

Expectation, Disappointment, and Exit: Reference Point Formation in a Marketplace*

Matthew Backus[†] Thomas Blake[‡] Dimitriy V. Masterov[§]
Steven Tadelis[¶]

29th December, 2016

Abstract

We study expectation-based reference point formation using data from an online auction marketplace. We hypothesize that exit from the marketplace is affected by disappointment from abruptly losing an auction after being the leading bidder. Expectation-based reference points that evolve over time imply that a bidder who spends more time in the lead prior to an abrupt loss will suffer a higher degree of disappointment. We find that for every additional day in the lead, bidders who lose abruptly are 6 percentage points more likely to exit. In contrast, losing bidders whose expectations are informed by early, competing bids, show no effect at all. Also, consistent with our theoretical model, more experienced bidders are less sensitive to time spent in the lead.

JEL Codes: D03, D47, D83.

*We thank Ned Augenblick, Stefano DellaVigna, Etan Green, Ulrike Malmendier, Conrad Miller, Don Moore, Collin Raymond, Alex Rees-Jones, Hersh Shefrin and Charlie Sprenger for helpful discussions, as well as Dan Ariely and George Lowenstein for comments on a previous draft.

[†]Columbia University, NBER, and eBay Research Labs, matthew.backus@columbia.edu

[‡]eBay Research Labs, thblake@ebay.com

[§]eBay Research Labs, dmasterov@ebay.com

[¶]UC Berkeley, NBER, and CEPR, stadelis@berkeley.edu

1 Introduction

The framework of prospect theory and reference-dependent preferences developed by the seminal work of Kahneman and Tversky (1979) has offered intuitive hypotheses for a number of so-called “behavioral” puzzles, from loss aversion to the endowment effect. To the canonical model of preferences it adds a *reference point*, a useful degree of freedom that distinguishes between a gain and a loss, or disappointment and elation. Moreover, if behavior in real markets mirrors findings from the lab, reference points matter for the way we make policy, manage firms, and design markets.

But where do reference points come from, and how can they be manipulated? This question is basic to both the positive and the normative implications of the theory, yet we have little empirical evidence on the point. Kahneman and Tversky (1979) are ambivalent; they mostly focus on the status quo as a reference point, but they also suggest that expectations may play a role. Following a line of thought that takes expectations as reference points (Bell, 1985; Loomes and Sugden, 1986), Kőszegi and Rabin (2006) propose an elegant and powerful solution in which reference points are given by rational expectations of equilibrium outcomes, implying that reference points arise naturally and endogenously within the context of the application.¹ A practical theory of reference-dependent preferences relies on the empirical challenge of understanding reference point formation in context, *i.e.*, in real market behavior.

This paper documents reference point formation and disappointment aversion using data from online consumer behavior. Unique data from eBay’s auction marketplace gives us leverage on the difficult empirical task of teasing out variation in reference points. We show how disappointment from losing an auction, mediated by variation in the *anticipated* probability of winning, affects a bidder’s likelihood of abandoning the auction platform. The data support two critical ingredients of the expectations-as-reference-points approach. First, a bidder should rationally expect a higher likelihood of winning the longer they are in the lead. Second, expecting a higher likelihood of winning makes an abrupt loss more disappointing, reducing the perceived benefit from participating in the auction platform, which in turn leads to exit.

¹Kőszegi and Rabin (2006) make clear that their adoption of rational expectations is not for the sake of dogmatism but modeling discipline: “...research on the nature of reference points themselves is quite limited. While we hope that experiments and other empirical work will shed light on this topic, our model makes the extreme assumption that the reference point is fully determined by the *expectations* a person held in the *recent past*” (p. 1141, emphasis original).

Using observational data to study reference point formation presents two salient measurement problems: finding an observable reference point “shifter” and measuring *ex post* disappointment. With those two challenges in mind, we use data from eBay’s hybrid Auction Buy-It-Now (ABIN) sales format. These auctions proceed similar to the standard ascending online auctions on eBay, but potential buyers may execute the Buy-It-Now (BIN) option and end the auction immediately to purchase the item at the pre-specified, publicly-listed BIN price set by the seller. Bidders who participated in the auction before a BIN event have no opportunity to counter or raise their bids in response. We focus on bidders who hold the standing high bid in an auction but lose abruptly when another bidder executes the BIN option. In our model, the time these bidders spend in the lead is a proxy for their expectations about the likelihood of winning: over time, the leading bidder places a higher probability on winning the auction, which puts them further in the domain of losses should they abruptly lose.

Our model formalizes the argument that, within this sample, bidders who have held the high bid for longer will suffer greater disappointment than those who held it only briefly. This is because their reference point shifts over time as they come to rationally anticipate a higher probability of winning the auction. Combined with a model of platform exit, our main hypothesis is that they are more likely to exit the platform altogether.² Our results confirm this: first-time bidders who spend more time in the lead before the BIN event are significantly more likely to exit.

To check that disappointment aversion is the mechanism and not, for instance, irritation in proportion to time squandered, we validate this against a placebo sample of bidders whose expectations are dampened by competition from other bidders—that is, they are outbid and ultimately lose, but they had ample time to raise their bid. An implication of our model is that, although these bidders have the same *material outcome*, they do not suffer disappointment because they had the option to win, and so anticipate loss when they chose not to re-bid. This is confirmed in our placebo sample.

²We are not the first to use platform exit as an outcome of interest. Ascarza et al. (2016) show that recommendations for consumers who make suboptimal bundle choices on a platform can cause exit and conjecture that the recommendations make expenditures more salient. Israel (2005) uses customer retention in auto insurance to identify learning events for an experience good, and Ho et al. (2015) use price changes and expensive medical events as price salience shocks to study substitution between health insurance plans and, in turn, the pricing of those plans. On the eBay platform Backus et al. (2015) showed that first-time bidders are substantially less likely to return to the platform when they are “sniped,” *i.e.*, lose to a last-minute bid, and Nosko and Tadelis (2015) and Masterov et al. (2015) use buyer exit to measure the quality of competing measures of seller reputation.

The explicit model of exit choice also allows us to make predictions about how BIN events affect more experienced bidders. Bidders return to the site if the expected surplus from doing so exceeds the value of their outside option. Experienced bidders have more data and therefore less variance in their prior beliefs about future surplus, so they update less severely on a negative experience. Moreover, they are selected on survival, which implies that they are likely to have a worse outside option. Therefore we predict that while time spent in the lead before a BIN event may affect their exit rates, they will be less sensitive than new bidders. The data confirms our hypothesis. This result is related to arguments that experienced market participants are less likely to display behavioral biases (List, 2003), but with an important difference in interpretation. Our model does not require that experienced traders have different *preferences* from inexperienced ones – rather, it explains why nonstandard preferences are less likely to be *expressed* in the case of experienced participants.

Barberis (2013), reviewing three decades of research on reference-dependent preferences, argues that outside the areas of choice under uncertainty, namely finance and insurance, a lack of empirical analyses casts doubt on the relevance of reference points, and prospect theory more generally. This is a gap that our paper helps fill, and closest to ours is Card and Dahl (2011), who use family violence as the outcome variable and find that upset losses strongly impact family violence, while losses in games that were expected to be close have small and insignificant effects. Our study uses platform exit as the outcome variable and expectations about winning as reference point shifters. While we employ a different reference point shifter and different measure of disappointment, our research design mirrors theirs in many ways. Our application, however, concerns market behavior, a domain in which documentation of reference points tends to be indirect, *e.g.*, the long series of papers concerning the labor supply of NYC taxi drivers (Camerer et al., 1997; Crawford and Meng, 2011; Farber, 2015). We believe therefore that we offer the clearest example of reference dependent preferences using observational data from real market behavior to date.

Our work is also related to the notion of a “quasi-” or “pseudo-endowment” effect, (Ariely and Simonson, 2003; Heyman et al., 2004; Wolf et al., 2005; Bramsen, 2008; Cotton, 2009). This notion posits that as a consumer holds the standing high bid in an auction, they become progressively attached.³ Accordingly, a losing bidder who has

³This is related to an argument of Strahilevitz and Loewenstein (1998), that the endowment effect

grown more attached should be *more* likely to return and bid higher when they do. This contrasts with our basic hypothesis – and our empirical findings – that such bidders will be less likely to return. Moreover, in Section 5.2 we show that bidders who do return do not make higher bids, and in Section 6.3 we show that our results are robust to focusing on “commodity” items, *i.e.*, those for which there are likely to be many equivalent listings on the platform that bidders could substitute to.

While our primary interest is the documentation of reference point formation, our results have implications for market design in online platforms. We show that employing allocation mechanisms that risk leaving consumers in the “domain of losses,” *i.e.*, disappointed *vis-à-vis* their reference point, can cause attrition from the platform. This is important for platforms because they face economies of scale via network effects.⁴ These effects are strongest for newer bidders to the platform, which is intuitive: these bidders are still engaged in learning to use the platform and forming expectations. Platforms should be wary of user experiences that encourage unrealistic expectations; disappointed consumers may leave the site entirely. It may be worthwhile to curate the experience of new users on a platform because loss alone does not cause disappointment and exit; consumers who expect to lose do not demonstrate the higher likelihood of exit, even though, in our sample, they share the same material outcome: they go home with nothing.

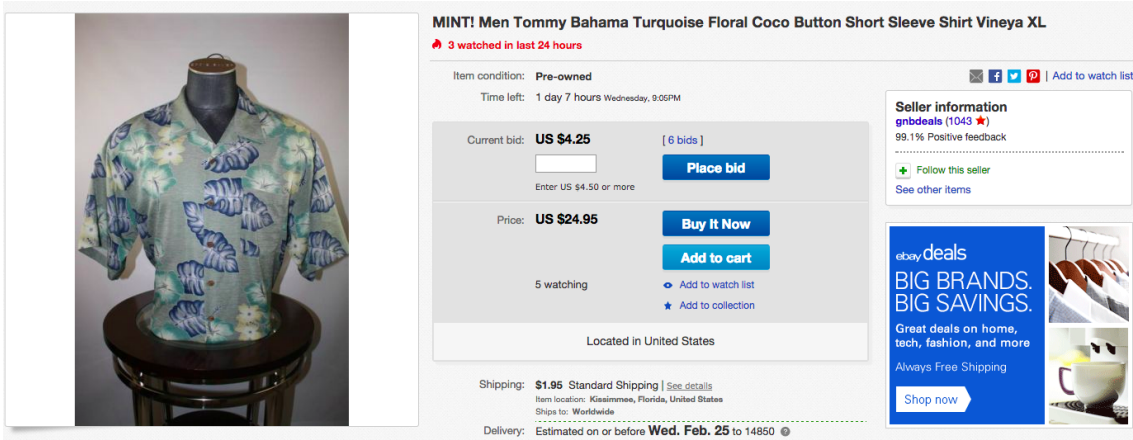
2 Auction Buy-It-Now Listings on eBay

We examine Auction Buy-It-Now (ABIN) listings on eBay.com. This is a subset of the auction listings in which a seller has elected to add the Buy-It-Now (BIN) feature. In a standard auction listing, the seller specifies the auction’s duration and the bidder with the highest bid at the end wins, paying the second highest bid plus an increment. In an ABIN listing, the auction proceeds similarly but potential buyers can execute the BIN option to immediately end the auction early and purchase the item at the BIN price. ABIN listings account for about one quarter of all auction listings on eBay. Figure 1 presents an example of an ABIN listing—buyers can either bid or to or click on the BIN button,

takes effect gradually as we incorporate possession into our reference point.

⁴In the platforms literature these are known as *direct network effects*, in which my payoff depends on the number of other participants (*e.g.*, social media) and *indirect network effects*, in which my payoff depends on the number of participants on the other “side” of the market (*e.g.*, buyers and sellers on eBay). See Rochet and Tirole (2003).

Figure 1: Example of an ABIN Listing



Notes: This is the listing page for an example ABIN listing. In addition to the “Place Bid” button, which is standard for auction-style listings and does not change the ending time, the buyer can use the “Buy-It-Now” option and purchase the good, closing the listing immediately. The current standing high bidder has no recourse when the BIN option is exercised.

in which case they will purchase the item at the BIN price of \$24.95.

Importantly, the BIN option may disappear in the course of the auction. For the product categories in our data, it disappears when the *standing price* exceeds 50% of the BIN.⁵ Ignoring bidding increments, this will occur only if there are two bidders who have each placed a bid that it at or above 50% of the BIN price. If only one bidder entered a bid that is higher than 50% of the BIN price then the standing price will be lower than 50% of the BIN price, and the BIN button will still appear for future buyers to click on. When the BIN option is exercised, the bidder with the standing high bid will be notified that the auction ended, and that their bid was unsuccessful. Therefore some bidders may anticipate winning and then abruptly lose without the ability to respond.

Sellers on eBay choose from a variety of mechanisms when they create a listing.⁶ Budish and Takeyama (2001) argue that offering a BIN option may allow sellers to extract more revenue from risk-averse buyers; relatedly, Akerberg et al. (2014) identify risk preferences as well as discounting preferences from the bid-or-BIN decision faced by bidders who arrive at a listing. Many other conjectures have been put forward to explain the use of

⁵In some other categories the BIN option disappears once a single bid is submitted that is above the secret reserve.

⁶These are auctions (with or without a BIN option) and fixed price (with or without a bargaining option). Minor options include duration, starting price, secret reserve price, end time, and shipping.

the ABIN mechanism, from reference dependent preferences to common-values concerns; see Bauner (2015) for a recent survey of this still-lively literature as well as a structural model of mechanism choice. For our application we take the set of ABIN listings as given, though we acknowledge that there may be substantial selection of listings into the ABIN mechanism. This is an external validity concern. We note, however, that buyers cannot explicitly condition on the availability of a BIN option among auctions when using the website’s built-in search feature.

3 A Model of Disappointment and Platform Exit

This section presents a simple behavioral model of the effect of unanticipated loss on the likelihood of platform exit. Our main hypothesis is that the longer a first-time bidder has been in the lead, the more severe the disappointment when they abruptly lose to the execution of a BIN option, and therefore the less likely he is to return and participate again. We also develop ancillary hypotheses for bidders who lose by other means, as well as for more experienced bidders who lose to the execution of a BIN option.

We develop the model in two steps. First, we characterize disappointment aversion for bidders with rational expectations reference points in an ascending price auction. Second, we model bidders’ decisions to revisit the auction platform as a function of past outcomes. Bidders weigh revisiting the platform against a random outside option, and update negatively on future expected surplus when they have negative experiences.

Our model is written to the end of generating meaningful empirical predictions. This guides us at several points. For instance, we focus exclusively on *losing* bidders because we can’t meaningfully compare platform exit among winners versus losers, given that winners have obtained the item they were seeking. Moreover, our theoretical model explicitly excludes many salient features of bidders and their bids, *e.g.*, when the bid was placed. We exclude these features from our model not because they are unimportant, but because in the empirical exercise they are *observable*, and therefore we can control for them to construct the empirical analogue of the simple model to which we now turn.

3.1 Disappointment and Reference Points

At any given point before an auction ends, there will be one bidder who is in the lead. A bidder whose standing high bid goes unchallenged for longer may rationally anticipate

that they are more likely to win the auction. As described below in Figure 2, we show that this anticipation is consistent with our observational data. This motivates our use of time spent in the lead as an instrument for movement in reference points. Namely, the longer a bidder is in the lead, the more they expect to win, and the greater the disappointment when they lose.

Consider a bidder who is currently in the lead, and normalize the difference between the consumption value of the item and the amount of their bid to 1 (we only consider losing bidders in our empirical exercise, so we ignore rebidding). We write the surplus realized from the auction as

$$\pi = y + \mu(y - p), \tag{1}$$

where $y \in \{0, 1\}$ is the bidder's allocation from the auction, p is the bidder's expectation of y at the time the auction ended, and μ represents the bidder's gain-loss utility as in Kőszegi and Rabin (2006). For simplicity let μ be piecewise linear, so that

$$\mu(y - p) = \begin{cases} \alpha(y - p) & \text{if } y - p \leq 0 \\ \beta(y - p) & \text{if } y - p > 0 \end{cases}$$

Disappointment aversion implies that $\alpha > 0$. In the empirical section we focus exclusively on losing bidders, *i.e.*, the region where $y = 0 \leq p$. This means we compare bidder with identical outcomes but different reference points.⁷

There are two circumstances that lead to $y \leq p$ at the end of the auction. First, the leading bidder can be “sniped” by another bidder who submitted a last-second bid, so that the leading bidder cannot possibly respond. Second, a buyer can exercise the BIN option, ending the auction abruptly. To formalize this, let $t = 0$ denote the beginning of an auction and $t = 1$ its end, and assume

⁷A more subtle distinction between disappointment aversion as formulated in Loomes and Sugden (1986) and reference-dependence as in Kőszegi and Rabin (2006) concerns whether the referent is the certainty equivalent or the entire distribution of outcomes. Sprenger (2015) explores this distinction experimentally and finds evidence for a risk endowment effect, consistent with the latter. Unfortunately our empirical setting does not offer any further leverage on this distinction.

ASSUMPTION 1 (Evolution of p):

$$p = \begin{cases} 0 & \text{if bidder } i \text{ is outbid and does not re-bid at } t < 1 \\ p_s \leq 1 & \text{if bidder } i \text{ holds the high bid at } t = 1 \\ p(t) < p_s & \text{if bidder } i \text{ loses to a BIN event at time } t < 1 \end{cases}$$

Moreover, $p(t)$ is strictly increasing in t .

The first scenario reflects bidders who are outbid early in the auction; at the end, they anticipate losing. However, bidders may also lose at the end of the auction. They face competition from “snipers,” and so we allow for the possibility that $p_s < 1$.⁸ Finally, the assumption that $p(t)$ is strictly increasing in t is consistent with any model in which the hazard rate of being outbid is strictly positive.


Assumption 1 is confirmed in the auctions that comprise our dataset (described in more detail in Section 4). Figure 2 presents a plot of the empirical likelihood that the bidder who holds the standing high bid at time t goes on to eventually win the auction. We see that $p(t)$ is monotone over t . It appears to be converging to approximately 0.7, before the slope increases sharply to arrive at 1, consistent with the presence of snipers.

In light of this, our assumption implies that bidders’ expectations about the likelihood of winning are qualitatively consistent with the data. This is related to the rational expectations reference points of Kőszegi and Rabin (2006), but somewhat weaker. We emphasize that for our analysis, bidders’ expectations about the path of $p(t)$ need not be *exactly* correct—this would be a particularly untenable assumption for first-time bidders—we only rely on monotonicity. Therefore we require bidders’ reference points to match the data qualitatively, rather than exactly. Our assumption formalizes the notion that a rational bidder who successfully maintains the standing high bid will also update positively on the likelihood that they will win the auction. This exposes them to greater potential disappointment if they lose.

A possible objection to our model concerns predictions for bidders who lose by being outbid, rather than to the BIN option. Because bidders evaluate surplus only at the *end* of the auction, they update their expectations before the auction ends and hence and

⁸See Roth and Ockenfels (2002) and Bajari and Hortaçsu (2003) for evidence on the causes and prevalence of sniping, and Backus et al. (2015) for evidence on its effects on platform exit. Much of this sniping happens in the last seconds of an auction, too late for any feasible response, and so we treat it as qualitatively different from being outbid at $t < 1$.

Figure 2: Empirical Likelihood that Current Leading Bidder Wins Auction



Notes: This figure presents the empirical likelihood that the bidder who holds the current high bid, at a time t ultimately wins the auction. Note that, as in the model, t is normalized such that an auction begins at $t = 0$ and ends at $t = 1$. To construct the plot we took, for a grid of points in t , the set of bidders who were currently in the lead in each of our auctions and computed the average value of a dummy variable for whether they were the *eventual* winner at the end of the auction. We use all auctions in the four ABIN categories during our sample period.

suffer no disappointment. But why should they suffer less disappointment when they are outbid? There is a substantive difference between these bidders and those who lose to a BIN option, which is that the former have *recourse*, *i.e.*, the option to re-bid. We exclude this from our model for the sake of parsimony, embedding it in a timing assumption, but in this sense our setup implicitly draws on a notion of regret. A more comprehensive formal treatment would require a unified theory of regret and disappointment, which to our knowledge does not exist.⁹ By assuming that surplus is evaluated at the end of the

⁹From Loomes and Sugden (1986), p. 281: “Because regret theory makes comparisons across actions but within states of the world, it can predict violations of the transitivity axiom but not violations of the sure-thing principle; whereas disappointment theory, which makes comparisons across states of the world but within actions can predict violations of the sure-thing principle, but not violations of transitivity. However, both theories generate many of the same predictions such as the common consequence and common ratio effects in the case of statistically independent prospects, simultaneous gambling and insurance, and the isolation effect. Given the similarity of the fundamental structure of both theories, there may be grounds for thinking that a more general theory of rational choice under uncertainty may encompass both regret and disappointment.”

auction, then bidders who lose to another bid suffer no disappointment because—unlike those to lose to a BIN event—conditional on losing, they anticipate loss. This is related to the finding of Filiz-Ozbay and Ozbay (2007) that bidders have no regret in standard second-price auctions.

3.2 Platform Exit

The second component of our model connects negative outcomes to platform exit. Bidders will update their beliefs about the value of participating in an auction and will chose to exit the platform after m auctions if the expected utility of participating again is less than that of the next-best substitute,

$$\mathbb{E}[\pi_{m+1}|\pi_1, \dots, \pi_m] \leq \theta_i. \quad (2)$$

In our notation, the expectation \mathbb{E} is taken with respect to bidder i 's subjective beliefs about the benefit of participating in another auction, and θ_i is an exogenous expected payoff to bidder i 's next-best substitute. θ_i is a bidder-specific random variable with distribution $F(\cdot)$. We assume that bidders have an astructural understanding of the surplus from platform participation: *i.e.*, their expectations about future values of π are monotone in prior realizations. We formalize this as

ASSUMPTION 2 (Beliefs): $\mathbb{E}[\pi_{m+1}|\pi_1, \dots, \pi_m]$ is strictly increasing in π_m .

This implies that a bidder who obtains less surplus in their last transaction is strictly less likely to revisit the site. It is not entirely innocuous: it could be violated if bidders are learning about mechanisms – *e.g.*, I lose to a late bidder today, but believe that I will get more surplus tomorrow now that I understand the practice. This is the sense in which bidders are astructural; we assume they are not learning about the structure of the market, so that their updates about future surplus are monotone in today's surplus. This is perhaps most reasonable for new bidders who may not fully apprehend the mechanisms on the platform.¹⁰

¹⁰For more rigorous foundations on misattribution and reference-dependent preferences, we refer the reader to recent work on this topic in Bushong and Gagnon-Bartsch (2016).

3.3 Empirical Predictions

Our main hypothesis combines disappointment aversion with platform exit:

HYPOTHESIS 1: The longer a first-time bidder has been in the lead, the more likely they are to exit the platform after experiencing a loss to a BIN event.

For first-time bidders who lose, the probability of exit is given by $F(\mathbb{E}[\pi_2|\pi_1])$, with $\pi_1 = \alpha(y - p)$. By Assumption 2, this is decreasing in p . For bidders who lose to the BIN option at time t , $p = p(t)$ by Assumption 1; moreover, $p(t)$ is increasing in t . Therefore the probability of exit is also increasing in t .

Note that Hypothesis 1 stands in contrast with prior work that has argued for a pseudo-endowment effect. The pseudo-endowment hypothesis suggests that bidders become more attached to the item over time, thereby making them more likely to return and to bid more aggressively when they do. There is a non-behavioral version of this hypothesis as well: if I anticipate winning the auction, perhaps I plan my consumption around the item, *e.g.*, purchase complementary goods. In either case, the empirical prediction is an increase in the bidder's willingness to pay for the product. To address this, in addition to offering evidence for Hypothesis 1, we test whether bidders who do return and bid on the same product place higher bids the second time.

Our model yields two additional ancillary hypotheses.

HYPOTHESIS 2: Among first-time bidders who expect to lose at the end of the auction, i.e. do not lose to a BIN event or a late bid, time spent in the lead is unrelated to exit.

For these bidders, $p = 0$ and therefore our model predicts that the probability of exit is invariant to time spent in the lead.

Hypothesis 2 helps to address alternative stories that depend on time spent in the lead but not disappointment. For instance, perhaps bidders update negatively in proportion to, and because of, time wasted on an unsuccessful bid. The comparison to bidders who had an opportunity to re-bid and chose not to, helps us isolate the role of reference points.

HYPOTHESIS 3: Among experienced bidders who lose to a BIN event, the relationship between time spent in the lead and exit is weaker than for inexperienced bidders.

This follows from selection on survival: repeat bidders are likely to have worse outside options. Formally, we have a long sequence of bounds on the value of their outside option. The probability that they return is given by $F(\min_{k=1,\dots,m}\{\mathbb{E}[\pi_{k+1}|\pi_1, \dots, \pi_k]\})$. Bidders who have recently had a series of good experiences become infra-marginal and are unlikely to exit even in the face of severe disappointment. Informally one might also expect, in Bayesian spirit, that with more data, bidders update less (and are therefore less likely to exit) with each additional data point.

In a baseball card trading field experiment, List (2003) showed that experienced traders are less subject to the endowment effect, and in that sense “less behavioral.” Anticipating selling the cards that they buy implies that their ownership is transient, making them less subject to the endowment effect. Our explanation, particular to our setting, is substantively different: experienced and inexperienced buyers on the platform need not have different preferences. Rather, the effect of disappointment aversion on exit behavior is endogenously determined by our model.

4 Sample Design


We begin constructing our sample with all new users who created an account on the U.S. eBay site between June 1, 2009 and October 31, 2013 and confirmed their account through e-mail verification. From this set, we selected users whose first non-browsing consumer action was casting a bid in an ABIN auction. We focus on first-time bidders because, as our model posits, these are the set for whom exit is most responsive to disappointment. We relax this restriction later to test Hypothesis 3.

We further limited the sample to auctions belonging to four specific product categories: cell phones and accessories; clothing, shoes and accessories; event tickets and experiences; and motor parts and accessories. These were chosen because therein the BIN option remains available until the standing price exceeds 50% of the BIN price. In other categories, the BIN option disappears once the first bid is placed.¹¹ We then restrict attention to successful auctions where there were at least two participants. For each auction that closed prematurely through the BIN option, we retained a *single* losing bidder: the one who was leading at the time the BIN option was exercised and was also a first-time bidder.¹² This

¹¹If there is a secret reserve price, the BIN option lingers until a bid exceeds it. In all categories other than vehicles, the BIN price must always be at least 30% higher than the starting price.

¹²We relax this restriction later to test Hypothesis 2.

Figure 3: Variation in Timing of the BIN Event



(a) Distribution of BIN Option Use

(b) Empirical CDF of Time in the Lead

Notes: Panel (a) presents a kernel density plot of the times at which the BIN option was exercised in our sample, normalized by the length of the auction to a $[0,1]$ scale. Panel (b) presents the empirical CDF of time spent in the lead by bidders in our sample in \log_2 hours. For example, the median bidder who loses to the execution of the BIN option spent six hours in the lead.

reduces the sample to 23,439 first-time bidders who lost due to the BIN event.

There is substantial variation in the timing of these BIN events, both in relation to the auction clock and with respect to the length of time that the losing bidder was in the lead. Figure 3a shows the distribution of when the BIN option was exercised. Because auctions vary in length, we normalized the timing by the intended duration so that 0 corresponds to the beginning and 1 to the scheduled end. Figure 3b shows the empirical CDF of time in the lead in hours. The median bidder is in the lead for approximately 6 hours, but there is a substantial right tail who lead for significantly longer periods of time. The average “time in the lead” for a bidder was just under 15 hours.

For each losing bidder, we construct our primary outcome variable, “exit,” as a dummy for whether the losing bidder ever participated in another auction within 365 days of their loss. 72% of new bidders who lost via BIN returned to bid in another auction within a year. This rate is substantially higher than in Backus et al. (2015), mainly because the bulk of these auctions are of phones and clothes.

We also obtain data on the timing and levels of their bids, demographic characteristics, as well as the characteristics of the auction and the seller.¹³ Auctions are scheduled to last

¹³Detailed summary statistics on these controls are in shown in Appendix Tables A-1 and A-2.

about 95 more hours at the time of the BIN event and are generally evenly distributed over the week. Items being auctioned are 42 percent new and are mostly split evenly between clothes and phones, with only 8 percent of the sample in auto parts and tickets. The joint coverage of the demographic data limits the sample to 12,068 bidders, which we use for some specifications that control for demographic characteristics. Bidders for whom we have demographics in this sample are 62 percent female and generally young with a wide distribution of incomes.¹⁴

5 Main Results: Testing Hypothesis 1

Here we test our main hypothesis, that time spent in the lead is correlated with bidder exit among first-time bidders who lose because of a BIN event. According to our model, they exit with probability $F(\mathbb{E}[\pi_2 | \alpha(y - p(t))])$, which is increasing in t .


5.1 Disappointment and Bidder Exit

Figure 4 plots the probability of bidder exit by hours spent in the lead as the highest bidder for a sample of first-time bidders who lost to a BIN event. We find a positive relationship between time in the lead and the likelihood of exit that is consistent with disappointment aversion and rational reference point formation.

Absent any confounds, Figure 4 is evidence for Hypothesis 1: the longer a bidder maintains the standing high bid before the BIN event, the more likely they are to exit. However, there are many omitted variables driving exit that could be correlated with variation in time in the lead. Mechanically, that variation comes from two sources: the time at which the bid is placed and the time of the BIN event. While the former is clearly endogenous, we posit that the latter source of variation is exogenous – we see no reason to believe that the buyer exercising the BIN option does so in a way that is correlated with the likelihood that the current high bidder will return to the site. We isolate this variation by introducing a large set of controls for the former source: the timing and the level of the standing high bid as well as attributes of the standing high bidder. This is our identification assumption – that conditional on this large set of controls, the timing of the arrival of the BIN event is exogenous. In Appendix A we consider the robustness of

¹⁴Note that this sample is not representative of the eBay users in general because it is for a select set of categories and conditional on a BIN event occurring and consists of new users only.

Figure 4: Exit Rate by Time in the Lead



Notes: This figure plots the empirical exit rate of bidders in our sample against time spent in the lead, in hours, prior to the BIN event which closed the auction. This does not control for any observable bid, auction, time, behavioral, or bidder attributes.

our results to omitted variables bias following the suggestions of Altonji et al. (2005) and Oster (2015). There we show that, given the robustness of our results to the large set of controls, one would need to hold very particular beliefs about the covariance structure of the observables and unobservable confounds in order to qualitatively overturn our result.

Table 1 shows the average marginal effects on exit from spending an additional 24 hours in the lead using six probit specifications. The first column presents the slope of Figure 4, with time rescaled to be in days rather than hours. The estimated marginal effect of an additional day in the lead is 7.3 percentage points over the baseline of 28%. That corresponds to a 26% increase in exit.

Our first robustness check is shown in Column 2, where we control for the number of spells as the high-bidder that our loser experienced in the course of the auction, as well as the decile of the ratio of the losing bid to the BIN price. These covariates are proxies for bidder attentiveness and willingness to pay. A bidder who has high willingness to pay may be more or less likely to exit because the high willingness to pay reflects that the bidder has limited options in other auctions, fixed price listings, or non-eBay merchants.

Table 1: AME of an Additional 24 hours in the Lead on Exit

	(1)	(2)	(3)	(4)	(5)	(6)
	AME	AME	AME	AME	AME	AME
Time in the Lead (24 Hours)	0.073*** (0.003)	0.069*** (0.003)	0.065*** (0.003)	0.067*** (0.003)	0.064*** (0.003)	0.064*** (0.005)
Bid-Bin Ratio and Spells		Yes	Yes	Yes	Yes	Yes
Auction Attributes			Yes	Yes	Yes	Yes
Time Variables				Yes	Yes	Yes
Site Visits					Yes	Yes
Demographics						Yes
<i>N</i>	23,439	23,439	23,439	23,439	22,390	12,068

Notes: This table presents the average marginal effect of time spent in the lead on the likelihood of exiting the auction marketplace for our sample of bidders who lose to a BIN event. Specification (1) corresponds to an approximation of the relationship in Figure 4, while specifications (2) - (6) introduce an extensive, nested set of controls. See text for details and discussion.

Their high bid is unlikely to be reached by other bidders, giving them a long spell as the winner. The mirror image of this is a “bargain hunter” who enters the auction at the very beginning with a low bid hoping to get lucky. We expect that bidder to enjoy a long spell as the top bidder since most bidding happens at the very end. This happens in part because the search ranking algorithm promotes auctions ending soon, a preference perhaps shared by eBay bidders. Leaving these variables out of the analysis may bias the effect in a theoretically ambiguous direction. The estimated marginal effect falls only slightly to 6.9 percentage points with these added controls.

In Column 3 of Table 1 we add variables that characterize the auction and the seller to absorb some of the variation in exit rates. We augment the specification with the product category of the item, the intended duration and ending day of the week for the auction, the natural log of the number of item page views for each day the listing was up, a set of dummies that describe the number of transactions that the seller had completed up to that point, the item condition (new, used, refurbished, unknown) and whether the item was in the product catalog (a commodity item). The exit rates clearly vary by product category, with clothes having the lowest and tickets the highest. Experienced sellers are more likely to provide better service and produce less exit from the platform, as shown in Nosko and Tadelis (2015). The duration variable is meant to capture bargain hunters, who are likely to return. Shorter auctions are typically used by sellers who are eager to sell, even at the cost of a lower price. The page views variable allows us to control

for characteristics unobservable to the econometrician that correlate with the listing’s desirability. For instance, a misspelling that makes a listing hard to find or high-quality photos, which influence search rank, would be captured by this metric. Item condition has a theoretically ambiguous effect on exit. With these, the estimated effect on exit drops to 6.5 percentage points, which is still economically and statistically significant.

Column 4 of Table 1 adds a set of dummy variables corresponding to deciles of scheduled time remaining when the BIN event occurs and a set of dummies corresponding to quintiles for auction part for the first bid of our loser. The motivation for this control is that time in the lead at the beginning of the auction has a qualitatively different effect than the same time in the lead near the end. This does not seem to be the case; the estimated effect is 6.7 percentage points.

The penultimate column adds the logged number of events and sessions divided by time between the first time that our buyer appears on the site and the losing auction start time. Events are essentially any actions on the site, such as a search query or a review of the bidding history of a listing. Whenever more than 30 minutes pass between events, a new session is started. Both of these variables proxy for interest in the site and such users are less likely to cease using auctions and eBay more generally. They also allow us to separate the effect of being more attentive from time in the lead. Controlling for these variables reduces the estimated effect to 6.4 percentage points. We are not controlling for site activity that takes place while the auction is unfolding since that is in part how the reference points are shifted.


In the final column of Table 1 we add demographics: annual household incomes of the users, their genders, and ages. These variables are missing for almost half the sample, but as they are important controls, we believe that their inclusion was a useful robustness check.¹⁵ The estimate of the marginal effects on exit from an additional 24 hours in the lead remains steady at 6.4 percentage points.

Figure 5 plots the expected probability of exit from the probit specification in Column 6 of Table 1 as function of time in the lead with a 95% confidence interval. Figure 6 shows the average marginal effect on exit from an additional 24 hours in the lead from specification 6 of Table 1 as a function of time in the lead with a 95% confidence interval.

The concave shape indicates that the marginal impact of time in the lead on exit increases with time in the lead, and drops off once someone has been winning for most

¹⁵Age is the main variable that is missing from our data.

Figure 5: Estimated Probability of Exit by Hours in the Lead



Notes: This figure presents a plot of the predicted probability of exit as a function of time in the lead using specification (6) of Table 1.


of the auction, yet it is always positive, and significantly different from zero. As such, the analysis provides strong evidence, both large in magnitude and highly statistically significant, that spending more time in the lead before losing the item abruptly cause the bidder to be more likely to exit the auction platform to avoid such future experiences. We view this as convincing evidence of our hypothesis.

5.2 Other Outcomes

Table 2 recreates the six specifications from Table 1 (by row) for a variety of other outcomes of interest. Note that each cell of the table presents an average marginal effect for a separate regression. The first column of Table 2 reproduces the Table 1 to facilitate comparisons. The second column takes as a dependent variable exit from eBay.com altogether, not merely the auction part of the market.¹⁶ Not surprisingly, the effect of time in the lead is somewhat weaker since some of the people who abandon auctions may substitute towards fixed price listings. In the third column, a successful outcome is an attempt

¹⁶Einav et al. (2013) show that fixed price sales are a large and growing portion of the website.

Figure 6: Marginal Effect on Probability of Exit of 24 Additional Hours in the Lead



Notes: This figure presents a plot of the average marginal effect on the predicted probability of exit as a function of time in the lead (*i.e.*, plugging in alternative values for time spent in the lead) using specification (6) of Table 1.

to purchase the same product catalog item, conditional on coming back to eBay. We find that time in the lead makes that less likely, and the effect is significant even in our richest specification despite the limited sample. This means that losing to a BIN event makes returning to find the same item less likely when the loser was in the lead for longer.¹⁷ This is a different result from both the experimental and field literature, which has found a heightening effect on rebidding within the same auction from additional time in the lead. Of course, it is possible that this attempt takes place on a different e-commerce site or offline, so that is unobservable to us. However, we anticipate this would attenuate rather than reverse the measurement of an attachment effect unless bidders who are susceptible to attachment are particularly likely to make purchases off-site. Finally, the fourth column shows the marginal effect on the change in willingness to pay (either the highest bid or fixed price), conditional on returning to eBay for the same product catalog item.

¹⁷We limited this regression to items where there were at least 10 listings available on eBay in the year after the loss to ensure this was not driven by stock-outs.

The scale is such that 100 would correspond to doubling the WTP and 0 to no change. We limited the regression to products with at least 10 listing within 1 year of the lost auction and removed any observations where the outcome exceeded 500. The estimates are generally positive, but noisy. While the estimated effects are typically not significant in the richer specifications, this is largely attributable to the fact that the data is pretty sparse as we move down the rows and to the right across the columns.

Table 2: AME of an Additional 24 hours in the Lead on Other Outcomes

	(1)	(2)	(3)	(4)
	Auct. Exit	eBay Exit	Same Product ID	WTP Change
Spec. (1)	0.073*** (0.003)	0.060*** (0.003)	-0.044*** (0.009)	-7.026 (5.018)
Spec. (2)	0.069*** (0.003)	0.057*** (0.003)	-0.050*** (0.010)	4.807 (4.573)
Spec. (3)	0.065*** (0.003)	0.056*** (0.003)	-0.054*** (0.011)	3.400 (5.332)
Spec. (4)	0.067*** (0.003)	0.058*** (0.003)	-0.055*** (0.011)	2.373 (5.478)
Spec. (5)	0.064*** (0.003)	0.055*** (0.003)	-0.056*** (0.011)	2.534 (5.564)
Spec. (6)	0.064*** (0.005)	0.054*** (0.004)	-0.043** (0.015)	0.000 (6.528)

Notes: This table presents the average marginal effect of time spent in the lead on other outcomes for our sample of bidders who lose to a BIN event. Each cell of the table corresponds to a separate regression. Column-wise, we consider alternative dependent variables. See text for further discussion. Row-wise, specifications (1) - (6) correspond to the nested control sets introduced in Table 1, where Specification 6 includes observables capturing bid and auction attributes, time variables, behavioral variables, and bidder demographics.

6 Additional Results: Testing Hypotheses 2 and 3

In the above we have shown that first-time bidders who lose abruptly following a BIN event are more likely to leave the site the longer they have been in the lead, consistent with disappointment aversion and expectations-based reference points. However, there are alternative stories consistent with Hypothesis 1. For instance, one might worry that time spent in the lead is a proxy for time and attention devoted to the purchase, that losing bidders may simply be annoyed at the effort wasted. To rule out this and other alternative hypotheses, we also test Hypotheses 2 and 3.

6.1 Expecting to Lose

Hypothesis 2 is a placebo test. It claims that the exit rate of bidders who have been outbid, and therefore anticipate loss at $t = 1$, should show no sensitivity to time spent in the lead. To test this, we augment our sample with first-time bidders who lost with more than 12 hours remaining in the auction. We impose the 12-hour restriction to guarantee that these bidders had ample time to respond.¹⁸ Intuitively, they had an opportunity to rebid and chose not to, and therefore anticipated their loss ($p = 0$).

Table 3: AME of 24 More Hours in the Lead on Exit By Loser Type

	(1)	(2)	(3)	(4)	(5)	(6)
	AME	AME	AME	AME	AME	AME
AME:						
Binned Loser	0.075*** (0.003)	0.076*** (0.003)	0.073*** (0.003)	0.061*** (0.003)	0.058*** (0.004)	0.057*** (0.004)
Non-Sniped Loser	0.011*** (0.001)	0.013*** (0.001)	0.009*** (0.001)	-0.010*** (0.001)	-0.011*** (0.002)	-0.005 (0.003)
Bid-Bin Ratio and Spells		Yes	Yes	Yes	Yes	Yes
Auction Attributes			Yes	Yes	Yes	Yes
Time Variables				Yes	Yes	Yes
Site Visits					Yes	Yes
Demographics						Yes
N	142,481	142,481	142,476	142,476	82,459	27,980

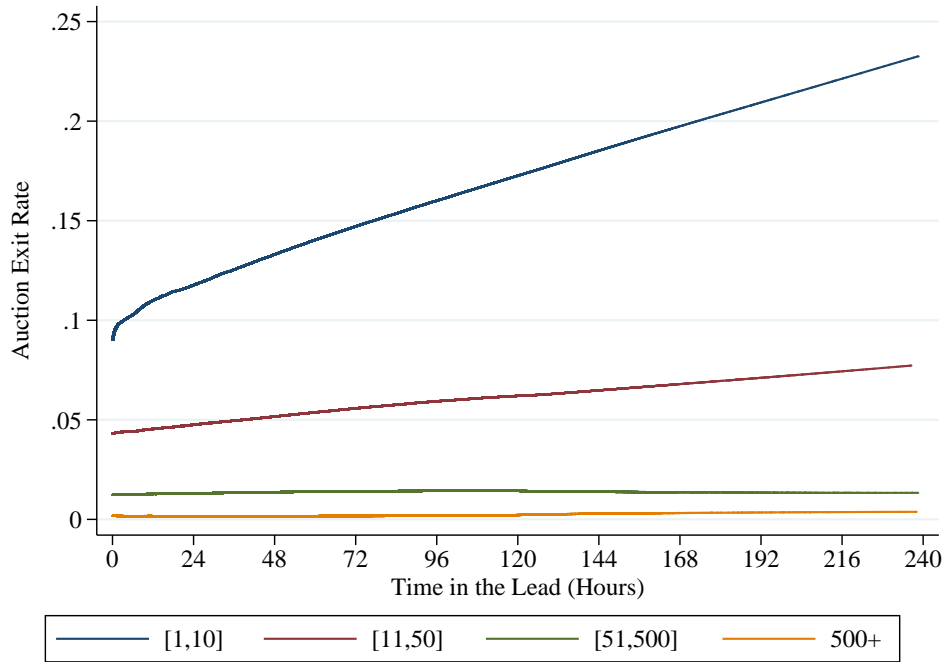
Table 3 recreates Table 1 for this augmented sample, where we allow for separate constant and slope parameters for the two samples. For specifications (1) - (5) we see smaller, non-robust effects for non-sniped losers that are sometimes negative and sometimes positive. However, the non-sniped losers who were outbid exhibit no disappointment aversion once we include the full set of controls in specification (6). Therefore we take this as confirmation of Hypothesis 2.

6.2 Experienced Losers

Hypothesis 3 claims that more experienced bidders will be less sensitive to time spent in the lead when they lose to a BIN event for two reasons. First, consistent with the model of Section 3, more experienced bidders are a selected sample with, in expectation, a worse

¹⁸The results are not sensitive to using a 6 hour or a 24 hour cutoff. The added observations come from a disjoint set of auctions from our main analysis. These are also losers, but the manner in which they were bested is different: they were outbid, rather than binned, with some time in which to respond.

Figure 7: Exit Rates by Time in the Lead for Experienced Bidders



Notes: This figure replicates Figure 4 by plotting exit rates against time spent in the lead for four samples of bidders who lose to a BIN event: bidders with experience in 1 to 10 prior transactions, 11 to 50 prior transactions, 51 to 500 prior transactions, and 500+ prior transactions.

outside option. Second, more informally but consistent with Bayesian updating, bidders should update their beliefs less with each additional experience.

Figure 7 recreates Figure 4 accounting for user experience. Consistent with Hypothesis 3, the effect of time in the lead on exit declines markedly as users accumulate additional experience. For bidders with 51 or more prior transactions, there appears to be no effect at all. These results are strongly consistent with Hypothesis 3.

6.3 Commodity Items

We noted earlier that an alternative story to disappointment aversion would be one of a quasi-endowment or attachment effect (Ariely and Simonson, 2003). It posits that as bidders become attached to a good in the course of bidding and are therefore *more* likely to return to the site upon losing, in order to ultimately obtain the item. Our results on exit are inconsistent with this hypothesis; our supplementary results in Table 2 showed

that bidders who do return do not bid on the same product, and bidders who do return *and* bid on the same product do not bid more. However, one might suspect that these bidders lack opportunities to bid on the exact same product. We therefore offer a sharper test by restricting attention to items that are in eBay’s product catalog, meaning that they are sufficiently common to have a UPC code and we expect that there are many such items offered, so that a bidder could return and buy exactly the same product. To further ensure this, we also limit the sample to items that had at least 10 listings in the year after the BIN event.

Table 4 recreates Table 1 using this restricted sample. If shifting reference points were creating an endowment effect, *i.e.*, changing valuations for an entire class of products, say *all* 64GB gold iPhone 3s rather than the particular one that was lost, we would expect the losing bidder to return to bid at a higher rate. Instead, we find a positive effect of time in the lead on exit probability. Moreover, we cannot reject the hypothesis that the slope is different for commodity and long-tail items.¹⁹

Table 4: AME of an Additional 24 Hours in the Lead on Exit for Commodities

	(1)	(2)	(3)	(4)	(5)	(6)
	AME	AME	AME	AME	AME	AME
Time in Lead (24 Hours)	0.091*** (0.008)	0.084*** (0.008)	0.067*** (0.009)	0.072*** (0.009)	0.072*** (0.009)	0.078*** (0.013)
Bid-Bin Ratio and Spells		Yes	Yes	Yes	Yes	Yes
Auction Attributes			Yes	Yes	Yes	Yes
Time Variables				Yes	Yes	Yes
Site Visits					Yes	Yes
Demographics						Yes
<i>N</i>	7,006	7,006	7,006	7,006	6,509	3,167

Notes: This table re-creates Table 2 using the following subsample: We restrict attention to the 7,006 bidders who made their first bid in an auction for a “commodity” item, *i.e.* an item for which there exists an entry in the eBay product catalog and had at least 10 listings in the year after the loss.

7 Discussion

We contribute to a growing literature that studies the role of reference points in decision making; see DellaVigna (2009) and Barberis (2013) for recent surveys. Barberis (2013)

¹⁹These results are available upon request.

argues that outside the areas of choice under uncertainty, namely finance and insurance, a lack of empirical analyses casts doubt on the importance of reference points, and of prospect theory more generally.

The main difficulty Barberis (2013) points to is the lack of guidance on what a gain or loss represent relative to well understood reference points. He states that the rational expectations solution offered by Kőszegi and Rabin (2006) “remains a hypothesis in need of more testing and, in any case, is unlikely to be completely correct.” He later argues that “Reference dependence is the most basic idea in prospect theory, and if any element of the theory finds a permanent place in economic analysis, it will surely be this one.” To the best of our knowledge, our analysis offers a first application with rational reference points that naturally shift with time to explain observed behavior in a large market setting.

Documenting rationally-formed, evolving reference points, involves two major measurement problems that have mostly confined the study of reference-dependent preferences to the lab. The first is finding a reference point shifter. In our setting, this is time spent as the winning bidder, which is consistent with rational expectations reference point formation in the spirit of Kőszegi and Rabin (2006). The data confirm that the longer one holds the winning bid, the more likely it is that they will win the auction. This parallels the reference point shifter in Card and Dahl (2011), who exploit the fact that rational expectations about NFL game outcomes are easy to measure using betting markets.

The second problem is finding a measure of *ex post* disappointment when the BIN option is exercised. Platform exit serves this purpose in our study, and it is also economically important. We believe that this is the clearest example of reference dependent preferences using observational data from real market behavior to date. Card and Dahl (2011) use family violence as the outcome variable and find that upset losses strongly impact family violence, while losses in games that were expected to be close have small and insignificant effects. Hence, both studies provide confirmation of rational reference point formation.

Even though Card and Dahl (2011) is the closest study to ours in many respects, we do find an interesting difference between the exit behavior of individuals in our marketplace and the violent behavior in their setting. NFL football games last for several hours, creating the opportunity for updating the reference point during the game. Card and Dahl (2011) test for such updating using the score at halftime and conclude that behavior is driven by the game outcome relative to expectations *at the start of the game*, with

no updating of reference points based on halftime information. We find that individuals appear to update their reference points in two important ways. First, time in the lead seems to continuously and rationally shift their beliefs about the likelihood of winning the object. Second, and more striking, if they lose and have recourse to re-bid, and choose not to, then they update their reference point in an extreme way. Namely, their behavior suggests that they seem to shed the expectation of winning the auction, regardless of how long they were in the lead before hand.

Our study points at the potential for future empirical work in identifying the magnitude and significance of behavioral biases using data from the exponentially growing datasets associated with online consumer behavior. As ours and other recent studies demonstrate, the breadth and increasing availability of online market data promises new directions for testing behavioral theories and developing an empirical behavioral research agenda that reaches far beyond lab environments.

References

- Akerberg, Daniel, Quazi Shahriar, and Keisuke Hirano**, “Identification of Time and Risk Preferences in Buy Price Auctions,” 2014. Working Paper.
- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber**, “Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools,” *Journal of Political Economy*, 2005, *113* (1), 151–184.
- Ariely, Dan and Itamar Simonson**, “Buying, Bidding, Playing, or Competing? Value Assessment and Decision Dynamics in Online Auctions,” *Journal of Consumer Psychology*, 2003, *13*, 113–123.
- Ascarza, Eva, Raghuram Iyengar, and Martin Schleicher**, “The Perils of Proactive Churn Prevention Using Plan Recommendations: Evidence from a Field Experiment,” *Journal of Marketing Research*, 2016, *53* (1), 46–60.
- Backus, Matt, Tom Blake, Dimitriy V. Masterov, and Steven Tadelis**, “Is Sniping A Problem For Online Auction Markets?,” 2015. NBER Working Paper 20942.
- Bajari, Patrick and Ali Hortaçsu**, “The Winner’s Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions,” *The RAND Journal of Economics*, 2003, *34* (2), 329–355.
- Barberis, Nicholas C.**, “Thirty Years of Prospect Theory in Economics: A Review and Assessment,” *Journal of Economic Perspectives*, February 2013, *27* (1), 173–96.
- Bauner, Christoph**, “Mechanism Choice and the Buy-it-Now Auction: A Structural Model of Competing Buyers and Sellers,” *International Journal of Industrial Organization*, 2015, *38*, 19–31.
- Bell, David E.**, “Disappointment in Decision Making Under Uncertainty,” *Operations Research*, 1985, *33* (1), 1–27.
- Bramsen, Jens-Martin**, “A Pseudo-Endowment Effect in Internet Auctions,” March 2008. Working paper.
- Budish, Eric B. and Lisa N. Takeyama**, “Buy Prices in Online Auctions: Irrationality on the Internet?,” *Economics Letters*, 2001, *72* (3), 325–333.

- Bushong, Benjamin and Tristan Gagnon-Bartsch**, “Learning with Misattribution of Reference Dependence,” 2016. Working Paper.
- Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler**, “Labor Supply of New York City Cabdrivers: One Day at a Time,” *Quarterly Journal of Economics*, 1997, *112*, 407–441.
- Card, David and Gordon B. Dahl**, “Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior,” *Quarterly Journal of Economics*, 2011, *126* (1), 103–43.
- Cotton, Christopher**, “Sniping to Avoid the Endowment Effect in Auctions,” 2009. Working Paper.
- Crawford, Vincent P. and Juanjuan Meng**, “New York City Cab Drivers Labor Supply Revisited: Reference-Dependent Preference with Rational Expectations Targets for Hours and income,” *American Economic Review*, 2011, *101*, 1912–1932.
- DellaVigna, Stefano**, “Psychology and Economics: Evidence from the Field,” *Journal of Economic Literature*, 2009, *47* (2), 315–372.
- Einav, Liran, Chiara Farronato, Jonathan D Levin, and Neel Sundaresan**, “Sales Mechanisms in Online Markets: What Happened to Internet Auctions?,” 2013. NBER Working Paper.
- Farber, Henry S.**, “Why you Cant Find a Taxi in the Rain and Other Labor Supply Lessons from Cab Drivers,” *Quarterly Journal of Economics*, 2015, *130* (4), 1975–2026.
- Filiz-Ozbay, Emel and Erkut Y. Ozbay**, “Auctions with Anticipated Regret: Theory and Experiment,” *American Economic Review*, 2007, *97* (4), 1407–1418.
- Heyman, James E., Yesim Orhun, and Dan Ariely**, “Auction fever: The effect of opponents and quasi-endowment on product valuations,” *Journal of Interactive Marketing*, 2004, *18* (4), 7–21.
- Ho, Katherine, Joseph Hogan, and Fiona Scott Morton**, “The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program,” November 2015. Working Paper.

- Israel, Mark**, “Services as Experience Goods: An Empirical Examination of Consumer Learning in Automobile Insurance,” *The American Economic Review*, December 2005, *95* (5), 1444–1463.
- Kahneman, Daniel and Amos Tversky**, “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, March 1979, *47* (2), 263–91.
- Kőszegi, Botond and Matthew Rabin**, “A Model of Reference-Dependent Preferences,” *The Quarterly Journal of Economics*, November 2006, *121* (4), 1133–1165.
- List, John**, “Does Market Experience Eliminate Market Anomalies?,” *Quarterly Journal of Economics*, 2003, *118* (1), 41–71.
- Loomes, Graham and Robert Sugden**, “Disappointment and Dynamic Consistency in Choice Under Uncertainty,” *Review of Economic Studies*, 1986, *53* (2), 271–282.
- Masterov, Dimitriy V, Uwe F Mayer, and Steven Tadelis**, “Canary in the e-Commerce Coal Mine: Detecting and Predicting Poor Experiences Using Buyer-to-Seller Messages,” in “Proceedings of the Sixteenth ACM Conference on Economics and Computation” ACM 2015, pp. 81–93.
- Nosko, Chris and Steven Tadelis**, “The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment,” 2015. NBER Working Paper 20830.
- Oster, Emily**, “Unobservable Selection and Coefficient Stability: Theory and Validation,” 2015. Working Paper.
- Rochet, Jean-Charles and Jean Tirole**, “Platform competition in two-sided markets,” *Journal of the European Economic Association*, 2003, *1* (4), 990–1029.
- Roth, Alvin E. and Axel Ockenfels**, “Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet,” *American Economic Review*, 2002, *92* (4), 1093–1103.
- Sprenger, Charles**, “An Endowment Effect for Risk: Experimental Tests of Stochastic Reference Points,” *Journal of Political Economy*, 2015, *123* (6), 1456–1499.

Strahilevitz, Michal A and George Loewenstein, “The Effect of Ownership History on the Valuation of Objects,” *Journal of Consumer Research*, December 1998, *25* (3), 276–89.

Wolf, James R., Hal R. Arkes, and Waleed A. Muhanna, “Is Overbidding in Online Auctions the Result of a Pseudo-Endowment Effect?,” 2005. Working Paper.

A Robustness Appendix

We conducted a series of robustness checks out of concern both for the exogeneity of our variation as well as the interpretation of the effects. First, we use a linear probability model with a heteroskedasticity-robust variance-covariance matrix. This is partly a specification check and partly meant to lay the ground for our second check, that applies Altonji et al. (2005) and Oster (2015) to calculate bounds on the omitted variable bias under proportional selection on observables and unobservables.

This robustness exercise relies on adding control variables which have substantial explanatory power on the outcome. To that end, we collected extensive controls on the timing and levels of bids, demographic characteristics of bidders, as well as the characteristics of the auction and the seller. Tables A-1 and A-2 display summary statistics for our sample.

A-1 Linear Probability Model

As a specification check and prelude to Section A-2, we replicate our results using a linear probability model (LPM) rather than the probit approach of Section 5. LPM estimates are reported in Table A-3. They are somewhat larger than the probit AMEs, but they have a very similar pattern of adjustment to covariates.

A-2 LPM Coefficient Stability

While our coefficients are stable as we add covariates, that is not a sufficient condition for omitted variable bias to be negligible. Figures A-1 and A-2 are based on the omitted variable bias bounding approach of Oster (2015). We assume that the unobservables share at least some of the covariance properties of the observables we selected. Without this reasonable assumption, very little can be said about the merits of our identification strategy.

The horizontal line in Figure A-1 shows the marginal effect on exit from an additional 24 hours in the lead from specification (6) in Table A-3. The bias of this estimate depends on two unknown in-sample parameters: the R^2 from a hypothetical regression of the outcome on all observed and unobserved determinants of exit, R^2_{MAX} , and the coefficient of selection proportionality δ . A δ of one indicates that the observed and unobserved

Table A-1: Summary Statistics: Auctions and Bids

	Mean	Std. Dev	Min	Max	N
Abandoned Auctions After Losing	0.28	0.45	0	1	23,439
Abandoned eBay After Losing	0.23	0.42	0	1	23,439
Number of Distinct-Auction Bids In the Year Since Losing	20.1	116.99	0	8770	23,439
Number of BINs In the Year Since Losing	3.15	10.32	0	367	23,439
Perc. Diff. B/W Subsequent Attempt and Losing Bid	1764.1	29300.02	-99.9	999900	2,093
Attempted to Buy Same Product ID	0.29	0.45	0	1	7,294
Time in the Lead (24 Hours)	0.61	0.91	0.000012	9.42	23,439
Time in the Lead Intervals	1.13	0.49	1	11	23,439
Losing Bid-BIN Price Ratio	34.6	20.13	0.0024	100.0	23,439
Seller's Previous Transaction Count (1Ks)	32.9	276.16	0	4183.5	23,439
Item Page Views Per Day Up	105.7	329.22	0.29	18078.3	23,439
Events Before Auction Was Up (Normalized)	176.2	11146.98	0.00010	1589760	22,390
Sessions Before Auction Was Up (Normalized)	2.04	69.76	0.00010	8640	22,390
Item in Product Catalog	0.31	0.46	0	1	23,439
Listings Within 1 Year for Losing Product ID	2541.2	9209.17	1	154947	7,294
First Bid Normalized By Duration	0.26	0.27	0.00014	1.00	23,439
Scheduled Time Remaining When Outbid (Hours)	95.7	54.38	0.0011	240.0	23,439
Intended Lost Auction Duration (Days)	5.52	2.22	1	10	23,439
Intended Auction End Day:					
Sun	0.15	0.35	0	1	23,439
Mon	0.15	0.36	0	1	23,439
Tue	0.15	0.36	0	1	23,439
Wed	0.15	0.36	0	1	23,439
Thu	0.14	0.35	0	1	23,439
Fri	0.13	0.34	0	1	23,439
Sat	0.13	0.33	0	1	23,439
Vertical:					
Clothes	0.46	0.50	0	1	23,439
Phones	0.46	0.50	0	1	23,439
Auto Parts	0.063	0.24	0	1	23,439
Tickets	0.022	0.15	0	1	23,439
Item Condition:					
New	0.42	0.49	0	1	23,439
Refurbished	0.017	0.13	0	1	23,439
Used	0.53	0.50	0	1	23,439
Unknown	0.032	0.18	0	1	23,439

Notes: This table presents summary statistics concerning auction and bid attributes from our main dataset of first-time bidders who participate in ABIN auctions and ultimately lose. See the text for further details on sample definition.

Table A-2: Summary Statistics: Bidders

	Mean	Std. Dev	Min	Max	N
Female User	0.62	0.49	0	1	18,836
User Age (2 Year Increments)	35.0	14.42	18	99	13,650
Annual Household Income:					
<15K	0.15	0.35	0	1	23,316
15-19K	0.066	0.25	0	1	23,316
20-29K	0.12	0.33	0	1	23,316
30-39K	0.11	0.32	0	1	23,316
40-49K	0.098	0.30	0	1	23,316
50-74K	0.21	0.40	0	1	23,316
75-99K	0.10	0.31	0	1	23,316
100-124K	0.055	0.23	0	1	23,316
125K+	0.086	0.28	0	1	23,316


Notes: This table presents summary statistics concerning bidder attributes from our main dataset of first-time bidders who participate in ABIN auctions and ultimately lose. See the text for further details on sample definition.

Table A-3: LPM Marginal Effect of an Additional 24 Hours in the Lead on Exit

	(1)	(2)	(3)	(4)	(5)	(6)
	ME	ME	ME	ME	ME	ME
Time in the Lead (24 Hours)	0.080*** (0.004)	0.076*** (0.004)	0.073*** (0.004)	0.075*** (0.004)	0.072*** (0.004)	0.071*** (0.005)
Bid-Bin Ratio and Spells		Yes	Yes	Yes	Yes	Yes
Auction Attributes			Yes	Yes	Yes	Yes
Time Remaining				Yes	Yes	Yes
Site Visits					Yes	Yes
Demographics						Yes
<i>N</i>	23,439	23,439	23,439	23,439	22,390	12,068
<i>R</i> ²	0.03	0.04	0.05	0.05	0.08	0.09

Notes: This table presents the LPM analogue of the profit specification employed in Table 2. As before, specifications (1) - (6) introduce an extensive, nested set of controls.

Figure A-1: Treatment Effect Under Assumptions About Selection Proportionality



Notes: This figure presents the treatment effect for the LPM results of specification (6) adjusted according to the method of Oster (2015). See text for a discussion of this method and definition of parameters δ and R^2_{MAX} . The dashed vertical line refers to the recommended value for comparison of $1.3 \cdot \tilde{R}^2$.

variables have an equally important effect on the coefficient of interest. A δ greater than one indicates that the unobserved variables are relatively more important.²⁰ Similarly, an R^2_{MAX} of 1 would indicate that the *potential* explanatory variables are highly predictive and there is very little measurement error or idiosyncratic variation driving a bidder's exit decision.

²⁰To fix the idea, consider our model


$$y = \alpha + \beta \cdot t + \sum_{j=1}^C \gamma_j \cdot x_j + \sum_{k=1}^U \eta_k \cdot z_k + \varepsilon = \alpha + \beta \cdot t + C + U + \varepsilon,$$

where y is the binary exit outcome, t is time in the lead, x_j are the controls and z_k are the unobserved variables. We assume that U and C are orthogonal. The proportionality of selection relationship is defined as

$$\delta = \frac{\frac{Cov(U,t)}{Var(U)}}{\frac{Cov(C,t)}{Var(C)}}.$$

Essentially, δ is the ratio of univariate regression coefficients of the linear index t on U and the linear index t on C . R^2_{MAX} comes from the infeasible regression of y on t and the variables that comprise C and U .

Figure A-2: Proportionality of Selection Coefficient δ s.t. $\beta = 0$



Notes: This figure plots the proportionality of selection coefficient (δ) required to generate adjusted estimates such that $\beta = 0$, as a function of R^2_{MAX} . We view this as a bound on the assumptions on unobserved heterogeneity required to qualitatively overturn our result. The dashed vertical line refers to the recommended value for comparison of $1.3 \cdot \tilde{R}^2$ from (Oster, 2015).

The green line shows the bias-adjusted marginal effect for the equal selection case ($\delta = 1$) for a range of plausible R^2_{MAX} values. Other downward-sloping lines make different assumptions about the selection proportionality. In a setting where the observed and unobserved variables explain a lot of the variation, the bias-adjusted effect turns negative. However, near the $1.3 \cdot \tilde{R}^2$ cutoff for R^2_{MAX} suggested by Oster (2015), the bias-adjusted estimates are generally near the unadjusted one and considerably above zero.²¹ This suggests that our estimated effect is robust and economically significant.

For each value of R^2_{MAX} , we also calculate the value of δ that would explain away the marginal effect and reduce it all the way to zero. That value is plotted in Figure A-2.

²¹The cutoff is based on a sample of high-quality randomized studies that considered the sensitivity of treatment effects to the inclusion of controls. The assumption is that since treatment is experimental, these estimates should generally “survive” the adjustments process. The cutoff of $1.3 \cdot \tilde{R}^2$, where \tilde{R}^2 is the R^2 from the regression with the controls, allows 90% of estimates to survive. This value corresponds to a bound on δ where the unobservables explain somewhat less than the observables.

Near our cutoff, the unobserved variables needs to be more than 15 times as important as observed ones to completely suppress the estimated effect. When the exit decision is more predictable, at least hypothetically, a smaller role for unobservables is sufficient to reduce the estimates below zero.

Another motivation for this robustness check is that it is possible for the perceived time in the lead to exceed the measured time since users are typically not monitoring the site closely and are inattentive to outbid notifications. We use the timestamp of when the BIN event occurs to define time in the lead, but some amount of time may pass before the losing bidder reads the e-mail or becomes cognizant of in-app notification, which are instantaneous. This gap between measured time in the lead and perceived time in the lead for the user means the effect we seek to estimate will be attenuated. Since this issue can be re-expressed as a kind of omitted variable problem, this bounding approach allows us to characterize the magnitude of this measurement error-induced bias in the behavioral response.