

Managing Bidder Learning in Retail Auctions

Paul Schäfer and Simon Schulten*

October 2022

[Click for most recent Version.](#)

When firms exploit behavioral biases it is natural to think that, eventually, consumers will learn to avoid their mistakes, limiting their exploitation. Profit maximizing firms, however, have an incentive to undermine such learning. We study the consumer learning dynamics and the firm's response in a multi-unit descending price auction with a simultaneous fixed price offer. In our panel of 8 million bids by 280.000 bidders, consumers often bid more than the fixed price. Depending on competing bids, an overbid can lead to paying more than the fixed price (overpaying). We argue overpaying increases the saliency of the consumers' mistake by making it payoff relevant, which is likely to affect consumer learning. Indeed, bidders who overpay subsequently overbid less often and are more likely to leave the market compared to bidders who similarly overbid but did not overpay. We show the resulting loss in future profits makes overpaying undesirable, and document a structural break in our data at which the firm eliminates such overpayments — and the resulting consumer learning — through changes in how it runs its auctions. Methodologically, we discuss identification of our treatment effects using causal graphs and show how these treatment effects identify a three-type structural model of bidder behavior with learning dynamics.

1. Introduction

Theory and evidence suggests that firms price strategically to exploit consumer biases (see, for example [Della Vigna and Malmendier, 2006](#); [Grubb, 2015](#); [Grubb and Osborne, 2015](#); [Heidhues and Kőszegi, 2018](#); [Malmendier and Szeidl, 2020](#)). Despite such evidence, a common critique is that consumers will learn to avoid being exploited. This argument, however, ignores the firm’s incentives to inhibit consumer learning. In this paper, we empirically study consumer learning in retail auctions and document how a firm improves at designing the sales environment to impede consumer learning.

Despite its potential importance, we know little about whether and how firms are able to manage consumer learning. A measurement challenge arises when a firm takes actions to prevent consumer learning. While we can in principle observe the firm’s actions, there is often little variation in those actions. Additionally, we cannot easily infer the prevented consumer learning from firm’s actions alone. When a firm improves at managing consumer learning over time, we can overcome the measurement challenge. In such a case, researchers can collect data on consumer learning and on the firm’s response.

In our case, the firm operates a televised multi-unit descending auction with uniform pricing and an online shop, where goods are sold at a fixed price. The auction starts at a high price that is lowered in increments over time. Bidders bid at the current price and the auction ends when all units are claimed. According to the uniform pricing rule, every winning bidder pays the lowest successful bid. Following the empirical literature, we call a bid that is higher than the fixed price an *overbid*.¹ When the auction price is higher than the fixed price, we call the auction *overpaid*. Crucially, overbidding does not imply overpaying, as overpaying requires that all bids in an auction are overbids (the lowest bid is higher than the fixed price).

We collect a bidder-level panel of the firm’s multi-unit descending auction. The data

⁰Schäfer: University of Bonn, Paul.Schaefer@uni-bonn.de, Schulten: DICE, University of Düsseldorf, Simon.Schulten@hhu.de, We would like to thank Paul Heidhues, Lorenz Götte, Florian Zimmermann, Axel Ockenfels, Joel Stiebale, Florian Heiß, Frauke Kreuter, Paul Hünermund, Hans-Peter Grüner, Matthew Backus, Regina Keller, seminar audiences at Bonn, Düsseldorf, Mannheim and Louvain-la-Neuve and the IT support of Mannheim School of Social Sciences for kindly providing a server. This research was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation= under Germany’s Excellence Strategy - EXC 2126/1- 390838866, Collaborative Reserch Center (SFB) 884 "Political Economy of Reforms" and Graduiertenkolleg GRK 1974 / 235577387.

¹For a descriptive analysis of overbidding in auctions run by the same firm see [Ocker \(2018\)](#).

spans over two years, detailing more than 8 million bids submitted by 280.000 bidders in approx. 70.000 auctions. The bidder level panel structure allows us to analyze consumer learning and the firm’s reactions to it. In line with lab and field evidence many customers, at the beginning of our sample, overbid (e.g. [Kagel and Levin, 2011](#); [Malmendier and Lee, 2011](#); [Gesche, 2022](#); [Ocker, 2018](#)).

In our auction format, overbidding is a necessary, but not a sufficient condition for overpaying. An auction ends in overpayment if and only if all bids are overbids. We exploit this fact to construct a natural treatment and control group design: treated are those bidders who overbid and — due to other consumers’ bids - ended up overpaying, whereas bidders who similarly overbid but did not overpay end up in the control group.

Overpaying makes the consumers’ mistake (overbidding) payoff relevant and arguably more salient. A salience effect of overpaying on future behavior compared to consumers who overbid but did not overpay is then naturally interpreted as evidence of consumer learning. We argue that theories explaining overbidding as optimal behavior are inconsistent with the observed behavioral changes following overpayment in our data. We find overpaying leads consumers to spend less in the future. More explicitly, it leads consumers to hand in fewer overbids and fewer bids overall, as well as increases the likelihood of refraining from bidding altogether.

We use a DAG as a convenient and precise way to codify causal knowledge about the data generating process. We derive our DAG from institutional knowledge about the way the firm plans and runs the auction, the auction rules (uniform pricing), and natural assumptions about bidder behavior. In addition, we discuss how such a DAG can be derived from, perhaps more familiar, structural equation models.²

Our analysis demonstrates a novel way to combine a traditional economic model, a directed acyclic graph (DAG) ([Pearl, 2009](#); [Imbens, 2020](#); [Hünermund and Bareinboim, 2019](#)), and the sufficient statistics approach ([Chetty, 2009](#)). We develop a three-type model of initial overbidders who may learn to become non-overbidders or dropouts. An overbidder who becomes a non-overbidder simply truncates her bidding function at the fixed price, so that she does not repeat her mistake. An overbidder who becomes a dropout ceases to participate in future auctions altogether. Learning not to overbid is a behavioral adjustment

²Assigning casual meaning to a structural equation model turns it into a structural causal model.

at the intensive margin, whereas dropouts represent the extensive margin. In our three-type model, conditional on bidding, learning not to overbid drives the observed behavioral change of bidders handing in fewer overbids. Dropouts, however, drive the negative effect on both overbids and non-overbids (fewer bids overall). We disentangle the two margins and find that an initial overbidder, who overpays, has a 4.2% chance of dropping out and a 7.2% chance of becoming a non-overbidder.

In the presence of consumer learning the firm faces a trade-off between extra overpaying revenues today and foregone revenue tomorrow. Back-of-the-envelope calculations demonstrate that the extraction of overpaying revenue is suboptimal in the beginning of our sample. In the second half of our sample we observe a structural break in the time series of overbidding and overpaying. Before the break, roughly 17% of all bids are overbids and 4% of auctions end in overpayment. After the break, overbidding prevalence is substantially reduced and practically no auction ends in overpayment. The structural break is accompanied by a small jump and a reversal in the trend of items sold per week, albeit statistically insignificant. Fixed prices remain unchanged at the break-point.

Management explains that they induced the structural break in overpaying with a relatively minor, but targeted quantity increase. Auctions are simultaneously run by an on-screen auctioneer and a director, who is off-screen and has access to real-time demand (e.g., number of people watching, revenue by the minute). The director uses this information to "steer" the auctions and may, if necessary, increase the quantity *during* the auction.³ Thus, the director is able to target quantity increases to auctions that are at risk of ending in overpayment, thus shutting down overpayment. Management also tells us of a new pricing policy implemented just after the end of our sample. This new pricing policy raises fixed prices and auctions routinely undercut these higher fixed prices with the starting bid, ruling out overbidding mechanically. Raising fixed prices (presumably) comes at the cost of lower online shop sales, so after an auction ends, fixed prices are lowered in a 24-hour sales discount to a small increment above the auction price. This gives customers who missed the auction the chance to purchase at a comparable price, rather than the very high regular

³In its terms and conditions the firm reserves the right to increase quantity after the auction has started. Quantity decreases are not mentioned in the terms and conditions, but there is an option to fully cancel an auction when demand is so low that the price would have to approach 0 for the auction to end.

fixed price.⁴

By documenting that our firm does not take into account consumer learning in the beginning of our sample, we contribute to the literature on non-optimizing firms (see, for example [Cho and Rust, 2010](#); [DellaVigna and Gentzkow, 2019](#); [Hortaçsu and Puller, 2008](#); [Hortaçsu, Luco, Puller and Zhu, 2019](#)). In contrast to these papers, our firm is adaptive as it gradually improves at managing consumer learning, by increasing its scope of maximization ([Doraszelski, Lewis and Pakes, 2018](#)). First, the firm optimizes its quantity choice in a targeted way, but leaves fixed prices unchanged. Later, it optimizes over quantities and fixed prices jointly. Interestingly, the last change was only implemented by a new CEO, in line with the idea of learning through noticing ([Hanna, Mullainathan and Schwartzstein, 2014](#)).

Our paper is related to the empirical literature that studies how consumers learn when they trigger a fee. This literature documents that (some) consumers learn to avoid triggering fees, but this learning effect depreciates over time ([Haselhuhn, Pope, Schweitzer and Fishman, 2012](#); [Agarwal, Driscoll, Gabaix and Laibson, 2013](#)). [Ater and Landsman \(2013\)](#) find consumers who switch their banking plans after paying an overage fee are more likely to switch to plans with larger allowances. In contrast to consumers who switch plans but did not pay overage fees, fee-switchers *increase* rather than decrease their monthly payments, suggesting that consumer learning can be detrimental to consumers.

Our finding that consumers give up on bidding rather than simply adjust their bid function to avoid overbidding also indicates a complex, and possibly non-rational, response to feedback. A transaction that leaves the bidders worse off than the reference point (i.e., overpaying) results in a negative transaction utility ([Thaler, 1999](#)) and may reduce future market participation through several channels: updated beliefs about the utility from market participation ([Backus, Blake, Masterov and Tadelis, 2022](#)), updated beliefs about their abilities ([Seru, Shumway and Stoffman, 2010](#)) and antagonizing consumes ([Anderson and Simester, 2010](#)). [Gesche \(2022\)](#) documents customer complaints in the eBay feedback system when they pay more at auction than the fixed price offered by the same seller. In our case, management confirms anecdotally that some consumers who overpay call the

⁴Another interpretation of setting a rebate to the auction price is that the auction serves as a mechanism to discover a "reasonable" fixed price.

hotline to complain.

To the best of our knowledge, we are first to illustrate how firms actively suppress or at least slow down consumer learning. If consumers bid less than their value of the good, trade is efficient and consumer attrition decreases consumer welfare. Under this assumption, the firm’s quantity targeting policy increases consumer welfare by minimizing attrition compared to the initial exploitation of overpaying. On the other hand, if consumers bid more than their value, they may pay an auction price above their value, particularly in overpaid auctions (for a survey of above value bidding see [Kagel and Levin, 2011](#), Section 1.2). Consumer attrition may then increase consumer welfare. In this case, shutting down consumer learning through targeted quantity increases may harm consumer welfare. In any case, targeted quantity increases are preferred in terms of consumer welfare over the pricing policy of increased fixed prices. The firm’s quantity targeting decreases the auction price - at least in auctions that would have ended in overpayment otherwise.

Our dynamic considerations complement the existing literature.⁵ In a static analysis, relatively few behavioral buyers suffice to generate a price impact in an auction compared to a fixed price market ([Malmendier and Szeidl, 2020](#)). Considering a dynamic setting, however, adds downsides to the exploitation incentive. In our application, exploiting overpaying causes some consumers to drop out of the market, so the firm loses those consumers’ lifetime value. The firm faces a trade-off between revenue maximization in a single auction (in a static sense) and customer retention across auctions. Controlling the learning opportunities that the customer has alleviates this trade-off for the firm. In our data, the firm can remove the learning stimulus altogether by changing fixed prices, thus resolving the trade-off.

We demonstrate the need to disentangle customer attrition (extensive margin) from strategic learning on the platform (intensive margin) and provide an empirical approach to that end. Customer attrition has recently been studied in the context of eBay auctions with buy-it-now option ([Backus et al., 2022](#)) and in the context of ending a session of chess on a won or lost game ([Avoyan, Khubulashvili and Mekerishvili, 2021](#)). Both papers analyze psychological reasons for consumer attrition, although they do not study behavioral

⁵[Heidhues and Kőszegi \(2018\)](#) discuss a lack of dynamic analysis in behavioral industrial organization, and call for more research on the topic.

mistakes.

The paper proceeds as follows. In Section 2 we discuss the rules of the multi-unit descending auctions and further institutional details. In Section 3 we report on data collection. In Section 4 we describe our data including the empirical evidence on firm behavior. The model of firm incentives is laid out in Section 5. Section 6 discusses our empirical strategy. We report estimates of the bidders' learning response in Section 7. Section 8 concludes.

2. The Multi-Unit Auctions

The seller uses a televised multi-unit descending price auction embedded in hour-long shows to sell consumer goods. Each show consists of several auctions for similar products such as home textiles, men's watches, or jewelry. The average auction lasts about 11 minutes. The seller broadcasts auction shows 20 hours a day, via TV (bids submitted by phone) or online (websites, several apps). At any given time only one auction is held and bids are submitted into the same auction through different channels.

Bidders can also purchase every product up for auction through an online shop at a fixed price. The online shop is available on the same website and the apps that also broadcast the live stream of the auction shows. We therefore view the fixed price as the relevant point of comparison for bids and auction price.

The auction rules ensure that only people who bid above the fixed price (overbid) also pay above the fixed price (overpay). At the beginning of each auction, the auctioneer announces the (initial) number of items to be sold and the auction's starting bid⁶. This starting bid is then gradually lowered over time in discrete increments. Bidders can enter the auction at the current bid and claim one or more units of the good. The auction ends when all units are claimed. All bidders pay the lowest successful bid (uniform pricing rule). Because of this uniform pricing rule, an auction is only overpaid if all bids are overbids.

Shipping costs apply to the fixed price and the auction in the same way. For that reason, we ignore shipping costs in our analysis. Additionally, bidders who bid by phone have to pay a flat fee of one Euro. Since research on shipping costs suggests that this fee is

⁶In their terms and conditions, the firm reserves the right to increase the number of units to be sold even after an auction has started.

likely (at least partially) ignored, we do not include this fee when we calculate overpaying (Hossain and Morgan, 2006). Furthermore, if customers actually internalise the phone fee, we erroneously assign some bidders to the control group and hence, our approach is conservative.

3. Data

We scrape data on bids and products from the seller’s website from October 20, 2016, to January 3, 2019. Since data is removed from the website after some time, we run the scraping script in hourly intervals.⁷

First, we access the schedule in the TV programming section of the website. This schedule gives us information on the show level, such as time and date, product category, and the auctioneer running the show. Second, we collect auction-level data by going through the list of all planned auctions. This list contains an auction ID that we use to scrape bids and bidder nicknames from a separate part of the website. Third, we collect product information from the online shopping section of the website. Most importantly, we collect the fixed price of each product at the time of the auction.⁸

This data collection yields a bidder level panel of 8.48 million bids in more than 69000 auctions spanning over 2 years and 2 months. We use this raw data to calculate several variables, including the auction price (the minimum of the bids), bidder history variables that capture typical behavior and past experience on the auction platform, and dummies indicating whether a bid is an overbid or leads to an overpayment (overpaid).

Table 1 reports summary statistics broken down to the show category level. Naturally, product categories differ with respect to price and quantity. For example, there are more than 7 times as many items sold in the household category than in the watches category. Nevertheless, the two categories are responsible for similar shares of the firms revenue at 20.5% and 18.3% of total revenue made in the auctions. Unfortunately, we could not obtain purchased quantity for the fixed price sales channel, but discussions with the firms management revealed that revenue is overwhelmingly made in the auctions.

⁷Due to a small coding error we did not collect auction shows at 6, 10 and 11 pm. Other than that, we observe all shows and within shows all auctions and bids that took place.

⁸We also collected product ratings, but those are quite sparse at this retailer, so we do not use them.

Show Category	Items Sold	Auction Price	Fixed Price	Revenue Share
Beauty & Wellness	2512517	10.646	15.863	0.170
Leisure & Collecting	20808	54.854	92.173	0.007
Household	2583827	12.497	17.625	0.205
DIY & Gardening	666292	13.534	21.569	0.057
Home Textiles	705398	11.917	16.718	0.053
Fashion & Accessory	987338	22.674	29.302	0.142
Jewelry	674742	42.865	56.709	0.183
Watches	338449	85.259	119.726	0.183

Table 1: Descriptives by Show Category

4. Break in Overbidding and Overpaying

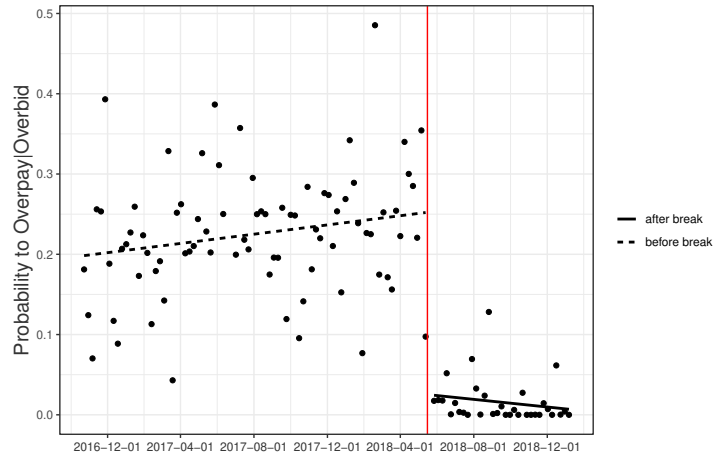
We observe a structural break in the empirical overbidding and overpaying rates in our data. In Panel A of Figure 1 we plot the probability to overpay conditional on overbidding aggregated to weekly averages. Initially, the probability to overpay given one has overbid is roughly 23%, so overbidders are likely to pay for their mistake. Subsequently, we observe a sharp decline in the conditional probability to overpay given one has overbid from roughly 23% to essentially 0%. We determine the exact date of the structural break with a QLR test (Kleiber and Zeileis, 2008).⁹ To illustrate, we add a linear trend on both sides of the structural break.

Table 2 reports summary statistics split by the structural break, since that is the defining feature of our data. Before the break, 17% of all bids are overbids. While the overall probability of overpaying is small at 4%, the probability that this behavioral mistake becomes payoff-relevant is substantially higher. After the break, the probability of overbidding collapses to just below 10%, which lowers the probability of overpaying to essentially 0%. Together, these statistics indicate that before the break consumer learning is a lot more likely than after the break.

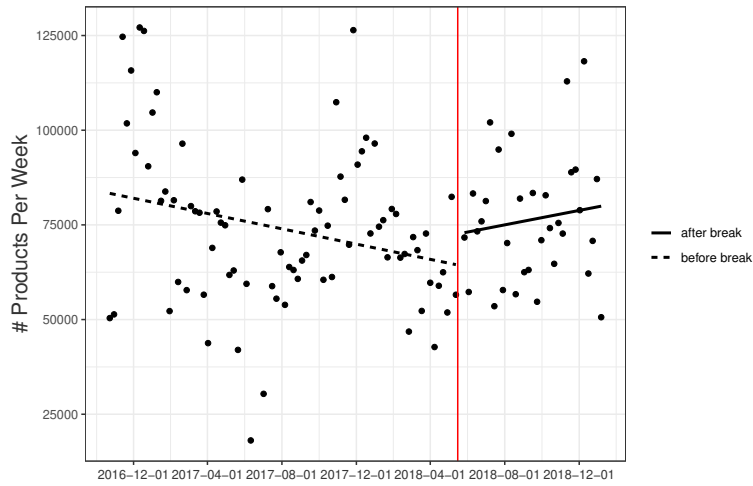
The sharp decline in overpaid auctions coincides with a discrete increase in the number of products sold in each week (Panel B in Figure 1). The number of products primarily increased because the seller conducted more auctions. Fixed prices do not change at the structural break (Figure 8 and 9 in Appendix G).

⁹The most probable break-point is the day with the highest individual test statistic, in our case, the 16th of May 2018. We plot the time series of test statistics in Figure 9 in Appendix G.

Figure 1: Structural Break



(a) Probability to overpay given one has overbid (weekly averages).



(b) Number of Products sold each week.

4.1. Back-of-the-Envelope Calculation

If overpaying causes a demand response, the firm faces a trade-off between extracting overpaying revenue today and foregone revenue tomorrow. Naturally, extracting overpaying revenue comes at little to no operational cost, so for the purpose of our back-of-the-envelope calculation we assume it is profit. Foregone revenue, however, does not equate foregone profit, so we need to adjust foregone revenues with the gross margin. Our back-of-the-envelope calculation suggests that extracting overpaying revenue is suboptimal if

$$\text{Foregone Revenue} \cdot \text{Gross Margin} > \text{Avg. Overpayment.}$$

Table 2: Estimated probabilities of overpaying and overbidding.

	before break (N = 4573854)	after break (N = 1960103)
overbid		
probability	0.17	0.093
average amount	3.6	4.4
median ammount	1.1	1.1
overpaid		
probability	0.039	0.0014
average amount	2.5	2.9
median amount	1.1	0.6
overpaid overbid		
probability	0.23	0.015
auctions		
average duration (minutes)	11.25	11.9
average product price	27.8	28.2
average auction price	21.3	20.8

In our data, overpaying revenue is small at 2.50€ on average (see Table 2). From the available balance sheets, we calculate gross margin at 0.256 and 0.268, in 2014 and 2015 respectively.¹⁰ To be conservative, we use gross margin at 0.256. Thus, if bidders drop out for one year, the foregone revenue threshold that renders extraction of overpaying revenues suboptimal is 9.77€.

Naturally, foregone revenue may be driven by the intensive and the extensive margin. We calculate a *lower* bound by only considering revenue lost from customers dropping out of the customer base. To quantify revenue lost from customer loss, we assume the customer abstains from bidding for at least one year. Thus, multiplying average annual spending with the probability of losing a customer due to overpaying, ϵ , gives the revenue loss due to overpaying attrition. In our data annual average spending is 360€ in 2017 and 383€ in 2018. This is in line with annual spending of over 300€ as claimed in investor presentations by the firm.¹¹ Thus, extracting overpaying revenue is suboptimal if

$$\text{Annual Spending} \cdot \text{Gross Margin} \cdot \epsilon > \text{Avg. Overpayment.}$$

¹⁰The balance sheet data needed to calculate gross margins for later years are unavailable due to a change in reporting format.

¹¹see https://www.1-2-3.tv/uploads/files/2013_06_123tv%20Company%20Profile.pdf
https://www.1-2-3.tv/uploads/files/2012_10_%20123tv%20Das%20Unternehmen.pdf
https://www.1-2-3.tv/uploads/files/PM_123tv_2014_07_01.pdf, accessed 12.01.2022

Plugging in the numbers yields an epsilon-threshold of $\epsilon > 0.027$. Note that we are using the conservative numbers for gross margin (0.256) and annual spending (360€) to arrive at this number.

We emphasize that the back-of-the-envelope calculations are conservative for two more reasons. First, using the annual spending number assumes a lost customer would have stopped purchasing after one year. Given the high rate of repeat purchases, however, it is unlikely that customers only stay one year. Second, assigning the full overpaying revenue to profit may be problematic. In our conversation with the firms management it was hinted that return rates may be higher for overpaid items: Customers can simply return their purchase at the overpaid auction price and - if they so choose - repurchase the item at the fixed price. This would undo the overpaying revenue for the seller and, even worse, impose costs on the seller who has to deal with the returned item. Unfortunately, we are neither able to calculate counterfactual auction prices, nor do we have data on average customer lifetime or return rates.¹²

5. Firm Incentives

The televised auctions are run jointly by an auctioneer, who is on-screen, and a director, who is off-screen. The director has access to rich data such as the number of viewers and revenue broken down by the minute. The director has the authority to "steer" the auctions progress in a number of ways, but crucially he may increase the quantity even after the auction has started.¹³ Management is also aware that overpaying may lead to unhappy customers, since those customers sometimes call the hotline to complain. We take this information seriously and use them in our model to capture the trade-off between overpaying today and customer retention.

We assume every bidder has a latent bid, which is the result of some mapping from values to latent bids. Whether a bidder hands in her latent bid depends on her behavioral type. We assume three types of bidders: overbidders, non-overbidders, and dropouts. An overbidder simply submits his latent bid. A non-overbidder, however, never submits an

¹²We note, however, that the fixed price is lagging the auction in terms of revenue: management told us that they make less than 10% of the revenue through the fixed price offering.

¹³Typical director tasks, for example, include sequencing of the camera feeds or when to show certain TV overlays.

overbid by truncating the bidding function at the fixed price. A dropout, as the name suggests, does not participate in the auction.

Consider an initial overbidder, whose overbid leads to overpayment and thus makes the mistake salient. In our model, learning happens through type changes. An initial overbidder may learn not to repeat her mistake by becoming a non-overbidder. We denote the probability of learning at this intensive margin by ι . Besides learning not to overbid, we also allow for learning at the extensive margin. There are multiple reasons why a bidder may cease to participate in the auctions. These include updated beliefs about the utility from market participation (Backus et al., 2022), updated beliefs about ability (Seru et al., 2010)¹⁴ and consumer antagonism (Anderson and Simester, 2010; Gesche, 2022). We denote the probability of learning at the extensive margin by ϵ

Figure 3 visualises how latent bids, $\beta_{i,t}$, map into actual bids. Even though our analysis does not rely on functional form assumptions, for ease of illustration we suppose that latent bids are uniformly distributed in Figure 3. An overbidder simply hands in her latent bid, which is indicated by the 45° line. Thus, the distribution of actual bids submitted by overbidders is also the uniform distribution. A non-overbidder, on the other hand, never hands in a latent bid larger than the fixed price p_t . Instead, we assume that non-overbidders exactly bid the fixed price in the auction, which is why the bid distribution has a mass point at p_t .

We model intensive margin learning narrowly, in the sense that bidders who experience the consequences of their mistake avoid said mistake in the future. The literature provides evidence for such a learning dynamic, both in controlled lab experiments on auctions as well as in field settings other than auctions. In the lab, bidders adjust their bid in the direction that would have been better in the past (Neugebauer and Selten, 2006). In the field, customers who pay a fee avoid the action that triggered the fee (Haselhuhn et al., 2012; Agarwal et al., 2013; Ater and Landsman, 2013).

Suppose there are $N = o_t + s_t$, bidders in the auction at time t , of which o_t are overbidders and s_t are non-overbidders. All bidders have unit demand and the same latent bid $\beta > p$ for simplicity. Then, the auction price is a function of the number of overbidders and

¹⁴In our context, bidders may think they are irredeemably bad or unlucky at bidding, so they should stop bidding altogether.

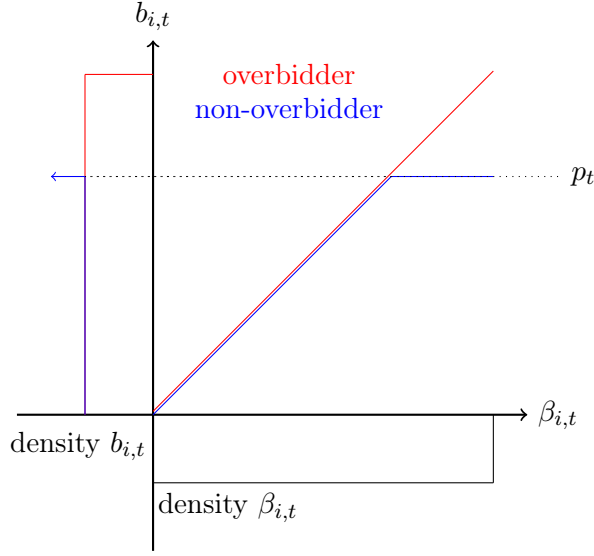


Figure 3: Bids as a function of latent bids. Marginal distribution of bids and latent bids for uniformly distributed latent bids.

non-overbidders and the auction quantity.

We present a simplified version of our empirical model in Section 6.1 in Appendix A to illustrate the seller’s incentives arising from changes in type from overbidder o_t to non-overbidder s_t or dropouts. We simplify the analysis by assuming there are $N = o_t + s_t$ bidders in the auction at time t and that all bidders have unit demand and the same latent bid β .¹⁵ To make the case interesting, we assume that the latent bid is larger than the fixed price $\beta > p$. Then, the auction price is a function of the number of overbidders and non-overbidders and the auction quantity.

$$p_a(q_t, o_t, s_t) = \begin{cases} \beta & q_t \leq o_t \\ p & o_t < q_t \leq o_t + s_t \\ 0 & o_t + s_t < q_t \end{cases}$$

Choosing a sufficiently small quantity, $q_t < o_t$, will ensure that all bids in the auction are overbids and thus the auction ends in overpayment at $p_a = \beta > p$. A larger quantity, $o_t < q_t \leq o_t + s_t$ will ensure that non-overbidders also bid in the auction and thus the

¹⁵We make the simplifying assumptions here to clearly state the main point, but we drop them in our empirical analysis.

auction price will realize exactly at the fixed price $p_a = p$. Finally, setting quantity larger than the number of participants leads to an auction price of 0, which is never optimal.

Appendix A presents this model in more detail. Here we restrict ourselves to discussing the cases where the firm views fixed prices as exogenously given and chooses quantity and when both variables are chosen simultaneously. This approach is motivated by the fact that fixed prices were only changed after a new CEO took office, whereas they remained unchanged at the structural break in our data. This is in line with a model of learning through noticing where a decision maker may not fully optimize because he fails to notice an important feature of the optimization problem (Hanna et al., 2014).

Policy 1: Choosing Quantities As discussed above, the director may increase quantity during the auction, while taking fixed prices as given. Offering extra units in the auction leads to a downward move along the demand curve in each auction, which lowers prices.¹⁶ Overpaying may also be reduced without increasing overall quantity, by shifting quantity from a non-overpaid auction to an overpaid auction, where both auctions sell the same good.

We find that in the case of exogenous fixed prices, the seller's optimal quantity choice depends on how large overpaying $\beta - p$ is compared to the discounted revenue lost due to extensive margin learning ϵ as given by Observation 1.¹⁷

Observation 1. *Suppose the seller can only choose quantity q and the fixed price p is exogenously given. If $\beta - p \geq \frac{\delta}{1-\delta} \cdot p \cdot \epsilon$ the profit maximizing quantity is $q_t = o_t, \forall t$. If $\beta - p < \frac{\delta}{1-\delta} \cdot p \cdot \epsilon$ the profit maximizing quantity is any $q_t \in (o_t, o_t + s_t], \forall t$.*

If overpaying is larger than the revenue loss due to extensive margin learning, then the seller chooses a small quantity and the auctions end in overpaying, $p_a = \beta > p$. If the reverse is true, the seller prefers not to extract the extra revenue through overpaying to preserve future demand and the auctions end at the fixed price. In this setting with exogenous fixed prices, the sellers can only shut down learning by lowering the auction price and thus foregoing the extra revenue. Note that the optimal decision does not depend on

¹⁶While we observe the number of units sold in each auction, we do not observe the number of units that were originally planned for the auction. Unfortunately, this means we do not know which auction increased supply dynamically during the auction.

¹⁷The proof is in Appendix B.

intensive margin learning as taking away the opportunity to exploit overbidding tomorrow is not an effective deterrent against exploitation today.¹⁸

Policy 2: Choosing Fixed Prices and Quantities Since the seller makes most revenue in the auctions, the fixed-price outside option primarily acts as a reference price. Consequently, the seller can use adjustments of this reference price as a second instrument to avoid overpaying.

After our sample ends, the seller raised fixed prices. Initially, the fixed price is set (presumably very) high and the auction high bid always undercuts the fixed price. After an auction ends the fixed price is lowered to the auction price plus a small increment for 24 hours. This strategy gives potential customers who missed the auction the chance to purchase the good at a price below the recommended retail price. For example, an item may be offered at its recommended retail price of € 30, while the auction starts at € 20. The auction price may realize at € 12, and the fixed price falls to € 15 for 24 hours.¹⁹ Since we did not collect data after the policy change, we cannot check if the fixed prices rose on average. By construction, the new policy, however, makes overbidding and overpaying impossible.

Observation 2. *If sellers can choose p and q , they maximize their profits by setting $p > \beta$ and $q_t \in (o_t, o_t + s_t]$, $\forall t$.*

Observation 2 states that the firm can circumvent revenue losses due to learning, by setting high fixed prices. Instead of addressing the cause of consumer learning (high auction prices) the firm can remove the stimulus (overpaid auctions) by adjusting the fixed price. This policy increases prices without changing quantity and, thus, surplus is redistributed to the firm.

6. Empirical Strategy

We take a sufficient statistics approach to quantify the extensive and intensive margin learning rates. We present a structural equation model of how consumers bid at auction

¹⁸Observation 1 closely resembles our back-of-the-envelope calculation in Section 4.1

¹⁹The short-term rebate on the fixed price may be seen as the price discovery aspect of an auction.

and how the firm runs the auction. We derive treatment effects that we can estimate from our data. We explain how to represent our structural equation model as a Directed Acyclic Graph (DAG)²⁰. We use the DAG representation of our model to show that our treatment effects are identified. The DAG representation allows us to derive the conditional independence assumption and the control variables required for identification from our empirical model using Pearl (2009)’s back-door criterion. We discuss interpretation of our treatment effects and the assumptions needed to recover extensive and intensive margin learning rates, ϵ and ι . The strength of this approach is that the treatment effects are valid under weaker assumptions and we only need additional assumptions on a part of our model to back-out the structural parameters of interest.

6.1. Empirical Model

We use a general version of our model from Section 5. This model includes bid heterogeneity and firm behavior as a function of exogenous shocks. We model bidder learning parametrically and avoid parametric assumptions on firm behavior and the latent-bid distribution.

We introduce new notation to describe bidder behavior. We focus on a specific bidder i whose first overbid is at time $t \in \{1, \dots, \infty\}$. We observe this bidder from their first overbid until the end of our sample. Since this time differs between bidders, we aggregate bidder’s outcomes over a standardized period of time. We assume that a bidder i ’s own bids today have no effect on bids she faces in the future, so that we can treat different auctions as independent observations. Thus, the stable unit treatment value assumption is satisfied and, consequently, the behavior of all other bidders is uncorrelated across auctions. To set notation, we collect all other bidders in the set $J_t = \{1, \dots, N_t\}$.

Our model uses exogenous shocks to model empirically relevant sources of heterogeneity. We assume that in each period, a fixed price shock \tilde{p}_t , an auction quantity shock \tilde{q}_t and an individual specific time-varying shock $v_{i,t}$ realize as independent draws from a continuous distribution. We assume that these shocks are independent from each other. We model individual specific unobserved heterogeneity with the time-constant variable u_i , which in our model realizes before any choices are made and is independent across individuals.

²⁰Sometimes also called Causal Graph

Auction-specific characteristics are denoted by A_t .

As before, we separate the bidding process into latent (unmodeled) bids and a bidder type (overbidder, non-overbidder, dropout) that determines the submission of these bids. We refrain from modelling the process of mapping valuations into latent bids because of the highly complicated nature of our dynamic auction. For example, bidding would depend on the observed pace of other bidders bids being submitted, which we neither observe, nor do we think it adds much to our analysis. We are mainly interested in whether bidders learn to avoid a specific mistake or drop out of the market. Our approach assumes that other differences in the bidding process between auctions are independent of the learning margins we describe.

We let the individual latent bid depend on auction-specific characteristics A_t , the individual specific time-varying shock $v_{i,t}$ and the time-constant unobserved heterogeneity u_i . These shocks are i.i.d. from a continuous distribution. We denote the latent bid of the individual in question by $\beta_{i,t} = \beta(A_t, u_i, v_{i,t})$ and the set of latent bids by the other bidders by $\beta_{-i,t} = \{\beta(A_t, u_j, v_{j,t})\}$, where $j \neq i$. Together, the latent bids by bidder i and competing bidders $j \neq i$, $\beta_{i,t}$ and $\beta_{-i,t}$, represent the the latent demand of all bidders in auction t .

The dependence of β_{it} on auction characteristics, A_t , models that different individuals might be interested in different auctions. A special case of this is that most individuals do not participate in an auction. In this case their latent bid is 0. The dependence on u_i models that different bidders might differ in the amount they usually bid. The time-varying individual specific shock v_{it} models the main source of heterogeneity for a specific bidder across auctions.

According to the model in Section 5 and our discussion with management, the firm targets its quantities and fixed prices to latent bidder demand. While our empirical analysis focuses on the period before the structural break when the firm likely does not behave optimally, we still allow for the fixed price p_t and the auction quantity q_t to depend on latent bids β_{it} and β_{-it} . Since we do not specify the parametric form of this dependence, we allow for optimal as well as non-optimal firm behavior in our empirical analysis. Further the firm might tailor fixed-prices to auction characteristics, e.g. the type of products on sale. We model this by letting the fixed price and the auction quantity depend on auction-characteristics A_t . Thus,

auction quantity (q_t) and fixed-price (p_t) may depend on these quantities as well as their specific exogenous shock (\tilde{q}_t and \tilde{p}_{st}).²¹

$$\begin{aligned} p_t &= c(\tilde{p}_t, A_t, \beta_{it}, \beta_{-it}) \\ q_t &= d(\tilde{q}_t, A_t, \beta_{it}, \beta_{-it}) \end{aligned}$$

As in Section 5, we assume that a bidder's bid depends on their type $\theta_{i,t}$ and their latent bid β_{it} . The bidder's type at time t is $\theta_{i,t} \in \{o, s, d\}$, where we denote overbidders by o , non-overbidders by s , and dropouts by d . Since we consider bidders after their first overbid, we only select overbidders. Overbidders always bid their latent bid $\beta_{i,t}$, while non-overbidders wait until the price drops below the fixed price, that is they bid $\min\{\beta_{i,t}, p_t\}$. Bidders who drop out always bid 0. We summarize this behavior in the following bid function.

$$b_{i,t} = f(p_t, \beta_{i,t}, \theta_{i,t}) = \begin{cases} \beta_{i,t} & \text{if } \theta_{i,t} = o \\ \min\{\beta_{i,t}, p_t\} & \text{if } \theta_{i,t} = s \\ 0 & \text{if } \theta_{i,t} = d \end{cases}$$

In our empirical model, we allow for heterogeneity in learning rates. Specifically, we assume that learning rates are time constant, but may differ between individual bidders. Overpaying turns overbidders into non-overbidders with probability ι_i (intensive margin), and makes them drop out with probability ϵ_i (extensive margin). We allow for dependence between these treatment effect parameters and bidder specific shocks u_i . In Section 6.2.2, we assume, for technical reasons, that learning rates are a function of individual characteristics, u_i . This assumption is natural as it still allows bidders of different individual characteristics, u_i , to exhibit different learning rates.

We express the auction's outcome from the perspective of bidder i in terms of two order statistics of all other bids (the set $\beta_{-i,t}$): the q_t -highest and the $q_t - 1$ -highest rival bid,

²¹Think of these shocks as supply shocks that are not captured by auction characteristics A_t .

which we denote by $b(q_t)$ and $b(q_t - 1)$, respectively. The q_t -highest rival bid determines if bidder i wins, and the $q_t - 1$ -highest rival bid influences the auction price. The auction ends when all products are sold, and the lowest successful bid determines the price. Bidder i loses the auction if all q_t units are sold to bidders in J_t , that is, bidder i loses if $b_{i,t} < b(q_t)$. Conversely, bidder i 's bid is successful if $b_{i,t} > b(q_t)$. In this case, there are $q_t - 1$ units that remain for the competing bidders included in J_t . The lowest successful bid is then either by the bidder in question or the lowest successful bid by the other bidders ($b(q_t - 1)$). If bidder i places a winning bid, the auction price is $\min\{b_{i,t}, b(q_t - 1)\}$. We define an overbid as a bid that is strictly larger than the fixed price $b_{i,t} > p_t$. Similarly, we call a bid a non-overbid when it is strictly smaller than the fixed price $b_{i,t} < p_t$. An auction is overpaid when all bids are overbids and that means that the auction price is higher than the fixed price, or $\min\{b_{i,t}, b(q_t - 1)\} > p_t$.

While we use our parametric assumptions on bidding behavior to interpret our treatment-effects, we do not need parametric assumptions to estimate these effects. For this purpose, we summarise our model as a system of non-parametric structural equations. Each structural equation expresses a left-hand side variable in terms of other variables and exogenous shocks. This model is non-parametric in the sense that we do not use any functional form assumptions on the right-hand side.

$$p_t = c(\tilde{p}_t, A_t, \beta_{it}, \beta_{-it}) \quad (1)$$

$$q_t = d(\tilde{q}_t, A_t, \beta_{it}, \beta_{-it}) \quad (2)$$

$$\beta_{i,t} = \beta(A_t, u_i, v_{i,t}) \quad (3)$$

$$b(q_t) = f(q_t, p_t, \beta_{-i,t}, \theta_{-i,t}) \quad (4)$$

$$overbid_{i,t} = g(\beta_{i,t}, b(q_t), b(q_t - 1), p_t, \theta_{i,t}) \quad (5)$$

$$non - overbid_{i,t} = v(\beta_{i,t}, b(q_t), b(q_t - 1), p_t, \theta_{i,t}) \quad (6)$$

$$overpaid_{i,t} = v(overbid_{i,t}, b(q_t), p_t) \quad (7)$$

We briefly go through each equation and relate it to the previous discussion. We explicitly introduced Equation 1 to 3 in the preceding section. Equation 1 and 2 describe

the information set of the firm when setting auction parameters. One model that fits these equations is the simplified model from Section 5. Since these equations do not assume any structure, they nest all (optimal and non-optimal) firm policies that condition fixed prices and quantities on a signal of latent demand, $\beta_{i,t}$ and $\beta_{-i,t}$, and auction-specific characteristics, A_t . Equation 4 summarizes the order statistics of the rival bidder's bids, that determines whether a bid is a winning bid. Equations 5 to 7 apply the auction's rule to this section's expression of bidder's successful bids ($b_{i,t}$). We will introduce parametric versions of Equations 5 to 7 in the next section. According to Equation 5, an overbid is an observed variable (as governed by the order statistic). We will refer to overbids regardless of observation status as latent overbids throughout the text.

6.2. Interpretation of Treatment Effects

Recall that we consider the first overbid for each bidder in our sample (if there is any). That means we look at overbidders and would like to quantify type changes to non-overbidder or dropout. A bidder who changes type from overbidder to dropout just leaves the market. Consequently, we should observe fewer overbids, as well as fewer non-overbids in this case. A bidder who changes from overbidder to non-overbidder type, however, just avoids overbids in the future and bids at the fixed price whenever the latent bid is a latent overbid. Hence, bidders who become a non-overbidder bunch at the fixed price. We exclude the fixed price by focusing on strictly defined non-overbids. Thus, a first starting point to test for extensive margin learning is to estimate the treatment effect of overpaying on strict non-overbids.

To see this in more detail, reconsider Figure 3 from Section 5, where we depict bids as a function of latent bids and bidder type. As an example we depict a marginal uniform distribution of latent bids underneath the x-axis. The figure also shows the resulting marginal density of bids to the left of the y-axis. Non-overbidders bid at the fixed price for all latent bids larger than the fixed price, so the density of bids has a mass point at the fixed price and has zero mass at bids above the fixed price. Strictly below the fixed price, the density is identical to that of overbidders. A dropout maps all latent bids to the zero bid (not depicted in Figure 3).

Figure 3 illustrates what we can learn about latent changes in type from changes in the observed bid distribution. Bids strictly below the fixed price (strict non-overbids) are

submitted by overbidders and non-overbidders and bids strictly above the fixed price (strict overbids) are only submitted by overbidders. Latent bids below the fixed price (p_t) directly translate into observed bids. Latent bids above the fixed price bunch at the fixed price for non-overbidders and directly translate into observed bids for overbidders. Bids at the fixed-price are composed of bunched overbids by non-overbidders and latent bids at the fixed price by both non-overbidders and overbidders.

Since strict overbids are only submitted by overbidders a decrease in these bids indicates a reduction in overbidders. Intensive as well as extensive margin learning can cause such a decrease. Since strict non-overbids are submitted by overbidders as well as non-overbidders, a decrease in these bids indicates extensive margin learning (overbidders leaving the auction). Bids directly at the fixed price increase when there are more non-overbidders and decrease when there are more overbidders. Consequently these bids increase with intensive margin learning and decrease with extensive margin learning. We focus on treatment effects of overpaying on strict overbids and strict non-overbids, in order to recover the extensive and intensive margin learning parameters, ϵ_i and ι_i .

The only way to observe a non-overbid is a latent non-overbid ($\beta_{i,t} < p_t$), which is successfully submitted ($\beta_{i,t} > b(q_t)$), either by an overbidder or a non-overbidder. Thus, the treatment effect of overpaying on non-overbids in the next period is the expected extensive margin learning parameter scaled by the probability of a successful strict non-overbid. We calculate this effect conditional on u_i . This conditioning renders ϵ_i and the latent bid independent and thus allows the factorization. Since we are interested in learning as a response to overpaying, we also look at the most narrow way to avoid overpaying, which is learning not to overbid. In our model only overbidders submit overbids. Thus, a type transition from overbidder to non-overbidder, as well as dropping out reduces overbidding. Consequently, the treatment effect of overpaying on overbidding in the next period is given by the sum of learning at both margins multiplied by the probability of a latent successful overbid.

Proposition 1. *The treatment effects of overpaying in t on the number of non-overbids*

and number of overbids in $t + 1$ take the following form.

$$E[TE_{non-overbid}^{t+1}|u_i] = -E[\epsilon_i|u_i]\mathbb{P}(p_{t+1} > \beta_{i,t+1} > b(q_{t+1})|u_i)$$

$$E[TE_{overbid}^{t+1}|u_i] = -E[\epsilon_i + \iota_i|u_i] \cdot \mathbb{P}(\beta_{i,t+1} > p_t \wedge \beta_{i,t+1} > b(q_{t+1})|u_i)$$

Proposition 1 reminds us that the treatment effect of overpaying on the number of overbids is a function of learning rates at the extensive margin and intensive margin. To disentangle the two margins it would be sufficient to divide the treatment effect by the probability of observing a latent overbid (non-overbid) and taking the difference of the two treatment effects. It turns out that the probability of observing a latent overbid (non-overbid), is simply the potential outcome of the untreated.²² In Subsection 6.2.2, we explain how the analysis is complicated by the fact that we need to pool observations in order to estimate the treatment effects. Before we get to pooling of observations we explain, however, how a shift in the latent bid distribution as a result of overpaying would change our results.

6.2.1. Shift in Latent Bid Distribution

Bidders may react to overpaying in more general ways than the two adjustments we impose, truncating the bid function at the fixed price or dropping out. Suppose, for example, that a bidder, who overpays, may, in addition to truncating their bid function at the fixed price, increase the shading of their bid. Shading one's bid as a reaction to overpaying is a natural adjustment: the auction price a bidder has to pay is weakly lower than her bid and so shading said bid reduces payments.

We show an example of such a bid shading adjustment in Figure 4. Instead of simply truncating, the additional bid shading shifts the bid function downward (in red). This is, of course, a departure from the behavioral type change to non-overbidder we impose (in blue), so it's worth reviewing how this would impact our results. The treatment effect of overpaying on future overbids is unchanged since the bidder truncates her bid function at

²²We can calculate this potential outcome of the untreated from our regression in Section 7 by setting the treatment dummy to 0 and all other variables to the sample mean.

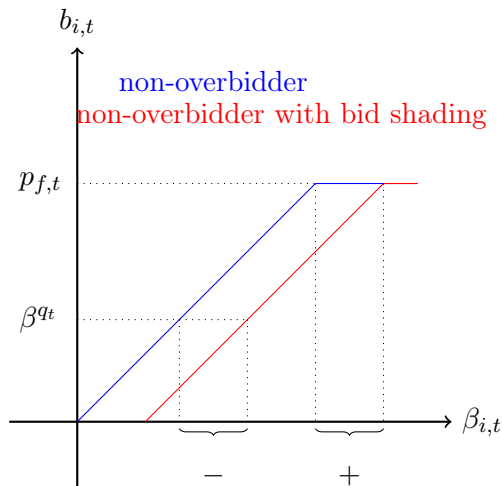


Figure 4: Truncating and shading bids as response to overpaying. Some latent bids are shaded out of the auction ($-$) and some that were otherwise mapped to the fixed price are now shaded in to the non-overbid region below the fixed price ($+$).

the fixed price, $p_{f,t}$, as before. A difference may arise when we consider the treatment effect on non-overbids, which are strictly smaller than the fixed price. Some latent bids that are mapped to the fixed price are now handed in below the fixed price due to the additional bid shading. This is the interval marked with $+$ in Figure 4. On the other hand, some bids close to the lowest competing bid in the same auction, β^{qt} are pushed below this lowest competing bid and thus are not handed-in in the auction — these bids are losing because of the bid shading. In Figure 4 these latent bids are in the interval marked with $-$.

How well our approach using two behavioral types to capture learning works in the presence of unmodeled bid shading depends on the relative size of the intervals marked with $+$ and $-$. In Figure 4, we depict the case of a linear bid function (and a uniform distribution of latent bids). As it turns out, the intervals $+$ and $-$ cancel exactly. In this case, imposing the assumed changes in behavioral types is a good approximation of this more general adjustment of truncating and shading the bid function.

Proposition 2 discusses the effect of bid shading on our estimation strategy for arbitrary latent bid distributions. Here, the difference in shaded, \mathbb{P}' and non-shaded bid distributions, \mathbb{P} , may be negative or positive. If bid shading shifts more bids below the fixed price than it pushes out of the auction, the difference of the shaded and non-shaded distributions is positive and thus we would underestimate the extensive margin using our two-type approximation. If, on the other hand, bid shading prices more latent bids out of the auction

than it pushes bids below the fixed price, the difference in bid distributions is negative, $\mathbb{P}'(\cdot) - \mathbb{P}(\cdot) > 0$. In this case, our approach mis-attributes some of the bid shading to dropouts and thus we would overestimate the extensive margin.

Proposition 2. *If non-overbidders shift their distribution of latent bids compared to overbidders, the treatment effect of overpaying on non-overbids in the next period is given by,*

$$E[TE_{non-overbid}^{t+1}] = -E[\epsilon_i|u_i] \cdot \mathbb{P}(p_{t+1} > \beta_{i,t+1} > b(q_{t+1})|u_i) \\ + E[l_i|u_i] \cdot \underbrace{[\mathbb{P}'(p_{t+1} > \beta_{i,t+1} > b(q_{t+1})|u_i) - \mathbb{P}(p_{t+1} > \beta_{i,t+1} > b(q_{t+1})|u_i)]}_{\text{Shift in latent bid distribution}},$$

where P' is a probability calculated from the latent bid distribution of non-overbidders and P is a probability calculated from the latent bid distribution of overbidders.

The proof of this result is in Appendix C.

In the eyes of the firm's management, these differences are somewhat muted. In one case the firm loses a sale because a consumer drops out of the market. In another case, the firm loses a sale because a consumer shades her bid so much that she does not bid in the auction anymore — effectively dropping out. In particular, a consumer who always bid very low, say 1€, is observationally indistinguishable from a consumer who does not take part in the auction.

6.2.2. Pooling of Observations Over a Period of Time

Our analysis, so far, focuses on treatment effects of overpaying in t for outcomes of interest in $t + 1$, i.e. for the next auction. Bidders do not generally take part in every auction, so we need to pool observations over a period of time in order to estimate treatment effects. To calculate a pooled treatment effect, we simply sum over the treatment effect in $t + k$, where k covers the period of time over which we are pooling observations. The derivation of the period $t + k$ treatment effect is largely similar to the treatment effect in $t + 1$, but for the possibility of subsequent treatments.

Subsequent treatments may occur in both the treatment and the control group. In the control group, subsequent treatments are likely since overbidders did not have the

opportunity to learn in period t (they are the control after all). In the treatment group, subsequent treatments may happen, whenever bidders do not learn from their initial treatment — either by chance or because their learning rate is small. The subsequent treatments in control and treatment group bias our results in opposite directions: including treated bidders in the control group biases downwards, while multiple treatments in the treatment group increases the likelihood of finding a learning effect.

In Proposition 3 we show that the treatment effects are attenuated rather than exacerbated by the possibility of subsequent treatments due to pooling over a period of time. The intuition behind Proposition 3 is that subsequent treatments are more likely to occur in the control group than in the treatment group. This is the case, as a bidder in the treatment group can only receive an additional treatment if he fails to learn from the first treatment. No such condition applies for subsequent treatments of the control group and, thus, our results are conservatively estimated.

Proposition 3. *Let $p_{l,k}$ denote the probability that a bidder changes his type from overbidder to dropout because of a treatment in subsequent periods $t + 1$ to $t + k$. Similarly, let $p_{s,k}$ denote the probability that a bidder changes type from overbidder to non-overbidder due to a treatment in periods $t + 1$ to $t + k$. Then, the treatment effect of the initial treatment in period t on non-overbids and overbids in period $t + k$ is given by the following expressions.*

$$E[TE_{non-overbid}^{t+k}|u_i] = E[-\epsilon_i E[p_{l,k}|\epsilon_i, \iota_i, u_i]|u_i] \mathbb{P}(p_{t+k} > \beta_{i,t+k} > b(q_{t+k})|u_i)$$

$$E[TE_{overbid}^{t+k}|u_i] = E[-(\epsilon_i + \iota_i) E[p_{l,k} + p_{s,k}|\epsilon_i, \iota_i, u_i]|u_i] \mathbb{P}(\beta_{i,t+k} > p_{t+k} \wedge \beta_{i,t+k} > b(q_{t+k})|u_i)$$

The proof of this result is in Appendix D. Note that the $t + k$ period treatment effects in Proposition 3 are identical to the corresponding $t + 1$ period treatment effects in Proposition 1, but for the scaling factor $E[p_{l,k}|\epsilon_i, \iota_i, u_i]$ and $E[p_{l,k} + p_{s,k}|\epsilon_i, \iota_i, u_i]$, respectively. Since these factors are conditional probabilities, they are (at least weakly) smaller than 1 and thus the treatment effects in later periods are attenuated rather than exacerbated by subsequent treatments due to pooling of observations.

To back out the extensive and intensive margin learning parameters, ϵ_i and ι_i , from the pooled treatment effects, we assume that individual heterogeneity in learning rates is

a function of the unobserved individual characteristics, u_i . This restricts learning rates to be homogeneous for individuals with the same unobserved individual characteristics. Formally, we denote $E[\epsilon_i|u_i] = \epsilon_u$ and $E[\iota_i|u_i] = \iota_u$ and we call it heterogeneous-learning-rates assumption. We need to make this technical assumption since our pooled treatment effects do not factorize otherwise. The heterogeneous-learning-rates assumption is natural, since it still allows heterogeneity in learning rates along the heterogeneity in individual characteristics, u_i .

Recall, that the treatment effect in Proposition 1 is just the conditional expectation of the extensive margin learning parameter, $E[\epsilon_i|u_i]$, scaled by the probability to observe a latent non-overbid. Untreated individuals are still participating in the auctions and submit their overbids, so dividing our treatment effect by the potential outcome of the untreated $E[\text{non-overbid}_t^{t+1}(0)|u_i]$, yields the conditional expectation of the extensive learning parameter.

$$\frac{E[TE_{\text{non-overbid}}^{t+1}|u_i]}{E[\text{non-overbid}_t^{t+1}(0)|u_i]} = \frac{-E[\epsilon_i|u_i]\mathbb{P}(p_{t+1} > \beta_{i,t+1} > b(q_{t+1})|u_i)}{\mathbb{P}(p_{t+1} > \beta_{i,t+1} > b(q_{t+1})|u_i)} \quad (8)$$

$$= -E[\epsilon_i|u_i] \quad (9)$$

It remains to account for the pooling of observations to estimate the effect. In Proposition 3, we provide expressions for the treatment effect in some subsequent auction $t+k$. Summing over these treatment effects gives us the treatment effect that we estimate from pooled data. Proposition 4 provides these sums over the treatment effects on non-overbids and overbids divided by the appropriate potential outcome as shown in Equation 8. The proof is in Appendix E.

Lemma 4. *Suppose individuals are treated at time $t \in \{1, \dots, \infty\}$ and we aggregate our treatment effects over the following $k \in \{1, \dots, \infty\}$ periods. Then the treatment effects divided by the potential outcomes are given by the following expressions:*

$$\begin{aligned} \frac{\sum_{m=0}^k E[TE_{\text{non-overbid}}^{t+m}]}{\sum_{m=0}^k E[\text{non-overbid}_t^{t+m}(0)|u_i]} &= \frac{\sum_{m=0}^k E[-\epsilon_i E[p_{l,m}|\epsilon_i, \iota_i, u_i]|u_i]}{\sum_{m=0}^k E[E[p_{l,m}|\epsilon_i, \iota_i, u_i]|u_i]} \\ \frac{\sum_{m=0}^k E[TE_{\text{overbid}}^{t+m}]}{\sum_{m=0}^k E[\text{overbid}_t^{t+m}(0)|u_i]} &= \frac{\sum_{m=0}^k E[-(\epsilon_i + \iota_i) E[p_{l,k} + p_{s,k}|\epsilon_i, \iota_i, u_i]|u_i]}{\sum_{m=0}^k E[E[p_{l,k} + p_{s,k}|\epsilon_i, \iota_i, u_i]|u_i]}. \end{aligned}$$

Unfortunately, the expressions in 4 do not immediately simplify because the probabilities of subsequent treatment ($p_{l,k}$ and $p_{s,k}$) are functions of the corresponding learning parameters (ϵ_i and ι_i). Consider, for example, the expression $E[\epsilon_i E[p_{l,m}|\epsilon_i, \iota_i, u_i]|u_i]$, where $E[\epsilon_i|u_i]$ does not factor out as ϵ_i and $E[p_{l,m}|\epsilon_i, \iota_i, u_i]$ are dependent.

Assuming that ϵ_i is a function of the individual characteristics u_i solves this problem and allows us to recover ϵ_i from the pooled treatment effects. This assumption effectively means that we are restricting individual heterogeneity in learning rates ϵ_i and ι_i to be captured by heterogeneity in individual characteristics u_i . In other words, we have to assume that learning rates are homogeneous for every value of individual characteristics, u_i . As a short hand for this assumption we write $E[\epsilon_i|u_i] = f(u_i) = \epsilon_u$. Corollary 5 shows that using this assumption our expressions from Proposition 4 simplify and we can recover the learning parameters ϵ_u and ι_u from our treatment effects.

Proposition 5. *Under the heterogeneous-learning-rates assumption, heterogeneity in learning is restricted to the heterogeneity in individual characteristics, that is, $E[\epsilon_i|u_i] = \epsilon_u$ and $E[\iota_i|u_i] = \iota_u$. Then, we can recover the learning parameters ϵ_u and ι_u from the pooled treatment effects.*

$$\frac{\sum_{m=0}^k E[TE_{non-overbid}^{t+m}]}{\sum_{m=0}^k E[non-overbid_t^{t+m}(0)|u_i]} = \epsilon_u$$

$$\frac{\sum_{m=0}^k E[TE_{overbid}^{t+m}]}{\sum_{m=0}^k E[overbid_t^{t+m}(0)|u_i]} = \epsilon_u + \iota_u$$

Note that the probabilities of subsequent treatments $p_{l,k}$ and $p_{s,k}$ depend on the time period we pool over. Thus, violations of the assumption in Corollary 5 should lead to incongruous results, when we pool over different time periods. Indeed, in Section 7 we show that our results do not depend on the pooling period used.

6.3. Identification of Treatment Effects

After clarifying the connection between estimable treatment effects and the learning parameters of our underlying model in the previous section, we now turn to the identification of

those treatment effects. Recall that a bidder is assigned to the treatment (control) group when his first overbid (did not) lead to overpayment and that whether an overbid leads to overpayment is entirely determined by the rival bidders in that auction. A remaining concern with this design is selection into treatment. This is an issue if, for example, bidders who learn well select into watch auctions, while bidders who do not learn well spread out evenly over all product categories. For a rigorous analysis of this logic, we represent our empirical model in Equation 1 to 7 as a causal Directed Acyclic Graph (DAG).²³

The DAG representation illustrates the causal relationships implied by the structural equations model and allows us to compute a set of control variables to satisfy the conditional independence assumption required for identification. In our case, identification depends in part on unobserved bidder characteristics, so we conclude with a discussion of how we can implement our identification strategy using past bidder behavior as a proxy for this unobserved variable.

In a DAG, a directed edge (an arrow) indicates a causal relationship, that is, the node where the arrow originates is a cause of the node that the arrow points to. For example, if we draw an arrow from $overpaid_{i,t}$ to $overbid_{i,t+1}$, we show that our model allows for a causal effect of overpaying on overbidding in a subsequent auction. In our context, the fact that DAGs do not contain any cycles has an economic interpretation: bidders are myopic. Otherwise, future auctions would influence bidding behavior in today's auction, which would lead to a cycle in our graph. This assumption is in line with other behavioral economics auction papers such as [Malmendier and Lee \(2011\)](#). We try to explain the theory on DAGs as we go along.²⁴

To generate a DAG from our empirical model in Section 6.1, we go through each equation and draw an edge from each right-hand side variable to each left-hand side variable.²⁵ We leave out exogenous shocks for ease of exposition²⁶ and draw boxes around variables that are observable. The procedure results in the DAG depicted in Figure 5.

We use $y_{i,t+1}$ as a stand-in for the outcomes we are interested in: revenue and number of

²³To put it precisely, we interpret our Structural Equation Model (SEM) in Section 6.1 as a Structural Causal Model (SCM).

²⁴For a gentle introduction, see chapter 3 of [Cunningham \(2021\)](#).

²⁵See, for example, [Peters, Janzing and Schölkopf \(2017\)](#) for a more complete treatment of the connection between DAGs and Structural Causal Models.

²⁶This is without loss of generality because these shocks are exogenous by assumptions, so no edges point to these nodes.

overbids and non-overbids in subsequent auctions. We focus on time period t and display arrows pointing from t to $t + 1$ only in a stylized way. In particular, the path $u_i \rightarrow y_{i,t+1}$ abstracts from the fact that this causal relationship is again channeled through the bidding process. This simplification is without loss of generality, since u_i is the only connection between behavior in t and $t + 1$. We also abstract from bidding behavior before t . We restrict our data set to behavior after the first overbid in t . This restriction selects only bidders who are overbidders in t and thus, there is no remaining variance in the bidders behavioral type, $\theta_{i,t}$, and we can omit it from the DAG.

In a DAG, the paths where all arrows point from the treatment to the outcome variable are called front-door paths, or causal paths. This is the causal relationship of interest, in our case $overpaid_{i,t} \rightarrow y_{i,t+1}$. There are also paths from the treatment to the outcome, where at least one arrow points in the opposite direction, called back-door paths. In our case, all other paths from treatment to outcome are back-door paths, since every path other than the causal path starts with an arrow pointing to $overpaid_{i,t}$ (instead of originating from it). The main idea of proving identification in a DAG is to select control variables to block all back-door paths.

Panel A of Figure 5 shows the origin of our causal graph. The arrows in blue encode institutional knowledge about our setting. For example, the director of the auction sees latent demand and can choose quantity accordingly, so there are arrows from the latent bids $\beta_{i,t}$ and $\beta_{-i,t}$ to quantity q_t . Similarly, the seller incorporates that auction characteristics will have an impact on demand, quantity and fixed prices when planning the auction so there are arrows from auction characteristics, $auction_t$, to latent demand, $\beta_{i,t}$ and $\beta_{-i,t}$, quantity, q_t , and the fixed price $p_{f,t}$.

The arrows in violet depict the auction rules, namely uniform pricing and our definitions of overbidding and overpaying. The order statistic $\beta_{-i,t}^{(q_t)}$ determines winning bids and what price winning bidders have to pay. As bidder i has to beat the q_t -highest rival bids to win the auction, the order statistic has arrows incoming from q_t and $\beta_{-i,t}$. We only observe winning bids and consider initial overbids for each bidder, so there is an arrow from the order statistic to $overbid_{i,t}$. Together with the fixed price it is determined whether the overbid leads to overpayment.

Finally, the arrows in orange depict substantial economic assumptions. Considering only

initial overbidders i means we can omit the behavioral type of those bidders from the DAG. Rival bidders, however, may be non-overbidders, so the fixed price has an influence in whether rival bidders hand in their latent bids. Thus, we draw an arrow from the fixed price $p_{f,t}$ to the order statistic of handed in bids $\beta_{-i,t}^{(q_i)}$. Finally, we assume that latent bids today and latent bids tomorrow are connected by individual characteristics. Since we abstract away the bidding process in $t + 1$, we end up with arrows from individual characteristics $ind.char_i$ to latent bids $\beta_{i,t}$ and our outcomes of interest $y_{i,t+1}$.

Our effect of interest is the black arrow from $overpaid_{i,t}$ to outcome $y_{i,t}$. Threats to identification are posed by, so-called, back-door-paths, which are paths that start with an arrow going into our treatment indicator $overpaid_{i,t}$ and go to the outcome $y_{i,t+1}$, but not through the the direct arrow. The back-door paths consist of two patterns: confounders (e.g. $\leftarrow auction_t \rightarrow$) and colliders (e.g. $\rightarrow overbid_{i,t} \leftarrow$).²⁷ A back-door path through a confounder is blocked if we control for that confounder. A back-door path through a collider is blocked if we do not control for that collider, instead it is undesirably opened (in DAG lingo) if we control for the collider (cf. bad control problem [Angrist and Pischke \(2009\)](#)).

In Panel B of Figure 5, an adjustment set that blocks confounder paths is highlighted in blue. That is, we block all confounder paths if we include the set $\{overbid_{i,t}, A_t, q_t, p_t\}$ as control variables in our regression. We control for auction price and quantity directly and we operationalize auction characteristics using fixed effects such as weekday, week, hour, product category and auctioneer fixed effects. Additionally, we condition on $overbid_{i,t}$ as we only look at bidders first overbids. This conditioning on the first overbid, however, is not without drawbacks. Indeed, Panel B in Figure 5 delineates the collider path that is opened by restricting the analysis to first overbids: $\beta_{-i,t}^{(q_i)} \rightarrow overbid_{i,t} \leftarrow \beta_{i,t}$. That is, by conditioning we on $overbid_{i,t}$ we leave open the possibility that bidders with high latent bids select into similar auctions and that this drives treatment.

Fortunately, the same graph also shows that the collider path also passes through individual bidder characteristics u_i . In fact, all back-door paths go through u_i , so we could block them all by simply conditioning on u_i . While this is an elegant solutions, it is complicated by the fact that u_i is unobserved. Thus, we have to rely on proxies that are, by definition, imperfect. The variable u_i mainly determines the height of a bidder's latent

²⁷Readers interested in graph theory will recognize these patterns as forks and inverted forks.

bid. Thus, variables such as the average amount of a bidder’s past behavior and experience in the auctions are very informative about individual characteristics. We calculate bidder history variables, both for bidder i and the rival bidders (see Appendix F).

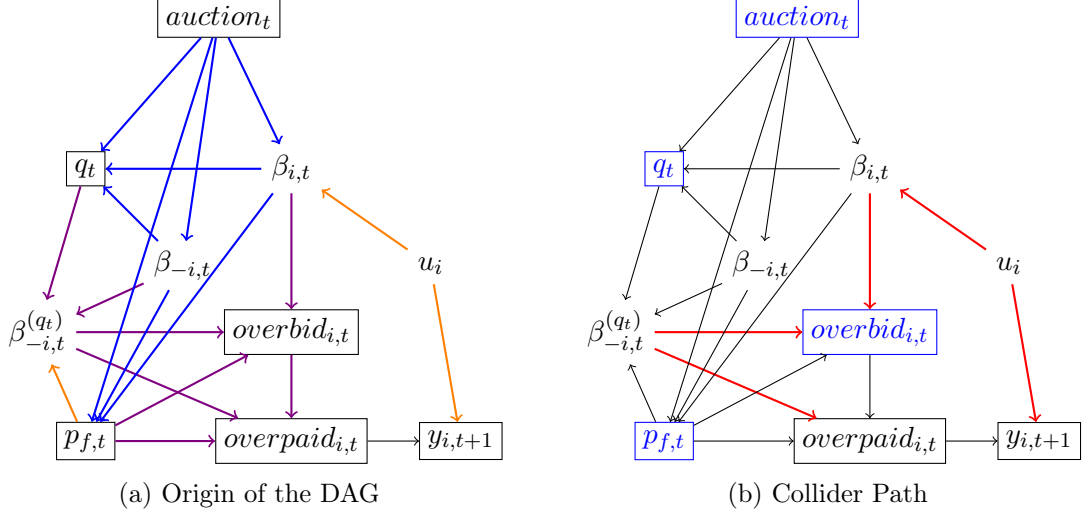


Figure 5: Panel (a) shows the origin of the DAG: Seller planning and running the auction (blue), Auction rules (uniform pricing, violet), Economic Assumptions (orange). Panel (b) shows that the collider path (red) opened by conditioning on $overbid_{i,t}$ (blue) also goes through $ind.char_i$.

We formalize our empirical strategy with the back-door criterion (Theorem 3.3.2 in Pearl (2009)). As we have shown $\{ind.char_i, auction_t, p_t, q_t, overbid_{i,t}\}$ or $\{ind.char_i\}$ block all back-door paths. Thus the causal effects of overpaying on future overbids and future non-overbids are identified and can be computed by controlling for these variables. This statement is equivalent to the statement that our potential outcomes are independent conditional on $\{ind.char_i, auction_t, p_t, q_t, overbid_{i,t}\}$ or $\{ind.char_i\}$.

7. Regression Results

We adjust for overbidding by restricting the sample to the first overbid for any customer. These initial overbids can be in an auction that ends below or above the fixed price. Bidders whose initial overbid was in an overpaid auction overpay and are in our treatment group. Bidders whose initial overbid was not in an overpaid auction do not overpay and form the control group. We follow these bidders for 90 days after their first overbid and count the number of overbids and non-overbids during that period. We exclude data after the

structural break, because overbids are much less likely due to firm policy after the break (see Figure 1).

To estimate the treatment effects in Proposition 4, we run the regression in Equation 7 on the sample of bidders first overbids.

$$Y_{i,t+k} = \beta_1 \text{overpaid}_{i,t} + \beta_2 p_t + \beta_3 q_t + H_{i,t} B_1 + H_{-i,t} B_2 + A_t + \eta_{i,t+k}$$

Here, $\text{overpaid}_{i,t}$ indicates treatment when the first overbid of bidder i in auction t lead to overpayment. Following our analysis of the causal graph in Section 6, we control for the auction price p_t , auction quantity q_t and a set of history controls as proxies of individual characteristics. We include bidder history variables, such as the average bid before the first overbid, both for bidder i , $H_{i,t}$, as well as the rival bidders in the same auction, $H_{-i,t}$. The full set of bidder history controls is described in detail in Appendix F. We use weekday, week, hour, product category, and auctioneer fixed effects to capture auction characteristics, A_t . Treatment is assigned at the auction level, so we cluster standard errors at the auction level (Abadie, Athey, Imbens and Wooldridge, 2017).

The variable $Y_{i,t+k}$ is a stand-in for revenue and the number of overbids and number of non-overbids in a k day long period. We use 0 to 90 and 90 to 180 days after the first overbid to provide estimates with a varying time frame. Finding similar results should reinforce confidence in the assumptions needed to recover the learning rates from the pooled treatment effects as argued in Corollary 5. We also report regressions excluding the history controls.

Table 7 shows the results of our revenue regressions. We find overpaying reduces revenue by roughly 9.95€ in the first 90 days after treatment. Compared to the revenue threshold we calculate in Section 4.1 this effect is sufficient to determine that the extraction of overpaying revenues was indeed suboptimal. The results for the period 90-180 days after the treatment are qualitatively similar, suggesting additional revenue loss of 7.31€, albeit statistically insignificant (see Appendix H).

We also consider the number of overbids and non-overbids in our regression analysis, as this allows us to back-out the underlying learning rates from our three-type model. Table 7 reports the results of these regressions. The first row reports the causal effect of overpaying

Table 3: Overpaying Reduces Future Revenue

	Revenue	Revenue
Overpaid	-5.079 (3.981)	-9.955** (4.852)
Num.Obs.	115 295	71 261
R2	0.043	0.100
Counterfactual Mean	138.206	160.407
Bidder History	No	Yes
Window	0-90	0-90

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

Overpaying reduces revenue in the 90 days after treatment by 9.95€. This revenue effect is larger than the revenue threshold we calculate in Section 4.1 of 9.77€ and we find further revenue effect for the period of 90-180 days after treatment (see Appendix H).

Table 4: Overpaying Reduces #Overbids and #Non-Overbids

	# Overbids	# Overbids	# Non-Overbids	# Non-Overbids
Overpaid	-0.171*** (0.032)	-0.190*** (0.034)	-0.185* (0.110)	-0.288** (0.127)
Num.Obs.	115 295	71 261	115 295	71 261
R2	0.074	0.131	0.071	0.154
Cf. Mean	1.433	1.62	5.796	6.854
Bidder History	No	Yes	No	Yes
Window	0-90	0-90	0-90	0-90

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

The negative impact of overpaying on number of overbids indicates treated bidders repeat their mistake less often than untreated. Larger negative effect on number of non-overbids is evidence of adjustment at extensive margin.

on future overbids and non-overbids. The row labelled counterfactual mean reports the fitted value for the regression with all variables set at their means and overpaid set to zero. With the full set of controls we find that overpaying decreases overbids in the following 90 days by -0.19 (compared to a counterfactual mean of 1.62) and non-overbids by -0.288 (compared to a counterfactual mean of 6.8). Results for the time period 90-180 days after treatment are similar (see Appendix H).

We assess our strategies of using bidder histories to proxy for u_i by looking at coefficient movements when adding these variables. Since our proxies have a good theoretical justification (high bids in the past are likely a good indicator of a tendency for high bids), our estimates should move closer to the truth when controlling for these proxies. Thus, if the magnitude of our estimates increases when we add the proxies, it should increase even more if we could actually control for u_i (see Oster, 2019, for a formalisation of this argument). Adding history controls (our proxies) increases the magnitude of our coefficient estimates. We take this as evidence that our identification strategy works well.

We recover the extensive and intensive margin learning rates as laid out in Corollary 5. Since we have to pool auctions to make estimation feasible, the expressions in the Corollary rely on the time period we aggregate over. Using two different time periods affords us a plausibility check: if the results are consistent it reinforces our confidence in the validity of our assumptions.

We report the backed out learning rates in Tables 5. We repeat the exercise using a Poisson regression model to complement the linear regression (the corresponding regression Tables are in Appendix I). The results are quite consistent between the two periods of time and the different regression methods, affirming our confidence in the heterogeneous-learning-rates assumption. We estimate the average overbidder has an extensive margin learning parameter of approximately 4% and an intensive margin learning parameter of approximately 7%. In other words, overpaying causes roughly 4% of overbidders to drop out of the market due to overpaying. This suggests extracting the extra overpaying revenue in the beginning of our sample was indeed suboptimal (the extensive margin threshold in Section 4.1 is 2.7%). This reinforces our conclusion from the revenue regressions that extracting overpaying revenue is indeed suboptimal.

Table 5: Backed out Learning Rates

Linear model	ϵ_u	ι_u	Poisson model	ϵ_u	ι_u
0-90	0.042	0.075	0-90	0.039	0.075
90-180	0.041	0.062	90-180	0.033	0.066

8. Conclusion

We find evidence for extensive as well as intensive margin learning as overpaying decreases future overbids and non-overbids. We use our model to recover extensive and intensive margin learning from these treatment effects. The causal effects imply that roughly 4% of bidders who overpay learn to drop out of the market (extensive margin), and about 7% of bidders who overpay learn not to overbid (the intensive margin).

A simple economic model teaches us how the firm should react to learning at these margins. We model an increase in firm sophistication by a broader scope of optimization: initially, the firm behaves sub-optimally, exploiting extra revenue from overpaying. Then, the firm behaves optimally along the quantity-adjustment margin, but treats fixed prices as exogenous, even though it has control over fixed prices. Finally, the firm chooses fixed prices and quantity optimally. If fixed prices are exogenous and extensive margin learning is high, the firm offers quantities that prevent overpaying. On the other hand, if the firm endogenously chooses fixed prices, it sets them high enough to prevent overbidding entirely, while still achieving a high price in the auction.

We make a number of observations that are in line with our model. First, we observe a period with overpaying in the market, followed by a sudden elimination of overpaying and increased quantity. This is in line with the initially suboptimal behavior of the firm and a policy change to target optimal quantity dynamically in the auctions, while keeping fixed prices unchanged. Second, we document a policy change after our sample ends. The new policy increases fixed prices and undercuts these higher prices with the auction’s starting bid, and thus, ruling out overbidding by definition. This policy is in line with our model with endogenous fixed prices.

We find that strategic learning and leaving the market are roughly equally likely. This finding unites the literature on learning (Haselhuhn et al., 2012; Agarwal et al., 2013; Ater and Landsman, 2013) and customer retention (e.g. Seru et al., 2010; Backus et al., 2022;

Anderson and Simester, 2010). From the perspective of the learning literature, consumers try to avoid the action that had negative consequences: they avoid overbidding because they overpaid. According to the literature on customer retention, they might also leave the market. Bidders might leave the market because they learn about their abilities as bidders or the value of participating in auctions. They can also become angry and leave the market.

That the firm is shaping this learning process is a novel reason for the persistence of consumer biases. The previous literature finds that firms can exploit consumer biases because consumers forget what they have learned (Agarwal et al., 2013), or new naive consumers replace experienced ones (Wang and Hu, 2009; Augenblick, 2016). We document that biases may also persist because firms make learning harder, when it is profit maximizing to do so. Our results on firms shaping consumer learning can explain market design choices and suggests possible avenues for regulating this behavior.

Complementing previous results, we find that when biased consumers learn, market-like institutions might be preferable to multiple single-unit auctions. Malmendier and Szeidl (2020) argue that firms want to sell several goods in individual auctions to fish for fools. In single-unit auctions, the highest bidder (likely upward biased) sets the price, whereas, in markets (and in the market-like auction we study), a larger share of biased buyers is needed to influence the price. According to Malmendier and Szeidl (2020) choosing individual auctions maximizes period profits. We show, however, that this may cost the firm customers because more individual auctions end overpaid. Consequently, sellers should be more likely to choose markets when bidders learn.

Firms can shape consumer learning in two ways: ways that benefit and ways that harm consumers.²⁸ According to our model, consumers are worse off when firms can change reference prices. In this case, the firm can remove the learning stimulus without benefiting the consumer. If the reference prices are exogenous, the firm prevents consumer learning through lower prices, which is in the interest of consumers.

In our setting, reference price regulation can constrain a firm's harmful ways of shaping consumer learning. For example, a regulator could mandate a minimum revenue share through sales at fixed prices. While the practical implementation of such a policy is

²⁸We do not model consumer preferences. Consequently, our only criterion for welfare analysis is that a lower price for the same quantity is good for consumers.

uncertain, it diminishes a firm's ability to raise fixed prices. Consequently, firms have to shape consumer learning through higher quantities, which benefits the consumer. There are already other types of reference price regulation. In Germany, for example, firms that advertise undercutting a reference price need to offer that reference price for a sufficient amount of time.²⁹

We provide a foundation for further research on customer retention and learning in platform markets. While we study policies specific to our context (higher quantities and higher fixed prices), these policies suggest a general pattern. Consumers learn from negative experiences. Consequently, the firm can reduce the number of negative experiences (higher quantities) or make existing negative experiences less salient (higher fixed prices). Further, more general research can build on our work and map features of existing markets into these two categories.

References

Abadie, Alberto, Susan Athey, Guido Imbens and Jeffrey Wooldridge (2017), When Should You Adjust Standard Errors for Clustering?, Technical report, National Bureau of Economic Research, Cambridge, MA.

URL: <http://www.nber.org/papers/w24003.pdf>

Agarwal, Sumit, John C. Driscoll, Xavier Gabaix and David I. Laibson (2013), 'Learning in the Credit Card Market', *SSRN Electronic Journal* .

Anderson, Eric T. and Duncan I. Simester (2010), 'Price stickiness and customer antagonism', *Quarterly Journal of Economics* **125**(2), 729–765.

Angrist, Joshua D and Jörn-Steffen Pischke (2009), *Mostly harmless econometrics: An empiricist's companion*, Princeton university press.

Ater, Itai and Vardit Landsman (2013), 'Do customers learn from experience? Evidence from retail banking', *Management Science* **59**(9), 2019–2035.

²⁹<https://www.frankfurt-main.ihk.de/recht/uebersicht-alle-rechtsthemen/wettbewerbsrecht/unlauterer-wettbewerb/irrefuehrende-werbung/mondpreise-5196206> accessed: 2.02.2022.

- Augenblick, Ned (2016), ‘The sunk-cost fallacy in penny auctions’, *Review of Economic Studies* **83**(1), 58–86.
- Avoyan, Ala, Robizon Khubulashvili and Giorgi Mekerishvili (2021), ‘Behavioral market design for online gaming platforms’, *Available at SSRN 3690620* .
- Backus, Matthew, Thomas Blake, Dimitriy Masterov and Steven Tadelis (2022), ‘Expectation, disappointment, and exit: evidence on reference point formation from an online marketplace’, *Journal of the European Economic Association* **20**(1), 116–149.
- Chetty, Raj (2009), ‘Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods’, *Annual Review of Economics* **1**(1), 451–488.
- Cho, Sungjin and John Rust (2010), ‘The flat rental puzzle’, *The Review of Economic Studies* **77**(2), 560–594.
- Cunningham, Scott (2021), *Causal Inference: The Mixtape*, Yale University Press, New Haven.
- Della Vigna, Stefano and Ulrike Malmendier (2006), ‘Paying Not to Go to the Gym’, *American Economic Review* **96**(3), 694–719.
URL: <https://pubs.aeaweb.org/doi/10.1257/aer.96.3.694>
- DellaVigna, Stefano and Matthew Gentzkow (2019), ‘Uniform pricing in us retail chains’, *The Quarterly Journal of Economics* **134**(4), 2011–2084.
- Doraszelski, Ulrich, Gregory Lewis and Ariel Pakes (2018), ‘Just starting out: Learning and equilibrium in a new market’, *American Economic Review* **108**(3), 565–615.
- Gesche, Tobias (2022), ‘Reference-price shifts and customer antagonism: Evidence from reviews for online auctions’, *Journal of Economics & Management Strategy* **31**(3), 558–578.
- Grubb, Michael D (2015), ‘Overconfident consumers in the marketplace’, *Journal of Economic Perspectives* **29**(4), 9–36.
- Grubb, Michael D and Matthew Osborne (2015), ‘Cellular service demand: Biased beliefs, learning, and bill shock’, *American Economic Review* **105**(1), 234–71.

- Hanna, Rema, Sendhil Mullainathan and Joshua Schwartzstein (2014), ‘Learning through noticing: Theory and evidence from a field experiment’, *The Quarterly Journal of Economics* **129**(3), 1311–1353.
- Haselhuhn, Michael P, Devin G Pope, Maurice E Schweitzer and Peter Fishman (2012), ‘The impact of personal experience on behavior: Evidence from video-rental fines’, *Management Science* **58**(1), 52–61.
- Heidhues, Paul and Botond Kőszegi (2018), Behavioral Industrial Organization, in ‘Handbook of Behavioral Economics: Applications and Foundations’, Vol. 1, Elsevier, Amsterdam and Boston, pp. 517–612.
URL: <https://linkinghub.elsevier.com/retrieve/pii/S2352239918300071>
- Hortaçsu, Ali, Fernando Luco, Steven L Puller and Dongni Zhu (2019), ‘Does strategic ability affect efficiency? evidence from electricity markets’, *American Economic Review* **109**(12), 4302–42.
- Hortaçsu, Ali and Steven L Puller (2008), ‘Understanding strategic bidding in multi-unit auctions: a case study of the texas electricity spot market’, *The RAND Journal of Economics* **39**(1), 86–114.
- Hossain, Tanjim and John Morgan (2006), ‘... plus shipping and handling: Revenue (non) equivalence in field experiments on ebay’, *Advances in Economic Analysis & Policy* **5**(2).
- Hünermund, Paul and Elias Bareinboim (2019), ‘Causal Inference and Data Fusion in Econometrics’, pp. 1–62.
URL: <http://arxiv.org/abs/1912.09104>
- Imbens, Guido W. (2020), ‘Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics’, *Journal of Economic Literature* **58**(4), 1129–1179.
- Kagel, John H and Dan Levin (2011), ‘Auctions: A survey of experimental research, 1995-2010’, *Handbook of experimental economics* **2**, 563–637.
- Kleiber, Christian and Achim Zeileis (2008), *Applied Econometrics with R*, Springer New

York, New York, NY.

URL: <http://link.springer.com/10.1007/978-0-387-77318-6>

Malmendier, Ulrike and Adam Szeidl (2020), ‘Fishing for fools’, *Games and Economic Behavior* .

Malmendier, Ulrike and Young Han Lee (2011), ‘The bidder’s curse’, *American Economic Review* **101**(2), 749–787.

Neugebauer, Tibor and Reinhard Selten (2006), ‘Individual behavior of first-price auctions: The importance of information feedback in computerized experimental markets’, *Games and Economic Behavior* **54**(1), 183–204.

Ocker, Fabian (2018), “‘Bid more, pay less’—overbidding and the Bidder’s curse in teleshop-
ping auctions’, *Electronic Markets* **28**(4), 491–508.

Oster, Emily (2019), ‘Unobservable Selection and Coefficient Stability: Theory and Evi-
dence’, *Journal of Business & Economic Statistics* **37**(2), 187–204.

URL: <https://doi.org/10.1080/07350015.2016.1227711>

Pearl, Judea (2009), *Causality*, Cambridge university press.

Peters, Jonas, Dominik Janzing and Bernhard Schölkopf (2017), *Elements of causal inference
: foundations and learning algorithms*, The MIT press, Cambridge (MA) ;London.

Seru, Amit, Tyler Shumway and Noah Stoffman (2010), ‘Learning by Trading’, *Review of
Financial Studies* **23**(2), 705–739.

URL: <https://academic.oup.com/rfs/article-lookup/doi/10.1093/rfs/hhp060>

Thaler, Richard H. (1999), ‘Mental accounting matters’, *Journal of Behavioral Decision
Making* **12**(September 1998), 183–206.

Wang, Xin and Ye Hu (2009), ‘The effect of experience on Internet auction bidding
dynamics’, *Marketing Letters* **20**(3), 245–261.

URL: <http://link.springer.com/10.1007/s11002-009-9068-3>

A. Seller's Problem

In this Section we present a simplified model used to derive Observations 1 and 2 in Section 5. We denote the number of overbidders by o_t and the number of non-overbidders by s_t . We assume that non-overbidders buy at the fixed price whenever they do not buy in the auction. To simplify, we assume that all bidders have the same latent bid $\beta > p$. We assume overbidders never buy at the fixed price. While this may sound restrictive, the seller will always supply each overbidder through the auction at a auction price weakly higher than the fixed price. All bidders have unit demand.³⁰

Definition 1. Seller's Problem

A profit-maximizing firm solves the following problem:

$$\max_{\{q_t\}_{t=0}^{\infty}} \sum_t \delta^t \pi(q_t, o_t, s_t),$$

where

$$\pi(q_t, o_t, s_t) = \begin{cases} \beta \cdot q_t + p \cdot s_t & \text{if } q_t \leq o_t \\ p(o_t + s_t) & \text{if } o_t < q_t \leq o_t + s_t \\ 0 & \text{if } o_t + s_t < q_t \end{cases}$$

subject to:

$$o_{t+1} = \begin{cases} o_t - (\epsilon + \iota)q_t & \text{if } q_t \leq o_t \\ o_t & \text{if } q_t > o_t \end{cases}$$

$$s_{t+1} = \begin{cases} s_t + \iota q_t & \text{if } q_t \leq o_t \\ s_t & \text{if } q_t > o_t \end{cases}.$$

B. Proof of Observation 1

Proof. We guess two policy functions (always choose $q_t = o_t + s_t$) and always choose $q_t = s_t$). Since the union of these conditions covers the parameter space the desired result follows.

³⁰Most of these assumptions are for illustrative purposes and are dropped in the empirical model in Section 6.1

Because o_0, s_0 are both larger than one and ϵ, ι are both smaller than one we guarantee that $o_t > 0 \wedge s_t > 0 \forall t$.

In the remainder of this proof we drop the time index to simplify our notation.

We can simplify the strategy space because some actions are dominated and some are outcome-equivalent. All actions with $q > o + s$ are dominated because profits are zero and we can get positive profits with $q = o + s$. Profits are constant over $o < q \leq o + s$. Thus we can eliminate this interval from the action space if we include its upper boundary $o + s$.

Having simplified the strategy space in this way we state the Bellman equation.

$$V(o, s) = \max_{q \in Q} \begin{cases} q \cdot \beta + sp + \delta V(o - (\epsilon + \iota) \cdot q, s + \iota \cdot q) & \text{if } q \leq o \\ (s + o)p + \delta V(o, s) & \text{if } q = o + s \end{cases},$$

where $Q = [0, o] \cup \{o + s\}$.

We guess and verify the policy $q = o + s$. The result of this policy is that the firm sells $o + s$ units each period at a price of p . This leads to the following value function

$$V(o, s) = \sum_{k=0}^{\infty} \delta^k (o_t + s_t)p = \frac{(o_t + s_t)p}{1 - \delta}.$$

We derive conditions under which this value function solves the Bellman equation

$$\frac{(o + s)p}{1 - \delta} = \max_{q \in Q} \begin{cases} q \cdot \beta + sp + \delta \frac{(o + s - \epsilon q)p}{1 - \delta} & \text{if } q \leq o \\ (s + o)p + \delta \frac{(o + s)p}{1 - \delta} & \text{if } q = o + s \end{cases},$$

where $Q = [0, o] \cup \{o + s\}$.

We need to check two cases, either the left arm ($q \leq o$) of the right-hand side of the Bellman equation rises or falls in q . It (weakly) rises if

$$\beta \geq \frac{\delta}{1 - \delta} \epsilon p.$$

In this case profits are either maximized at $q = o + s$ or at $q = o$. They are maximized at $q = o + s$ and our guess is true if

$$o \cdot \beta + sp + \delta \frac{(s + (1 - \epsilon)o)p}{1 - \delta} \leq (o + s)p + \delta \frac{(s + o)p}{1 - \delta} \quad (10)$$

$$\Leftrightarrow \frac{\beta - p}{p} \leq \frac{\delta}{1 - \delta} \epsilon. \quad (11)$$

If

$$\frac{\beta}{p} < \frac{\delta}{1 - \delta} \epsilon \quad (12)$$

the left arm ($q < o$) of the right-hand side of the Bellman falls in q .

In this case profits are either maximized at $q = 0$ or at $q = o + s$. They are maximized at $q = o + s$ if

$$sp + \delta \frac{(s + o)p}{1 - \delta} < (o + s)p + \delta \frac{(s + o)p}{1 - \delta}$$

$$\Leftrightarrow 0 < op,$$

which is true. Since condition 11 is strictly stronger than 12, we can verify our guess of no overbidding if condition 11 holds.

We guess that the seller wants all auctions to end in an overpay. Then the seller derives a profit of $p \cdot s$ from the initial non-overbidders in perpetuity. They derive a profit of β per overbidder in each period from a steadily declining stock of overbidders. This results in $o_t(1 - \epsilon - \iota)^k \beta$ in each future period k . In each future period a fraction ι of the current overbidders is transformed into non-overbidders $o_t(1 - \epsilon - \iota)^{p-1} \iota p$. Consequently, in period k there are $\sum_{p=1}^k o_t(1 - \epsilon - \iota)^{p-1} \iota p$ that were generated through intensive margin learning. The discounted sum of these period profits yields the value function under the conjecture that the seller ends all auctions in an overpay

$$\begin{aligned}
V(o, s) &= \sum_{k=0}^{\infty} \delta^k (o(1-\epsilon-\iota)^k \beta + sp) + \sum_{k=1}^{\infty} \delta^k \sum_{p=1}^k o(1-\epsilon-\iota)^{p-1} \iota p \\
&= o\beta \sum_{k=0}^{\infty} \delta^k (1-\epsilon-\iota)^k + sp \sum_{k=0}^{\infty} \delta^k + \frac{o\iota p}{1-\epsilon-\iota} \sum_{k=1}^{\infty} \delta^k \sum_{p=0}^k (1-\epsilon-\iota)^p - 1 \\
&= \frac{o\beta}{1-\delta(1-\epsilon-\iota)} + \frac{sp}{1-\delta} + \frac{o\iota p}{1-\epsilon-\iota} \sum_{k=1}^{\infty} \delta^k \left(\frac{1-(1-\epsilon-\iota)^{k+1}}{\epsilon+\iota} - 1 \right) \\
&= \frac{o\beta}{1-\delta(1-\epsilon-\iota)} + \frac{sp}{1-\delta} \\
&+ \frac{o\iota p}{(1-\epsilon-\iota)(\epsilon+\iota)} \sum_{k=1}^{\infty} \delta^k (1-\epsilon-\iota) - (1-\epsilon-\iota) \delta^k (1-\epsilon-\iota)^k \\
&= \frac{o\beta}{1-\delta(1-\epsilon-\iota)} + \frac{sp}{1-\delta} + \frac{o\iota p}{\epsilon+\iota} \sum_{k=1}^{\infty} \delta^k - \delta^k (1-\epsilon-\iota)^k \\
&= \frac{o\beta}{1-\delta(1-\epsilon-\iota)} + \frac{sp\delta}{1-\delta} + \frac{o\iota p}{\epsilon+\iota} \left[\frac{\delta}{1-\delta} - \frac{\delta(1-\epsilon-\iota)}{1-\delta(1-\epsilon-\iota)} \right].
\end{aligned}$$

We look for conditions under which this conjecture for the value function solves the seller's Bellman equation (equation 10). If

$$\beta < \delta(\epsilon + \iota) \left(\frac{\beta}{1-\delta(1-\epsilon-\iota)} + \frac{\iota p}{\epsilon + \iota} \left[\frac{\delta}{1-\delta} - \frac{\delta(1-\epsilon-\iota)}{1-\delta(1-\epsilon-\iota)} \right] \right) - \frac{\delta p}{1-\delta} \quad (13)$$

there is no overpaying because the left arm of profits fall in q . Then we have to compare $q = 0$ with $q = o + s$. Since the latter leads to higher period profits and both lead to the same future profits the firm prefers $q = o + s$, which refutes our conjecture.

If condition 13 does not hold the left-arm of the values function rises in q and the seller ends every auction in overpaying if he prefers that to $q = o$. This is the case if

$$\begin{aligned}
\beta - p &\geq \delta(\epsilon + \iota) \left(\beta(1 - \delta(1 - \epsilon - \iota))^{-1} \right. \\
&\quad \left. + \iota p(\epsilon + \iota)^{-1} \left[\frac{\delta}{1 - \delta} - \frac{\delta(1 - \epsilon - \iota)}{1 - \delta(1 - \epsilon - \iota)} \right] \right) - \iota \delta p(1 - \delta)^{-1} \\
&\Leftrightarrow \\
\beta - p &\geq \delta(\iota + \epsilon) \beta(1 - \delta(1 - \epsilon - \iota))^{-1} + \iota p \frac{\delta^2}{1 - \delta} - \iota p \frac{\delta}{1 - \delta} \\
&\quad - \iota p \frac{\delta^2(1 - \epsilon - \iota)}{1 - \delta(1 - \epsilon - \iota)} \\
&\Leftrightarrow \\
\beta - p &\geq \delta(\iota + \epsilon) \beta(1 - \delta(1 - \epsilon - \iota))^{-1} + \iota p \frac{\delta^2 - \delta}{1 - \delta} - \iota p \frac{\delta^2(1 - \epsilon - \iota)}{1 - \delta(1 - \epsilon - \iota)} \\
&\Leftrightarrow \\
\beta - p &\geq \delta(\iota + \epsilon) \beta(1 - \delta(1 - \epsilon - \iota))^{-1} + \iota p \frac{\delta^2 - \delta}{(1 - \delta)(1 - \delta(1 - \epsilon - \iota))},
\end{aligned}$$

where the last step follows if since $\frac{1}{d} > 1 - \epsilon - \iota$, which is always true since ϵ, ι and d are all between zero and one. Having simplified the condition so far we can collect terms and solve for a condition on ϵ

$$\begin{aligned}
(1 - \delta(1 - \epsilon - \iota))(\beta - p) &\geq \delta(\iota + \epsilon)\beta + \iota p \frac{\delta^2 - \delta}{1 - \delta} \\
\Leftrightarrow (1 - d)\beta - \frac{1 - 2\delta + \delta^2 + \delta\epsilon - \delta^2\epsilon}{1 - d}p &\geq 0 \\
\Leftrightarrow (1 - \delta)\beta - (1 - \delta)p - \delta\epsilon p &\geq 0 \\
\Leftrightarrow \frac{\beta - p}{p} \frac{1 - \delta}{\delta} &\geq \epsilon.
\end{aligned}$$

This condition covers all cases in which the other strategy is not optimal. Consequently, the seller either sets $q = o$ or $o < q \leq o + s$.

□

C. Proof of Proposition 2

Proof. The potential outcome for the the untreated (people that did not overpay) is the probability that an overbidder submits a strict non-overbid,

$$E[\text{non-overbid}_t^{t+1}(0)|u_i] = \mathbb{P}(p_{t+1} > \beta_{i,t+h} > b(q_{t+h})|u_i).$$

If we exogenously assign a bidder to the treated status they either stay an overbidder, become a non-overbidder or leave. In the cases in which they become a non-overbidder they also change their latent bid distribution. This leads to a change in probabilities which we denote by switching from P to P' . We calculate the potential outcome of a bidder treated in t and observed in $t + 1$ as

$$\begin{aligned} E[\text{non-overbid}_t^{t+1}(1)|u_i] &= E[(1 - \epsilon_i - \iota_i)|u_i] \mathbb{P}(\beta_{i,t+1} < p_{t+1} \wedge \beta_{i,t+1} > b(q_{t+1})|u_i) \\ &\quad + E[\iota_i|u_i] P'(\beta_{i,t+1} < p_{t+1} \wedge \beta_{i,t+1} > b(q_{t+1})|u_i). \end{aligned}$$

Adding an intelligent zero and taking the difference of potential outcomes yields the following expression for the treatment effects

$$\begin{aligned} E[TE_{\text{non-overbid}}^{t+1}] &= E[-\epsilon_i|u_i] \mathbb{P}(p_{t+1} > \beta_{i,t+1} > b(q_{t+1})|u_i) \\ &\quad + \iota_i (P'(p_{t+1} > \beta_{i,t+h} > b(q_{t+1})|u_i) - \mathbb{P}(p_{t+1} > \beta_{i,t+1} > b(q_{t+1})|u_i)) \end{aligned}$$

□

D. Proof of Proposition 3

Proof. We are interested in the treatment effect of overpaying in period t , but we can only estimate it from pooling observations over a time period. Previously, we have written down the treatment effect from overpaying in t on the next auction, i.e. in auction $t + 1$. If we add up these treatment effects, we do not account for the possibility that there could be double treatments. To rule out that effects come from double treatments we first write down the probability of a type change from treatments in periods $t + 1$ to $t + k - 1$. Using this probability we can calculate the effect from treatment in t excluding subsequent treatments. In Proposition 4 we sum over these treatment effects to get an expression for the treatment effect from pooled observations, that we can then use to recover learning rates.

Let $p_{l,k}$ denote the probability that an overbidder changes his type to dropout because of a treatment in periods $t + 1$ to $t + k - 1$. Notice that a change in type is irreversible in our model, so we simply need to sum over the probability to change type at a specific point in time, but not before that point in time. The probability of a type change in period $t + m$ is the probability of treatment, $overpaid_{i,t+m}$ times the probability of changing type due to treatment, ϵ_i . The converse probability is the probability of neither changing to dropout nor to non-overbidder before $t + m$. Thus, we get the following expression for $p_{l,k}$ (the argument for $p_{s,k}$ is analogous).

$$\begin{aligned} E[p_{l,k}|\epsilon_i, \iota_i, u_i] &= E \left[\sum_{m=1}^{k-1} overpaid_{i,t+m} \epsilon_i (1 - (\epsilon_i + \iota_i) overpaid_{i,t+m-1})^m | \epsilon_i, \iota_i, u_i \right] \\ E[p_{s,k}|\epsilon_i, \iota_i, u_i] &= E \left[\sum_{m=1}^{k-1} overpaid_{i,t+m} \iota_i (1 - (\epsilon_i + \iota_i) overpaid_{i,t+m-1})^m | \epsilon_i, \iota_i, u_i \right]. \end{aligned}$$

Next we use these probabilities to characterize the potential outcomes in period $t + k$ for a bidder who we assign exogenously to be either untreated ($E[non-overbid_t^{t+k}(0)|u_i]$) or treated ($E[non-overbid_t^{t+k}(1)|u_i]$) in period t . Note that the term $(1 - E[p_{l,k}|\epsilon_i, \iota_i, u_i])$ captures that the type change did not occur in periods $t + 1$ to $t + k - 1$.

$$\begin{aligned}
E[\text{non-overbid}_t^{t+k}(0)|u_i] &= E[(E[p_{l,k}|\epsilon_i, \iota_i, u_i])\mathbb{1}(\beta_{i,t+k} < p_{t+k} \wedge \beta_{i,t+k} > b(q_{t+k})|u_i)] \\
&= E[E[p_{l,k}|\epsilon_i, \iota_i, u_i]|u_i] \cdot \mathbb{P}(\beta_{i,t+k} < p_{t+k} \wedge \beta_{i,t+k} > b(q_{t+k})|u_i) \\
E[\text{non-overbid}_t^{t+k}(1)|u_i] &= E[(1 - \epsilon_i)(E[p_{l,k}|\epsilon_i, \iota_i, u_i])\mathbb{1}(\beta_{i,t+k} < p_{t+k} \wedge \beta_{i,t+k} > b(q_{t+k})|u_i)] \\
&= E[(1 - \epsilon_i)(E[p_{l,k}|\epsilon_i, \iota_i, u_i])|u_i] \cdot \mathbb{P}(\beta_{i,t+k} < p_{t+k} \wedge \beta_{i,t+k} > b(q_{t+k})|u_i).
\end{aligned}$$

The last step in each follows because conditional on u_i , ϵ_i and $\beta_{i,t+k}$ are independent. If we take the difference of potential outcomes we get the treatment effect.

$$E[\text{TE-overbid}_t^{t+k}|u_i] = E[-\epsilon_i \cdot E[p_{l,k}|\epsilon_i, \iota_i, u_i]|u_i]\mathbb{P}(\beta_{i,t+k} < p_{t+k} \wedge \beta_{i,t+k} > b(q_{t+k})|u_i).$$

The calculation for the treatment effect on observed overbids is analogous. □

E. Proof of Proposition 4

Proof. We take the expression for the potential outcome of the untreated and the treatment effect from the proof of 3 and divide one by the other.

$$\begin{aligned}
&\frac{-\sum_{m=0}^k E[TE_{\text{non-overbid}}^{t+m}]}{\sum_{m=0}^k E[\text{non-overbid}_t^{t+m}(0)|u_i]} \\
&= \frac{\sum_{m=0}^k E[\epsilon_i(1 - E[p_{l,m}|\epsilon_i, \iota_i, u_i])|u_i]\mathbb{P}(p_{t+m} > \beta_{i,t+m} > b(q_{t+m})|u_i)}{\sum_{m=0}^k E[(1 - E[p_{l,m}|\epsilon_i, \iota_i, u_i])|u_i]\mathbb{P}(p_{t+m} > \beta_{i,t+m} > b(q_{t+m})|u_i)} \\
&= \frac{\sum_{m=0}^k E[\epsilon_i(1 - E[p_{l,m}|\epsilon_i, \iota_i, u_i])|u_i]}{\sum_{m=0}^k E[(1 - E[p_{l,m}|\epsilon_i, \iota_i, u_i])|u_i]}.
\end{aligned}$$

The proof for the treatment effect on strict overbids is analogous. □

F. Control Variables

We calculate two sets of bidder history variables. First, we calculate histories for the bidders in our control and treatment group. Given that we restrict attention to the first overbid

of each bidder our bidder history variables only capture behavior and experience in the auctions of this seller before that first overbid. Hence our variables do not control for a previous overbid as there was none by construction. When a bidder overbids on the first bid we do not observe a history before that, because there was none. In this case we substitute the average from treatment and control bidders in the same auction. This substitute may not be available if all control and treatment bidders were new bidders. In this case we keep the NA and exclude these observations in regression that include bidders histories.

The bidder history variables roughly fall into two categories. First, there are variables that measure the average previous behavior. For example, the average difference to the high bid measures whether a bidder usually bids early in the auction and the share of bids by phone measures whether a bidder usually bids by telephone or online. Second, some variables refer more to the experience that the bidder had in the previous auctions. For example, the time in the market measures how many hours have past since the first observed bid for that bidder in our sample and total savings measures how much money the bidder has saved compared to the fixed price.

We calculate the same set of bidder history variables also for the other bidders in the auction, even if they are not in the treatment or control group. Referring back to Section 6.3, this controls for the other bidders individual characteristics u_i , that were left out of the DAG for simplicity.

Table 6 gives summary statistics for all history variables that we calculate. It is evident that there are differences between the treatment and control groups, which reassures us that it is helpful to control for this set of proxy variables.

Table 6: Average value of bidder history variables, and fixed price (p_t) and number of products (q_t) at the first overbid, split by overpaid.

	0		1	
	Mean	Std. Dev.	Mean	Std. Dev.
fixed price	29.63	67.20	30.78	85.86
quantity	282.88	275.17	272.61	257.81
new bidder	0.39	0.49	0.35	0.48
own number of bids	6.74	6.95	7.25	7.63
own average savings, logged	3.50	1.07	3.63	1.09
own average bid, logged	3.56	0.62	3.60	0.66
own time in market (hours)	1479.14	2491.12	1756.69	2690.80
own share of bids by phone	0.82	0.32	0.79	0.34
own average difference to the high bid	11.44	18.44	12.19	24.16
others average number of bids	44.56	39.27	47.42	39.18
others logged total savings	4.99	1.33	5.13	1.30
others logged average bid	3.35	0.47	3.43	0.49
others time in market (hours)	2502.92	2471.15	2701.00	2404.57
others fraction of new bidders	0.13	0.14	0.10	0.11
others share of bids by phone	0.58	0.15	0.60	0.15
others average difference to the high bid	7.73	3.75	8.21	4.63

Category	Share Overpaid	n
Heimwerken & Garten	0.32	6246
Mode & Accessoires	0.28	13269
Beauty & Wellness	0.26	15257
Uhren	0.25	8059
Schmuck	0.22	7950
Haushalt	0.20	16786
Möbel & Heimtextilien	0.16	6981
Freizeit & Sammeln	0.08	157

Table 7: Average probability of a bidder to be treated (overpay) at their first overbid by show category.

Weekday	Share Overpaid	n
Sunday	0.21	13909
Monday	0.17	9706
Tuesday	0.27	10209
Wednesday	0.28	9329
Thursday	0.24	9856
Friday	0.25	9581
Saturday	0.26	12115

Table 8: Average probability of a bidder to be treated (overpay) at their first overbid by day of the week.

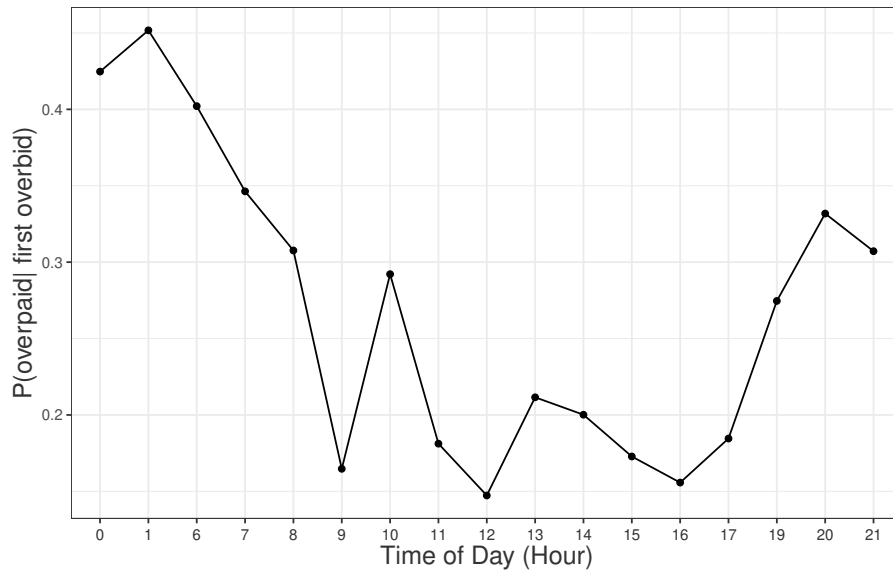


Figure 6: Average probability of a bidder to be treated (overpay) at their first overbid by time of day. Averages are by hour. The time between 18:00 and 19:00 is missing from hour data because of a coding error.

G. Structural Break

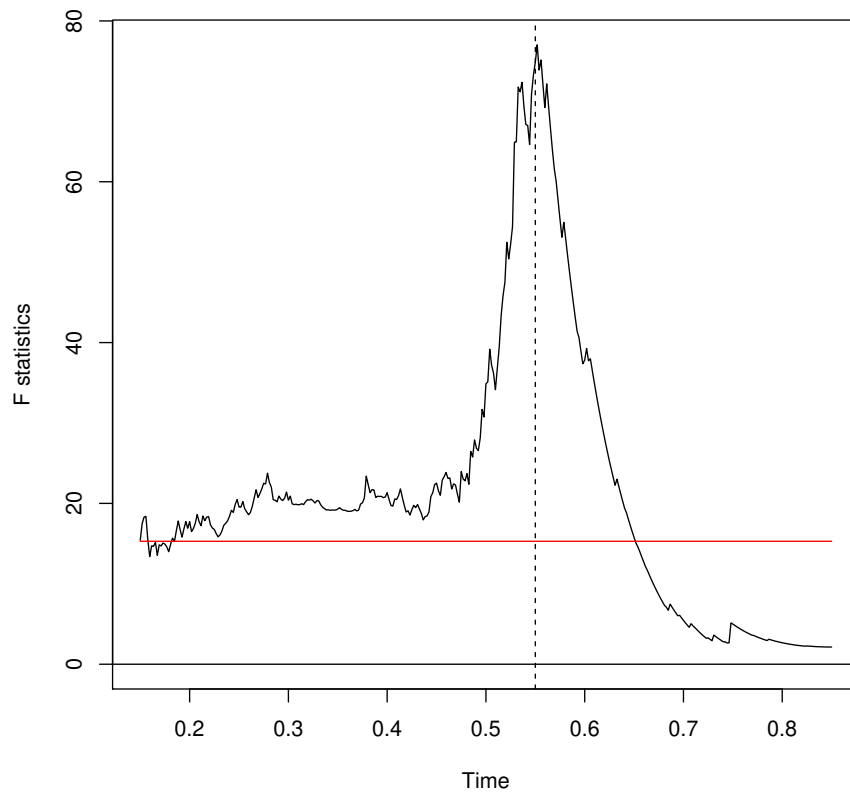
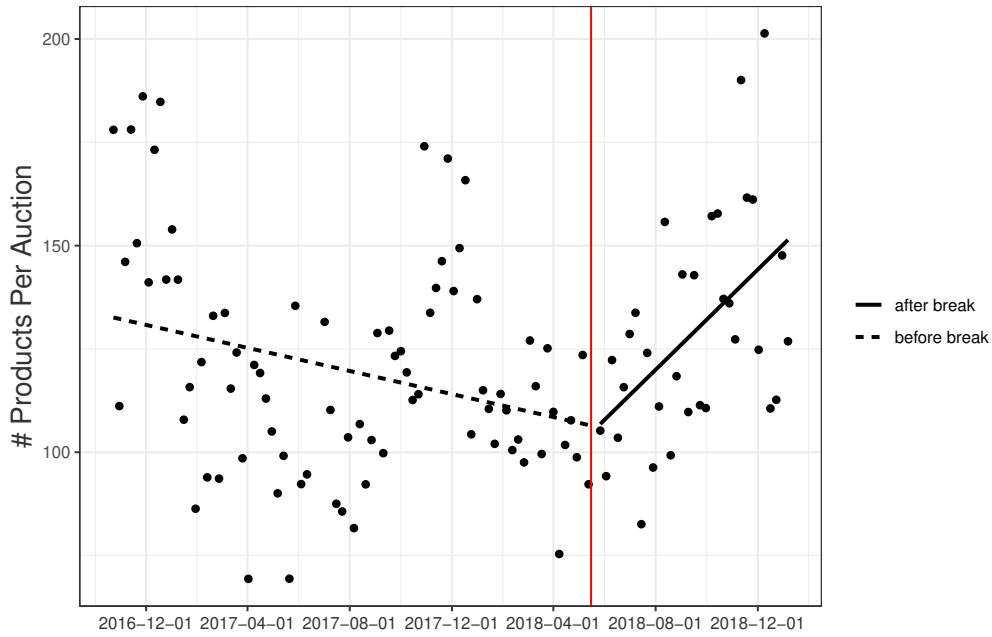
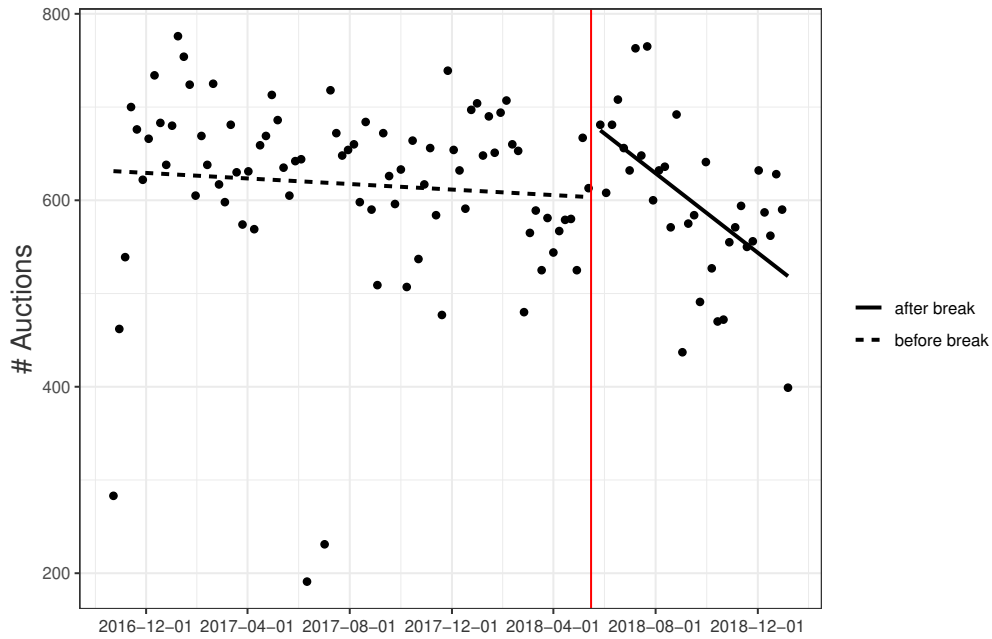


Figure 7: Time series of F statistics for a single shift hypothesis, fitted at every day in our sample. The red line gives the critical value at the 1 percent significance level. We accept the most probable break-point at the dashed line, 16th of May 2018.



(a) Number of products per auction.



(b) Number of auctions each week.

Figure 8: Changes in number of auctions and number of products per auction to both sides of the structural break. We use weekly averages and fit a linear trend at both sides of the structural break.

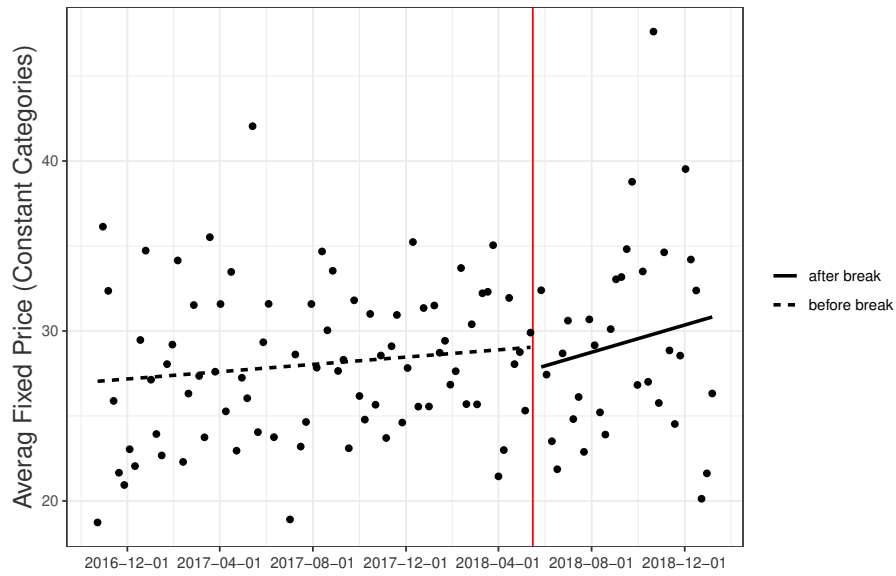


Figure 9: Changes in the fixed price of the auctioned products. We use weekly averages and fit a linear trend at both sides of the structural break.

H. Regression Using Period 90-180 Days After Treatment

Table 9: Overpaying Reduces Future Revenue (90-180)

	Revenue	Revenue
Overpaid	-0.598 (4.865)	-7.314 (6.005)
Num.Obs.	124 136	77 029
R2	0.051	0.126
Cf. Mean	193.165	224.008
Bidder History	No	Yes
Window	90-180	90-180

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

Table 10: Overpaying Reduces #Overbids and #Non-Overbids (90-180)

	# Overbids	# Overbids	# Non-Overbids	# Non-Overbids
Overpaid	-0.173*** (0.035)	-0.204*** (0.041)	-0.213 (0.150)	-0.422** (0.171)
Num.Obs.	124 136	77 029	124 136	77 029
R2	0.060	0.121	0.095	0.204
Cf. Mean	1.755	1.995	8.886	10.405
Bidder History	No	Yes	No	Yes
Window	90-180	90-180	90-180	90-180

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

I. Poisson Regression

Table 11: Overpaying Reduces #Overbids and #Non-Overbids (Poisson)

	# Overbids	# Overbids	# Non-Overbids	# Non-Overbids
Overpaid	-0.133*** (0.024)	-0.121*** (0.022)	-0.034* (0.020)	-0.040** (0.019)
Num.Obs.	115 295	71 261	115 295	71 261
R2				
R2 Pseudo	0.132	0.206	0.127	0.257
Cf. Mean	1.431	1.616	5.795	6.846
Bidder History	No	Yes	No	Yes
Window	0-90	0-90	0-90	0-90

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

Table 12: Overpaying Reduces #Overbids and #Non-Overbids (90-180, Poisson)

	# Overbids	# Overbids	# Non-Overbids	# Non-Overbids
Overpaid	-0.103*** (0.021)	-0.105*** (0.022)	-0.022 (0.018)	-0.034** (0.016)
Num.Obs.	124 136	77 029	124 136	77 029
R2				
R2 Pseudo	0.101	0.205	0.162	0.345
Cf. Mean	1.752	1.991	8.878	10.383
Bidder History	No	Yes	No	Yes
Window	90-180	90-180	90-180	90-180

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

J. First Five Overbids

We redo the empirical exercise separately using the first five overbids of each bidder. Learning may turn an initial overbidder into a non-overbidder or dropout already after the first overbid if it was overpaid. Accordingly, we observe fewer second overbids than first overbids as is depicted in Figure 10. Learning after the first overbid is likely to be skewed to well-learning overbidders and, hence, we expect to back-out smaller learning rates using the second overbid for each bidder compared to the first overbid. Figure 11 reports the backed-out extensive margin learning rates for the first five overbids. The pattern is the same across methods (ols and poisson regression) and time periods of aggregation (90 days after overbid and 90-180 days after overbid). Subsequent overbids are associated with smaller epsilon values, bar some outliers.

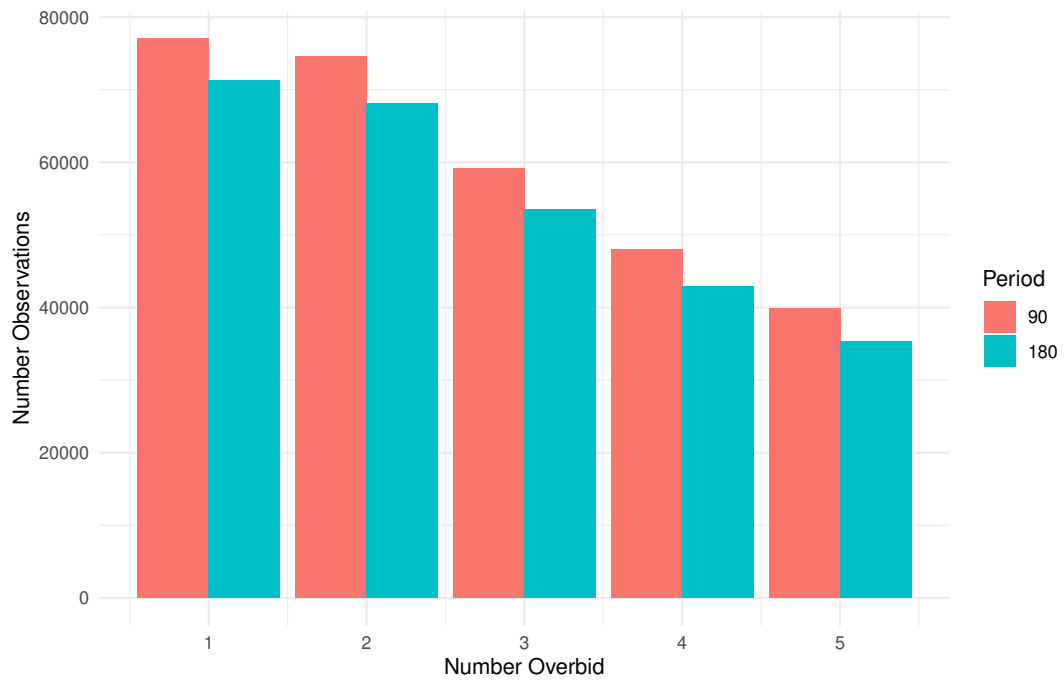


Figure 10: Number of Observations First Five Overbids

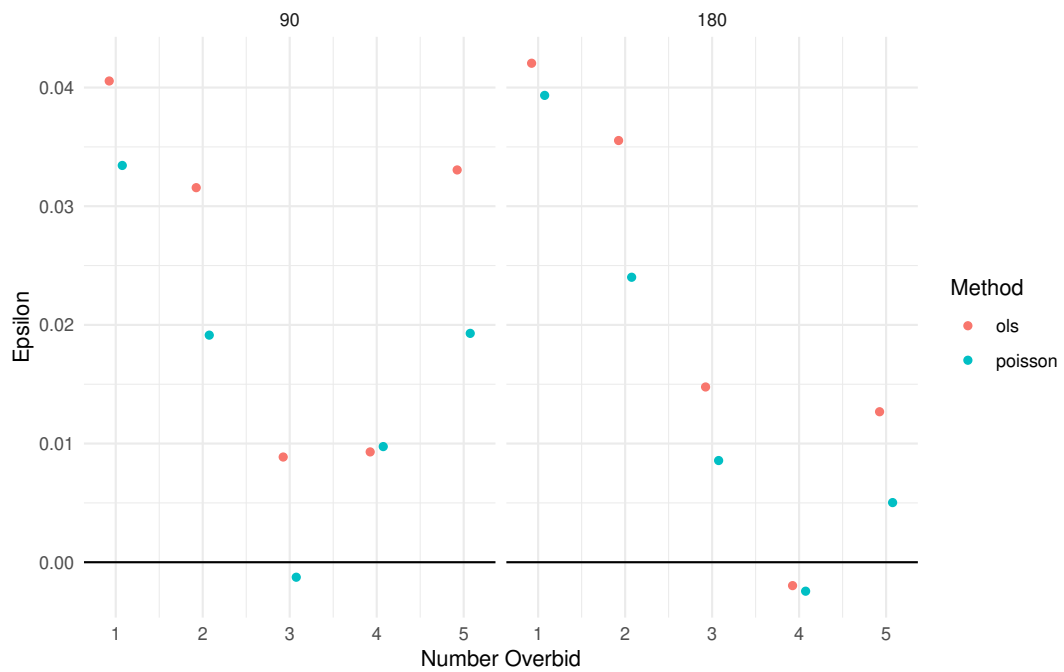


Figure 11: Epsilon First Five Overbids