VATT Working Papers 44

Effects of Minimum Bid Increments in Internet Auctions: Evidence from a Field Experiment

Janne Tukiainen

VATT INSTITUTE FOR ECONOMIC RESEARCH

VATT WORKING PAPERS

44

Effects of Minimum Bid Increments in Internet Auctions: Evidence from a Field Experiment

Janne Tukiainen

Valtion taloudellinen tutkimuskeskus
VATT Institute for Economic Research
Helsinki 2015

Janne Tukiainen, Valtion taloudellinen tutkimuskeskus ja Helsinki Center of Economic Research, janne.tukiainen@vatt.fi, tel. +358295519451

I thank Estelle Cantillon, Phil Haile, Brent Hickman, Tim Hubbard, Ari Hyytinen, Klaus Kultti, Laurent Lamy, Emmi Martikainen, Harry Paarsch, Heikki Pursiainen, Matti Sarvimäki, Tanja Saxell, Otto Toivanen, Alexander White, Hanna Willman-Iivarinen and two anonymous referees for their comments, Niklas Jahnsson for the excellent research assistance and Toni Ruuska at Huuto.net for providing industry expertise. Conference and seminar participants at Helsinki, EARIE, FEAAM, Toulouse, London, Paris and Copenhagen provided useful discussions. I am grateful to the Yrjö Jahnsson Foundation (YJF) for funding this experiment and to the Emil Aaltonen Foundation for funding in general. I thank STICERD at the LSE for their hospitality and the Jenny and Antti Wihuri Foundation and YJF for funding that.

This version (September 17, 2015) replaces the February 15, 2013 version.

ISBN 978-952-274-057-1 (PDF)

ISSN 1798-0291 (PDF)

Valtion taloudellinen tutkimuskeskus
VATT Institute for Economic Research
Arkadiankatu 7, 00100 Helsinki, Finland

Helsinki, October 2015

Effects of Minimum Bid Increments in Internet Auctions: Evidence from a Field Experiment

VATT Institute for Economic Research
VATT Working Papers 44/2013

Janne Tukiainen

Abstract

I study the role of minimum bid increments (MBI) in internet auctions using natural field experiment data from a Finnish online auction site Huuto.net. Internet auctions are typically viewed as secondprice, implying truthful bidding. However, due to the presence of the MBI, equilibrium bidding behaviour involves bid shading. I sell otherwise identical gift cards while varying the MBI and test between truthful bidding and bid shading. Truthful bidding is rejected. Consistently with equilibrium behavior, bid shading increases as the MBI increases. This implies that real-world bidders are capable of reacting rationally even to seemingly small variations in their strategic environment.

Key words: Field experiment; internet auctions; minimum bid increment; revenue; strategic bidding

JEL classes: C52, C72, C93, D44, L81

Effects of Minimum Bid Increments in Internet Auctions: Evidence from a Field Experiment

Janne Tukiainen

VATT Institute for Economic Research

Arkadiankatu 7, P.O.Box 1279, FI-00101 Helsinki, Finland.

Email: janne.tukiainen@vatt.fi; Tel. +358295519451.

September 17, 2015

Abstract

I study the role of minimum bid increments (MBI) in internet auctions using natural field experiment data from a Finnish online auction site Huuto.net. Internet auctions are typically viewed as second-price, implying truthful bidding. However, due to the presence of the MBI, equilibrium bidding behavior involves bid shading. I sell otherwise identical gift cards while varying the MBI and test between truthful bidding and bid shading. Truthful bidding is rejected. Consistently with equilibrium behavior, bid shading increases as the MBI increases. This implies that real-world bidders are capable of reacting rationally even to seemingly small variations in their strategic environment.

Keywords: Field experiment; internet auctions; minimum bid increment; revenue; strategic bidding.
JEL Classification: C52; C72; C93; D44; L81.

1 Introduction

Internet auctions (also referred to as online or electronic auctions) account for a large and increasing share of C2C trading, retail trade (B2C) and B2B trading. The defining feature of most internet auctions is the so called proxy-bidding, where bidders reveal to the auction server the maximum bid that the server can submit on their behalf. The server then acts as a proxy and increases the current bid in response to competitive bids up to the submitted maximum, but where possible, to equal the second-highest bid plus a minimum bid increment (MBI). The MBI also defines the minimum amount by which a new bid must exceed the current price for it to be accepted. I argue that in internet auctions, the MBI is an important but empirically largely overlooked feature of the auction design. The received theory (Bapna et al. 2003, Hickman 2010 and Rogers et al. 2007) predicts that the MBI may be an important determinant of internet auction bidding strategies and revenue, and that the revenue maximizing MBI may be higher than zero. I use a novel natural field experiment conducted in a Finnish online auction site Huuto.net, where a seller can set the MBI freely, to show that these theoretical predictions are supported by the data. I document that increasing the MBI increases bid-shading and also seller revenue, but only up to a point limited by an entry deterring effect.

In internet auctions, the MBI affects not only whether a new bid is accepted, but also how the price is determined. Most importantly, the larger the MBI the larger the likelihood is that the pricing rule is based on the first-price rather than the second-price mechanism. If the highest and the second-highest proxy bids are within one MBI of each other, the first-price rule is used, because the current price cannot exceed the highest bid. Because the MBI adjusts the probability that the pricing rule in a given auction is either a second-price or a first-price rule, it should have a significant impact on bidder behavior. Accordingly, Hickman (2010) shows that in equilibrium, bid functions converge towards first-price auction bid functions as the MBI increases. The main contribution of my paper is that I show by applying a test to my field experiment data that bidders do indeed account rationally for the MBI in their bidding strategies. It is generally thought that the equilibrium in internet auctions is equivalent to the equilibrium in sealed bid second-price auctions (e.g. Bajari and Hortacsu 2003). The main implication of the results is that internet auctions should not be modeled as second-price auctions as most of the literature does. To my knowledge, this is the first field experimental or causal evidence on the topic in the literature.[1]

The existing gap in both the theoretical and empirical literature concerning the analysis of MBIs could be due to the seller not being able to set the MBI on large internet auction sites like eBay. Another explanation is that the MBI has been regarded only as a minor detail of these auctions, since they are typically small relative to the value of the objects on sale. While their small relative size limits the possible revenue effects, even relatively small MBIs may have significant implications for the analysis of these auctions. In particular, Hickman et al. (2011) show that not accounting for the MBI in estimation leads to statistically significantly different structural estimates compared to using a correctly specified model, and the differences are large in magnitude. In particular, the estimated valuation distributions differ substantially even with relatively small MBIs. Besides the analysis of online auctions, my experimental results contribute more generally to the field experimental evidence on the rationality and strategic behavior of real economic agents in real auction settings (e.g. Reiley 2006, List and Lucking-Reiley 2000). From this perspective, the possibility that MBIs may be perceived only as a minor detail of the auction design makes my experimental results more compelling. In particular, I show that real-world bidders are capable of reacting according to the economic theory even to seemingly small variations in their strategic environment in ways that can be observed in the data.

[1] Zeithammer and Adams (2010) provide novel model validation tests and empirical evidence that the second-price sealed bid assumption is not the correct one in eBay auctions. They present evidence that while some bidders may submit sealed bids, a significant number of bidders conduct incremental bidding, where the current price is raised by one MBI at a time repeatedly up to a valuation. Although they utilize information on MBIs in their tests, they overlook the implications of MBIs on entry and strategic behavior. With incremental bidding, the top two bids in these auctions no longer reflect the true bidder valuations. Accordingly, Hortacsu and Nielsen (2010) argue for the importance of incorporating richer models in structural empirical analysis of online auctions. Hickman et al. (2011) indeed provide a structural estimation technique that accounts for MBI in a setting with independent private values and exogenous entry.

2 The experiment

Most internet auction sites, including Huuto.net, are a variant of an ascending auction with a minimum bid increment and a proxy-bidding system. Proxy refers to the algorithm that will conduct the bidding in the place of the real person. Submitting only one bid to the proxy is called proxy-bidding. Huuto.net advises buyers to submit only their true valuation once and let the proxy do the rest, which would result in a sealed-bid second-price mechanism absent the role of MBI. To understand how the MBI may affect strategies and revenue, we need to understand how the proxy-bidding system works. For a bid to be accepted, it must exceed the current bid (or the reservation price for the first bidder) plus the MBI. This is the minimum amount that the proxy accepts. It is also possible to set any value higher than this to the proxy. The value entered does not need to follow the grid imposed by the MBI. This implies that even if sellers for some reason prefer even numbers, the MBI cannot be used to achieve them. Whenever a new bidder informs his proxy of his valuation, the current bid immediately advances to the minimum of the highest price entered so far and the second-highest price plus the MBI. This current price formula (1) also implies that the current price and therefore the selling price need not be restricted to the integer multiples of the MBI.

$$CP_t = min(HB_t, SHB_t + MBI), \qquad (1)$$

where CP_t denotes the current price at time t and therefore also the revenue for the auctioneer at time $t = T$, where T denotes the closing time. HB_t is the highest bid submitted in the auction and SHB_t the second-highest bid at time t. A higher MBI level implies that the current price will more often be determined as the highest bid submitted to the proxy. If the MBI is zero, the price is determined as in a second-price sealed-bid auction. If the MBI is very high, then the price is determined as in a first-price sealed-bid auction. With intermediate levels of MBI, an internet auction is a hybrid between these two auction formats.

In Huuto.net, the seller has some control over the auction mechanism. The seller has to set some starting price, which is equivalent to setting a public reservation price. It is also possible to set a secret reservation price. In addition to these parameters, the site offer sellers a wide variety of marketing options. Most importantly, Huuto.net is ideal for testing theory related to MBIs, because on Huuto.net, it is possible for the sellers to choose the MBI level. Huuto.net sellers typically set the MBI higher than the smallest possible level. MBIs of 0.5, 1 and 2 euro are the most common and very small MBIs are rare.

I conduct a field experiment to study how the MBI affects bidding behavior. As a secondary objective, I also look at how it affects the selling price and the number of observed bidder identities. In this experiment, I sell Stockmann gift cards. Stockmann is the largest department store chain in Finland. The gift cards are valid in all of Stockmann's seven large department stores in Finland. These seven stores are located in the six largest cities in Finland that together have 1.6 million inhabitants or about 30% of the entire population. These stores sell millions of different commodities and services. These cards can also be spent in the company's subsidiary stores like Seppälä, which sells clothes in 90 Finnish cities and towns. Seppälä is thus easily accessible to most of the Finnish population. These gift cards were chosen for this experiment mainly because there is a large demand for them. It is very likely that the potential demand for the last card sold is about the same as for the first card sold. However, this assumption is not necessary for the internal validity of the experiment; the object was chosen simply to guarantee that each auction generates bid data.

In the first experiment, I sell a total of 72 identical 15 euro Stockmann gift cards. 24 are sold at a 1 cent MBI. This constitutes the control group. There are two treatment groups, the 33 cent MBI (corresponding to the eBay level, see Table A1 in the Appendix) and the 50 cent MBI. The second experiment is run in exactly the same way, with the exception that the cards now have a nominal value of 50 euro, and the MBI levels are now 1 cent, 66 cent (corresponding to the eBay level) and 100 cent. It would have been interesting to have one treatment in the experiment at the optimal MBI, but unfortunately this was not possible. To be able to calculate the optimal MBI for the objects sold in the experiment, I would need to know many unobservable factors, such as the number of potential bidders and their value distributions, the nature of the entry process and bidder strategies, quite apart from the fact that the relevant theory for this calculus does not exist yet either. Therefore, absent a better benchmark, it seems that a natural set-up is to look at the MBI level of the largest internet auction site eBay and compare it with the minimum level and a higher level.

The experiment is conducted with objects that have both a reservation price (the starting price plus the MBI) and a maximum selling price (the nominal value of the card) within range where eBay would have kept the MBI the same throughout the auction. For the 15 euro cards, the starting price level was 8 euro (about $12 at the time of the experiment) in the experiment and the maximum value of the object to buyers was 15 euro (about $22.5). In practice, the 1 cent MBI cards have a 7.99 euro starting price level, the 33 cent MBI cards 7.67 euro and the 50 cent MBI cards 7.5 euro. Since the first accepted bid is the starting price level plus the MBI, all the cards have a limit of 8 euro for the first bid to be accepted. In the 50 euro card experiment, I followed the same logic and set the level of acceptance of the first bid at 30 euro. Within these price ranges, the MBI level for the 15 euro cards would have been $0.50 (about 33 cents given the exchange rates around the experiment date) and $1 (about 66 cents) for the 50 euro cards. These reserve prices are low enough compared to the nominal values of the cards not to be binding. Thus, none of the forthcoming analysis needs to address the reservation prices. They are simply imposed to make comparison with eBay MBI levels possible, that is, to keep the bidding always within a single eBay MBI bracket.

I sell all the cards in separate auctions that each last 5 days, from Tuesday afternoon to Sunday afternoon. 12 cards are auctioned at the same time. Both the experiments last six weeks and constitute each of six 12-card batches (clusters) with each batch including 4 cards of each of the different MBIs. I randomize the order in which the cards are placed on the auction within each batch. Note that this set-up creates balanced variation in the treatment within each cluster. Therefore, even if the auctions within each cluster are not independent, for example due to the presence of decreasing average transaction costs per card bought or substitution between the cards, the experiment achieves good power with relatively few observations, when compared to randomization at the cluster level only. Typical issues impairing experiments with within-cluster variation include risks of contamination, and ethical, political, administrative or financial reasons (e.g. Moerbeek 2005). None of these problems are present in my experiment. Moreover, my reputation as a seller increased during the experiment, but with within-cluster variation in the treatment this causes no problems, since this potential effect is easily controlled by adding weekly (cluster) fixed effects to the regressions. These fixed effects also control for any other unobserved weekly changes.

This experiment is what List (2011), for example, calls a natural field experiment. In these experiments, real economic agents operate in the real and same environment where they would usually operate, without knowing that they are participating in an experiment. Therefore the observed treatment effect cannot be a result of subjects reacting to simply the fact that they are in an experiment and being observed by researchers. Typically, natural field experiments are very reliable both in the sense of estimating internally valid causal effect and being broadly generalizable (Al-Ubaydi and List 2013).

One property that gift cards have is that bidders are likely to have private valuations for them, since there is probably no significant common uncertainty about the value of the objects. Since everyone knows the nominal value of the cards, the uncertainty is related more to individual bidder characteristics such as how close the other bidders live to a Stockmann store. What this means is that bidders would not update their valuations if they learned their competitors' valuations or signals. Here, common values are not to be confused with the likely possibility that overall, the uncertainty over the value of gift cards is likely to be small. However, at least most of the existing uncertainty regarding valuations is essentially private. For example, bidders may have different transaction costs, may discount the future at different rates and have different uses for the cards. They would not update their valuations if they learned their competitors' transaction costs, for example.[2] On the other hand, the perceived trustworthiness of the seller, for example, may be a potential common value component. Moreover, reselling the cards is possible, which would also imply common values. However, the difference between selling prices and nominal values will turn out to be quite small, and therefore the potential gains from reselling are likely to be lower than the transaction costs.

[2] Because at least some of these private value components are likely to be present, the appropriate model is also very unlikely to be a one of pure common values without any uncertainty. If this were the case, all bidders would submit the true value of the object in both the first and second price cases. The data does not support such behavior either, because because bidder fixed effects turn out to be significant predictors.

It is important to note that the information regime is not relevant for the internal validity of the experiment, although it may play a role in interpreting the results in light of the received theory and in evaluating the external validity. For example, with independent private values, it could be that a higher MBI would lead to higher revenue gains than under some other information assumptions such as the affiliated values framework.

From the perspective of testing for truthful bidding, Huuto.net has also other attractive properties than being able to set the MBI. First, Huuto.net applies a soft stopping rule. The soft closing rule on Huuto.net means that if a bid is submitted when the auction is about to close in less than 5 minutes, 5 minutes are added to the time that the auction is open. This implies very limited incentives for sniping (Ockenfels and Roth 2006). Second, Huuto.net also gives preferential treatment to winners who submit their bid to the proxy earlier than the second-highest bidder. In that case, the winner pays only the second-highest bid, but not the MBI. Therefore, in the case of early-placed winning bids, Huuto.net auctions are really just second-price auctions. Both the soft stopping rule and this preferential treatment implies that second-price sealed bid abstraction would seem to be natural in Huuto.net, which also makes it an ideal market from the perspective of conservatively testing for truthful bidding. Another but less ideal feature of Huuto.net is that the bidders are allowed to jump bid. If a jump bidder wins, she pays her bid. Because jump bids are likely to involve bid shading, auctions won by jump bids are omitted from the truthful bidding testing procedure. Moreover, jump bidding wins are equally common in all the different MBI treatment groups used in the experiment (see Table 6), which implies that the results concerning the MBI should not be much affected by the possibility of jump bidding.

On the one hand, these unique features limit the external validity of some of the results, for example how the results apply to larger sites such as eBay, which differ to some extent from Huuto.net in the exact auction mechanism that they use. On the other hand, it is important to point out that these differences are not likely to hurt the external validity of my results concerning strategic bidding as opposed to truthful bidding. If bidders who are normal people are strategic and rational on one site, there is no reason to suspect that similar bidders would not be rational on other sites, even if the equilibrium bidding strategies may differ due to the different mechanisms. However, these differences may question the external validity of the revenue results here. Therefore, I analyze the external validity of especially the revenue results in the Appendix C, for example by using different kinds of subsample analyses.

3 Results

3.1 Revenue and entry effects

One contribution of this paper is the analysis of the revenue effects of MBIs. Because the MBI affects the probability that the pricing rule in a given auction is either a second-price or a first-price rule, it could have a significant impact on revenue in situations where the revenue equivalence theorem (Myerson 1981 and Riley and Samuelsson 1981) does not hold (e.g., risk-averse (Maskin and Riley 1984), budget-constrained (Che and Gale 1998) or asymmetric bidders (Krishna 2002) or bidders with affiliated valuations (Milgrom and Weber 1982)). In particular, the revenue maximizing MBI may well be higher than zero in cases where first-price auctions generate more revenue than second-price auctions.[3] However, since these results on revenue and

[3] Bapna et al. (2003) provide the only other empirical attempt to study the association between the MBI and seller revenue. They find it to be important both theoretically and empirically when using observational data from B2C online auctions. However, Bapna et al. (2003) do not address the potential problems regarding unobserved heterogeneity. Moreover, their sampling procedure may introduce a selection bias since they exclude auctions with low participation from their data. Because low participation may be caused by a high MBI, this selection may result in overestimating the effect of the MBI on revenue. Rogers et al. (2007) also analyze the MBIs from a theoretical perspective. According to their model and simulations, the seller can maximize revenue by setting the reservation price to zero and the MBI to an optimal level that is higher than zero. According to these studies, the MBI seems to be an even more important mechanism parameter than setting the reservation price optimally, which has been at the center of auction research ever since Myerson's (1981) seminal contribution. However, these results hinge on behavioral assumptions regarding bidder behavior. In other words, they are not equilibrium results. The results in Rogers et al. (2007), for example, require that all bidders submit a single truthful bid.

entry do not directly test theory, they should be seen more as stylized facts, despite being causal effects identified by the experiment.

The main questions of interest in this section are the effects of the MBI on the selling price and the number of actual bidders. I observe the selling price and bidding history in the data. The bidding history is limited to the current price and the current highest bidder identity after each accepted new bid. Therefore, it does not reveal the actual bids submitted to the proxy-bidding system. Thus, I cannot use all the submitted bids in the analysis but rather only the winning bids. Nor does the history reveal the true number of actual bidders, since a new entrant may place an accepted bid, but I would only observe that particular new bidder identity if her bid was the highest at the time. Thus, the variable for the number of observed bidder identities, which I call the number of bidders, is a lower bound or a downwards-biased proxy for the number of actual bidders. However, there is little reason to assume that the measurement error in this variable is correlated with the MBI. Therefore, even if the observed levels may be too low, regressing the number of bidder identities on the MBI treatments should identify the treatment effect of the MBI on entry correctly.

In Table 1, I describe the variables of interest in the first experiment. The mean price is 13.24 euros. The average price is lowest for the 1 cent MBI and highest for the 33 cent MBI. All the auctions receive bids from at least two different bidders. The average number of bids is lower the higher the MBI is. There is enough variation in all the response variables to warrant meaningful regression analysis.

[Table 1 about here]

In Table 2, I describe the variables of interest in the second experiment. The mean price is 44.46 euros. The conditional descriptive statistic patterns look exactly as in the first experiment. The average price is lowest for the 1 cent MBI and highest for the 66 cent MBI. All the auctions receive bids from at least two different bidders and the average number of bidders is lower the higher the MBI is. Again there is much variation in all the variables.

[Table 2 about here]

The objects sold are identical in all but two dimensions. Since I randomize the order in which the cards are placed on the auction within each patch, only the week that they are entered in the auction and the MBI differ systematically. Therefore I regress the variables of interest on the week fixed effects and the treatment dummies. Another limitation of my experiment compared to an ideal experiment is that I cannot randomize over the participating bidders. Unfortunately, it turned out that only 13 different bidders participated and a single bidder won about half of the auctions. This is potentially problematic both for the internal and external validity of the results due to possible selection of the potential and actual bidders. I will deal with this issue in the analysis. Firstly, I add the winning-bidder fixed effects in addition to the week fixed effects.[4] Secondly, I repeat the analysis using a leave-one-out strategy, where I drop all the auctions won by a given bidder from the data, one bidder at the time (see Table C3 in the Appendix C). Appendix C provides a detailed analysis of bidder participation (see Tables C1 and C2).

Since the week fixed and winning bidder fixed effects may not remove all correlations between the error terms, for example if different bidders select into different weeks and different MBI levels (although they do not seem to based on Table C2 in the Appendix C), I also report the results clustered either at the week level or the winning-bidder level or both simultaneously. When I analyze the 15 and 50 euro experiment data separately, I do not report clustered standard errors due to very small number of clusters. When I pool the data, I report both one-way clustering at week or winning bidder level and two-way clustering on both levels simultaneously. However, the number of clusters may be too small for the standard clustering methods to work properly. Therefore, I have checked that clustering using a wild bootstrap procedure (Cameron et al. 2008) results in even smaller p-values than those reported in Table 6.

The regression results for the 15 euro experiment are presented in Table 3. The auctions with the 1 cent MBI attracted more bidders than the auctions with a higher MBI, but the difference is not statistically significant in most cases. Both the eBay level and the highest MBI give the seller more revenue than the

[4] Using random effects instead of fixed effects does not change the results substantially.

1 cent MBI and the difference is statistically significant at the 5% level in the specifications with control variables included and thus smaller residual variance. There is no statistically significant difference between the two higher-level treatments. The revenue point estimate is stable to adding controls, which suggests that the randomization works as intended.

[Table 3 about here]

The results for the 50 euro experiment are presented in Table 4. The auctions with a 1 cent MBI receive more bids than the auctions with a higher MBI. The difference is typically statistically significant, especially for the highest MBI. The eBay level gives the seller more revenue than the 1 cent MBI, but the difference is statistically significant only at the 10% level and only in the richest specification. There is a statistically significant difference between the two higher-level treatments in one specification. The results are again very stable to adding controls.

[Table 4 about here]

In order to get more power out of the experiments, I pool the data. To run the pooled regressions on price, I construct a new variable called "discount". This is calculated by dividing the selling price by the nominal value of the card. It is very interesting to note that this discount is very much the same in both the experiments. The average value of the "discount" is 0.88 in the 15 euro and 0.89 in the 50 euro experiment. One interpretation of this evidence is that transaction costs are not very important, since a fixed transaction cost should make the average of the "discount" variable higher for the low-value than for the high-value cards. It also makes sense now to run a series of pooled regressions, since especially the mean but to a lesser extent the variance (0.013 for 15 euro and 0.026 for 50 euro data) of the response variable are about the same in both the experiments. In the pooled regressions, I treat both mid-level MBIs (33 cents for the 15 euro experiment and 66 cents for the 50 euro experiment) as one single treatment dummy and both high-level MBIs (50 cents for the 15 euro experiment and 100 cents for the 50 euro experiment) as another treatment dummy. If e.g. 33 cents in 15 euro experiment is not the same thing as 66 cents in the 50 euro experiment, this could result in attenuation bias due to measurement error. Given significant results, this does not impose an issue for qualitative interpretations of the results.

According to the pooled results in Table 5, higher MBI levels decrease entry. Moreover, I find a stable revenue effect both in terms of statistical significance and the level, again suggesting valid randomization. Using the eBay level increments increases seller revenue compared to using the 1 cent MBI by an average of 0.84% of the nominal value (or by 1% of the average selling price) in the richest specification. This result is statistically significant at least at the 5% level, except in specification (1), where I do not yet control for anything. In summary, my experiment reveals that the MBI is a relevant determinant of online auction revenues, both in statistical and economic terms. Because in most Internet auctions sellers cannot set the MBI, these revenue results may not be of much interest to a typical real-world seller. However, they could be of some interest to some auctions sites contemplating their designs. The revenue effects estimated here are joint effects of MBIs in a sense that I cannot quantify the role of possible different causal channels. I do not have independent variation to separate between e.g. the negative revenue effects of higher MBI due to less entry, the direct positive revenue effects caused by higher MBI as in Eq. (1) and the effects due to strategic reactions to MBI.

Being a randomized trial with identical objects, the internal validity of these results is very strong. However, the external validity may be less strong for three reasons: The results may not generalize in respect of the characteristics of the bidders, the objects sold and the auction site. I discuss these external validity issues extensively in the Appendix C and conclude that the results may be also of wider relevance.

[Table 5 about here]

We learn three important things from these revenue and entry results from the perspective of understanding which theories are able to explain the evidence. First, the results are not consistent with incremental bidding. Second, the results are consistent with truthful bidding with sequential entry. Third, the results are also consistent with bid-shading under situations where the first-price auction revenues are larger than second-price auction revenues.

Rogers et al. (2007) show that under incremental bidding, the revenue should be decreasing in MBI. Therefore, the result that a low level MBI does worse in terms of revenue than higher levels, is not consistent with incremental bidding. Moreover, the incremental model advocated by Zeithammer and Adams (2010) for eBay cannot explain why sellers on Huuto.net often choose a high MBI. Neither does the bid history that I observe support incremental bidding as the typical behavior, because it is rare to observe the same bidder identity many times in the bid history of a single object.[5] Ruling out incremental bidding is relevant later on when I construct a test for bid-shading. Zeithammer and Adams (2010) show that the MBI together with incremental bidding imply that the first and second bids may be downward-biased from true first and second valuations. Therefore, telling bid-shading apart from incremental bidding is challenging empirically.

The formula for the current price (Eq. 1) reveals immediately how the expected current price and therefore the expected selling price is strictly increasing in MBI assuming truthful bidding. With a sufficiently high MBI, the pricing rule is always first-price, and with truthful proxy-bidding, the auctioneer can collect all the rents. Therefore the analysis of revenue under truthful proxy-bidding is interesting only with an additional feature, namely entry. Rogers et al. (2007) introduce sequential entry to the model. In their model, bidders come to the auction site randomly and sequentially. The higher the MBI, the higher the chance is that a new arrival is not willing to submit any bid, even if he has the highest valuation, because the current price plus the MBI is more likely to be higher than his valuation when the MBI is higher. In this case, the highest-value bidder is deterred and revenue (as well as efficiency) is lost. With truthful proxy-bidding and sequential entry, the expected revenue is higher than the second-highest valuation but by an amount less than the MBI. In this case, there is an optimal value for the MBI that is higher than zero. The entry results

[5] The average of (observed bids/observed bidder identities) over all the auctions is 1.4, which is way too small for incremental bidding to be an often applied strategy. Moreover, even if only a single bid is submitted by one bidder, many bids from her may appear in the bidding history (see previous discussion on the entry data).

reveal that potential efficiency and revenue loss due to entry is a relevant concern for the auctioneer and the revenue results alone cannot reject truthful bidding.

In the strategic bidding model, increasing the MBI may increase expected seller revenue in the cases where the first-price auction revenue is known to dominate the second-price auction revenue, such as the cases of risk-averse or budget constrained bidders, because by increasing the MBI the seller increases the probability that an internet auction behaves as a first-price auction instead of a second-price auction. Thus, the revenue results are also consistent with strategic bid shading.

3.2 Testing truthful bidding

In this section, I test for truthful bidding. Moreover, I analyze whether the results are consistent with bid shading increasing in MBI as the equilibrium result by Hickman (2010) requires. For the main contribution of this paper, which is testing whether the bidders are rational in a sense of responding to MBI in a way that is consistent with an equilibrium, this is enough. Moreover, if truthful bidding is rejected, we can conclude that the second-price auction (whether sealed-bid or ascending) framework is not the correct model for the analysis of internet auctions. The test hypothesis is formulated as:

H_0: Truthful bidding

H_1: MBI affects bidding strategies

I use the bid history data (current price and current winner identity) to infer which auctions used which of the four the possible price rules. In Table 6, I show these pricing rules, how a given pricing rule is identified based on the bid history and what we know about the valuations under each pricing rule when we assume H_0. Moreover, Table 6 reports how the use of different pricing rules is distributed over auctions with different values of MBI. I denote the winning or selling price by WP, the highest bid by HB and the second-highest bid by SHB. t_{HB} and t_{SHB} denote the time when these bids are submitted. V denotes valuation and $N : N$

is the highest order statistic among N realizations of a random variable and $(N-1):N$ is the second-highest order statistic.

[Table 6 about here]

It is possible to detect all the different pricing rules that are used based on the bid history. Firstly, the second-price rule with an MBI was used if the last observed price increase is exactly the MBI. Secondly, the first-price rule was used if the last price increase is smaller than the MBI.[6] In my data, only 2% of all the auctions and 5% of the auctions that ended with the pricing rule determined by equation (1) use this pricing rule. This is a low number even when accounting for the fact that in a third of my auctions (those with the 1 cent MBI), the first-price rule is not possible. Moreover, 2% is a very low number compared to the emphasis I put on the pricing equation (1). However, its low realization in my data does not mean that the MBI is not accounted for in the bidding strategies, because bidding is based on the expectations over the pricing rule rather than its realizations. Thirdly, the last price increase can be higher than the MBI only if a jump bid wins the auction. Fourthly, the last price increase can be zero only if an early proxy bidder wins the auction. Due to the use of this preferential treatment of early bidders, it is harder to reject the null hypothesis than it would be if the same experiment was run without this preferential treatment. In an alternative analysis, I use the early winner auctions as a separate group and test whether those auctions are absent of bid-shading as they should.

The main idea of the test is that MBI cannot affect the valuations that bidders have for the sold objects. If the calculated value distributions differ for the different MBIs, it means that the MBI has to have affected the bidding strategies. To achieve this, I compare the valuation distributions under different MBIs by utilizing

[6] In Hickman's (2010) eBay data, about 22% of the auctions used the first-price rule. This may seem to be surprising, because on eBay, the MBIs are quite small compared to the current prices. However, this is not the correct comparison. The ex-ante probability of the first-price rule is determined by the difference between the highest valuation and the second-highest valuation relative to the MBI. This probability may be high even if the MBI is low relative to the current price.

order statistics as in Table 6. Only the winning price data can be used, because I do not observe and cannot infer the other proxy bids. This calculation strategy assumes H_0. I should emphasize that this assumption does not need to be correct for the purposes of testing. On the contrary, the test will reveal whether this assumption made in the calculation is correct. Auctions with a jump bid winner are omitted because it is not realistic to assume H_0 for them. To simplify testing, I also exclude the cases that used the first-price rule. The test hypotheses for the two experiments are:

H_0: $F_V(.|MBI = 1, NP = 15) = F_V(.|MBI = 33, NP = 15) = F_V(.|MBI = 50, NP = 15)$

H_1:$F_V(.|MBI = 1, NP = 15) \neq F_V(.|MBI = 33, NP = 15)$ or $F_V(.|MBI = 1, NP = 15) \neq F_V(.|MBI = 50, NP = 15)$ or $F_V(.|MBI = 33, NP = 15) \neq F_V(.|MBI = 50, NP = 15)$, and

H_0: $F_V(.|MBI = 1, NP = 50) = F_V(.|MBI = 66, NP = 50) = F_V(.|MBI = 100, NP = 50)$

H_1: $F_V(.|MBI = 1, NP = 50) \neq F_V(.|MBI = 66, NP = 50)$ or $F_V(.|MBI = 1, NP = 50) \neq F_V(.|MBI = 100, NP = 50)$ or$F_V(.|MBI = 66, NP = 50) \neq F_V(.|MBI = 100, NP = 50)$.

Alternatively, I can again pool the data from the separate experiments together and test for:

H_0: $F_V(.|MBI = 1) = F_V(.|MBI = \text{eBay level}) = F_V(.|MBI = \text{high})$

H_1:$F_V(.|MBI = 1) \neq F_V(.|MBI = \text{eBay level})$ or $F_V(.|MBI = 1) \neq F_V(.|MBI = \text{high})$ or $F_V(.|MBI = \text{eBay level}) \neq F_V(.|MBI = \text{high})$.

Whether these distribution are the same can be tested with a number of standard tests. I will conduct two types of testing. First, and most importantly, I will conduct a nonparametric test that does not require any assumptions concerning the distribution of valuations or the number of potential bidders N. Here, I will simply use the Kolmogorov-Smirnov test to test whether the empirical cumulative distribution functions (ecdfs) of the observed $V_{(N-1):N}$ (as calculated in Table 6) are the same for different levels of MBI. Before testing, I divide the calculated valuations by the nominal value $NP = 15$ or 50 for the respective experiment. The minor limitation of the Kolmogorov-Smirnov test is that it is designed for continuous distributions, whereas I am testing between discrete distributions. In the case of discrete distributions, the Kolmogorov-Smirnov test is conservative (see e.g. Noether 1968), and therefore, I am comfortable with using it. Second, as a robustness check, I conduct a parametric test that requires more assumptions. This test is discussed and reported in the Appendix B. It largely echoes the findings of the main test.

In Figure 1, I show the distributions for different levels of the MBI for the 15 euro and the 50 euro experiment separately. The distribution for the high MBI auctions is mainly to the left of the others in both the experiments, as one would expect with rational bidding. The same pattern is observed in Figure 2 for the pooled experiment results. These differences are present even without any controls. The differences become even more striking when I control for the week and winner fixed effects by comparing the ecdfs of the residuals resulting from regressing the $V_{(N-1):N}$ to the week and winner fixed effects (Figure 2). With these controls, also the difference between the mid-range MBI and the low MBI becomes visible. I use the specification with controls also to conduct the test in a subgroup of auctions that excludes those won by the single dominant bidder (Figure 3). This is done in order to confirm that the test results are not driven be the single strong bidder.

The results of testing for differences in the ecdfs of $V_{(N-1):N}$ (Figure 1) or the residuals (Figures 2 and 3) formally are presented in Table 7. In line with the Figures, the null of truthful bidding is typically rejected when the highest MBI level is compared to the lowest or the mid-range level. When the mid-range level is compared to the lowest level, the difference is significant only when excluding the auctions won by the largest bidder from the sample. While this indicates that the results are slightly mixed, the similarities between the distributions are rejected way too often and occasionally with so low p-values, that these results are highly unlikely be simply due to chance. Moreover, the direction of the deviations from the null is in all the cases consistent with strategic rational bidding, even if the difference is not always significant. Furthermore, I show in Figure 4 that the empirical cumulative distribution for the 1 cent MBI auctions and those won by preferentially treated early bidders are practically identical. This is further comforting since one should not expect bid-shading in either case. To summarize, truthful bidding is rejected, whereas the results are consistent with bid shading, and thus, the equilibrium behavior as proposed by Hickman (2010). This is the main result of the paper.

Overall, this section implies that most of the previous literature that assumes that internet auctions are second-price auctions is overly simplistic. The result is fundamental for future theoretical work, because it is not clear how the bidders formulate first-price bidding strategies in this context. Since the number of potential bidders may be unknown to everyone, and unlike in second-price auctions or incremental bidding, first-price auction (or Hickman's (2010) hybrid auction) equilibrium strategies are a function of the number of potential bidders, it is not immediately clear how to go on analyzing online auction equilibrium bidding.

[Table 7 about here]

[Figure 1 about here]

[Figure 2 about here]

[Figure 3 about here]

[Figure 4 about here]

4 Conclusions

In this study, I argue that the MBI is an important yet often overlooked feature of internet auctions. I conduct a field experiment to study the effects of the MBI especially on bidding strategies, but also on seller revenue and entry in internet auctions. The institutional set up of the Finnish internet auction site Huuto.net allows for a novel field experiment. I sell otherwise identical objects with different MBIs. To my knowledge, this is the first experimental study on this subject. I find that it is optimal for the seller to set the minimum bid increment level higher than the lowest possible level. I also find that the number of actual bidders is decreasing in the MBI. Since my experiment reveals that the MBI is a relevant determinant of both of these outcomes, the effects of the MBI on bidder behavior and auction outcomes should be incorporated into further analysis of internet auctions.

Importantly, I use a novel econometric test to distinguish between truthful bidding and bid shading as a strategic equilibrium response to MBI. I reject the truthful bidding hypothesis. This implies that these auctions cannot be modeled as second-price auctions. This observation challenges the assumptions of many of the previous studies on internet auctions. It is very important to stress that despite the typical small size of MBIs relative to the values of auctioned object, the bid-shading imposed by MBI in equilibrium may have substantial theoretical and empirical relevance. In particular, Hickman et al. (2011) show that if bidders account for MBIs in bidding strategies, structural empirical estimated based on the truthful bidding abstraction are substantially biased. Since I show that bidders account for MBIs in bidding strategies, my study shows that a serious structural empirical analysis of online auctions needs to correctly account for the MBIs. Finally, the results show that real-world normal people are able to respond rationally and according to the economic theory on seemingly small changes in the auction design.

5 Ackowledgements

I thank Estelle Cantillon, Phil Haile, Brent Hickman, Tim Hubbard, Ari Hyytinen, Klaus Kultti, Laurent Lamy, John List, Emmi Martikainen, Harry Paarsch, Heikki Pursiainen, Tuukka Saarimaa, Matti Sarvimäki, Tanja Saxell, Otto Toivanen, Alexander White, Hanna Willman-Iivarinen and anonymous referees for their comments, Niklas Jahnsson for the excellent research assistance and Toni Ruuska at Huuto.net for providing industry expertise. Conference and seminar participants at Helsinki, EARIE, FEAAM, Toulouse, London, Paris, Copenhagen and CEPR-IO provided useful discussions. I am grateful to the Yrjö Jahnsson Foundation (YJF) for funding this experiment and to the Emil Aaltonen Foundation for funding in general. I thank STICERD at the LSE for their hospitality and the Jenny and Antti Wihuri Foundation and YJF for funding that visit.

6 References

Al-Ubaydi O., List J. (2013) "On the Generalizability of Experimental Results in Economics". In Frechette G. & Schotter A., Methods of Modern Experimental Economics, Oxford University Press.

Bajari P., Hortacsu A. (2003) "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions". The RAND Journal of Economics 34, 329-355.

Bapna R., Goes P., Gupta A. (2003). "Analysis and Design of Business-to-Consumer Online Auctions". Management Science 49, 85-101.

Brown J., Morgan J. (2009). "How Much Is a Dollar Worth? Tipping versus Equilibrium Coexistence on Competing Online Auction Sites". Journal of Political Economy 117, 668-700.

Cameron A., Gelbach J., Miller D. (2006). "Robust Inference with Multi-Way Clustering". NBER Technical Working paper 327.

Cameron A., Gelbach J., Miller D. (2008). "Bootstrap-based improvements for inference with clustered errors". Review of Economics and Statistics 90, 414-427.

Che Y.-K., Gale I. (1998). "Standard Auctions with Financially Constrained Bidders". Review of Economic Studies 65, 1-21.

Harrison G., List J. (2008). "Naturally Occuring Markets and Exogenous Laboratory Experiments: A Case Study of the Winner's Curse." The Economic Journal 118, 822–843.

Hickman B. (2010). "On the Pricing Rule in Electronic Auctions". International Journal of Industrial Organization 28, 423-433.

Hickman B., Hubbard T., Paarsch H. (2011). "Investigating the Importance of Pricing-Rule Mis-Specification in Empirical Models of Electronic Auctions". Mimeo.

Hortacsu A., Nielsen E. (2010). "Do Bids Equal Values on eBay". Marketing Science 29, 994-997.

List J. (2011). "Why Economists Should Conduct Field Experiments and 14 Tips for Pulling One Off". Journal of Economic Perspectives 25, 3-16.

List J., Lucking-Reiley D. (2000). "Demand Reduction in Multi-Unit Auctions: Evidence from a Sportscard Field Experiment." American Economic Review 90, 961-972.

List J., Lucking-Reiley D. (2002). "Bidding Behavior and Decision Costs in Field Experiments." Economic Inquiry 40, 611-619.

Krishna V. (2002). "Auction Theory". Academic Press.

Maskin E., Riley J. (1984). "Optimal Auctions with Risk Averse Buyers". Econometrica 52, 1473-1518.

Milgrom P., Weber R. (1982). "A Theory of Auctions and Competitive Bidding". Econometrica 50, 1089-1122.

Moerbeek M. (2005). "Randomization of Clusters Versus Randomization of Persons Within Clusters: Which Is Preferable?" The American Statistician 59, 72-78.

Myerson R. (1981). "Optimal Auction Design". Mathematics of Operations Research 6, 58-73.

Noether G. (1968). "Note on the Kolmogorov Statistic in the Discrete Case". Metrika 7, 115-116.

Ockenfels A., Roth A. (2006). "Late and Multiple Bidding in Second Price internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction". Games and Economic Behavior 55, 297-320.

Reiley D. (2006). "Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet." RAND Journal of Economics 37, 195-211.

Riley J., Samuelsson W. (1981). "Optimal Auctions". American Economic Review 71, 381-392.

Rogers A., David E., Jennings N., Schiff J. (2007). "The Effects of Proxy Bidding and Minimum Bid

Increments within eBay Auctions". ACM Transactions on the Web 1, No. 2, Article 9.

Song U. (2004). "Nonparametric Estimation of an eBay Auction Model with an Unknown Number of Bidders". University of British Columbia Mimeo.

Zeithammer R., Adams C. (2010). "The Sealed Bid Abstraction in Online Auctions". Marketing Science 29, 964-987.

Table 1. Descriptive statistics for the 15 euro experiment.

Sample	Variable	Obs	Mean	Std. Dev.	Min	Max
All	price	72	13.24	0.20	13	13.65
MBI=1	price	24	13.18	0.20	13	13.51
MBI=33	price	24	13.27	0.18	13	13.65
MBI=50	price	24	13.26	0.22	13	13.52
All	bidders	72	3.06	0.98	2	5
MBI=1	bidders	24	3.25	0.99	2	5
MBI=33	bidders	24	3.04	1.08	2	5
MBI=50	bidders	24	2.88	0.85	2	5

Notes: "price" denotes the transaction price at which the object is sold. "bidders" means the number of different bidder identities that are observed to submit bids.

Table 2. Descriptive statistics for the 50 euro experiment.

Sample	Variable	Obs	Mean	Std. Dev.	Min	Max
All	price	72	44.46	1.30	41	46.64
MBI=1	price	24	44.31	1.33	41	46
MBI=66	price	24	44.69	1.36	41	46.64
MBI=100	price	24	44.40	1.23	41	45.4
All	bidders	72	4.08	0.99	2	7
MBI=1	bidders	24	4.46	0.98	3	7
MBI=66	bidders	24	3.92	1.10	2	6
MBI=100	bidders	24	3.88	0.80	2	5

Notes: "price" denotes the transaction price at which the object is sold. "bidders" means the number of different bidder identities that are observed to submit bids.

Table 3. Results of the 15 euro experiment.

	Price		
	(1)	(2)	(3)
MBI33	0.090*	0.090**	0.073**
s.e.	0.054	0.044	0.035
MBI50	0.081	0.081*	0.085**
s.e.	0.060	0.043	0.035
R^2	0.04	0.50	0.77
	Number of Bidders		
	(4)	(5)	(6)
MBI33	-0.208	-0.208	-0.115
s.e.	0.299	0.209	0.196
MBI50	-0.375	-0.375*	-0.250
s.e.	0.266	0.218	0.218
R^2	0.02	0.55	0.68
Week FE	no	yes	yes
Winner FE	no	no	yes

Notes: N=72. The reference (control) group for the treatment groups is the 1 cent MBI. "MBI33" and "MBI50" denote dummies for the 33 cent and 50 cent MBI treatments. All standard errors are robust to heteroskedasticity. * denotes 10% significance level, ** denotes 5% level and *** 1% level. The null hypothesis for mbi33=mbi50 is not rejected in any of the specifications.

Table 4. Results of the 50 euro experiment.

		Price	
	(1)	(2)	(3)
MBI66	0.380	0.380	0.323*
s.e	0.390	0.239	0.163
MBI100	0.092	0.092	-0.071
s.e	0.371	0.243	0.133
R^2	0.02	0.61	0.88
		Number of Bidders	
	(4)	(5)	(6)
MBI66	-0.542*	-0.542**	-0.548**
s.e	0.300	0.223	0.229
MBI100	-0.583**	-0.583***	-0.639***
s.e	0.257	0.196	0.198
R^2	0.07	0.55	0.57
Week FE	no	yes	yes
Winner FE	no	no	yes

Notes: N=72. The reference (control) group for the treatment groups is the 1 cent MBI. "MBI66" and "MBI100" denote dummies for the 66 cent and 100 cent MBI treatments. All standard errors are robust to heteroskedasticity. * denotes 10% significance level, ** denotes 5% level and *** 1% level. The null hypothesis for mbi66=mbi100 is rejected only in specification (3) at the 5% level.

Table 5. Pooled results of both experiments.

	Discount					
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment 1	0.0068	0.0068**	0.0084***	0.0084**	0.0084***	0.0084**
s.e	0.0043	0.0028	0.0023	0.0032	0.0020	0.0026
Treatment 2	0.0036	0.0036	0.0043*	0.0043	0.0043	0.0043
s.e	0.0042	0.0028	0.0023	0.0037	0.0026	0.0026
R^2	0.02	0.59	0.77	0.77	0.77	0.77
	Number of Bidders					
	(7)	(8)	(9)	(10)	(11)	(12)
Treatment 1	-0.375	-0.375**	-0.404***	-0.404**	-0.404**	-0.404**
s.e	0.236	0.153	0.146	0.177	0.184	0.180
Treatment 2	-0.479**	-0.479***	-0.505***	-0.505**	-0.505	-0.505
s.e	0.216	0.146	0.144	0.210	0.298	0.292
R^2	0.03	0.64	0.69	0.69	0.69	0.69
Week FE	no	yes	yes	yes	yes	yes
Winner FE	no	no	yes	yes	yes	yes
Clustering	no	no	no	week	winner	two-way

Notes: N=144. The reference (control) group for both the treatment groups is the 1 cent MBI. "Treatment 1" and "Treatment 2" denote dummies for the mid-level and high-level MBI treatments respectively. All the standard errors are robust to heteroskedasticity. In the last three columns, the standard errors are clustered either at the week level, at the winning bidder identity level or at both levels simultaneously using two-way clustering (the method of Cameron et al. (2006) is implemented with ivreg2 in STATA). The clustered significance levels are corrected for the small number of clusters. Moreover, using the wild bootstrap procedure to account for the small number of clusters produces even smaller p-values than reported here (this is checked with the Cameron et al. (2008) method as implemented by cgmwildboot in STATA, but not reported here). * denotes 10% significance level, ** denotes 5% level and *** 1% level. The null hypothesis for Treatment 1 = Treatment 2 is rejected only in specifications (5) and (6) at the 5% level.

Table 6. Pricing rules and valuation distribution estimation.

Pricing rule	Timing	Bid history	Valuation calculus
$WP = SHB + MBI$ (eq. (1))	$t_{HB} > t_{SHB}$	$WP - SHB = MBI$	$WP - MBI = V_{(N-1):N}$
$WP = HB$ (eq. (1))	$t_{HB} > t_{SHB}$	$WP - SHB < MBI$	Not conducted
$WP = HB$ (jump bid)	not relevant	$WP - SHB > MBI$	Not conducted
$WP = SHB$ (early winner)	$t_{HB} < t_{SHB}$	$WP - SHB = 0$	$WP = V_{(N-1):N}$

Shares	MBI All	$MBI = 1$	$MBI = 33\|66$	$MBI = 50\|100$
$WP = SHB + MBI$ (eq. (1))	48%	30%	56%	57%
$WP = HB$ (eq. (1))	2%	0%	5%	2%
$WP = HB$ (jump bid)	23%	21%	23%	24%
$WP = SHB$ (early winner)	28%	49%	16%	19%

Table 7. Testing for differences in the empirical cumulative distribution functions

Testing for differences in the ecdf's						
	15 euro experiment, no controls			50 euro experiment, no controls		
MBI	1 vs. 33	1 vs. 50	33 vs. 50	1 vs. 66	1 vs. 100	66 vs. 100
p-value, KS-test	0.1813	0.0050	0.4614	0.6821	0.0041	0.0063
	Pooled, no controls			Pooled, week and winner FE as controls		
MBI	1 vs. T1	1 vs. T2	T1 vs. T2	1 vs. T1	1 vs. T2	T1 vs. T2
p-value, KS-test	0.8217	0.0096	0.0192	0.2453	0.0008	0.0404
	Pooled, excluding the largest bidder					
	no controls			week and winner FE as controls		
MBI	1 vs. T1	1 vs. T2	T1 vs. T2	1 vs. T1	1 vs. T2	T1 vs. T2
p-value, KS-test	0.0240	0.0000	0.0010	0.0832	0.0005	0.0282

Notes: "1" refers to the 1 cent MBI. "T1" and "T2" denote the mid-level and high-level MBI treatments respectively.

Figure 1. Empirical cumulative distributions for the valuations calculated under H0 for 15 and 50 euro experiments.

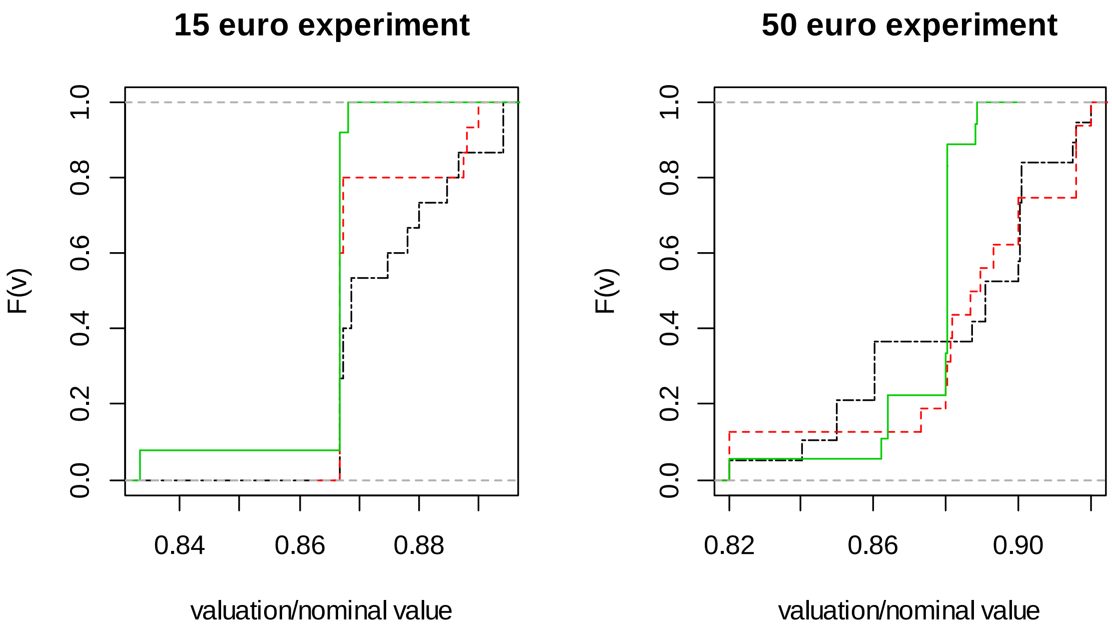

Notes: The green connected line represents the high MBI auction, the red dotted line the mid-range MBI auctions and the black semi-dotted line the low MBI auction.

Figure 2. Empirical cumulative distributions for the valuations calculated under H0 or residuals of these valuation (given week and winner fixed effects) for the pooled data.

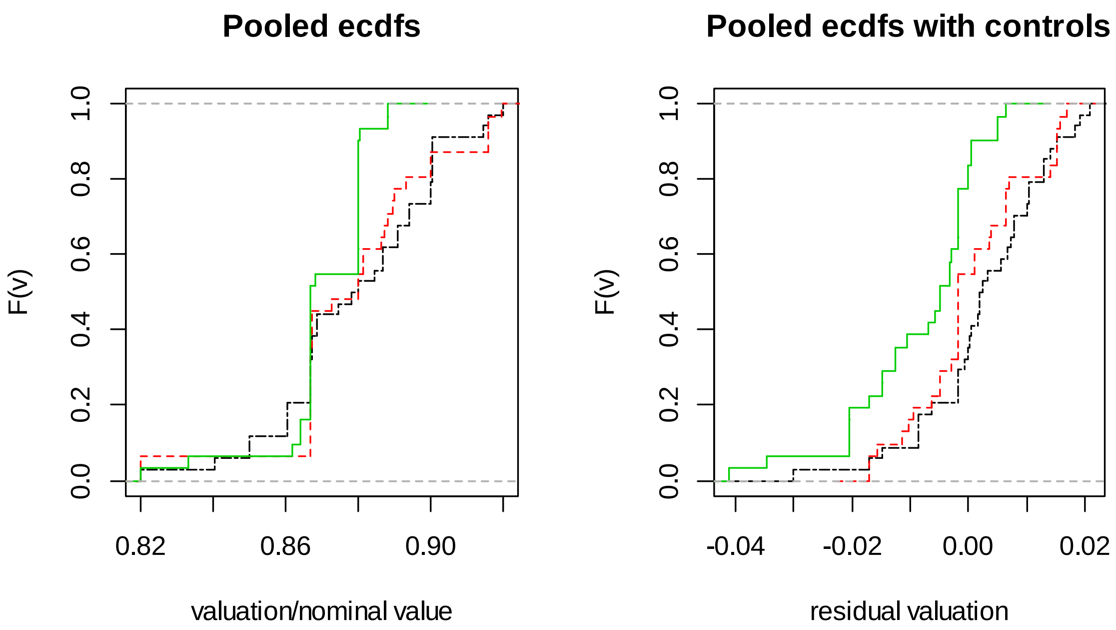

Notes: The green connected line represents the high MBI auction, the red dotted line the mid-range MBI auctions and the black semi-dotted line the low MBI auction. The right hand graph represents empirical cumulative distribution functions of residuals from regressing (valuation $V_{(N-1):N}$/nominal value) on week and winner fixed effects. The left hand graphs represents empirical cumulative distribution functions of these valuations.

Figure 3. Empirical cumulative distributions for the valuations calculated under H0 or residuals of these valuation (given week and winner fixed effects) for the pooled data, excluding auctions won by the largest bidder from the sample.

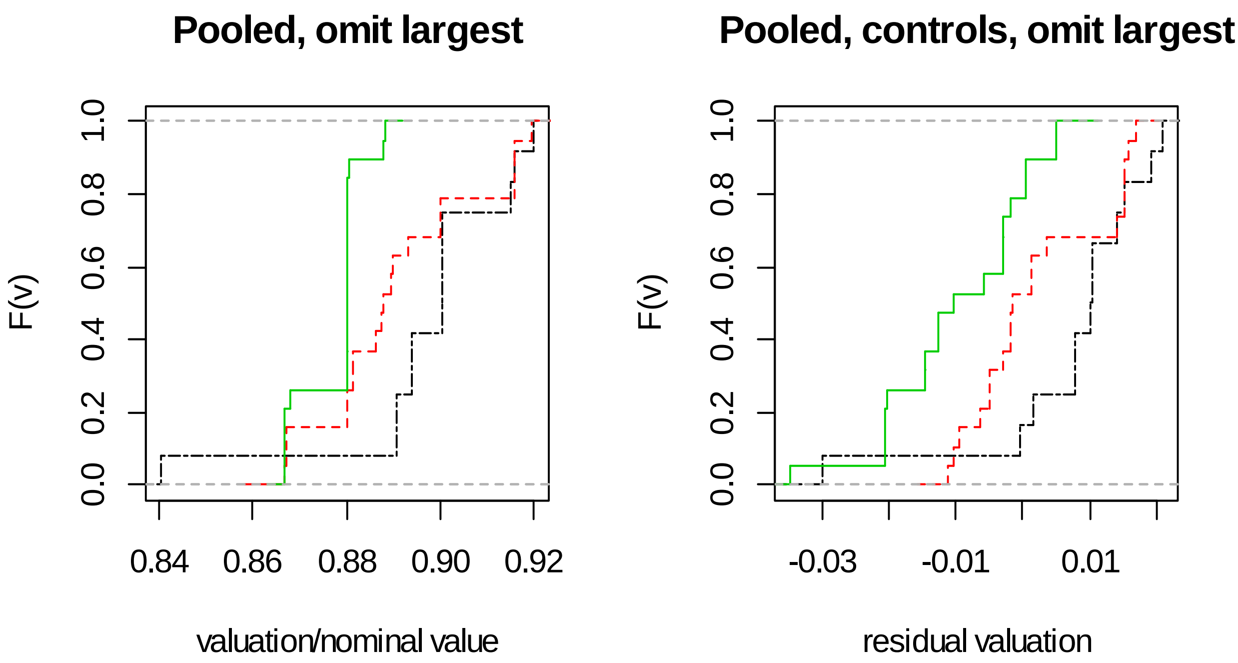

Notes: The green connected line represents the high MBI auction, the red dotted line the mid-range MBI auctions and the black semi-dotted line the low MBI auction. The right hand graph represents empirical cumulative distribution functions of residuals from regressing (valuation $V_{(N-1):N}$/nominal value) on week and winner fixed effects. The left hand graphs represents empirical cumulative distribution functions of these valuations. The auctions won by the bidder who won most auctions in the entire experiment are omitted.

Figure 4. Empirical cumulative distributions for the valuations calculated under H0 for the pooled data for the auctions where an early bidder won or MBI=1.

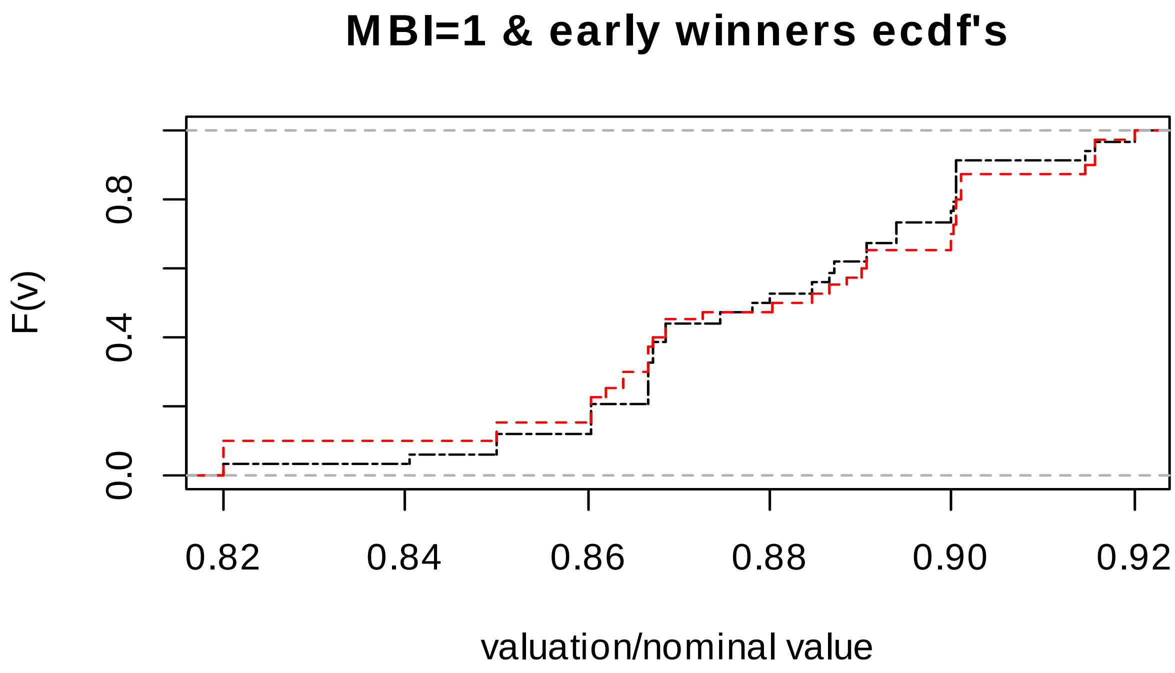

Notes: In the left-hand-picture, the green connected line represents the high MBI auction, the red dotted line the mid-range MBI auctions and the black semi-dotted line the low MBI auction. In the right-hand-picture the red dotted line represents the auctions won by an early bidder and the black semi-dotted line the low MBI auction.

Appendix A: Additional Tables

Table A1: MBI schedule in eBay.

Current Price	Bid Increment
$ 0.01 - $ 0.99	$ 0.05
$ 1 - $ 4.99	$ 0.25
$ 5 - $ 24.99	**$ 0.50**
$ 25 - $ 99.99	**$ 1.00**
$ 100 - $ 249.99	$ 2.50
$ 250 - $ 499.99	$ 5.00
$ 500 - $ 999.99	$ 10.00
$ 1000 - $ 2499.99	$ 25.00
$ 2500 - $ 4999.99	$ 50.00
$ 5000 and up	$ 100.00

Appendix B: Parametric test of truthful bidding

In this Appendix, I will use a t-test to analyze whether the structurally estimated parameters of the underlying valuation distributions are the same for different levels of MBI. The estimation of these parameters requires more assumptions than the nonparametric Kolmogorov-Smirnov test and should thus be seen only as a robustness check. First, estimating these model primitives requires information on the number of potential bidders, N. N is unknown, but it can be assumed to be the same, certainly within each batch, and perhaps also within each experiment. I use the total number of different bidder identities observed in all the auctions in the entire experiment as N. This happens to be 8 for both the experiments. Given balanced batches in respect of different MBI levels, the assumption regarding N should not influence the results to a large extent. To check this I also confirmed that results are robust to some alternatives values of N.[7]

Given N and the pricing rule, I can use the order statistic formula and estimate the distribution of the valuations F_V either with non-parametric or parametric techniques.[8] I use parametric approach due to the small number of observations to get more power. The assumption on the functional form is the second main assumption for this testing approach. I estimate the parameters of a Weibull distribution that is upper-truncated at 1. Upper-truncation is warranted, because I have divided the calculated valuations $V_{(N-1):N}$ (see Table 6) by the nominal value of the sold cards. The Weibull specification is chosen because it is fairly flexible. Experimental data is useful in this exercise, because there are no observable or unobservable characteristics that need to be controlled for, with the possible exception of the week and winner fixed effects. However, these were shown not to affect the coefficients of the revenue results and only managed to reduce the standard errors. Therefore, not including week or winner fixed effects should work against rejecting H_0. Nonetheless, I also report in Table B1 results that include these controls. I control for the week and winner

[7] I checked that the results are qualitatively the same as in Table B1 for $N = 6$ and $N = 15$ for the pooled experiment with controls. These results are not reported, but they are available from the author.

[8] A more sophisticated approach to estimate the model primitives also without the need to assume N is provided by Song (2004). However, the bid history in Huuto.net does not contain the necessary information for such analysis. I would need to observe two valuations with known ranks which is not possible without ambiguity.

fixed effects by estimating the Weibull distribution of the residuals resulting from regressing the $V_{(N-1):N}$ to the week and winner fixed effects. I also scale these residuals by adding the regression model constant to the residuals to avoid having negative values when estimating the parameters of the Weibull functions.

In Table B1, I present the testing results. The top panel is simply showing the estimated parameters of the Weibull distributions and, somewhat trivially, whether they differ from zero for the individual experiments separately. The results for testing between the differences between the estimated parameters of these Weibull distributions are presented in the second panel. Likewise, the parameter estimates for the pooled data with (Figures B2) and without controls (Figures B1) are presented in the third panel, and the main test results for them in the fourth. The evidence is fairly consistent in showing that the distributions for the highest MBI levels differ from the other two distributions. Therefore, I reject the H_0 of truthful bidding also with this test. Moreover, a further out-of-sample piece of evidence supporting this result to follow from Hickman (2010)-type bid-shading as a reaction to the MBI is that the observed pattern is similar to what Hickman et al. (2011) find in their Monte Carlo analysis. They compare estimations based on truthful bidding in the presence of an MBI to the true value distributions (and estimates that account for the MBI). They find that truthful estimates are located to the left and have a higher peak than the true distribution. This pattern is replicated in my experiment because a higher MBI leads either to a higher peak or a location to the left of the distributions with a lower MBI, although these patterns are not always statistically significant.

There is a potential risk of over-interpretation of the data involved in this testing approach, because I have treated each auction as an iid observation in the estimation. There are various realistic reasons for the observations not to be iid. First, it is possible that a new draw of a given bidder is closer to her own old draw than the draw of some other bidder. This seems to be likely, given for example the large share of auctions won by bidder 1. This would be the case of asymmetric bidders, where a separate valuation distribution should be estimated for each bidder. However, in online auctions the possible differences between the bidders are not likely to be commonly observed, thus making the bidder symmetry the appropriate model. Moreover, the symmetry assumption is standard in empirical online auction literature. Furthermore, even if the symmetry assumption was unrealistic, it is not likely to bias the comparison of the treatments, because each bidder wins about the same amount of auctions in each treatment (see Table C2). A worse, yet possible case is that bidders draw valuations only once. In that case, I have been replicating data. However, if that is really the case, there should have been less variation in prices under truthful bidding than I observe in the data. Furthermore, the distribution from which valuations are drawn may change over time, for example as a function of items won in the past. However, if this was likely to confound the test here, it should also have affected the revenue results earlier, which it did not based on the week fixed effects results. Moreover, week (and winner) fixed effects are controlled for in some of the results in Table B1.

Table B1. Alternative test results.

	15 euro experiment			50 euro experiment		
	Location and shape parameters of the Weibull distribution					
MBI	**1**	**33**	**50**	**1**	**66**	**100**
Location	0.862***	0.861***	0.854***	0.85***	0.86***	0.86***
se	0.0042	0.0032	0.0038	0.0084	0.0092	0.0042
Shape	30.0***	39.1***	40.0***	13.8***	13.9***	30.0***
se	5.70	9.94	9.67	2.35	2.66	5.70
	Testing for differences in the Weibull parametes					
MBI	**1 vs. 33**	**1 vs. 50**	**33 vs. 50**	**1 vs. 66**	**1 vs. 100**	**66 vs. 100**
Location	0.0013	0.0080	0.0067	-0.0087	-0.0039	-0.0048
se	0.0053	0.0056	0.0049	0.0094	0.0125	0.0101
Shape	-9.08	-9.97	-0.90	-0.16	-16.27**	-16.1**
se	8.99	11.22	11.90	3.55	6.17	6.29
	Pooled			**Pooled with controls**		
	Location and shape parameters of the Weibull distribution					
MBI	**1**	**T1**	**T2**	**1**	**T1**	**T2**
Location	0.856***	0.855***	0.858***	0.891***	0.890***	0.881***
se	0.0051	0.0054	0.0031	0.0024	0.0024	0.0026
Shape	17.0***	16.76***	30.1***	37.1***	39.4***	36.9***
se	2.12	2.19	4.15	4.67	5.07	4.97
	Testing for differences in the Weibull parametes					
MBI	**1 vs. T1**	**1 vs. T2**	**T1 vs. T2**	**1 vs. T1**	**1 vs. T2**	**T1 vs. T2**
Location	0.0009	-0.0017	-0.0026	0.0016	0.011***	0.0094***
se	0.0074	0.0059	0.0062	0.0034	0.0035	0.0035
Shape	0.273	-13.0***	-13.31***	-2.26	0.17	2.43
se	3.04	4.66	4.69	6.89	6.82	7.10

Notes: *** denotes statistical significance at <1%, ** denotes significance at <5%, and * at <10%.

Figure B1. Estimated cumulative Weibull distributions and Weibull densities for the valuations calculated under H0 for the pooled data.

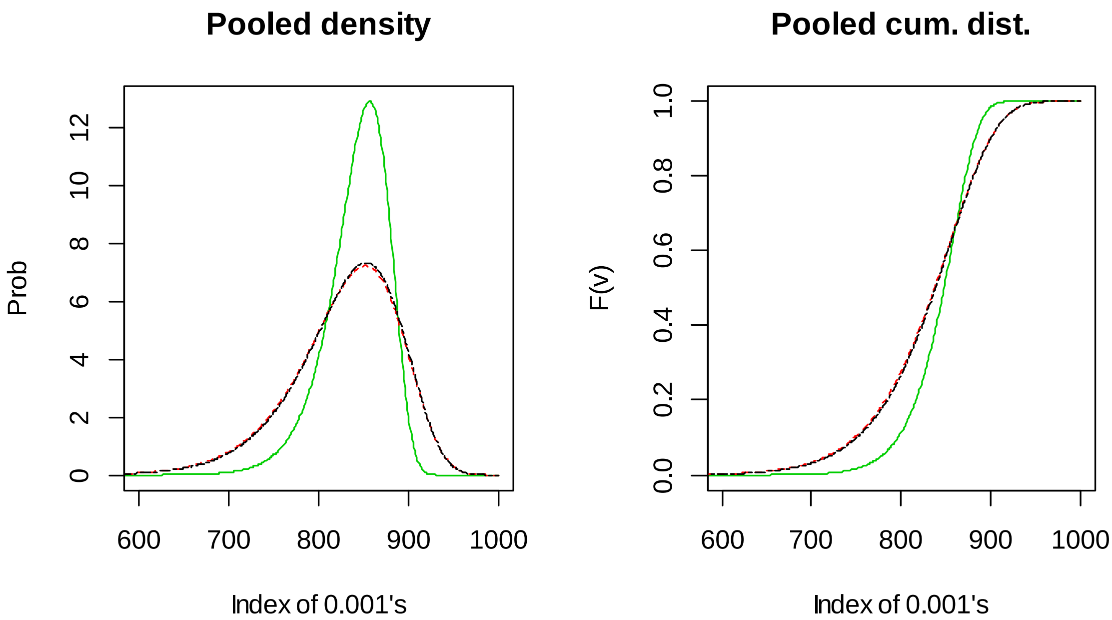

Notes: The green connected line represents the high MBI auction, the red dotted line the mid-range MBI auctions and the black semi-dotted line the low MBI auction.

Figure B2. Estimated cumulative Weibull distributions and Weibull densities for the valuations calculated under H0 for the pooled data with controls.

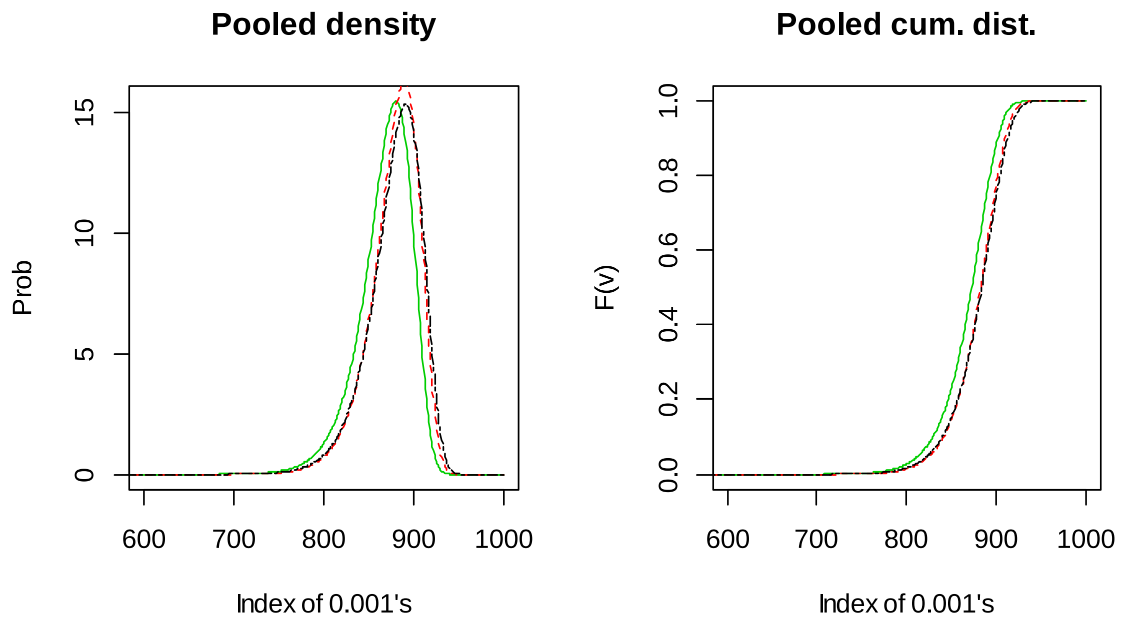

Notes: The green connected line represents the high MBI auction, the red dotted line the mid-range MBI auctions and the black semi-dotted line the low MBI auction.

Appendix C: External validity

In this Appendix, I study to what extent the results may generalize in respect of the characteristics of the bidders, the objects sold and the auction site. The first reason is related to whether the bidders represent a typical set of bidders in internet auctions. For example, more experienced bidders may behave in a more sophisticated way and more in line with equilibrium than less experienced bidders (see e.g. Harrison and List 2008). I analyze the set of participants in Table C1. I describe how many auctions each of the observed bidders won. Altogether 13 different bidder identities are observed in the data. 8 different identities won auctions in the first experiment and 8 in the second, of which 3 were also winners in the first experiment. A large share of the auctions are won by bidder 1. She dominates especially the first experiment, which may be a concern for the external validity of these results. Besides winning most of the objects, this bidder pays on average the lowest prices. The issue in the first experiment is that, although there are many different bidders that participate in these auctions, one bidder wins most of the auctions. It is tempting to argue that the results could have been different if that particular bidder was not present. On the other hand, this particular bidder does not dominate the second experiment and yet the results are very similar in both the first and the second experiment. Moreover, the price is not determined by the winning bidder alone, as in standard first-price auctions, but rather as a joint function of both the highest and the second-highest bid, as can be seen from equation (1). In these experiments, there is much more variation in the sets of participants than there is in the winner identities. For these reasons, the external validity of these results with respect to bidders should be fairly strong.

Table C1. How many auctions each bidder won.

bidder	w1	w2	w3	w4	w5	w6	w7	w8	w9	w10	w11	w12	total	price 15	price 50
1	11	12	6	12	5	9	4	11	8				78	13.16	43.1
2	1											1	2	13.35	45
3			4										4	13.45	NA
4			1			1							2	13.49	NA
5			1										1	13.5	NA
6					5				4				9	13.51	45.02
7					2								2	13.5	NA
8						2							2	13.42	NA
9							7						7	NA	45.13
10							1	1					2	NA	43.52
11										8		9	17	NA	44.87
12										4			4	NA	45.72
13											12	2	14	NA	45.39

Notes: "w1-w6" denotes weeks 1 to 6 , which constitute the first experiment. "w7-w12" are the weeks of the second experiment. "price 15" denotes the average winning bid of each bidder in the first experiment and "price 50" in the second experiment.

To evaluate the previous arguments concerning the robustness of the results in respect of the set of bidders, I report two separate pieces of evidence in Tables C2 and C3. In Table C2, I show the number of auctions that each bidder wins in each MBI category. There is no systematic evidence in Table C2 implying that some bidders prefer some MBI levels over others. In Table C3, I show the results of the leave-one-out type revenue regressions. To be more specific, in Table C3, I report results where I omit all auctions won by a single bidder from the data, one winner at a time. The six smallest winners are dropped out together. I include the week fixed effect, but not the winner fixed effect, because here the winner effects are analyzed by dropping out data. All the results are qualitatively similar to the full sample results. None of the Treatment 1 effects are statistically significantly different from the comparable full sample effect (0.068) and all of them are significantly different from zero, at least at the 10% level.

Table C2. Distribution of wins over treatments for each bidder

bidder	statistics	MBI=1	treatment 1	treatment 2
1	obs	27	26	25
	mean	0.870	0.877	0.873
6	obs	5	1	3
	mean	0.901	0.900	0.901
9	obs	1	3	3
	mean	0.900	0.909	0.897
11	obs	6	5	6
	mean	0.896	0.896	0.900
13	obs	5	5	4
	mean	0.910	0.911	0.900
rest	obs	4	8	7
	mean	0.885	0.902	0.901

Notes: "obs" denotes the number of auctions that a given bidder wins in each treatment group. "mean" denotes the mean of price/(nominal value). "rest" denotes bidders 2,3,4,5,7,8,10 and 12 together.

Table C3. Results when data won by a given bidder is omitted.

Omitted auctions won by bidder:	1	2,4,5,7,8,10	3	6
Treatment 1	0.0042*	0.0061*	0.0057**	0.0081***
s.e.	0.0020	0.0026	0.0024	0.0025
Treatment 2	-0.0009	0.0032	0.0031	0.0032
s.e.	0.0038	0.0018	0.0017	0.0027
N	66	133	140	135
Omitted auctions won by bidder:	**9**	**11**	**12**	**13**
Treatment 1	0.0065**	0.0072**	0.0064**	0.0075**
s.e.	0.0028	0.0030	0.0029	0.0030
Treatment 2	0.0039*	0.0035	0.0037*	0.0053***
s.e.	0.0020	0.0023	0.0020	0.0010
N	137	127	140	130

Notes: The response variable is "discount". All the specifications include week fixed effects but exclude winner fixed effects. The standard errors are clustered both at the week and winning bidder identity levels using two-way clustering. The clustered significance levels are corrected for the small number of clusters. * denotes 10% significance level, ** denotes 5% level and *** 1% level.

The second potential concern for the external validity of the results is related to whether the objects sold are a good representation of typical objects sold in internet auctions. One particular concern in that one may expect more sophisticated bidding when stakes are high because making sophisticated decisions may involve cognitive decision costs (see e.g. List and Lucking-Reiley 2002). The values of the gift cards sold here do represent quite typical stakes. More generally, although there were many other types of gift cards for sale at the same time, only a few other Stockmann gift cards were on sale. Thus the object that I sell is not perhaps a typical object. On the other hand, there is large heterogeneity in the objects sold in internet auctions, and thus, there is no such thing as a representative object of sale. Moreover, the presence of identical or similar objects sold by other sellers may potentially add unnecessary noise to the experiment since they may affect bidder behavior in the experiment's auctions. Furthermore, the effect of the MBI depends more on the bidder characteristics and their strategies, their risk attitudes and the nature of the entry process than the object characteristics. There should be a huge number of markets where the entry behavior and bidder characteristics are similar enough to the Stockmann gift card markets to make these results of external interest. However, one limitation to the generalizability is that the amount of uncertainty over the valuations may be smaller for gift cards than for many other objects thus making the first-price rule more common in expectation. On the other hand, in practise, the first-price rule was applied surprisingly rarely in these data. For these reasons, the external validity of these results also in respect of the objects sold should be fairly strong.

The main concern for external validity is whether the results apply to any other sites than Huuto.net. There are several differences between Huuto.net and the other internet auction sites, most importantly eBay. One difference between the auction mechanisms used by Huuto.net and eBay is that the former applies a soft stopping rule and the latter a strict stopping rule. The soft closing rule on Huuto.net means that if a bid is submitted when the auction is about to close in less than 5 minutes, 5 minutes are added to the time that the auction is open. On eBay the closing time is strict. Amazon auctions, for example, also used a soft closing rule. The stopping rule has implications for equilibrium bidding. In particular, the strategic advantages of so-called late bidding or sniping, which is often observed on eBay, are severely attenuated in auctions that apply an automatic extension rule (Ockenfels and Roth 2006). In Yahoo! auctions, the seller could set the closing rule. Brown and Morgan (2009) use this feature to construct a field experiment on the effects of the closing rule. They find that prices and bidder counts are unaffected by the auction ending rule. Therefore it is fair to assume that the generalizability of the revenue results obtained from a field experiment set-up on Huuto.net is not limited by the particularities closing rule.

There are two other differences between the two auction sites that are more problematic. Firstly, Huuto.net gives preferential treatment to winners who submit their bid to the proxy earlier than the second-highest bid. In that case, the winner pays only the second-highest bid. Therefore, in the case of early-placed winning bids, Huuto.net auctions are really just second-price auctions instead of hybrid auctions that use the pricing rule (1). In my experiment, 28% of the auctions were won by an early proxy bidder. This set-up creates further incentives for early bidding beyond the soft closing rule. In my experiment, no last-minute bidding was observed. Due to this rule, I probably estimate the lower bound for the revenue effects of the MBI compared to conducting this experiment on eBay, because equation (1) is only relevant for auctions where a later proxy bidder wins. Therefore, qualitatively speaking my findings should be robust for this difference in the pricing rule. Moreover, when testing for bid-shading, the existence of preferential treatment can be turned into an advantage, since it provides an additional set of auctions where one should not expect bid-shading to take place.

The second large concern is that bidders are allowed to jump bid. If a jump bidder wins, she pays her bid. Unlike on eBay, bidders are not forced to use the proxy machine, which is more of an option. 23% of the auctions in my experiments were won by a jump bid. This is a non-negligible share. Unfortunately, jump bidding creates many complexities for an analysis of bidding strategies which are clearly beyond the scope of this paper. In this respect, the external validity is limited. However, jump bidding wins was equally common in all the three different MBI treatment groups, which may imply that the results concerning the MBI should not be much affected by the possibility of jump bidding. Moreover, I repeat the analysis by excluding observations that result in a jump bid winning in Table C4.[9] The effects of the MBI on revenue are slightly smaller in the richest specification and the effect on entry is smaller for this sample than for the full sample. Therefore it seems that on average jump bidders may be able to deter entry with a high MBI. However, the differences with the full sample and the sample excluding jump bids are not statistically significant. On the other hand, it is worrying that the revenue result is less stable for this subset than for the full sample. I also account for these different rules and strategies in the bidding strategy test by analyzing only the subsample of data without the auctions won by jump bids.

[9] The sample in Table C4 also excludes 14 auctions for which I do not observe the pricing rule. This is due to a human error. Originally, I did not record the information from my experiments on the pricing rule that was used. I added that information later. Between these two recording events I lost the back-up paper copies of 14 of my auctions while moving office. This lost information is not easily recoverable because Huuto.net keeps information on its auctions for 3 months only. However, I have verified that omitting these 14 observations changes nothing of practical importance relative to Table 5.

Table C4. Pooled estimation results when auctions won by jump bids are omitted

	Discount					
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment 1	0.0136**	0.0082*	0.0049	0.0049	0.0049*	0.0049*
s.e	0.0056	0.0043	0.0031	0.0041	0.0023	0.0024
Treatment 2	0.0115**	0.0081*	0.0024	0.0024	0.0024	0.0024
s.e	0.0051	0.0044	0.0034	0.0062	0.0035	0.0037
	Number of Bidders					
	(7)	(8)	(9)	(10)	(11)	(12)
Treatment 1	-0.067	-0.064	-0.195	-0.195	-0.195	-0.195
s.e	0.262	0.196	0.199	0.287	0.322	0.325
Treatment 2	-0.325	-0.283	-0.367*	-0.367	-0.367	-0.367
s.e	0.246	0.194	0.187	0.272	0.340	0.327
Week FE	no	yes	yes	yes	yes	yes
Winner FE	no	no	yes	yes	yes	yes
Clustering	no	no	no	week	winner	two-way

Notes: N=99. The reference (control) group for the treatment groups is the 1 cent MBI. "Treatment 1" and "Treatment 2" denote dummies for the mid-level and high-level MBI treatments. All the standard errors are robust to heteroskedasticity. In the last three columns, the standard errors are clustered either at the week level, at the winning bidder identity level or at both levels using two-way clustering. The clustered significance levels are corrected for the small number of clusters. * denotes 10% significance level, ** denotes 5% level and *** 1% level.