

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

December 12, 2018

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

Abstract: Prior theoretical research has shown that a Name-Your-Own-Price (NYOP) seller can profit from charging each prospective buyer a non-refundable fee for the opportunity to place a bid, akin to an entry fee to the seller's store. We examine the profitability of such two-part tariffs in NYOP markets using incentive-compatible laboratory experiments. Our results show that a two-part tariff can be profitable for a NYOP seller, but the profitability is strongly moderated by the buyers' ability to jointly optimize their entry and bidding decisions. We find that when buyers are provided with a decision aid that partly offsets their cognitive limitations by calculating the payoff consequences of different candidate bids for them, the profitability of using a two-part tariff vanishes. Overall, our results suggest two-part tariffs increase NYOP profit in a standard information-poor setting in the short run, but they are not as profitable as theory would suggest when the bidders get more information or experience. We also propose an individual-level non-parametric test of the risk-averse expected utility model, and the test results suggest that this canonical model for decision making under uncertainty is not a good fit to the behavior of a substantial proportion of our subjects.

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

Contact: Robert Zeithammer, rzeitham@ucla.edu

1. Introduction

Name Your Own Price (NYOP) refers to a pricing mechanism pioneered by Priceline in 1997 to sell travel products, and has been 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. Retailers continue to innovate and experiment with NYOP. For example, eBay has been using NYOP in its “Best Offer” bilateral mechanism for over a decade, and it is currently considering a “Name Your Price” mechanism to facilitate an interaction between one buyer and multiple providers of the same item (Steiner 2018). Despite this innovation and the attractive features of NYOP, such as its ability to empower buyers and facilitate fast price negotiations, NYOP remains a niche pricing strategy. The theoretical literature has proposed that one of the reasons NYOP has not diffused more widely is that its existing implementations leave potential seller profits on the table (Spann, Zeithammer, and Häubl 2010; Zeithammer 2015). Using incentive-compatible laboratory experiments, we test one proposed improvement—a two-part tariff. Instead of providing the bidding opportunity free of charge, an NYOP seller who uses a two-part tariff charges an upfront non-refundable fee for the opportunity to submit a bid.² 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). Theoretically, the same ability to extract consumer surplus that makes two-part tariffs profitable for a posted-price seller also operates within NYOP as long as buyers are risk neutral (Spann, Zeithammer,

¹ Examples include event tickets at scorebig.com, electronics at greentoe.com, a variety of things on eBay.com under its “Best Offer” mechanism, and, until recently, restaurant vouchers on chiching.com.

² 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.

and Häubl 2010). This paper presents the first empirical test of a two-part tariff in NYOP selling with real bidders who may not be risk neutral, may not like to pay fees for a mere chance of getting a good deal, and may find the task of NYOP bidding unfamiliar and difficult.

Despite the clear theoretical advantage of two-part tariffs over providing the bidding opportunity for free³, real-world NYOP sellers have not implemented two-part tariffs yet. One candidate explanation for this absence of a potentially profitable pricing scheme is that NYOP sellers fear bidders may exhibit “fee aversion” (i.e., a decision rule to never pay for anything other than the product itself; Amir and Ariely 2007). We use very small bidding fees to test for fee aversion in bidder behavior, and we do not find evidence of fee aversion. Another reason why real-world sellers do not use two-part tariffs is that such mechanisms are too complex, and the associated decision-making too difficult for consumers. To examine this potential explanation, we test the impact of a decision aid on bidder behavior and NYOP seller profit. Decision aids are a common online tool to assist consumers in their purchase decisions (Häubl and Trifts 2000) and are already offered by some NYOP sellers.⁴ Our 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 upon bid acceptance.

We find two-part tariffs can be profitable for NYOP sellers but only (1) when the potential buyers do not have access to our decision aid and (2) when they do not have too much experience with our laboratory NYOP market. The novelty of NYOP two-part tariffs and cognitive difficulty associated with buyers’ joint optimization of their entry and bidding decisions are thus catalysts of two-part-tariff profitability in NYOP, not inhibitors, as one might initially hypothesize. Although decision aids can help

³ The advantage of a two-part tariff in NYOP settings stems from an increase in consumer surplus captured by the seller. When the bidding opportunity is free, all buyers submit a bid. Raising the first part of the tariff (the “bidding fee”) from zero makes all but the lowest-value buyers pay more in expectation, and only detracts few lowest-value buyer from participating. Since buyers with the lowest value generate no profit for the seller, raising the first part of the tariff from zero is always profitable.

⁴ For example, greentoe.com uses a tachometer-style decision aid to inform about the chance of success of a bid.

reduce the complexity associated with the NYOP channel for buyers (Fay and Lee 2015), our results suggest sellers may have a strategic incentive to prevent buyers from having complete information about their chances and payoffs. Indeed, even real-world NYOP sellers who do not use two-part tariffs tend to limit the amount of publicly available information needed for bid optimization; for example, Priceline does not offer information about the chances of success at different bid levels.

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 less 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 subjects enter excessively because they find the joint optimization of their entry and bidding decisions difficult, we designed the above-described decision aid that helps buyers calculate the potential consequences of their planned actions. Our second study finds that the decision aid reduces buyers' willingness to pay bidding fees, which in turn reduces the profitability of using a two-part tariff. Mere experience also reduces seller profits and primarily affects the entry decision, but it also affects it differently (more crudely) than the decision aid, suggesting a different underlying mechanism.

The behavior of most of our subjects is clearly inconsistent with the risk-neutral model most analytical work assumes. In the last section of the paper, we provide novel evidence that the behavior of a substantial proportion of our subjects is also inconsistent with the canonical risk-averse Expected Utility Theory model (Bernoulli 1738, von Neumann and Morgenstern 1944). Our evidence is non-parametric, and consists of an unreasonably high lower bound on the errors that would be needed to rationalize the observed behavior of some subjects under the theory.

In summary, we find two-part tariffs do not increase NYOP profit as existing theory predicts they should, and we provide meaningful constraints for future models of NYOP markets with bidding fees.

2. Literature Review

This paper is related to two strands of literature: (1) the literature on NYOP selling and (2) the experimental auction literature. 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, 2015) that a two-part tariff can be profitable for an NYOP seller, and by documenting strong moderating roles of novelty and cognitive difficulty not previously addressed by theoretical models.

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 levels (Hinz and Spann 2008, Ding et al. 2005), 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.

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 and is not affected by cognitive difficulty of the joint entry and bidding decision.

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 phenomenon. Ertaç et al. (2011) propose a model that combines risk aversion and the joy of winning to explain it. They show that a model incorporating the joy of winning together with risk aversion better matches the observed entry behavior than a model lacking those components. We at least partially explain over-entry as a result of heuristic thinking in the face of cognitive difficulties. We do not find evidence of the joy of winning because even low-valuation bidders bid within their meagre means, and hence choose a low chance of winning. Perhaps “winning” in NYOP bidding is not as desirable because it merely arises from the seller accepting the bid; there are no competitors to beat.

3. Model

We model the NYOP setting as follows: a buyer wants to buy a specific indivisible object (e.g. a room in a four-star hotel in a given part of a city on particular dates on Priceline.com). The object is available in an outside posted-price spot-market for a commonly-known price p . There also exists one NYOP seller, who can procure the same object for a cost c ex-ante distributed uniformly on $[0,p]$, but privately revealed to the seller after the buyer's bid is received. The seller charges a non-refundable bidding fee $f \geq 0$.

The buyer i 's utility u_i of buying the object depends only on his surplus s , i.e. his valuation v minus the total cost incurred, and WLOG $u_i(0)=0$ and $u_i(1)=1$.⁵ So buying from the posted-price spot-market yields a utility of $u_i(v-p)$, buying from the NYOP seller who accepted a bid b yields a utility of $u_i(v-f-b)$, and paying the fee but getting the bid rejected yields $u_i(-f)$. Given these preferences, the buyer solves the following bidding problem:

$$bid(v, f) = \arg \max_{b \geq 0} \Pr(b \text{ accepted})u_i(v-f-b) + [1 - \Pr(b \text{ accepted})]u_i(-f) \quad (1)$$

The buyer tries to buy from the NYOP seller whenever his expected utility exceeds zero, and stays out of the market otherwise. In the case their bids are rejected, buyers with $v > p$ buy from the posted-price market.

Our experiments are motivated by Spann, Zeithammer, and Häubl (2010), who consider risk-neutral buyers $u_i(s) = s$ with valuations drawn from a 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 $f^*(p) = 4M^2/49p$ as long as $p > 4M/7$. Given the assumed distribution of seller costs, the buyers face a bid-acceptance probability $\Pr(b \text{ accepted}) = \Pr(c < b) = bid / p$. Risk-

⁵ To see how this model includes standard expected utility theory, let the consumer i 's utility over wealth ω be $\tilde{u}(\omega)$, let ω_i be the consumer i 's initial wealth, and define $u_i(s) \equiv \frac{\tilde{u}(\omega_i + s) - \tilde{u}(\omega_i)}{\tilde{u}(\omega_i + 1) - \tilde{u}(\omega_i)}$.

neutrality implies the following buyer behavior: for any $f < p/4$, 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 $b(v) = v/2$, and “high” bidders with $v \geq p$ who mimic the bidder with $v = p$ and bid $p/2$ because they have a real option of buying in the outside market should their NYOP bid not be successful.

We do not attempt to estimate the precise shape of $u_i(s)$ in our subject population, and most of our results are model-free. Specifically, we first test the managerially relevant qualitative prediction from the $u_i(s) = s$ model that two-part tariffs are more profitable than providing the bidding opportunity for free, and then document a strong moderator of such profitability. An obvious alternative hypothesis to $u_i(s) = s$ is that $u_i(s)$ are concave, and so the buyers are risk-averse under the Expected Utility Theory. We conclude the paper with a novel individual-level test of $u_i(s)$ being concave. The test results are mixed, and a more detailed measurement of $u_i(s)$ is beyond the scope of this paper.

4. Experimental paradigm: Within-subject manipulation of buyer’s valuation and seller’s bidding fee

Throughout this paper, we use incentive-compatible laboratory experiments to study the 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-subjects 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 within-subjects 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, {5, 20, 35, 50, 65, 80, 95} in Study 1 and {5, 20, 35, 50, 65} in Study 2. 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, and 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 ($p=70$ 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 $[0,p]$.⁶ 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

⁶ 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.

measurement of entry behavior (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 points between 0, which everyone should be willing to pay, to 18,⁷ 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 our 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 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 the 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.

5. Experiment 1: Test of the Risk-Neutral Benchmark Model

Our first study is motivated by the predictions of Spann, Zeithammer, and Häubl (2010) discussed in Section 3. In experiment 1, we set $M=100$ and $p=70$, resulting in $f^*(p)=11.7$, and we approximate the

⁷ We could have selected fees randomly from a continuum between 0 and 18. By contrast, selecting only a few specific values allows for a clean full-factorial within-subjects design.

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 explain the experimental procedure in detail.

5.1 Experiment 1: Experimental Procedure

Each round carries out the following experimental procedure: First, 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 (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 $[0, p]$ and accept the bid when it exceeds w . When 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. When the buyer purchases the product from the posted-price store right away, her payoff in that round is her valuation minus the posted price. When the NYOP store accepts a buyer's bid, the buyer's payoff in that round is her valuation minus the bid submitted and minus the bidding fee. If the NYOP store rejects a buyer's bid, the final payoff is contingent on her subsequent decisions. If she decides not to use her second chance to buy from the posted-price store, her final payoff from the round is 0 minus the fee. On the other hand, if she buys from the posted-price store, her final payoff is her valuation minus the posted price *and* minus the fee. Each round is incentivized; subjects are shown their payoff after each round but not their cumulative profit to limit potential wealth effects.

A pilot study found some subjects used the first few rounds to explore actual consequences of seemingly irrational behavior, such as bidding above one's valuation. Once they incurred an avoidable loss, most subjects refrained 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: the number of “safe” choices in the paired lottery choice task by Holt and Laury (2002); the number of rejected risky lotteries in the lottery choice task by Gächter et al. (2010); the subjective risk-taker scale by Dohmen et al. (2012); a question about attitudes toward bidding fees; and the frequency of participation in lab experiments to date. All scales were administered as hypotheticals, not separately incentivized beyond a flat payment.

5.2 Experiment 1: 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 excluded from our analysis the five subjects who bid over valuation even once after the training rounds. Everything that follows is based on the 91 remaining subjects.

5.3 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 $\pi(\text{bid})$ as follows:

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

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. The contribution of each subject-round to the seller's profit is thus

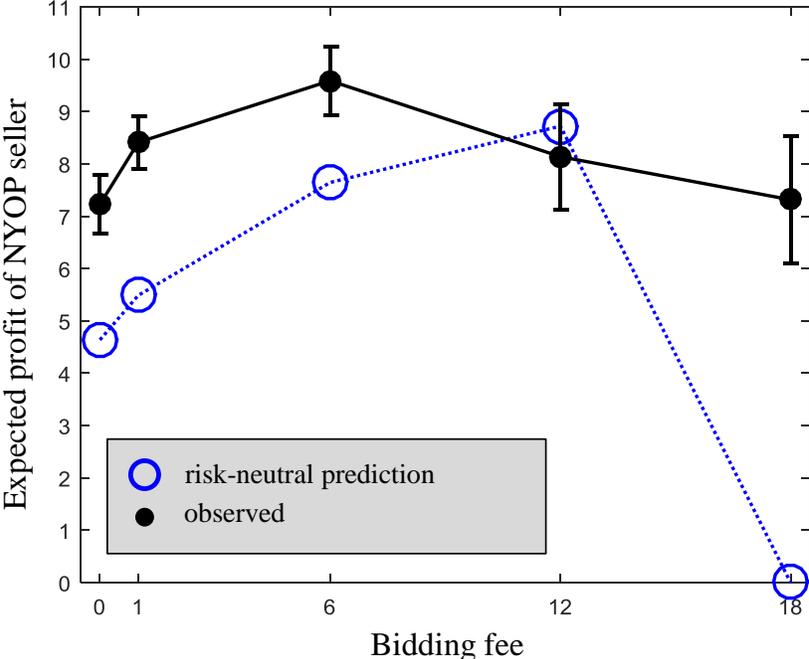
$$\text{contribution}_{i,v,f} = \mathbf{1}(\text{bid}_{i,v,f} > 0) \left[f + \pi(\text{bid}_{i,v,f}) \right]. \quad (3)$$

Then we calculate the expected NYOP seller profit when the fee is f as the average over i and v of $\text{contribution}_{i,v,f}$. Figure 1 plots this expected seller profit and compares it with the theoretical profit expected under risk neutrality.

We start our discussion of Figure 1 by considering the main prediction of the literature on two-part

tariffs in NYOP, 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).⁸ The expected profit with $f=6$ is also significantly greater than the adjacent fee levels in our design, so the empirically 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:

Figure 1: NYOP Seller Profit



Note to figure: The solid (black) 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.

⁸ 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-subjects differences in expected profit to zero. The test using within-subjects 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.

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

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

5.4 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 underpredict profits is that it underpredicts 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.

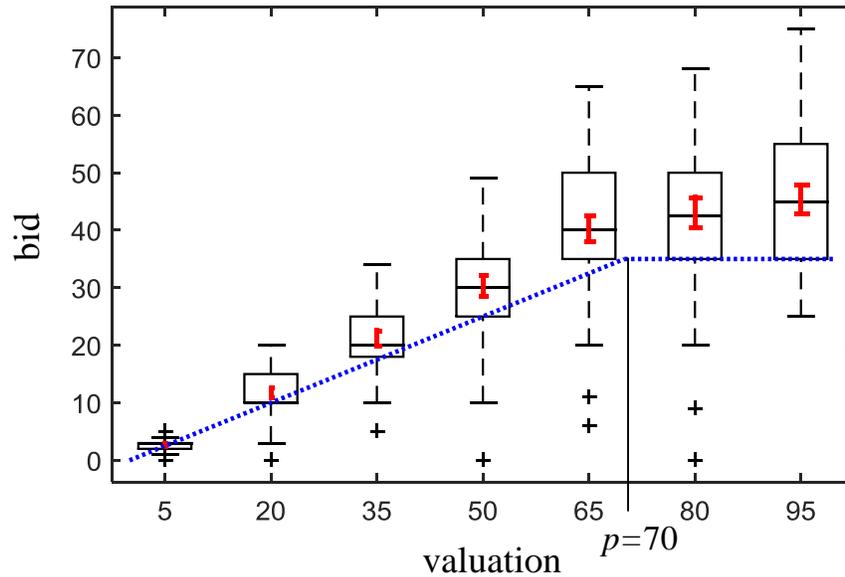
The first two columns of Table 1 also provide evidence against the fee-aversion hypothesis: When the fee increases from 0 to 1, entry does decline, but only in the low-valuation conditions ($v \leq 20$). By contrast, a “knee-jerk” fee aversion would predict a decline in entry irrespective of valuation. Not only does fee aversion 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).⁸

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).

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 bid 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].

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 is 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 find no 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 the 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 subjects in assessing the payoff consequences of their potential bids. We now turn to the details of Experiment 2.

6. 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 in return. In Experiment 2, we examine whether the excessive entry observed in Experiment 1 can be attributed to the cognitive difficulty of solving the optimal bidding problem, by using an intervention that allows buyers to better 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 are required to click a button that activated the decision aid, which instantaneously displays 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's output).

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-subjects 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 $\{5, 20, 35, 50, 65\}$. Due to the

already large number of tasks per subject, we also omitted the “training” tasks in Experiment 2.

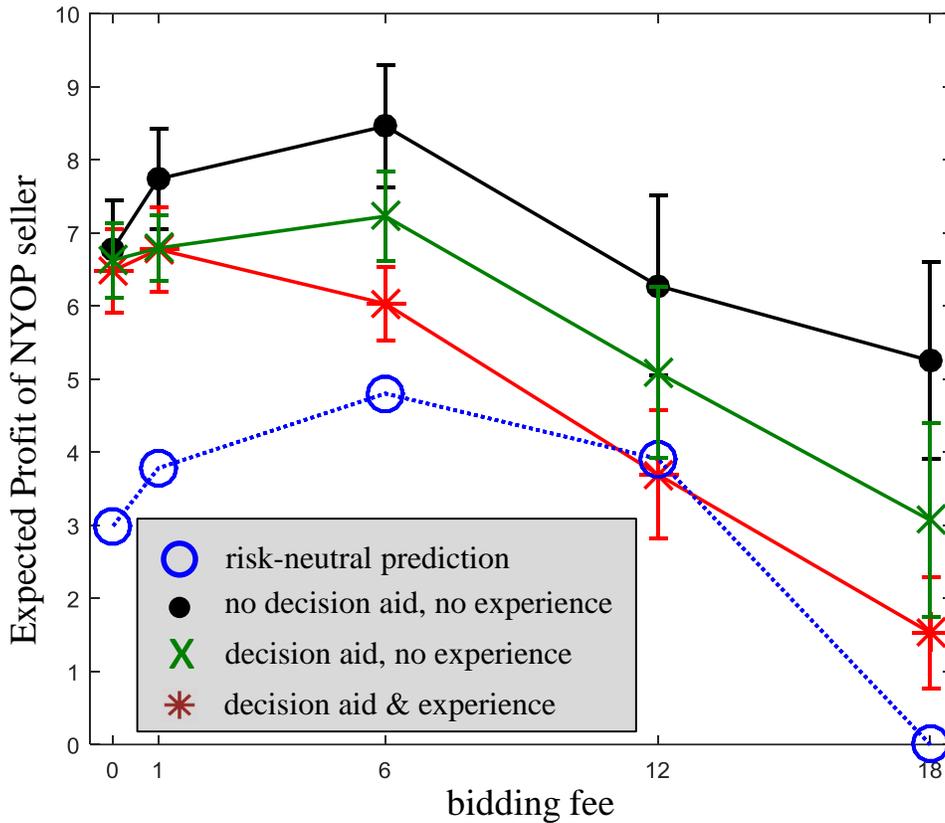
The data collection was analogous to that in Experiment 1, with four sessions of 24 subjects. We eliminated four subjects who bid more than their valuation more than once. All the analyses that follow are based on the remaining 92 subjects (46 in each cross-over condition).

6.1 Experiment 2: Effect of Decision Aid

Comparing the first-block behavior of subjects with the decision 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 profit 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 the number of observations to be the number of subjects, analogously with the approach in Study 1. The only exception is the between-subjects difference at $f=12$. In contrast to the positive fees, neither the decision aid nor its combination with experience has 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 (almost all bidders enter when the fee is zero, regardless of condition). In addition to comparing the different conditions in our study, Figure 3 also clearly shows the main puzzling 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.

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.

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 (i.e., replicating⁹ the situation of Experiment 1), the seller's profits significantly differ from charging the empirically optimal bidding fee of 6; relative to zero fee, profits increase 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

⁹ Note that restricting the range of valuations to below the posted price of 70 changes the shape of the risk-neutral prediction relative to Study 1, making a lower fee of 6 optimal in Study 2 under risk-neutrality.

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.

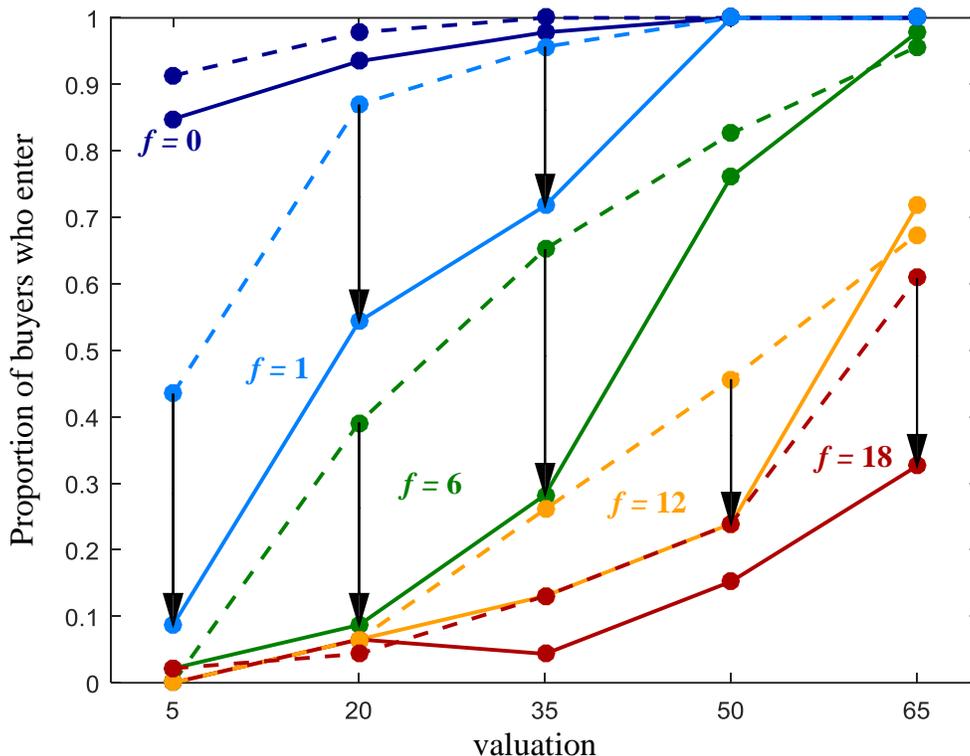
We summarize the takeaway from Figure 3 in our next finding:

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 significant evidence of the profitability of a two-part tariff 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-subjects effect of the decision aid on entry. 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. We do not report the within-subjects effect of the decision aid and experience, because it is qualitatively very similar to the between-subjects 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. See Figure A3 in the Appendix for a plot of the within-subjects 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 primarily affects entry into the NYOP store.

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.

6.2 Experiment 2: Effect of Experience

Because each subject experiences the different (v, f) conditions in random order, we can isolate the effect of mere experience between subjects by focusing only on the subjects without the decision aid in their first block (their first 25 rounds), and analyzing their behavior as a function of round. To assess the impact of experience on seller profit, we regress the person-round contribution defined in equation (2) on round, controlling for valuation, fee, and their interactions. Table 2 shows the detailed regression results, which imply seller profit falls by about 0.10 per round during the first 25 rounds, *ceteris paribus*. We interpret this effect as a result of buyer experience. Interestingly, the same regression on the subjects who have the

decision aid in their first block does not yield a statistically significant relationship between round and contribution, so learning from experience seems to be stronger in the cognitively harder condition.

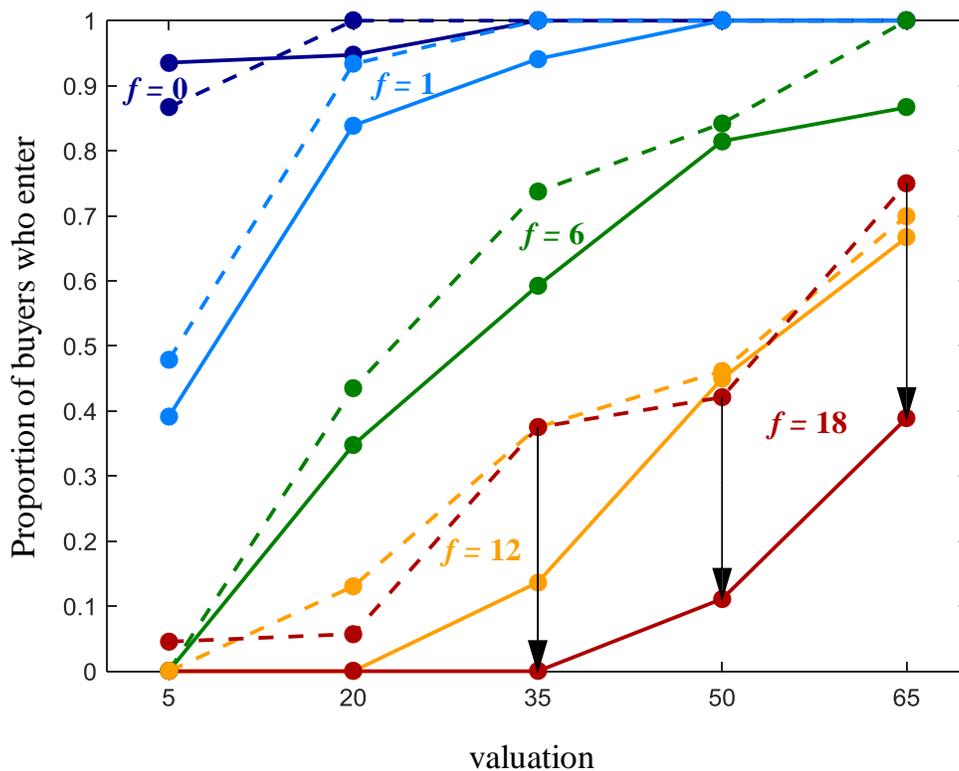
Table 2: Effect of Round on Seller Profit, First Block without Decision Aid

Variable	Coefficient	t-statistic based on clustered standard errors	
		no fixed effects	individual fixed effects
constant	1.86	2.62	-
individual fixed effects	-	-	included
round	-0.10	-3.11	-3.05
valuation 20	0.82	1.74	1.70
valuation 35	3.16	4.79	4.69
valuation 50	10.39	11.61	11.38
valuation 65	16.51	13.56	13.29
fee 1	-0.15	-0.34	-0.33
fee 6	-0.15	-0.41	-0.40
fee 12	-0.49	-1.14	-1.12
fee 18	-0.34	-1.57	-1.54
valuation20 X fee1	0.98	1.88	1.84
valuation20 X fee6	1.36	2.34	2.29
valuation 20 X fee12	-0.13	-0.15	-0.15
valuation 20 X fee18	-0.61	-0.99	-0.97
valuation 35 X fee1	2.25	3.17	3.11
valuation 35 X fee6	2.63	3.30	3.24
valuation 35 X fee12	0.34	0.30	0.30
valuation 35 X fee18	-0.69	-0.56	-0.55
valuation 50 X fee1	1.39	1.70	1.67
valuation 50 X fee6	1.90	1.50	1.47
valuation 50 X fee12	-1.78	-1.00	-0.98
valuation 50 X fee18	-5.34	-3.18	-3.11
valuation 65 X fee1	1.15	1.36	1.33
valuation 65 X fee6	3.50	3.36	3.29
valuation 65 X fee12	1.93	0.80	0.79
valuation 65 X fee18	0.38	0.14	0.13

Note to table: $N=1150$ observations (person-rounds). Standard errors clustered at the individual level (t-statistic of the round coefficient is even higher without clustering). The 45 individual fixed effects, where applicable, are omitted. R^2 is 0.55 without fixed effects and 0.59 with fixed effects. The coefficients on the valuation, fee, and round variables do not depend on the inclusion of the individual fixed effects.

The reduction of profit due to experience is mainly driven by a reduction in entry in later rounds, *ceteris paribus*. To demonstrate the reduction in entry at different fee levels, we compare the entry behavior of subjects who experience a given (v, f) condition early in the block with those who experience it later. This test has less power than the one considered in the previous subsection, because it only involves half the subjects ($N=46$). Figure 5 plots the result of this analysis analogously to Figure 4, illustrating 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 $p=0.06$).

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 13-25; the dashed lines do the same thing during periods 1-12. Arrows indicate significant differences at the 5% level using the two-proportion z -test; all are also significant at the 1% level.

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 interpret the different patterns of entry reduction from the decision aid and experience is to conjecture that the decision aid helps the subjects realize the normative tradeoff between maybe winning $v-f$ and losing f for sure, whereas mere experience focuses learning on fees and valuations separately. In other words, mere experience seems to teach subjects to avoid high fees, whereas the decision aid teaches subjects to avoid *combinations* of valuations and fees that offer only small expected surplus.

So far, all of our results have been essentially model-free in that the risk-neutral benchmark is not needed to establish the strong moderating role of the decision aid, nor is it needed to rule out fee aversion. It is clear, however, that the risk-neutral model does not capture the observed behavior well, and other models may perform better. In the next section, we explore the possibility that a canonical model of risk-aversion might fit well enough.

7. Excess Entry Through the Lens of Expected Utility: A Lower Bound on Error

As summarized in Result 1 and Figures 1-3, the risk-neutral model with $u_i(s) = s$ does not capture the observed behavior well. An obvious alternative hypothesis to $u_i(s) = s$ is that $u_i(s)$ are concave, and so the buyers are risk-averse under the assumptions of Expected Utility Theory model (Bernoulli 1738, von Neumann and Morgenstern 1944, hereafter "RAEU"). This section presents a novel non-parametric test of a RAEU hypothesis. The test results show that our decision aid improves the ability of RAEU to capture the entry behavior of some of our subjects, but a substantial segment remains for whom RAEU is not a good fit even with decision aid and experience.

Assume that person i enters the NYOP store in round t whenever

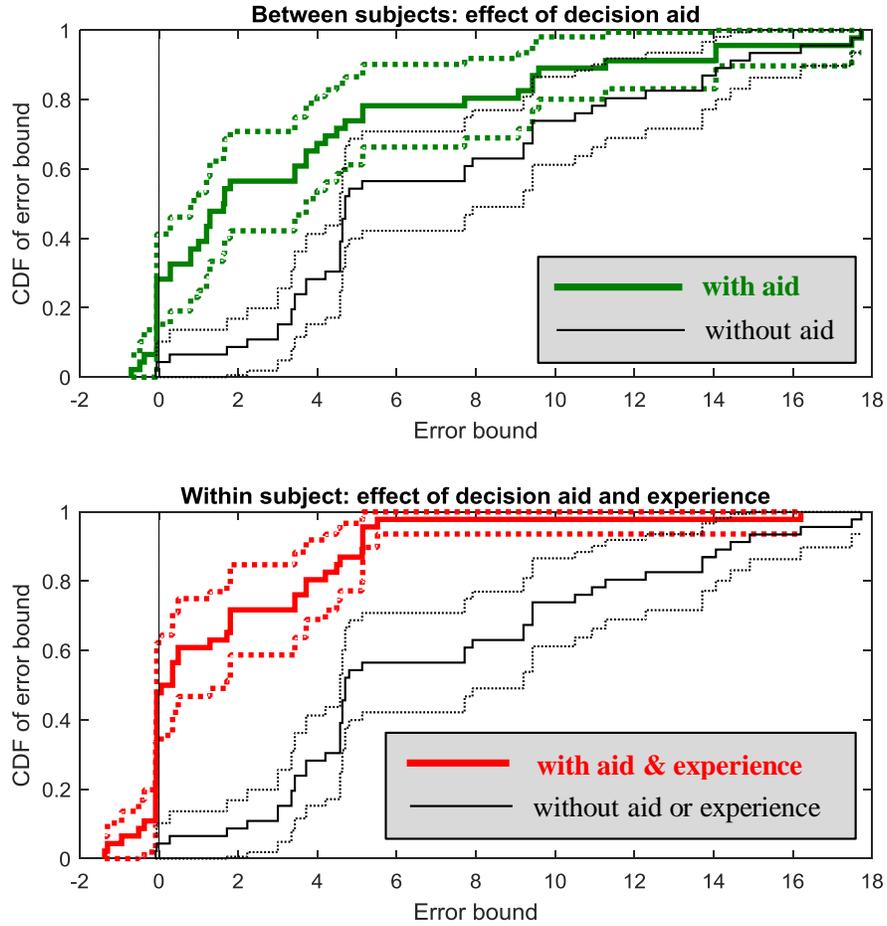
$$\underbrace{\left(\frac{bid_{i,t}}{p}\right)u_i(v_{i,t} - f_{i,t} - bid_{i,t}) + \left(1 - \frac{bid_{i,t}}{p}\right)u_i(-f_{i,t})}_{= \text{Expected utility}(v_{i,t}, f_{i,t}, bid_{i,t})} + \varepsilon_{i,t} > 0 \quad (4)$$

where $\varepsilon_{i,t}$ is an error in the spirit of random utility models (e.g. logit or probit). We now explain how RAEU implies a lower bound on $\varepsilon_{i,t}$. Under Expected Utility theory, risk-aversion is equivalent to concavity of the utility function u_i . Given our $u_i(0)=0$ and $u_i(1)=1$ normalization, any globally concave utility function is weakly dominated by its argument - namely the consumer surplus - everywhere other than on the $[0,1]$ interval. Since all our subjects bid whole numbers and the valuations and fees we used were also whole numbers, we do not observe any potential surpluses in the $[0,1]$ interval. Therefore, concavity or u_i (i.e., risk-aversion) implies $u_i(s) < s$ in our data. The $u_i(s) < s$ relationship allows us to construct a lower bound on the additive error $\varepsilon_{i,t}$ each of our subjects in Experiment 2 made in entering the NYOP store:

$$\varepsilon_{i,t}^{entry} > -\text{Expected utility}(v_{i,t}, f_{i,t}, bid_{i,t}) \stackrel{u_i(s) < s}{>} f_{i,t} - \left(\frac{bid_{i,t}}{p}\right)(v_{i,t} - bid_{i,t}) = -\text{Expected surplus}_{i,t} \quad (5)$$

In words: the error must exceed the negative of the expected surplus. Since the expected surplus facing each entrant is directly observable to us via a simple calculation, we can define a person-specific error bound as the maximum of all the bounds in equation 5: $\bar{\varepsilon}_i = \max_{t \text{ when } i \text{ entered}} (-\text{Expected surplus}_{i,t})$. The bound $\bar{\varepsilon}_i$ ignores non-entry and does not interpret the observed bid as optimal in any way. It relies only on equation 4 and the concavity of u_i . It thus allows us to examine the feasibility of RAEU at the individual level without making any assumptions about the precise shape of u . Figure 6 plots the population distribution of $\bar{\varepsilon}_i$ in three conditions of Experiment 2 examined before.

Figure 6: Effect of Decision Aid on Errors Needed to Explain Observed Entry Behavior by a Risk-Averse Expected Utility Theory



Note to figure: The dashed lines represent the 95% confidence intervals calculated by the Greenwood's Formula. Color coding is the same as in Figure 3. The thinner (black) lines are the same in both panels, and correspond to the subjects in the first block without decision aid. The thicker (colored) lines in the top panel correspond to subjects in the first block with the decision aid. The thicker (colored) lines in the bottom panel correspond to subjects in the second block with the decision aid.

Figure 6 clearly shows that the distribution of $\bar{\varepsilon}_i$ in the population depends on the availability of the decision aid: the behavior of more subjects is consistent with RAEU when the subjects have access to the decision aid in that the errors needed are closer to zero. More precisely, the distribution of $\bar{\varepsilon}_i$ without the decision aid first-order stochastically dominates the distribution with the aid, and the median $\bar{\varepsilon}_i$ drops from 4.70 to 1.64.

While the decision aid clearly improves the fit of RAEU in the population on average, the key take-away from Figure 6 is that even with the decision aid and the benefit of experience, only about half of the subjects can be captured by RAEU with only a small error (near zero). The question of what constitutes a small error is philosophical. We consider errors over 5 to be large in that an $\varepsilon_{i,t}$ so large (in absolute magnitude) justifies *any* entry behavior in the most informative half of our (v,f) conditions – the half along the diagonal of the (v,f) space where some but not all of our subjects enter. Table 3 shows the maximum achievable expected surplus, and highlights the absolute expected surpluses below 5.

Table 3: Expected Surplus of Entering the NYOP Store and Bidding Optimally Under Risk-Neutrality

Valuation	Bidding Fee				
	0	1	6	12	18
5	0.1	-0.9	-5.9	-11.9	-17.9
20	1.4	0.4	-4.6	-10.6	-16.6
35	4.4	3.4	-1.6	-7.6	-13.6
50	8.9	7.9	2.9	-3.1	-9.1
65	15.1	14.1	9.1	3.1	-2.9

We conclude that RAEU cannot provide a good fit to the behavior we observe for a significant portion of our subjects: in the baseline condition without decision aid, 46 percent of subjects have $\bar{\varepsilon}_i > 5$, a proportion that reduces to 26 percent among subjects with the decision aid, and to 13 percent among subjects with both decision aid and experience. More research is needed to model the behavior of these subjects adequately. It is already clear from our results that preferences are wildly heterogeneous, and any future work will need to account for the differences in preferences across subjects.

8. General Discussion

We test the profitability of two-part tariffs in an NYOP setting. Instead of providing the bidding opportunity free of charge, an NYOP seller who uses a two-part tariff charges a non-refundable fee for the opportunity to bid. 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 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 subjects enter excessively because they find the joint optimization of their entry and bidding decisions difficult, we designed a tool to help 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. In the condition without the decision aid, buyer experience alone also reduces the profitability of two-part tariffs. We show that experience works differently from the decision aid in that experience alone merely teaches subjects not to enter when the fee is large, whereas the decision aid reduces entry whenever (*valuation, fee*) combination implies a marginal expected surplus.

In addition to documenting and contrasting the effects of decision difficulty and experience on entry and bidding in a NYOP setting, we also provide non-parametric evidence that the canonical model for decision making under uncertainty is not a good fit to the behavior of a substantial proportion of our subjects. Specifically, we derive a lower bound on the error magnitude in an individual-level random utility model of entry, with the deterministic part of the utility following risk-averse expected utility

theory. While our decision aid reduces the average (in the population) magnitude of errors needed to justify observed entry decisions under this model, there remains a substantial proportion of subjects whose behavior can only be rationalized under risk-averse expected utility theory with unreasonably large errors.

In summary, we find that two-part tariffs can be profitable for NYOP sellers, but only when (1) the potential buyers do not have access to our decision aid and (2) when they do not have too much experience with our laboratory NYOP market. Novelty of NYOP two-part tariffs and cognitive difficulty associated with buyers' joint optimization of their entry and bidding decisions are thus catalysts of two-part-tariff profitability in NYOP. We do not find evidence of fee aversion. In terms of implications for modeling, we reject risk-neutrality and risk-averse expected utility theory for a substantial segment of buyers.

We extend previous findings of excessive entry into 1PSB 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 is at least partly driven by the cognitive difficulty of trading off the probability of winning and the surplus contingent on winning.

In a post-hoc analysis, we also examined individual differences that correlate with excessive entry, defined as entering when a risk-neutral agent should not. Table A1 in the Appendix reports a linear regression of each subject's observed probability of excess entry 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 that correlate significantly: gender and the Holt-Laury risk-aversion measure. The sign of 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 A1 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.

Our findings also have important practical implications. Although assuming buyers welcome the opportunity to influence their transaction price is reasonable, coming up with a numerical bid amount is an involving and complex task. Consumers may find using the NYOP channel difficult, wishing to avoid the (perceived) hassle costs of taking part in the bidding process. To ease this task and help prospective buyers generate appropriate bids, sellers can add a decision aid to their NYOP bid-elicitation interface, a practice already employed by some NYOP sellers (e.g., Greentoe.com). Our results provide evidence that sellers charging no fee can follow this strategy of making the NYOP pricing format potentially more attractive to their customers without hurting their profits. However, we also show that sellers interested in increasing their profitability via a two-part tariff should not offer a decision aid. Finally, because additional buyer experience with the task reduces profitability, our results predict that short-run field experiments with two-part tariffs may overestimate the long-run profitability of the mechanism.

We acknowledge some limitations that provide avenues for future research. Our laboratory analysis is well suited to identify causal effects and to understand why—as well as how—the profitability of two-part tariffs in NYOP markets comes about. However, it does not tell us much about the magnitude of the observed effects in real markets. Therefore, additional field experiments would be very interesting. In addition, we have examined only one form of a decision aid. Decision aids can be designed in many different ways, which may moderate the effect on the profitability of two-part tariffs in an NYOP setting.

References

- 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.
- 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.
- Bernoulli, Daniel (1738); translated by Dr. Louise Sommer. (January 1954). "Exposition of a New Theory on the Measurement of Risk". *Econometrica* 22 (1): 22–36.
- 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.
- 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.
- Ertuç, 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 (2009), "Competitive reasons for the name-your-own-price channel," *Marketing Letters* 20 (3), 277-293
- 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-Ozbay, Emel and Erkut Y. Ozbay (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.
- Häubl, Gerald and Valerie Trifts (2000), "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids," *Marketing Science*, 19 (1), 4-21.
- 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.
- 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.

- 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, Vernon L. (1976), "Experimental Economics. Induced Value Theory," *American Economic Review*, 66 (2), 274–279.
- 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.
- Steiner, Ina (2018), "eBay Heralds Name Your Price Shopping at Recent Conference", <https://www.ecommercebytes.com/C/blog/blog.pl?pl/2018/8/1533516748.html>
- 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.
- von Neumann, John and Morgenstern, Oskar (1944). *Theory of Games and Economic Behavior* (Third ed.). Princeton, NJ: Princeton University Press.
- 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

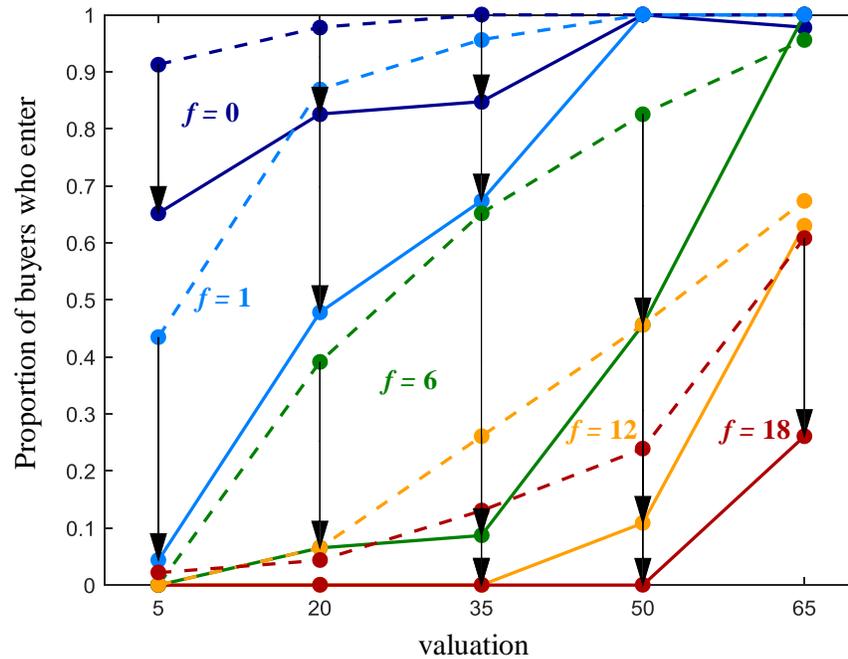
Figure A1: Experimental Interface of Experiment 1

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 following options.		
<p style="text-align: center;">Store A</p> <p>The bidding fee is: xx Please enter your bid: <input type="text"/></p> <p>In case your bid is not successful, you can still purchase from store B.</p> <p style="text-align: center;">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;">Buy</p>	<p style="text-align: center;">Don't Buy On This Period</p> <p>In case you don't want to buy on this period, please press the button "Don't Buy" below.</p> <p style="text-align: center;">Don't Buy</p>

Figure A2: Experimental Interface of Experiment 2

Period x of x	
<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 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;">Submit Bid</p>	<p style="text-align: center;">Don't Buy On This Period</p> <p>In case you don't want to buy on this period, please press the button "Don't buy" below.</p> <p style="text-align: center;">Don't Buy</p>

Figure A3: 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 have the decision aid (second block); the dashed lines do the same thing when the subjects do not have the decision aid (first block). Arrows indicate significant differences at the 5% level.

Table A1: 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: $R^2=0.196$. A unit of observation is one subject. The dependent variable is the probability of entering in the 13 cells of the (v,f) design that are shaded in Table 1 and involve $v < 70$. 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 of 70, matching Experiment 2's (v,f) design. Risk aversion is measured as in Holt and Laury (2002), loss aversion as in Gächter et al. (2010), subjective risk-taking as in Dohmen et al. (2012), namely, "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*), 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 agreement scale. All four scales are standardized within the entire population, and hence are measured in population standard deviations. All other variables are dummies. Number of observations = 185 (93 subjects from Experiment 1, 92 from Experiment 2. Specifically, 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).