What Brings a Consumer Back for More? Evidence from Quantifiable Gain and Loss Experiences in Penny Auctions

Jaimie W. Lien† Minbo Xu‡ Jie Zheng±

Current Version: May 25th, 2021
Initial Draft: April 30th, 2015

Abstract:[1]

Can loss aversion explain consumer retention behavior for new products and experiences? We utilize detailed data from penny auctions, which first appeared as an internet phenomenon in the late 2000's, to uncover how consumers' prior experiences predict their willingness to try a relatively new experience again. The penny auction setting allows us to quantify consumers' positive and negative experiences, avoiding the measurement challenges typical in other consumer experience settings. Using several empirical approaches, we find the following patterns with respect to nominal gains and losses, which bear intuitive resemblance to reference-dependent behavior: First, consumers' tendency to return is at least weakly increasing in the outcome from their prior experience; Second, consumers' return rate on marginal loss experiences drops more steeply than it increases for marginal gain experiences; Third, consumers have an aversion to loss at any level, making them discontinuously less likely to return after any loss compared to after a gain. Developing a theoretical model of consumers' decisions to return to a service after prior experiences, we derive a reasonable set of sufficient conditions on utility functions and reference point formation which can account for these seemingly intuitive empirical findings. Finally, we examine the bracketing tendencies of consumers in calculating their gains and losses after repeated experiences. Consumers tend to bracket narrowly, with the above patterns strengthening, the more recently the experience occurred. The bracketing result is consistent with the intuition that a recent negative experience tends to drive consumers away permanently in spite of previous positive experiences.

JEL Codes: D03, D11, D12

Keywords: consumer experience, reference-dependence, purchase mechanisms, penny auctions

---

†Lien: Department of Decision Sciences and Managerial Economics, The Chinese University of Hong Kong, Shatin, Hong Kong, jaimie.academic@gmail.com ;
‡ Xu: Department of Economics, Business School, Beijing Normal University, Beijing, China, mbxu@bnu.edu.cn ;
± Zheng: Department of Economics, School of Economics and Management, Tsinghua University, Beijing, China, jie.academic@gmail.com ;
We are grateful to Vincent Crawford, Daniel Houser, Wei-Torng Juang, Hanjoon Michael Jung, Kamhon Kan, Botond Koszegi, Zhen Lei, David K. Levine, Mengyu Liang, Jacopo Magnani, Juanjuan Meng, Lionel Page, Stéphane Robin, Sen Tian, Neil Thakral, Chunlei Yang, Jipeng Zhang, and participants in the 3rd Tsinghua Conference on Theoretical and Behavioral Economics, the 6th SEBA-GATE Workshop, Asia Meeting of the Econometric Society 2016, Chengdu Economic Theory Workshop 2017, and seminar participants at Institute of Economics, Academia Sinica and School of Economics, Renmin University for helpful comments and encouragement. We gratefully acknowledge research funding from National Science Foundation of China, Hong Kong Research Grants Council, Tsinghua University, Beijing Normal University and The Chinese University of Hong Kong. All errors are our own.

1. Introduction

With the prevalence of internet-based commerce in modern society, consumers are provided with an ever-expanding set of new purchase methods and products to experience. In such a marketplace, the initial experiences of consumers are particularly important to firms, since the likelihood of earning a repeat customer may heavily depend on the consumer's prior experience. While conventional wisdom on customer service tells us that consumers with negative experiences are much less likely to return than consumers with positive experiences, precisely how customers' return tendencies vary by the actual degree of positive or negative experience obtained, has not typically been tested systematically.

Much of the challenge in studying the above question rigorously arises from a difficulty in assessing the actual values or surpluses associated with each consumer experience, and in collecting the data on a new consumer experience at the outset of its introduction to the marketplace. For example, in most settings (hotel stays, for example), consumers have multi-dimensional experiences (ex. check-in, cleaning service, food and beverage, etc.), and these factors weigh differently for different individuals into their overall experience. In addition, self-reported quality of experience can be difficult to compare quantitatively across individuals – for example, a 'fine' rating by one customer, may be equivalent in true satisfaction levels to a 'very good' rating for another customer.

These measurement and aggregation issues make it challenging to assess very precisely how important consumer behaviors such as retention, are determined by past experiences. Recent works such as Huang and Zhang (2020), and Qiu and Rao (2020) theoretically analyze how refund policies and cash back opportunities can be used by retailers in different settings to optimize consumer retention. These studies show the importance of the consumer retention problem for firms, and suggest strategies that firms may use in order to overcome consumers' potentially negative or loss experiences. In this paper, we utilize detailed bidding data from a penny auction website to address the underlying question of how prior experiences affect consumers' willingness to return to a relatively new purchase mechanism.

Penny auctions are an attractive setting for studying this question for two key reasons: Firstly, penny actions were a new phenomenon introduced to internet shoppers in the late 2000's. A large consumer base of internet shoppers were thus faced with the option of trying this new purchase mechanism, and subsequently had to decide whether to return and participate in the penny auction mechanism again. Second, as in most consumer experiences, participation in penny auctions entails a degree of effort and investment by the consumer, while yielding uncertain benefits or returns. For example, in the recent platform-based business models Airbnb and Uber, consumers exert some effort and cost in trying these services for the first time. One question is whether consumers will return for Airbnb or Uber experiences conditional on their experience in their first encounter (for example, the cleanliness of room, friendliness of host, and so on). However, in contrast to the examples of Airbnb and Uber, as well as numerous other types of consumer experiences, the experience in penny auctions can be objectively quantified using the monetary gain or loss which they experience in the auction process. We can thus determine how positive and negative consumer experiences correspond to the likelihood of consumers returning, across individual experiences.

We find apparent evidence of reference-dependence in consumers' propensity to return to the penny auction website, conditional on the monetary value of their prior experience – the empirical pattern of return rates seems to closely mirror that predicted by loss aversion and other reference-dependent models. That is, there is a discontinuity in the probability of returning at the nominal zero gains/losses value. In the framework of reference-dependence, small differences in outcomes can substantially change consumers' behavior, depending on whether the consumer experiences a loss or a gain. The shape of the return function has important implications for new enterprises seeking to earn repeat customers. Specifically, while larger loss magnitudes significantly reduce the consumer's likelihood of returning, larger gain magnitudes do not significantly increase the likelihood of return. We find this pattern robustly across nearly all product categories, across moderately experienced customers and brand new customers alike, as well as in the domain of bidding intensity. Our findings imply that to earn repeat customers, it is more important to avoid the consumer's feeling of loss than to increase their feeling of satisfaction through increased gains. This pattern contrasts with the evidence on loss aversion in purely gambling contexts (for example, Post, van den Assem, Baltussen and Thaler, 2008; Eil and Lien, 2014; Lien and Zheng, 2015) in which clear evidence of break-even effects are found, and suggests that penny auctions and gambling activities differ in crucial aspects from the consumer's point of view.

While our empirical findings hold an intuitive and seemingly obvious connection to the loss aversion model, we then rigorously derive the theoretical conditions under which a loss aversion model can deliver the result detected in the data. A key contribution of our model is to shed light on the conditions on the reference point updating function needed to produce the intuitive empirical result. It is shown that the crucial conditions apart from the standard reference-dependence formulation are: 1. discontinuity in consumer utility at the reference point if the reference point is negative and 2. a continuous reference point formation function, for which the reference point itself shares the same sign as the previous outcome, and the reference point adjustment after a negative outcome is less severe than that for a corresponding positive outcome. These conditions are psychologically plausible and consistent with an adaptively updating expectations-based reference dependence. Our model thus advances the theoretical literature on reference-dependence by applying it to consumer experiences, potentially one of the most prevalent domains where reference-dependent behavior can be found in real world settings.

Our study contributes to the literature on consumer applications of reference-dependent preferences and empirical evidence on the importance of considering the psychology of gains and losses in consumer decision-making contexts. Heidhues and Koszegi (2014) study the pricing strategy of a monopolist selling to a loss averse consumer with a rational expectations reference point. They find that consumer loss aversion introduces an incentive for the monopolist to introduce risk into his pricing strategy, specifically by implementing fluctuating low sales prices and high regular prices. Camerer, Ray and Shum (2015) examine internet purchases and find that once a consumer knows that a product has previously been offered at a lower price, he or she is much less likely to purchase that item.

Our study is one of few to examine the impact of gains and loss magnitudes in a consumer's experience on their likelihood of becoming a repeat consumer. Backus, Blake, Masterov and Tadelis (2017) examines eBay auctions and finds that bidders who lose auctions after having more "virtual attachment" via being the leading bidder for a longer time period, are more likely to subsequently exit the platform completely. Our study complements and differs from theirs in several key respects. Firstly, our analysis focuses on the return likelihood of consumers as a function of monetary gains and losses, whereas Backus et al. (2017) focuses on attrition from the eBay auction platform due to psychological losses. The difference in auction format between our field setting (all-pay) and theirs (non all-pay) allows appropriate testing of each phenomenon, respectively.[2] Secondly, their model and empirical findings support expectations-based reference dependence, where expectations are raised in response to being a leading bidder for a longer period of time. As the penny auction is all-pay in nature, we focus on monetary loss and gain rather than purely psychological measures, and find evidence of a kink and discontinuity at the lagged status quo or nominal zero monetary level. Our theoretical analysis in turn shows that in order to support our empirical finding, some conditions

[2] In particular, the second price auction format in eBay, combined with the Buy It Now option which can end the auction directly, creates a situation where all non-winning bidders in the auction will experience zero monetary losses and gains, but may experience psychological or "virtual" losses. In the penny auction setting by contrast, non-winning bidders do experience direct monetary losses.

on reference point formation are needed. Our paper shares with Backus et al. (2017) in documenting important consumer retention implications of loss aversion in online purchase platforms. Chang, Chen and Hsu (2010), uses survey methods to assess the relationship between service quality and post-dining return intentions in the restaurant sector, and also reports evidence of loss aversion around consumers' expectations, in an empirical pattern notably similar to what we find in the penny auction data. While their study lends support to our findings, our study improves upon their line of inquiry by utilizing an objective measure of gains and losses, and further shows the theoretical conditions needed to produce such empirical patterns in a reference-dependent model of consumer return behavior.

Several studies have highlighted the importance of gains versus losses as motivators in non-consumer purchase settings. Berger and Pope (2011) and Pope and Schweitzer (2009) find that professional sports players in basketball and golf respectively, respond with higher performance to being marginally below a reference point compared to being above it. Camerer, Babcock, Lowenstein and Thaler (1997) and Crawford and Meng (2011) find evidence for reference-dependent loss aversion in the labor supply of taxi drivers. Abeler, Falk, Goette and Huffman (2011) find evidence for reference-dependent effort provision around an expectations-based reference point in an experimental setting, while Gill and Prowse (2012) find evidence for disappointment aversion in an experimental real effort tournament. Eil and Lien (2014) and Lien and Zheng (2015) find that poker and slot machine players respectively, are most prone to keep playing while they are posting a loss. Behavioral finance literature examines the effects of loss aversion on asset purchase behavior, including Meng and Weng (2017) which shows that expectations-based reference-dependence with lagged expected final wealth as the reference point can predict the disposition effect. Using a rational bubble framework, Zhang and Zheng (2017) show that reference-dependence can make an asset bubble more robust than under classical preferences.

Our study contributes to the empirical evidence on reference-dependence in three main respects. First, we find that consistently with Berger and Pope (2011) and Lien and Zheng (2015), the nominal break-even level appears to serve as an influential reference point, around which decisions are substantially different at the margin. Second, we find that in the domain of consumer experiences, losses do not encourage the consumer to come back again, to try harder, or perform better, in contrast to the previous literatures on reference-dependence in gambling, labor supply or effort settings. Third, we find that while consumer gains at the margin of nominally breaking even does lead to a higher likelihood of becoming a repeat consumer, marginal gains farther from the reference point do not substantially increase the chance of coming back as one would expect in a house-money effect. This is somewhat surprising given the existing evidence that individuals tend to be more 'care-free' with their money when experiencing a gain. It implies that in the domain of consumer experience, preventing the consumer's negative experience seems to be the most important factor in retaining consumers.

We further investigate the bracketing tendencies of consumers in their decision to return to play penny auctions based on initial decisions. We find evidence that a consumer's most recent experience tends to be most influential in their decision on whether or not to return. That is, replicating all our aforementioned findings in terms of the general shape of the customer return function, consumers are most *sensitive* to gains and losses which have just occurred in the most recent penny auction compared to any previous penny auctions they participated in, including their very first experience. In other words, the weighting of previous experience tends to be monotonically decreasing over time. This differs from some previous findings on reference point formation which find a U-shaped pattern of weights over time (see for example, Baucells, Weber and Welfens, 2011). It also differs from the intuition that first impressions matter most (Rabin and Schrag, 1999), while supporting the idea of Recency Bias (Fudenberg and Levine, 2014). This empirical finding suggests that firms should be very concerned about even single incidents of bad consumer experiences, as consumers tend to have short-term memories, and are likely to be driven away by such incidents. Our research question and findings contrast with List (2003, 2004), which focuses on the question of whether reference-dependent behavior fades over longer-term experiences in the marketplace, and finds supporting evidence.

Finally, our research contributes to the small but emerging literature on penny auctions as an auction mechanism, but differs from this literature in our focus on across-auction behavior rather than within-auction behavior. Augenblick (2016) finds evidence of the sunk cost fallacy in within-auction bidding behavior in penny auctions. Platt, Price and Tappen (2013) rationalize bidding behavior in penny auctions using bidders' risk preferences which are potentially risk-loving. Wang and Xu (2016) find that in order for the penny auctioneer to sustain large profits, he must keep attracting new and naïve consumers who lose money in the auction. In contrast to these studies, our research focuses on the likelihood of consumers returning to play again, as a function of their previous gains and losses, and we do not theoretically or empirically analyze the penny auction mechanism itself. Instead, we exploit the penny auction as a prototype for consumer experience of a new purchase mechanism, where the benefit of doing so is that positive and negative experiences can be easily and objectively measured in monetary terms.

The remainder of the paper is organized as follows: Section 2 describes the penny auction mechanism and our data; Section 3 presents our empirical results and robustness checks; Section 4 provides a model of reference-dependence in consumer return decisions, deriving conditions under which the empirical patterns can be accounted for by reference-dependent consumers; Section 5 explores how consumers bracket their gains and losses over time; Section 6 concludes.

2. Penny Auctions and Data

Penny Auction websites were first introduced in the United States in 2008 by the German penny auction firm Swoopo. The penny auction industry in the US subsequently expanded to as many as 125 different websites by 2010 (see Wang and Xu, 2016, for details). Sometimes controversial as a consumer purchasing mechanism, penny auction websites are often labeled as "entertainment shopping", and often advertise the ability of consumers to purchase high value items for very low prices. Indeed, penny auctions offer a wide range of good and bad outcomes for participants: in the best case scenario, receiving a remarkably great deal on a desired item, and in the worst case scenario, spending large amounts of non-refundable bids and not winning the desired item. In penny auctions, the winning chance is typically very small, thus the worst case scenario of losing one's bids is highly likely to occur.

Our data is from the penny auction website BigDeal.com, which was in operation from November 2009 until August 2011. During its operation period, BigDeal was one of the most popular penny auction websites, and was one of few of the penny auction websites to be rated by the Better Business Bureau (BBB) as a reliable business endeavor. We collect detailed consumer-level bidding data from each auction of BigDeal during the entire time horizon of its operation, which allows us to construct a panel dataset of individual bidders over the entire course of their participation on the website.

As a typical penny auction website, the rules of BigDeal.com are specified as follows. First, the website releases the item-for-sale one or two days in advance with detailed product descriptions and the suggested retail price. Items typically sold include computers, TVs, video games and consoles, Apple products, non-Apple electronics, gift cards, handbags, jewelry, and so on. Second, any potential consumer needs to buy bid-tokens before bidding. Each bid-token costs $0.75 and the bid-tokens are sold in packs.[3] Third, bidders place their bid-tokens one-by-one to bid for the item.

[3] BigDeal sold bid-tokens in packs of 30, 50, 100, 200, 300, and 500 tokens. Unused tokens are refundable at the price of $0.75 per token.

As a bidder bids a token, he or she becomes the current "leader" in the auction and the bid token is non-refundable. Fourth, the auction price begins at zero, and the increment of price increase is specified and fixed as $0.01, $0.05 or $0.15. Fifth, a bidder wins if a commonly-observed countdown timer hits zero, and he or she is the last bidder to place a bid. Each bid placed extends the countdown timer to an additional 30 seconds, during which any consumer can bid. In addition, the website offers a "buy-it-now" option to losing bidders in many of the auctions. This option allows losing bidders to purchase the target product at the retail price, where the sunk cost of tokens bid during that particular auction can be used as credit towards the item.[4]

BigDeal operated in its standard manner until April 30, 2011. When the website decided to shut down, it drastically reduced the daily supply of auctions. To avoid this regime change which occurred in the late stages of the website operation, we limit our sample period for this study from the starting date of the website, November 19, 2009, through April 30, 2011. During this period, there were 207,069 unique consumers who participated in 107,219 auctions.[5] On average, each consumer participated in 8.32 auctions, and the average probability of loss in each auction is 93%. Table 1 describes the summary statistics for variables in the sample.

Apple products are one of the most popular items on BigDeal, with 4.27% of all auctions selling Apple products, such as iPad, iPhone, iPod, and so forth. BigDeal also sold packs of bid-tokens via penny auction, and 42.64% of all auctions on the website were bid-token auctions. BigDeal offered a special type of auction to new bidders, called "beginner auctions", which were open exclusively to bidders who had not yet won their first auction. Most of the beginner auctions were selling a 10 bid-token pack or 20 bid-token pack.[6]

To focus on the decision to return or quit, we construct a panel data set of individual bidders' data over all the auctions they took part in. Each bidder's experience is ordered based on the ending time of the auctions he or she participated in. For each bidder-auction pair, we observe the information about the auction, including the name of the product, retail price, bid increment, auction price, the winner's screen name, ending time, buy-it-now availability, and individual bidder information including the number of tokens placed. We group the products into 16 categories based on the type of product: Apple products, Bid-Token packs, Gift Cards, Housewares, Non-Apple computer, other electronics, TV, Video Games, Handbags, Health and Beauty products, Jewelry, Movies, Non-Apple mobile phone, Sunglasses, Toys, and Watches.

---

[4] Sunk tokens cannot be accumulated across auctions towards any item's buy-it-now price. This option can only be exercised for a specific auction and the tokens an individual bidder placed within that auction.

[5] We use screen names to identify players. Each screen name is linked to a unique email account and user account on the website.

[6] The website may have offered the beginner auctions in order to attempt to bring new customers into the penny auction experience, or to bring back previously unsuccessful bidders. Bidders may perceive less competition due to the restriction of the auctions to players who have not yet won an auction. In addition, since most of the beginner auctions were for bid-tokens, winners may be likely to return to the website to participate in another auction in order to use them. However, as our later analysis shows, all else equal, participants in beginner auctions were actually less likely to return.

Table 1: Summary Statistics

Variable	Mean	Standard deviation
Number of auctions	107,219	--
Number of unique bidders	207,069	--
Bidder's duration (No. of auctions)	8.32	23.22
Auction price	$7.06	31.26
Retail price	$113.00	203.84
Average probability of loss	0.93	0.25
Percent of Apple products	4.27%	--
Percent of Token auctions	42.64%	--
Percent of Beginner auctions	26.72%	--
Percent of buy-it-now	88.49%	--

3. Empirical Strategy and Results

We test two main issues in our empirical analysis. The first question is: Do consumers respond differently to gains and losses in their decision to become a repeat consumer, and what is the approximate shape of the return likelihood? The second question we address is about the strength of the main empirical pattern as consumers become slightly more experienced: How do repeat consumers bracket their sequence of previous gains and losses in order to make a decision about whether to return?

We use two approaches in order to test the robustness of results for our first question. Since we are interested in the issue of what makes consumers return, we first examine the participation decision in the first ten rounds of penny auction participation round-by-round in a logit regression discontinuity framework. We then control for the underlying tendency for individuals to stop playing in the penny auctions using hazard analysis on the entire dataset, to check whether these patterns persist on average, regardless of the number of auctions participated in.

We argue that heterogeneity of consumer preferences and selection effects cannot explain the patterns found in the data. The reason is that bidding more in a penny auction and/or participating in more penny auctions is generally not sufficiently correlated with winning the auction, round by round. This is consistent with the perception that the penny auction has a large element of luck involved. Given the structure of the penny auction mechanism, placement in the gains vs. loss domain can be reasonably ascribed to chance in the sense that winning or losing the auction itself can be due to random factors (ex. the persistence and attention characteristics of opponents, and so on), however, one's exact placement *within* the range of losses or gains will be determined by the intensity of bidding.[7] In this sense, each observation in the loss domain has a counterpart in the gain domain conditional on a consumer's bidding amount, and the more strictly valid comparison is between any of these pairs of points. By empirically verifying both a significant discontinuity at the nominal break-even point which distinguishes winning from losing (in the absence of overbidding beyond the product's value), as well as a significant slope in loss domain as compared to the gains domain, our empirical approach serves effectively as test of all points in the loss domain collectively,

[7] Meaning that a bidder who bids a total of b for an item worth v, obtains a monetary amount of v-b if they win, but obtains -b if they lose.

against those in the gain domain. We further confirm the limited role for selection effects in the data by conducting robustness checks and investigation on the role of experiences and losses, provided in Section 6. Furthermore, our clearest evidence of the loss aversion pattern always comes from the most recent auction that consumers participate in, indicating that the effects of prior experiences are strongest for consumers when they consider their gains and losses in the most recent experience. This suggests that any effects of heterogeneous consumers in terms of preferences, skill, or learning are dominated by the impact of the most recent monetary loss and gain.

To answer our questions about bracketing, we arbitarily consider the decisions of consumers who have played 5 auctions, and those who have played 10 auctions, respectively. Including each prior auction's outcome in terms of experienced nominal gains and losses, shows us which auction's losses are most influential in the consumer's decision regarding whether to return. The results for other numbers of auctions are very similar and are provided in the Appendix.

Throughout the empirical analysis, references to "gains" and "losses" denote nominal gains and losses around the zero benchmark of consumer surplus, and not the gains and losses around particular reference points, as discussed in our theoretical model. Recall that the theoretical model, which includes formation of the reference point based on prior experiences, provides the necessary and sufficient conditions for the exact empirical pattern we find, including the kink and discontinuity in the return ratio at the zero consumer surplus benchmark.

3.1 Round-by-Round Analysis

Our empirical analysis focuses on the nominal gain or loss based on the consumer surplus implied by the market value of the product. We calculate the nominal gain or loss as the difference of product's market value and bidder's payment.

For Bid-Token packs, we set the monetary value of a single token as defined by the website, at \$0.75. For all other products, we search for the identical item on Amazon.com in June of 2011 to assess the market value. We find about 61.7% of the non-token BigDeal items available at Amazon. For these products, the Amazon prices are the proxy of the product value, and the Amazon prices are on average about 78.0% of the retail price listed on BigDeal. Therefore, for those non-token products which we cannot find a match on Amazon.com, we assume the product value to be 78.0% of the retail price posted on BigDeal.[8]

If a bidder wins the auction by being the last bidder to place a token in the auction, his or her total payment is the auction price plus the cost of his or her bid-tokens used. The difference between the product value and the total payment made defines the gain or loss. If a bidder loses the auction and the buy-it-now option is not available, her loss is simply calculated as the cost of bid-tokens, and it is well defined.[9]

A potential complication arises when the losing bidder faces the buy-it-now option. From the data, we cannot observe whether the losing bidder actually exercises the buy-it-now option. Assume that the losing bidder's cost of tokens bid is c, and the value of the product being auctioned is v. Given the Big Deal retail price p, we assume the losing bidder exercises the buy-it-now option if $v \geq p - c$. If a losing bidder does not exercise the buy-it-now option, her loss is the cost of bid-tokens. If a losing bidder exercises the buy-it-now option, her loss is the difference between the buy-it-now retail price and the product value.

8 Bid Tokens can be exchanged back for cash on the BigDeal website, at \$0.75 per Token.

9 A small fraction of auctions (about 4,394/107,219) were won by bidders who 'overbid' for the item in question, in the sense that the value of their tokens bid exceeded the value of the product. Thus, we interpret our results on the loss side as being driven by auction losses, rather than by auction winners accompanied by monetary losses.

To check the sensitivity of our calculation of gain or loss, we also try using the buy-it-now retail price to proxy for the product value rather than the price on Amazon, and we assume various different thresholds for losing bidders to exercise the buy-it-now option at 10%, 20%, or 30% of the posted retail price. We find the results are highly correlated with our benchmark calculation, with the Spearman rank order correlation coefficients among all these measures lying between 94% to 98%. Thus, we interpret our main calculation as well-defined, and utilize it in the remainder of our analysis. The results under the robustness checks are available upon request.

Our measure of a consumer returning is whether he or she placed any bid in a subsequent auction. While this may be an imperfect proxy for whether the consumer was 'interested' in participating in another auction, it is the measure which tends to matter most for the website's business model.[10] Furthermore, since this measure in fact indicates whether the consumer *ever* returned to BigDeal over the entire life of the website, it seems to be a reasonable proxy. In this setting, a consumer who was interested in several subsequent auctions but did not bid in any of them, is of limited benefit to the company.

Figure 1: Bidders' Return Ratio after the First Try by Nominal Gains and Losses

(x-axis shows nominal gain and loss subgroups, by 3 dollar increments)

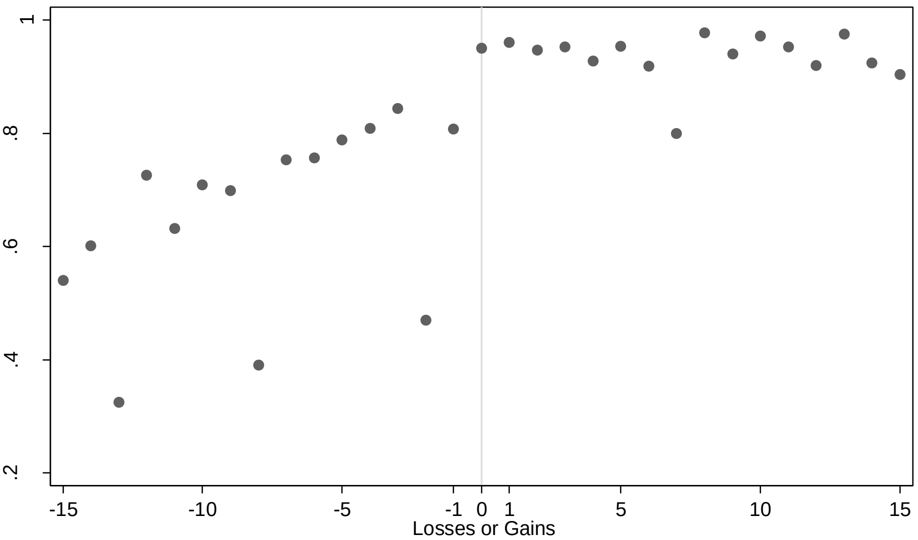

Figure 1 shows the relationship between bidder's return ratio after the first auction and his gains or losses in the first auction. If a bidder participated in two or more auctions, he is defined as a returning bidder after the first experience. To smooth the return ratio, we organize the gains and losses into thirty-one subgroups: group = 0 if gain or loss equals zero; group = 1 if gain is from 0 to 3 dollars; group = 2 if gain is from 3 to 6 dollars;…; group = 14 if gain is from 39 to 42 dollars; group = 15 if gain is above 42 dollars; group = -1 if loss is from 0 to 3 dollars; group = -2 if loss is from 3 to 6 dollars;…; group = -15 if loss is above 42 dollars.

From the return ratio figure, we can clearly see that losing bidders are more likely to quit permanently compared to winning bidders. For losing bidders, the return ratio decreases over the magnitude of losses. For winning bidders, the return ratio is quite high and may not significantly increase in trend with incremental gains. Figure 1 aggregates the return ratio across all product categories. Additional figures on the return ratio as a function of gains and losses separated by product category are provided in Appendix A.[11]In order to estimate the shape of the return ratio

[10] For example, some users may keep watch on an auction without actually placing any bids.

[11] While in the aggregate shown in Figure 1, the average trend line in the gains domain does not fully reach and remain at the return ratio ceiling, the figures in Appendix A show that different categories of products are heterogeneous in this

function more accurately, we implement a logit regression discontinuity estimation with linear segments, following the empirical approach in Berger and Pope (2011) in detecting loss aversion in NBA and NCAA basketball games.[12] Since our main interest is in estimating the entire return likelihood function across surplus levels, we present the estimation results with unrestricted sample in the current section. However, as implemented in Berger and Pope (2011), to test specifically whether the discontinuity at the nominal zero mark is robust to potential non-linearities in the gain and loss domains, we conduct robustness checks restricting the sample to much narrower gain and loss windows ($10, $50, $100, and $200). The loss coefficient remains significantly negative throughout and the results are presented in Appendix E.

Our empirical specification for the decision-maker deciding whether to participate in the auction for the t^{th} time is as follows:

$$return_{i,t} = \alpha + \beta \cdot loss_{i,t-1} + \gamma \cdot (loss_{i,t-1} \cdot magnitude_{i,t-1}) + \delta \cdot magnitude_{i,t-1} + \eta \cdot X_{i,t-1} + \varepsilon_i$$

where $return_{i,t}$ is a binary variable equal to 1 if bidder i comes back to participate for the t^{th} time, $loss_{i,t-1}$ is a binary variable equal to 1 if bidder incurs a net loss in the previous auction played, $magnitude_{i,t-1}$ is the absolute value of bidder's net gain or loss.

We also control for some prior features faced by the individual bidder denoted by the vector , $X_{i,t-1}$, which contains total profit prior to the most recent auction to control for possible "income" effects, log retail price as a proxy for value, an indicator variable for Apple product, an indicator for whether the previous auction was a Token auction, an indicator for beginner auctions, and an indicator for late night auctions (i.e., those auctions scheduled to start during late-night hours, from 1 a.m. Pacific Standard Time (PST) through 4 a.m. PST).[13] We control for the value of the object by including the retail price in the regression. Due to the popularity of Apple brand products on the website, we include a dummy variable for whether the previous auction the customer participated in was for an Apple product. We also include a dummy variable for whether the previous auction participated in was an auction for bid tokens, which could substantially affect the decision to come back, given that bid tokens are the currency needed for future play. We control the effect of beginner auctions, which provides higher probability of winning due to less competition. Finally, we include a dummy variable for whether the previous auction occurred at night, which could influence the behavioral traits of the consumer through selection effects. The variables in $X_{i,t-1}$ are included incrementally to potentially help us in better understanding possible selection effects in our coefficients on gains and losses, which are our main results of interest.

We focus on the most recent gains and losses in this particular specification, because we are primarily interested in understanding consumers' return tendencies after initial experiences (i.e., a relatively small number of rounds). We run the aforementioned round-by-round regression for 10

regard. This raises a feature regarding the upper limit of the response variable: when considering only a single discrete return decision (return or not), consumers have a limited action space over which to express their reaction to gains and losses, truncated in the aggregate at ratios of 1 and 0, respectively. While this essentially the nature of the consumer return question, later on in Section 6.2, we examine the behavior on the intensive (bidding) margin, and still find significant effects of loss experiences.

[12] We consider linear relationships between monetary loss/gain and the return ratio, for simplicity, and because in our return ratio plots, the relationship seems well-characterized without higher-order polynomials. However, specifications with higher order polynomial terms could be implemented without drastically changing the direction and magnitudes of our results.

[13] We classify these hours as "night auctions" since they capture the hours during which the entire continental United States is typically sleeping. The website also offered substantially fewer auctions during these hours, compared to during the daytime.

rounds to check how long the effects persist. In subsequent sections, we estimate single coefficients for all rounds together in a hazard framework, and explore richer possible bracketing dynamics for gains and losses across rounds.

Tables 2 through 4 show our estimation results for the round-by-round analysis as specified above. Each column shows the results of a logit estimation for consumers' likelihood of returning after the t^{th} round. Table 2 shows the unconditional results without any control variables, Table 3 controls for the cumulative profits of the individual bidder, and Table 4 controls for the five additional variables regarding the most recent auction played: retail price, whether the auction is for an Apple product, whether it is a Token Auction, whether it is a beginner auction, and whether it is a late-night auction.

Table 2 shows that for any number of return experiences, "losses loom larger than gains". The coefficient on the interaction term $loss_{i,t-1} \cdot magnitude_{i,t-1}$ shows that as losses are larger, the likelihood of returning is decreasing on the loss side. The corresponding effect on the gains side on the other hand, is indistinguishable from zero. In addition, the coefficient on *loss* shows that there is a discontinuity at the zero mark. A previous monetary loss discontinuously reduces the likelihood of returning to purchase again via penny auction. We can also notice that the magnitudes of these effects are diminishing with the number of auctions participated in, suggestive of selection effects via behavioral heterogeneity. In other words, bidders who stay for several auctions may have other considerations which more heavily influence their decisions, compared to the immediate gains and losses experienced.

Table 3 shows that the results in Table 2 are robust to the accumulation of previous gains and losses of the bidder. While the accumulation of previous profits of the consumer has a significant positive effect on the likelihood of returning, the coefficients on the gains and losses of the most recent auction experience are of similar magnitude as in Table 2.

Table 4 shows the analogous results when including several control variables. We gain several insights from the addition of these control variables. First, the coefficients on price of the item, and whether the product was an Apple brand item, are consistently negatively significant on the decision to return. We believe this has to do with the appeal of penny auctions in carrying the possibility of obtaining high value items such as Apple products among others, at an especially low price. While many consumers are initially drawn to this feature in their participation decision, they are less likely to return for future auctions. Second, the inclusion of these two variables as controls tends to reduce the magnitude of the coefficient on the loss indicator, while making the coefficient on magnitude as it applies to gains, significantly positive with a very small slope. As in Table 2, the coefficients still show that losses loom larger than gains for penny auction consumers. Third, the effect of beginner auctions is negative on the return ratio, especially for very early experiences. This implies that the success of the beginner auction as a tool for acquiring or retaining consumers was limited.

Table 2: Logit Model, Returning After t^{th} Experience

Number of Auctions Participated, N	Dependent variable: Return = 1 if the bidder returned after the t^{th} time									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Loss	-0.188***	-0.145***	-0.125***	-0.102***	-0.090***	-0.075***	-0.074***	-0.071***	-0.072***	-0.074***
	(0.061)	(0.047)	(0.038)	(0.034)	(0.031)	(0.029)	(0.028)	(0.028)	(0.019)	(0.029)
Loss*magnitude of gain or loss/100	-0.275***	-0.174***	-0.139***	-0.144***	-0.133***	-0.147***	-0.125***	-0.107***	-0.111***	-0.072***
	(0.065)	(0.027)	(0.030)	(0.039)	(0.025)	(0.024)	(0.035)	(0.022)	(0.026)	(0.026)
Magnitude of gain or loss/100	-0.016	0.035	0.007	0.019	0.015	0.020	0.030	0.014	0.023	0.014
	(0.015)	(0.026)	(0.016)	(0.026)	(0.020)	(0.020)	(0.027)	(0.014)	(0.021)	(0.016)
Pseudo-R^2	0.0235	0.0125	0.0110	0.0115	0.0104	0.0136	0.0115	0.0100	0.0114	0.0068
Observations	207,069	155,188	123,465	101,334	86,028	73,555	64,585	57,114	50,441	44,848

Notes: Marginal-effects reported, standard errors clustered by product category; *significant at 10%; **significant at 5%; ***significant at 1%. The column (t) shows the results of returning after the t^{th} auction. Loss and gain refers to nominal losses and gains around zero consumer surplus.

Table 3: Logit Model, Returning After t^{th} Experience

Number of Auctions Participated, N	Dependent variable: Return = 1 if the bidder returned after the t^{th} time									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Loss	-0.188***	-0.144***	-0.124***	-0.101***	-0.089***	-0.073***	-0.072***	-0.069**	-0.071***	-0.073***
	(0.061)	(0.046)	(0.038)	(0.033)	(0.031)	(0.028)	(0.027)	(0.027)	(0.018)	(0.028)
Loss*magnitude of gain or loss/100	-0.275***	-0.169***	-0.135***	-0.139***	-0.130***	-0.146***	-0.120***	-0.104***	-0.102***	-0.068***
	(0.065)	(0.027)	(0.029)	(0.037)	(0.025)	(0.024)	(0.035)	(0.023)	(0.026)	(0.026)
Magnitude of gain or loss/100	-0.016	0.033	0.008	0.019	0.014	0.023	0.029	0.013	0.019	0.010
	(0.015)	(0.026)	(0.015)	(0.025)	(0.020)	(0.020)	(0.027)	(0.014)	(0.020)	(0.016)
Cumulative profit/100	--	0.104***	0.063***	0.056***	0.040***	0.042***	0.033***	0.032***	0.023***	0.026***
		(0.012)	(0.005)	(0.007)	(0.008)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)
Pseudo-R^2	0.0235	0.0147	0.0133	0.0152	0.0134	0.0193	0.0163	0.0155	0.0155	0.0118
Observations	207,069	155,188	123,465	101,334	86,028	73,555	64,585	57,114	50,441	44,848

Notes: Marginal-effects reported, standard errors clustered by product category; *significant at 10%; **significant at 5%; ***significant at 1%. The column (t) shows the results of returning after the t^{th} auction. Loss and gain refers to nominal losses and gains around zero consumer surplus.

Table 4: Logit Model, Returning After t^th Experience

Number of Auctions Participated, N	Dependent variable: Return = 1 if the bidder returned after the t^th time									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Loss	-0.047**	-0.034*	-0.047***	-0.045**	-0.037**	-0.017	-0.027**	-0.031**	-0.035**	-0.036*
	(0.020)	(0.019)	(0.017)	(0.021)	(0.017)	(0.018)	(0.013)	(0.013)	(0.014)	(0.019)
Loss*magnitude of gain or loss/100	-0.169***	-0.189***	-0.117***	-0.121**	-0.111***	-0.140***	-0.118**	-0.091***	-0.101***	-0.058**
	(0.058)	(0.032)	(0.043)	(0.058)	(0.035)	(0.048)	(0.051)	(0.033)	(0.026)	(0.029)
Magnitude of gain or loss/100	0.076***	0.163***	0.076***	0.068*	0.057***	0.071**	0.071*	0.042**	0.054***	0.036**
	(0.027)	(0.040)	(0.021)	(0.037)	(0.023)	(0.031)	(0.037)	(0.017)	(0.021)	(0.018)
Cumulative profit/100	--	0.064***	0.042***	0.039***	0.028***	0.031***	0.025***	0.025***	0.018***	0.021***
		(0.008)	(0.008)	(0.004)	(0.002)	(0.004)	(0.004)	(0.002)	(0.002)	(0.002)
Log retail price	-0.073***	-0.058***	-0.047***	-0.040***	-0.036***	-0.031***	-0.029***	-0.028***	-0.026***	-0.033***
	(0.010)	(0.008)	(0.006)	(0.006)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Apple product	-0.086***	-0.077***	-0.059***	-0.051***	-0.047***	-0.044***	-0.044***	-0.039***	-0.031***	-0.032***
	(0.019)	(0.012)	(0.007)	(0.006)	(0.004)	(0.005)	(0.006)	(0.005)	(0.005)	(0.006)
Token auction	0.042*	0.013	0.018*	0.016	0.027***	0.009	0.009	0.015**	0.012*	0.016**
	(0.021)	(0.016)	(0.011)	(0.010)	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)	(0.007)
Beginner auction	-0.034**	-0.031**	-0.024*	-0.055***	-0.057***	-0.003	-0.011	-0.028**	-0.010	-0.076***
	(0.016)	(0.013)	(0.012)	(0.014)	(0.013)	(0.010)	(0.009)	(0.012)	(0.006)	(0.011)
Night auction	-0.014	-0.007	-0.009	-0.006	-0.010	0.002	-0.005	0.001	-0.0004	0.003
	(0.013)	(0.006)	(0.011)	(0.005)	(0.009)	(0.004)	(0.006)	(0.006)	(0.006)	(0.006)
Pseudo-R^2	0.1157	0.0862	0.0702	0.0661	0.0636	0.0653	0.0612	0.0572	0.0527	0.0549
Observations	207,069	155,188	123,465	101,334	86,028	73,555	64,585	57,114	50,441	44,848

Notes: Marginal-effects reported, standard errors clustered by product category; *significant at 10%; **significant at 5%; ***significant at 1%. The column (t) shows the results of returning after the t^th auction. Loss and gain refers to nominal losses and gains around zero consumer surplus.

3.2 Hazard Analysis

While the round-by-round results convince us that gains and losses play a critical role in consumers' decisions on whether to return or not, we are interested in conducting a systematic analysis on the whole sample, extending beyond just the initial 10 rounds. A hazard analysis is a convenient way to summarize the effects of gains and losses, while accounting for consumers' natural rate of leaving the penny auction website. Eil and Lien (2014) use a hazard estimation in evaluating how poker player's gains and losses affect their decision to stop playing.

Since we restricted the sample period to BigDeal's main operating dates, it raises a potential concern regarding the censoring issue. That is, bidders might indeed come back later, but do not appear in our main sample period. In our data, 90% of bidders' interactions on BigDeal last less than a month, 80% of bidders never come back after more than a week, and 50% of bidders only show up on the website within a single day. Given these relatively short time horizons, we presume that the censoring issue is not crucial for our main analysis. First, non-parametric analysis of the survival rates show that consumers experiencing a loss, exit from the website more rapidly than when experiencing a gain. Figure 2 plots the Kaplan-Meier survival function using the penny auction data, with number of auctions participated in on the horizontal axis, showing that the survival function indeed drops off much more sharply for the loss group. Similarly, the figures show that the cumulative hazard function is substantially higher for losses as compared to gains. Restricting the sample to 100 auctions participated or less (bottom row), the differences are even more pronounced.

Figure 2: Kaplan-Meier Survival Estimates and Nelson-Aalen Cumulative Hazard Estimates
Top row: All auctions; Bottom row: sample restricted to 100 or fewer auctions for each bidder

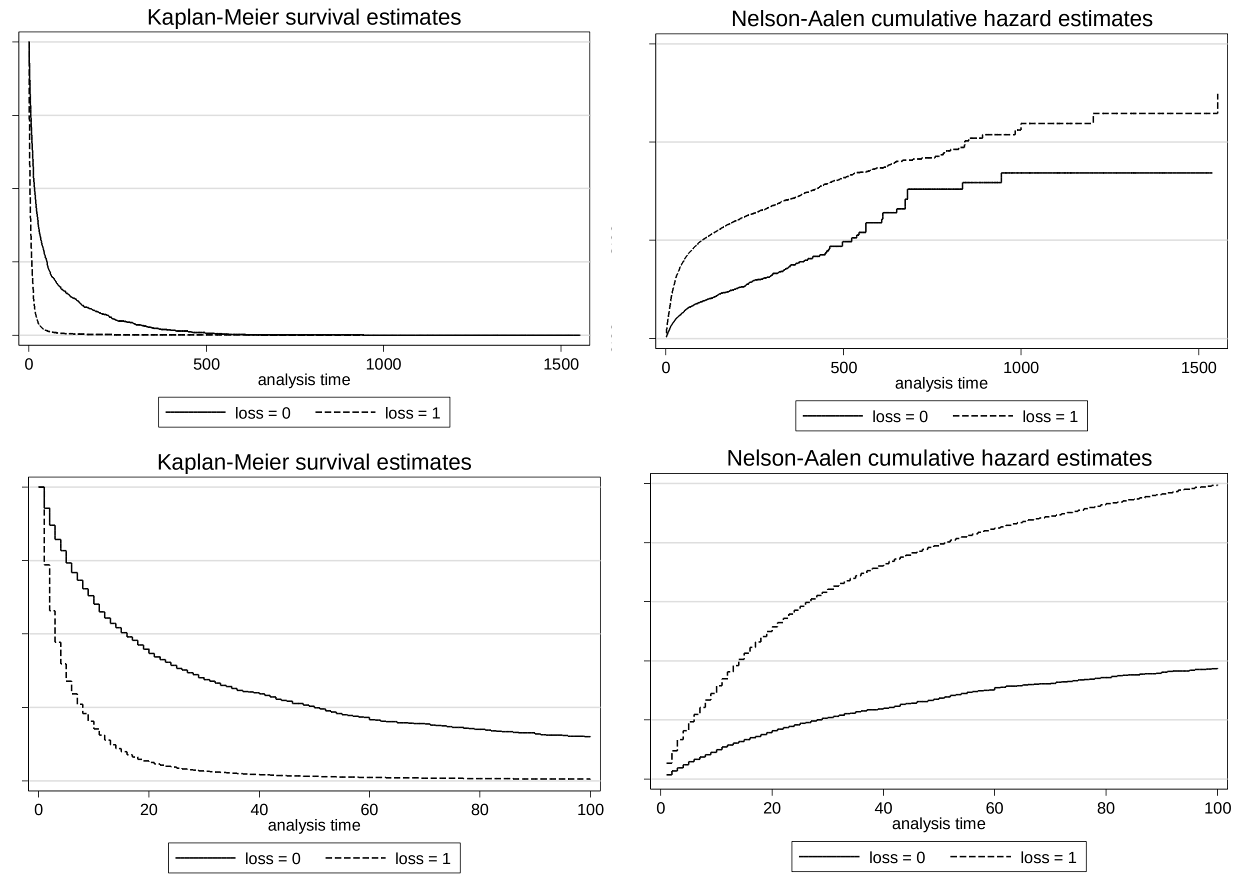

Although the non-parametric analysis can tell us the relative hazard functions for losses compared to gains at every level of experience, it cannot efficiently tell us about the precise shape of the hazard as a function of loss and gain magnitudes. We next consider the semi-parametric Cox proportional hazard model, as well as a parametric version under a Weibull distribution assumption, as a robustness check. The parametric version of the model has the drawback of specifying a particular distributional form for the baseline hazard, while having the benefit of allowing for time-varying hazard rates. The potential drawback to the Cox model is that potential time trends are unaccounted for, while the benefit is that no particular distributional assumption on the baseline hazard function is required.[14] Our main coefficients of interest are the variables on gains and losses, so as long as these estimates do not vary widely between the two models, our formulation is well-specified.

The hazard ratios have the interpretation of marginal effects in the corresponding covariates. For instance, if the hazard ratio for loss is 3.15, then the change from gain to loss increases the hazard by 215%. The results for both models are very similar in magnitude and significance levels as shown in Table 5. In particular, the coefficient on the loss indicator remains significant throughout, with a large effect of experiencing a loss on quitting. The marginal impact of loss magnitudes also remains positive and significant in consumers' decision not to return. The coefficients on the other variables are consistent with what we find in the round-by-round analysis. We note that the Weibull distribution parameter estimated is less than one, which accords with our intuition that the hazard rate decreases over the number of rounds, with the largest proportions of consumers exiting after short experiences. To check the robustness of these results in Table 5, we also examine the analyses restricted to the sub-sample in which all bidders' experience less than 100 auctions, and the results are close to those displayed in Table 5.

Table 5: Proportional Hazard Estimation

	Dependent variable: Return = 1 if the bidder came back					
	Cox	Cox	Cox	Weibull	Weibull	Weibull
Loss	3.152***	3.079***	1.485***	3.462***	3.189***	1.606***
	(0.049)	(0.052)	(0.031)	(0.055)	(0.058)	(0.037)
Loss*magnitude		1.406***	1.706***		1.528***	1.830***
of gain or loss/100		(0.034)	(0.057)		(0.041)	(0.074)
Magnitude of gain		1.003	0.711***		0.867***	0.597***
or loss/100		(0.017)	(0.023)		(0.021)	(0.024)
Log retail price			1.331***			1.361***
			(0.003)			(0.004)
Apple product			1.409***			1.490***
			(0.006)			(0.007)
Token auction			0.854***			0.867***
			(0.006)			(0.007)
Beginner auction			1.094***			1.392***
			(0.017)			(0.023)
Night auction			1.016			0.989
			(0.010)			(0.011)
Rho (Weibull)				0.7923	0.7901	0.8172
Observations	1,697,192	1,697,192	1,697,192	1,697,192	1,697,192	1,697,192

Notes: Hazard ratio reported, standard errors clustered at the consumer level; *significant at 10%; **significant at 5%; ***significant at 1%. Loss and gain refers to nominal losses and gains around zero consumer surplus.

4. Theoretical Framework

In this section, we provide a theoretical framework which can generate the patterns we find in the data with respect to realized gains, losses, and the propensity to participate in the penny auction again. Although the literature has hinted that a Prospect Theory value function may drive consumers' return decisions and such hypothesis does hold attractive intuition, to our knowledge, prior research has not derived the conditions of the loss averse framework needed to account for consumers' return

[14] See Cox (1972) for details.

behavior.

Thus, the objective of this section is to pin down what types of additional assumptions are needed in order to generate the return patterns found in our data. In fact, we show that a loss aversion framework with a fixed reference point alone, is insufficient to generate the return patterns we find in the data. In particular, additional assumptions on the formation process of the reference point, as well as a discontinuity at the reference point such as in Aspiration Level Theory (Diecidue and Van de Ven, 2008) are needed to justify the empirical findings.

We first summarize our results from the previous section. The propensity to return for another penny auction experience depends distinctly on the monetary gains and losses experienced in the following ways: First, the propensity to return is at least weakly increasing in the value of the surplus obtained; Furthermore, the marginal propensity to return is significantly steeper as a function of losses as compared to gains; Finally, the propensity to return is discontinuous at the break-even point where gains and losses intersect, with a marginal loss significantly reducing the likelihood of return compared to the effect of a marginal gain.

Now we turn to our theoretical model of how a consumer updates his expected reference-dependent utility after having a new experience. To summarize our theoretical results, we show that although the return function itself bears resemblance to a loss averse utility function, in a model in which consumers form expectations on their future utility being a repeat customer, further behavioral assumptions are needed in order to generate the kink and discontinuity in the empirical result. The shape of the return function that we find in our data can be accounted for using a standard reference-dependent loss averse framework with two additional key intuitive features incorporated into the model: 1. discontinuity in consumer utility at the reference point as proposed in Diecidue and Van de Ven (2008), if the reference point is negative and 2. a continuous reference point formation function, for which the reference point itself shares the same sign as the previous outcome, and the reference point adjustment after a negative outcome is less severe than that for a corresponding positive outcome. We explain the intuitive psychological justifications of such a reference point adjustment in the consumer experience setting.

We now describe the model, which is derived independently of the penny auction setting and can be applied to other consumer experience settings in general.

A consumer decides whether or not to participate in a relatively new activity again, which involves uncertainty over outcomes and will generate a reference-dependent expected utility to the consumer. Following the gain-loss utility in Koszegi and Rabin (2006), the consumer has preferences represented by $u(c \mid r) = m(c) + \mu(m(c) - m(r), r)$, where $m(\cdot)$ is the classical utility of consumption which is differentiable, monotone and normalized such that $m(0) = 0$, $\mu(\cdot)$ is the generalized version of the reference-dependent utility function as proposed in Koszegi and Rabin (2006), $c \in R$ is the realized outcome in net wealth levels after the consumer's experience, and $r \in R$ is the reference point. For simplicity, we assume piecewise linearity of the gains and loss segments, which is typically sufficient to generate the characteristic predictions of loss aversion apart from the diminishing sensitivity additionally proposed in the original Prospect Theory (Kahneman and Tversky, 1979).

More specifically, we assume $\mu(x,r) = \begin{cases} \eta x & if \quad x \geq 0 \\ \lambda \eta x & if \quad x < 0, r \geq 0 \\ \lambda \eta x - \delta & if \quad x < 0, r < 0 \end{cases}$, where $\eta > 0$, $\lambda > 1$ is the loss-aversion parameter measuring the disutility of marginal losses compared to the utility of marginal gains, and $-\delta$ where $\delta \geq 0$ denotes the extra disutility from experiencing any loss magnitude with a negative reference point. Our specification of the utility function reduces to the traditional gain-loss utility framework of Koszegi and Rabin (2006) if $\delta = 0$. For positive δ, conditional on the reference point being negative, utility additionally incorporates the discontinuous utility feature of aspiration level theory as proposed in Diecidue and Van de Ven (2008).[15]

The psychology of this utility function can be understood as follows: If the consumer has a reference point in a nominally positive range such as the case of having a positive expectation or target, he or she has standard loss averse utility. However, if the consumer holds a nominally negative reference point such as having a negative expectation or target, he or she feels even worse from an outcome marginally below that reference point, represented by the discontinuity parameter δ.

The consumer's perception or belief of the underlying quality of the website, purchase mechanism, and overall consumer experience of this relatively new activity is represented by the variable $q \in (0, +\infty)$. $F(c \mid q)$ is then the cumulative distribution function of a consumer's (subjective) belief on his or her net wealth outcomes conditional on the previous experience obtained by the consumer utilizing this purchase mechanism or service. It is easy to see that the expected utility of participating in the activity for an individual who has a quality perception of q and a reference point of r is given by $U(q,r) = \int_c u(c \mid r) dF(c \mid q)$.

It is intuitive that a consumer with a higher quality perception of a particular website, service or purchase platform expects a higher likelihood for better wealth outcomes. We describe this insight by the following assumption, which is commonly assumed in the literature.

Assumption 1: $\forall q, q' \in R_+, q > q' \Rightarrow F(c \mid q) \quad FOSD \quad F(c \mid q')$.

Since the experience in question is new to the individual, there is an uncertainty regarding the underlying quality q. Suppose that consumers are ex-ante identical in terms of their prior about the quality, represented by probability distribution function $G(q)$ over the range $(0, +\infty)$. For an individual who has participated in the experience t times in the past and has a history of previous outcomes given by $y \in R^t$, his updated belief about his future experience is represented by $G(q \mid y)$. Therefore, the expected utility of participating in the activity for an individual who has a reference point of r and a history of y can be written as

$$U(r, y) = \int_q \int_c u(c \mid r) dF(c \mid q) dG(q \mid y).$$

For simplicity, in the remainder of the analysis we focus on the case where $t = 1$. That is, the individual only has participated in this particular consumer activity once before. The same reasoning we describe below applies to multiple rounds of the activity.

It is intuitive to assume that a consumer with a better outcome in their previous experience would perceive a more positive distribution on his future experience, described by the following straightforward assumption:

Assumption 2: $\forall y, y' \in R, y > y' \Rightarrow G(q \mid y) \quad FOSD \quad G(q \mid y')$.

Understanding how an individual forms their reference point is key to studying the reference-dependent behavior. It is often thought that the way in which the reference point is formed

[15] Diecidue and Van de Ven (2008) consider their aspiration level theory model an "extreme" type of loss aversion. Thus, it is natural for us to find that the discontinuity in utility they propose is needed to explain our empirical patterns.

should be continuous, and in most cases smooth. It is also reasonable to assume that the reference point is zero (ie. no change from the original normalized reference point) when the previous outcome is zero, and naturally that the reference point should share the same sign as any previous nonzero outcome. The final feature we would like to capture regarding reference point formation is that reference point is "confidently" positive (ie. larger adjustment) in response to prior positive outcomes and "resiliently" negative (ie. smaller adjustment) in response to prior negative outcomes. We characterize these insights in the following assumption:

Assumption 3: The reference point formation function $r(y): R \to R$ satisfies the following conditions:

(3.1) $r(y)$ is continuous;

(3.2) $r(y)$ is differentiable almost everywhere with the potential exception at $y = 0$;

(3.3) $r(0) = 0$;

(3.4) $\forall y \neq 0$, $r(y) \cdot y > 0$;

(3.5) $\lim_{y \to 0^+} r'(y) > \lim_{y \to 0^-} r'(y)$.

An example of a reference point formation function satisfying all the conditions 3.1-3.5 in Assumption 3, is provided in Diagram 1 below. Note that the monotonicity of $r(y)$ as shown in the figure is a sufficient condition, not necessarily needed for any conditions in Assumption 3 to hold.

Diagram 1: Example reference point formation function

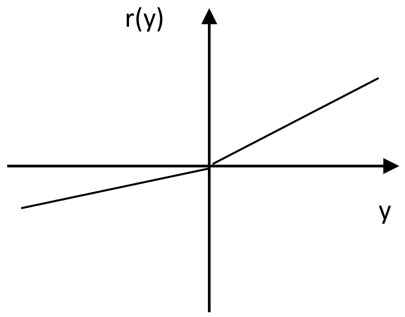

Under the above model set up and assumptions, we derive three propositions regarding the properties of the return likelihood across the possible realized experience outcomes. The proofs are provided in the Appendix. First, we have the following result which is consistent with the first feature of the empirical pattern in the penny auction experience data: A better previous experience outcome leads to a higher expected utility, and hence a higher likelihood to return.

Proposition 1 (Monotonicity): Under Assumptions 1 and 2, $\forall r \in R, \forall y, y' \in R$,

$$y > y' \Rightarrow U(r, y) > U(r, y').$$ [16]

The proof is provided in Appendix D.

So far, our model provides us with a consumer return function which is monotonic in prior experience, as we have found in the data. We have two remaining empirical features to account for: the asymmetric slopes of the return function for the nominal gains and loss segments, and the significant discontinuity in the return function at the nominal break-even level.

Due to the ease of explaining the assumptions needed to generate these features, we first introduce the proposition which generates the discontinuity.

[16] Here for simplicity we assume a fixed reference point in Proposition 1 without imposing further conditions on reference point formation. We also show in Appendix D that Monotonicity is guaranteed under some additional technical conditions when the reference point is not fixed.

Proposition 2 (Discontinuity): Under Assumptions 3.1, 3.3 and 3.4,

$$\delta > 0 \Rightarrow U(r(y), y = 0) > \lim_{y \to 0^-} U(r(y), y).$$

The proof is provided in Appendix D.

We denote the mapping between the previous outcome y and the new reference point r as the reference point formation function $r(y)$. Proposition 2 states that under some conditions on $r(y)$, there will be a gap between the consumer's break-even expected utility and the limit of his expected utility from direction of negative outcomes approaching the zero mark. The conditions are that for positive previous outcomes, the reference point is positive, and for negative previous outcomes, the reference point is negative.

The psychological justification of the condition on the reference point formation function arises from the idea that reference points are adjusted downward after prior negative experiences, but upward after prior positive experiences. Such an updating process is intuitive in many situations and includes case of adaptively adjusting expectations-based reference points.

The intuition of the proposition is that after receiving any negative experience consumers not only have their gain-loss utility component reset around a lower reference point, but they adopt a discontinuity δ incurred by the negative reference point, as specified in our gain-loss utility function. Combined with the low consumption utility which is tied to the overall assessment of the activity's attractiveness, the prospect of participating in the activity again is decidedly less attractive.

The final return pattern we are interested in modeling is the asymmetric return behavior in response to marginal gains and losses, as seen in the data. We note that similarly to the finding that a discontinuity in the utility function alone is insufficient to generate a discontinuity in the return function, the kink in the reference-dependent loss averse utility function alone is insufficient to generate an asymmetric reaction to gains and losses, due to the 'smoothing' effect of the uncertainty over future experiences.

However, if we incorporate an additional assumption regarding how individuals form their reference point given their previously experienced outcome (Assumption 3.5), we can obtain a prediction which shares the asymmetric effects of marginal gains and losses in the empirical return pattern:

Proposition 3 (Kink): Under Assumptions 3.1-3.3 and 3.5,

$$\lim_{y \to 0^-} \frac{U(r(y), y < 0) - \lim_{y \to 0^-} U(r(y), y)}{y} > \lim_{y \to 0^+} \frac{U(r(y), y > 0) - U(r(0), y = 0)}{y}.$$

The proof is provided in Appendix D.

Proposition 3 states that in order to generate a kink in the consumer's return function at the nominal zero level, a reference point formation function is needed which has higher slope on the gains side than on the loss side. In other words, the consumer forms his reference point more 'enthusiastically' after a nominally good experience compared to the downward formation after a corresponding bad experience. This assumption can be understood psychologically as having a reference point which is confidently positive in response to prior positive outcomes, and resiliently negative in response to prior negative outcomes, recalling again that a negative reference point induces the discontinuity δ in the gain-loss utility function.

Our propositions on expected utility can be straightforwardly translated into the likelihood of returning, under the additional standard assumption that higher expected utility levels correspond directly to ordinally higher likelihoods of returning. One method of transitioning our model from the propositions about expected utility to predictions about return likelihood is through imposing a unit mass of consumers with loss aversion parameters distributed uniformly over a range of values.

A key insight on reference dependent behavior in this setting is that when consumers are deciding whether to be a repeat customer, the kink and discontinuity at what "appears" to be a reference point are actually driven by reference point formation behavior which is highly contingent on the previous experience. Notably, the formation of the reference point rarely leaves the reference point remaining at the nominal break-even level, even though the kink in the consumer return function remains at that level.

To summarize, although the empirical pattern we find in consumers' likelihoods of returning to penny auction experiences is highly suggestive of reference-dependent loss averse utility, two key further assumptions are needed in the standard reference-dependent loss aversion model in order to account for this pattern. First, a gain-loss utility function which has a discontinuity as in aspiration level theory if the reference point is negative. Second, a smooth reference point formation function which updates downward in response to negative outcomes, and updates upward in a more enthusiastic manner in response to corresponding positive outcomes, with a reference point of the same (positive/negative) sign as the previous outcome. We interpret these additional assumptions as being realistic with regard to consumers' psychology after initial experiences with new purchase methods and products, and they can generate our intuitive empirical findings which we describe in detail in subsequent sections.

5. Additional Evidence on Bracketing

Another important question in reference-dependent models is how decision-makers calculate their gains and losses. Our data allow us to examine this question along the domain of the number of auctions played. Given a sequence of penny auctions played, how does the sequence of gains and losses experienced, factor into the consumer's feeling of gain and loss? Do early experiences, recent experiences, or the accumulation of all prior experiences matter the most?

In an experimental asset market setting, Baucells, Weber and Welfens (2011) find a U-shaped function of previous experiences in terms of the weights applied to the current period's reference point. Such a pattern also seems intuitive in our setting, since one can imagine that both first experiences and recent experiences might matter most to the consumer.

In order to address this question, we run our basic logit specification (for simplicity, without the control variables), tracing the marginal impact of gains and losses on the return decision, for *each* previous round of penny auction played. Thus, for the decision after the 5^{th} auction, on whether to return to play a 6^{th} time, we estimate coefficients on the loss indicator, gain or loss magnitude, and interaction term, for the 1^{st} through 5^{th} experience outcomes. This controls for consumers' prior positive and negative experiences and their magnitudes, during their entire history of play. We are then interested in comparing the magnitude of coefficients on the gains and loss variables across time periods or "experiences". Due to the large number of coefficients in each regression, we arbitrarily choose the decision after the 5^{th} auction and the decision after the 10^{th} auction for analysis. We also run the analogous specifications for all other rounds, 2 through 9, and find very similar patterns in each case.

In Table 6, the column on the left, labeled "After 5^{th} experience" shows the coefficients for each previous round of play, through the 5^{th} round. As the coefficients show, the effect of loss aversion is stronger for recent experiences. In contrast to a U-shape, the coefficient on the loss indicator is growing more negative with each more recent round, while the coefficient on the interaction term also tends to get more negative.

A similar story holds for the decision taking place for consumers after the 10th experience, split into two columns due to the large number of coefficients. While not quite as monotonic as the 5th experience regression, the coefficients for the regression after the 10th experiences do follow a similar trend. That is, the coefficient on the loss indicator tends to become more negative as the experience was more recent. The marginal impact of larger losses is not as precisely and uniformly estimated across rounds, and tends to remain constant. This suggests that when considering a longer history of experiences, it may be the loss itself rather than the magnitude of it, which is increasingly relevant in consumers' minds.

By plotting the return ratio as a function of calculated losses or gains, as in Figure 1, we can clearly observe that the most recent experience is also more influential compared to the accumulation of prior experiences. Figure 3 shows the plot of the return ratio over gains and losses calculated over all previous played (left panel), compared to the same information for gains and losses of the most recent auction played only (right panel). While both figures show a trend in larger losses leading to lower return rates, the structural break at the break-even level is more clearly apparent from the graph which only uses the most recent experience. Very similar patterns hold for other numbers of auctions played, and some of these are provided in Appendix C or are otherwise available upon request.

Table 6: Logit Regressions: Bracketing

	Dependent variable: Return = 1 if the bidder returned			
	After 5th experience	After 10th experience (1st through 5th)	**After 10th experience (6th through 10th)**	
Loss (1)	-0.029***	0.007	**-0.044*** **	**Loss (6)**
	(0.002)	(0.008)	**(0.012)**	
Loss(1)*magnitude of gain or loss(1)/100	-0.031**	-0.093***	**0.019**	**Loss(6)*magnitude of gain or loss(6)/100**
	(0.013)	(0.017)	**(0.013)**	
Magnitude of gain or loss(1)/100	0.045***	0.078***	**0.020**	**Magnitude of gain or loss(6)/100**
	(0.013)	(0.019)	**(0.016)**	
Loss (2)	-0.052***	-0.017**	**-0.021*** **	**Loss (7)**
	(0.009)	(0.008)	**(0.006)**	
Loss(2)*magnitude of gain or loss(2)/100	0.008	-0.028	**-0.021**	**Loss(7)*magnitude of gain or loss(7)/100**
	(0.015)	(0.030)	**(0.026)**	
Magnitude of gain or loss(2)/100	0.003	0.022	**0.030**	**Magnitude of gain or loss(7)/100**
	(0.009)	(0.019)	**(0.022)**	
Loss (3)	-0.055***	-0.026***	**-0.033*** **	**Loss (8)**
	(0.011)	(0.006)	**(0.010)**	
Loss(3)*magnitude of gain or loss(3)/100	-0.045	0.029	**-0.082*** **	**Loss(8)*magnitude of gain or loss(8)/100**
	(0.032)	(0.021)	**(0.016)**	
Magnitude of gain or loss(3)/100	0.034*	0.012**	**0.054*** **	**Magnitude of gain or loss(8)/100**
	(0.018)	(0.006)	**(0.018)**	
Loss (4)	-0.071***	-0.019***	**-0.049*** **	**Loss (9)**
	(0.014)	(0.006)	**(0.015)**	
Loss(4)*magnitude of gain or loss(4)/100	-0.048**	-0.080*	**-0.078*** **	**Loss(9)*magnitude of gain or loss(9)/100**
	(0.021)	(0.046)	**(0.023)**	
Magnitude of gain or loss(4)/100	0.011	0.113**	**0.024***	**Magnitude of gain or loss(9)/100**
	(0.019)	(0.050)	**(0.013)**	
Loss (5)	-0.084***	-0.020***	**-0.062***	**Loss (10)**
	(0.029)	(0.003)	**(0.025)**	
Loss(5)*magnitude of gain or loss(5)/100	-0.134***	-0.021	**-0.080*** **	**Loss(10)*magnitude of gain or loss(10)/100**
	(0.024)	(0.021)	**(0.022)**	
Magnitude of gain or loss(5)/100	0.009	0.036***	**0.003**	**Magnitude of gain or loss(10)/100**
	(0.019)	(0.013)	**(0.012)**	
Pseudo-R^2	0.0215		0.0245	
Observations	86,028		44,848	

Notes: Marginal effects reported, standard errors clustered at the product category level; *significant at 10%; **significant at 5%; ***significant at 1%. Loss and gain refers to nominal losses and gains around zero consumer surplus.

Figure 3: Return Ratios after the 5th and 10th Experience

(Gains and losses calculated over 5 rounds (left) and from 5th auction only (right))

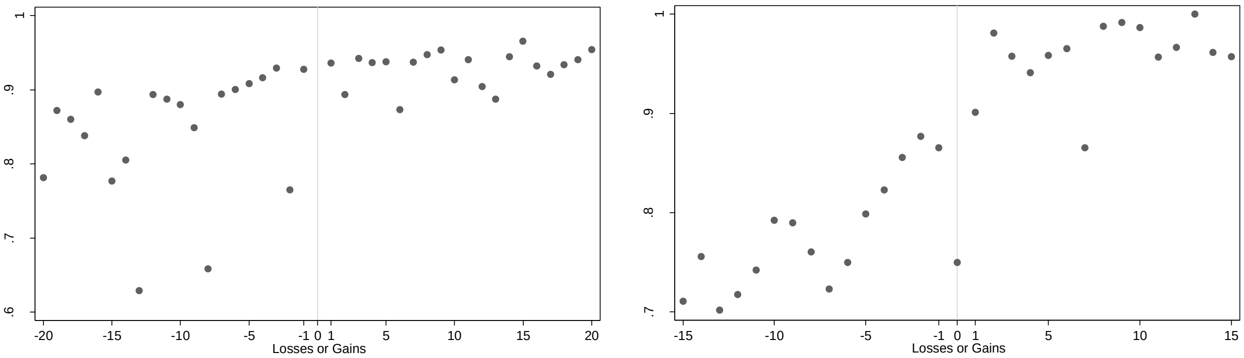

(Gains and losses calculated over 10 rounds (left) and from 10th auction only (right))

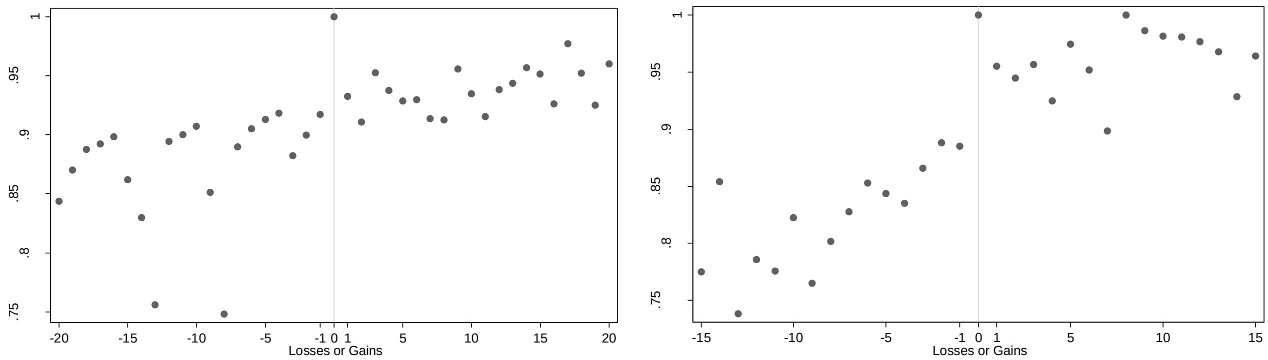

The implication of our findings for customer retention is that recent experiences seem to matter the most in the calculation of gain and loss experiences. Since our regressions control for all other previous experiences, it suggests that even an early very positive experience can be 'erased' by a recent negative one, which also accords with anecdotal evidence of long-time customers ending their relationship with a service provider after a single dissatisfactory incident. The results suggest the importance of avoiding *any* feeling of loss by the consumer, if the goal is to increase the likelihood of the consumer returning.

6. Further Remarks and Robustness Checks

6.1. Effectiveness of learning from previous experiences

Given that we focus on the effect of gain and loss experiences on consumer's returning, it is natural to ask whether playing more auctions might increase consumer's winning chances. In other words, is there any significant room for a consumer to learn to play the penny auctions more effectively? We use our panel data to analyze the relationship between individual bidders' performance and experiences using the following simple model:

$$\pi_{in} = c + \theta_1 \cdot n_i + \theta_2 \cdot n_i^2 + \varphi_i + \mu_i$$

where π_{in} is bidder *i*'s net gain in his nth auction, n_i denotes the number of auctions participated by bidder *i* in total, and φ_i is the individual fixed effect.

Wang and Xu (2016) find that experience has small positive effect for highly experienced, sophisticated bidders who played more than 200 auctions. Augenblick (2014) also suggests that only those bidders with very high experience levels can enjoy positive benefits from experience, and that experience has an insignificant or even a negative effect for relatively inexperienced bidders. In this paper, we are interested in the effect of early experience, and thus we examine the relationship between performance and experience during the first 10 auctions for bidders who played at least 10 auctions.

Table 7 reports the results from the above specification. In column (1), the coefficient on experience is negative and significant, which indicates that bidders' performance is on average, getting worse over time within the range of their first 10 auctions. As mentioned in previous section, BigDeal offered special auctions to new bidders, in which bidders are more likely to perform well. With this concern in mind, column (2) reports results which exclude bidders' first experiences. The regressions show that performance tends to have a downward or neutral relationship with experience, at least for the low number of auction experiences we are interested in. In other words, it is not easy to learn and improve individual's bidding skill in penny auctions within a small number of rounds. To examine the robustness of these findings, we also report the results in their first 5 and first 15 auctions, for bidders who played at least 5 and 15 auctions respectively, and the results are similar.

Table 7: Relationship between Experience and Performance

	Dependent variable: Bidder's net gain					
	First 10 auctions		First 5 auctions		First 15 auctions	
	(1)	(2)	(3)	(4)	(5)	(6)
Experience	-0.419***	-0.027	-0.884***	0.080	-0.326***	-0.067
	(0.071)	(0.087)	(0.144)	(0.287)	(0.056)	(0.064)
Experience^2	0.010	-0.020***	0.045**	-0.063	0.006*	-0.008**
	(0.006)	(0.007)	(0.023)	(0.041)	(0.003)	(0.004)
Constant	-0.153	-1.308***	-0.370*	-1.734***	0.344*	-0.740***
	(0.167)	(0.227)	(0.191)	(0.463)	(0.192)	(0.237)
R^2	0.001	0.001	0.001	0.001	0.001	0.001
Observations	452,028	406,631	431,890	345,252	375,682	350,484

Notes: Column (2) (4) and (6) exclude the first experience. Individual fixed effects included. Robust standard errors reported in parentheses. *significant at 10%; **significant at 5%; ***significant at 1%.

6.2. Effect on the intensive margin

Conditional on winning bidders being more likely to return, we may be interested in the effect of prior experiences on bidding expenditures in the next auction played. In other words, do previously winning bidders subsequently place more bids in future auctions, similar to a house money effect? This also serves as a robustness check on the reference-dependent behavior, since one can straightforwardly extend the implications of the model in Section 3 for the average bidding intensity, all else equal. Furthermore, the analogous results on the intensive margin, provide further confirmation that the effects found in the main results are driven by individual preferences rather than heterogeneity or selection.

To test the relationship between prior experiences on bidding intensity, we examine the relationship between bidders' total number of bids in their second auction (conditional on having returned) and their outcomes in their first auction. Table 8 reports these results.

In the simple regression of column (1), we see that indeed, a negative outcome in the first auction tends to significantly discourage bidders' intensity of bids in the next auction. As we control for the magnitude of loss in column (2) and the (log) retail price of the second auction, the negative effect of loss in the first auction remains. Meanwhile, as we can see, the magnitude of the prior outcome plays a strong role on the intensive margin in the second auction, no matter whether the first outcome was a gain or loss. In addition, if a bidder participates in an auction for an expensive product, he tends to place more tokens. We also implement the same specification for the relationship between outcomes in the second auction and bidding in the third auction, and find that this pattern is robust (omitted for space considerations, but available upon request).

As may be expected, the impact of gain and loss differs in shape compared to that shown in the domain of returning behavior (extensive margin). Here, although greater magnitudes in either direction increase bid intensity, and for losses more so than for gains, a discontinuity at the break-even point remains. Compared to the domain of returning behavior, consumers' bidding patterns more closely reflect the break-even and house money effects found in previous literature in the pure gambling domain, although the coefficient on *loss* shows that consumers in this setting still react in a discontinuous manner to any type of marginal loss.

Table 8: Effect of Prior Gains and Losses on Bid Intensity

	Dependent variable: Total number of bids placed in the 2nd auction		
	(1)	(2)	(3)
Loss	-1.802***	-1.095***	-2.007***
	(0.303)	(0.321)	(0.318)
Loss*magnitude of gain or loss		0.117***	0.104***
		(0.022)	(0.022)
Magnitude of gain or loss		0.065***	0.061***
		(0.013)	(0.013)
Log(price of the 2nd auction)			2.293***
			(0.047)
Constant	9.637***	8.047***	-2.566***
	(0.295)	(0.307)	(0.338)
R-squared	0.0003	0.0095	0.0254
Observations	155,188	155,188	155,188

Notes: Robust standard errors reported in parentheses; *significant at 10%; **significant at 5%; ***significant at 1%. Loss and gain refers to nominal losses and gains around zero consumer surplus.

6.3. Persistence of Product Demand

While penny auctions provide a platform and purchase mechanism for consumers, one question is not just the return behavior to the website itself, but the persistence of demand for a particular product on the website. If a bidder comes to the website specifically to buy a particular item, and he lost in the first auction, then we might expect that the bidder will return and target purchase of same product. We focus our analysis on this question to non-token auctions, since participation in token auctions can be interpreted as an intention to bid for another non-token product in the future.

We present two pieces of evidence on the persistence of bidders' product demand. First, we examine the number of distinct products among bidders' first 10 auctions by those bidders who played at least 10 non-token auctions. There are 24,936 unique bidders in our data who played at least 10 non-token auctions. Table 9 displays the distribution of the number of distinct products played by them during those 10 auctions. More than 70% of them played for a single product at least twice or more. In particular, 38 of them stick to the identical product at least 10 times.

Table 9: Distribution of the Number of Distinct Products

Bidders playing 10 or more auctions, first 10 auctions

	Number of distinct products									
	1	2	3	4	5	6	7	8	9	10
Number of bidders	38	161	296	585	1,086	1,759	3,093	4,825	6,575	6,518
Percentage	0.15	0.65	1.19	2.35	4.36	7.05	12.40	19.35	26.37	26.14

Secondly, we check the probability that consumers return to the penny auction website for the same product directly just after the experience in the first auction. Our intuition is that a losing bidder in the first auction is more likely to play the same product in the second auction, and our evidence data confirms this. There are 8.27% of bidders who lost in their first auction that came back for the same product, while the analogous probability for winners is just 6.56%.

7. Conclusions

In this study we have explored an issue which is important to any firm providing a new experience to consumers: How do positive and negative experiences influence the consumer's likelihood of coming back for another experience with the same new retailer or service provider?

To address this question, we utilize data from online penny auctions during their introduction to the online marketplace as a means for testing new consumer experiences, where a key advantage of this setting compared to other new consumer experiences is that consumer surplus experiences can be objectively quantified in terms of monetary gains and losses. Consumers react more strongly to marginal losses as compared to marginal gains by reducing the likelihood of returning, and there is a discontinuity in the likelihood of returning associated with having a losing outcome in the penny auction experience.

We construct a theoretical model of consumer experiences under reference-dependence, and derive the conditions under which the customer return pattern follows the empirical findings with respect to nominal gains and losses. Since the decision problem in the consumer experience setting requires that consumers form expectations over the true quality or value of the product or service, conditions on reference point formation conditional on gains and losses experienced are required in order to generate the kink (asymmetry in slopes) and discontinuity in the return function. In particular, a previous negative outcome induces a negative reference point, and a previous positive outcome induces a positive reference point, which accords with an adaptively formed expectation-based reference point. The reference point formation function is steeper for positive outcomes than negative ones, meaning that the consumer 'updates' more enthusiastically in response to prior gains. An additional condition on the utility function is that a negative reference point induces a discontinuity in the utility function at the reference point.

In terms of how consumers bracket their gains and losses over multiple experiences, we find that the most recent gain or loss experience weighs most heavily in the decision regarding whether to return or not. This finding sheds light on how consumers psychologically aggregate their histories of positive and negative experiences over multiple interactions, and is consistent with conventional wisdom of customer service oriented firms. This is consistent with the intuition that a single bad experience which has occurred most recently, can permanently drive away a previously loyal customer.

We implement a number of robustness checks to support our main empirical finding and gain further insight into consumers' behavior in response to surplus gains and losses. Given the highly random nature of penny auction wins, the retention behavior we detect here can be attributed to the gains and losses acquired by consumers in their previous auction played. In particular, experience over a low number of auctions has a neutral or even negative impact on future successes. In other words, it is quite difficult in practice for consumers to increase their "skill" in penny auctions, especially over the small number of auctions we are interested in here. We also find evidence for the same reference-dependent patterns over bidding intensity that we find for the return rate. Previously losing bidders bid substantially less than their winning counterparts, in subsequent auctions, which is indicative of their lower willingness to spend. This provides a robustness check against purely selection-based explanations of the empirical return function, since customer's future intensity of bidding and not merely their future participation is shown be highly responsive to marginal losses. Finally, we find that there is substantial persistence in product demand among new consumers, with about 70% of moderate customers (at least 10 auctions participated) returning for the same product at least once in a short frame, and that this persistence is slightly higher for auction losers than winners.

Our findings also reveal insights regarding the study of penny auctions as a selling mechanism. While the rules of the auction are somewhat complicated, involve strategic factors, and may require learning processes by consumers, the consumers' return behavior is to large degree, predictable by a loss aversion framework. Thus, even when the problem faced by consumers is relatively complicated and potentially strategic, a decision model of preferences can capture consumer decisions well. Another noteworthy insight pertains to the risk involved in penny auctions. This auction format carries risk often thought to be similar to many gambling mechanisms, in the sense that the 'winning' return a consumer obtains may occur with fairly low probability, which is largely out of his or her own control. Even consumers who selected into the penny auction mechanism, and are apparently more risk-seeking than average, tend to be loss averse in their decision whether to continue or not. Furthermore, they do not treat the experience as purely a gambling one, since we do not observe the signature break-even effect in the return function typically associated with purely risky settings.

While our findings here resonate with conventional wisdom about customer service and consumer retention practices, and thus should be generalizable across consumer settings other than penny auctions, there are limitations to our study which we should mention. In particular, the *magnitudes* of the effects measured here may not be specific to other settings. One potential reason is that penny auctions are a relatively rare method of purchasing items, and the types of consumers drawn to penny auctions may have different utility function parameters compared to consumers of other settings. As is the case for empirical research in general, further studies in a variety of settings are needed to test the robustness of the quantitative aspects of our findings. Further, customers' experiences on the penny auction website were often relatively short lived, and it cannot be discerned from our data alone whether we would still find such exact patterns, especially with regard to the bracketing results, for other consumer experiences which stretch over a longer horizon.

We also see several additional directions for future research. First, as mentioned earlier, a variety of evidence is needed from other new consumer experiences in other domains to check that the general magnitudes of effects we find here are robust to different settings. Second, it will be useful to pursue survey-based work in pinpointing consumers' qualitative and emotional reactions after gain and loss experiences, while simultaneously checking using real data regarding whether they follow up on their future purchase intentions. Our study contributes to the literature on consumer retention through the ability to measure monetary surplus, but cannot provide any emotion or psychology-based validation of these monetary amounts. Additionally, survey work can help determine the extent to which consumers comprehend the auction mechanism. Last but not least, future research may consider more precisely how firms can best respond to consumers with the behavioral patterns observed in our study.

References:

Abeler, J., A. Falk, L. Goette, and D. Huffman, "Reference Points and Effort Provision." *The American Economic Review*, Vol. 101 (2011), Issue 2, p. 470-492.

Augenblick, N., “The Sunk-Cost Fallacy in Penny Auctions”, *Review of Economic Studies,* Vol. 83 (2016), No. 1, p. 58 – 86.

Backus, M., Blake, T., Masterov, D.V., and Tadelis, S., “Expectation, Disappointment, and Exit: Reference Point Formation in a Marketplace”, *NBER Working Paper 23022* (January 2017).

Baucells, M., Weber, M., and Welfens, F., “Reference-Point Formation and Updating”, *Management Science,* Vol. 57 (2011), Issue 3, p. 506 – 519.

Berger, J. and Pope, D., “Can Losing Lead to Winning?” *Management Science,* Vol. 57 (2011), No. 5, p. 817 – 827.

Camerer, C., L. Babcock, G. Loewenstein, and R. Thaler, “Labor supply of New York City cabdrivers: One day at a time,” *The Quarterly Journal of Economics,* 1997, 112 (2), 407-441.

Camerer, C., Shum, M., and Ray, D., “Loss-Aversion in Consumer Reactions to Sales-Price Changes”, *American Economic Review Papers and Proceedings,* Vol. 105 (2015), p.376 - 380.

Chang, K., Chen, M., and Hsu, C., “Applying loss aversion to assess the effect of customers’ asymmetric responses to service quality on post-dining behavioral intentions: An empirical survey in the restaurant sector”, *International Journal of Hospitality Management,* Vol. 29 (2010), p. 620 – 631.

Cox, D.R., “Regression models and life-tables”, *Journal of the Royal Statistical Society, Series B (Methodological),* Vol. 34 (1972), No. 2, p. 187 – 220.

Crawford, V.P. and J. Meng, “New York City Cabdrivers' Labor Supply Revisited: Reference-Dependence Preferences with Rational-Expectations Targets for Hours and Income,” *American Economic Review,* 2011, 101(5): 1912 – 32.

Diecidue, E., and Van de Ven, J., “Aspiration Level, Probability of Success and Failure, and Expected Utility”, *International Economic Review,* Vol. 49, No. 2 (2008), p. 683 – 700.

Eil, D., and Lien, J.W., “Staying Ahead and Getting Even: Risk Attitudes of Experienced Poker Players”, *Games and Economic Behavior,* Vol. 87 (2014), p. 50 – 69.

Fudenberg, D., and D.K. Levine, “Learning with Recency Bias”, *Proceedings of the National Academy of Sciences,* Vol. 111 (2014), p. 10826 – 10829.

Gill, D. and Prowse, V., “A Structural Analysis of Disappointment Aversion in a Real Effort Competition”, *American Economic Review,* Vol. 102, No. 1, p. 469 – 503.

Heidhues, P. and Koszegi, B., “Regular prices and sales”, *Theoretical Economics,* Vol. 9 (2014), p. 217 – 251.

Huang, X. and Zhang, D., “Service Product Design and Consumer Refund Policies”, *Marketing Science,* Vol. 39 (2020), No. 2, p. 366 – 381.

Kahneman, D., and Tversky, A., “Prospect Theory: An Analysis of Decision under Risk”, *Econometrica,* Vol. 47 (1979), No. 2, p. 263 – 292.

Koszegi, B., and Rabin, M., “A Model of Reference-Dependent Preferences”, *Quarterly Journal of Economics,* Vol. 121 (2006), Issue 4, p. 1133 – 1165.

Koszegi, B., and Rabin, M., “Reference-Dependent Risk Attitudes”, *American Economic Review,* Vol. 97 (2007), No. 4, p. 1047 – 1073.

Lien, J.W., and Zheng, J., “Deciding When to Quit: Reference-Dependence over Slot Machine Outcomes”, *American Economic Review Papers and Proceedings*, Vol. 105 (2015), p.366 - 370.

List, J.A., “Does Market Experience Eliminate Market Anomalies?” *Quarterly Journal of Economics*, Vol. 118 (2003), No. 1, p. 41 – 71.

List, J.A., “Neoclassical Theory Versus Prospect Theory: Evidence from the Marketplace”, *Econometrica*, Vol. 72 (2004), No. 2, p. 615 – 625.

Meng, J. and Weng, X., “Can Prospect Theory Explain the Disposition Effect? A New Perspective on Reference Points”, *Management Science,* (2017), forthcoming.

Platt, B.C., Price, J., and Tappen, H., “The Role of Risk Preferences in Pay-to-Bid Auctions”, *Management Science*, Vol. 59 (2013), No. 9, p. 2117 – 2134.

Qiu, Y., and Rao, R.C., “Increasing Retailer Loyalty Through the Use of Cash Back Rebate Sites”, *Marketing Science*, Vol. 39 (2020), No. 4, p. 743 – 762.

Rabin, M., and Schrag, J.L, “First Impressions Matter: A Model of Confirmatory Bias”, *Quarterly Journal of Economics*, Vol. 114 (1999), No. 1, p. 37 – 82.

Thaler, R.H. and Johnson, E.J. “Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice,” *Management Science*, 1990, 36 (6), 643-660.

Wang, Z., and Xu, M., “Selling a Dollar for More Than a Dollar? Evidence from Online Penny Auctions”, *Information Economics and Policy,* Vol. 36 (2016), p. 53 – 68.

Zhang, M., and Zheng, J. “A Robust Reference-Dependent Model for Speculative Bubbles,” *Journal of Economic Behavior and Organization*, Vol. 137 (2017), p. 232-258

Appendix A: Return Ratio after the First Try Plots by Product Category

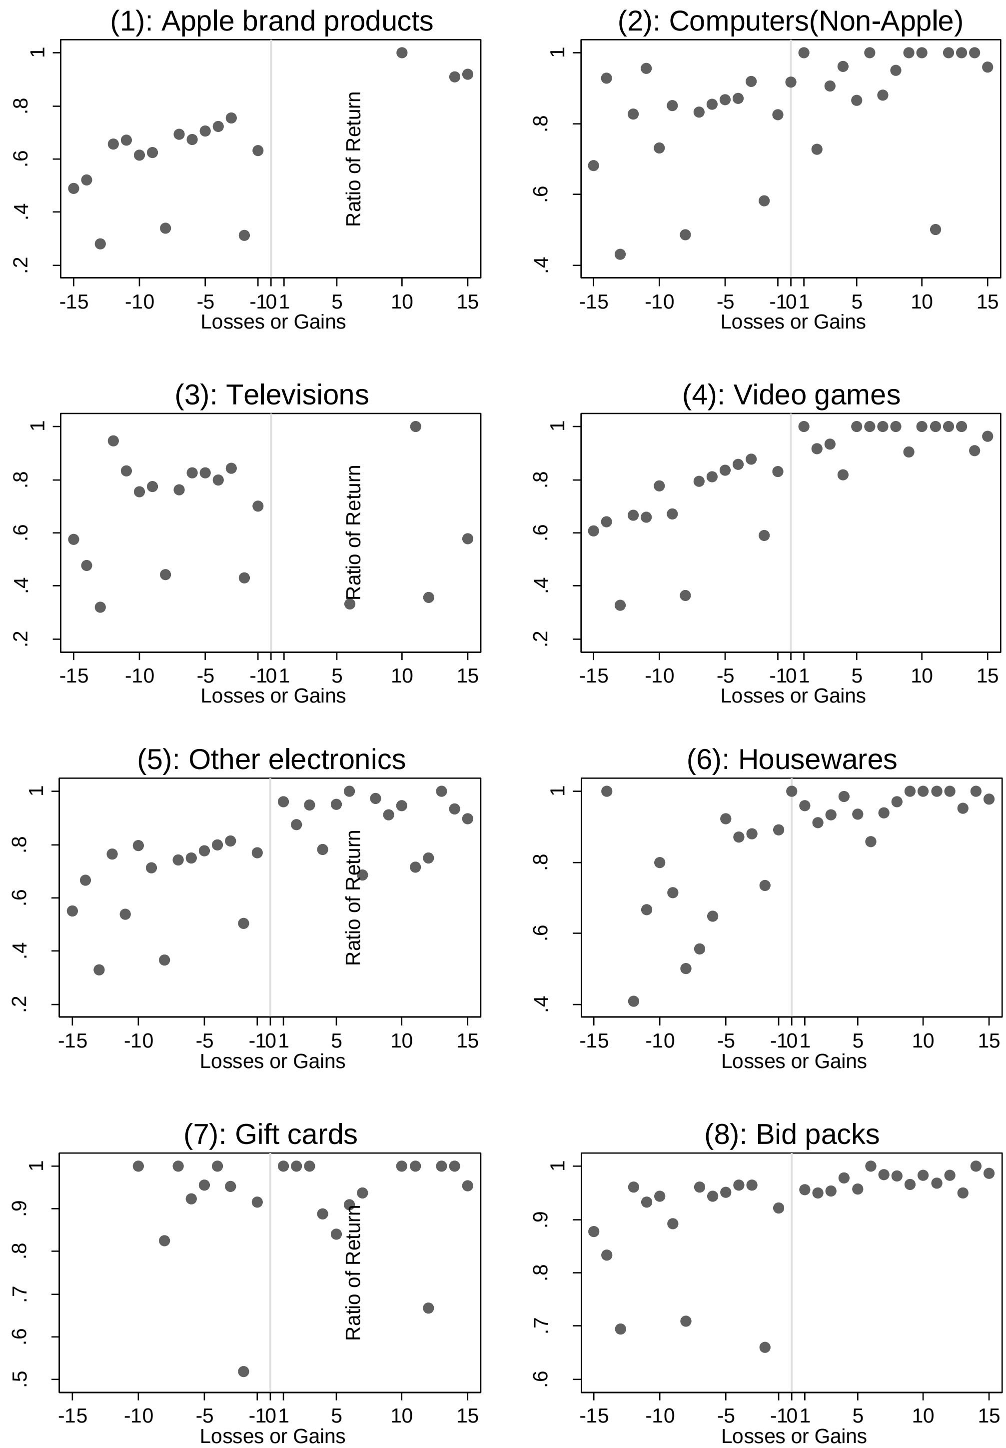

Notes: Only 8 categories with a lot of auctions reported. Apple brand products include iPad, iPhone, MacBook, iPod, etc. Computers (Non-Apple products) include laptop, desktop, notebook, netbook, flash drive, laser mouse, and other accessories made by non-Apple producers, such as HP, Samsung, Toshiba, Gateway, Acer, Compaq, Logitech, Microsoft, PNY, etc. Televisions include all kinds of TVs made by Samsung, Panasonic, Sharp, etc. Video games include those games for Wii, XBOX 360, PS3, etc. Other electronics include products like headphone, camera, Kindle, GPS, speaker, etc. Housewares include coffee maker, ice cream maker, grill, cookware set, kettle, etc. Gift cards include $25, $50, $100 Visa Gift cards. Bid packs include 10 BigDeal Bid Tokens, 20 BigDeal Bid Tokens, 30 BigDeal Bid Tokens, etc.

Appendix B: Number of Auctions Played by Bidders

Table B1: Distribution of Participation Density of Individual Bidders

	1 auction	2 auctions	3 auctions	4 auctions	5 auctions
Num. of bidders	54,731	33,388	23,370	16,284	13,235
	6 auctions	7 auctions	8 auctions	9 auctions	10 auctions
Num. of bidders	9,530	7,905	7,004	5,900	5,921
	11 auctions	12 auctions	13 auctions	14 auctions	15+ auctions
Num. of bidders	4,906	4,000	3,412	2,921	25,198

Appendix C: Additional Bracketing Results

Figure C1: Return Ratio after the 2nd Experience

(Gains and losses calculated over 2 rounds (left) and from 2nd auction only (right))

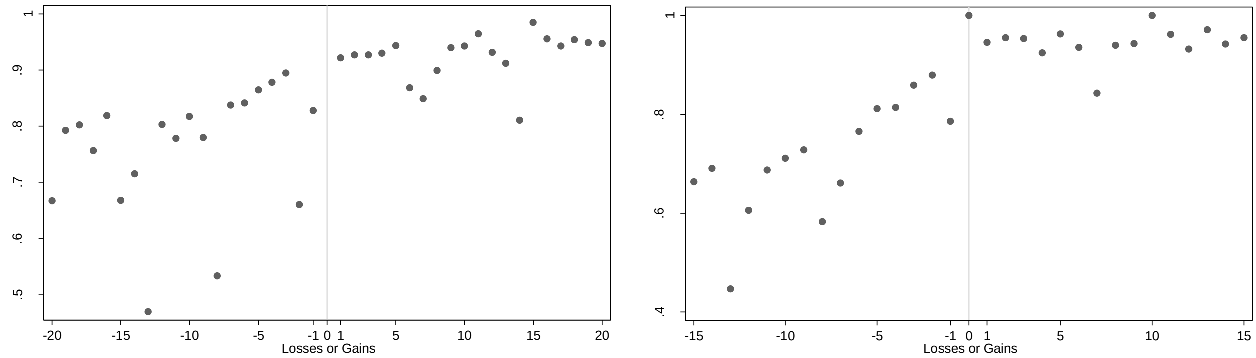

Table C2: Logit Regressions: Bracketing

	Dependent variable: Return = 1 if the bidder came back
	After 2nd experience
Loss (1)	-0.101***
	(0.014)
Loss(1)*magnitude of gain or loss(1)/100	-0.091***
	(0.023)
Magnitude of gain or loss(1)/100	0.045**
	(0.021)
Loss (2)	-0.142***
	(0.045)
Loss(2)*magnitude of gain or loss(2)/100	-0.168***
	(0.026)
Magnitude of gain or loss(2)/100	0.029
	(0.025)
Pseudo-R^2	0.0198
Observations	155,188

Notes: Marginal effects reported, standard errors clustered at the product category level; *significant at 10%; **significant at 5%; ***significant at 1%. Loss and gain refers to nominal losses and gains around zero consumer surplus.

Regression results for the return likelihood after the 3rd, 4th, 6th, 7th, 8th, and 9th experience follow similar patterns and are available upon request. The results for the return likelihood after the 5th and 10th experiences are provided in Section 5 of the paper.

Appendix D: Proofs of Propositions

Proposition 1 (Monotonicity): Under Assumptions 1 and 2, $\forall r \in R, \forall y, y' \in R$,

$$y > y' \Rightarrow U(r,y) > U(r,y').$$

Proof:

$\forall r \in R, \forall y, y' \in R$ such that $y > y'$, we have $U(r,y) - U(r,y')$

$$= \int_q \int_c u(c|r)dF(c|q)dG(q|y) - \int_q \int_c u(c|r)dF(c|q)dG(q|y') = \int_q \int_c u(c|r)dF(c|q)d(G(q|y) - G(q|y'))$$
$$= \int_q U(q,r)d(G(q|y) - G(q|y')) = U(q,r)\cdot(G(q|y) - G(q|y'))|_{q\to 0^+}^{q\to+\infty} - \int_q (G(q|y) - G(q|y'))dU(q,r)$$
$$= \int_q (G(q|y') - G(q|y))dU(q,r).$$

First, we would like to show $\forall r \in R$, $U(q,r)$ is strictly increasing in q.

Note that $\forall q, q' \in R_+$ such that $q > q'$, we have $U(q,r) - U(q',r) = \int_c u(c|r)d(F(c|q) - F(c|q'))$

$= u(c|r)\cdot(F(c|q) - F(c|q'))|_{c\to-\infty}^{c\to+\infty} - \int_c (F(c|q) - F(c|q'))du(c|r) = \int_c (F(c|q') - F(c|q))du(c|r)$. By Assumption 1, $q > q'$ implies $F(c|q)$ *FOSD* $F(c|q')$, which means $F(c|q') \geq F(c|q)$ for all $c \in R$ and $F(c|q') > F(c|q)$ for some $c \in R$. Since $u(c|r)$ is increasing and differentiable in c, we have $\int_c (F(c|q') - F(c|q))du(c|r) > 0$, indicating that $U(q,r)$ is strictly increasing in q.

Second, by Assumption 2, $G(q|y)$ *FOSD* $G(q|y')$, implying $\forall q \in R_+$, $G(q|y') - G(q|y) \geq 0$, and $\exists q^* \in R_+$, $G(q^*|y') - G(q^*|y) > 0$.

The above two results guarantee that $\forall r \in R, \forall y, y' \in R$ such that $y > y'$, we have $\int_q (G(q|y') - G(q|y))dU(q,r) > 0$, or $U(r,y) > U(r,y')$.

This completes the proof.

Proposition 2 (Discontinuity): Under Assumptions 3.1, 3.3 and 3.4,

$$\delta > 0 \Rightarrow U(r(y), y=0) > \lim_{y\to 0^-} U(r(y), y).$$

Proof:

$\forall r \in R, \forall \delta > 0$, $U(r(y), y=0) = \int_q \int_c u(c|r(0))dF(c|q)dG(q|y=0) = \int_q \int_c u(c|r=0)dF(c|q)dG(q)$

$$= \int_q \int_c m(c)dF(c|q)dG(q) + \eta \int_q \int_{c\geq 0} (m(c) - m(0))dF(c|q)dG(q) + \lambda\eta \int_q \int_{c<0} (m(c) - m(0))dF(c|q)dG(q),$$

and

$$\lim_{y\to 0^-} U(r(y), y) = \lim_{y\to 0^-} \int_q \int_c u(c|r(y))dF(c|q)dG(q|y)$$
$$= \lim_{y\to 0^-} \int_q \int_c m(c)dF(c|q)dG(q|y) + \eta \lim_{y\to 0^-} \int_q \int_{c\geq r(y)} (m(c) - m(r(y)))dF(c|q)dG(q|y)$$
$$+ \lambda\eta \lim_{y\to 0^-} \int_q \int_{c<r(y)} (m(c) - m(r(y)))dF(c|q)dG(q|y) - \delta \lim_{y\to 0^-} \int_q \int_{c<r(y)} dF(c|q)dG(q|y)$$

$$= \int_q \int_c m(c) dF(c|q) dG(q) + \eta \lim_{y \to 0^-} \int_q \int_{c \geq r(y)} (m(c) - m(r(y))) dF(c|q) dG(q)$$
$$+ \lambda \eta \lim_{y \to 0^-} \int_q \int_{c < r(y)} (m(c) - m(r(y))) dF(c|q) dG(q) - \delta \lim_{y \to 0^-} \int_q \int_{c < r(y)} dF(c|q) dG(q).$$

Since the reference point formation function $r(\cdot)$ is continuous, we have $\lim_{y \to 0^-} r(y) = r(0) = 0$, and the above expression can be rewritten as $\lim_{y \to 0^-} U(r(y), y)$

$$= \int_q \int_c m(c) dF(c|q) dG(q) + \eta \int_q \int_{c \geq 0} (m(c) - m(0)) dF(c|q) dG(q)$$
$$+ \lambda \eta \int_q \int_{c < 0} (m(c) - m(0)) dF(c|q) dG(q) - \delta \int_q \int_{c < 0} dF(c|q) dG(q).$$

Thus, we have $U(r(y), y = 0) - \lim_{y \to 0^-} U(r(y), y) = \delta \int_q \int_{c < 0} dF(c|q) dG(q) = \delta \int_q F(0^- | q) dG(q) > 0$, since $\forall q \in R_+$, $F(0^- | q) \geq 0$ and $G(q) \geq 0$ and for a positive measure of q, $F(0^- | q) > 0$ and $G(q) > 0$.

This completes the proof.

Proposition 3 (Kink): Under Assumptions 3.1-3.3 and 3.5,

$$\lim_{y \to 0^-} \frac{U(r(y), y < 0) - \lim_{y \to 0^-} U(r(y), y)}{y} > \lim_{y \to 0^+} \frac{U(r(y), y > 0) - U(r(0), y = 0)}{y}.$$

Proof:

$$U(r(y), y > 0) - U(r(0), y = 0) = \int_q \int_c u(c|r(y)) dF(c|q) dG(q|y > 0) - \int_q \int_c u(c|r(0)) dF(c|q) dG(q|y = 0)$$
$$= \int_q \int_c u(c|r(y)) dF(c|q) d[G(q|y > 0) - G(q|y = 0)] + \int_q \int_c [u(c|r(y)) - u(c|r(0))] dF(c|q) dG(q|y = 0)$$

Note that $\lim_{y \to 0^+} \frac{\int_q \int_c u(c|r(y)) dF(c|q) d[G(q|y > 0) - G(q|y = 0)]}{y} = \lim_{y \to 0^+} \frac{\int_q \int_c u(c|r(0)) dF(c|q) d[G(q|y > 0) - G(q|y = 0)]}{y}$

$$= \lim_{y \to 0^+} \int_q \int_c u(c|r(0)) dF(c|q) \frac{\partial g(q|y)}{\partial y} dq = \int_q \int_c u(c|r(0)) dF(c|q) \lim_{y \to 0^+} \frac{\partial g(q|y)}{\partial y} dq.$$

Also note that $\forall y > 0$, $\int_q \int_c [u(c|r(y)) - u(c|r(0))] dF(c|q) dG(q|y = 0)$

$$= \eta \int_q \int_{c \geq r(y)} (m(c) - m(r(y))) dF(c|q) dG(q) + \lambda \eta \int_q \int_{c < r(y)} (m(c) - m(r(y))) dF(c|q) dG(q)$$
$$- \eta \int_q \int_{c \geq 0} (m(c) - m(0)) dF(c|q) dG(q) - \lambda \eta \int_q \int_{c < 0} (m(c) - m(0)) dF(c|q) dG(q)$$
$$= \eta \int_q \left[\int_{c \geq r(y)} (m(c) - m(r(y))) dF(c|q) - \int_{c \geq 0} (m(c) - m(0)) dF(c|q) \right] dG(q)$$
$$+ \lambda \eta \int_q \left[\int_{c < r(y)} (m(c) - m(r(y))) dF(c|q) - \int_{c < 0} (m(c) - m(0)) dF(c|q) \right] dG(q)$$

Assuming $m(\cdot)$, $r(\cdot)$, and $F(\cdot|q)$ are all differentiable almost everywhere, we have

$$\lim_{y\to 0^+}\frac{\eta\int_q\left[\int_{c\geq r(y)}(m(c)-m(r(y)))dF(c\mid q)-\int_{c\geq 0}(m(c)-m(0))dF(c\mid q)\right]dG(q)}{y}$$

$$=\eta\int_q\left[\lim_{y\to 0^+}\frac{\int_{c\geq r(y)}(m(c)-m(r(y)))dF(c\mid q)-\int_{c\geq 0}(m(c)-m(0))dF(c\mid q)}{y}\right]dG(q)$$

$$=\eta\int_q\left[\lim_{y\to 0^+}\frac{\left[\int_{c\geq 0}m(0)-\int_{c\geq r(y)}m(r(y))-\int_{0\leq c<r(y)}m(c)\right]dF(c\mid q)}{y}\right]dG(q)$$

$$=\eta\int_q\left[\lim_{y\to 0^+}\frac{m(0)\left(1-F(0\mid q)\right)-m(r(y))\left(1-F(r(y)\mid q)\right)-\int_{0\leq c<r(y)}m(c)dF(c\mid q)}{y}\right]dG(q)$$

$$=\eta\int_q\lim_{y\to 0^+}\left[-m'(r(y))r'(y)\left(1-F(r(y)\mid q)\right)+m(r(y))f(r(y)\mid q)r'(y)-m(r(y))f(r(y)\mid q)r'(y)\right]dG(q)$$

$$=-\eta\int_q\left[\lim_{y\to 0^+}m'(r(y))r'(y)\left(1-F(r(y)\mid q)\right)\right]dG(q)$$

$$=-\eta\cdot m'(0)\lim_{y\to 0^+}r'(y)\int_q\left[\left(1-F(0\mid q)\right)\right]dG(q)$$

$$=-\eta\cdot m'(0)\lim_{y\to 0^+}r'(y)\left[1-\int_q F(0\mid q)dG(q)\right]$$

Similarly, we have $\lim_{y\to 0^+}\frac{\lambda\eta\int_q\left[\int_{c<r(y)}(m(c)-m(r(y)))dF(c\mid q)-\int_{c<0}(m(c)-m(0))dF(c\mid q)\right]dG(q)}{y}$

$$=\lambda\eta\int_q\left[\lim_{y\to 0^+}\frac{\int_{c<r(y)}(m(c)-m(r(y)))dF(c\mid q)-\int_{c<0}(m(c)-m(0))dF(c\mid q)}{y}\right]dG(q)$$

$$=\lambda\eta\int_q\left[\lim_{y\to 0^+}\frac{\left[\int_{c<0}m(0)-\int_{c<r(y)}m(r(y))+\int_{0\leq c<r(y)}m(c)\right]dF(c\mid q)}{y}\right]dG(q)$$

$$=\lambda\eta\int_q\left[\lim_{y\to 0^+}\frac{m(0)F(0\mid q)-m(r(y))F(r(y)\mid q)+\int_{0\leq c<r(y)}m(c)dF(c\mid q)}{y}\right]dG(q)$$

$$= \lambda\eta \int_q \lim_{y\to 0^+} \left[-m'(r(y))r'(y)F(r(y)\,|\,q) - m(r(y))\,f(r(y)\,|\,q)r'(y) + m(r(y))\,f(r(y)\,|\,q)r'(y) \right] dG(q)$$

$$= -\lambda\eta \int_q \left[\lim_{y\to 0^+} m'(r(y))r'(y)F(r(y)\,|\,q) \right] dG(q)$$

$$= -\lambda\eta \cdot m'(0) \lim_{y\to 0^+} r'(y) \int_q F(0\,|\,q)dG(q)$$

Thus, $\lim_{y\to 0^+} \dfrac{U(r(y), y>0) - U(r(0), y=0)}{y}$

$$= \int_q \int_c u(c\,|\,r(0))dF(c\,|\,q) \lim_{y\to 0^+} \frac{\partial g(q\,|\,y)}{\partial y} dq - \eta \cdot m'(0) \lim_{y\to 0^+} r'(y) \left[1 - \int_q F(0\,|\,q)dG(q) \right] - \lambda\eta \cdot m'(0) \lim_{y\to 0^+} r'(y) \int_q F(0\,|\,q)dG(q)$$

$$= \int_q \int_c u(c\,|\,r(0))dF(c\,|\,q) \lim_{y\to 0^+} \frac{\partial g(q\,|\,y)}{\partial y} dq - \eta \cdot m'(0) \lim_{y\to 0^+} r'(y) \left[1 - \int_q F(0\,|\,q)dG(q) + \lambda \int_q F(0\,|\,q)dG(q) \right]$$

By conducting similar algebra, we get $\lim_{y\to 0^-} \dfrac{U(r(y), y<0) - \lim_{y\to 0^-} U(r(y), y)}{y}$

$$= \int_q \int_c u(c\,|\,r(0))dF(c\,|\,q) \lim_{y\to 0^-} \frac{\partial g(q\,|\,y)}{\partial y} dq - \eta \cdot m'(0) \lim_{y\to 0^-} r'(y) \left[1 - \int_q F(0\,|\,q)dG(q) + \lambda \int_q F(0\,|\,q)dG(q) \right]$$

In this paper we assume $g(q\,|\,y)$ is differentiable everywhere (hence there is no belief asymmetry around $y=0$), which implies $\lim_{y\to 0^+} \dfrac{\partial g(q\,|\,y)}{\partial y} = \lim_{y\to 0^-} \dfrac{\partial g(q\,|\,y)}{\partial y} = \lim_{y\to 0} \dfrac{\partial g(q\,|\,y)}{\partial y}$.

Therefore, $\lim_{y\to 0^-} \dfrac{U(r(y), y<0) - \lim_{y\to 0^-} U(r(y), y)}{y} - \lim_{y\to 0^+} \dfrac{U(r(y), y>0) - U(r(0), y=0)}{y}$

$= \eta \cdot m'(0) \left[\lim_{y\to 0^+} r'(y) - \lim_{y\to 0^-} r'(y) \right] \left[1 - \int_q F(0\,|\,q)dG(q) + \lambda \int_q F(0\,|\,q)dG(q) \right] > 0$, as $\eta > 0$, $m'(0) > 0$, $\lim_{y\to 0^+} r'(y) > \lim_{y\to 0^-} r'(y)$, and $1 - \int_q F(0\,|\,q)dG(q) + \lambda \int_q F(0\,|\,q)dG(q) > 0$.

This completes the proof.

Appendix E: Robustness Checks: Narrow Windows on Nominal Gains and Losses

Table E1: Logit Model, Returning After t^{th} Experience (Magnitude of gain or loss < 10 dollars)

Number of Auctions	Dependent variable: Return = 1 if the bidder returned after the t^{th} time									
Participated, N	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Loss	-0.200**	-0.173***	-0.121**	-0.052	-0.091***	-0.062**	-0.067*	-0.081***	-0.080***	-0.093**
	(0.102)	(0.063)	(0.049)	(0.038)	(0.034)	(0.030)	(0.040)	(0.028)	(0.023)	(0.037)
Loss*magnitude of gain or	-1.145	1.628	0.372	-1.484	-0.169	-0.643	-0.166	0.644	0.556	1.342
loss/100	(2.854)	(1.797)	(0.886)	(0.965)	(0.835)	(0.679)	(1.291)	(0.951)	(1.102)	(1.367)
Magnitude of gain or loss/100	-1.086	-1.145	0.624	1.866**	0.049	0.018	-0.411	-1.202	-1.092	-1.513
	(2.628)	(1.665)	(0.847)	(0.940)	(0.840)	(0.660)	(1.360)	(0.980)	(1.066)	(1.410)
Pseudo-R^2	0.0224	0.0069	0.0075	0.0042	0.0027	0.0050	0.0041	0.0036	0.0039	0.0011
Observations	176,229	133,453	105,472	86,092	73,116	62,538	55,043	48,829	42,956	37,908

Notes: Marginal-effects reported, standard errors clustered by product category; *significant at 10%; **significant at 5%; ***significant at 1%. The column (t) shows the results of returning after the t^{th} auction. Loss and gain refers to nominal losses and gains around zero consumer surplus.

Table E2: Logit Model, Returning After t^{th} Experience (Magnitude of gain or loss < 50 dollars)

Number of Auctions	Dependent variable: Return = 1 if the bidder returned after the t^{th} time									
Participated, N	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Loss	-0.205***	-0.157***	-0.131***	-0.098***	-0.087***	-0.062***	-0.058**	-0.052**	-0.058**	-0.060*
	(0.067)	(0.049)	(0.039)	(0.032)	(0.029)	(0.024)	(0.023)	(0.023)	(0.025)	(0.031)
Loss*magnitude of gain or	-0.052	-0.027	-0.106	-0.291*	-0.291***	-0.454***	-0.467***	-0.471***	-0.422***	-0.312***
loss/100	(0.387)	(0.190)	(0.173)	(0.155)	(0.104)	(0.139)	(0.080)	(0.102)	(0.131)	(0.084)
Magnitude of gain or loss/100	-0.560*	-0.317*	-0.183	-0.020	0.016	0.062	0.147**	0.150*	0.122	0.083
	(0.339)	(0.178)	(0.183)	(0.166)	(0.097)	(0.131)	(0.063)	(0.087)	(0.117)	(0.084)
Pseudo-R^2	0.0262	0.0127	0.0103	0.0111	0.0102	0.0186	0.0142	0.0138	0.0142	0.0073
Observations	202,971	152,388	121,082	99,145	84,125	71,854	62,976	55,737	49,121	43,585

Notes: Marginal-effects reported, standard errors clustered by product category; *significant at 10%; **significant at 5%; ***significant at 1%. The column (t) shows the results of returning after the t^{th} auction. Loss and gain refers to nominal losses and gains around zero consumer surplus.

Table E3: Logit Model, Returning After tth Experience (Magnitude of gain or loss < 100 dollars)

Number of Auctions Participated, N	Dependent variable: Return = 1 if the bidder returned after the t^{th} time									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Loss	-0.199***	-0.152***	-0.128***	-0.101***	-0.091***	-0.071***	-0.069***	-0.067**	-0.070***	-0.065*
	(0.065)	(0.049)	(0.039)	(0.033)	(0.033)	(0.027)	(0.026)	(0.029)	(0.024)	(0.035)
Loss*magnitude of gain or loss/100	-0.162	-0.115	-0.158*	-0.220**	-0.187	-0.257***	-0.246**	-0.226**	-0.207*	-0.221*
	(0.165)	(0.101)	(0.091)	(0.101)	(0.117)	(0.096)	(0.110)	(0.098)	(0.119)	(0.134)
Magnitude of gain or loss/100	-0.283**	-0.132	-0.070	-0.019	-0.032	0.014	0.035	0.007	0.011	0.062
	(0.121)	(0.083)	(0.078)	(0.096)	(0.110)	(0.086)	(0.096)	(0.094)	(0.116)	(0.137)
Pseudo-R^2	0.0266	0.0131	0.0117	0.0135	0.0110	0.0163	0.0145	0.0145	0.0141	0.0084
Observations	205,701	154,307	122,686	100,625	85,346	72,950	63,996	56,612	49,945	44,352

Notes: Marginal-effects reported, standard errors clustered by product category; *significant at 10%; **significant at 5%; ***significant at 1%. The column (t) shows the results of returning after the t^{th} auction. Loss and gain refers to nominal losses and gains around zero consumer surplus.

Table E4: Logit Model, Returning After tth Experience (Magnitude of gain or loss < 200 dollars)

Number of Auctions Participated, N	Dependent variable: Return = 1 if the bidder returned after the t^{th} time									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Loss	-0.194***	-0.146***	-0.128***	-0.100***	-0.090***	-0.073***	-0.070***	-0.068**	-0.071***	-0.067**
	(0.062)	(0.047)	(0.038)	(0.034)	(0.031)	(0.028)	(0.025)	(0.027)	(0.021)	(0.031)
Loss*magnitude of gain or loss/100	-0.233**	-0.196***	-0.136**	-0.202***	-0.164**	-0.190***	-0.194***	-0.170***	-0.151**	-0.167***
	(0.093)	(0.059)	(0.063)	(0.071)	(0.067)	(0.041)	(0.059)	(0.060)	(0.072)	(0.063)
Magnitude of gain or loss/100	-0.126**	-0.002	-0.037	0.020	-0.001	0.012	0.046	0.015	0.017	0.057
	(0.060)	(0.049)	(0.050)	(0.065)	(0.057)	(0.039)	(0.046)	(0.054)	(0.063)	(0.063)
Pseudo-R^2	0.0252	0.0133	0.0114	0.0129	0.0107	0.0150	0.0131	0.0130	0.0125	0.0082
Observations	206,614	154,921	123,242	101,111	85,810	73,368	64,402	56,951	50,284	44,690

Notes: Marginal-effects reported, standard errors clustered by product category; *significant at 10%; **significant at 5%; ***significant at 1%. The column (t) shows the results of returning after the t^{th} auction. Loss and gain refers to nominal losses and gains around zero consumer surplus.