

Freeman, David; Kimbrough, Erik O.; Reiss, J. Philipp

Working Paper

Opportunity cost, inattention and the bidder's curse

KIT Working Paper Series in Economics, No. 101

Provided in Cooperation with:

Karlsruhe Institute of Technology (KIT), Institute of Economics (ECON)

Suggested Citation: Freeman, David; Kimbrough, Erik O.; Reiss, J. Philipp (2017) : Opportunity cost, inattention and the bidder's curse, KIT Working Paper Series in Economics, No. 101, Karlsruher Institut für Technologie (KIT), Institut für Volkswirtschaftslehre (ECON), Karlsruhe, <https://doi.org/10.5445/IR/1000068511>

This Version is available at:

<https://hdl.handle.net/10419/155326>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Opportunity cost, inattention and the bidder's curse

by David Freeman, Erik O. Kimbrough,
J. Philipp Reiss

No. 101 | FEBRUARY 2017

WORKING PAPER SERIES IN ECONOMICS



Impressum

Karlsruher Institut für Technologie (KIT)
Fakultät für Wirtschaftswissenschaften
Institut für Volkswirtschaftslehre (ECON)

Kaiserstraße 12
76131 Karlsruhe

KIT – Die Forschungsuniversität in der Helmholtz-Gemeinschaft

Working Paper Series in Economics
No. 101, February 2017

ISSN 2190-9806

econpapers.wiwi.kit.edu

Opportunity cost, inattention and the bidder's curse*

David Freeman[†] Erik O. Kimbrough[‡] J. Philipp Reiss[§]

February 10, 2017

Abstract

Recent research suggests that auction winners sometimes fall prey to a “bidder's curse”, paying more for an item at auction than they would have paid at a posted price. One explanation for this phenomenon is that bidders are inattentive to posted prices. We develop a model in which bidders' inattention, and subsequent overbidding, is driven by a rational response to the opportunity cost of acquiring information about the posted price. We test our model in a laboratory experiment in which subjects bid in an auction while facing an opportunity cost of looking up the posted price. We vary the opportunity cost, and we show that information acquisition decreases and consequently overbidding increases with opportunity cost as predicted.

JEL classification: C72, C92, D44

Keywords: Auctions, Bidder's Curse, Limited Attention, Experiments, Rational Ignorance

*The authors would like to thank Maastricht University's METEOR research organization and SSHRC (Kimbrough) for funding. We are also grateful to Henry Schneider, Bart Wilson, participants in a brown bag workshop at Simon Fraser University, and conference participants at the 2012 World Meetings of the Economic Science Association for useful comments.

[†]Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, BC Canada, V5A 1S6; e-mail: david.freeman@sfu.ca

[‡]Corresponding Author: Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, BC Canada, V5A 1S6; e-mail: ekimbrough@gmail.com

[§]Karlsruhe Institution of Technology; e-mail: philipp.reiss@kit.edu

There is no ‘imperfection’ in a market possessing incomplete knowledge if it would not be remunerative to acquire (produce) complete knowledge: information costs are the costs of transportation from ignorance to omniscience, and seldom can a trader afford to take the entire trip. \sim (Stigler, 1967)

1 Introduction

Overbidding in auctions has been observed in a wide range of scenarios ranging from private-value second-price auctions to common-value auctions of many formats in both the lab and the field (Kagel et al., 1987; Kagel and Levin, 1993; Cooper and Fang, 2008; Kagel and Levin, 1986; Klemperer, 2002; Levin and Reiss, 2017). These pathological cases frequently result in a “winner’s curse” in which the winner of the auction pays more for the item than it is actually worth. A different form of overbidding has been documented among eBay bidders who sometimes overpay in auctions relative to a simultaneously available fixed-price (Malmendier and Lee, 2011; Jones, 2011). Bidders subject to the “bidder’s curse” fail to maximize their gains from exchange; however, paying more than a simultaneous fixed-price need not imply that the final price exceeds the buyer’s value.

In the case of overpayment relative to a simultaneous fixed price on eBay, the proffered explanation is that bidders “fail to pay sufficient attention to their outside options, especially when rebidding” (Malmendier and Lee, 2011, p. 776). However, it is still unclear *why* individuals fail to pay ‘sufficient’ attention. Malmendier and Lee (2011) suggest that the source of limited attention is non-rational, but Schneider (2016) suggests that some overbidding might be explained as a rational response to search costs – rational ignorance. In the context of eBay, the time and effort cost of researching the posted price, which we will refer to as opportunity cost, seems to be a plausible explanation. Search results are quite sensitive to the particular search terms used and frequently return a large number of items that vary in relevance, condition, and along other dimensions (Schneider, 2016). In field data, the opportunity costs of research are uncontrolled and unobservable, and thus what looks like insufficient attention may instead be (at least in part) rational ignorance. To address this limitation, we design a laboratory experiment to directly test whether the opportunity cost of looking up a fixed price can account for the bidder’s curse in auctions where opportunity costs are controlled and partially observed.

In our setting, subjects participate in a sequence of induced private value, ascending clock auctions for goods that are subsequently available for purchase at a fixed price drawn

from a known distribution.¹ Each subject demands a single unit and values the unit sufficiently that she will strictly prefer to purchase at the fixed price if she is unsuccessful in the auction. To learn the actual fixed price before bidding, subjects must forgo the opportunity to earn money in an individual effort task. Thus in deciding whether to research the fixed price, subjects face a tradeoff between this opportunity cost and the expected savings from making an informed bid.²

We provide a benchmark theoretical model in which rational ignorance is an equilibrium outcome when the opportunity cost of becoming informed is sufficiently high. In the model, rational ignorance generates behavior consistent with the bidder’s curse because uninformed agents sometimes bid above the fixed price. By varying the piece rate of our experimental effort task, our design allows us to observe the information acquisition strategies of subjects with varying opportunity costs and their bidding behavior conditional on whether they have observed the fixed price.

We find evidence that buyers who face higher opportunity costs of learning the fixed price are less likely to learn, and hence more likely to overbid, the fixed price. However, when the opportunity cost of researching the fixed price is high, auction winners who choose *not* to research the fixed price receive significantly higher payoffs than those who *do* since their earnings from the effort task more than offset their occasional losses from overbidding. Our analysis shows that opportunity cost-driven overbidding is not harmful on net.

Existing field studies offer suggestive evidence to support theoretical claims that opportunity cost plays an important role in bidding behavior (e.g. Backus et al., 2014; Schneider, 2016), albeit without observing or controlling opportunity cost. In laboratory auctions in which the item is alternatively available at a fixed price any time until the auction starts, Durham et al. (2013) observe bids in excess of the posted price in only 4 out of 359 observations, suggesting that limited attention is a weak constraint in a zero-opportunity cost environment.³ Other evidence from internet auctions suggests a role for more careful attention to quasi-rational explanations as opposed to behavioral biases (e.g. Podwol and Schneider, 2016).

¹For a discussion of these “English clock” auctions, see Cassady (1980); McCabe et al. (1990).

²Because the fixed price is common to all bidders in an auction and because bidders have single unit demand, our auctions are equivalent to second-price common-value auctions, with the common value being the fixed price, so long as bidders are fully rational. Informed bidders observe this common value while uninformed bidders do not.

³Data on overbidding collected via personal communication with the authors.

A separate bidding puzzle in ascending auctions is the observation that bidders sometimes increase their bids by more than the minimum increment, a phenomenon known as “jump bidding” (see e.g. Cramton, 1997; Christopher, 1998). Our proposed explanation for the bidder’s curse parallels attempts to explain jump bidding via time and effort opportunity costs (Easley and Tenorio, 2004; Isaac et al., 2007). For example, Easley and Tenorio (2004), studying data from online auctions, provide a model in which the opportunity cost of repeatedly logging into an auction website to increase one’s bid leads bidders to make jump bids in equilibrium. Similarly, rational ignorance could account for “shipping-fee neglect” in online purchases (see e.g. Hossain and Morgan, 2006; Brown et al., 2010).

2 Theory

In our setting there are $n \geq 2$ players, who each demand a single unit of a good. The good is sold in both a second-price ascending clock auction and an aftermarket. One unit of the good is available in the auction, but an unlimited quantity of the good is available in an aftermarket at a fixed price. The aftermarket price, q , is not initially observed by players, but is commonly known to be drawn from cumulative distribution function G which admits a probability density function and has support on $[\underline{q}, \bar{q}]$ where $\underline{q} \geq 0$. A given player i values one unit of the good at v_i , which she privately observes, but v_i is commonly known to satisfy $v_i \geq \bar{q}$.

Each player first bids in the auction, and if she does not win the auction, she can instead buy the good in the aftermarket. However, each player may choose to privately observe the aftermarket price before the auction starts, but doing so for agent i is costly. In particular, looking up the price requires forgoing a task with individual reward, c_i . This reward is privately observed by i and is known to be drawn independently according to the cumulative distribution function F . A player’s lookup decision is also her private information.

Our environment closely approximates the situation faced by an eBay bidder who has recently been outbid and must choose whether to continue to increase her bid.⁴ When eBay bidders are outbid, they receive a notification by email that they are no longer the high bidder. They then must choose whether to increase their bid and also whether

⁴Ariely et al. (2005) and Roth and Ockenfels (2002) argue that second-price auctions can be considered stylized versions of those run at eBay, where initial data on overbidding fixed price options was collected.

to devote some of their valuable time to researching other purchase options (like eBay's Buy-It-Now price, or fixed price options from other sellers, which are equivalent to our aftermarket price). Here we ask whether it is ever rational *ex ante* to overbid these fixed price options.

2.1 Preview of the Analysis

Formally, a strategy for player i prescribes her information acquisition as a function of her realized v_i and c_i ; the price at which she drops out of the ascending auction (her bid) conditional on her information acquisition, v_i , and c_i ; and her decision of whether or not to acquire the good in the aftermarket conditional on losing the auction and on v_i , c_i , her information acquisition, and her bid.

Below, we assume that all players are risk neutral and construct an essentially unique symmetric perfect Bayesian equilibrium (PBE) of this game. This PBE is described by an information acquisition function ι and an uninformed bid b^* such that each player follows the strategy given by (A)-(D) below:

- (A) Each player acquires information with probability $\iota(c_i)$ when her type is (v_i, c_i) .
- (B) Each player, conditional on becoming informed that the aftermarket price is q , bids q in the auction.
- (C) Each player, conditional on remaining uninformed, bids b^* in the auction.
- (D) All types always buy in the aftermarket conditional on losing the auction.

Characterizing a PBE involves finding the restrictions on ι and b^* that make the strategies described by (A)-(D) consistent with PBE. We construct the PBE by first finding the optimal bid of a player, given her information about the aftermarket price and conditional on the information acquisition strategies of others. We then characterize the value of information and derive a subject's information acquisition strategy. In the PBE, agents acquire information according to a cutoff strategy in opportunity cost, becoming informed only when the expected value of being informed exceeds the opportunity cost of acquiring information. Then given their information, agents bid either the expected aftermarket price (if uninformed) or the actual aftermarket price (if informed). When some bidders are uninformed, this equilibrium may generate both over- and under-bidding in the auction stage, relative to the aftermarket price. Thus, with positive opportunity

costs of information acquisition, rationally ignorant bidders exhibit behavior that looks like the “bidder’s curse”.

2.2 Bidding strategies conditional on information

A player’s optimal bid in the auction depends on whether she is informed about the aftermarket price, q . Note that since $v_i \geq \bar{q} \forall i$, players will always prefer to purchase the good at a price q in the aftermarket if they do not buy during the auction. Thus, for an *informed* player, a bid of q in the auction weakly dominates all other bids. Bidding q both ensures that the player will never overpay relative to the aftermarket price and maximizes the probability of winning the item at a price below q , should other players underbid. Since the price paid by the auction winner is unaffected by her own bid, it follows that this strategy minimizes her expected price paid. Thus, we look for a PBE in which informed bidders always bid q .

Knowing this, when we define the optimal bid for *uninformed* bidders, it must be a best response when all *informed* bidders bid q . Suppose for now that all other players become informed with probability ρ , and further suppose they bid q when informed and b^* when uninformed. The *ex post* price paid by player i from bidding b_i then depends on whether she wins the auction and whether the second-highest bid was made by an informed or uninformed bidder:

Case I: $\max b_{-i}$ is made by an informed player.

Here i always pays q *ex post*, whether in the auction or the aftermarket. If $b_i < \max b_{-i}$ she loses the auction and pays q in the aftermarket. If $b_i > \max b_{-i} = q$, she wins the auction and pays q .

Case II: $\max b_{-i}$ is made by an uninformed player.

Here i will either pay q or b^* *ex post*, depending on whether she wins the auction. If she submits any bid $b'_i < b^*$ she loses the auction and pays q in the aftermarket. If she submits any bid $b''_i > b^*$, she wins the auction and pays b^* . If instead, she bids $b_i = b^*$, the tie-breaking rule ensures that she sometimes wins and sometimes loses the auction, paying b^* if she wins and q in the aftermarket if she loses.

Using the *ex post* prices in Cases I and II, the expected price paid by any uninformed player is equal to the probability of Case I times the expected value of q plus the probability of Case II times her expected price in Case II, which is a convex combination of b^* and

the expected value of q . Using this, we can evaluate the expected prices that would result from bids greater than and less than b^* . When player i bids b^* , her distribution over *ex post* prices is a convex combination of the price resulting from bidding $b'_i < b^*$ and $b''_i > b^*$ and so is her *ex ante* expected price. For bidding b^* to be a best-reply for i , she must be indifferent between always winning and always losing to an uninformed bid; otherwise, she would strictly prefer to bid above or below b^* . This provides the restriction for b^* to be the PBE bid of all uninformed players.

Uninformed bidders pay an expected price of $\mathbb{E}q$ since bidding below b^* always leads to an *ex post* price of q either in the auction or aftermarket. Deriving the corresponding value for bidding above b^* and equating the expected prices faced by i when bidding any $b'_i < b^*$ and when bidding any $b''_i > b^*$ gives the restriction that:

$$\mathbb{E}q = (1 - \rho)^{n-1}b^* + \rho^{n-1}\mathbb{E}q + (1 - \rho^{n-1} - (1 - \rho)^{n-1}) \left[\int_{\underline{q}}^{\bar{q}} \max[b^*, q] dG(q) \right].$$

Rearranging the above gives:

$$(1 - \rho)^{n-1} [\mathbb{E}q - b^*] = (1 - \rho)^{n-1} - (1 - \rho)^{n-1}) \left[\int_{\underline{q}}^{\bar{q}} \max[0, b^* - q] dG(q) \right]. \quad (1)$$

When $n = 2$, the right-hand side of (1) equals zero, and thus $b^* = \mathbb{E}q$. But when $n > 2$, (1) implicitly defines the uninformed bid b^* given ρ , the probability of becoming informed.⁵ Since the implied b^* exceeds \underline{q} , it follows that uninformed players will sometimes overbid (and sometimes underbid) the aftermarket price. The bidding strategy of an uninformed player thus generates behavior consistent with the “bidder’s curse”. It remains to be shown that some bidders will choose to be uninformed in equilibrium.

2.3 Information acquisition

Given what we have shown so far about bidding strategies, we can now consider the costs and benefits of becoming informed before bidding in the auction. Information acquisition is costly because player i must forgo c_i to become informed. Acquiring information is beneficial because it eliminates the possibility of overpaying relative to the aftermarket

⁵See Lemma 1 in Appendix A for the full argument.

price and raises the likelihood of getting a bargain in the auction against uninformed bidders whenever $b^* < q$. We characterize the PBE information acquisition strategy in which a player becomes informed whenever these expected benefits exceed the costs, given the expected information acquisition strategies of others.

Let ι represent the information acquisition strategy of $-i$. If i has rational expectations about others' behavior, then i 's *ex ante* assessment of the probability of each other player becoming informed, ρ , is given by:

$$\rho = \int \iota(c) dF(c) \tag{2}$$

Using the bidding strategies of $-i$ derived above, an informed bidder pays a price q either in the auction or the aftermarket whenever at least one other bidder is informed, but pays the minimum of q and b^* when no other player is informed. Thus the expected price paid, conditional on choosing to lookup the aftermarket price, is given by

$$(1 - \rho)^{n-1} \int \min[q, b^*] dG(q) + p^{n-1} \mathbb{E}q + (1 - \rho^{n-1} - (1 - \rho)^{n-1}) \mathbb{E}q.$$

As noted earlier, in equilibrium, uninformed bidders pay an expected price of $\mathbb{E}q$. The expected savings from acquiring information is then given by

$$(1 - \rho)^{n-1} \int \max[0, q - b^*] dG(q) \tag{3}$$

Notice that, conditional on b^* , the value of becoming informed is declining in ρ ; that is, players' information acquisition decisions are strategic substitutes. This is because when other players are likely to be informed, it is more likely that all other bids in the auction will equal the aftermarket price, which eliminates the benefit of being informed. In contrast, when no other bidders are informed, then a given bidder who looks up the price can save $q - b^*$ whenever $q > b^*$, though the probability of getting these savings decreases as ρ increases.

In equilibrium, each player acquires information when the value of information, (3), is greater than c_i and remains uninformed whenever it is less than c_i . Thus i follows ι if it is consistent with the cutoff strategy:

$$\bar{c} = (1 - \rho)^{n-1} \int \max[0, q - b^*] dG(q) \quad (4)$$

$$\iota(c) = 1 \text{ if } c > \bar{c} \quad (5)$$

$$\iota(c) = 0 \text{ if } c < \bar{c} \quad (6)$$

These equations together with (2) fully characterize an ι consistent with symmetric PBE, given b^* .⁶ Thus agents weigh the benefits of avoiding overpaying (and potentially underpaying), relative to q , against their opportunity costs, and some will find it rational to remain ignorant.

If $\rho = 1$, then an uninformed bidder pays an expected price of $\mathbb{E}q$ (regardless of her bid in the auction), as does an informed bidder. But then if $c_i > 0$, a player is strictly better off by remaining uninformed. Thus so long as $F(0) < 1$, some bidders will choose to remain uninformed in equilibrium. Since opportunity costs are rarely zero, our theory predicts that rationally ignorant agents will be present in both lab and field markets and will influence auction prices.⁷

2.4 Equilibrium

The equilibrium described is summarized in Proposition 1 below.

Proposition 1. Existence of Symmetric PBE

The game described above has a symmetric PBE which can be described by an information acquisition function ι and an uninformed bid b^ that satisfy (1)-(6) such that each player follows the strategy given by (A)-(D).*

In this equilibrium, if $F(0) < 1$ then $\rho < 1$, and if $F(\int \max[0, q - \mathbb{E}q] dG(q)) > 0$ then $\rho > 0$.

⁶Note that if F has mass points, (4) - (6) also pin down the probability of acquiring information at \bar{c} , $\iota(\bar{c}) \in [0, 1]$, when necessary. The right-hand side of (4) is weakly positive, strictly decreasing in ρ , and takes all values between 0 and $\int \max[0, q - b^*] dG(q) \geq 0$. To find the cutoff value \bar{c} , substitute in $\rho = F(\bar{c})$ into the right-hand side, and starting at $\bar{c} = 0$, increase \bar{c} on the right-hand side of the equation until the first point at which the right-hand side is weakly less than the left-hand side. This point is the desired \bar{c} . If the right-hand side is strictly below \bar{c} at this point, we can find a $\rho = \in [\lim_{c \uparrow \bar{c}} F(c), F(\bar{c})]$ to make the desired equality hold, pinning down $\iota(\bar{c})$.

⁷If $\rho = 0$, then by (1), uninformed agents would bid $\mathbb{E}q$. Then by (3) would find it worthwhile to become informed if $\int \max[0, q - \mathbb{E}q] dG(q) > c_i$. Thus so long as $F(\int \max[0, q - \mathbb{E}q] dG(q)) > 0$, some bidders will choose to become informed in equilibrium.

Uniqueness of the Symmetric PBE In Appendix A, we show that in every symmetric PBE, types who become informed bid the aftermarket price, q , with a probability of 1; types who remain uninformed pay an expected price of $\mathbb{E}q$. Thus, information acquisition decisions are given by a cutoff strategy in c_i consistent with (4)-(6).

Below we describe an experiment that allows us to test the behavioral implications of symmetric PBE, in particular on overbidding (relative to q) and rational ignorance.

3 Experimental Design, Procedures, and Predictions

The experiment consists of two parts, the “Productivity Measurement” part and the “Market” part. In the “Productivity Measurement” part, each subject participates in three, two-minute periods of a real-effort ‘slider task’ due to Gill and Prowse (2012) with a break of one minute in between periods. Subjects observe a screen with 48 sliders, each representing a scale from 0 to 100 (See Figure B5 in Appendix B below). The sliders are initially set to “0”, and subjects receive a payment $k^* = 1$ Experimental Currency Units (ECU) for each slider that they set to “50” by the end of the two minutes. Observing behavior in this part allows us to control for individual skill heterogeneity when attempting to understand the role of opportunity cost in influencing bidding behavior.

After the Productivity Measurement part, subjects participate in the Market part, consisting of $T = 16$ unrelated markets in which they attempt to buy a single unit of a fictitious item. In each market, one unit of supply is available in an ascending auction, and – after the auction ends – unlimited supply is available at a fixed price for those who did not obtain the unit in the auction and still wish to purchase it. At the beginning of each market, subjects observe their induced value for the item, and the distribution from which the fixed price is drawn. Importantly, if they want to learn the actual fixed price before bidding in the auction, they must choose to forgo the opportunity to earn additional income in another 30 seconds of the slider task, where they receive k ECU per correctly placed slider. Hence, by varying the value of k in the slider task, we vary the opportunity cost c_i of looking up the fixed price, and we can observe the effect of opportunity cost on the probability of overbidding in the auction.

Next, those subjects who choose not to lookup the fixed price participate in 30 seconds of the slider task in which each correctly placed slider yields a return of k , while those who choose to forgo the slider task can observe the fixed price displayed on their screens for 30 seconds. At the end of the 30 seconds, subjects submit their bids in the clock

auction with the highest bidder winning the auction and paying the second-highest bid. To ensure that we observe the high bid, bidders did not observe the dropping behavior of their competitors and the auction clock stopped when the last bidder dropped from the auction. Then, each subject learns whether they were able to purchase the good at auction and the auction price, and those subjects who were unsuccessful are then given the opportunity to purchase the item at the fixed price. Just before the decision to purchase is made, the fixed price is revealed to all remaining buyers at no cost, and subjects simply choose whether or not to buy at the revealed price. This concludes the first market. At the end of the first market, the second market begins, and so on until all 16 markets have closed.

Each market consists of $n = 2$ bidders chosen from a matching group of size $N = 8$. For each market, each subject receives a randomly drawn integer value denominated in Experimental Currency Units, $v \sim U[50, 100]$, and we draw the fixed price, $q \sim U[0, 50]$, to ensure that the possibility of overbidding (almost) always exists. Over the sequence of 16 markets, the value of k in the slider task varies randomly. Specifically, k is drawn independently for each bidder from the set $\{0, 0.25, 3, 6\}$ with equal probability of each value.⁸ We randomly rematch within groups so that each bidder faces each value of k four times, on average. Hence, we are able to capture the benefits of both a between and within subjects comparison of the treatment effect.

Subjects are rematched across markets and receive no additional information about the other bidders in their market. The bidder’s role is framed as that of a seller facing the opportunity of buying several commodities in the markets for resale to the experimenters, one commodity per market. One advantage of this framing is that it better motivates participation in multiple markets. In the auction stage of the market, the clock starts at 0, and increases in unit increments until a price of 105 ECU. At any time, each bidder can press a button labeled “Stop Bidding” to drop out of the auction; their bid is given by the price at which they stop bidding. Once both bidders have dropped out, each bidder is informed whether they won the auction and the price at which the earliest bidder dropped out, which gives the price paid by the winning bidder (as in a second-price auction).

We employ an ascending auction because this has been shown to facilitate comprehension and reduce overbidding relative to sealed-bid alternatives (Kagel et al., 1987).

⁸We also ran two unreported pilot sessions with $k \in \{0, 0.25, 1, 5\}$. We obtained qualitatively similar results but less overbidding of the posted price with these parameter values, which is consistent with our predictions.

Eliciting drop-out prices from all bidders represents a deviation from the typical clock auction, but it allows us to observe the entire distribution of bids, so that we can test the full set of predictions from our theory.⁹ See Appendix B below for screenshots and a complete set of instructions.

At the end of the experiment, we randomly select one market for each subject for payment. ECUs are converted to Canadian dollars at a rate of 10 ECU = \$1. Then subjects receive private cash payments including a \$7 payment for arriving to the experiment on time, their earnings from the Productivity Measurement stage, and their earnings from the randomly selected market (including what they earned in the slider task, if they participated). Prior to beginning the incentivized markets, subjects participated in two practice markets that were unpaid to familiarize themselves with the environment.

In total we ran 8 sessions of 8 subjects drawn from the undergraduate population of a mid-sized North American university (Average age = 22, 38% Male). Subjects from all disciplines are invited to sign up for our online recruitment system. We emailed a subset of eligible students invitations at random. The only restriction on participation was that the subjects could not have any prior experience in this experiment. On average, subjects earned \$19 for a 90-minute session, including show-up payment.

3.1 Predictions

Here, we detail the behavioral implications of the symmetric PBE for our experiment where $n = 2$. The model introduced in section 2 shows how the opportunity cost of researching the aftermarket price leads to behavior that looks like the bidder's curse. In our model, all informed bidders bid q and therefore never overbid the aftermarket price. Uninformed bidders, however, always bid $\mathbb{E}q$, and when $\mathbb{E}q > q$, they will overbid the aftermarket price (which occurs half of the time in equilibrium.) The presence of uninformed bidders is entirely driven by the opportunity cost of looking up the aftermarket price. When the cost of acquiring this information is sufficiently high, bidders find it in their interest to remain uninformed, and as a result, they may overbid the aftermarket price. We state four testable predictions of our model below:

1. Overbidding Overbidding relative to the aftermarket price is increasing in the cost of acquiring information.

⁹See also Engelmann and Wolfstetter (2009) who provide evidence that the standard clock auction conceals overbidding due to selection bias. As we show below, this is consistent with our observations.

2. Information Acquisition Each player follows a cutoff information acquisition strategy: there exists a cutoff \bar{c}_i such that the player looks up the price whenever $c_i > \bar{c}_i$ and does not look it up whenever $c_i < \bar{c}_i$.

3. Rationality of Inattention Uninformed bidders achieve higher payoffs, on average, than they would have by acquiring information, even when their bid exceeds the aftermarket price.

4. Bidding Strategies Informed bidders bid q , while uninformed bidders bid $\mathbb{E}q$ and are, hence, equally likely to over- and under-bid the aftermarket price.

4 Experimental Results

4.1 Opportunity Cost and Overbidding

Our data consist of 1024 price lookup and bidding decisions from 512 markets. Overall, we observe 510 ($\sim 50\%$) instances in which bidders lookup the aftermarket price and 463 (45%) instances of overbidding the aftermarket price. On average, when informed, subjects bid 27 (median 25), and when uninformed they bid 32 (median 30). Among the overbids, 294 (63%) are made by uninformed bidders, and of the 344 sizable overbids ($b_{i,t} - q_t > 5$), the number made by uninformed bidders is 244 (71%). Figure 1 displays the relationship between bids and fixed prices conditional on the decision to lookup the aftermarket price for each value of k . The figure also displays equilibrium predictions and OLS fits to the data.

As expected, bids often correspond to the fixed price when a bidder is informed, and we observe substantial overbidding, particularly when bidders are uninformed. Under the theory, the probability of remaining uninformed (and hence overbidding) is increasing in the expected value of the individual effort task. Summary statistics reported in Table 1 indicate both that the probability of looking up the fixed price is decreasing in k and that the probability of overbidding is higher when bidders choose not to lookup the fixed price. Overall, 23% of the auctions ended up with prices greater than the aftermarket price.¹⁰

¹⁰Recall that we conducted auctions with two bidders. With more than two bidders, the probability of overpaying may increase.

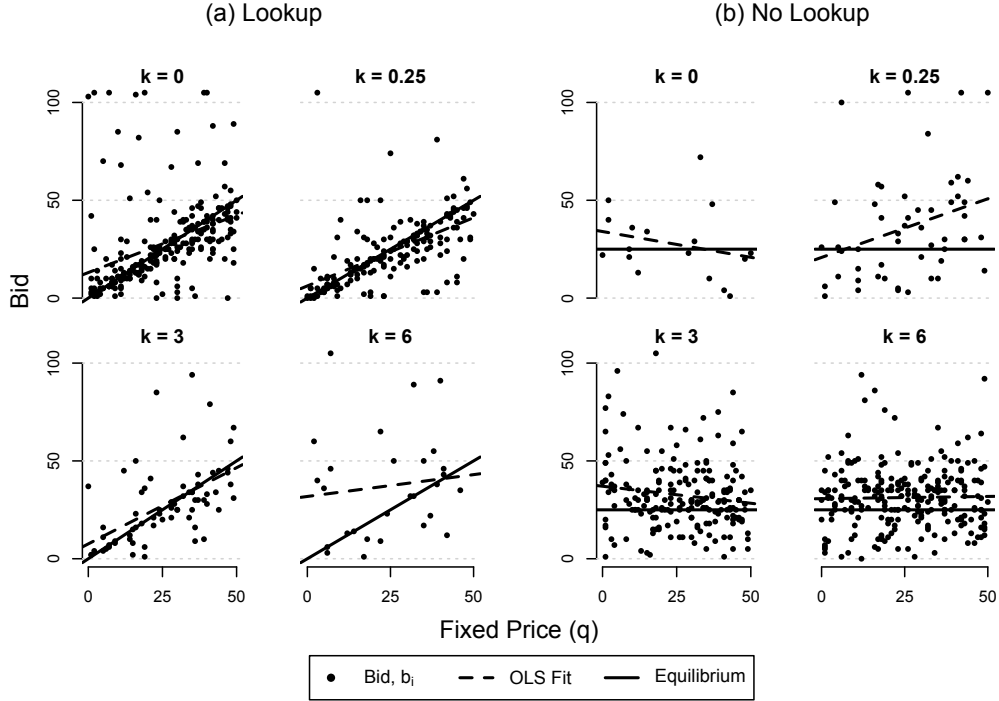


FIGURE 1: The relationship between bids and fixed prices by opportunity cost level and the decision to lookup the fixed price

k	$Pr(\text{Lookup})$	$Pr(\text{Overbid} \text{Lookup})$	$Pr(\text{Overbid} \text{NoLookup})$	$Pr(\text{Overbid})$
0	.93	.35	.59	.37
0.25	.77	.26	.60	.34
3	.26	.36	.54	.49
6	.10	.54	.59	.59
Total	.50	.33	.57	.45
N	1024	510	514	1024

TABLE 1: Summary statistics on the probability of overbidding, conditional on k .

To support these observations statistically and to explain both the decision to lookup the fixed price and the probability of overbidding (relative to the fixed price) in the auction, we estimate a system of two seemingly unrelated probit equations. In the first equation, the dependent variable λ takes a value of 1 if the bidder chose to lookup the fixed price and a value of 0 otherwise; the independent variables include a constant term, and the expected earnings of the slider task (computed as the product of the piece rate, k_t , and the number of sliders a subject correctly placed in part I (Productivity Measurement),

e_i^I , which measures individual productivity at the slider task). In the second equation, the dependent variable θ takes a value of 1 if a bidder bid more than the fixed price in the auction and 0 otherwise. The independent variables are a dummy variable indicating whether the subject chose to lookup the fixed price and a subject's value for the item. We also include a control for the fixed price, since overbidding should be less likely, the higher is the fixed price. In both equations, standard errors are clustered at the subject level to control for heteroskedasticity.

$$\lambda_i = \alpha_0 + \alpha_1 k_{i,t} e_i^I + \epsilon \quad (7)$$

$$\theta_i = \beta_0 + \beta_1 \lambda_{i,t} + \beta_2 v_{i,t} + \beta_3 q_t + \epsilon \quad (8)$$

Table 2 displays the output of the regression; one session is excluded from the regression because a software error left us without an estimate of productivity for those subjects (e_i^I).¹¹ As expected, in equation (7) the estimated coefficient on $k_t \times e_i^I$ is negative and significant indicating that the probability of looking up the fixed price decreases with a subject's opportunity cost. Moreover, in equation (8), a negative and significant estimated coefficient on the decision to lookup the fixed price indicates that overbidding is significantly driven by opportunity cost. When bidders choose to lookup the price, they are significantly less likely to overbid than when they instead choose to participate in the slider task.

Finding 1: Overbidding relative to the aftermarket price increases as the opportunity cost of becoming informed increases.

4.2 Information Acquisition and the Rationality of Inattention

Under the theory, individuals acquire information only when doing so is worthwhile. In particular, theory predicts that individuals follow a cutoff strategy, acquiring information whenever their opportunity cost is sufficiently low – that is, whenever $k_{i,t}$ is below a subject-specific threshold that depends in the experiment on a subject's productivity in the slider task. We test whether subjects play cutoff strategies using lookup decisions for each subject at each level of k . We find that 50 out of 64 subjects (78%) use a

¹¹Our results are robust to including that session and replacing $k_{i,t} e_i^I$ with $k_{i,t}$ in the selection equation. Estimates available upon request.

	Estimates
Eq. 7: Lookup	
<i>Expected Earnings</i>	-0.029***
$(k \times e_i^I)$	(0.003)
<i>Constant</i>	0.952***
	(0.131)
Eq. 8: Overbid	
<i>Lookup</i> ($\lambda_{i,t}$)	-0.679***
	(0.180)
<i>Value</i> ($v_{i,t}$)	0.010***
	(0.003)
<i>Posted Price</i> (q_t)	-0.037***
	(0.004)
<i>Constant</i>	0.356
	(0.249)
Log Lik.	-872.059
N	896
Clustered standard errors in parentheses	
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$	

TABLE 2: Bivariate probit regression explaining the probability of looking up the after-market price and the probability of overbidding.

cutoff strategy; that is, there is a cutoff value \bar{k}_i such that subject i always looks up the aftermarket price when $k_{i,t} < \bar{k}_i$, but never looks it up when $k_{i,t} > \bar{k}_i$.

While adherence to cutoff strategies is perfect for a large majority of subjects, a weaker test reported in Table 2 strongly confirms the comparative static prediction that subjects are less likely to lookup the aftermarket price when their opportunity cost of doing so, as measured by their expected earnings in the slider task, is higher.

Finding 2: Subjects generally acquire information about the aftermarket price in a manner consistent with a cutoff strategy.

A disadvantage of a cutoff strategy in which a bidder sometimes remains uninformed

is that it can lead to the bidder’s curse, where the auction winner pays a price exceeding the aftermarket price. In equilibrium, this can occur if both bidders remain uninformed. Thus, under the theory, it is individually rational to occasionally fall prey to the bidder’s curse, because avoiding it in any circumstance is too costly.

To see the empirical earnings impact of remaining uninformed, we estimate a linear panel regression, with heteroskedasticity-robust standard errors and random effects for each subject to control for individual differences, where the dependent variable is the total net gain of an auction winner i in market t ($\pi_{i,t}$), relative to obtaining the good in the aftermarket. Let p_t represent the price paid by the winning bidder in the auction during market t (which is always less than or equal to the winning bid), and let q_t represent the fixed price at which the item can be purchased in the aftermarket. The winning bidder’s earnings from the auction are then $q_t - p_t$. However, this is not the only component of earnings during a market since bidders may also participate in 30 seconds of the effort task to earn additional money. Let e_t^{II} be the number of sliders correctly adjusted in market t during part II; then total earnings from the effort task are $k_t e_t^{II}$, where k_t is the value of correctly adjusting a single slider. Thus, we define total net gains in market t as $\pi_t = k_t e_t^{II} + (q_t - p_t)$.¹² The independent variables of the regression include a constant term, a dummy variable that takes a value of 1 when the bidder was uninformed and 0 otherwise ($!\lambda_{i,t}$), and an interaction between the uninformed dummy variable and the slider task earnings ($k_t e_{i,t}^{II}$).

$$\pi_{i,t} = \gamma_0 + \gamma_1 !\lambda_{i,t} + \gamma_2 !\lambda_{i,t} k_t e_{i,t}^{II} + \eta_i + \epsilon \quad (9)$$

Table 3 reports the output of this regression. A positive and significant coefficient on the constant term indicates that informed bidders stand to gain approximately 9 ECU in the auction. Uninformed bidders sacrifice approximately 4 ECU because they tend to overbid the aftermarket price; however, we can reject the hypothesis that an uninformed bidder who earned nothing in the slider task would receive a total payoff of 0 (i.e. a Wald test that $\gamma_0 + \gamma_1 = 0$ yields $\chi^2 = 7.51$, p -value < 0.01). In this auction, an uninformed bidder can still get a deal when her counterpart underbids but will never overpay the aftermarket price when her counterpart overbids her. The coefficient on the interaction term (mechanically) equals 1, and evaluated at the mean of $k_t e_{i,t}^{II}$, which is 13, we can reject the null hypothesis that uninformed bidders and informed bidders receive the same

¹²Our analysis in this section is in terms of experimental earnings and ignores any utility/disutility of the real effort task captured in the d_i term of the model in Section 2.

mean payments in favor of the alternative that uninformed bidders' earnings are greater (a Wald test that $\gamma_1 + 13 * \gamma_2 = 0$, yields $\chi^2 = 28.53$, p -value < 0.01).

	Total Net Gains ($\pi_{i,t}$)
<i>No Lookup</i> , ($!\lambda_{i,t}$)	-4.35** (1.88)
<i>Slider Task Earnings</i> ($!\lambda_{i,t} \times k_t \times e_{i,t}^{II}$)	0.97*** (0.04)
<i>Constant</i>	8.84*** (1.14)
R^2	0.62
N	512

Clustered standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE 3: Panel linear regression explaining the total payoffs earned by auction winners.

Finding 3: On average, uninformed bidders earn more than they would have earned by becoming informed about the aftermarket price.

Moreover, as noted above, overbidding relative to an aftermarket price (as opposed to one's value) need not have negative welfare implications. If the opportunity cost of looking up the fixed price is sufficiently high, the gains from bidding correctly in the auction may be outweighed by the forgone gains from alternative actions. To identify the impact of overbidding on earnings, we compute the total net gains ($\pi_{i,t}$) for the winner of each auction in which the winning *bid* was greater than the fixed price.

Figure 2 displays box-and-whisker plots of total net gains ($\pi_{i,t}$) of auction winners who submitted bids greater than the fixed price for each level of opportunity cost k (i.e. those bidders subject to the bidder's curse). In general, overbidding the fixed price had a *positive* impact on monetary earnings after taking into account earnings from the slider task; auction bidders who did not lookup the aftermarket price, overbid, and then won the auction earned an average of 27 ECU. Table 4 decomposes these earnings and shows that they would have earned only 4 ECU on average had they instead become informed

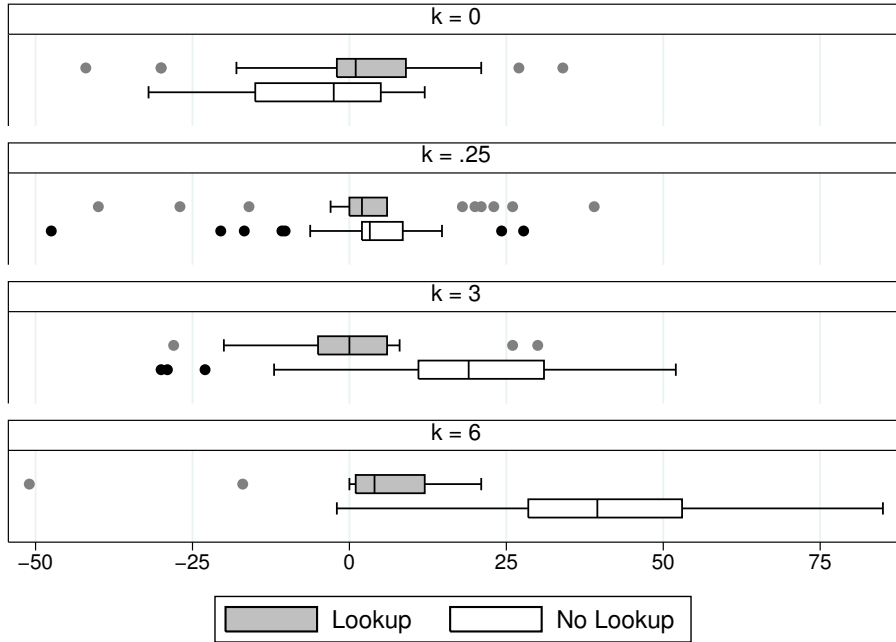


FIGURE 2: Box-and-whisker plots of total net gains (auction earnings + slider earnings) in ECU, for auction winners with $b_{i,t} > q_t$, conditional on the decision to lookup the posted price.

and bid the aftermarket price (q_t). Gains from the slider task swamp any losses from overbidding. Cf. Figure 3, which displays empirical cumulative distribution functions of observed differences between the auction price and the fixed price, conditional on whether the bidder chose to lookup the fixed price.

Finding 4: Bidders subject to the bidder’s curse (i.e. $b_{i,t} > q_t$) earn positive payoffs on average and earn more than they would have earned by becoming informed.

	Uninformed (Actual)	Informed (Counterfactual)
Net Auction Savings ($q_t - p_t$)	0.48	4.11
Slider Task Earnings	26.65	0
Total Payoff ($\pi_{i,t}$)	27.13	4.11

TABLE 4: Earnings among overbidders who win the auction.

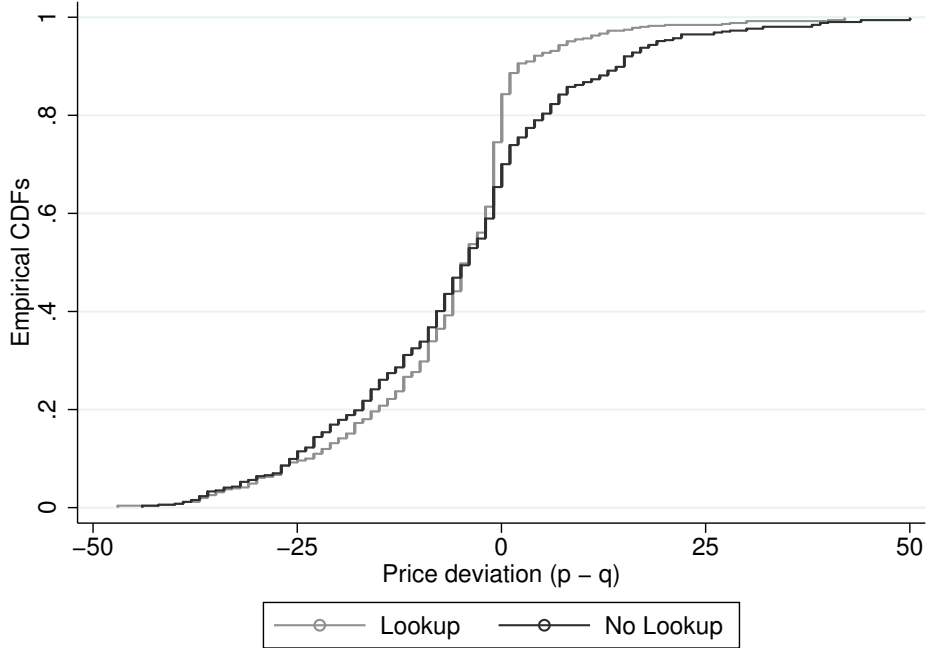


FIGURE 3: CDFs of net auction savings ($p_t - q_t$) for auction winners conditional on the decision to lookup the fixed price and the value of k

4.3 Bidding Strategies

In addition to the general implication that the opportunity cost of acquiring information about the aftermarket price drives overbidding, the model makes specific testable predictions about subjects' bidding strategies, conditional on information acquisition. In equilibrium, informed subjects bid the aftermarket price, and uninformed subjects bid the expected aftermarket price, $\mathbb{E}q$, which may be greater or lower than the actual aftermarket price. With our parameters, on average, they neither overbid nor underbid the aftermarket price.

To test these predictions, we report GLS panel regressions where the dependent variable is the difference between the bid and the aftermarket price, and the independent variables include a constant term and a dummy that takes a value of 1 if the bidder chose to lookup the price and 0 otherwise. We include random effects for each subject to control for repeated measures and we cluster standard errors to control for heteroskedasticity. Column (1) of Table 5 reports the regression results. A positive and significant coefficient on the Constant term indicates that uninformed subjects overbid the aftermarket price by 6 ECU on average, so that bids are biased upward relative to the theory. In contrast, bids made by informed bidders seem to be unbiased: The negative and significant coefficient

on the Lookup dummy indicates that informed bidders reduce their bids significantly, and a Wald test cannot reject the null hypothesis that the Constant term and the Lookup term sum to 0, meaning that overbidding among informed subjects is not significantly different from 0.

	Bid - Aftermarket Price	Bid
	$b_{i,t} - q_t$	$b_{i,t}$
<i>Lookup</i> , ($\lambda_{i,t}$)	-4.766*** (1.735)	-20.782*** (2.299)
<i>Lookup</i> \times <i>Aftermarket Price</i>		0.597*** (0.066)
Constant	6.375*** (1.817)	32.398*** (1.604)
R^2	0.012	0.125
N	1024	1024

Clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE 5: GLS regression analysis of bidding behavior.

The second regression in Table 5 allows us to compare observed behavior directly to the bidding strategies predicted by the theory. We report GLS panel regressions in which the dependent variable is the bid, and the independent variables include a constant term, a dummy that takes a value of 1 if the bidder chose to lookup the price and 0 otherwise, and an interaction between the aftermarket price and the lookup dummy, which allows us to test the extent to which informed subjects' bids reflect the aftermarket price. We include random effects for each subject to control for repeated measures, and we cluster standard errors to control for heteroskedasticity.

Column (2) of Table 5 displays the regression output. Informed bidders should bid exactly the aftermarket price, which implies that the coefficient on the *Lookup* \times *Aftermarket Price* interaction will equal 1 and that the sum of the coefficients on the Constant term and the *Lookup* dummy will equal to 0. Uninformed bidders, on the other hand, should bid $\mathbb{E}q$, which implies that the coefficient on the Constant term will equal 25 ECU, given our distribution of aftermarket prices. The regression coefficients reveal that un-

informed bidders overbid substantially, and although informed subjects' bids respond significantly to the aftermarket price, and the lookup dummy is substantially negative, we still observe a tendency to overbid the aftermarket price on average. A Wald test rejects the joint hypothesis that all of the predictions are true ($\chi^2 = 22.81$, p -value < 0.01). We summarize the regression results as follows.

Finding 5: Subjects on average tend to overbid relative to the theoretical predictions whether they inform themselves about the aftermarket price or not, although upward bias is pronounced only among uninformed bidders.

In addition to being slightly too high on average, bids in the experiment are also more diffuse than the equilibrium (point) predicts, especially among uninformed bidders (see Figure 4). For informed subjects, 13% of all bids are exactly equal to the aftermarket price, while 31% are within 1 ECU of the aftermarket price. Nearly half of all informed bids are within 5 ECU of the aftermarket price and 70% of all such bids are within 10 ECU. Uninformed subjects also do not bid exactly as predicted by the theory. In fact, only 7% of uninformed bids are exactly equal to $\mathbb{E}q$, i.e. 25 ECU, while only 12% are within 1 ECU and 51% within 10 ECU.

Finding 6: Bids are more diffuse than predicted by the theory, particularly when subjects are uninformed.

4.4 Behavioral Explanations of Bidding

Overbidding and excess bid dispersion by uninformed subjects remains a puzzle. One candidate explanation comes from the Malmendier and Lee (2011) notion of limited attention. We note that our data are inconsistent with bidders having limited attention in the sense described in their Proposition 3. If subjects simply ignore the aftermarket fixed price option, then it is optimal for them to bid $b_i = v_i$ as in Vickrey (1961). Only 4 out of 1024 bids are equal to v_i and 9% of bids are greater than 50 (the minimum of v_i), so this cannot explain our data.

Another possibility is that Finding 5 may be related to existing evidence of *on average* overbidding in second price auctions without an aftermarket (Kagel and Levin, 1993; Ariely et al., 2005; Cooper and Fang, 2008; Georganas et al., 2017).¹³ We observe no

¹³Unlike a standard clock auction that ends when only one bidder remains, our auction does not have

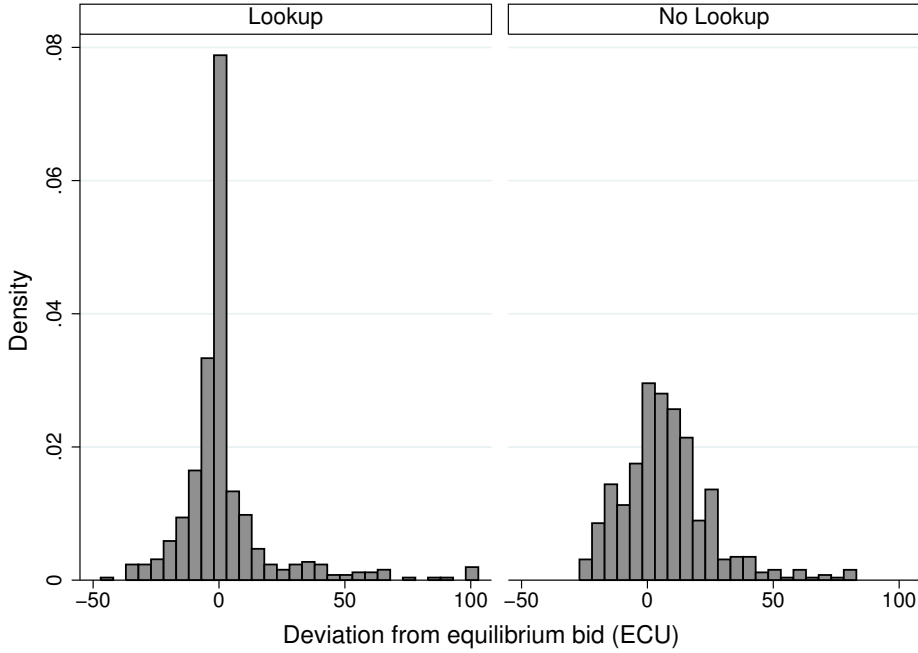


FIGURE 4: Histograms of deviations from equilibrium bidding ($b^* - b_{i,t}$)

consistent trends in the likelihood of overbidding across auctions; for example, informed (uninformed) subjects overbid with probability 0.36 (0.57) in the first five auctions and probability 0.32 (0.55) in the final five auctions. This is consistent with prior evidence, in which even substantial experience does not eliminate overbidding (Kagel and Levin, 1993); however, we cannot rule out the possibility that additional experience would eventually eliminate overbidding.

Nevertheless, our data cannot be explained by many leading models that have been used to account for overbidding. For instance, introducing an additively-separable utility of winning would counterfactually predict a similar degree of overbidding by both informed and uninformed bidders (Cox et al., 1988). Furthermore, while the level-k model has been successful in explaining overbidding in other settings (e.g. Crawford and Iriberri, 2007), it cannot explain overbidding in a second-price auction under the standard assumption that the level-0 type uniformly randomizes over bids, since informed bidders have a weakly dominant strategy to bid q and uninformed bidders of level-1 and higher would have a best response of bidding $\mathbb{E}q$.

a solution in obviously-dominant strategies in the sense of Li (2015). Li suggests this might explain why past studies have tended to find less overbidding in ascending clock auctions than in otherwise equivalent second-price sealed-bid auctions. Our finding of overbidding lends support to his claim.

One possible explanation relies on noisy behavior, as captured by quantal response equilibrium (Goeree et al., 2016). QRE assumes that each player behaves stochastically, but is more likely to choose actions that yield a higher expected value, given the stochastic behavior of other players. Thus the model can account for bid dispersion relative to the PBE. Moreover, in our setting QRE can generate more dispersion in bids by uninformed than by informed bidders (consistent with Finding 6) and overbidding on average for both informed and uninformed bidders (consistent with Finding 5).¹⁴

To see the intuition for higher dispersion among the uninformed, consider the incentives faced by agents when others play according to the PBE strategies given in section 2. Notice that an uninformed bidder will be indifferent across all bids. When combined with QRE assumptions, the presence of this “flat maximum” generates dispersion in bids for the uninformed.¹⁵ Also note that a given informed bidder strictly prefers to place a winning bid when the aftermarket price exceeds 25 and to place a losing bid when the aftermarket price is below 25. This intuition also extends to the QRE model where play is noisy, and players anticipate the noisy play of others. Under QRE informed bidders have incentives to bid close to q , while uninformed bidders have relatively weaker incentives to bid close to 25. Together these forces will tend to generate more dispersion in the bids of uninformed bidders relative to informed bidders.

To see the intuition for why QRE generates overbidding, note that the expected payoff consequences of deviations are roughly symmetric around the PBE bid. Since possible deviations are truncated below by 0 and above by 105, there are more possible deviant overbids than underbids. As a result, noise in the QRE will tend to generate overbidding on average and at the median. Then, since uninformed bidders exhibit more dispersion, they will tend to overbid more in the model.

In Appendix C we explore this intuition quantitatively by estimating a QRE model to our bidding data and describing how the fit model can help explain some features, as well as some features of our data that the model gets wrong. The main results of this exercise give a quantitative sense of the extent to which the intuition above can explain bid dispersion and overbidding, and are summarized in Finding 7.

¹⁴Another implication of QRE is that there is a positive probability that a subject will not buy in the aftermarket after losing in the auction, even though doing so is always payoff-maximizing. Surprisingly, this implication of QRE finds some qualitative support in the data. 12/64 subjects failed to purchase in the aftermarket after losing in the auction at least once. One subject did so 7 times, and in total there were 29 instances out of 512 in which auction losers failed to purchase in the aftermarket.

¹⁵The flat maximum problem is well-known in the context of second-price auctions (Harrison, 1989).

Finding 7: A model of noisy behavior can plausibly account for both observed overbidding and excess dispersion.

5 Discussion

Our data suggest that overbidding a fixed price option becomes more likely as the opportunity cost of acquiring price information increases - consistent with the comparative statics of our model. We show that this drives much of the observed *ex-post overbidding*, which is commonly called the bidder’s curse.¹⁶ While the data are not perfectly consistent with the theory, they nevertheless support an interpretation in terms of rational ignorance.

While we were motivated by an online pricing puzzle, we believe our results may be useful for the study of anomalies in markets more generally. One broader implication of this paper is that it can be important to account for opportunity costs of attention in constructing a theoretical benchmark against which we might define behavior as anomalous. When opportunity cost is uncontrolled and unobservable, as in the field, behavior that appears consistent with the bidder’s curse may instead reflect rational ignorance. The presence of such tradeoffs is not unique to auction environments, and the opportunity cost frugality of many human decision-making processes is well-known (Gigerenzer and Goldstein, 1996). In this sense, our results are analogous to previous findings that rationalize price dispersion via heterogeneous search costs, for example, the finding that unemployed people and retirees are more likely to buy items at a discount and use coupons (Narasimhan, 1984; Aguiar and Hurst, 2007; Kaplan and Menzio, 2015). In that setting as in ours, accounting for opportunity cost makes specific predictions about who will and won’t “overpay”, and if we ignore opportunity cost, higher prices paid by some buyers might be mistaken for a more general “buyer’s curse”.

¹⁶As noted earlier, a similar mechanism can explain the apparently aberrant phenomenon of jump bidding (Easley and Tenorio, 2004). In a multi-round auction in which submitting an additional bid is (opportunity) costly, jump bidding reduces incurred costs. We conjecture that a hypothetical experiment similar to ours, in which subjects must forgo a potential gain each time they submit a new bid would also likely demonstrate that the frequency and magnitude of jump bidding is increasing in the opportunity cost of the marginal bid.

References

- AGUIAR, M. AND E. HURST, “Life-cycle Prices and Production,” American Economic Review 97 (2007), 1533–1559.
- ARIELY, D., A. OCKENFELS AND A. E. ROTH, “An experimental analysis of ending rules in internet auctions,” RAND Journal of Economics 36 (2005), 890–907.
- BACKUS, M., J. PODWOL AND H. S. SCHNEIDER, “Search Costs and Equilibrium Price Dispersion in Auction Markets,” European Economic Review 74 (2014), 173–192.
- BROWN, J., T. HOSSAIN AND J. MORGAN, “Shrouded attributes and information suppression: Evidence from the field,” Quarterly Journal of Economics 125 (2010), 859–876.
- CASSADY, R., Auctions and Auctioneering (Univ of California Press, 1980).
- CHRISTOPHER, A., “Strategic Jump Bidding in English Auctions,” Review of Economic Studies 65 (1998), 185–210.
- COOPER, D. J. AND H. FANG, “Understanding overbidding in second price auctions: An experimental study,” Economic Journal 118 (2008), 490–529.
- COX, J. C., V. L. SMITH AND J. M. WALKER, “Theory and individual behavior of first-price auctions,” Journal of Risk and Uncertainty 1 (1988), 61–99.
- CRAMTON, P., “The FCC spectrum auctions: An early assessment,” Journal of Economics & Management Strategy 6 (1997), 431–495.
- CRAWFORD, V. P. AND N. IRIBERRI, “Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner’s Curse and Overbidding in Private-Value Auctions?,” Econometrica 75 (2007), 1721–1770.
- DURHAM, Y., M. R. ROELOFS, T. A. SORENSEN AND S. S. STANDIFIRD, “A Laboratory Study of Auctions with A Buy Price,” Economic Inquiry 51 (2013), 1357–1373.
- EASLEY, R. F. AND R. TENORIO, “Jump bidding strategies in internet auctions,” Management Science 50 (2004), 1407–1419.
- ENGELMANN, D. AND E. G. WOLFSTETTER, “A Proxy Bidding Mechanism that Elicits all Bids in an English Clock Auction Experiment,” Humboldt University Working Paper, 2009.

- GEORGANAS, S., D. LEVIN AND P. MCGEE, “Optimistic irrationality and overbidding in private value auctions,” Experimental Economics (2017), 1–21.
- GIGERENZER, G. AND D. G. GOLDSTEIN, “Reasoning the fast and frugal way: Models of bounded rationality.,” Psychological Review 103 (1996), 650 – 669.
- GILL, D. AND V. PROWSE, “A structural analysis of disappointment aversion in a real effort competition,” American Economic Review 102 (2012), 469–503.
- GOEREE, J., C. HOLT AND T. PALFREY, Quantal Response Equilibrium: A Stochastic Theory of Games (Princeton University Press, 2016).
- HARRISON, G., “Theory and misbehavior of first-price auctions,” American Economic Review (1989), 749–762.
- HOSSAIN, T. AND J. MORGAN, “... Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay,” The BE Journal of Economic Analysis & Policy 5 (2006), 1–30.
- ISAAC, R. M., T. C. SALMON AND A. ZILLANTE, “A theory of jump bidding in ascending auctions,” Journal of Economic Behavior & Organization 62 (2007), 144–164.
- JONES, M. T., “Bidding fever in eBay auctions of Amazon.com gift certificates,” Economics Letters 113 (2011), 5–7.
- KAGEL, J. H., R. M. HARSTAD AND D. LEVIN, “Information impact and allocation rules in auctions with affiliated private values: A laboratory study,” Econometrica 55 (1987), 1275–1304.
- KAGEL, J. H. AND D. LEVIN, “The winner’s curse and public information in common value auctions,” American Economic Review 76 (1986), 894–920.
- , “Independent Private Value Auctions: Bidder Behavior in First-, Second- and Third-Price Auctions with Varying Number of Bidders,” Economic Journal 103 (1993), 868–879.
- KAPLAN, G. AND G. MENZIO, “The Morphology of Price Dispersion,” International Economic Review 56 (2015), 1165–1206.
- KLEMPERER, P., “What really matters in auction design?,” Journal of Economic Perspectives 16 (2002), 169–189.

- LEVIN, D. AND J. P. REISS, “Could we overcome the winner’s curse by behavioral auction design?,” Working Paper (2017).
- LI, S., “Obviously strategy-proof mechanisms,” Working Paper, 2015.
- MALMENDIER, U. AND Y. H. LEE, “The bidder’s curse,” American Economic Review 101 (2011), 749–787.
- MCCABE, K. A., S. J. RASSENTI AND V. L. SMITH, “Auction Institutional Design: Theory and Behavior of Simultaneous Multiple-Unit Generalizations of the Dutch and English Auctions,” American Economic Review 80 (1990), 1276–1283.
- NARASIMHAN, C., “A Price Discrimination Theory of Coupons,” Marketing Science 3 (1984), 128–147.
- PODWOL, J. U. AND H. S. SCHNEIDER, “Nonstandard bidder behavior in real-world auctions,” European Economic Review 83 (2016), 198–212.
- R DEVELOPMENT CORE TEAM, R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria (2015), ISBN 3-900051-07-0.
- ROTH, A. E. AND A. OCKENFELS, “Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet,” American Economic Review 92 (2002), 1093–1103.
- SCHNEIDER, H. S., “The Bidder’s Curse: Comment,” American Economic Review 106 (2016), 1182–1194.
- STIGLER, G. J., “Imperfections in the capital market,” Journal of Political Economy 75 (1967), 287–292.
- VICKREY, W., “Counterspeculation, Auctions and Competitive Sealed Tenders,” Journal of Finance 16 (1961), 8–37.

Appendix (for online publication)

A Theoretical Results

Lemma 1. *Given ρ , a there exists a value of b^* that solves the equation (1).*

Proof. The left-hand side of (1) is decreasing in b^* and the right-hand side increasing, both sides continuously since G has a continuous density. The right-hand side achieves a value of 0 at $b^* = 0$ and $(1 - \rho^{n-1} - (1 - \rho)^{n-1})(\bar{q} - \mathbb{E}q) > 0$ at $q = \bar{q}$. The left-hand side achieves a value of $(1 - \rho)^{n-1}\mathbb{E}q > 0$ at $b^* = 0$ and a value of $(1 - \rho)^{n-1}[\mathbb{E}q - \bar{q}] < 0$ at $b^* = \bar{q}$. \square

Lemma 2. *In any symmetric PBE, informed bidders bid q with probability 1.*

Proof. Consider a candidate set of symmetric perfect Bayesian equilibrium strategies. Suppose that each player becomes informed with a strictly positive probability ex-ante. Without loss of generality, a given player i can view other players' strategies as inducing a single (aggregate) mixed continuation bidding strategy over bids used by all informed players when q is the aftermarket price. Let σ denote the cumulative distribution function induced by this mixed strategy. Similarly, we can view the continuation strategies used when players choose to be uninformed as a inducing a single (aggregate) mixed continuation bidding strategy over bids used by all uninformed players.

For this set of candidate strategies to be supportable in a symmetric equilibrium, a player who is informed that the aftermarket price is q must be indifferent between any bids in the support of σ . The verbal argument below establishes that if σ puts a strictly positive probability on bids not equal to q , then bidding q achieves a strictly lower expected price than some bid in the support of σ .

So suppose that σ puts a strictly positive probability on bids strictly below q . Then if i bids q , the case where $-i$ all bid less than q occurs with strictly positive probability; in this case, i would pay the highest bid of $-i$, which is strictly less than q . But if i had instead made a bid $b < q$ in the support of σ for which $\sigma(q) > \sigma(b)$, i would pay the same price as if she had bid q when $-i$ all bid strictly below b . She must sometimes lose to bids equal to or exceeding b , in which case she would pay q in the aftermarket. If instead she had bid q , she would pay a price strictly below q . When $-i$ have at least one bid exceeding q , bidding q and bidding b will both lose and result in paying the same aftermarket price. Thus, bidding q must lead to a strictly lower interim expected price

than bidding b in this case. It follows that σ cannot induce bids strictly below q with positive probability.

Similarly, if σ prescribes that bids strictly exceed q with strictly positive probability, we can apply the same logic to show that bidding q would always lead to a weakly lower price paid, and would with strictly positive probability lead to a strictly lower price paid than some bid exceeding q in the support of σ . Thus σ must prescribe bidding q with a probability of one. \square

Lemma 3. *In any symmetric PBE, uninformed bidders pay an expected price of $\mathbb{E}q$.*

Proof. Consider candidate symmetric PBE strategies. Without loss of generality, view continuation strategies used conditional on remaining uninformed as inducing a single aggregate mixed continuation bidding strategy which generates a cumulative distribution function σ for uninformed types; let Q denote its support.

In a symmetric PBE, each uninformed player must be indifferent between bidding any point in the support of σ . If $\sigma(\bar{q}) < 1$, then any player i could achieve a higher ex-ante payoff than any mixed strategy that prescribes bidding above \bar{q} with positive probability by modifying the continuation strategy to change bids above \bar{q} to \bar{q} and this would contradict our assumption that σ forms part of a PBE. Thus $\sigma(\bar{q}) = 1$.

Let $\hat{b} = \sup [b : \sigma(b) = 0]$. If $\sigma(\hat{b}) = 0$, then there must be a sequence b^n in Q with $b^n \rightarrow \hat{b}$. Since each $b^n \in Q$, each must yield the same expected price. Along this sequence, the equilibrium probability of winning the auction with a bid of b^n goes to zero; when a bid of b^n wins the auction, the ex-post payment is bounded by $\max[b^n, q]$ while the payment is q when the bid of b^n loses the auction. Thus, as $b^n \rightarrow \hat{b}$, bidder i 's expected payment for the good converges to $\mathbb{E}q$. Thus the expected price paid conditional on becoming uninformed must be $\mathbb{E}q$ for all bids made by uninformed bidders in any symmetric PBE in which $\sigma(\hat{b}) = 0$.

Now suppose $\sigma(\hat{b}) > 0$. Then (due to the tie breaking rule) the expected price paid when bidding \hat{b} is a convex combination of the expected prices paid in the limit of sequences of bids converging to \hat{b} from above and from below. But when making any bid below \hat{b} , a bidder always loses when an uninformed type is present, always pays q when an informed type is present, and thus has an expected price of $\mathbb{E}q$. Thus for her to be willing to bid \hat{b} (rather than below or slightly above), she must be indifferent between bidding below \hat{b} and the limit of bids converging to \hat{b} from above. Thus the expected price for an uninformed player who bids \hat{b} must be $\mathbb{E}q$ in a symmetric equilibrium; thus an uninformed player's

expected price given any bid in Q must be $\mathbb{E}q$. \square

Lemma 4. *When $n = 2$, in any symmetric PBE, uninformed bidders bid $\mathbb{E}q$ with probability of 1.*

Proof. Consider an arbitrary symmetric PBE. Summarize the probability distribution over other players' bidding strategy when uninformed by cumulative distribution function H , and let ρ denote the ex-ante probability that each other player becomes informed. Then, the expected price to a given uninformed bidder of bidding b_i is given by:

$$(1 - \rho) \left[\int_{\underline{q}}^{b_i} b dH(b) + (1 - H(b_i))\mathbb{E}q \right] + \rho\mathbb{E}q$$

This is weakly decreasing in b_i when $b_i < \mathbb{E}q$ and weakly increasing in b_i when $b_i > \mathbb{E}q$ and when compared to bidding $b_i = \mathbb{E}q$, any bid $b'_i < \mathbb{E}q$ must yield a strictly higher expected price whenever H assigns positive probability to the interval $[b'_i, \mathbb{E}q)$. Thus H must assign a probability of zero to bids below $\mathbb{E}q$. By a similar argument, H must assign a probability of zero to bids above $\mathbb{E}q$. Thus, uninformed bidders must bid $b^* = \mathbb{E}q$ with a probability of one in any symmetric PBE. \square

B Screenshots and Instructions

B.1 Screenshots

Below we include screenshots of the sequence of decisions made by subjects. At the end of the market, a new market started over and went through an identical sequence until all 16 markets were complete.

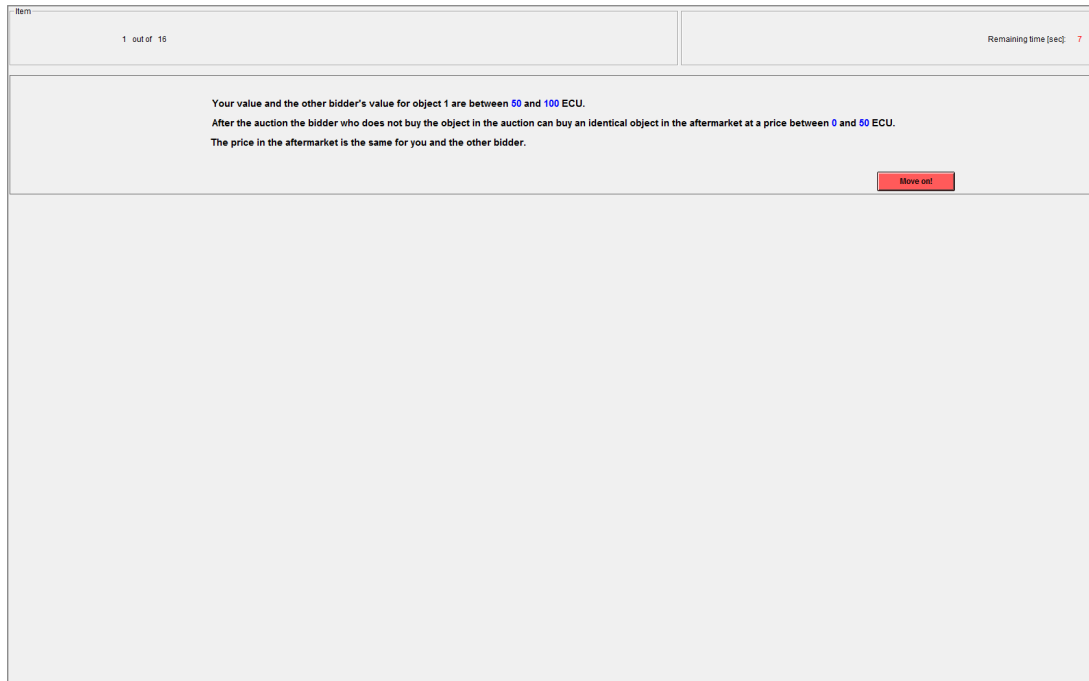


FIGURE B1: Screenshot of basic information

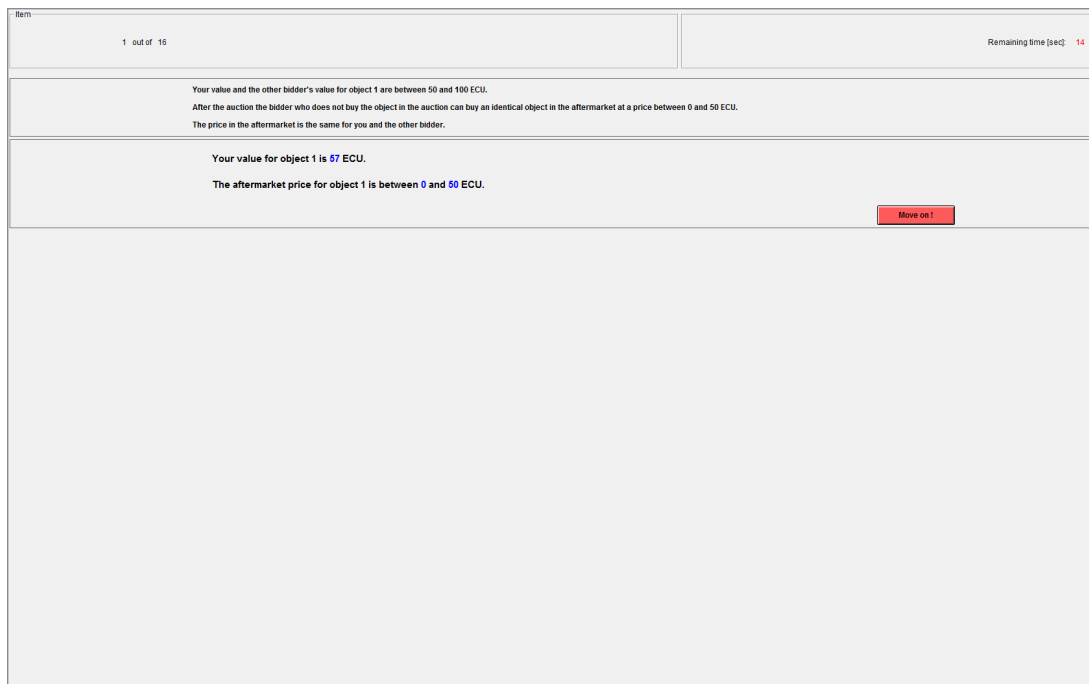


FIGURE B2: Screenshot of bidder value information

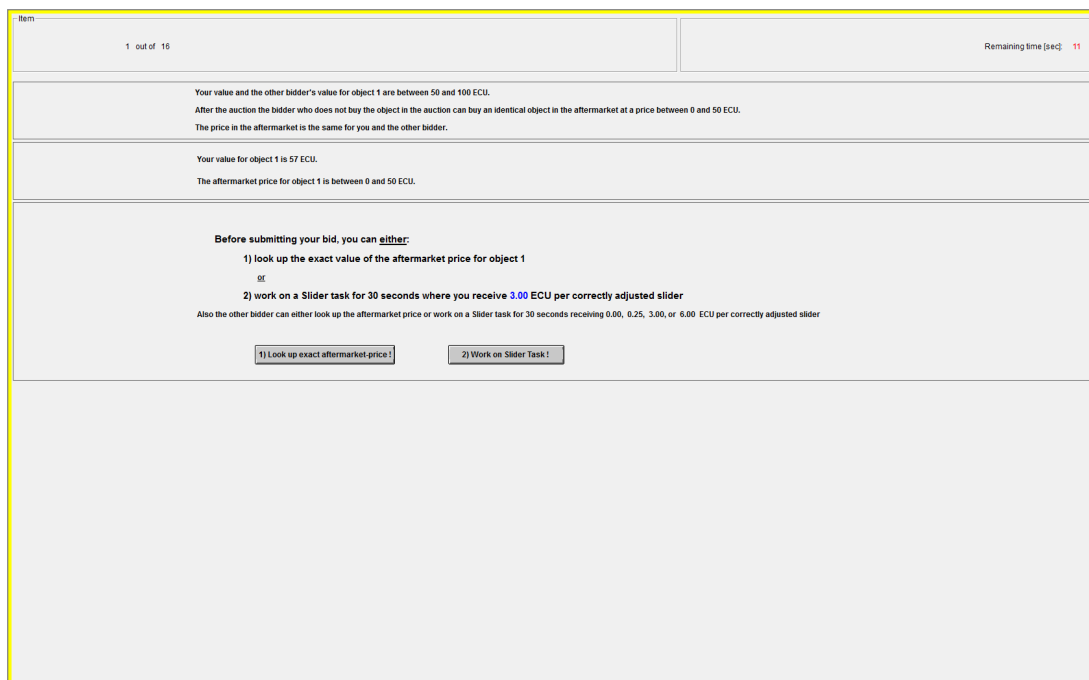


FIGURE B3: Screenshot of information acquisition decision

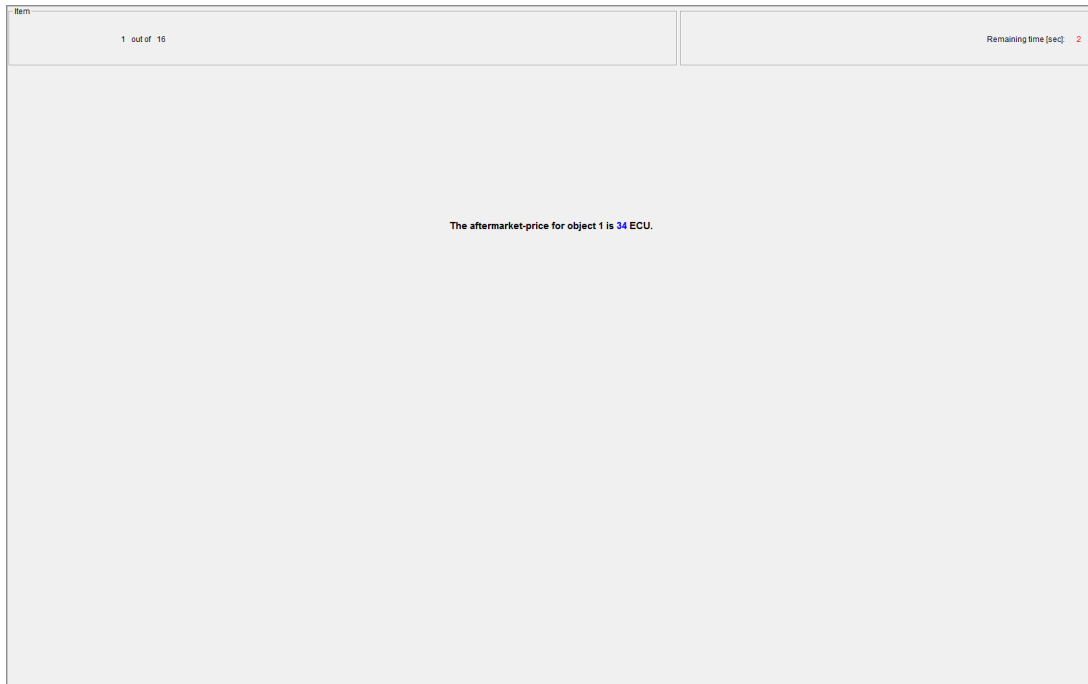


FIGURE B4: Screenshot for bidders who lookup the aftermarket price




FIGURE B5: Screenshot for bidders who choose the slider task

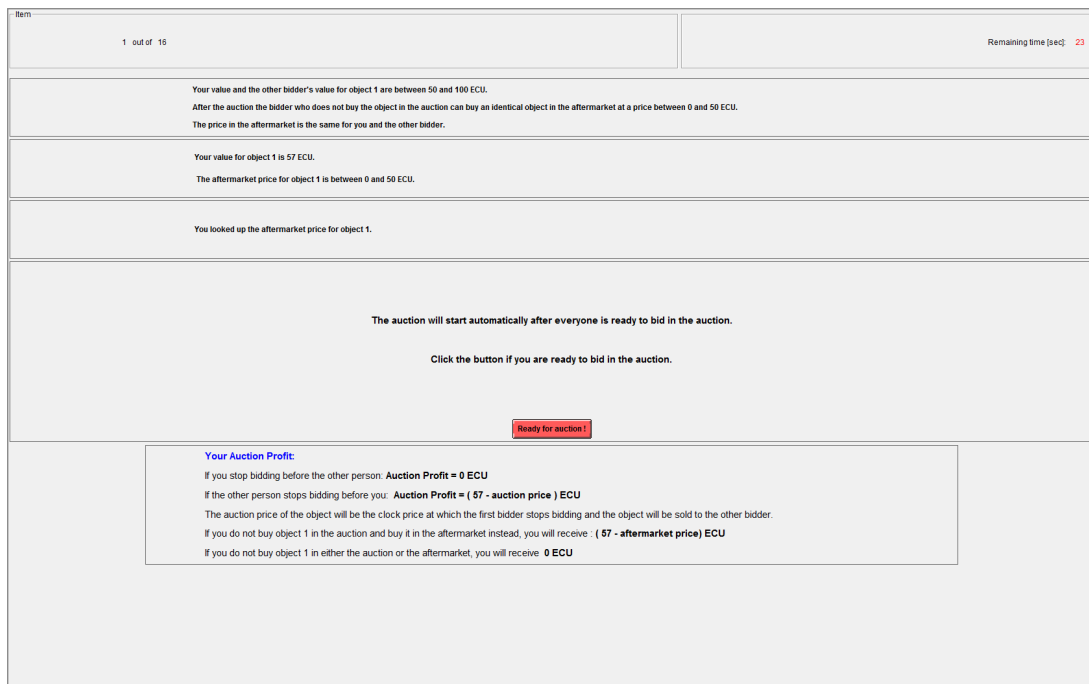


FIGURE B6: Screenshot of preliminary auction information

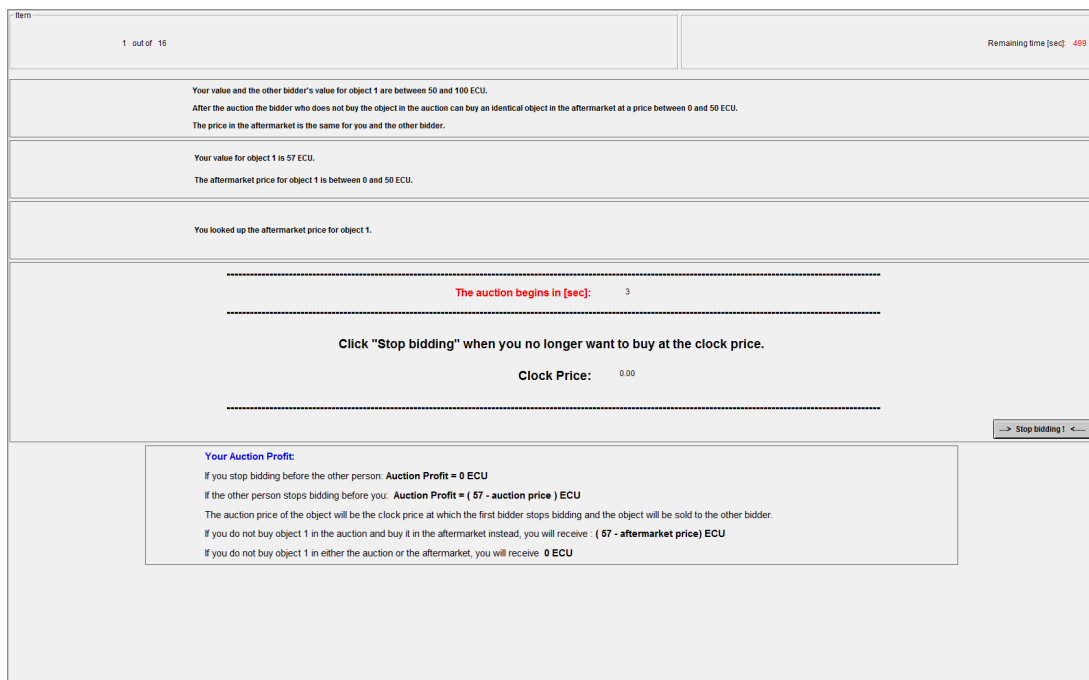


FIGURE B7: Screenshot before the clock starts ticking

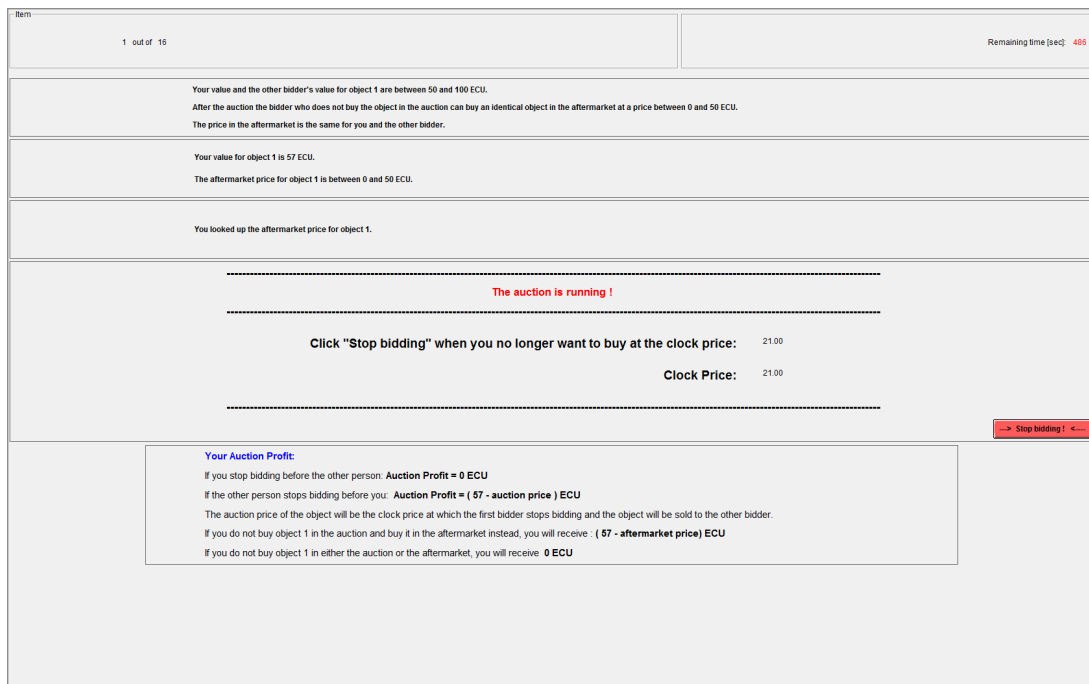


FIGURE B8: Screenshot of an active bidder during the auction

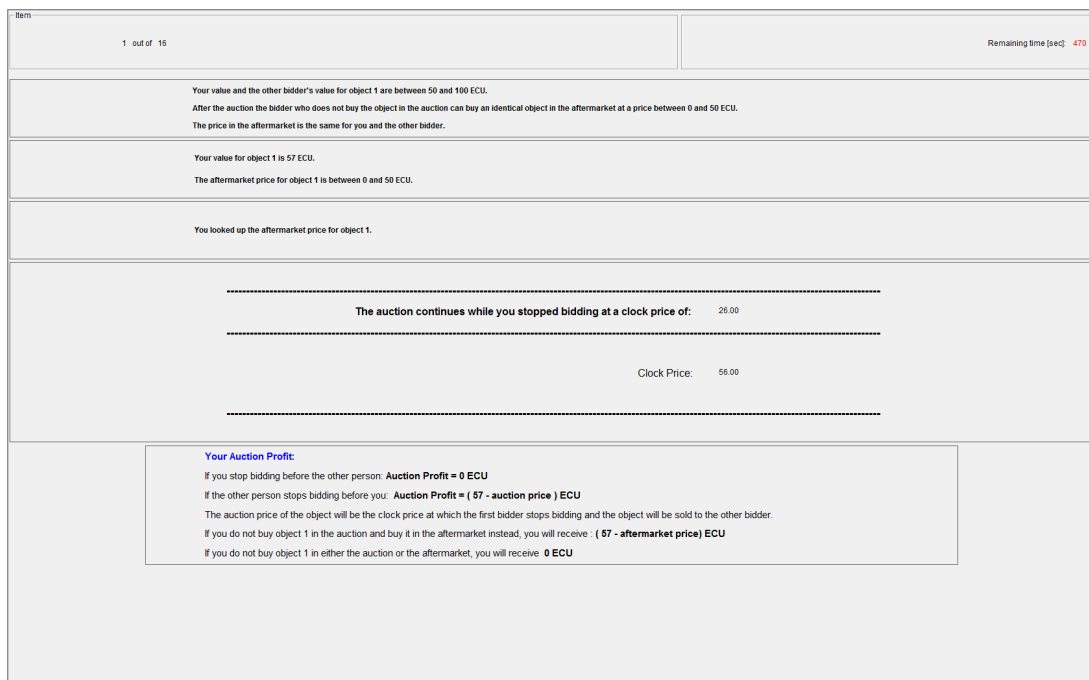


FIGURE B9: Screenshot after dropping out of the auction

Item 1 out of 16 Remaining time [sec]: 34

Your value and the other bidder's value for object 1 are between 50 and 100 ECU.
After the auction the bidder who does not buy the object in the auction can buy an identical object in the aftermarket at a price between 0 and 50 ECU.
The price in the aftermarket is the same for you and the other bidder.

Your value for object 1 is 57 ECU.
The aftermarket price for object 1 is between 0 and 50 ECU.

You looked up the aftermarket price for object 1.

The auction for object 1 finished, you stopped bidding at a price of: 26
The auction price of the item is: 26

stopped bidding at	Item values
28.00	57
80.00	97

The auction finished. You stopped bidding at a price of 26 ECU while someone else continued bidding so that you do not obtain the auctioned object.
In this period, your object value is 57 ECU.
The price of the object is 26 ECU. This is the price at which the next-to-last bidder stopped bidding.
Your profit is 0 ECU.
Would you like to buy an identical object in the aftermarket and earn your value of 57 minus the aftermarket-price of 34 = 23 ECU ?

Yes, buy at aftermarket-price!
No, do not buy at aftermarket-price!

FIGURE B10: Screenshot of the aftermarket for the auction loser

Item 1 out of 16 Remaining time [sec]: 0 Please reach a decision!

Your value and the other bidder's value for object 1 are between 50 and 100 ECU.
After the auction the bidder who does not buy the object in the auction can buy an identical object in the aftermarket at a price between 0 and 50 ECU.
The price in the aftermarket is the same for you and the other bidder.

Your value for object 1 is 97 ECU.
The aftermarket price for object 1 is between 0 and 50 ECU.

You worked on a Slider task for 1 minute where you received 3 ECU per correctly adjusted slider

The auction for object 1 finished, you stopped bidding at a price of: 80
The auction price of the item is: 26

stopped bidding at	Item values
28.00	57
80.00	97

The auction finished since everyone else except of you stopped bidding so that you obtain the auctioned object.
In this period, your object value is 97 ECU.
The price of the object is 26 ECU. This is the price at which the next-to-last bidder stopped bidding.
Your profit = object value - price = 97 - 26 = 71 ECU.
Since you received the object in the auction you cannot buy an identical object in the aftermarket that would have earned you your value of 97 minus the aftermarket-price of 34 = 63 ECU.

More De!

FIGURE B11: Screenshot summarizing the outcome for the auction winner

B.2 Instructions

General information

You are now participating in a decision making experiment. If you follow the instructions carefully, you can earn a considerable amount of money depending on your decisions, the decisions of the other participants, and an element of chance.

This set of instructions is for your private use only. **During the experiment you are not allowed to communicate with other participants.** If you have any questions, please raise your hand, and we will come to your seat and answer privately. Any violation of this rule excludes you immediately from the experiment.

During the experiment the outcome of your decisions will be measured in ECU (Experimental Currency Units) instead of Dollars. We will convert your total earnings into Dollars at a rate of 10 ECU = 1 Dollar at the end of the experiment and pay you in cash privately.

Outline of the Experiment

This experiment consists of 16 periods in which you can buy a fictitious object in either an **auction market** or an **aftermarket**. Each fictitious object will have some “value” to you – **you can think of your value as the amount of money (in ECU) that the experimenter will pay you for the object if you buy it.** So, if you have a value of 80 ECU for the object, and you buy it at a price P , you will earn 80 minus the price ($80 - P$) ECU.

In each period, your value will be a random whole number between 50 and 100 ECU. Each value is equally likely. **You will have a new random value in each period.**

The period starts with an auction. There are 2 bidders in each auction, you and someone else. **The other person bidding with you in each auction will change randomly.**

Other people also receive random values between 50 and 100 ECU for each object (with each value being equally likely), and all values are determined independently of one another. **You will only know your own value, and the other person in your auction will only know his/her own value.**

At the end of each auction, the **person who submitted the highest bid receives the object.** If there is a tie for the high bid, the object is assigned to one of the bidders by chance.

If you do not submit the highest bid, then you will not receive the object in the auction. However, anyone who did not buy the object in the auction **can buy an identical object after the auction at a fixed price that is the same for you and the other person in your auction. This price is called the aftermarket price, and it will be any number between 0 and 50 ECU where each number is equally likely.** Since the object for sale at the aftermarket price is identical to the auctioned object, **you have the same value for it as for the auctioned object.**

So, if your value is 80 ECU, and the aftermarket price is 42 ECU, and you choose to buy the object after the auction, you would receive $80 - 42 = 38$ ECU.

Before these 16 periods you will participate in 3 rounds of a “Slider Task” in which you earn money by manipulating objects on the computer screen. You will be paid for your performance in the 3rd round of the “Slider Task”.

You start off the experiment with an endowment of 70 ECU, which corresponds to your show up payment. At the end of the experiment, we will pay you the total of your earnings from the Slider Task plus your earnings from 1 *randomly chosen* period. This plus your Slider Task earnings will be added to your show-up payment and paid to you in **cash** at the end of the experiment.

Stage 1: The Slider Task Stage

The *Slider Task* stage allows you to earn money from a computerized task.

You will now participate in 3 rounds of an identical task **lasting 120 seconds per round** with a break of 60 seconds between each round. Rounds 1 and 2 allow you to practice the task and will not affect your earnings. The score that you will have obtained in round 3 will be used for determining your total earnings from the Slider Task Stage.

The task will consist of a screen with **48** sliders. Each slider is initially positioned at **0** and can be moved as far as **100**. Each slider has a number to its right showing its current position.

You can use the mouse or the keyboard's arrow keys in any way you like to move each slider. To use the arrow keys you have to select the slider that you want to adjust first by clicking it with the mouse. You can readjust the position of each slider as many times as you wish. You adjust the sliders to earn points. Specifically, **your "score" in the task will be the number of sliders positioned at exactly 50 at the end of the 120 seconds.**

Your score in round 3 will be converted to ECU at a rate of **1 point = 1 ECU**. For example, if after 120 seconds you have correctly placed 6 sliders at 50, then you have earned 6 points, which is converted to 6 ECU and will be paid to you at the end of the experiment.

Thus your Slider Earnings are equal to (Your Score in the third round x 1 ECU)

Are there any questions?

Stage 2: The Market Stage

Each period consists of two parts, the *Auction* and the *Aftermarket*.

In the *Auction*, each person **bids** for the fictitious object. Before bidding, you will learn your **value**, and you can *either*:

- 1) Look up the **aftermarket price** for 30 seconds before submitting your bid
- OR
- 2) Participate in a 30-second round of the **Slider Task** to earn ECUs.

If you choose to look up the aftermarket price, you **cannot** participate in the Slider Task, but for 30 seconds you will see the aftermarket price of an identical object that you can buy in the *Aftermarket* if you do not receive the object in the auction.

If you choose to **participate in the Slider Task**, you will earn additional ECU based on your score from correctly adjusting sliders to 50, but *cannot* look up the **aftermarket price** before **bidding**. The amount of ECU that you receive per correctly adjusted slider will change over the 16 auctions, and you will be informed about it before you choose whether to participate in the task. Before each auction, the slider rate is chosen independently for each bidder with equal probability from 0, 0.25, 3, or 6 ECU per slider.

The other person in your auction can also either look up the aftermarket price or participate in the Slider Task, but they may earn more, less or the same amount per correctly adjusted slider as you.

At the end of 30 seconds, the bidding will begin. Your **value** will be displayed again. Below, in the auction panel, you will see a price called the **clock price**. The clock price will start at **0 ECU**, and will increase by **1 ECU** with each tick of the clock. The bidding starts immediately, and **your bid is equal to the clock price**. You can stop bidding and leave the auction by clicking the “Stop Bidding” button next to the clock price. Similarly, the other bidder can stop bidding and leave the auction. You remain actively bidding as long as you have not yet clicked “Stop Bidding”.

Neither bidder will observe whether the other bidder has stopped bidding. **If the other bidder leaves the auction while you continue bidding, the market price of the object is the clock price at which the other bidder stopped bidding**, although you continue to see the clock price increasing until you stop bidding. Similarly, if you stop bidding while the other bidder continues, the market price of the object is the clock price at which you stopped bidding.

After you and the other person in your auction both click “Stop Bidding”, you will see the outcome of the auction. The bidder who continued to bid longer and higher in the auction (the high bidder) gets the object and receives **Auction Earnings** equal to the difference between the object’s value and the price of the object; the other bidder receives no Auction Earnings, that is:

- if you are the high bidder:
 - o Your Auction Earnings = (Your **value** – **market price**) ECU

If this difference is negative because the price exceeds the object’s value, then the Auction Earnings of the high bidder are negative, which represents a loss.

- if you are not the high bidder:
 - o Your Auction Earnings = 0 ECU

If you stop bidding at the same price as the other bidder, chance decides who gets the object, and the price is equal to the clock price at which you both stopped bidding. If no one stops bidding before the clock price reaches 105 ECU, then the object will be allocated by chance at a price of 105 ECU.

If you are the high bidder so that you receive the auctioned object, you **do not** participate in the *Aftermarket*. However, if you are not the high bidder, you can buy an **identical object** in the *Aftermarket*.

In the *Aftermarket*, you will be reminded of your **value** for the object (which is exactly equal to your value for the auctioned object) and the **aftermarket price**. If you want to buy the object at the aftermarket price, click the button labeled "Buy". Your Aftermarket Earnings are determined as follows:

- if you choose to buy the object in the aftermarket:
 - o Your Aftermarket Earnings = (your **value** – **aftermarket price**) ECU
- if you choose *not* to buy the object in the aftermarket:
 - o Your Aftermarket Earnings = 0 ECU

At the end of the *Aftermarket*, one period is complete.

Your Stage 2 Earnings

Your **Earnings** for an individual period can be determined in one of the following six ways:

- 1) You receive the object in the auction, and you participate in the Slider Task
Earnings = (your **value** – **market price** + Slider Task Earnings) ECU
- 2) You receive the object in the auction, and you *do not* participate in the Slider Task
Earnings = (your **value** – **market price**) ECU
- 3) You do not receive the object in the auction, you participate in the Slider Task, and you choose to buy an object in the *Aftermarket*.
Earnings = (your **value** – **aftermarket price** + Slider Task Earnings) ECU
- 4) You do not receive the object in the auction, you *do not* participate in the Slider Task, and you choose to buy an object in the *Aftermarket*.
Earnings = (your **value** – **aftermarket price**) ECU
- 5) You do not receive the object in the auction, you participate in the Slider Task, and you choose *not* to buy an object in the *Aftermarket*.
Earnings = Slider Task Earnings ECU
- 6) You do not receive the object at auction, you *do not* participate in the Slider Task, and you choose *not* to buy an object in the *Aftermarket*.
Earnings = **0** ECU

Total Earnings

Now we will do 2 practice auctions, to familiarize you with the interface. These will not count for payment. After the practice, we will restart the experiment. At the end of the experiment, we will randomly draw **1 out of the 16 periods** for payment. If your stage 2 earnings are positive, then your ECU balance increases by this amount. If your earnings are negative, you made a loss and your ECU balance decreases by the amount of the loss. If your earnings are 0 ECU, then your ECU balance does not change. Finally, we will add your earnings from the *Slider Task* stage (Stage 1) to your ECU balance and we will convert it to Dollars at a rate of 10 ECU = 1 Dollar. When the experiment is over, please wait quietly until we call you for payment.

C Estimates of a QRE model of bidding behavior

To investigate the extent to which noisy behavior can explain observed dispersion in bids and overbidding, we estimated a QRE model of bidding behavior, taking information acquisition decisions as a given. In our QRE model, each agent noisily chooses among strategies, with a greater tendency to choose strategies with higher expected payoffs, given the behavior of other agents. We model a player’s behavior in the aftermarket as governed by a separate and independent agent (as in “Agent-Based QRE”; see section 3.2 in Goeree et al., 2016). Behavior is assumed to follow a logit form, with the parameter α capturing the degree of noisiness. Our model assumes that heterogeneity in behavior, given an agent’s information arises only due to random variation, rather than underlying individual differences – which gives us a one-parameter generalization of the standard theory in section 2.¹⁷

In the aftermarket, an agent with valuation v who faces aftermarket price q buys the object with probability

$$p(v, q) = \frac{\exp(\alpha(v - q))}{\exp(\alpha(v - q)) + \exp(0)} \quad (10)$$

and otherwise does not buy.

Let $i \in \{50, \dots, 100\} \times (\{0, \dots, 50\} \cup \{\text{uninformed}\})$ denote a player’s information, which consists of her valuation of the object and either knowing the aftermarket price or being uninformed. Let ρ denote the fraction of the time that each agent becomes informed, which we take as exogenous for the purpose of this exercise.¹⁸ Let σ', p' denote the stochastic response functions in the auction and aftermarket played by other agents, and let $U(b|i)(\sigma', p')$ denote the expected utility of bidding b given the behavior of other agents. Then, assume a player’s probability of making bid b given information i is given by:

$$\sigma(b, i) = \frac{\exp(\alpha U(b, i)(\sigma', p'))}{\sum_{b'=0}^{105} \exp(\alpha U(b', i)(\sigma', p'))} \quad (11)$$

The value of $U(b|i)(\sigma', p')$ is given by:

¹⁷An extension of the model that allows α to vary between informed and uninformed agents yielded very similar estimates of α for both.

¹⁸Note that, due to the logit functional form, this model predicts that earnings from the slider task will not affect subsequent behavior in the auction and aftermarket.

$$U(b|(v, q))(\sigma', p') =$$

$$\begin{aligned} & \rho \left(\frac{1}{51} \sum_{v'=50}^{100} \left[\sum_{b'=0}^{b-1} \sigma'(b'|v', q) (v - b') + \frac{1}{2} \sigma'(b|(v', q)) (v - b) \right. \right. \\ & \left. \left. + \left(\frac{1}{2} \sigma'(b|(q, v')) + \sum_{b'=b+1}^{106} \sigma'(b'|v', q) \right) p'(v, q) (v - q) \right] \right) \\ & + (1 - \rho) \left(\sum_{b'=0}^{b-1} \sigma'(b'|v', \text{uninformed}) (v - b') + \frac{1}{2} \sigma'(b|(v', \text{uninformed})) (v - b) \right. \\ & \left. + \left(\frac{1}{2} \sigma'(b|(q, v')) + \sum_{b'=b+1}^{106} \sigma'(b'|v', \text{uninformed}) \right) p'(v, q) (v - q) \right) \end{aligned} \quad (12)$$

$$U(b|(v, \text{uninformed}))(\sigma', p') =$$

$$\begin{aligned} & \rho \left(\frac{1}{51} \sum_{q=0}^{50} \frac{1}{51} \sum_{v'=50}^{100} \left[\sum_{b'=0}^{b-1} \sigma'(b'|v', q) (v - b') + \frac{1}{2} \sigma'(b|(v', q)) (v - b) \right. \right. \\ & \left. \left. + \left(\frac{1}{2} \sigma'(b|(q, v')) + \sum_{b'=b+1}^{106} \sigma'(b'|v', q) \right) p'(v, q) (v - q) \right] \right) \\ & + (1 - \rho) \left(\sum_{b'=0}^{b-1} \sigma'(b'|v', \text{uninformed}) (v - b') + \frac{1}{2} \sigma'(b|(v', \text{uninformed})) (v - b) \right. \\ & \left. + \left(\frac{1}{2} \sigma'(b|(q, v')) + \sum_{b'=b+1}^{106} \sigma'(b'|v', \text{uninformed}) \right) \frac{1}{51} \sum_{q=0}^{50} p'(v, q) (v - q) \right) \end{aligned} \quad (13)$$

A strategy σ^*, p^* is an equilibrium when p^* is given by (10), and given $U(b, i)(\sigma^*, p^*)$ computed according to (12 and 13), σ^* is consistent with (11).

We constructed a program in R (2015) to find the QRE strategies σ^α and p^α corresponding to α as follows. By Proposition 3.1 in Goeree et al. (2016), a QRE exists in our problem. Our program starts with random behavior, and then iteratively solves for the response functions using (2) until strategies have converged. Notice that choice probabilities in this model are continuous in utility and in α , and utility is in turn continuous in the strategy of other agents. As a result, the program will find strategies arbitrarily close to a QRE.¹⁹

¹⁹We have not proven that the QRE is unique; by starting this process from uniform random behavior,

QRE model fit	
$\hat{\alpha}$.46
s.e.	.02

TABLE C1: Estimated α from the QRE Model

Given σ^α , we can construct the contribution of a given bid to the likelihood function:

$$ll(b, i|\alpha) = \sigma^\alpha(b, i)$$

and given p^α , the contribution of an aftermarket decision to the likelihood function is given by:

$$ll(\text{buy}, v, q|\alpha) = p^\alpha(v, q)$$

$$ll(\text{don't buy}, v, q|\alpha) = 1 - p^\alpha(v, q)$$

We estimate α by maximum likelihood using the `mle` command.²⁰ Our results are reported in Table C1.

Table C2 compares the bidding predictions generated by our estimated QRE model against the core features of our bidding data that were reported in our results section. The estimates suggest that QRE can help account for some of our observations such as excess bid dispersion and on-average overbidding, especially among the uninformed. However, in some ways QRE predicts “too much” noise and fails to capture the relatively close adherence to dominant strategy bidding by the informed. We summarize the evidence below.

Notice that overbidding is substantially more pronounced by uninformed bidders relative to informed bidders, and uninformed bidders also appear noisier relative to the PBE predictions: informed bidders are substantially more likely to bid close to the PBE point prediction. In the model, this can only arise out of weaker payoff incentives for uninformed than for informed bidders to adhere to the PBE strategy. It appears that overbidding our approach will select the “logit solution”, which is common practice in estimating the QRE model (see Goeree et al., 2016, p. 152).

²⁰Note that the standard errors displayed assume that each observation (i.e. each bid and each aftermarket decision) is independent. Since we are not primarily interested in inference about α , we do not make corrections to account for interdependencies among observations.

is driven by the conjunction of noisy behavior and a strategy space that always contains more bids above the PBE point prediction than below. Since noise in the QRE model generates a positive probability of making each possible bid, it tends to produce more overbidding than underbidding both on average and at the median; since uninformed bidders have weaker incentives and make more dispersed bids, they tend to overbid more than informed bidders; see Figure C1.

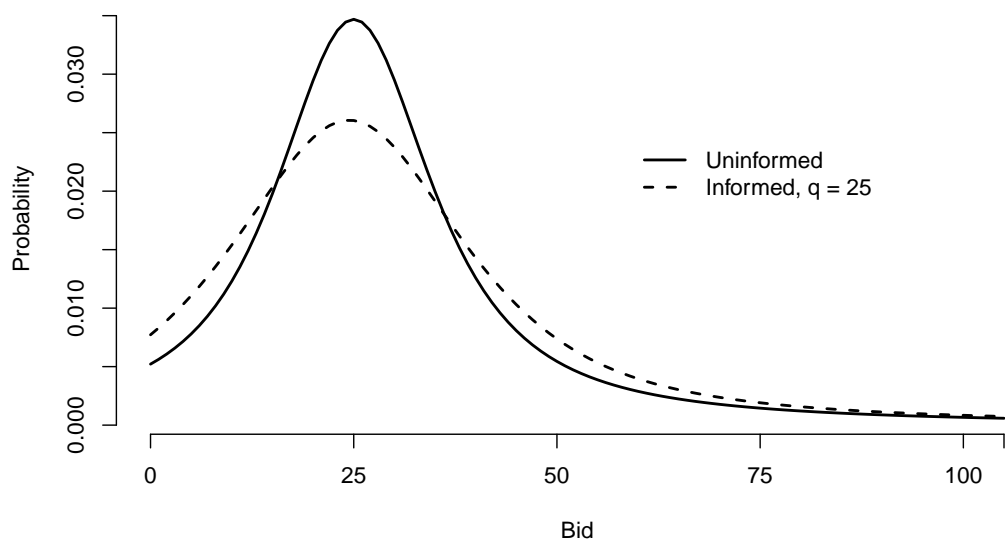


FIGURE C1: The distribution of bids by informed and uninformed bidders under QRE

The model does not capture the very-close adherence to the dominant strategy by a large fraction of informed bids. Instead it predicts noisy behavior more continuously about those predictions. One possible explanation is that bidding around the aftermarket price is salient and/or cognitively easy to identify for some subjects when informed.

We also note that the estimated model actually underestimates the frequency at which subjects fail to buy when in the aftermarket, predicting that players who reach the aftermarket only fail to buy 0.2% of the time (while players fail to buy 5.7% of the time in our data). This suggests that, in the estimated model, the anticipation of potentially failing to buy in the aftermarket does not drive much overbidding (nor differential overbidding).

Deviation of bids from point predictions	Uninformed	Informed
+1	7.8% (11.7%)	11.8% (30.8%)
+5	27.8% (30.4%)	38.7% (54.7%)
+10	49.4% (51.2%)	61.1% (70.2%)
median	27 (30)	26 (25)
average	30.2 (32.4)	29.9 (27.0)

TABLE C2: Model-predictions of the % of bids for a given information type that fit within specified bounds of the PBE point predictions; corresponding figures from data in brackets

Working Paper Series in Economics

recent issues

- No. 101** *David Freeman, Erik O. Kimbrough, J. Philipp Reiss*: Opportunity cost, inattention and the bidder's curse, February 2017
- No. 100** *Andranik S. Tangian*: Selection of questions for VAAs and the VAA-based elections, January 2017
- No. 99** *Dominik Rothenhäusler, Nikolaus Schweizer and Nora Szech*: Guilt in Voting and Public Good Games, November 2016
- No. 98** *Eckhardt Bode, Stephan Brunow, Ingrid Ott and Alina Sorgner*: Worker personality: Another skill bias beyond education in the digital age, November 2016
- No. 97** *Clemens Puppe*: The single-peaked domain revisited: A simple global characterization, November 2016
- No. 96** *David Burka, Clemens Puppe, Laszlo Szepesvary and Attila Tasnadi*: Neural networks would 'vote' according to Borda's rule, November 2016
- No. 95** *Vladimir Korzinov and Ivan Savin*: Pervasive enough? General purpose technologies as an emergent property, November 2016
- No. 94** *Francesco D'Acunto, Daniel Hoang and Michael Weber*: The effect of unconventional fiscal policy on consumption expenditure, October 2016
- No. 93** *Andranik S. Tangian*: The third vote experiment: VAA-based election to enhance policy representation of the KIT student parliament, September 2016
- No. 92** *Clemens Puppe and Arkadii Slinko*: Condorcet domains, median graphs and the single-crossing property, June 2016
- No. 91** *Markus Höchstötter, Mher Safarian, Anna Krumetsadik*: Analysis of stochastic technical trading algorithms, June 2016