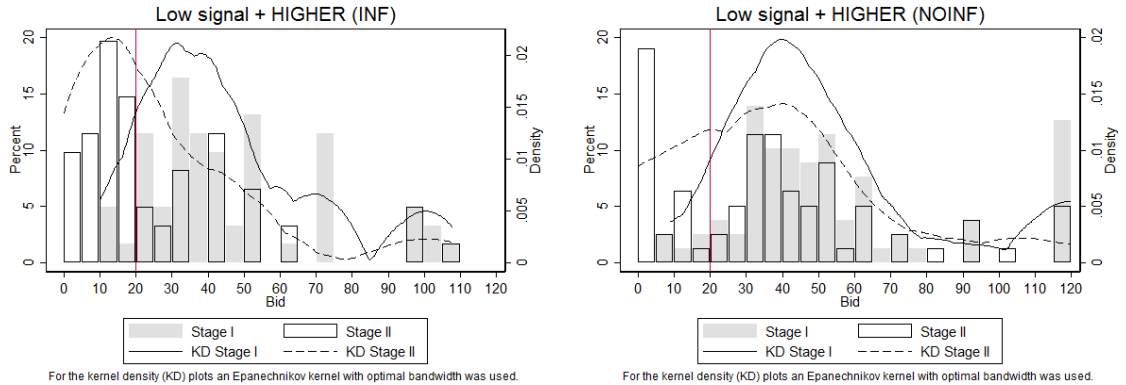


Figure 3: Constellation: Low signal + HIGHER (ls^H)




Figure 4: Constellation: Low signal + LOWER (ls^L)




Figure 5: Constellation: High signal + HIGHER (hs^H)




Figure 6: Constellation: High signal + LOWER (hs^L)

Most noticeable are the changes in bids in constellations ls^H and hs^L for the subjects in treatment INF (Figures 3 and 6). We can see that in those constellations, the mass of bids in stage II is closer to the Nash prediction than before. For the other constellations in treatment INF (ls^L and hs^H), the effect is reverse, albeit the changes are less distinct (see Figures 4 and 5). Here the mass of bids is shifted to the right for low signals and to the left for high signals - hence, in both cases further away from the Nash prediction. The bids in the NOINF treatment were also changed, although the subjects received no information about their bids, but here the pattern is less distinct and uniform.

In the left part of Figure 3 we can see that there is an accumulation of bids below 20 in stage II for low signals, when receiving the information HIGHER (72.92% of the newly selected bids, which are lower than before, are below or equal to 20).²² This in fact indicates that the subjects indeed realized that if they win with a low signal, the other player has most likely a low signal as well and, hence, they bid exactly for this case (if both players have low signals, the *maximal* value of the good 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, a slight accumulation is still noticeable which is, however, less distinct. Here only 42.53% of the newly selected bids, which are higher than before, are above or equal to 100.

This leads to the following conclusions. If we only look at those constellations in which the respective information provides a relatively unambiguous hint about whether the opponent's signal is low or high (ls^H and hs^L), the *information* treatment (INF) is indeed helpful for the bidders to adjust their bids. For the other constellations we can observe a reverse effect. Possible reasons for this behavior will be discussed in the further course of this paper.

5.3 Profits

The profits of the bidders are calculated as defined by equation (1) in Section 3. The change in profits from stage I to stage II for individual i in round m is then given by

$$\Delta\pi_{im} = \pi_{im}^{II} - \pi_{im}^I \quad (4)$$

Hence, a positive value of $\Delta\pi_{im}$ implies that the subject's payoff increased from stage I to stage II, whereas a negative value implies the opposite. Table 3 reports OLS regressions with clustered standard errors, testing whether the *information*

²²Newly selected bid means that the bid in stage II is different from the one in stage I.

treatment (INF) has a significant effect on the changes in profits from stage I to stage II for different constellations of signal and information.

	Change in profits						
	ls^H	ls^L	hs^H	hs^L	H	L	<i>Overall</i>
INF	3.773 (4.907)	-1.971 (1.327)	-1.414* (0.758)	5.115* (2.644)	-0.469 (1.482)	0.246 (1.284)	-0.158 (1.001)
Constant	2.063 (4.534)	-1.793* (0.982)	-0.205 (0.207)	0.566 (1.388)	0.498 (1.396)	-1.046 (0.796)	-0.251 (0.856)
Observations	140	376	391	173	531	549	1080
Subjects	49	70	70	50	70	70	72
R^2	0.014	0.010	0.015	0.043	0.001	0.000	0.000

Notes: Cluster-robust standard errors (on the subject-level) are in parentheses. Constellation ls^H stands for low signals paired with the information HIGHER. Constellation L stands for all signals paired with the information LOWER. The other constellations are defined analogously.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Regression Table - Change in profits

Overall, we can see that the change in average profits follows the same pattern as the deviation from the Nash bid: if the subjects bid closer to the Nash prediction in stage II, their profits increase, and vice versa. However, the effects, although pointing into the right direction, are only significant for constellations hs^H and hs^L and here only at the 10%-level. The coefficient in constellation hs^L indicates that the subjects in treatment INF increased their profits per auction from stage I to stage II by 5.115 ECU, on average, when they learned that they lost an auction with a high signal, compared to the subjects in treatment NOINF who did not get the information of losing. Conversely, in constellation hs^H , we can see that the profits per auction decreased from stage I to stage II by 1.414 ECU, on average, when learning that the bid for a high signal was the winning bid.

A major problem in second-price auctions, however, is that the magnitude of the profits is affected by the behavior of the opponents. Hence, the very same winning bid can lead to different profits, depending on the bid of the opponent.²³ For this

²³Consider the following example. Player 1 has a signal of 10 and player 2 has a signal of 5. Player 1 submits a bid of 40 in the stage I and player 2 submits a bid of 35. Player 1 wins and his payoff would be -20 in this scenario. If player 1 decreases his bid to 20 in stage II, he receives a payoff of 0 and improves his payoff by 20 ECU. Now suppose player 2 would have chosen a bid of 25 in stage I, with everything else exactly as before. Now player 1 improves his payoff by only 10 ECU, when he decreases his bid to 20 in stage II. Hence, in the first scenario the very same decreasing strategy is twice as profitable.

reason, I also used an alternative and more robust 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. Henceforth, I will call this *adjusted profit*. In contrast to the actual profit, the magnitude of the adjusted profit is not affected by the opponent’s bid which makes it a cleaner measure for sophisticated bidding. The adjusted profit has also a second interpretation: switching auctions with negative payoff into auctions with zero payoff in stage II refers to overcoming a winner’s curse. Switching auctions with zero payoff into auctions with positive payoff in stage II refers to overcoming a loser’s curse. Thus, the change in the adjusted profit from stage I to stage II can actually be interpreted as the (net) rate of switching unfavorable auctions into favorable auctions. Table 4 reports OLS regressions with clustered standard errors, testing whether the *information* treatment (INF) has a significant effect on the changes in adjusted profits from stage I to stage II, for different constellations of signal and information.

	Change in <u>adjusted</u> profits						
	ls^H	ls^L	hs^H	hs^L	H	L	<i>Overall</i>
INF	0.306** (0.152)	-0.071* (0.041)	-0.050** (0.025)	0.190** (0.092)	0.025 (0.041)	0.011 (0.040)	0.016 (0.030)
Constant	0.038 (0.115)	-0.061** (0.029)	-0.006 (0.006)	0.026 (0.051)	0.008 (0.035)	-0.033 (0.024)	-0.012 (0.022)
Observations	140	376	391	173	531	549	1080
Subjects	49	70	70	50	70	70	72
R^2	0.074	0.012	0.019	0.039	0.001	0.000	0.000

Notes: Cluster-robust standard errors (on the subject-level) are in parentheses. Constellation ls^H stands for low signals paired with the information HIGHER. Constellation L stands for all signals paired with the information LOWER. The other constellations are defined analogously.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Regression Table - Change in adjusted profits

Here we can see a significant effect for all constellations except for the last three, where low and high signals are pooled, and for the whole sample. For constellations ls^H and hs^L there is a significant and positive effect of the *information* treatment (INF) on the change in adjusted profits. For example, the coefficient of 0.306 in the first column can be interpreted in two ways. (i) In constellation ls^H , the subjects in treatment INF have an overall net rate of switching unfavorable auctions into favorable auctions in stage II, which is by 30.6 percentage points larger than the one

of the subjects in treatment NOINF. Or equivalently, (ii) in constellation ls^H , the subjects in treatment INF increased their adjusted profits per auction from stage I to stage II by 0.306 units, on average, compared to the subjects in treatment NOINF. The other coefficients can be interpreted analogously. For constellations ls^L and hs^H , there is a negative effect of the *information* treatment (INF) on the change in adjusted profits, making the subjects in treatment INF worse off compared to those in treatment NOINF. However, the positive effects of the *information* treatment (INF) are greater in absolute terms than the negative ones and the coefficient for constellation ls^L is only significant at the 10%-level. But since constellations ls^L and hs^H occur, by construction of the game, more often than ls^H and hs^L , the overall effect of treatment INF is not significantly different from the one of treatment NOINF. This indicates that, in total, the subjects who received information about their bids are not better off than those who received no information.

5.4 Differentiation between low and high signals

By construction of the game and the observed bidding behavior, it is profitable for the bidders to increase the initial bid for high signals and to decrease the initial bid for low signals, in most of the cases. 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 treatment INF strongly differentiate between low and high signals. The decrease rate for low signals is 78.7% and only 42.3% for high signals ($p = 0.000$, Fisher’s exact test). For the information LOWER, the subjects in treatment INF 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 ($p = 0.168$, Fisher’s exact test). 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.²⁴ 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 increases his initial bid. The problem

²⁴An alternative explanation for this behavior might be a general misunderstanding of second-price auctions. However, this point seems less likely, since in the instructions the mechanism of a second-price auction was explained in detail and the control questions could only be answered when this mechanism was understood correctly.

is that the new bid is not chosen conditional on winning - this can result in an even more frequent appearance of the winner’s curse for low signals.

Figure 7 reports the fractions of decreasing, increasing and retaining the initial bids in stage II for different constellations of signal and information for the subjects in treatment INF.




Figure 7: Changing of bids in stage II (treatment INF)

We have already seen that most of the bidders in treatment INF, in constellation ls^H , seem to realize that winning implies that the opponent also has a low signal and hence they select a bid below 20 in stage II. Combined with the pattern that those subjects seem to differentiate between low and high signals when receiving the information HIGHER, it seems to be very plausible to assume that the changing behavior in stage II is to a substantial 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 treatment INF answered that the information they received helped them indeed to get a better estimate of the opponent’s signal.

5.5 Cursed equilibrium

A further conclusion that can be derived from the previous results is that the overall behavior of the subjects in stage II, especially of those in treatment INF, cannot be explained by cursed equilibrium since a “cursed” bidder assumes by definition that the bids and signals of the opponent are not correlated. Hence, such a bidder should 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. Additionally, we have the best-response assumption, inherent in cursed equilibrium, stating that the bid of a “cursed” player is already evaluated conditional

on winning in stage I. Hence, there would be no need to change the bid in stage II for such a bidder. As already explained in Section 1, this reasoning holds for both fully and partly cursed bidders.




Figure 8: These Figures report the fractions of subjects who changed their bids in at least X rounds (left) and in at least X rounds by more than 10 units (right).

I observed that 97.44% of the subjects in treatment INF changed their bids at least once and that 66.67% of them changed their bids in more than 10 rounds. In treatment NOINF the changing rate is slightly smaller: 90.91% of the subjects changed their bids at least once and only 51.52% of them changed their bids in more than 10 rounds. If we only look at bids that were changed by at least 10 units, the difference between INF and NOINF is more distinct (see Figure 8).

So the best-response assumption is violated for the majority of subjects in both treatments. If the bid in stage I would have been an optimal bid, given the subject’s beliefs, there would be no need to change the bid in stage II, since the decision problem in a given auction is the same in stage II. For the subjects in treatment INF, we can additionally use the even stronger argument that those subjects actually recognized the informational content of winning or losing an auction. This pattern casts doubts whether the bidding behavior of the majority of subjects can be explained by cursed equilibrium, unless one assumes that the bidders are suddenly less cursed in stage II or 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 (2017) are 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, Levin, and Niederle (2010) who stated that bidders in common value auctions might act “as if” they have cursed beliefs.²⁵

²⁵Ivanov, Levin, and Niederle (2010) were the first to claim that bidders in common value

6 Conclusion

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.

First, the overall results show that the majority of subjects indeed reacted on the respective information they received and changed their bids in stage II. From a theoretical point of view this is surprising, since we have to assume that a bid, which was chosen in stage I, was considered as optimal in this context. So in the framework of a second-price auction, there would be no need to change such a bid in stage II, even when learning whether this bid is the winning bid or not. In other words, if a bid was chosen conditional on winning in stage I, there would be no need to change it in stage II.

Second, we have seen that at least in some constellations (ls^H and hs^L), the subjects profited from the information they received. In other constellations, however, this effect is reverse and, overall, there is no difference between the subjects in treatment INF and NOINF when we look at average profits. Therefore, additional information can be even negative for the bidders.²⁶ A crucial problem is that subjects still often used simple heuristics instead of making strategic choices after receiving feedback about their bid.

Finally, there is evidence in the data that at least for the information HIGHER, the changing of bids is not just a rule of thumb, but it rather occurs due to Bayesian updating since there is a much higher decreasing rate for low signals than for high signals. This claim is also supported by the pattern that, for low signals, after receiving the information HIGHER, most of the newly selected bids are below 20. This suggests that the bidders, who won with a low signal, indeed realized that the other bidder has most likely also a low signal.

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.

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.

²⁶In related studies like [Charness and Levin \(2009\)](#) and [Koch and Penczynski \(2017\)](#) 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 more 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.

References

- Adam, M. T., J. Krämer, and M. B. Müller (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., J. Krämer, and C. Weinhardt (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. *Quarterly Journal of Economics* 84(3), 488–500.
- Astor, P. J., M. T. Adam, C. Jähmig, and S. Seifert (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 J. H. Kagel (1997). Second-price auctions with asymmetric payoffs: An experimental investigation. *Journal of Economics & Management Strategy* 6(3), 573–603.
- Bazerman, M. H. and W. F. Samuelson (1983). I won the auction but don’t want the prize. *Journal of Conflict Resolution* 27(4), 618–634.
- Camerer, C., S. Nunnari, and T. R. Palfrey (2016). Quantal response and nonequilibrium beliefs explain overbidding in maximum-value auctions. *Games and Economic Behavior* 98, 243–263.
- Capen, E. C., R. V. Clapp, and W. M. Campbell (1971). Competitive bidding in high-risk situations. *Journal of Petroleum Technology* 23(6), 641–653.
- Charness, G. and D. Levin (2009). The origin of the winner’s curse: A laboratory study. *American Economic Journal: Microeconomics* 1(1), 207–236.
- Costa-Gomes, M. A. and M. Shimoji (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 N. Iriberry (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 E. Vespa (2014). Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics* 6(4), 180–202.

- Esponda, I. and E. Vespa (2016). Contingent preferences and the sure-thing principle: Revisiting classic anomalies in the laboratory. *working paper*.
- Evans, J. S. B. (2007). *Hypothetical thinking: Dual processes in reasoning and judgement*, Volume 3. Psychology Press.
- Eyster, E. and M. Rabin (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. (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 R. Sherman (1994). The loser’s curse. *American Economic Review* 84(3), 642–652.
- Ivanov, A., D. Levin, and M. Niederle (2010). Can relaxation of beliefs rationalize the winner’s curse?: An experimental study. *Econometrica* 78(4), 1435–1452.
- Kagel, J. H., R. M. Harstad, and D. Levin (1987). Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica* 55(6), 1275–1304.
- Kagel, J. H. and D. Levin (1986). The winner’s curse and public information in common value auctions. *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 S. P. Penczynski (2017). The winner’s curse: Conditional reasoning & belief formation. *JET forthcoming*.
- Levin, D., J. Peck, and A. Ivanov (2016). Separating bayesian updating from non-probabilistic reasoning: An experimental investigation. *American Economic Journal: Microeconomics* 8(2), 39–60.
- Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review* 107(11), 3257–87.
- Louis, P. (2013). The barrel of apples game: Contingent thinking, learning from observed actions, and strategic heterogeneity. *working paper*.

- Ngangoue, K. and G. Weizsäcker (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.
- Singmann, H., K. C. Klauer, and S. Beller (2016). Probabilistic conditional reasoning: Disentangling form and content with the dual-source model. *Cognitive Psychology* 88, 61–87.
- Thaler, R. H. (1988). Anomalies: The winner’s curse. *The Journal of Economic Perspectives* 2(1), 191–202.
- Van den Bos, W., J. Li, T. Lau, E. Maskin, J. D. Cohen, P. R. Montague, and S. M. McClure (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 Additional Figures

Decision screens

Teil 1 - Runde 1 von 15

Der Inhalt Ihres Umschlags beträgt 52 Taler.

Wie viele Taler bieten Sie dafür, sowohl Ihren Umschlag als auch den Umschlag des anderen Spielers zu erhalten?

Gebot:

Bieten

Figure 9: 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?*

Teil 2 - Runde 1 von 15

Ihre Rolle: A

Der Inhalt Ihres Umschlags beträgt 50 Taler.
Ihr Gebot in Teil 1 war 99 Taler.
Ihr Gebot war **HÖHER** als das des anderen Spielers.
Der andere Spieler wird das gleiche Gebot wie in Teil 1 erneut abgeben.
Wie viele Taler bieten Sie dafür, sowohl Ihren Umschlag als auch den Umschlag des anderen Spielers zu erhalten?

Gebot:

Bieten

Figure 10: 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?*

Teil 2 - Runde 1 von 15

Ihre Rolle: B

Der Inhalt Ihres Umschlags beträgt 8 Taler.

Ihr Gebot in Teil 1 war 99 Taler.

Ihr Gebot war ??? als das des anderen Spielers.

Der andere Spieler wird das gleiche Gebot wie in Teil 1 erneut abgeben.

Wie viele Taler bieten Sie dafür, sowohl Ihren Umschlag als auch den Umschlag des anderen Spielers zu erhalten?

Gebot:

Bieten

Figure 11: 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?*

B Proofs

Proof of Proposition 2. 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 $\alpha > 1$. For $\alpha = 2$ we have the unique symmetric equilibrium. Note that this argumentation does not rely on the distribution of the signals x_1 and x_2 . \square

Proof of Proposition 3. Any bid $b_i(x_i)$ outside the interval $[x_i, x_i + 60]$ 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 $x_i + 60$ is weakly dominated by bidding $x_i + 60$. If player 1 would have won the auction with $b_1(x_1) = x_1 + 60$, he receives either a positive, negative or zero payoff. Deviating to a bid above $x_1 + 60$ ($b'_1(x_1) > b_1(x_1)$) would not change the result because $b'_1(x_1) > b_1(x_1) \geq b_2(x_2)$ and the price to pay remains $b_2(x_2)$. If player 1 would have lost the auction with $b_1(x_1) = x_1 + 60$, he receives a payoff of 0. Deviating to a bid above $x_1 + 60$ can lead to winning the auction, but the payoff is at best 0, because the value of the object is at most equal to $x_1 + 60$ and the price to pay remains $b_2(x_2)$ with $b_2(x_2) \geq b_1(x_1) = x_1 + 60$. So if player 1 deviates to $x_1 + 60 + \delta$ and player 2 bids $x_1 + 60 + \epsilon$ with $\delta > \epsilon \geq 0$, player 1 faces a loss of at least $\epsilon \geq 0$. \square

Proof of Proposition 4. Suppose two players bid according to $b_1(x_1) = \alpha \cdot x_1$ and $b_2(x_2) = \frac{\alpha}{\alpha-1} \cdot x_2$. For any $\alpha \neq 2$ one player has a slope strictly greater than 2. For $\alpha > 2$ the proof is trivial. For $\alpha < 2$ we have the following equivalence: $\frac{\alpha}{\alpha-1} > 2 \Leftrightarrow \alpha > 2 \cdot (\alpha - 1) \Leftrightarrow \alpha > 2 \cdot \alpha - 2 \Leftrightarrow 2 > \alpha$.

Now suppose, without loss of generality, that player 1 uses the bidding function $b_1(x_1) = \alpha \cdot x_1$ with $\alpha = 2 + \epsilon$ for $\epsilon > 0$. For $x_1 = 60$ the bid is then given by $b_1(x_1) = \alpha \cdot 60 = 2 \cdot 60 + \epsilon \cdot 60 = 120 + \epsilon \cdot 60$. It follows that $120 + \epsilon \cdot 60 > x_1 + 60 = 120$ for all $\epsilon > 0$. By definition of Proposition 3, this bid is weakly dominated. So we have found at least one signal for which the resulting bid is weakly dominated for all $\epsilon > 0$. Note that such a bid would not even be feasible, because the bidding range is restricted to the interval $[0, 120]$. \square

Proof of Proposition 5. If a player bids according to $b_i(x_i) = x_i + 30$ in stage I and the other player does not use any weakly dominated bids, he can predict in stage II whether the signal of the other player is low or high when (i) he wins with a low signal or (ii) he loses with a high signal.

Consider two players, indexed by $i = 1, 2$. Suppose, without loss of generality, player 1 follows the bidding rule $b_1(x_1) = x_1 + 30$ and player 2 does not use any weakly dominated bids. This means for each x_2 , player 2 may choose any bid in a range of $[x_2, x_2 + 60]$.

(i) Player 1 has a low signal ($x_1 \in \{0, \dots, 10\}$) and wins: Player 1 wins with a bid of $b_1(x_1) = x_1 + 30$. This means his bid can be in a range of $[30, 40]$. If player 2 has a low signal, his bid can be in a range of $[0, 70]$. If player 2 has a high signal, his bid can be in a range of $[50, 120]$. So player 1 wins only if player 2 has a low signal.

(ii) Player 1 has a high signal ($x_1 \in \{50, \dots, 60\}$) and loses: Player 1 loses with a bid of $b_1(x_1) = x_1 + 30$. This means his bid can be in a range of $[80, 90]$. If player 2 has a low signal, his bid can be in a range of $[0, 70]$. If player 2 has a high signal, his bid can be in a range of $[50, 120]$. So player 1 loses only if player 2 has a high signal. \square

C 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. First, **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

$$10 \text{ Taler} = 1 \text{ EURO}$$

At the beginning 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:

$$\text{Final Payout} = 4 \text{ EUR} + \text{experimental credit from part 1 and part 2 (in EUR)}$$

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 learns 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 only. 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 only 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 and the contents of the envelopes are equal, both players receive a payoff of 0. If your bid is lower than the other player's 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 you have to pay 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 would 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 learn 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 learn 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 neither know who your opponent is nor does your opponent know this. 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 screens, 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 learn your final payoff only at the very end of the experiment (**i.e., only when part 2 is finished**). 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

You will now bid for the envelopes once again. For this purpose you and your previous opponent will receive the **same amounts** as in part 1 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, you will also see "**LOWER**" on your screen (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** you only see "???" on your screen.

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 learn 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.

D Control questions

The following control questions were asked before the subjects started with the actual bidding task. The subjects were only allowed to continue with the experiment after all control questions were answered correctly.

1. You have 50 Taler in your envelope and the other player has 10 Taler in his envelope. You bid 110 Taler and the other player bids 30 Taler. What would be your payoff? (*Correct answer: 30 Taler*)

2. You have 50 Taler in your envelope and the other player has 10 Taler in his envelope. You bid 50 Taler and the other player bids 30 Taler. What would be your payoff? (*Correct answer: 30 Taler*)
3. You have 5 Taler in your envelope and the other player has 50 Taler in his envelope. You bid 25 Taler and the other player bids 90 Taler. What would be your payoff? (*Correct answer: 0 Taler*)
4. You have 5 Taler in your envelope and the other player has 50 Taler in his envelope. You bid 100 Taler and the other player bids 90 Taler. What would be your payoff? (*Correct answer: -35 Taler*)
5. You have 10 Taler in your envelope and the other player has 5 Taler in his envelope. You bid 45 Taler and the other player bids 35 Taler. What would be your payoff? (*Correct answer: -20 Taler*)
6. You have 10 Taler in your envelope and the other player has 5 Taler in his envelope. You bid 60 Taler and the other player bids 35 Taler. What would be your payoff? (*Correct answer: -20 Taler*)
7. You have 55 Taler in your envelope and the other player has 50 Taler in his envelope. You bid 90 Taler and the other player bids 90 Taler. What would be your payoff? (*Correct answer: 15 Taler*)
8. You have 55 Taler in your envelope and the other player has 50 Taler in his envelope. You bid 80 Taler and the other player bids 90 Taler. What would be your payoff? (*Correct answer: 0 Taler*)