

Paying for a Chance to Save Money: Two-Part Tariffs in Name-Your-Own-Price Markets

September 26, 2017

Robert Zeithammer, Lucas Stich, Martin Spann, and Gerald Häubl

Abstract

We examine the profitability of two-part tariffs in name-your-own-price (NYOP) markets using incentive-compatible laboratory experiments. An NYOP seller who uses a two-part tariff charges an upfront non-refundable bidding fee akin to an entrance fee into his store. We find two-part tariffs can be profitable to NYOP sellers, and we identify a strong moderator of this profitability: a decision aid that calculates the payoff consequences of candidate bids. To compare the profitability of two-part tariffs with that of other profit-enhancing strategies proposed in the literature, as well as to interpret the effect of our decision-aid manipulation, we propose and estimate at the individual level a model of entry and bidding that nests both Expected Utility Theory and Prospect Theory as special cases. Counterfactual simulations based on the estimated model reveal two-part tariffs outperform the minimum-bid strategy, and the optimal bidding fee depends on the presence of our decision aid and on subjects' experience with our bidding task. We estimate that optimal two-part tariffs can increase NYOP seller profits by about 9% in the long run.

Keywords: Pricing, Auctions, Behavioral Economics, Incentive Compatible Laboratory Experiment

Contact: Robert Zeithammer, rzeitham@ucla.edu

INTRODUCTION

Name your own price (NYOP) refers to a pricing mechanism pioneered by Priceline in 1997 to sell travel products, and adopted by a wide variety of sellers since then to sell other products and services.¹ An NYOP buyer “bids” a binding price-offer to the seller, and the seller decides whether to accept or reject the bid. When the seller accepts a buyer’s bid, the buyer pays her bid as the purchase price, and the seller keeps the difference between the bid and his procurement cost. Despite the attractive features of NYOP, such as its ability to empower buyers and find market-clearing prices for many different products, NYOP remains a niche strategy. We believe one of the reasons NYOP has not diffused more widely is that its existing implementations leave potential seller profits on the table. Using incentive-compatible laboratory experiments, we test whether a two-part tariff can increase NYOP’s profitability. Instead of providing the bidding opportunity free of charge, an NYOP seller who uses a two-part tariff charges an upfront non-refundable fee akin to an entrance fee into his store. Note our two-part tariff (i.e., fee + bid) is different from the “bidding fee” in all-pay auctions empirically studied, for example, by Platt, Brennan, and Tappen (2013). We study a single bidder who pays the fee only once, and who then submits just one binding sealed bid. By contrast, all-pay auctions are not two-part tariffs but dynamic games among multiple bidders who pay a small amount for every bid increment if they want to retain a chance of winning.

We find two-part tariffs can be profitable for NYOP sellers, but the magnitude of the profit increase is strongly moderated by the buyer’s access to a decision aid that helps her calculate the payoff consequences of different potential bid levels. Specifically, the decision aid takes a potential bid as input, and shows the probability of the bid being accepted by the seller as well as the surplus due to the bidder

¹ Examples include event tickets at scorebig.com, electronics at greentoe.com, and a variety of products on eBay.com via its “Best Offer” option.

upon bid acceptance. We take the decision aid's moderating influence on profits as evidence that the profitability of a two-part tariff in NYOP depends on buyers' cognitive ability to solve her optimal entry and bidding problem. The unique bidding data produced by our experiment allow us to estimate a very flexible model of preferences at the individual level. The model estimates shed light on the process behind the decision aid's profit-moderating effect. We also use the estimated model to answer counterfactual questions about profitability of fee levels not included in the experimental design, and profitability of alternative profit-increasing strategies proposed in the literature.

A two-part tariff is a popular posted-pricing model in service industries (e.g., internet access, fitness clubs) as well as warehouse clubs (e.g., Costco). By contrast, we are not aware of a real-world application of a two-part tariff to NYOP. Theoretically, the same efficiency-enhancing property that makes two-part tariffs profitable for a posted-price seller also operates within NYOP as long as buyers are risk neutral (Spann, Zeithammer, and Häubl 2010). This paper presents the first empirical test of a two-part tariff in NYOP with real bidders who may not be risk neutral.

We present the results of two studies. In our first study, we find a two-part tariff can be profitable for the NYOP seller, but the optimal bidding fee is smaller than what a model with risk-neutral buyers suggests. Surprisingly, given ample evidence of risk aversion among our subjects, the risk-neutral model tends to *under*-predict the profitability of NYOP selling—a deviation we trace to excessive buyer entry.

To analyze whether this excessive entry is caused by bidders' biased assessment of the probabilistic consequences of their potential bids, we designed a tool to help the buyers with this assessment. Specifically, in our second study, we show that a "decision aid" that helps buyers calculate the probability and payoff consequences of their planned actions reduces their willingness to pay bidding fees, which in turn reduces the profitability of using a two-part tariff. Additional buyer experience with NYOP bidding compounds the negative effect of our decision aid on the profitability of bidding fees. As

a result, when buyers have experience and access to the decision aid, none of the few positive fee levels used in our experimental design is significantly more profitable than the zero level. But what would be the expected profit of an NYOP seller facing such buyers and charging other intermediate fee levels? How does the two-part tariff strategy compare with other profit-enhancing strategies suggested in the literature, such as the minimum-bid strategy? And what aspect of buyer preferences does our decision aid influence? To answer these counterfactual questions, we propose a flexible semi-parametric model of preferences, and estimate it at the individual level using the data from our second study.

Bidding behavior has long been recognized as a potentially rich source of information about preferences, because bids contain not only ordinal information about whether the bidder prefers the good on offer to the outside option, but also cardinal information about the strength of that preference. Because most markets that give rise to bidding data are governed by known and strictly enforced auction rules, empirical researchers should be able to extract the preference information from bidding data (Sutton 1993). Indeed, assuming a particular shape of the bidder's utility function (of monetary surplus) and "inverting" the bidding problem to obtain the bidder's underlying valuation of the good being sold is possible. For example, Guerre, Perrigne, and Vuong (2000) assume the utility function is an identity, that is, that the bidder is risk neutral, and perform the inversion on data from a first-price sealed-bid auction of which NYOP is a special case.² However, the joint identification of the utility function and the underlying distribution of valuations is impossible from bidding data alone (Guerre, Perrigne, and Vuong 2009), so the existing approaches for analysis of bidding data need to make strong assumptions about the shape of the bidders' utility functions. In our experimental paradigm, valuations are induced and controlled by the experimenter, and bids arising from these valuations are measured, leaving only the

² From the perspective of the buyer—the only perspective that matters for interpreting bids—NYOP is a first-price sealed-bid auction with a single bidder and a random hidden reserve price.

utility function as the unknown. Therefore, our model can estimate the shape of the utility function instead of assuming it. Because we only use a few levels of the bidding fee, we can model the utility function nonparametrically in the loss domain. In other words, we only need to specify the utility function at a few negative-surplus points to fit our data. In the gain domain, that is, for positive surplus levels, we use a standard functional form from the literature. Overall, we thus estimate a semi-parametric model of utility, and we estimate it at the individual level.

For almost all of our subjects, our results reject not only risk neutrality, but also a commonly assumed model of risk aversion arising from Expected Utility Theory with a globally concave utility function. A notable special case of our semiparametric utility model is Prospect Theory (Kahneman and Tversky 1979) with gains and losses defined relative to the natural reference point of earning zero surplus, and with a single-parameter model of subjective probability weighting as proposed by Prelec (1989). We find the preferences of about half our bidders are consistent with such a version of Prospect Theory in that they are convex in the loss domain and concave in the gain domain. However, our data do not suggest a prospect-theoretic alternative model can be imposed on all subjects: the preferences of half of our subjects cannot be classified in that they do not exhibit a consistent global curvature pattern.

We now briefly preview the answers to our three counterfactual questions. Regarding the effect of our decision aid on buyer preferences, we find the decision aid makes the utility functions steeper in the loss domain, especially for very small losses. The same effect occurs within subject, but the crossover design of our study confounds the within-subject measurement of the decision aid's effect with the effect of additional experience. We find that, within subject and jointly with experience, the decision aid also reduces both the extent and the incidence of prospect-theoretic probability weighting, and the model fits the data better. Overall, the average preferences of experienced buyers with the decision aid can thus be captured with a simple prospect-theoretic model with an s-shaped value function and no probability

weighting. Regarding the profitability of two-part tariffs and other selling strategies, the counterfactual simulations suggest the optimal bidding fee depends on the decision aid and experience of the buyers: when the buyers have no aid or experience, the optimal fee is sizeable and raises expected profits by 36% relative to a zero fee (less than the 67% lift predicted by the risk-neutral theory). Providing buyers with some experience and the decision aid reduces the optimal bidding fee, and reduces its profit lift to only 9%. In either case, the optimal two-part tariff outperforms the optimal minimum bid, contrary to the risk-neutral predictions of Zeithammer (2015). However, we also find the minimum-bid strategy is more robust in that the superiority of bidding fees depends on selecting the optimal fee carefully while the profits as a function of minimum bid are relatively “flat at the top.”

The remainder of this paper is structured as follows. After a brief review of the literature, we describe our experimental paradigm shared by both studies. The next section then presents the design and results of Experiment 1. The key finding of Experiment 1 is the excess buyer entry—a behavior our decision-aid intervention examined in Experiment 2 was design to moderate. We present the results of Experiment 2 in two separate sections, starting with the model-free evidence. After proposing our utility model and explaining our estimation strategy, we return to the analysis of Experiment 2 through the lens of model parameters. Finally, we present the results of our counterfactual simulations and conclude the paper with a general discussion.

LITERATURE REVIEW

This paper is related to three strands of literature: (1) the literature on NYOP selling, (2) the experimental auction literature, and (3) the literature on prospect-theoretic preferences. We discuss our contribution to the three strands in turn.

The majority of prior research on NYOP pricing is analytical and focuses on sellers' design decisions such as responding to repeat bidding (Fay 2004), facilitating joint bidding for multiple items (Amaldoss and Jain 2008), charging bidding fees or committing to minimum markups (Spann, Zeithammer, and Häubl 2010), or committing to the optimal bid-acceptance schedule (Zeithammer 2015). Another stream of research gives reasons for the emergence of the NYOP channel, including its ability to soften competition (Fay 2009), exploit buyer risk aversion (Shapiro 2011), achieve price discrimination based on haggling friction costs (Terwiesch, Savin, and Hann 2005), and adapt to uncertain demand (Wang, Gal-Or, and Chatterjee 2009). All but one of the papers mentioned above rely on the assumption of buyer risk neutrality. The one exception is Shapiro (2011), who assumed risk-averse buyers. We contribute to the theoretical literature by directly testing the predictions of Spann, Zeithammer, and Häubl (2010) that a two-part tariff can be profitable for an NYOP seller, and by developing and empirically validating a realistic structural model of bidding in NYOP settings.

We also contribute to the relatively smaller literature on laboratory tests of analytical model predictions regarding particular NYOP seller strategies, such as different threshold-setting strategies (Hinz, Hann, and Spann 2011), different modes of information diffusion about sellers' threshold level (Hinz and Spann 2008), or the opacity of the NYOP offering (Shapiro and Zillante 2009). The most related paper in this literature is the work by Bernhardt and Spann (2010), who study the effects of transaction fees on buyer behavior. In contrast to our setting, Bernhardt and Spann (2010) analyze fees that accrue only in the event of a successful bid, and they do not consider the NYOP seller's competition with the

outside posted-price market. They find such transaction fees can increase seller profit, because they make consumers bid by higher increments.

By proposing a behaviorally realistic structural model, we bridge the gap between the analytical literature rooted in economics and the behavioral literature rooted in psychology. Our model considers the impact of emotions (Ding et al. 2005), menus of possible bids (Chernev 2003; Spann et al. 2012), and expectations about changes in sellers' threshold level (Fay and Laran 2009; Fay and Lee 2015) on buyer behavior. Relatedly, our proposed model also contributes to the literature that uses NYOP selling as a convenient measurement tool of consumer characteristics such as risk aversion (Abbas and Hann 2010), frictional costs of online transactions (Hann and Terwiesch 2003), willingness to pay (Spann, Skiera, and Schäfers 2004), and buyer haggling costs (Terwiesch, Savin, and Hann 2005).

Second, our paper is related to the large literature in experimental economics on consumer behavior in first-price sealed-bid (1PSB) auctions. The decision of an NYOP bidder is simpler than that of a 1PSB bidder, because an NYOP bidder does not compete with other potential buyers. Therefore, NYOP bidding provides a clearer empirical setting for studying the impact of preferences on bidding behavior by avoiding the need for both subjects and the analyst to understand the equilibrium of the bidding game. Nevertheless, our results are consistent with two major findings of the empirical literature on 1PSB bidding: overbidding and over-entry relative to a risk-neutral model. Overbidding is one of the consistent findings in this literature, and thus a large body of work has focused on explaining that phenomenon (Cox, Roberson, and Smith 1982; Cox, Smith, and Walker 1988). The most common explanation has been risk aversion (Cox, Smith, and Walker 1988; Filiz-Ozbay and Ozbay 2007). However, other factors that have been considered include the misperception of winning probabilities (Dorsey and Razzolini 2003), anticipated regret (Filiz-Ozbay and Ozbay 2007), and the joy of winning (Ertaç, Hortaçsu, and Roberts 2011). We find overbidding can occur in the absence of equilibrium considerations, is not affected by a decision aid that computes the

probability of winning, and can arise in a prospect-theoretic model.

Another consistent finding in 1PSB auctions is over-entry. Palfrey and Pevnitskaya (2008) provide an excellent review of the literature on over-entry, and propose an entertainment value of bidding in an auction as an explanation of the over-entry phenomenon. Ertaç et al. (2011) propose a model that combines risk aversion and the joy of winning to explain both bidding behavior and endogenous entry behavior. They show a model incorporating joy of winning together with risk aversion better matches the observed entry behavior than a model lacking that component. We do not find evidence of joy of winning (recall that “winning” NYOP bidding merely arises from the seller accepting the bid; there are no other buyers to beat), and our model accommodates over-entry by allowing the value function to be relatively flat in the loss domain. In other words, our data are consistent with over-entry arising from buyers’ insensitivity to the magnitude of losses.

The third strand of the literature our paper is related to is work in marketing and economics on risk aversion (e.g., Charness, Gneezy, and Imas 2013, Dohmen et al. 2012, DellaVigna 2009; Köszegi and Rabin 2007) and loss aversion (e.g., Kahneman and Tversky 1979; Tversky and Kahneman 1991, Andersson et al. 2014). We contribute to this literature by proposing and estimating a flexible model that backs out parameters of prospect theory from bidding data at the individual level. As far as we can tell, we are the first to estimate a prospect-theoretic model of bidding. By estimating at the individual level, we can shed light on the population distribution of the model parameters, documenting both their heterogeneity and their mutual correlations. We find substantial and heterogeneous curvatures of the value function and the probability weighting function. Interestingly, we do not find evidence of loss aversion.

***EXPERIMENTAL PARADIGM: WITHIN-SUBJECT MANIPULATION OF
BUYER'S VALUATIONS AND NYOP SELLER'S BIDDING FEES***

Throughout this paper, we use incentive-compatible laboratory experiments to study entry and bidding behavior of buyers facing an NYOP seller, who uses a two-part tariff. All of our studies share several design elements: induced buyer valuations, a stylized exogenous marketplace motivated by theoretical models of NYOP retailing, and within-subject manipulation of both buyer valuations and NYOP bidding fees. In this section, we describe and justify all of these shared elements in turn.

We want to focus on understanding the entry and bidding strategies in the simplest possible setting of a single buyer with unit demand for one particular object. To gather sufficient data about each subject for individual-level analysis, we repeat the simple setting in a series of rounds. In our laboratory paradigm, the object in each round is always a virtual token with induced value, allowing us to vary and control buyer valuation while abstracting away from specific product categories (Smith 1976). The token has “induced value” in that we pay the buyer a pre-announced amount of experimental currency whenever she owns the token at the end of an experimental round, and pay her nothing when she does not. We have control over the buyer’s valuation because we set the valuation amount at the beginning of each round. To allow for subsequent pooling of the data at different valuation levels, we draw the individual buyer valuations in different rounds from several equispaced discrete points, for example, {5, 20, 35, 50, 65, 80, 95} in Experiment 1. All of the existing analytical results focus on the theoretically tractable uniform distribution of valuations, so we draw each point with the same probability. Note the distribution of buyer valuations only influences the expected profit of the seller, not the individual incentives of buyers.

Each experimental subject is assigned to the role of a buyer and faces a stylized computer-simulated marketplace. We designed the computer-simulated marketplace to implement the supply-side modeling assumptions of existing analytical models of NYOP retailing (e.g., Spann, Zeithammer, and

Häubl 2010, Shapiro 2011, or Zeithammer 2015) as follows. Two stores exist in the marketplace. The object is readily available from one of the stores for a posted price p (e.g., $p=70$ in both our studies). The other store in the marketplace uses NYOP selling, procures the object for a privately known procurement cost w (for “wholesale cost”), and accepts all buyer bids above his cost. This bid-acceptance strategy is thoroughly explained to the subjects upfront, along with the fact that the seller’s cost w is distributed *iid* uniformly between rounds on $[0,p]$.³ Note that for measurement of buyer behavior, it is sufficient for subjects to believe the resulting bid-acceptance probabilities—they do not need to understand anything about the relationship between the bid-acceptance probability and the seller’s wholesale cost. Finally, the supply side also always includes a “Don’t Buy In This Round” option to allow for a meaningful measurement of entry behavior (please see Figure A1 in the Appendix for the screen layout, and see the Web Appendix for our complete instructions and procedures).

Every round represents an independent market with one buyer on the demand side and the two stores described above on the supply side. The only aspect of the supply side that varies between rounds is the bidding fee at the NYOP seller. We vary the fee among a few discrete points⁴ between 0 that everyone should be willing to pay to 18, which, according to our calculations, no risk-neutral buyer with a valuation below 100 should be willing to pay when $p=70$. We include three theoretically motivated intermediate levels between 0 and 18: The first intermediate fee level we use is 1—the smallest possible positive fee given the experimental currency. One possible reason against using bidding fees is that human buyers may exhibit “fee aversion” and use a decision rule to never pay for anything other than the product

³ Realistically, the outside posted price is thus a public upper bound on the NYOP seller’s procurement cost. The wholesale cost uncertainty arises from the producer’s (e.g., a hotel’s) opportunity cost of not filling its full capacity using standard posted pricing (Belobaba 1989). Such an opportunity cost varies over time and is specific to details of the product sold.

⁴ We could have selected fees randomly from a continuum between 0 and 18. By contrast, selecting only a few specific values not only allows for a clean full-factorial within-subject design, but also enables the flexible nonparametric utility specification in our structural model.

itself (Amir and Ariely 2007). Evidence for such a fee aversion would arise in our experiment if high-valuation buyers who face a fee of 1 enter less frequently than buyers with the same high valuation who face a fee of zero. The second intermediate fee level we use is 12—the level approximately optimal under risk neutrality when $p=70$ and $M=100$ (see Experiment 1 for details of the calculation). Overwhelming evidence suggests people are risk averse and loss averse. Either aversion should intuitively reduce the optimal bidding fee, so we add one more intermediate fee level of 6 as an intermediate, non-negligible value. We thus test the following five fee levels: $\{0, 1, 6, 12, 18\}$.

Each subject experiences all possible combinations of valuations and fees (in random order across subjects), resulting in a full-factorial within-subject design. The goal of this design is to eventually analyze buyer behavior at the individual level, allowing a non-parametric approach to preference heterogeneity.

EXPERIMENT 1: TEST OF THE RISK-NEUTRAL BENCHMARK MODEL

Our first study is motivated by the theoretical predictions of Spann, Zeithammer, and Häubl (2010), who consider risk-neutral buyers with valuations drawn from the uniform distribution and find positive bidding fees can increase the expected profit of an NYOP seller in a marketplace described above. Specifically, they show the optimal bidding fee is $f^*(p) = \frac{4M^2}{49p}$ as long as $p > \frac{4M}{7}$. In experiment 1, we assume $M=100$ and $p=70$, resulting in $f^*(p) = 11.7$, and we approximate the Uniform $[0,M]$ distribution by drawing the buyer valuations from $\{5, 20, 35, 50, 65, 80, 95\}$. Therefore, we use a $5(\text{fee levels}) \times 7(\text{valuation levels})$ within-subject design, measuring each subject's entry and bidding behavior in 35 conditions.

We now briefly summarize the predicted behavior of a risk-neutral buyer facing a bidding fee f charged by an NYOP seller who accepts bids that exceed his wholesale cost drawn from Uniform $[0,p]$

distribution; see Spann, Zeithammer, and Häubl (2010) for details. For any $f < \frac{p}{4}$, buyers with $v > 2\sqrt{pf}$ pay the bidding fee and submit a bid to the NYOP seller. Two types of bidders emerge: “low” bidders with $v < p$ who cannot afford the outside option bid $b(v) = \frac{v}{2}$, and “high” bidders with $v \geq p$ who mimic the bidder with $v = p$ and bid $\frac{p}{2}$ because they have a real option of buying in the outside market should their NYOP bid not be successful. We now describe the detailed experimental procedure to test these predictions.

Experimental Procedure

Each round carries out the following experimental procedure: the buyer is informed about her private valuation, the bidding fee at store A (NYOP store), and the posted price at store B (posted-price store). The buyer then has to decide whether to bid in store A, buy from store B, or skip the round. If she chooses to bid, she enters the bid amount into a box and presses “Submit Bid,” automatically deducting the bidding fee (please see Figure A1 in the Appendix for the screen layout). To decide whether a bid is accepted, we draw the secret threshold price w from a uniform distribution on $[0, p]$. If a bid is rejected, the buyer is given a second chance to buy from the posted-price store.

If the buyer decides not to buy in a round, she receives a payment of 0 points and the round ends. The income of a buyer in a round in which she purchased the product from the posted-price store is her valuation minus the posted price. If the NYOP store accepts a buyer’s bid, the buyer’s payoff on this round is her valuation minus the bid submitted and minus the bidding fee. If the NYOP store rejects a buyer’s bid, the buyer’s tentative payoff on this round is 0 minus the fee. The final payoff, however, is contingent on her subsequent decisions. If she decides not to use this second chance to buy from the posted-price store, her final payoff from this round is 0 minus the fee. On the other hand, if she wants to make use of

the second chance, her final payoff is her valuation minus the posted price and minus the fee. Subjects are shown their profit after each round but not their total profit to limit potential wealth effects.

A pilot study found some subjects use the first few rounds to explore actual consequences of seemingly irrational behavior, such as bidding above one's valuation. Once they incur an avoidable loss, most subjects refrain from such behavior in later rounds. To give the subjects an opportunity for such an exploration without compromising our experimental design, we included five "training" rounds in the beginning of the session. The rounds were not marked in any way to the subjects, who simply experienced them as the first five rounds of the experiment, and we discarded the data. To encourage bidding, we kept the fees low during the training rounds. To expose the subjects to the second chance should their bid not be accepted, we also included a valuation above p . Specifically, the five rounds exposed the subjects to the following (*fee, valuation*) pairs: (1, 65), (0, 5), (6, 80), (6, 5), (0, 50).

After completing the five training rounds and the subsequent 35 experimental rounds, subjects were asked to answer an exit survey for additional credit. The exit survey focused primarily on measures of individual differences we hypothesized to be related to entry and bidding behavior: number of "safe" choices in the paired lottery choice task by Holt and Laury (2002); number of rejected risky lotteries in the lottery choice task by Gächter et al. (2010); subjective risk-taker scale by Dohmen et al. (2012): "How do you see yourself: Are you generally a person who is fully prepared to take risks or do you try to avoid taking risks?" (1=*completely unwilling to take risks*, 7=*completely willing to take risks*); a question about attitude toward bidding fees, "I want to pay for the actual product only, I refuse a bidding fee in general" (1=*strongly disagree*, 7=*strongly agree*); a question about mood, "How would you rate your general mood today?" (1=*very bad*, 5=*very good*); and frequency of participation in lab experiments to date. All scales were administered as hypotheticals, not separately incentivized beyond a flat payment.

Data Collection

We conducted four sessions of the experiment at the Munich Experimental Laboratory for Economic and Social Sciences (MELESSA) of the University of Munich in 2015. Subjects were mainly undergraduate students from the University of Munich and the Technical University of Munich studying a wide range of majors. We used the software z-Tree (Fischbacher 2007) and ORSEE (Greiner 2015) to program and conduct the experiments. Subjects took about 45 minutes to complete the main part of the study—a little over a minute per task. Subjects earned on average about 16.90 EUR (USD 21.70 at the time of the experiments), which included a show-up fee of 4 EUR (USD 5.10) and another 4 EUR for taking the exit survey. With 24 subjects per session, a total of 96 subjects participated in Experiment 1. We found the following evidence of learning during the five training rounds: 15 subjects bid over their valuation at least once during the five training rounds, but only five subjects did so in the subsequent 35 rounds. The per-round incidence of such seemingly irrational behavior was thus sharply reduced but not entirely eliminated in the subsequent rounds. We exclude the five subjects who continued bidding over valuation even after the training rounds from our analysis, because we question whether they understood the experimental procedure. Everything that follows is based on the 91 remaining subjects.

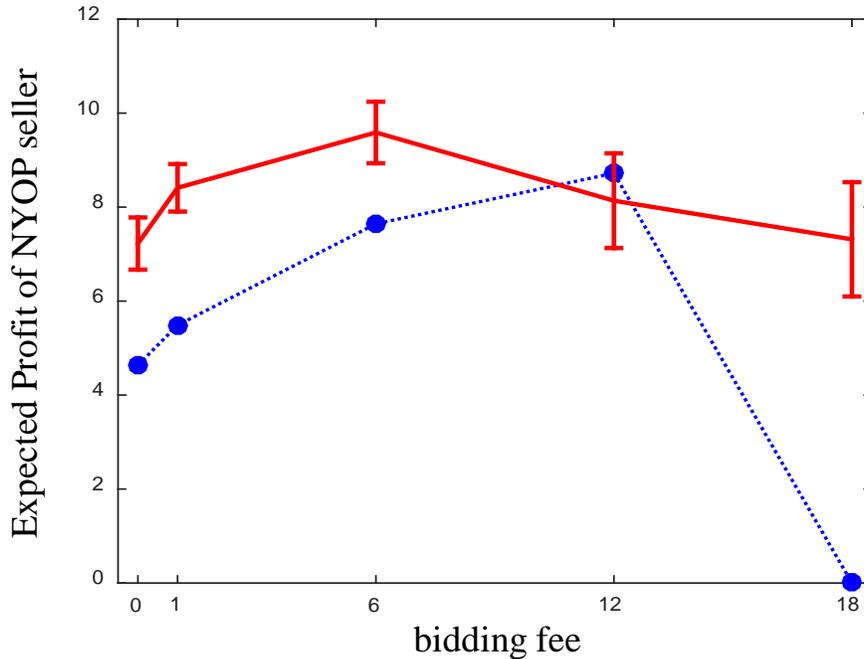
Experiment 1 Results: Seller Profits

We begin with an analysis of the key managerially relevant statistic, namely, the expected NYOP seller profit conditional on the observed bids. The NYOP seller's profit consists of the bidding fee paid by the buyers who enter the NYOP store plus the difference between the wholesale cost w and bid whenever a bid gets accepted. In the experiment, we draw an actual w randomly in each round to determine bid acceptance, but we average over this “noise” in our analysis below by computing the expected seller profit from each observed bid, denoted $\pi(\textit{bid})$ as follows:

$$\pi(\text{bid}) \equiv \int_0^{\text{bid}} (\text{bid} - w) \left(\frac{1}{p} \right) dw = \frac{\text{bid}^2}{2p} \quad (1)$$

In words, we take the bid as given, and average the unit contribution over all possible wholesale cost realizations. Let $\text{bid}_{i,v,f}$ submitted by subject i with induced valuation v when the bidding fee is f , and let $\text{bid}_{i,v,f} = 0$ when that subject does not enter. Then we calculate the expected NYOP seller profit when the fee is f as the average over i and v of $\pi(\text{bid}_{i,v,f})$, generating an estimate of $E_{i,v}[\pi(\text{bid}_{i,v,f})]$. Figure 1 plots this expected seller profit, and compares it to the theoretical profit expected under risk neutrality.

Figure 1: NYOP Seller Profit



Note to Figure: The solid (red) line connects the observed expected profits, displayed as error bars. Each error bar represents the 95% confidence interval with one observation defined as the expected profit from one subject, so $N=91$. The dotted (blue) line connects the expected profits predicted by risk-neutral buyers.

We start our discussion of Figure 1 by considering the main prediction of the literature on two-part tariffs, namely, that positive bidding fees generate higher profits than zero fees. Figure 1 shows the bidding fee of 6 generates 2.3 units more expected profit than no fee—an increase of 33% ($p < 0.001$ regardless of

how the relevant t-test is set up).⁵ The expected profit with $f=6$ is also significantly greater than the adjacent fee levels in our design, so the optimal fee to charge from the set $\{0, 1, 6, 12, 18\}$ is $f=6$. Whereas the main qualitative prediction of the two-part tariff literature thus holds in our data, the exact quantitative prediction based on a risk-neutral model does not: the empirically optimal fee is lower than the theoretically predicted level of $f=12$, and the observed profit exceeds the predicted profit for all fee levels other than 12. We summarize these findings in our first result:

Result 1: *The optimal fee to charge is positive but smaller than that suggested by the model with risk-neutral buyers, which tends to under-predict the profitability of NYOP selling.*

To understand why the risk-neutral model tends to under-predict profits and over-predict the level of optimal fee, we next decompose the profits into entry and bidding.

Experiment 1 Results: Entry and Bidding

Table 1 lists the percentage of subjects who enter under the different valuation-fee conditions. The shaded cells delineate the conditions under which risk-neutral buyers should not enter. Clearly, one of the reasons the risk-neutral model tends to under-predict profits is that it under-predicts entry by high-valuation buyers facing relatively high bidding fees. The discrepancy is the starkest at $f=18$, whereby the risk-neutral model predicts no entry, but about half our subjects enter when their valuations are relatively high.

The first two columns of Table 1 also provide evidence against the fee-aversion hypothesis: when the fee increases from 0 to 1, entry does decline, but only in the low-valuation conditions ($v \leq 20$). By contrast, fee aversion would predict a decline in entry irrespective of valuation. Not only does fee aversion

⁵ We consider both a two-sample t-test with one observation defined as the expected profit from one subject resulting in $N=91$, and a t-test comparing the 91 within-subject differences in expected profit to zero. The test using within-subject differences is definitive because fees vary within subject in our design. The two-sample test incorrectly assumes independence between the two samples, but facilitates visualization using error bars (e.g., in Figure 1) and ends up being more conservative because the bids are positively correlated within subject.

not seem to operate, but the impact of a very small fee is also positive in terms of seller profit, rising significantly from 7.2 to 8.4 ($p < 0.01$ regardless of how the relevant t-test is set up).⁶

Table 1: Proportion Of Subjects Who Enter The NYOP Store

valuation	bidding fee				
	0	1	6	12	18
5	91%	31%	0%	0%	0%
20	96%	84%	16%	5%	4%
35	97%	97%	53%	19%	14%
50	100%	98%	84%	43%	25%
65	99%	100%	96%	66%	48%
80	95%	95%	84%	63%	45%
95	85%	85%	75%	58%	48%

Figure 2 plots the submitted bids when the fee is zero and almost all buyers enter the NYOP store (and hence we are the least concerned about selection into the observed sample). It is immediate that at all valuation levels other than the lowest one, bids exceed the risk-neutral prediction. The difference is not only significant, but also large in that the entire inter-quartile range lies above the predicted level. We summarize our comparisons of the observed entry decisions and bids with the risk-neutral prediction in our next result:

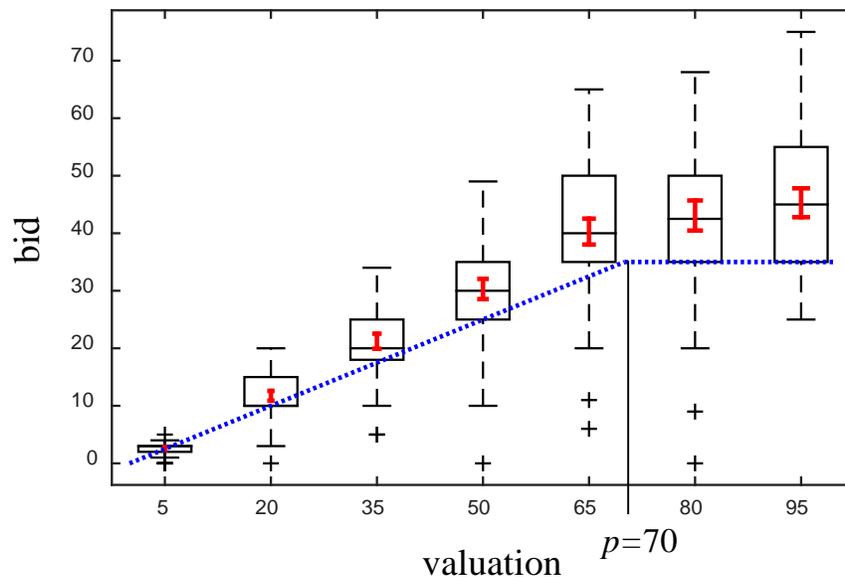
Result 2: *We do not find evidence of fee aversion. Buyers enter more often and submit higher bids upon entry than the risk-neutral model predicts.*

The higher-than-predicted profits summarized in Result 1 are thus not merely an outcome of excessive entry, because observed bids also exceed the risk-neutral benchmark. This finding presents us with a puzzle because existing expected utility models consistent with higher (than risk-neutral) bids imply we should see less rather than more entry: in standard expected utility models, risk aversion is well known to increase bids in 1PSB auctions, both in theory (Riley and Samuelson 1981) and in the laboratory (Cox,

Smith, and Walker 1988). Risk-averse bidders bid more because they experience diminishing marginal utility in surplus, and so (compared to risk-neutral bidders) they prefer increased chances of winning associated with higher bids. Therefore, the bidding behavior documented by Figure 2 is consistent with most of our subjects being risk averse. However, risk-averse buyers should enter in fewer (v, f) conditions than risk-neutral buyers, so the entry behavior documented by Table 1 is not consistent with risk aversion. In other words, our subjects enter as if they were risk seeking, but they bid as if they were risk averse.

We conjecture that the above puzzle of excessive entry by seemingly risk-averse bidders may arise from the cognitive difficulty of solving the optimal bidding problem. To test this conjecture, we designed Experiment 2 to explicitly help some of the subjects with assessing the payoff consequences of their potential bids. We now turn to the details of Experiment 2.

Figure 2: Observed Bids When Bidding Fee Is Zero



Note to Figure: The boxplots illustrate the distribution of bids at each valuation level. The thicker (red) error bars represent 95% confidence intervals with one observation defined as the expected profit from one subject, so $N=91$. The dotted (blue) line shows the optimal bidding function by risk-neutral buyers. The case of $v=5$ is difficult to discern from the figure—the risk-neutral prediction is 2.5, and the 95% confidence interval is [2.35, 2.85].

EXPERIMENT 2: EFFECT OF DECISION AID AND EXPERIENCE

Finding one's optimal NYOP bid involves solving the tradeoff between the probability of acceptance (increasing with the bid amount) and the utility of the monetary payoff (declining with the bid amount), with the utility of the payoff evaluated relative to the (dis)utility of paying a bidding fee and not getting anything back. In Experiment 2, we examine whether the excessive entry observed in Experiment 1 was due to the cognitive difficulty of solving the optimal bidding problem, by using an intervention that allows buyers to anticipate the consequences of submitting a particular bid amount. Specifically, for some of their bidding decisions, we provide subjects with a decision aid that, for any candidate bid amount of their choice, informs them about the two aspects of the tradeoff: (1) the probability of acceptance and (2) the contingent monetary payoff if the bid is accepted. Subjects can use the decision aid as much as they want before they finalize their bidding decision. Upon entering a candidate bid amount but before being able to submit it, subjects were required to click a button that activated the decision aid, which instantaneously displayed both the probability of acceptance and the contingent monetary payoff for that bid amount (see Figure A2 in the Appendix for a sample screenshot of the bidding interface with the decision aid).

To measure the effect of our decision aid both between subjects and within subject, we employ a balanced crossover design whereby half the subjects start with the decision aid and end without it, and the other half, vice versa. Each subject thus experiences two blocks of the same 25 decisions, with each block corresponding to a $5(\text{fee}) \times 5(\text{valuation})$ within-subject design analogous to Experiment 1 whereby each subject experiences all possible combinations of valuations and fees in random order. Half the subjects have the decision aid in their first block, and the other half have the decision aid in their second block. The only difference relative to Experiment 1 is that we simplify the situation by only using valuations below the posted-market price, resulting in only five valuation levels {5, 20, 35, 50, 65}.

The data collection was analogous to that in Experiment 1, with four sessions of 24 subjects. Due

to the already large number of tasks per subject, we omitted the “training” tasks in Experiment 2. To nevertheless remain within the spirit of eliminating subjects who persistently bid more than their valuations, we eliminated four subjects who bid more than their valuation more than once. All the analyses that follow are based on the remaining 92 subjects (46 in each cross-over condition).

Effect of Decision Aid: Model-free Evidence

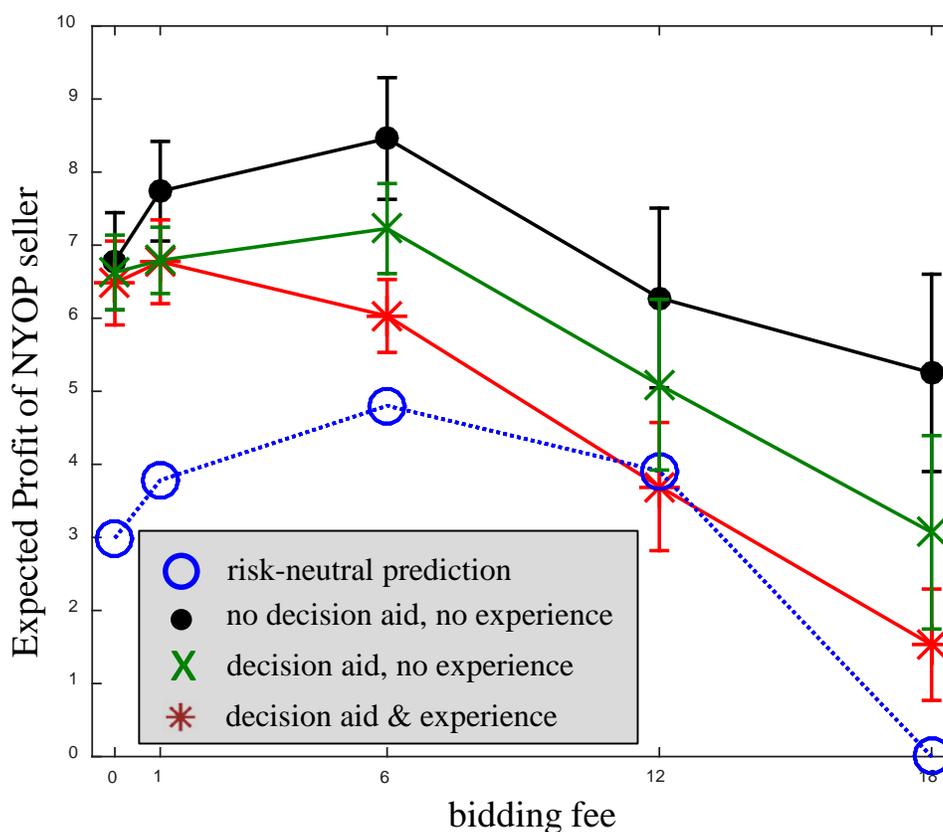
Comparing the first-block behavior of subjects with the decision in aid with subjects without it yields a clean measurement of the decision aid’s effect between subjects. Comparing the second-block behavior of subjects who start without the decision aid with their first-block behavior yields a within-subject measurement of the combination of the decision aid and experience. We first consider the effect of the decision aid (combined with experience in the within-subject case) on the seller’s profit, and then decompose the effect into an effect on entry and an effect on bidding given entry.

Figure 3 shows the decision aid reduces seller profits for all positive fees, and the decision aid combined with experience reduces it even more. For every positive fee level, all but one of the profit reductions relative to baseline are significant at the 5% level when we conservatively consider each subject’s average (over valuations) contribution as a single observation. The only exception is the between-subject difference at $f=12$. In contrast to the positive fees, neither the decision aid nor its combination with experience have a significant impact on profits when the bidding fee is zero. This finding suggests the decision aid and experience influence the subjects’ entry decisions but have little effect on their associated bidding strategies. In addition to comparing between the different conditions in our study, Figure 3 also clearly shows the main finding of Experiment 1 replicates in Experiment 2: for most fee levels, our subjects are more profitable to the seller than a risk-neutral theory would suggest.

Another notable feature of Figure 3 is the difference between the shapes of the three profit functions: in the baseline condition without the decision aid or experience, the seller profits significantly

from charging the empirically optimal bidding fee of 6: relative to zero fee, profits rise by 25%, $p < 0.01$. However, the 9% increase from the same strategy when the subjects are given the decision aid is not only much smaller, but also not statistically significant (in the relevant t -test, $p = 0.15$). Finally, providing the subjects with the decision aid and experience reduces the optimal fee level to 1, but the profit increase from zero fee is again not significant.

Figure 3: Expected Profit Of The NYOP Seller, By Condition



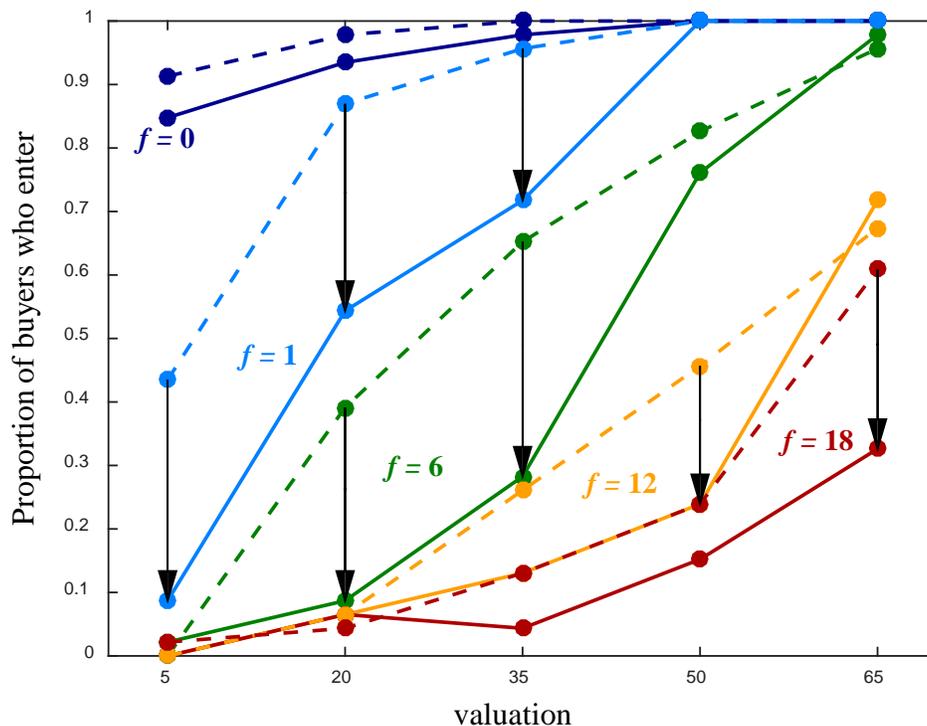
Note to Figure: The solid lines connect the observed expected profits, displayed as error bars. Each error bar represents the 95% confidence interval with one observation defined as the expected profit from one subject, so $N=46$. Solid (**black**) dots indicate the baseline condition without experience (i.e., in the first block) or decision aid. The (**green**) X markers indicate the condition with the decision aid but without experience. The (**red**) star markers indicate the condition with experience (i.e., in the second block) and with the decision aid. The dotted (**blue**) line connects the expected profits predicted by risk-neutral buyers.

We summarize the takeaway from Figure 3 in our next finding:

Result 3: *The decision aid reduces the expected profit of the NYOP seller who charges a positive fee, and additional buyer experience reduces the profits further. However, the decision aid has no effect on the profitability of a seller who charges zero fee. We only find model-free experimental evidence of the profitability of bidding fees when the buyers have no experience and no access to the decision aid.*

Having analyzed the main dependent value of managerial interest—the expected seller profit—we decompose it to its behavioral antecedents: entry and bidding. Figure 4 shows the between-subject effect of the decision aid on entry.

Figure 4: Effect Of Decision Aid On Entry, Between Subjects Without Experience



Note to Figure: Only first-block observations included. The solid lines connect the observed entry probabilities when the buyers have the decision aid ($N=46$), the dashed lines do the same thing for buyers who do not have the decision aid ($N=46$). Arrows indicate significant differences at the 5% level using the two-proportion z -test. All of them except ($v=50, f=12$) are also significant at the 1% level.

The decision aid reduces entry, but it does not have a simple main effect. Instead, the decision aid reduces entry along the diagonal of the (v,f) space—precisely in the conditions where the expected net utility of entry is likely near zero for most subjects. One way to interpret the effect pattern in Figure 4 is that the decision aid makes buyers more conservative whenever they are presented with a positive fee *and* their valuation is low enough that the fee looms relatively large. In the structural model described in the next section, we will interpret the effect of the decision aid on entry as an increase in the disutility of losing the fee amount. We do not report the within-subject effect of the decision aid and experience, because it is qualitatively very similar to the between-subject effect, also occurring along the diagonal of the (v,f) space. If anything, the combined effect of the decision aid and experience is stronger than the effect of the decision aid alone. Please see Figure A4 in the Appendix for a plot of the within-subject effect analogous to Figure 4.

In contrast to its large effect on buyer entry, the effect of the decision aid on bid magnitudes is negligible. Comparing bids at zero fee (when most buyers enter) in the first block, buyers bid about one unit of the experimental currency lower when they have the decision aid, but the difference is only significant for $v=5$. We conclude the decision aid affects mostly entry into the NYOP store.

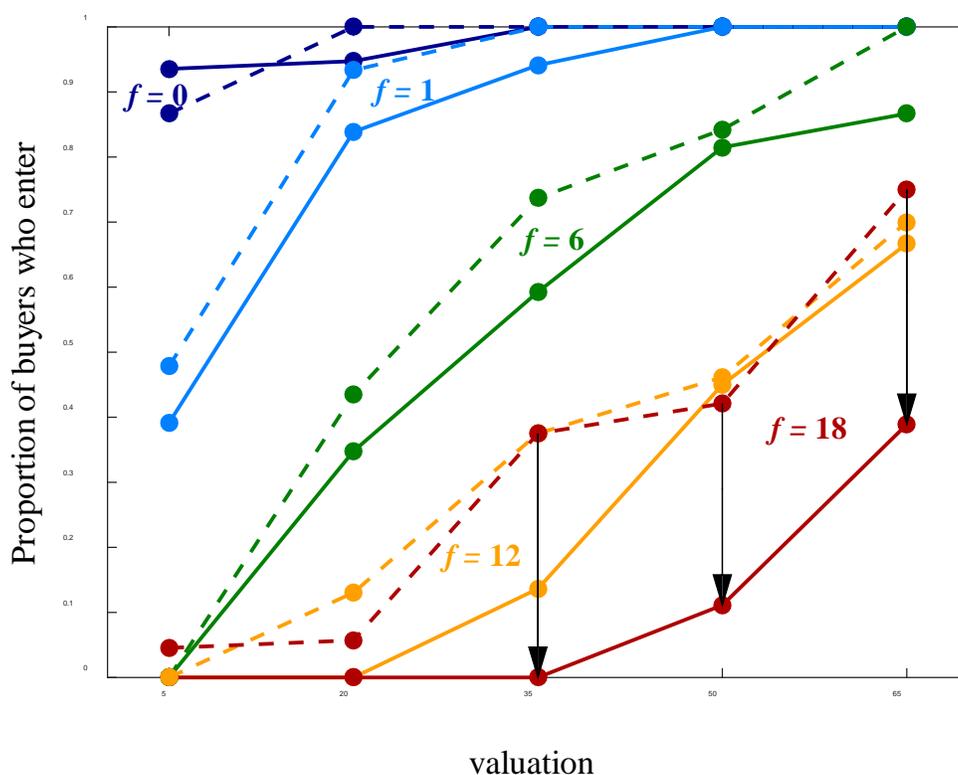
Effect of Experience: Model-free Evidence

Because each subject experiences the different (v,f) conditions in random order, we can isolate the effect of experience between subjects by focusing only on the subjects without the decision aid in their first block, and comparing the entry behavior of those who experience a given (v,f) condition early in the block with those who experience it later. Figure 5 plots the result of this analysis analogously to Figure 4.

The analysis presented in Figure 5 has less statistical power than that in Figure 4, but we still find a large-enough effect to reach significance at the 1% level: experience alone teaches subjects not to enter when the fee is large (all significant differences involve $f=18$, and the two sizeable differences for $f=12$

are marginally significant with $p=0.06$). Unlike the decision aid, experience alone does not reduce marginal entry for moderate fee levels: Figure 5 does not exhibit Figure 4's marked spread between the solid and dashed lines corresponding to fee levels of 1 and 6. One way to speculatively interpret the different patterns of entry reduction from the decision aid and experience is to conjecture that subjects treat the two dimensions of the (v,f) space independently when learning from experience alone: they can learn high fees are undesirable and high valuations are desirable. By contrast, by guiding the subjects to the tradeoff between winning $v-f$ and losing f , the decision aid helps them focus on the normatively correct combination of the two dimensions, that is, on the marginal-entry diagonal of the (v,f) space.

Figure 5: Effect Of Experience On Entry, Between Subjects Without The Decision Aid



Note to Figure: Only the first-block behavior of subjects who start without the decision aid ($N=46$) is considered. The solid lines connect the observed entry probabilities during periods 1-12, the dashed lines do the same thing during periods 13-25. Arrows indicate significant differences at the 5% level using the two-proportion z -test; all are also significant at the 1% level.

WHO ENTERS TOO OFTEN? A POST-HOC ANALYSIS

Table 1 (Experiment 1) and the dashed lines in Figure 4 (Experiment 2) establish that at least some subjects enter too often; that is, they enter when a risk-neutral agent should not (the shaded area in Figure 1). In this section, we examine individual differences that correlate with this type of excessive entry. Between Studies 1 and 2, we have 185 useful⁶ subjects for whom we know age, gender, frequency of prior participation in lab experiments, and a battery of psychological scales related to their risk and bidding-fee attitudes. Table 2 reports a linear regression of each subject's observed probability of excess entry⁷ on the available individual differences, as well as on controls for the two different studies and the presence of the decision aid.

We find two individual differences: gender and the Holt-Laury risk-aversion. The gender effect is surprising given Croson and Gneezy (2009), who survey the experimental literature and find women are more risk averse, and hence tend to avoid risky situations more than men. By contrast, we find women are 7% more likely than men to enter “too often.” This effect is large: Table 2 implies a young man in the baseline condition who is not a frequent subject and who is average in terms of all our psychological scales enters excessively about 15% of the time, whereas an otherwise identical female enters excessively about 22% of the time—an almost 50% increase. One way to reconcile our finding with that of Croson and Gneezy (2009) is that we control for both revealed and stated risk aversion, and the gender effect is robust to the exclusion of all these scale variables. In other words, in contrast to much of the research reviewed by Croson and Gneezy (2009), gender is not merely a proxy for a particular risk preference in our analysis.

⁶ For the purposes of this section, we use Experiment 2's definition of a seemingly irrational subject, eliminating subjects who bid above their valuation more than once. Hence, we have 93 “useful” subjects from Experiment 1.

⁷ For subjects in Experiment 2, we only consider behavior in block 1. For subjects in Experiment 1, the dependent variable considers only observations with valuations below the posted price, matching Experiment 2's (v,f) design. Regardless of the study, the probability of “excessive entry” is the probability of entering in the 13 cells of the (v,f) design that are shaded in Table 1 and involve $v < 70$.

Table 2: Who Enters Too Much? Linear Regression Of Individual Probability Of Excess Entry

Variable	Estimate	SE	<i>t</i> -stat	<i>p</i> -value	Estimate	SE	<i>t</i> -stat	<i>p</i> -value
Intercept	15.24%	2.59%	5.90	<0.001	16.69%	2.56%	6.51	<0.001
Decision aid	-11.80%	3.36%	-3.51	0.001	-13.71%	3.43%	-4.00	<0.001
Experiment 2	3.74%	3.00%	1.24	0.215	5.36%	3.00%	1.78	0.076
Age > 25	2.17%	2.81%	0.77	0.441	1.47%	2.89%	0.51	0.611
Female	7.27%	2.41%	3.02	0.003	6.53%	2.43%	2.69	0.008
Frequent subject	1.01%	2.62%	0.38	0.702	-0.95%	2.64%	-0.36	0.719
Risk aversion	-3.53%	1.28%	-2.76	0.006				
Loss aversion	-1.75%	1.33%	-1.31	0.191				
Subjective risktaker	1.42%	1.24%	1.14	0.254				
Fee aversion	0.52%	1.20%	0.44	0.663				

Note: Risk aversion is measured as in Holt-Laury (2002), loss aversion as in Gächter et al (2010), subjective risk taking as in Dohmen et al. (2012), and fee aversion is an agreement with “I want to pay for the actual product only; I refuse to pay a bidding fee in general” on a 7-point scale. All four scales are standardized within the entire population, and hence measured in population standard deviations. All other variables are dummies. Number of observations = 185 (93 subjects from Experiment 1, 92 from Experiment 2). $R^2=0.196$.

In addition to the above analysis, an investigation of how individual differences moderate the effect of decision aid in Experiment 2 would be interesting. For example, one could ask what kind of subjects who start without the decision aid tend to reduce their excessive entry more when they are given the decision aid in their second block of the session. Unfortunately, the small sample size ($N=46$ subjects) prevents reliable interpretation of such an analysis unless the effects are very large. We have regressed the within-subject entry reduction on the same variables as shown in Table 2, and we did not find any effects large enough to reach significance.

MODEL OF NYOP BIDDING WITH A SEMIPARAMETRIC UTILITY FUNCTION

Participating in NYOP bidding is risky for buyers because their offer may not be accepted. We assume that when considering entry into the NYOP store with a bidding fee of f , the buyer with valuation v considers the two potential outcomes of offering a bid b as follows. First, if b is accepted, she combines the relevant experimental currency amounts into a net surplus $v-b-f$, and represents her utility of this potential surplus as $u(v-b-f)$. Second, if her bid is not accepted, she represents her utility of spending f without getting anything back as $u(-f)$. To decide how much to bid and whether to enter, she solves

$$\max \left\{ \underbrace{u(0)}_{V_{\text{not enter}}}, \underbrace{\max_b W\left(\frac{b}{p}\right) u(v-b-f) + W\left(1-\frac{b}{p}\right) u(-f)}_{V_{\text{enter}}} \right\} \quad (2)$$

where W is a probability-weighting function and V represents the buyer's choice-relevant values of entering or not. Without loss of generality, we can set the intercept and scale of the utility function as $u(0)=0$ and $u(1)=1$. Given this normalization, the general model in equation (2) nests both classical Expected Utility Theory (with W an identity function and u a non-linear function—concave for risk-averse bidders and convex for risk-seeking bidders⁸) and Prospect Theory (with a non-linear W and a u “value function” concave in gains and convex in losses⁹). Our goal is to specify a parsimonious yet flexible model of W and u , and let the data determine which of the above two major theories of risk preferences fits our data better, at the individual level.

⁸ Given the normalization, u represents the *change* in the underlying utility from some initial wealth level w_0 to a new wealth level, either $w_0 + v-b-f$ in the case of an accepted bid, or $w_0 -f$ in the case of a rejected bid. Since we do not observe the initial wealth level, estimation at the individual level is imperative to preserve the Expected Utility interpretation of equation (2).

⁹ Prospect Theory evaluates uncertain outcomes as gains or losses relative to a reference point. In our NYOP setting, it is natural to let the income state in the beginning of the game be the reference point as implicitly assumed by the $u(0)=0$ normalization: if the buyer does not pay any bidding fees and does not buy any products, she remains at this reference point. When she buys a product, she experiences a net gain of $v-b-f$. When her bid is rejected, the buyer experiences a loss of the bidding fee f . Also note that compared to the fully general model of Tversky and Kahneman (1992), equation (2) assumes the probability weighting function is the same for gains and losses in order to reduce the number of parameters required.

Both Expected Utility Theory and Prospect Theory postulate a curved utility function in the domain of gains. Let the shape of u in the domain of gains be captured by $u(s) = s^r$. To interpret r , consider the case of zero fee and no probability weighting: then buyers are risk averse in the entire gain domain, and $1-r$ is the coefficient of relative risk aversion, $-s \frac{\partial^2 u}{\partial s^2} / \frac{\partial u}{\partial s} = 1-r$. In other words, lower $r < 1$ implies greater risk aversion, $r=1$ implies risk neutrality, and $r > 1$ implies risk-seeking preferences.

Expected Utility Theory and Prospect Theory differ in two respects: the shape of W and the curvature of u in the loss domain. To allow a prospect-theoretic W while nesting the identity function, we follow Prelec's (1989) axiomatically derived single-parameter form:

$$W(z) = \exp\left[-(-\ln(z))^\alpha\right] \quad (3)$$

where $\alpha = 1$ corresponds to an identity function, and $0 < \alpha < 1$ corresponds to a prospect-theoretic function that overweights small probabilities and underweights large ones. Because our data only include the four possible losses corresponding to the four levels of the bidding fee, we can model the shape of u in the loss domain nonparametrically as

$$u(-f) = -\sum_{i \leq f} \exp(\delta_i) \quad (4)$$

where $\delta \equiv \{\delta_1, \delta_6, \delta_{12}, \delta_{18}\}$ are four scalars, and the cumulative sum of exponentials ensures u is increasing.

In summary, we need to estimate the six structural parameters $\{r, \alpha, \delta\}$ for each subject in each decision-aid treatment.

To complete the specification of the empirical model, we now describe the econometric error structure that allows observed bids and entry decisions to randomly deviate from theoretical predictions.

Let $b(v, f | r, \alpha, \delta)$ be the solution to the inner maximization problem in equation (2), and denote the buyer's n -th observed bid bid_n . We model the deviation between bid_n and $b(v_n, f_n | r, \alpha, \delta)$ as

$$bid_n = b(v_n, f_n | \theta) \exp(\varepsilon_n) \quad (5)$$

where ε_n is distributed normally with a mean of 0 and variance σ^2 , truncated such that

$bid_n < v_n \Leftrightarrow \varepsilon_n < \log\left(\frac{v_n}{b(v_n, f_n | r, \alpha, \delta)}\right)$. In words, we introduce a positive multiple $\exp(\varepsilon_n)$ to ensure

the observed bids are positive, and we restrict the observed bid to be below valuation as implied by our definition of basic rationality in cleaning the data.

To model the entry decision, let $V_{enter}(v_n, f_n | r, \alpha, \delta)$ be the choice-relevant value of entering defined in equation (2). We model the deviation between a subject's n -th entry choice and the predicted behavior, using an additive normal error term η_n with mean zero and variance τ^2 :

$$\Pr_{enter}(v_n, f_n | r, \alpha, \delta, \tau) = \Pr(V_{enter}(v_n, f_n | r, \alpha, \delta) - \eta_n > 0) = \Phi\left(\frac{V_{enter}(v_n, f_n | r, \alpha, \delta)}{\tau}\right) \quad (6)$$

where Φ is the standard normal cumulative distribution function.

Model: Parameter Identification and Estimation

The entry and bidding behavior of each experiment participant is thus captured by eight parameters: $\{r, \alpha, \delta, \sigma, \tau\}$. We now discuss how the entry and bidding data from our experiment identify these parameters. The curvature parameter in the gain domain r is identified by the magnitude of bids given valuations (larger bids imply smaller r) as well as by the entry decisions (more risk-averse buyers enter less). Consider again the special case of zero fee and no probability weighting. It is well known¹⁰

that the optimal bidding function in this special case is $b(v, 0 | r, 1, \delta) = \frac{v}{1+r}$, so the ratio of bids to

¹⁰ In standard expected-utility models, it is well known that risk aversion increases bids in 1PSB auctions, both in theory (Riley and Samuelson 1981) and in the laboratory (Cox, Smith, and Walker 1988). Our model's concavity of the utility function in the gains domain leads to the same result. Bidders bid more because they experience diminishing marginal utility in surplus; therefore (compared to risk-neutral bidders), they prefer increased chances of winning associated with higher bids.

valuations is negatively correlated with the relative risk-aversion parameter $1-r = 2 - \frac{v}{b(v,0|r,1,\delta)}$.

The δ parameters capturing the shape of u in the loss domain are also identified both from bids (for a fixed v and f , a larger $\sum_{i \leq f} \exp(\delta_i)$ implies a larger bid) and from entry decisions (smaller $\sum_{i \leq f} \exp(\delta_i)$ implied more entry when the fee is f). Note that δ_f is identified only if the buyer entered at least once with a fee weakly exceeding f : for example, the behavior of a buyer who never paid a fee of 18 is perfectly captured by any sufficiently large δ_{18} . To restrict the δ vector to the empirically identified range, we find the highest fee each buyer paid, and only estimate δ 's up to that level while ignoring the consistent non-entry for higher fee levels. For subsequent simulation purposes, we set each person's non-identified δ 's to very large positive numbers, ensuring the person would never enter at a fee level he or she has never paid.

The probability-weighting parameter α is identified from curvature of the empirical bidding function. The bidding function under $\alpha=1$ is either a straight line (when $f=0$) or a slightly concave curve (when $f>0$), whereas the bidding function is convex for low valuations when $\alpha<1$: intuitively, overestimation of small probabilities of winning reduces the optimal bids for small valuations. Finally, the two standard deviations σ and τ of the errors are obviously identified by departures from predictions.

To estimate the $\{r, \alpha, \delta, \sigma, \tau\}$, we use all 25 observations of each participant in each condition to maximize the likelihood implied by equations (5) and (6) detailed in equations (A1) and (A2) in the Appendix. The computation is straightforward except for the need to solve the bidding problem in equation (2) at every step, which we accomplish by searching for the optimal bid on a fine grid. We conduct a constrained search for the maximum-likelihood estimates, with δ unconstrained, σ and τ positive, $r \in [0.1, 1.5]$ to allow both extreme risk aversion and moderate risk-seeking preferences, and $\alpha \in [0.1, 1]$ to enforce overestimation of small probabilities and underestimation of large ones.

Model: Parameter Estimates

We estimated two specifications of the model—with and without probability weighting—for each subject-condition combination. The likelihood ratio test indicated overwhelming support for the specification with probability weighting (i.e., with the α parameter estimated rather than restricted to 1).¹¹ Therefore, we restrict attention to the model with probability weighting from this point onward.

The left side of Table 3 shows the mean, median, and standard deviation of each model parameter (with the δ parameters exponentiated before averaging) in the first 25 rounds, broken down by the availability of the decision aid. The last four columns then compare the two groups of subjects to assess the between-subjects effect of the decision aid. Figure 6 shows the utility functions implied by the $\{r, \alpha, \delta\}$ parameter estimates.

The main takeaway from Table 3 and Figure 6 is that Prospect Theory captures the preferences of the *average* bidder better than Expected Utility Theory: the average utility function is convex in the loss domain in both conditions. The convexity is obvious in Figure 6, and clearly holds for all levels of the highest fee ever paid. In Table 3, convexity around zero is evidenced by $\exp(\delta_1) < 1 = u(1)$, and convexity around a small loss of 1 is evidenced by $5\exp(\delta_1) > \exp(\delta_6)$ —inequalities that hold for the average parameters robustly in both decision-aid conditions. Table 4 breaks down the curvature measurement by condition and maximum fee paid, showing people who pay higher fees tend to have more convex utility functions in the loss domain. Although the average utility function is convex in the loss domain, we also find substantial population heterogeneity: Table 4 shows only about half our subjects exhibit convexity at

¹¹ Because we estimate separate parameters for each subject and condition, a population-wide restriction of one of the parameters to a constant (e.g., $\alpha=1$) can be tested with a likelihood ratio test with the same degrees of freedom as the number of subjects, and the log likelihood added up across subjects. We have 46 subjects in each condition, the $p=0.01$ critical value of χ^2_{46} is 71.2, and the three statistics we get are 144 (no decision aid, no experience), 172 (decision aid, no experience), and 76 (decision aid & experience). Therefore, we reject the model with $\alpha=1$ at the 0.01 level throughout.

both 0 and -1, only one person exhibits concavity at both points, and the remainder exhibit a mix of curvatures. Therefore, we can rule out the concavity-in-losses of expected utility, but the data do not suggest a convex prospect-theoretic model can be imposed on all subjects.

Table 3: Parameter Estimates, By Condition, Between Subjects

Parameter	No decision aid			With decision aid			Difference: effect of aid			
	mean	median	population. std. dev	mean	median	population. std. dev	mean	Mann-Whitney U	median	population. std. dev
r	0.27	0.22	0.19	0.29	0.22	0.22	0.02	0.55	0.00	0.03
σ	0.24	0.17	0.21	0.23	0.20	0.15	-0.01	0.57	0.02	-0.06
τ	0.25	0.14	0.43	0.12	0.03	0.16	-0.12	0.17	-0.11	-0.26
$\exp(\delta_1)$	0.30	0.06	0.43	0.68	0.57	0.65	0.38	0.00	0.51	0.22
$\exp(\delta_6)$	0.75	0.24	1.60	1.27	0.43	2.80	0.52	0.26	0.19	1.19
$\exp(\delta_{12})$	0.69	0.56	0.80	0.46	0.14	0.65	-0.22	0.44	-0.42	-0.15
$\exp(\delta_{18})$	0.22	0.00	0.53	0.28	0.00	0.59	0.06	0.86	0.00	0.06
α	0.68	0.62	0.30	0.73	0.96	0.32	0.05	0.45	0.35	0.02
MAD entry	9%			5%			-3%			
MAD profit	2.00			1.42			-0.57			
MAD profit/AVG profit	29%			25%			-4%			

In addition to the convexity of the utility function in the loss domain, Prospect Theory also predicts loss aversion near the reference point whereby people feel the disutility of a small loss more than the utility from an equal-sized gain. Table 4 shows both the incidence and average magnitude of loss-aversion ratio, by condition and maximum fee paid. We clearly do not find loss aversion: only 29% (with decision aid) and 7% (without decision aid) of subjects exhibit loss aversion at a surplus of 1, with the incidence rising slightly to 40% and 17% at the surplus of 6. The reason we do not find loss aversion is the over-entry relative to the risk-neutral benchmark: the model rationalizes a subject's willingness to pay a high fee by reducing the disutility of that fee.

Table 4: Curvature And Loss Aversion Of The Utility Function, By Condition (Between Subjects)

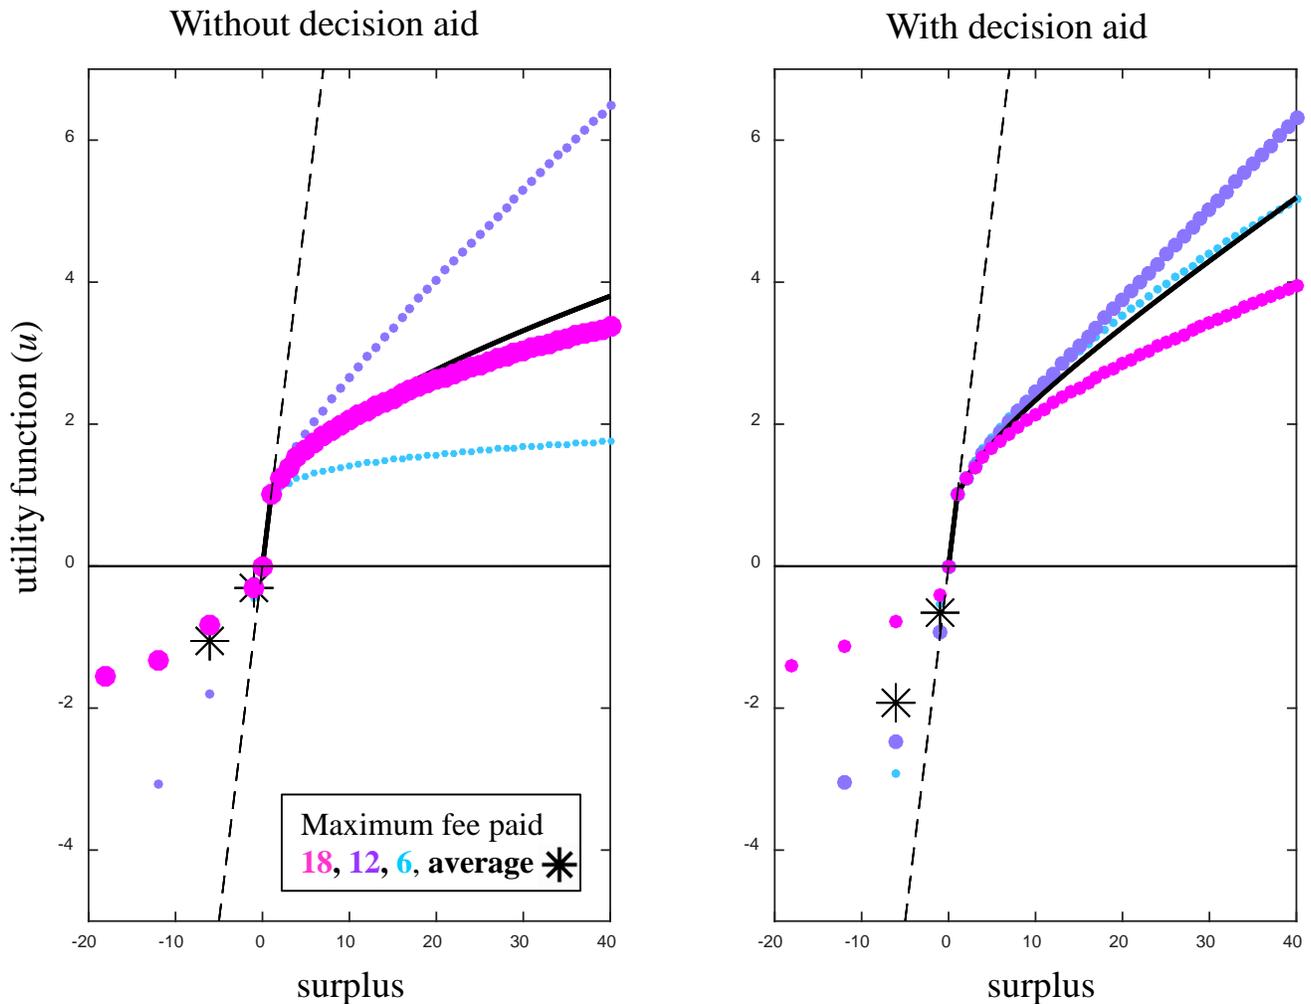
decision aid	Max fee paid	# subjects	convex at 0 and -1	concave at 0 and -1	concave at 0 = loss-averse at 1	loss-aversion ratio at 1	loss-averse at 6	loss-aversion ratio at 6
no	≥ 6	46	52%	2%	7%	30%	17%	54%
yes	≥ 6	45	42%	0%	29%	66%	40%	97%
no	6	7	43%	0%	14%	46%	29%	62%
	12	10	40%	10%	10%	20%	20%	62%
	18	29	59%	0%	3%	30%	14%	49%
yes	6	9	22%	0%	33%	53%	67%	150%
	12	19	21%	0%	47%	94%	53%	115%
	18	17	76%	0%	6%	41%	12%	49%

Note to Table: Incidence of loss aversion at a given surplus level s is measured by the proportion of subjects for whom $|u(-s)| > u(s)$. Magnitude of loss aversion at a given surplus level s is measured by the loss-aversion ratio $|u(-s)|/u(s)$. We have only 45 useful subjects in the decision-aid condition because one subject in that condition never paid any fee of 6 or more.

Model: Effect of Decision Aid on Parameter Estimates, between Subjects

We now turn to the effect of our decision-aid treatment. In comparing the model parameters (Table 3), we worry that relying on normal asymptotics of the maximum likelihood estimation (MLE) would be too heroic in our small-sample setting (with 15-25 observations for estimating 6-8 parameters), so we use the nonparametric Mann-Whitney-Wilcoxon test for comparing the point estimates between the group of subjects that was given the decision aid and the group that was not. One clear difference arises: the decision aid increases the disutility of a small fee. In Figure 6, we see the average utility in the loss domain (denoted by black star symbols) is indeed steeper with the decision aid. The increased slope of the utility curve is driven by significantly fewer people paying high fees with the decision aid (see previous section on model-free results). Although insignificant given our testing strategy, note the decision aid makes the entry model fit better—with the average τ parameter and MAD of about half the magnitude with the decision aid.

Figure 6: Utility Functions Without Prior Experience, Averaged By Highest Fee Paid



Note: Color indicates the maximum fee paid, with the overall average in black. The size of the points indicates the number of subjects in each maximum-fee group.

Model: Effect of decision aid and experience on parameter estimates, between subjects

Table 5 and Figure 7 are the within-subject analogues of Table 3 and Figure 6, respectively. They focus on the subjects who started the experiment without the decision aid, and compare their parameter estimates and utilities in the first block of 25 rounds with those in the second block of 25 rounds in which these subjects were provided with the decision aid.¹² Of course, our within-subject design confounds experience

¹² Recall our estimation strategy treats a person in the first block completely separately from the same person's behavior in the second block, so our analysis can remain the same as in the between-subjects analysis.

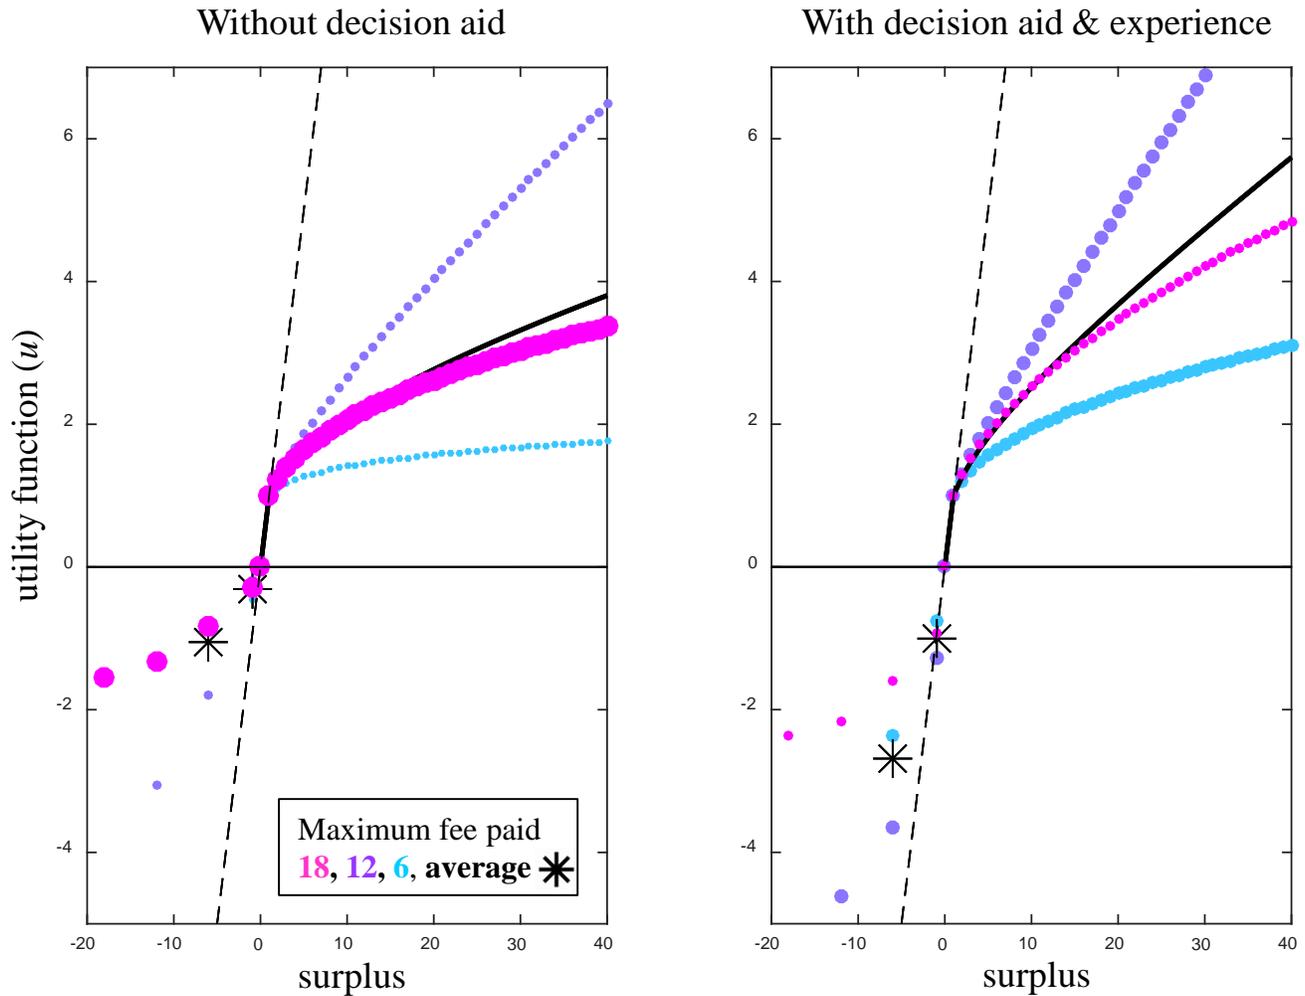
and the availability of the decision aid, so the within-person difference in behavior between the two blocks cannot be attributed solely to the presence of the decision aid.

Table 5: Parameter Estimates, By Condition, Within Subject

Parameter	No decision aid			With decision aid & experience			Difference			
	mean	median	population. std. dev	mean	median	population. std. dev	mean	Signed rank test p	median	population. std. dev
r	0.27	0.22	0.19	0.32	0.27	0.23	0.05	0.09	0.05	0.04
σ	0.24	0.17	0.21	0.20	0.13	0.22	-0.04	0.07	-0.04	0.01
τ	0.25	0.14	0.43	0.34	0.05	1.16	0.09	0.99	-0.08	0.73
$\exp(\delta_1)$	0.30	0.06	0.43	1.01	0.80	1.24	0.71	0.00	0.74	0.81
$\exp(\delta_6)$	0.75	0.24	1.60	1.68	0.98	2.82	0.93	0.00	0.74	1.21
$\exp(\delta_{12})$	0.69	0.56	0.80	0.83	0.00	2.28	0.14	0.55	-0.56	1.48
$\exp(\delta_{18})$	0.22	0.00	0.53	0.19	0.00	0.34	-0.03	1.00	0.00	-0.19
α	0.68	0.62	0.30	0.85	1.00	0.27	0.17	0.02	0.38	-0.03
MAD entry	9%			5%			-3%			
MAD profit	2.00			1.05			-0.94			
MAD profit/AVG profit	29%			21%			-7%			

The general pattern of results is the same as in the between-subjects comparison, but all of the effects of the decision aid become stronger. The average utility function is significantly steeper not only at -1, but also at -6, and the model fits substantially better both in terms of entry and in terms of profit. Finally, the combination of decision aid and experience significantly reduces both the extent and the incidence of probability weighting: the average α increases from 0.68 to 0.85, and the proportion of people with an estimated $\alpha=1$ (i.e., the upper bound on α binding in the estimation procedure) increases from 41% to 67%.

Figure 7: Utility Functions Of The Same Group of Subjects, Averaged by Highest Fee Paid



Note: Color indicates the maximum fee paid, with the overall average in black. The size of the points indicates the number of subjects in each maximum-fee group.

Figure A3 in the Appendix shows the main construct of interest, namely, the expected profit of the NYOP seller implied by equation (1), under the three experimental conditions considered above. The main takeaway from Figure A3 is that each condition's model captures the expected profits well, with the fit improving as subjects receive decision aid and gain experience. Note the expected profit did not enter estimation, so the model's ability to predict it is not automatically assured. Finding the model predicts profits well gives credence to our counterfactual exercises discussed in the next section.

COUNTERFACTUAL SIMULATIONS

We conduct counterfactual simulations to answer the following questions:

- 1) What is the optimal fee an NYOP seller facing uniformly distributed buyers should charge and how much does the optimal fee improve profits relative to charging no fee?
- 2) What is the optimal minimum bid an NYOP seller facing uniformly distributed buyers should set and how much does the optimal minimum bid improve profits relative to setting the minimum at zero?

Simulation 1: Profitability of Bidding Fees

Model-free evidence suggests the profitability of bidding fees depends on the availability of the decision aid. When subjects have access to the decision aid, neither the fee of 1 nor the fee of 6 significantly outperform a fee of zero. But what about an intermediate fee level, for example, 3 or 4? Also, the model-free results are based on buyer valuations distributed uniformly on a small discrete set {5, 20, 35, 50, 65}—a simplification for experimental purposes with limited external validity. What if the valuations were distributed more continuously?

We can use the calibrated structural model to predict profitability of all fee levels between 0 and 18 while also allowing for an arbitrary distribution of buyer valuations. To predict buyer response to fees not included in the experiment, we interpolate linearly between the relevant δ parameters. To consider a more continuous distribution of valuations, we allow valuations to be uniformly distributed on [0,70].

Consider a particular buyer with a particular valuation facing a particular bidding fee. To estimate that subject's profitability to the seller charging that fee, we simulate 1,000 draws of the bidding error ε (taking care to truncate its distribution as specified in our econometric specification), generate the corresponding draws of bids using equation (5), apply equation (1) to each bid draw to average over the possible seller acceptance thresholds, and average over the draws. Given the resulting estimate of the expected profitability of the focal buyer entering the NYOP store, we then calculate the overall expected

profitability of the buyer by calculating the probability of entry given that buyer's estimated τ . In simulation jargon, we thus generate a plug-in estimate of the expected NYOP profit at the MLE parameter estimates. We perform the above for every possible valuation, and average to obtain the expected profitability of the focal buyer by fee level. Averaging over all buyers yields the expected profitability of the entire market by fee level.

We calculate such plug-in estimates separately for each experimental block, focusing on the 46 subjects who started without the decision aid (baseline condition) and then received the decision aid in the second block (decision aid & experience condition). We propose the subjects in the decision aid & experience condition represent the closest estimate we have of long-run behavior in a market that provides the decision aid.

Table 6: Optimal Selling Strategies Based On Counterfactual Simulations

Selling strategy	Bidding fee		Minimum bid	
	optimal level	expected profit	optimal level	expected profit
risk-neutral theory	6	4.92	47	5.11
model estimates: baseline without aid or experience	6	8.57	30	6.46
model estimates: aid & experience	2	6.86	28	6.36

Table 6 and Figure 8 summarize the results of both counterfactual simulations. The left side of Table 6 shows that both the level and expected profitability of the optimal bidding fee depends on the decision aid and experience of the buyers: when the buyers have no aid or experience, the optimal fee of 6 raises expected profits by 36% relative to a zero fee—much less than the 67% lift predicted by risk-neutral theory. Providing buyers with some experience and the decision aid reduces the optimal bidding fee to 2, and reduces its profit lift to only 9%. Therefore, the model-free results may have been too pessimistic regarding the optimality of fees by not including other fee levels between 1 and 6. We conclude that bidding fees can be profitable for NYOP sellers even in the long run and even when buyers have access to decision aids, but the optimal fee levels and their expected profit lift are much smaller than risk-neutral models would predict.

Simulation 2: Profitability of Minimum Bids

In a recent paper, Zeithammer (2015) argues an NYOP seller analogous to the one modeled here can profit from setting a minimum bid (equivalent to a public reserve price in an auction). Under the assumption of buyer risk neutrality, Zeithammer (2015) also shows the minimum-bid strategy approximates the optimal mechanism, and it is more profitable than the entry-fee strategy for a wide variety of distributions of buyer valuations and seller wholesale costs. In the uniform example of Table 5, the superiority of minimum bids over bidding fees manifests itself by the optimal minimum bid of 47 outperforming the optimal bidding fee of 6 by about 4%.

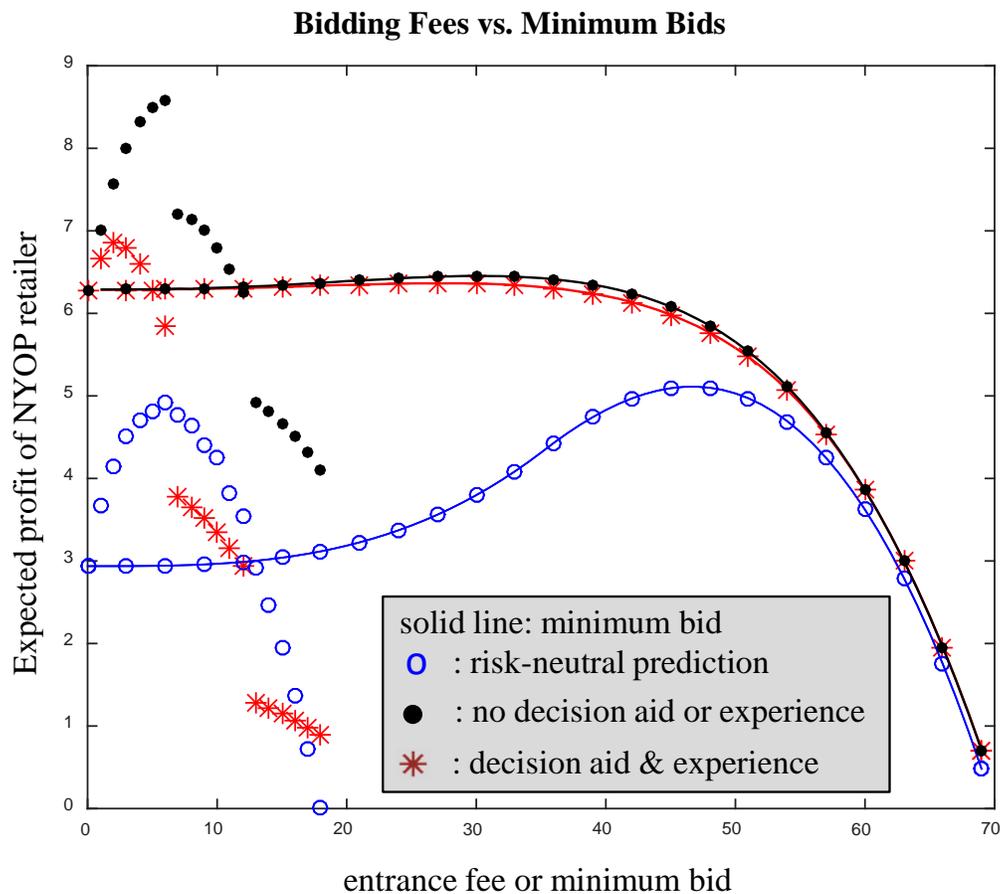
Given these theoretical predictions, simulating how our bidders would react to a minimum bid is interesting. To simulate their behavior under such circumstances, we assume bidders with valuations below the minimum bid do not enter, and bidders with valuations above the minimum bid submit either the minimum bid or their optimal bid suggested by equation (5), whichever is greater. Specifically, we simulate the econometric error terms from the truncated lognormal distribution as described above, and censor the draws below at the minimum bid.

After trying all possible minimum bids between 0 and 70, we find the optimal minimum bid is around 30, and it increases profits relative to no minimum bid by less than 2% regardless of bidder experience or the presence of a decision aid. Figure 8 reveals simulated profits are almost flat for small minimum bids, in contrast to the more pronounced peak of the risk-neutral model. The reason profits are both higher and flatter in the simulation than in the risk-neutral model is the relatively high level of bidding we observe. Recall the optimal risk-neutral strategy is to bid 50% of valuation, whereas our high-valuation (most profit-relevant) bidders bid about 75% of their valuations when the bidding fee is zero (as it is under the minimum-bid strategy).

Comparing the maximal achievable profitability of the two selling strategies is straightforward in

Figure 8, which plots the results of both simulations. Contrary to the risk-neutral predictions of Zeithammer (2015), the minimum-bid strategy does not outperform the entry-fee strategy in our simulations. Instead, the entry-fee strategy results in a higher lift relative to the common status-quo baseline. However, this superiority of bidding fees is fragile in that it depends on selecting the optimal fee carefully, and the profit improvement over minimum bids is sharply diminished by the presence of a decision aid and by buyer experience. In this sense, the minimum-bid strategy is more robust.

Figure 8: Impact Of Decision Aid And Experience On The Profitability Of Selling Strategies:



Note to Figure : Solid lines connect the profit estimates for different levels of the minimum bid; markers not connected by solid lines represent the profit estimates for different levels of bidding fee. Empty (blue) circles mark the risk-neutral prediction. Full (black) circles mark the prediction of the model calibrated on decisions by subjects without experience or the decision aid. (Red) stars mark the prediction of the model calibrated on those same subjects in the second block of their session, that is, when they have had experience with NYOP bidding and when they were given access to the decision aid.

GENERAL DISCUSSION

Suppose consumers bidding on an object have private valuations in mind. One of the core questions in the analysis of bidding data concerns the mapping between those secret valuations and the bids analysts observe. When bidder entry into the bidding market is costly, the mapping question is further complicated by some bidders deciding to abstain from bidding altogether, forcing analysts to predict the entry decisions as well as the subsequent bids. We answer this mapping question in the domain of NYOP selling by inducing valuations and entry costs in a laboratory setting, measuring entry and bidding decisions, and estimating an individual-level structural model with a flexible specification of utility. Unlike auctions, NYOP selling is a particularly clean setting for measuring underlying bidder preferences (i.e., the shape of the utility function) because each market interaction involves only one bidder, and so no equilibrium considerations arise—modeling the bidder’s best response to the seller’s bid-acceptance strategy is sufficient for analysis. As far as we know, this paper is the first to use bidding data to estimate a behaviorally realistic model of preferences that includes not only the Expected Utility Theory, but also Prospect Theory as special cases.

Our proposed model assumes the bidder considers the surplus levels associated with all three possible outcomes,¹³ maps the surplus levels into their associated utility levels, and takes an expectation, possibly using subjective probabilities in the sense of “probability weighting.” In contrast to Expected Utility Theory, we find evidence of reference dependence in that the estimated preferences tend to exhibit diminishing marginal utility in both the gain and the loss domain, where the reference point is defined as zero surplus. Although the estimated individual utility functions are almost all concave in the gains domain, only about half our subjects exhibit a consistent curvature in the loss domain—all of them convex.

¹³ The three possible outcomes of the bidder’s interaction with the market are (1) bid accepted, (2) bid rejected and outside posted price too high, and (3) bid rejected and object purchased for the outside posted price.

Thus, the preferences of about half our subjects can be captured by Prospect Theory, and half do not conform easily to any established curvature pattern. Subjects who conform to the prospect-theoretic curvature pattern do not necessarily exhibit other trademark aspects of Prospect Theory. For example, we find very little loss aversion, with most subjects disliking losses less than liking equivalent gains. The details of underlying preferences depend on bidder experience and the presence of a decision aid—a strong moderator of bidder entry we document: when bidders have experience and the decision aid (a situation we feel best approximates the long run), their preferences exhibit more loss aversion and less probability weighting, with the latter completely disappearing among a majority of the bidders.

The model helps us interpret our answer to the main managerial question of this paper regarding the profitability of two-part tariffs in the NYOP setting. Instead of providing the bidding opportunity free of charge, an NYOP seller who uses a two-part tariff charges an upfront non-refundable fee akin to an entrance fee into his store. The existing theory that advocates charging bidding fees is based on buyer risk neutrality and ambivalence to the very concept of a bidding fee. By contrast, real buyers are likely to be risk averse and may have an a priori negative reaction to paying a fee—an attitude we call “fee aversion.” We find bidding fees can be profitable for the NYOP seller, partly because our subjects do not exhibit fee aversion. However, the optimal fee to charge is smaller than that suggested by the model with risk-neutral buyers, which tends to under-predict the profitability of NYOP selling. The risk-neutral model under-predicts the profitability of NYOP selling with a two-part tariff, because most of our subjects bid more than risk-neutral people would, and yet also enter the NYOP store and pay fees risk-neutral people would not. Our model reconciles the resulting puzzle by rejecting Expected Utility Theory and estimating a robust pattern of insensitivity to loss magnitude.

We extend the common finding of excessive entry into first-price sealed-bid auctions to the NYOP domain, effectively showing excessive buyer entry does not arise only because of their inability to think

strategically about their competition against other bidders. Instead, we provide evidence that the excessive entry we observe is at least partly driven by the cognitive difficulty of trading off the probability of winning and the surplus contingent on winning: when we give our subjects a decision aid that merely displays the two parts of the tradeoff for any candidate bid the bidder enters into the bidding interface, excessive entry is sharply reduced, and the profitability of two-part tariffs with it. Using our model, we are able to interpolate between the few bidding-fee levels included in the experiment, and we find that when buyers have both some experience and the decision aid, the optimal bidding fee is substantially lower than in the baseline condition, and its expected profit lift is only 9% versus the 36% available in the baseline condition. Both profit-lift numbers should be contrasted to the 67% lift predicted by risk-neutral theory. We conclude two-part tariffs are a viable option for increasing NYOP profits even in the long run, but the profit improvements are much smaller than one would expect based on simple theoretical assumptions.

The model-based interpolation between experimental fee levels is an important contribution for practice, because a real-world experiment with bidding fees in NYOP will likely be restricted to only a few fee levels in order not to confuse or upset existing customers. Using an easy modification of the current specification, performing a counterfactual search for the optimal fee level should be possible based on a field experiment with just a single positive fee level: one only needs to switch the specification of utility in the loss domain from the current non-parametric form to some parametric formula. Another practical contribution we make is a comparison of a two-part tariff and the minimum-bid strategy—another NYOP-profit-enhancing strategy proposed in the literature. Contrary to some theoretical predictions, we find the two-part tariff outperforms minimum bids. Overall, we thus propose real-world sellers should focus on two-part tariffs as their primary method for increasing NYOP profits.

REFERENCES

- Abbas, Ali E. and Il-Horn Hann (2010), “Measuring Risk Aversion in a Name-Your-Own-Price Channel,” *Decision Analysis*, 7 (1), 123–136.
- Amaldoss, Wilfred and Sanjay Jain (2008), “Joint Bidding in the Name-Your-Own-Price Channel. A Strategic Analysis,” *Management Science*, 54 (10), 1685–1699.
- Amir, On and Dan Ariely (2007), “Decisions by Rules. The Case of Unwillingness to Pay for Beneficial Delays,” *Journal of Marketing Research*, 44 (1), 142–152.
- Andersson, Ola, Håkan J. Holm, Jean-Robert Tyran and Erik Wengström (2014), “Deciding for Others Reduces Loss Aversion,” *Management Science*, 141223041315002.
- Belobaba, Peter P. (1989), “Application of a Probabilistic Decision Model to Airline Seat Inventory,” *Operations Research*, 37 (2), 183–197.
- Bernhardt, Martin and Martin Spann (2010), “An Empirical Analysis of Bidding fees in Name-your-own-price Auctions,” *Journal of Interactive Marketing*, 24 (4), 283–296.
- Charness, Gary, Uri Gneezy and Alex Imas (2013), “Experimental methods. Eliciting risk preferences,” *Journal of Economic Behavior & Organization*, 87, 43–51.
- Chernev, Alexander (2003), “Reverse Pricing and Online Price Elicitation Strategies in Consumer Choice,” *Journal of Consumer Psychology*, 13 (1/2), 51–62.
- Cox, James C., Bruce Roberson and Vernon L. Smith (1982), “Theory and Behavior of Single Object Auctions,” in *Research in Experimental Economics*, Vol. 1 Smith, ed. Greenwich, CT: JAI Press, 1–43.
- Cox, James C., Vernon L. Smith and James M. Walker (1988), “Theory and individual behavior of first-price auctions,” *Journal of Risk and Uncertainty*, 1 (1), 61–99.
- Croson, Rachel, and Uri Gneezy. 2009. "Gender Differences in Preferences." *Journal of Economic Literature*, 47(2): 448-74.
- DellaVigna, Stefano (2009), “Psychology and Economics. Evidence from the Field,” *Journal of Economic Literature*, 47 (2), 315–372.
- Ding, Min, Jehoshua Eliashberg, Joel Huber and Ritesh Saini (2005), “Emotional Bidders—An Analytical and Experimental Examination of Consumers’ Behavior in a Priceline-Like Reverse Auction,” *Management Science*, 51 (3), 352–364.
- Dohmen, Thomas, Armin Falk, David Huffman and Uwe Sunde (2012), “The Intergenerational Transmission of Risk and Trust Attitudes,” *The Review of Economic Studies*, 79 (2), 645–677.
- Dorsey, Robert and Laura Razzolini (2003) *Experimental Economics*, 6 (2), 123–140.
- Ertaç, Seda, Ali Hortaçsu and James W. Roberts (2011), “Entry into auctions. An experimental analysis,” *International Journal of Industrial Organization*, 29 (2), 168–178.

- Fay, Scott (2004), "Partial Repeat Bidding in the Name-Your-Own-Price Channel," *Marketing Science*, 23 (3), 407–418.
- Fay, Scott and Juliano Laran (2009), "Implications of Expected Changes in the Seller's Price in Name-Your-Own-Price Auctions," *Management Science*, 55 (11), 1783–1796.
- Fay, Scott and Seung H. Lee (2015), "The role of customer expectations in name-your-own-price markets," *Journal of Business Research*, 68 (3), 675–683.
- Filiz-Ozby, Emel and Erkut Y. Ozby (2007), "Auctions with Anticipated Regret. Theory and Experiment," *American Economic Review*, 97 (4), 1407–1418.
- Fischbacher, Urs (2007), "z-Tree. Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 10 (2), 171–178.
- Gächter, Simon, Eric J. Johnson and Andreas Herrmann (2010), "Individual-Level Loss Aversion in Riskless and Risky Choices," CeDEx Discussion Paper.
- Greiner, Ben (2015), "Subject pool recruitment procedures. Organizing experiments with ORSEE," *Journal of the Economic Science Association*, 1 (1), 114–125.
- Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong (2009), "Nonparametric Identification of Risk Aversion in First-Price Auctions under Exclusion Restrictions." *Econometrica* 77, no. 4: 1193-227.
- Hann, Il-Horn and Christian Terwiesch (2003), "Measuring the Frictional Costs of Online Transactions. The Case of a Name-Your-Own-Price Channel," *Management Science*, 49 (11), 1563–1579.
- Hinz, Oliver, Il-Horn Hann and Martin Spann (2011), "Price Discrimination in E-Commerce? An Examination of Dynamic Pricing in Name-Your-Own Price Markets," *MIS Quarterly*, 35 (1), 81–98.
- Hinz, Oliver and Martin Spann (2008), "The Impact of Information Diffusion on Bidding Behavior in Secret Reserve Price Auctions," *Information Systems Research*, 19 (3), 351–368.
- Holt, Charles A. and Susan K. Laury (2002), "Risk Aversion and Incentive Effects," *American Economic Review*, 92 (5), 1644–1655.
- Kahneman, Daniel and Amos Tversky (1979), "Prospect Theory. An analysis of decision under risk," *Econometrica*, 47, 263–291.
- Kőszegi, Botond and Matthew Rabin (2007), "Reference-Dependent Risk Attitudes," *American Economic Review*, 97 (4), 1047–1073.
- Palfrey, Thomas R. and Svetlana Pevnitskaya (2008), "Endogenous entry and self-selection in private value auctions. An experimental study," *Journal of Econ. Behavior & Organization*, 66 (3-4), 731–747.
- Platt, Brennan C., Joseph Price and Henry Tappen (2013), "The Role of Risk Preferences in Pay-to-Bid Auctions," *Management Science*, 59 (9), 2117–2134.
- Prelec, Drazen (1998), "The Probability Weighting Function," *Econometrica*, 66 (3), 497.

- Riley, John G. and William F. Samuelson (1981), "Optimal Auctions," *American Economic Review*, 71 (3), 381–392.
- Shapiro, Dmitry (2011), "Profitability of the Name-Your-Own-Price Channel in the Case of Risk-Averse Buyers," *Marketing Science*, 30 (2), 290–304.
- Shapiro, Dmitry and Arthur Zillante (2009), "Naming your own price mechanisms. Revenue gain or drain?," *Journal of Economic Behavior & Organization*, 72 (2), 725–737.
- Smith, James L. and Dan Levin (1996), "Ranking Auctions with Risk Averse Bidders," *Journal of Economic Theory*, 68 (2), 549–561.
- Smith, Vernon L. (1976), "Experimental Economics. Induced Value Theory," *American Economic Review*, 66 (2), 274–279.
- Sommer, Adam (2014), "Top 5 Reasons to Use Priceline," [available at <http://www.visitflyovercountry.com/2013/02/top-5-reasons-to-use-priceline.html>].
- Spann, Martin, Gerald Häubl, Bernd Skiera and Martin Bernhardt (2012), "Bid-Elicitation Interfaces and Bidding Behavior in Retail Interactive Pricing," *Journal of Retailing*, 88 (1), 131–144.
- Spann, Martin, Bernd Skiera and Björn Schäfers (2004), "Measuring individual frictional costs and willingness-to-pay via name-your-own-price mechanisms," *Journal of Interactive Marketing*, 18 (4), 22–36.
- Spann, Martin, Robert Zeithammer and Gerald Häubl (2010), "Optimal Reverse-Pricing Mechanisms," *Marketing Science*, 29 (6), 1058–1070.
- (2015), "Erratum to "Optimal Reverse-Pricing Mechanisms" by Martin Spann, Robert Zeithammer, and Gerald Häubl," *Marketing Science*, 34 (2), 297–299.
- Sutton, John (1993), "Much Ado about Auctions: Some Introductory Remarks," *European Economic Review*, 37, 319–320.
- Terwiesch, Christian, Sergei Savin and Il-Horn Hann (2005), "Online Haggling at a Name-Your-Own-Price Seller. Theory and Application," *Management Science*, 51 (3), 339–351.
- Tversky, Amos and Daniel Kahneman (1991), "Loss Aversion in Riskless Choice. A Reference-Dependent Model," *The Quarterly Journal of Economics*, 106 (4), 1039–1061.
- (1992), "Advances in prospect theory. Cumulative representation of uncertainty," *Journal of Risk and Uncertainty*, 5 (4), 297–323.
- Wang, Tuo, Esther Gal-Or and Rabikar Chatterjee (2009), "The Name-Your-Own-Price Channel in the Travel Industry. An Analytical Exploration," *Management Science*, 55 (6), 968–979.
- Zeithammer, Robert (2015), "Optimal selling strategies when buyers name their own prices," *Quantitative Marketing and Economics*, 13 (2), 135–171.

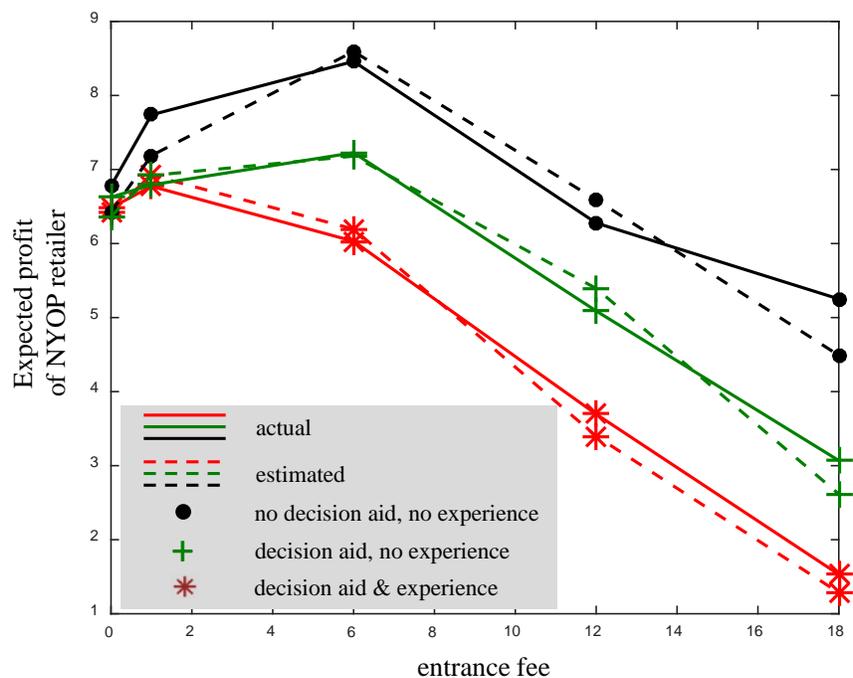
APPENDIX

Figure A1: Experimental Interface Of Experiment 1

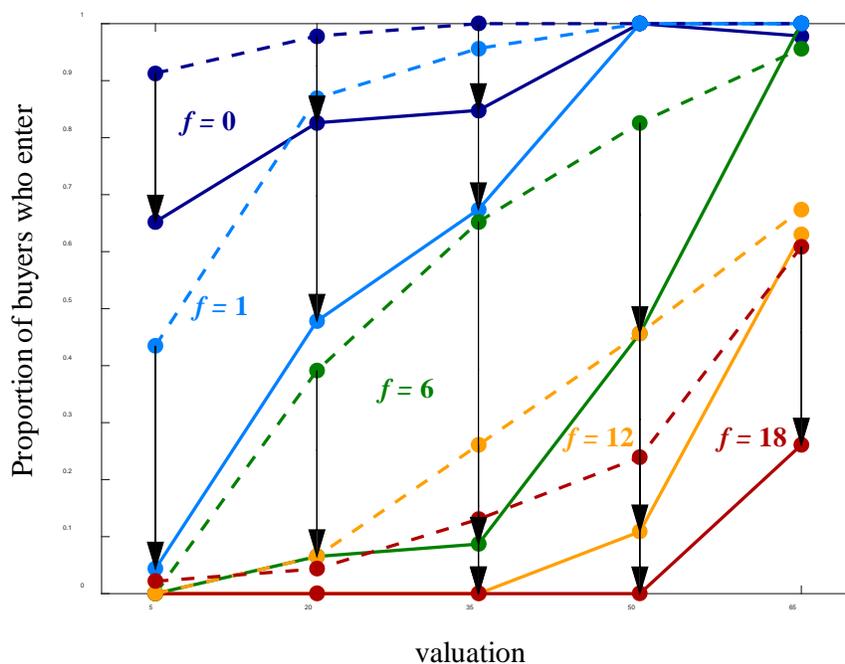
Period x of x		
<p>A new period begins. Your valuation is: xx The bidding fee in store A is: xx The posted price in store B is: xx</p>		
Please choose from the below options.		
<p style="text-align: center;">Store A</p> <p>The bidding fee is: xx Please enter your bid: <input style="width: 80px; height: 20px;" type="text"/></p> <p>In case your bid is not successful, you can still purchase from store B.</p> <p style="text-align: center;"><input type="button" value="Submit Bid"/></p>	<p style="text-align: center;">Store B</p> <p>The posted price is: xx Do you want to buy at the posted price?</p> <p style="text-align: center;"><input type="button" value="Buy"/></p>	<p style="text-align: center;">Don't Buy In This Period</p> <p style="text-align: center;">In case you don't want to buy in this period, please press the button "Don't buy" below.</p> <p style="text-align: center;"><input type="button" value="Don't buy"/></p>

Figure A2: Experimental Interface Of Experiment 2

<p>A new period begins. Your valuation is: 65 The bidding fee is: 1 The fixed posted price for the product is: 70 Because of your valuation, buying at the posted price would be too expensive for you.</p>	
Please choose from the following options.	
<p style="text-align: center;">Submit a bid</p> <p>The bidding fee is: 1 Please enter your bid: <input style="width: 50px; height: 20px; text-align: center;" type="text" value="40"/></p> <p>The chance your bid is accepted is: 57 out of 100 If your bid is accepted, your payoff will be: 24</p> <p style="text-align: center;"><input type="button" value="Submit Bid"/></p>	<p style="text-align: center;">Don't Buy On This Period</p> <p style="text-align: center;">In case you don't want to buy in this period, please press the button "Don't buy" below.</p> <p style="text-align: center;"><input type="button" value="Don't buy"/></p>

Figure A3: Model Fit: Predicted Vs. Actual Seller Profit, By Condition

Note: The observed lines are the same as in Figure 3. The dashed lines are estimated from the model.

Figure A4: Effect Of Decision Aid And Experience On Entry, Within Subject

Note: Only subjects who start without the decision aid ($N=46$) are considered. The solid lines connect the observed entry probabilities when the subjects do not have the decision aid (first block); the dashed lines do the same thing when the subjects do have the decision aid (second block). Arrows indicate significant differences at the 5% level.

Log-likelihood specification

We first derive the likelihood of a bid observation. Taking logs of both sides of equation (5) yields a truncated normal log-likelihood of the logarithm of observed bid data:

$$\begin{aligned} \log L_{n,bid}(\theta, \sigma | v_n, f_n) = & -\log(\sigma) - \frac{(\log \text{bid}_n - \log b(v_n, f_n | \theta))^2}{2\sigma^2} \\ & - \log \left[\Phi \left(\frac{\log(\min(v_n, p)) - \log(b(v_n, f_n | \theta))}{\sigma} \right) \right], \end{aligned} \quad (\text{A1})$$

where the last term accounts for the truncation. Overall, the likelihood of the data combines the entry information with the observed bid information as follows:

$$\begin{aligned} \log L(\theta, J, \sigma, \tau) = & \sum_{n \text{ where entry}} \log \left[\Pr_{enter}(v_n, f_n | \theta, J, \tau) \right] + \log L_{n,bid}(\theta, \sigma | v_n, f_n) \\ & + \sum_{n \text{ where no entry}} \log \left[1 - \Pr_{enter}(v_n, f_n | \theta, J, \tau) \right]. \end{aligned} \quad (\text{A2})$$

Note this likelihood is not defined when the observed bid is zero or when the bid exceeds the valuation minus the fee. We find a handful of such observations within the entire dataset, and omit them from estimation.