Consumer Behavior in Bidding Markets with Participation Costs

Robert Zeithammer Lucas Stich Martin Spann Gerald Häubl

October 21, 2019

Abstract: Consumer bidding is common in a wide variety of markets. An important source of friction in many markets with bidding is the cost of participation. We investigate the impact of participation costs on bidder entry and bidding behavior using incentive-compatible laboratory experiments with participation costs in the form of "bidding fees." Instead of a standard auction, our experiments utilize the more tractable name-your-own-price (NYOP) setting. We find that both entry frequencies and bid magnitudes exceed the predictions of a risk-neutral benchmark model. Although the bids we observe can be rationalized with a standard risk-averse expected utility model, the entry decisions cannot, because more than half of our subjects enter when a risk-neutral bidder would not. We identify two important moderators of such "excess entry"—experience with the task and reduced task difficulty due to a decision aid. In contrast to their effect on entry decisions, experience and difficulty have no effect on bid magnitude, which suggests figuring out the bid amount is easier than deciding whether to enter. We also apply our data to examine the profitability of two-part tariffs in NYOP settings, and we find such a pricing strategy can only be profitable when consumers are relatively inexperienced and when they do not have access to an effective decision support tool.

Keywords: Pricing, Auctions, Behavioral Economics, Laboratory Experiment

Contact: Robert Zeithammer, rzeitham@ucla.edu

Auctions and other participative pricing mechanisms, such as Name Your Own Price (NYOP) or Pay What You Want (PWYW), are widely used for pricing both goods and services in consumer markets (Spann et al., 2018; Haruvy and Popkowski Leszczyc, 2018). A unifying feature of participative pricing mechanisms is bidding, that is, the submission of a binding price offer by a prospective buyer. An important source of friction present in many markets with bidding is the cost of participation. Prior literature has identified several sources of such participation costs: the mental and physical effort associated with bid preparation (e.g., Samuelson, 1985; Krasnokutskaya and Seim, 2011), the hassle of submitting the bid and waiting for the outcome (e.g., Hann and Terwiesch, 2003; Fay, 2009), and the accounting cost arising from various fees and commissions charged by marketplace intermediaries (e.g., Palfrey and Pevnitskaya, 2008; Bernhardt and Spann, 2010; Moreno and Wooders, 2011).

Understanding consumer behavior in the presence of participation costs is important to both market designers and consumer researchers. In this paper, we examine the impact of participation costs on consumer behavior using incentive-compatible laboratory experiments that measure both key components of consumer behavior in bidding markets—whether to submit a bid and, if so, the bid amount. We show that considering both of these behavioral components provides important new insights about bidder preferences compared to considering submitted bids alone. In contrast to prior experimental work on bidding with participation costs, our experiments use the more tractable name-your-own-price setting, which we motivate next.

NYOP is 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 goods and services.[1] Conceptually,

[1] Examples include event tickets at scorebig.com, electronics at greentoe.com, a variety of products on ebay.com under its "Best Offer" mechanism, and software from ashampoo.com.

NYOP is a single-bidder first-price sealed-bid auction with a hidden reserve price. Instead of having to compete with other bidders as would be the case in a typical auction, an NYOP bidder submits a binding price offer to the seller, and the seller's hidden reserve price determines whether the bid is accepted. As a single-agent decision problem, NYOP is thus an ideally tractable mechanism to gain insight into consumer decision-making in markets with bidding and costly entry, because the analysis does not need to consider strategic interaction among multiple endogenously determined bidders as it would in a multi-bidder auction with entry costs analyzed by Samuelson (1985) or Menezes and Monteiro (2000).

In our laboratory implementation of NYOP with participation costs, subjects are buyers with induced valuations, allowing us to vary and control buyer valuation of the object sold while abstracting away from specific product categories (Smith, 1976). The NYOP seller is a computer program that faces a privately known procurement cost randomized across experimental rounds, and accepts bids above its current draw of this cost (i.e., it sets the hidden reserve to its cost). The stylized marketplace consists of the NYOP seller and a premium posted-price seller who sells the same product for a higher price than the NYOP seller's hidden threshold. To participate in bidding, the buyers face a participation cost composed of both the inherent hassle cost of bidding and a non-refundable "bidding fee" paid upon entry to the NYOP store. We vary both the valuations and the bidding fees within subject using a full-factorial experimental design. The optimal NYOP entry and bidding strategy in our stylized marketplace has been analyzed by Spann et al. (2010) under the theoretically tractable demand-side assumptions of buyer ambivalence to bidding fees, risk neutrality, and perfect computational ability. We present the results of two experiments that were designed to examine the implications for consumer behavior of three behaviorally realistic departures from these assumptions: Bidders may be averse to paying for the opportunity to bid as a matter of principle, may not be risk neutral, and may have difficulty making the entry and bidding decisions. Next, we elaborate on these three behavioral deviations that guided our experimental designs, and briefly summarize our key findings.

Instead of carefully computing the expected utility of participation and weighing it against the bidding fee in each task, consumers may utilize a decision rule to never pay for anything other than the product itself (Amir and Ariely, 2007)—a property of preferences we refer to as "fee aversion." We use very small bidding fees in combination with high valuations to test for fee aversion in bidder behavior, and we find no evidence of it. High-valuation bidders submit bids just as often when the bidding opportunity is free as when it costs a small amount.

Buyers may be risk averse, and hence rationally unwilling to pay a fee for a mere chance to save money through NYOP bidding. We find surprisingly mixed evidence of buyer risk aversion. On one hand, our buyers bid as if they were risk averse (i.e., above the risk-neutral benchmark), and risk aversion is at least partially driving bids (subjects with higher risk-aversion—measured in a separate task—bid more, ceteris paribus). On the other hand, our buyers are willing to pay the bidding fees more often than the tractable risk-neutral benchmark predicts. Thus, their entry behavior is not consistent with even mild risk aversion. As we discuss in the next section, such excess entry is a common finding in prior empirical work on auctions. In addition to extending it to the simpler NYOP setting (and thus eliminating potential explanations that involve strategic considerations among multiple bidders), we identify two strong moderators of excess entry that jointly link the phenomenon to the cognitive difficulty associated with making entry and bidding decisions. We discuss these two moderators next.

First, we find that experience with the task reduces subjects' excess entry while leaving their bids unaffected. This pattern of experience effects is consistent with at least two explanations:

(1) Bidders might find the task of evaluating the expected benefit of bidding difficult (as proposed by Engelbrecht-Wiggans and Katok, 2005) and enter more initially in order to explore and learn, or (2) bidders might find the task novel and interesting initially, but enter less as it gradually becomes more of a routine task. To differentiate between the two interpretations of the moderating effect of experience on excess entry, we designed a decision aid to simplify the task for some subjects without imposing any particular preferences. Our decision aid takes a potential bid as input and displays the two key components of expected utility from bidding: the probability of the bid being accepted by the seller, and the surplus due to the bidder upon bid acceptance. The decision aid thus provides the key information to evaluate the expected benefit of bidding and should therefore compensate for the cognitive limitations of human bidders with assessing the payoff consequences of their bids.

We find our decision aid reduces excess entry and moderates the impact of experience on it, so at least a part of the experience effect on excess entry we find is due to bidder learning. However, experience and the decision aid do not reduce excess entry the same way: We find that mere experience (in the absence of the decision aid) affects entry behavior more crudely than the decision aid: whereas the decision aid reduces entry in all valuation-fee combinations that likely involve an expected utility near zero, mere experience seems to teach subjects to avoid high fees. The latter pattern is also inconsistent with the novelty explanation – early rounds of the experiment are novel regardless of the fee amount.

Our results are also inconsistent with several additional alternative explanations of excess entry: We find that bids are unaffected by the decision aid, which rules out explanations based on misperception of probabilities. The fact that bids are systematically affected by bidding fees rules out explanations based on two-stage heuristic thinking with the entry cost sunk at the time of bidding. Finally, the fact that bidders with very low valuations do not overbid rules out explanations based on additive joy of winning entering the utility function.

The effect of the decision aid is managerially interesting in its own right because such tools are commonly deployed in online shopping environments to assist consumers in making their purchase decisions (Häubl and Trifts, 2000) and are indeed already offered by some NYOP sellers.[2] Although decision aids can clarify bidders' expectations about their chances and thus increase NYOP sellers' profits (Fay and Lee, 2015), our results also reveal that sellers may have a strategic incentive to prevent buyers from having complete information about their chances and payoffs, because such information reduces buyer entry. Our findings also have immediate managerial implications for the profitability of two-part tariffs in NYOP settings. 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.[3] In the absence of any friction, the same ability to extract consumer surplus that makes two-part tariffs profitable for a posted-price seller also operates within NYOP (Spann, Zeithammer, and Häubl, 2010). We find that a two-part tariff can be profitable for an NYOP seller, but only when the buyers have no experience with it and no decision aid to help them. When buyers are provided with a decision aid that compensates for their cognitive limitations and lack of experience, the profitability of using a two-part tariff vanishes. This finding may explain why real-world NYOP retailers do not charge bidding fees.

[2] For example, greentoe.com selling electronics and most major airlines selling seating upgrades use a tachometer-style decision aid to inform bidders about the chance of success of a bid.

[3] Note that our two-part tariff (i.e., fee + bid) is different from "bidding fees" in "all-pay" (or "penny") auctions (e.g., Platt, Brennan, and Tappen, 2013), which do not resemble two-part tariff schemes, but rather dynamic games among multiple bidders who all pay a small amount for every bid increment if they want to retain a chance of winning.

LITERATURE REVIEW

This research is related to two bodies of prior work: (1) the literature on NYOP selling and (2) the experimental auction literature. The majority of prior research on NYOP selling 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 examines 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 assumes risk-averse buyers. We contribute to the NYOP literature by documenting a systematic departure of buyer entry and bidding behavior from the risk-neutral predictions, and we argue that at least the entry patterns we observe are not consistent with a standard risk-averse model either. In our application, we also directly test Spann, Zeithammer, and Häubl's (2010, 2015) predictions that a two-part tariff can be profitable for an NYOP seller, and we document strong moderating roles of cognitive difficulty and experience 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, the current research 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. We are the first to propose the NYOP bidding paradigm with endogenous entry as a clearer empirical setting for studying the impact of preferences on consumer behavior by avoiding the need for both subjects and the analyst to understand the equilibrium of the entry and bidding game. Nevertheless, our results are consistent with two major findings of the empirical literature on 1PSB bidding: over-bidding and over-entry relative to a risk-neutral model. Over-bidding is one of the consistent findings in this literature, and thus a large body of work has focused on explaining that phenomenon (e.g., Cox, Roberson, and Smith, 1982; Cox, Smith, and Walker, 1988; Feng, Fay, and Sivakumar, 2016). The most common explanation has been risk aversion (Cox, Smith, and Walker, 1988; Filiz-Ozbay and Ozbay, 2007, Chakravarty et al. 2011). 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 that over-bidding can occur in the absence of equilibrium considerations and is not affected by the cognitive difficulty of the joint entry and bidding decisions.

In a closely related paper, Bajari and Hortaçsu (2005) use laboratory bidding data in 1PSB auctions with induced valuations to test which of four candidate structural models is best at recovering bidder valuations from observed bids. The risk-averse model performed best, contributing to the consensus explanation of over-bidding relative to the risk-neutral benchmark as first proposed by Cox, Smith, and Walker (1988). Unlike Bajari and Hortaçsu (2005), who study auctions with exogenous participation by multiple bidders, the current research endogenizes bidder entry to provide a more realistic characterization of the phenomenon. Although we do also find over-bidding relative to the risk-neutral predictions, the entry behavior of our buyers is not consistent with the dominant risk-averse 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 the entertainment value of bidding in an auction as an explanation of the phenomenon. Aycinena and Rentschler (2018) find over-entry relative to a risk-neutral benchmark in a variety of auction formats and information settings. Ertaç et al. (2011) propose a model that combines risk aversion and the joy of winning to explain over-entry. 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. By contrast, we at least partially explain over-entry as a result of cognitive difficulty of the decision problem involved in assessing the expected utility from bidding and comparing it to the cost of entry. We find no evidence of an additive joy of winning, because even low-valuation bidders bid within their meager 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. Since we find that cognitive difficulty does not impact bids, we can also rule out explanations based on biased probability perceptions. Finally, the over-entry we find is by construction of our NYOP task not a result of overconfidence in a strategic downstream game as in Camerer and Lovallo (1999).

Another closely related paper is Davis, Katok, and Kwasnica (2014), who study a two-bidder ascending auction with entry costs, and also find buyer entry into auctions deviates from theoretical predictions. In contrast to our setup, their buyers only learn valuations after entry, and risk aversion may result in mixed-strategy equilibria. In summary, our results are consistent with the prior literature on over-entry, and add additional insights about its underlying causes.

MODEL

We model the market as follows: A buyer wants to buy a specific indivisible object. The object is available in an outside posted-price spot market for a commonly known price p. One NYOP seller also exists, who can procure the same object[4] for a varying cost c ex-ante distributed uniformly on $[0,p]$, and unknown to the buyer.[5] The buyer can buy from the posted price spot market or make a binding offer to the NYOP seller by submitting a bid b for it. If the bid is rejected, the buyers can still buy from the spot market. Buyer participation in NYOP bidding is costly for two reasons: First, the marketplace charges a non-refundable bidding fee $f \geq 0$ to participate. Second, the buyer may also experience mental and hassle costs of entry, the total of which we denote h.

We assume throughout and implement in the experiment that the NYOP seller accepts bids above her cost c, because he lacks commitment to any other bid-acceptance policy. Thus, the buyer (correctly) believes her chance of a bid b being accepted is $\Pr\left(b \text{ acccepted}\right) = \Pr\left(c < b\right) = b / p$.

[4] Whereas some NYOP sellers (e.g., Priceline) make the product opaque, others (e.g., greentoe.com) do not. Opacity is thus not an intrinsic feature of NYOP selling, just as transparency is not an intrinsic feature of posted-price selling (e.g., hotwire.com sells opaque products for posted prices). Our model assumes the product sold by the NYOP seller is as opaque as that sold by the outside posted-price market. For example, the NYOP seller is priceline.com and the outside posted-price market is hotwire.com. Or the NYOP seller is greentoe.com and the outside posted-price market is amazon.com.

[5] One way to motivate this assumption is that the NYOP seller is a retailer with access to a wholesale market, so the posted price charged to consumers naturally bounds the wholesale cost from above. The uniform distribution is a simplification for expository and tractability purposes.

Throughout our laboratory experiments, the seller is hardwired as a computer program and the buyers are educated about the bid-acceptance probability. Therefore, we abstract from potential seller real-world behavior that might deviate from the following simple model of bid acceptance.

Buyer i's utility u_i of buying the object depends only on her surplus s, that is, her valuation v minus the total cost incurred, and WLOG $u_i(0)=0$ and $u_i(1)=1$, which is consistent with standard expected utility theory.[6] We assume buying from the posted-price spot market involves no hassle cost, so buying from it yields a utility of $u_i(v-p)$. By contrast, buying from the NYOP seller who accepted a bid b yields a utility of $u_i(v-f-h-b)$, and paying the fee but getting the bid rejected yields $u_i(\max(0,v-p)-f-h)$ to account for the second chance to buy from the posted-price market. Given these preferences and the above probability of bid acceptance, buyer i solves the following bidding problem:

$$bid(v,f)=\underset{b\geq 0}{\arg\max}\left(\frac{b}{p}\right)u_i(v-f-h-b)+\left(1-\frac{b}{p}\right)u_i(\max(0,v-p)-f-h) \tag{1}$$

The buyer takes the NYOP chance to save money whenever her expected utility of doing so exceeds $u_i(0)=0$, and stays out of the NYOP store otherwise, yielding this entry problem:

$$\max_{enter,not}\left\{\begin{matrix}\left(\frac{bid(v,f)}{p}\right)u_i(v-f-h-bid(v,f))+ \\ +\left(1-\frac{bid(v,f)}{p}\right)u_i(\max(0,v-p)-f-h), & \underbrace{0}_{u_i(\text{not enter})}\end{matrix}\right\} \tag{2}$$

If their bids are rejected, buyers with $v>p$ buy from the posted-price market. Buyers with lower

---

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

valuations behave as if the posted-price option did not exist because they cannot afford it.

Our implementation of a marketplace with an NYOP seller is motivated by Spann, Zeithammer, and Häubl (2010), who consider risk-neutral buyers $\left(\text{with utility functions } u_i(s)=s\right)$ and valuations drawn from a uniform distribution on [0,M] for some maximum valuation $M > p$. Risk 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.

In terms of the model introduced above, we do not attempt to estimate the precise shape of $u_i(s)$ in our subject population. Instead, we measure deviations in entry and bidding behavior from the benchmark $u_i(s)=s$ model, and interpret these deviations as either consistent or inconsistent with the dominant alternative hypothesis that $u_i(s)$ are concave, and so the buyers are risk averse. Bidder risk aversion has been the consensus explanation of over-bidding in 1PSB auction settings (e.g., Cox, Smith, and Walker, 1988; Filiz-Ozbay and Ozbay, 2007; Bajari and Hortaçsu, 2005), so testing for it in a richer context with endogenous entry is important. The directional predictions of an alternative model with heterogeneous risk aversion are straightforward: The bids should exceed the risk-neutral level of $\min(p/2, v/2)$, and buyers in a given (v, f) condition should not enter when risk-neutral buyers in the same condition would not. Note these predictions for behavior in NYOP settings do not necessarily hold in a multi-bidder auction setting, where the natural assumption of heterogeneity in risk aversion implies very complicated equilibrium behavior due to the resulting asymmetries among bidders.

A fee-averse bidder who exhibits a knee-jerk negative reaction to any positive fee may be well captured by a $u_i\left(s-\alpha\mathbf{1}(f>0)\right)$, with a positive *α* large enough to overwhelm even high *v*. Such a buyer would participate in zero-fee NYOP bidding but stay out in the face of a small fee even with a high *v*. She should also be insensitive to the magnitude of the fee.

EXPERIMENTAL PARADIGM: WITHIN-SUBJECT MANIPULATION OF BUYER'S VALUATION AND SELLER'S BIDDING FEE

Both our experiments share several design elements: induced buyer valuations, a stylized marketplace motivated by the model, and within-subject manipulation of both buyer valuations and NYOP bidding fees. In this section, we describe and justify all of these 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-subject 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 induce the individual buyer valuations in different rounds from several equispaced discrete points, {5, 20, 35, 50, 65, 80, 95} in Experiment 1 and {5, 20, 35, 50, 65} in Experiment 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; 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 experiments). The other store in the marketplace uses NYOP selling, procures the object for a privately known procurement cost c, and accepts all buyer bids above cost. This bid-acceptance strategy is thoroughly explained to the subjects upfront, along with the fact that the seller's cost c is distributed *iid* uniformly between rounds on $[0,p]$.[7] 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 measurement of entry behavior (see Figure A1 in Appendix A for the screen layout; the Web Appendix contains 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 NYOP bidding fee and the random procurement cost c. We vary the fee among a few discrete points between 0, which everyone should be willing to pay, and 18, which no risk-neutral buyer with a valuation below 100 should be willing to pay when p=70. We include

[7] 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.

three theoretically motivated intermediate levels between 0 and 18 (selecting only a few specific values allows for a clean full-factorial within-subject design). The first intermediate fee level we use is 1—the smallest possible positive fee given our experimental currency. As discussed in the previous section, reduced entry by high-valuation buyers as the fee rises from 0 to 1 would be evidence of fee aversion. The second intermediate fee level we use is 12—the level approximately optimal under risk neutrality when p=70 and M=100.[8] Overwhelming evidence suggests people are risk averse when they bid on products, which should reduce the optimal bidding fee, so we add the fee level of 6 as an intermediate, non-negligible value. In summary, we test the following 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.

EXPERIMENT 1: TEST OF THE RISK-NEUTRAL BENCHMARK MODEL

In Experiment 1, we set M=100 and p=70, 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(*fee levels*) × 7(*valuation levels*) within-subject design, measuring each subject's entry and bidding behavior in 35 conditions. We now explain the experimental procedure in detail.

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 Appendix A for the screen layout).

---

[8] Spann et al. (2010) show that as long as $p > 4M/7$, the optimal bidding fee is $f^*(p) = 4M^2/49p$.

To decide whether a bid is accepted, we draw the secret threshold price c 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 that 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 subjective risk-taker scale by Dohmen et al. (2012), a question about attitudes toward bidding fees, the frequency of participation in lab experiments to date, and demographics. None of the scales were separately incentivized beyond a flat payment.

Data Collection

We conducted four sessions of the experiment at the laboratory of a large European public university. Subjects were mainly undergraduate students 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. The following results are based on the 91 remaining subjects, and their behavior in the $5 \times 7=35$ rounds involving all possible combinations of valuations and fees presented in random order.

Results: Entry and Bidding

Table 1 lists the percentage of subjects who enter under the different valuation-fee conditions (recall that each subject experienced each condition exactly once). The shaded cells delineate the conditions under which risk-neutral buyers should not enter. Our main finding evident from Table 1 is that the risk-neutral model dramatically underpredicts entry. The discrepancy is the starkest at f=18, where the risk-neutral model predicts no entry, but about half of our subjects enter when their valuations are relatively high.

Result 1 (Excess entry): *At least half of our subjects enter when a risk-neutral buyer should not, and so a risk-averse buyer should not enter either.*

Table 1: Proportion of Subjects Who Enter the NYOP Store

	Bidding Fee				
Valuation	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%

The first two columns of Table 1 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 ($v \leq 20$) when such a decline can be rationalized by sufficient risk aversion. By contrast, a "knee-jerk" fee aversion would predict a decline in entry irrespective of valuation. Note that our evidence against fee aversion does not amount to small fees not mattering at all—as the fee increases from zero to 1, we observe a dramatic reduction in entry when the valuation is at its lowest level.

Result 2 (No fee aversion): *We find no evidence of fee aversion, because the entry behavior of high-valuation bidders is unaffected by a small increase of the bidding fee from zero.*

Having highlighted the two main findings about entry behavior, we now briefly comment on two other surprising patterns in Table 1. The top left corner of Table 1 suggests fewer than 10% of our subjects experience hassle costs h that are large enough to make entry unattractive even when the fee is zero. For most of these subjects, the implied hassle cost is small because they do enter when their valuation rises. To get a sense of the hassle costs required to explain this behavior, consider the risk-neutral model: Under risk neutrality, a hassle cost greater than 0.09 is sufficient to make entry unattractive with a valuation of 5, and the corresponding bound rises to 1.43 and 4.38 when valuations increase to 20 and 35, respectively.[9]

Small hassle costs cannot explain the second surprising pattern visible in the lower left corner of Table 1. We find that 14 (15%) of our subjects simply purchase at the posted price instead of submitting a bid below p=70 and thus effectively gambling for a bigger surplus. This behavior cannot be explained by small hassle costs, because the payoff from such gambling is very large. Risk aversion also cannot explain the behavior when f=0: Risk-averse bidders may want to bid more than risk-neutral ones, but they should all enter. We suspect this under-entry arises from an individual-level preference for paying a fixed price instead of bidding whenever the posted price is low-enough relative to valuation, akin to the analogous effect in Gneezy et al. (2012). The behavior seems to be a consequence of individual differences because of the 14 subjects who go straight to the posted-price store when v=95 and the NYOP fee is zero, 11 also forego bidding when the NYOP bidding fee increases to 1 but the valuation stays at 95. We also find no effect of

[9] Based on equation 2, the minimum hassle-cost needed to explain non-entry solves $h = \max_{b \geq 0} \left(b/p\right)\left(v-b\right) = v^2/\left(4p\right)$.

experience on entry in the relevant cells of Table 1: The pattern looks the same among subjects who experienced those cells early in the experiment and subjects who experienced them late. All of these observations are consistent with some subjects finding the evaluation of the bidding opportunity difficult, and choosing not to even explore it when the posted-price alternative presents an easy risk-free gain. We leave a definitive explanation of the underlying causes of this "bid avoidance" for future work.

Having discussed the entry behavior, we now turn to the magnitude of submitted bids. Table 2 shows the results of the following linear regression of bid by subject *i* when the subject had the induced valuation *v* and faced a bidding fee of *f*:

$$bid_{i,v,f} = \alpha_i + \beta \min\left(\frac{v}{2}, 35\right) + \gamma_f dummy(f) + \delta\, round_{i,v,f} + \varepsilon_{i,v,f}\,, \tag{3}$$

where $\min\left(\frac{v}{2}, 35\right)$ is the optimal bidding strategy under the risk-neutral model,[10] called the "risk-neutral bid" in Table 2, and $round_{i,v,f}$ is the round of the experiment in which subject *i* had the induced valuation *v* and faced a bidding fee of *f*. The fee dummies flexibly control for the average effect of bidding fees on bidding.

Because the coefficient β of the risk-neutral bid is significantly and substantially greater than 1, we conclude that buyers submit higher bids upon entry than the risk-neutral model predicts. Figure 1 illustrates this result by plotting both the risk-neutral prediction and a boxplot of the observed bids in the condition 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 1 reveals

[10] A risk-neutral bidder maximizes expected surplus to find $\dfrac{\min(v,p)}{2} = \arg\max_b \left(\dfrac{b}{p}\right)(v-b) + \left(1-\dfrac{b}{p}\right)(v-p)$.

that at all valuation levels other than the lowest one, bids exceed the risk-neutral prediction. The difference is large in that the entire inter-quartile range lies above the predicted level.

Table 2: Linear Regressions of Observed Bids on Risk-Neutral Predictions

	Basic model		**Fixed effects with clustered standard errors**	
Variable	Estimate	t-stat	Estimate	t-stat
Risk-neutral bid	**1.32**	**64.99**	**1.32**	**39.14**
Fee = 1	0.79	1.56	**0.82**	**2.16**
Fee = 6	-0.30	-0.53	0.32	0.68
Fee = 12	**-2.60**	**-3.81**	-1.47	-1.93
Fee = 18	**-5.01**	**-6.57**	**-3.45**	**-3.57**
Round	0.01	0.44	0.01	0.34
Constant	**-2.23**	**-3.23**	**-2.65**	**-2.92**
Subject FEs	-		**included**	
Number of observations	1907		1907	
R^2	0.71		0.78	

Note to table: The t-statistic of the coefficient on the risk-neutral bid in the model with fixed effects and clustered standard errors is the smallest among several other permutations of the model specification: without clustering, with robust standard errors, and without fixed effects.

Another qualitative prediction of the risk-neutral model that can be tested with our data is that the fee paid should not influence bidding, because it is effectively sunk. To see why the fee is sunk in the risk-neutral model, let $u_i(s) = s$ in equation 1, and observe that the bidding problem becomes additively separable from the fee paid.[11] In contrast to this prediction, Table 2 clearly shows that bids exhibit an inverse-U shape as a function of fee magnitude, first rising when fee increases from 0 to 1, and then gradually falling as the fee increases further. The finding of bids changing as a function of the fee level also rules out other models that assume fees to be sunk. For example, Smith and Levin (1996) propose a model with risk-averse preferences that additively

[11] In words, the argument goes as follows: Because the fee enters the objective function of risk-neutral bidders (i.e., the expected surplus) additively, considering the bidding problem as "first figure out how much to bid, and then see if the fee is small enough" is identical to "figure out how much to bid given that an extra fee is subtracted from the payoff when I win, and the same fee is also subtracted when I lose."

separates the fee f from the bidding surplus s in the buyer's utility function: $U\left(f, s\right) = s^r - f^r$. The relationship between fees and bids documented in Table 2 is not consistent with such a model. More generally, the relationship also rules out any heuristic two-stage model whereby potential buyers first decide whether to enter, and only then set the bid magnitude. Instead, our buyers seem to consider the fee in formulating their bidding strategy.

Figure 1: Observed Bids When Bidding Fee Is Zero

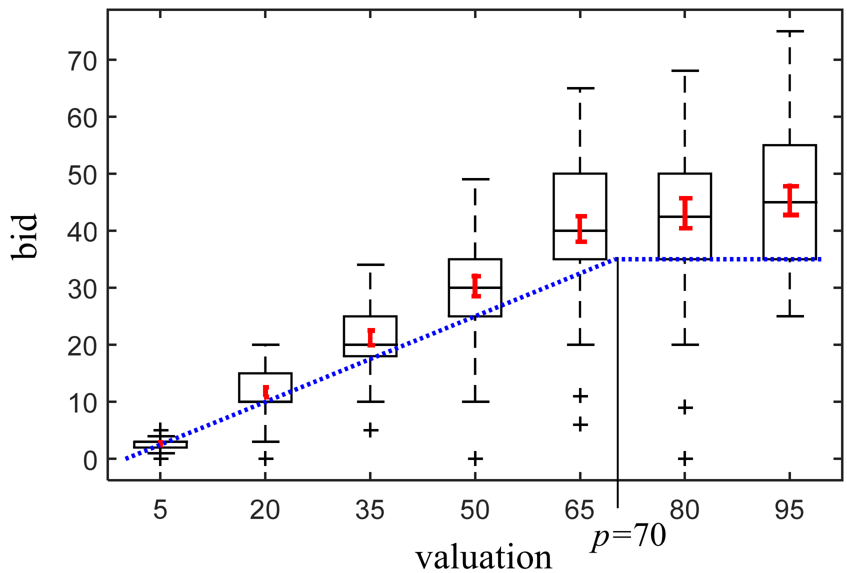

Note to figure: The boxplots illustrate the distribution of bids at each valuation level. The thicker (red) error bars represent 95% confidence intervals. The dotted (blue) line shows the optimal bidding function by risk-neutral buyers. The case of v=5 is difficult to discern from the figure—the risk-neutral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

We summarize our comparisons of the observed bids with the risk-neutral prediction as follows:

Result 3 (Over-bidding and impact of fees on bids): *Buyers submit higher bids upon entry than the risk-neutral model predicts. On average, bidders bid 32% higher than the risk-neutral model predicts. Buyers consider the bidding fee in formulating their bidding strategy, which is inconsistent with any model in which fees are effectively sunk, including the risk-neutral model.*

The over-bidding and the relationship with fees we find are both consistent with prior findings and easily rationalized: First, 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. Second, it can be shown that the concave relationship between bids and fees is consistent with at least some specifications of risk-averse utility, for example, the constant relative risk-aversion (CRRA) model $u_i(s) = s^r$.

To summarize the findings so far, our subjects enter as if they were risk seeking, but they bid as if they were risk averse. We conjecture that the excessive entry by otherwise risk-averse bidders may arise from the cognitive difficulty of first solving the optimal bidding problem and then correctly inducting backward to the entry decision. One piece of evidence supporting our conjecture is the effect of experience on excess entry. If subjects are unsure about the tradeoff between the fee and the benefit of bidding, they can enter a lot in the initial rounds to learn about the tradeoff, and then enter more judiciously in later rounds after gaining experience with the task. Because we randomized the order of the 35 different (v,f) combinations across subjects, we can interpret a regression of excess entry on round as evidence of experience.

Table 3 defines "excess entry" as entry when a risk-neutral buyer should not (the grey cells of Table 1), and documents a substantial downward effect of round on it via a logistic regression of excess entry on round, controlling for the within-subject (v,f) condition and subject fixed effects. The experience effect we find is large in that predicted excess entry probabilities are about half of their initial values by the last round of the experiment. For example, the coefficients of the basic regression (without individual fixed effects) imply the predicted probability of a buyer with v=5 facing f=1 entering the NYOP store declines from 48% in round 1 to 23% in round 35. The analogous excess entry reduction for a buyer with v=65 facing f=18 is from 64% to 36%.

Table 3: Effect of Experience on Excess Entry and Bid Magnitude

	Logistic regression of entry when a risk-neutral buyer should not enter				**Linear regression of observed bid conditional on entry (R^2=0.71)**			
Variable	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat
Round	**-0.033**	**-4.39**	**-0.053**	**-5.34**	-0.00093	-0.05	-0.00281	-0.11
Constant	**0.711**	**2.58**	1.260	1.57	**8.612**	**2.26**	**43.06**	**35.99**
Subject FEs	-		**included**		-		**included**	
Controls	15 dummy variables, one for each gray cell in Table 1 (coefficient estimates not reported)				31 dummy variables, one for each cell of Table 1 where at least one bidder entered (coefficient estimates not reported)			
Observations	1365 = 91 subjects x 15 design cells				1907 submitted bids			

Note to table: Both regressions include all person-round observations after the training rounds. The linear regression includes only the person-round observations in which the person entered. In the linear regression, standard errors are clustered at the individual level (91 subjects in each regression).

Table 3 also shows the absence of any effect of round on the magnitude of observed bids. If buyers are learning over time, they are learning to enter less whenever a risk-neutral buyer would not. However, they do not seem to be learning how much to bid upon entry, at least not in a way that would result in a monotonic shift of bid magnitude. The results can be summarized as follows:

Result 4 (Experience reduces excess entry but does not affect bids): *As subjects gain experience with the task during the multiple rounds of the experiment, they reduce their excess entry (i.e., entry when a risk-neutral buyer would not enter), but they do not systematically change the bids they submit.*

A moderating role of experience on entry does not necessarily imply our subjects are learning to solve the difficult entry decision. An alternative explanation of the experience effect could be the novelty of NYOP bidding wearing off over time. In this case, subjects would simply enter less later on, because NYOP bidding becomes less novel over time and therefore less interesting for them. To distinguish between these two potential explanations of excess entry, as well as to test the task-difficulty explanation directly, we designed a decision aid to help subjects assess the payoff consequences of their potential bids. Of the two alternative explanations, only the learning explanation predicts that making the task easier should moderate the experience effect. Experiment 2 described next examined the effect of such a decision aid on behavior.

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) shown in equation 1. Deciding whether to enter the NYOP store for a positive fee involves comparing the expected utility of the payoff from bidding to the (dis)utility of paying a bidding fee and not getting anything in return as shown in equation 2. Engelbrecht-Wiggans and Katok (2005) found the entry decision can be difficult for people. In Experiment 2, we examine whether the excessive entry and experience effects discovered in Experiment 1 can be attributed to the cognitive difficulty of both bid optimization (equation 1) and the entry decision (equation 2) 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 in bidding: (1) the probability of acceptance and (2) the contingent monetary payoff if the bid is accepted. Note this information is not new to the subjects, and it does not recommend any particular decision. Instead, it merely helps with calculations needed to appreciate the consequences of a given bid amount. It provides the inputs to equation 1, but does not combine them into expected utility (unknown to us, by construction) or any other objective function. Displaying the two aspects of the bidding tradeoff is a natural choice of a decision aid - other researchers have used the same idea to reduce the cognitive difficulty of their task (e.g., Chakravarty et al. 2011).

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 activates the decision aid, which instantaneously displays both the probability of acceptance and the contingent monetary payoff for that bid amount (see Figure A2 in Appendix A for the general layout 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 employed a balanced crossover design whereby half the subjects started with the decision aid and ended without it, and the other half, vice versa. Each subject thus experienced two blocks of the same 25 (*valuation,fee*) combinations presented in random order within the block as in Experiment 1. The only difference relative to Experiment 1 was that we simplified the situation by using only valuations below the posted-market price, i.e. only {5, 20, 35, 50, 65}, resulting in 25 = 5(*fee*) × 5(*valuation*) combinations per block. 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). We now turn to the results.

Comparing the first-block behavior of subjects with the decision aid with the behavior of subjects without it yields a clean measurement of the decision aid's effect. Table 4 shows the effect of the decision aid on entry and bidding, while controlling for the effect of experience discovered in Experiment 1. It is immediate that the decision aid reduces entry on average and strongly moderates the effect of experience (round). By contrast, the decision aid has no effect on bid magnitude and (replicating Experiment 1) neither does experience. The fact that the decision aid does not affect bid magnitude rules out all explanations of excess entry based on biased probability perceptions: If subjects without the aid entered too often because they overestimated their chances of winning compared to subjects with the aid, they would also bid less.

Table 4: Effects of Decision Aid and Experience on Entry and Bidding

	Logistic regression of entry (N=2300)				**Linear regression of observed bid (N=1217, R^2=0.81)**			
	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat
Round	**-0.045**	**-4.665**	**-0.063**	**-4.811**	0.056	1.596	0.085	1.411
Decision aid	**-1.029**	**-7.342**	**-1.544**	**-5.358**	-0.199	-0.220	0.577	0.634
Round X aid			**0.039**	**2.060**			-0.061	-0.885
The constant term, four valuations dummies, four fee dummies, and their interactions are included in the specification, but the respective coefficient estimates are not shown in this table.								

Note to table: The logistic regression includes all person-round observations in the first block (25 rounds, so number of observations is 2300 = 92 subjects x 25 rounds). The linear regression includes only the person-round observations in which the person entered. In the linear regression, standard errors are clustered at the individual level (92 subjects).

Table 5 documents the effect of experience on *excess* entry by aid condition, that is, the same analysis as shown in Table 3. The analysis of behavior without the decision aid replicates the presence of an experience effect found in Experiment 1, and the analysis of behavior with the aid sharpens the conclusion from Table 4 about the moderating effect of the decision aid on the experience effect from all entry to excess entry.

Table 5: Effect of Experience on Excess Entry, by Aid Condition (Logistic Regression)

Dependent variable: entry (binary)	No decision aid				Decision aid			
	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat
Constant	0.577	1.618	-1.038	-0.845	**-1.864**	**-3.258**	-1.985	-1.658
Round	**-0.068**	**-4.514**	**-0.102**	**-5.213**	**-0.039**	**-1.974**	**-0.075**	**-2.334**
12 (*v,f*) dummies	included		included		included		included	
Subject FEs	-		included		-		included	

Note to table: Excess entry is defined as entry when a risk-neutral bidder should not enter. Logistic regression with N=598=46(subjects)x13(*v,f* cells) observations in each aid condition. The regression specification includes all person-round observations in the first block (25 rounds) whenever the (*v,f*) combination was such that a risk-neutral buyer would not enter (the 13 cells of the grey area of Table 1 with v<70). The constant is coded as the v=5 & f=1 condition, and the 12 (*v,f*) dummies capture the remaining 12 cells relative to the constant.

Table 6 conducts the analysis analogous to Table 5 for bid magnitude, mostly replicates the finding from Experiment 1 that experience does not affect bid magnitude (though there is a small significant positive effect not robust to exclusion of individual fixed effects here), and extends it to the setting with the decision aid.

Table 6: Effect of Experience on Observed Bids (Linear Regression)

Dependent variable: bid (when observed)	No decision aid (N=666)				Decision aid (N=551)			
	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat	Estimate	t-stat
Constant	**3.396**	**2.573**	**5.971**	**2.641**	**2.396**	**2.204**	-2.763	-1.603
Round	0.062	1.441	**0.081**	**2.243**	0.020	0.526	0.039	1.221
22 (v,f) dummies	included		included		included		included	
Subject FEs	-		included		-		included	

Note to table: Linear regression that includes all the person-round observations in which the person entered. The number of observations is the number of observed bids in each condition. Valuation dummies, fee dummies, and their interactions provide a separate intercept for each of the 22 cells of Table A1 in Appendix A in which at least one bidder entered. Standard errors are clustered at the individual level (46 subjects in each regression).

We summarize the implications of Tables 4-6 in our next result:

Result 5 (Decision difficulty drives excess entry and moderates the effect of experience): *The difficulty of assessing the expected benefit from bidding amplifies entry. Excess entry declines and the experience effect discussed in Result 4 is moderated when subjects have a decision aid that simplifies the task.*

To go beyond average effects documented so far, Figure 2 unpacks the effect of decision aid on entry by analyzing it separately in each (v,f) condition. From Figure 2, we see that whereas the decision aid reduces entry and excess entry on average (Tables 4 and 5), it does not have a simple main effect (i.e., it does not reduce entry by the same amount in all (v,f) conditions). 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. Please see Table A1 in the Appendix A for the data behind Figure 2 and the crucial diagonal. Another interesting feature of Figure 2 is the lack of a significant difference in entry when the fee is zero and valuations are small. This finding implies our decision aid does not significantly increase the hassle cost h of participating in NYOP.

Figure 2: Effect of Decision Aid on Entry, between Subjects without Experience

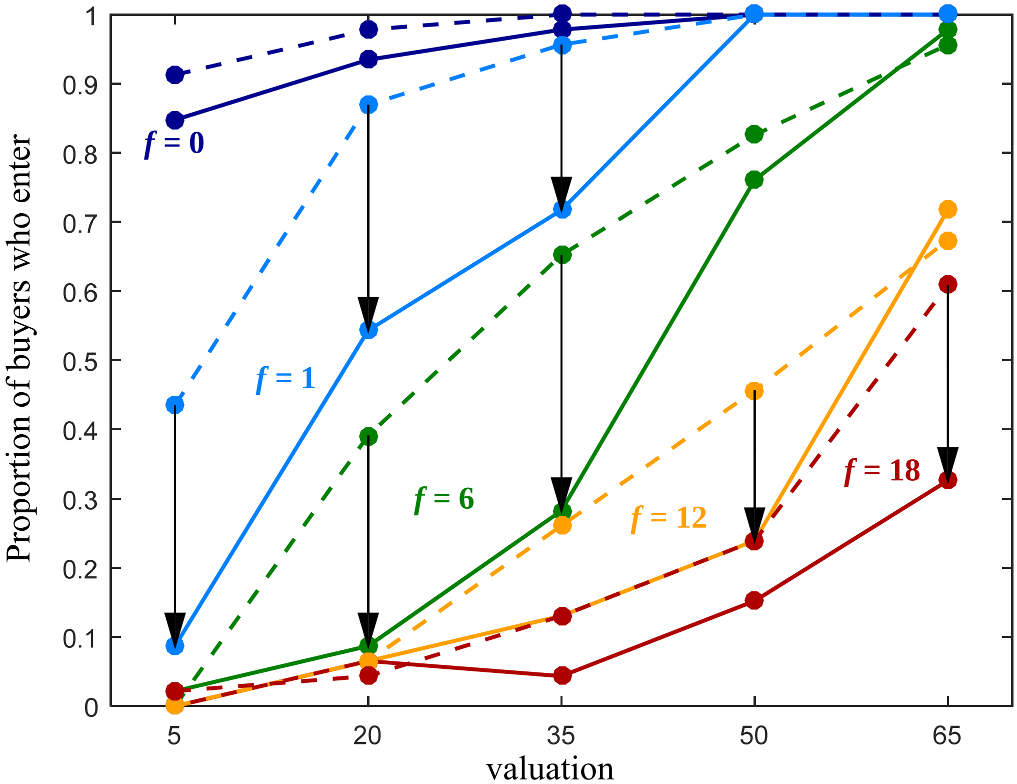

Note to figure: Only first-block observations are included. The solid lines connect the observed entry probabilities when the buyers have the decision aid (N=46); the dashed lines do the same for buyers who do not have the decision aid (N=46). Arrows indicate significant differences at the 5% level using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

Figure 3 unpacks the average experience effect shown by the above regressions by focusing on each (v,f) condition separately, analogously to the approach in Figure 2. Specifically, Figure A3 compares the entry behavior of subjects without the decision aid who experience a given (v,f) condition early in the block with those who experience it later. The entry-reduction pattern in Figure A3 suggests 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). But unlike the decision aid, experience alone does not reduce marginal entry for moderate fee levels. One way to interpret the different patterns of entry reduction from the decision aid and experience is that 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. The fact that reduction due to experience does not happen for a wide range of valuations and fees is further evidence against the novelty explanation of the experience effect.

Figure 3: Effect of Experience on Entry, between Subjects without the Decision Aid

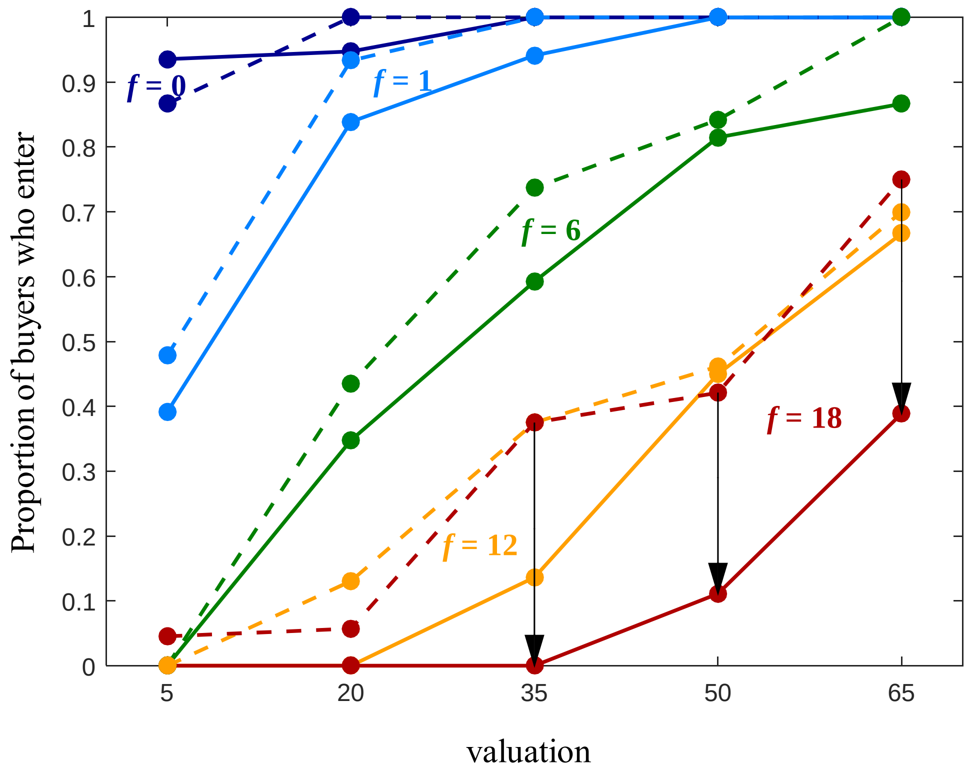

Note to figure: Only the first-block behavior of subjects who start without the decision aid (N=46) is considered. The solid lines connect the observed entry probabilities during periods 13-25; the dashed lines do the same during periods 1-12. Arrows indicate significant differences at the 5% level using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

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. The within-subject effect of the decision aid and full-block experience 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 Appendix A for a plot of the within-subject effect analogous to Figure 2.

In Appendix B, we analyze the demographic and psychographic correlates of bid magnitude and excess entry. We find that stated risk preferences correlate with entry and bidding as expected in an obvious alternative model with heterogeneous risk aversion. By contrast, attitude toward the bidding fee does not have a systematic effect on behavior, and gender only affects entry. Much heterogeneity remains after controlling for the demographics and stated preferences.

APPLICATION: PROFITABILITY OF TWO-PART TARIFFS IN NYOP SETTINGS

Nothing about the consumer-behavior patterns documented so far depends on what happens to the NYOP bidding fee collected by our experimental marketplace. Spann, Zeithammer, and Häubl (2010) propose that when the NYOP seller can keep such a fee as revenue, the resulting two-part tariff is more profitable than providing the bidding opportunity for free as long as the potential buyers are risk neutral. In this section, we apply our findings to test the profitability of two-part tariffs in NYOP settings (2PNYOP for short) empirically. We focus on data from Experiment 2. Corresponding figures for Experiment 1 are available from the authors upon request.

A 2PNYOP seller's profit consists of the bidding fee paid by the buyers who enter the NYOP store plus the difference between the seller's wholesale cost c and buyer's bid b whenever a bid gets accepted. In our experiments, we draw an actual c 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 $\pi(b)$ as follows:

$$\pi(b) \equiv \int_0^b (b-c)\left(\frac{1}{p}\right)dc = \frac{b^2}{2p}. \tag{4}$$

In words, we take the bid as given, and average the unit contribution over all possible wholesale-cost realizations. Let $bid_{i,v,f}$ be the bid submitted by subject *i* with induced valuation *v* when the bidding fee is *f*, and let $bid_{i,v,f} = 0$ when that subject does not enter. The contribution of each subject-round to the seller's profit is thus

$$contribution_{i,v,f} = \mathbf{1}\left(bid_{i,v,f} > 0\right)\left[f + \pi\left(bid_{i,v,f}\right)\right]. \tag{5}$$

Finally, the expected 2PNYOP seller profit when the fee is *f* can be calculated as the average over *i* and *v* of $contribution_{i,v,f}$. Figure 4 plots this expected seller profit in three different conditions of Experiment 2, and compares it with the theoretical profit expected under risk neutrality. We discuss three key take-aways from Figure 4 next.

The first noticeable feature of Figure 4 is that for most fee levels, our subjects are more profitable to the seller than a risk-neutral theory would suggest. This surprising fact arises from the combination of excess entry and over-bidding documented in the previous sections. Second, Figure 4 shows the decision aid reduces seller profit for all positive fees, and the decision aid combined with prior experience reduces it even more. For every positive fee level except 12, both 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. When the fee is 12, only the difference between decision aid + prior experience and baseline is significant. 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 pattern of profit reductions for positive fees resonates with the observation that the decision aid and experience influence the subjects' entry decisions under positive fees, but have little effect on their associated bidding strategies.

Figure 4: Expected Profit of the NYOP Seller, by Condition (Experiment 2)

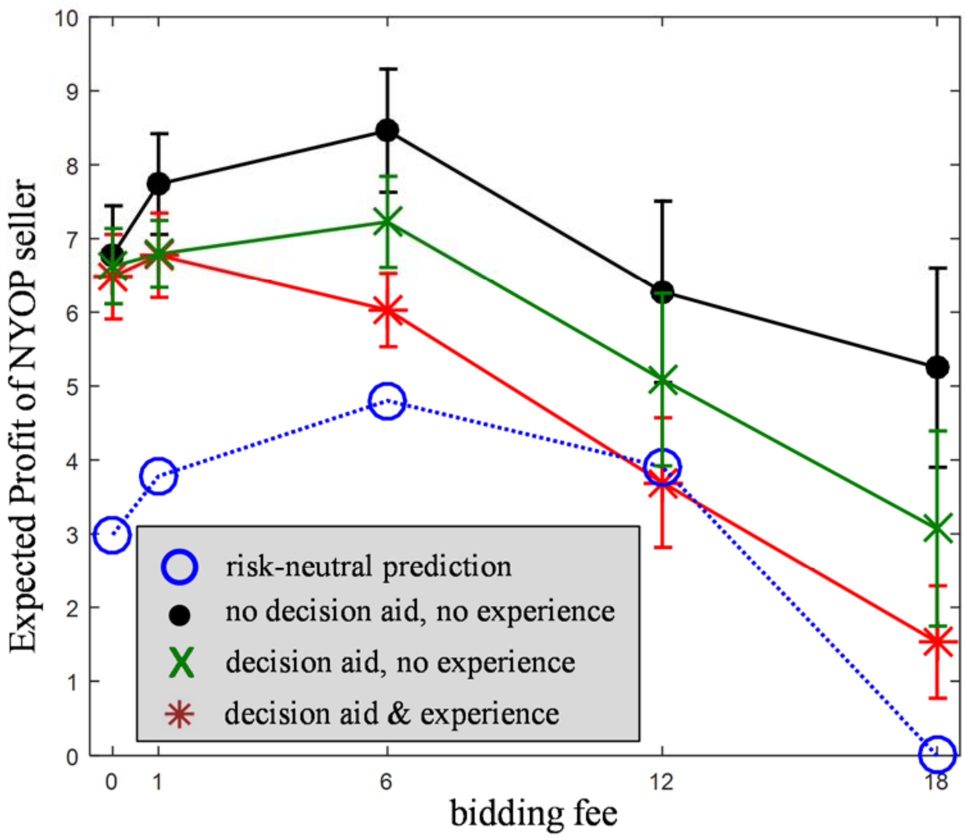

Note to figure: The solid lines connect the observed expected profits, displayed as error bars. Each error bar represents the 95% confidence interval, with one observation defined as the expected profit from one subject, so N=46. Solid (**black**) dots indicate the baseline condition without prior experience (i.e., in the first block) or decision aid. The (**green**) X markers indicate the condition with the decision aid but without prior experience (i.e., in the first block). The (**red**) star markers indicate the condition with prior 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.

The third notable feature of Figure 4 is the difference between the shapes of the three profit functions: In the baseline condition without the decision aid or experience, 2PNYOP is clearly profitable; the seller's profits increase by 25% (p<0.01) when she increases the fee from zero to the empirically optimal level of 6. However, the 9% increase from the same strategy when the subjects are given the decision aid is not only much smaller, but also statistically not 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 the zero fee is again not significant.

We summarize the three takeaways from Figure 4 in our next result:

Result 6: *The risk-neutral model under-predicts seller profit. The decision aid reduces the expected profit of the NYOP seller who charges a positive fee, and additional buyer experience reduces the profits further. A two-part tariff is significantly more profitable at our experimental fee levels only when the buyers have no experience and no access to the decision aid.*

The last sentence of Result 6 leaves room for other fee levels, such as 3, to be potentially more profitable than a fee of zero. However, we only experimented with a limited set of fees.

GENERAL DISCUSSION

We use a laboratory implementation of a tractable single-bidder NYOP setting to study consumer behavior in markets with bidding and costly participation. In both of our experiments, the submitted bids are more stable across rounds and less susceptible to task-difficulty manipulations than the entry decisions of the buyers. Although the observed bids can be rationalized with a standard risk-averse model, the entry probabilities exceed those in any risk-averse expected utility model. This "excess entry" finding resonates with prior empirical work on auctions. We identify two moderators of excess entry: experience and reduced task difficulty.

One underlying mechanism consistent with our findings can be based on the entry decision being more difficult than the bidding decision nested in it: Because bids do not exhibit experience and difficulty effects, our subjects seem to be solving the tradeoff between winning more often and making more surplus upon winning (equation 1). But because their entry decisions are moderated by experience and difficulty, our subjects seem to find the comparison of the bidding fee with the expected utility from bidding (equation 2) difficult. This difficulty results in both higher entry and a downward experience effect on entry because subjects try to learn about the key tradeoff by entering a lot, especially in the early rounds of the experiment. Our contribution is an extension of previous findings of excessive entry into first-price sealed-bid (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. We also 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. Finally, we rule out explanations based on novelty of the task, joy of winning, misperception of probabilities, and overconfidence in competitive settings.

We now discuss the two moderators of excess entry in more detail. Experience reduces excess entry in the sense that subjects in later rounds of both our experiments enter less whenever a risk-neutral subject would not, ceteris paribus. A closer examination of the entry reduction reveals that subjects behave as if they learned to categorically avoid high bidding fees over time. However, subjects do not seem to learn to also avoid medium-sized fees when their valuations are low enough that a risk-neutral subject would not enter. Thus, subjects with mere experience do not seem to be learning how to compare the fee with the expected utility of bidding. Because the entry reduction due to experience does not happen for a wide range of valuations and fees, we can also rule out novelty of the task as an alternative explanation of the experience effect. Another reason we propose the experience effect represents learning is that it is strongly attenuated by our decision aid, which we discuss next.

Hypothesizing that our 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 experiment, we show that a decision aid that helps buyers calculate the probability and payoff consequences of their planned actions reduces excess entry and attenuates the effect of experience. Because the decision aid does not have any significant impact on bids, our results suggest the entry part of our subjects' decision is particularly error prone, whereas the bidding part is more consistently driven by underlying preferences. A detailed analysis reveals the decision aid reduces expected entry along the entire curve in the (*valuation*,*fee*) space along which a risk-neutral buyer is indifferent between entering and staying out for the round. Therefore, our decision aid is more effective in helping buyers decide whether the bidding fee is worth the expected utility from bidding—a more thorough and integrative change in behavior than that produced by experience alone.

Our experiments can be used to directly test the profitability of two-part tariffs in an NYOP setting. We find a two-part tariff can be profitable for the NYOP seller, 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. The cognitive difficulty associated with buyers' joint optimization of their entry and bidding decisions are thus catalysts of two-part-tariff profitability in NYOP. We find no evidence of fee aversion that has been hypothesized to inhibit profitability of two-part tariffs in an NYOP setting. 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 when and why excess entry and over-bidding occur in bidding markets with costly entry. However, it does not tell us much about the magnitude of the observed effects in real markets. Therefore, additional field experiments with decision aids such as the one we proposed would be very interesting.

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.

Aycinena, Diego and Lucas Rentschler (2018), "Auctions with Endogenous Participation and an Uncertain Number of Bidders: Experimental Evidence," Experimental Economics, 21 (4), 924–949.

Bajari, Patrick and Ali Hortaçsu (2005), "Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data," Journal of Political Economy, 113 (4), 703–741.

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.

Camerer, Colin and Dan Lovallo (1999), "Overconfidence and Excess Entry: An Experimental Approach," American Economic Review, 89 (1), 306–318.

Chakravarty, Sujoy, Glenn W. Harrison, Ernan E. Haruvy, and Elisabet E. Rutström (2011), "Are You Risk Averse over Other People's Money?" Southern Economic Journal, 77 (4), 901–913.

Cox, James C., Bruce Roberson, and Vernon L. Smith (1982), "Theory and Behavior of Single Object Auctions," In: Vernon L. Smith, ed., Research in Experimental Economics, Vol. 2. Greenwich, CT: JAI Press.

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), 448–474.

Davis, Andrew M., Elena Katok, and Anthony M. Kwasnica (2014), "Should Sellers Prefer Auctions? A Laboratory Comparison of Auctions and Sequential Mechanisms," Management Science, 60 (4), 990–1008.

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), "Explaining Overbidding in First Price Auctions Using Controlled Lotteries," Experimental Economics, 6 (2), 123–140.

Engelbrecht-Wiggans, Richard and Elena Katok (2005), "Experiments on Auction Valuation and Endogenous Entry," In: John Morgan, ed., Experimental and Behavioral Economics. Bingley: Emerald Group Publishing Limited.

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 (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.

Feng, Cong, Scott Fay, and K. Sivakumar (2016), "Overbidding in Electronic Auctions: Factors Influencing the Propensity to Overbid and the Magnitude of Overbidding," Journal of the Academy of Marketing Science, 44 (2), 241–260.

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.

Gneezy, Ayelet, Uri Gneezy, Gerhard Riener, and Leif D. Nelson (2012), "Pay-What-You-Want, Identity, and Self-Signaling in Markets," Proceedings of the National Academy of Sciences, 109 (19), 7236–7240.

Greiner, Ben (2015), "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE," Journal of the Economic Science Association, 1 (1), 114–125.

Haruvy, Ernan and Peter T. L. Popkowski Leszczyc (2016), "Measuring the Impact of Price Guarantees on Bidding in Consumer Online Auctions," Journal of Retailing, 92 (1), 96–108.

Haruvy, Ernan and Peter T. L. Popkowski Leszczyc (2018), "A Study of Bidding Behavior in Voluntary-Pay Philanthropic Auctions," Journal of Marketing, 82 (3), 124–141.

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.

Krasnokutskaya, Elena and Katja Seim (2011), "Bid Preference Programs and Participation in Highway Procurement Auctions," American Economic Review, 101 (6), 2653–2686.

Menezes, Flavio M. and Paulo K. Monteiro (2000), "Auctions with Endogenous Participation," Review of Economic Design, 5 (1), 71–89.

Moreno, Diego and John Wooders (2011), "Auctions with Heterogeneous Entry Costs," RAND Journal of Economics, 42 (2), 313–336.

Palfrey, Thomas R. and Svetlana Pevnitskaya (2008), "Endogenous Entry and Self-Selection in Private Value Auctions: An Experimental Study," Journal of Economic 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.

Samuelson, William F. (1985), "Competitive Bidding with Entry Costs," Economics Letters, 17 (1–2), 53–57.

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.

Smith, James L. and Dan Levin (1996), "Ranking Auctions with Risk Averse Bidders," Journal of Economic Theory, 68 (2), 549–561.

Spann, Martin, Robert Zeithammer, Marco Bertini, Ernan Haruvy, Sandy D. Jap, Oded Koenigsberg, Vincent Mak, Peter T. L. Popkowski Leszczyc, Bernd Skiera, and Manoj Thomas (2018), "Beyond Posted Prices: The Past, Present, and Future of Participative Pricing Mechanisms," Customer Needs and Solutions, 5 (1-2), 121–136.

Spann, Martin, Robert Zeithammer and Gerald Häubl (2010), "Optimal Reverse-Pricing Mechanisms," Marketing Science, 29 (6), 1058–1070.

Spann, Martin, Robert Zeithammer and Gerald Häubl (2015), "Erratum to "Optimal Reverse-Pricing Mechanisms" by Martin Spann, Robert Zeithammer, and Gerald Häubl," Marketing Science, 34 (2), 297–299.

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.

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 A

Figure A1: Layout of Experimental Interface of Experiment 1

Figure A2: Layout of Experimental Interface of Experiment 2

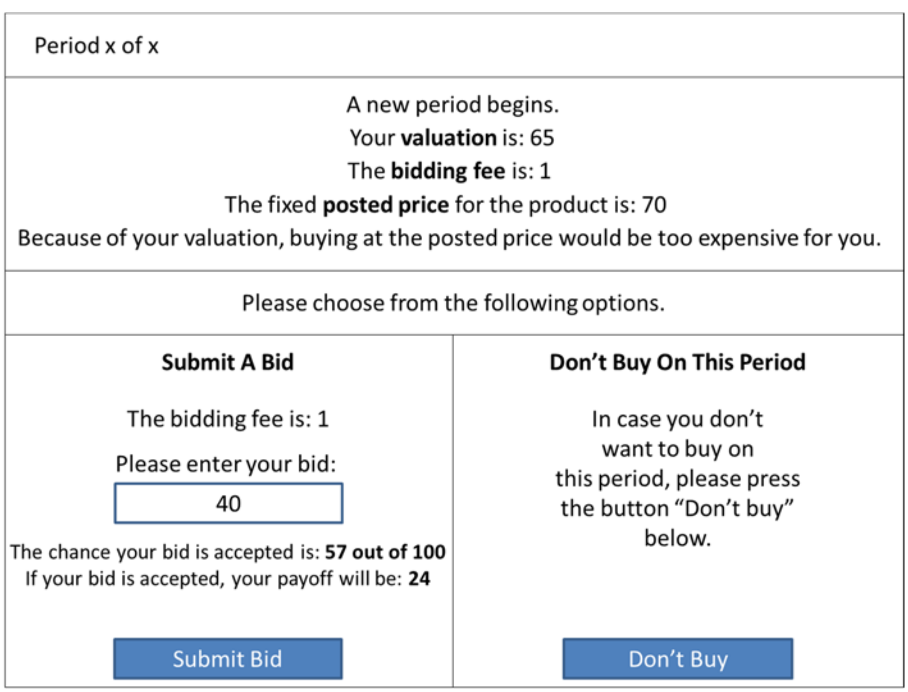

Figure A3: Effect of Decision Aid and Prior Experience on Entry, within Subject

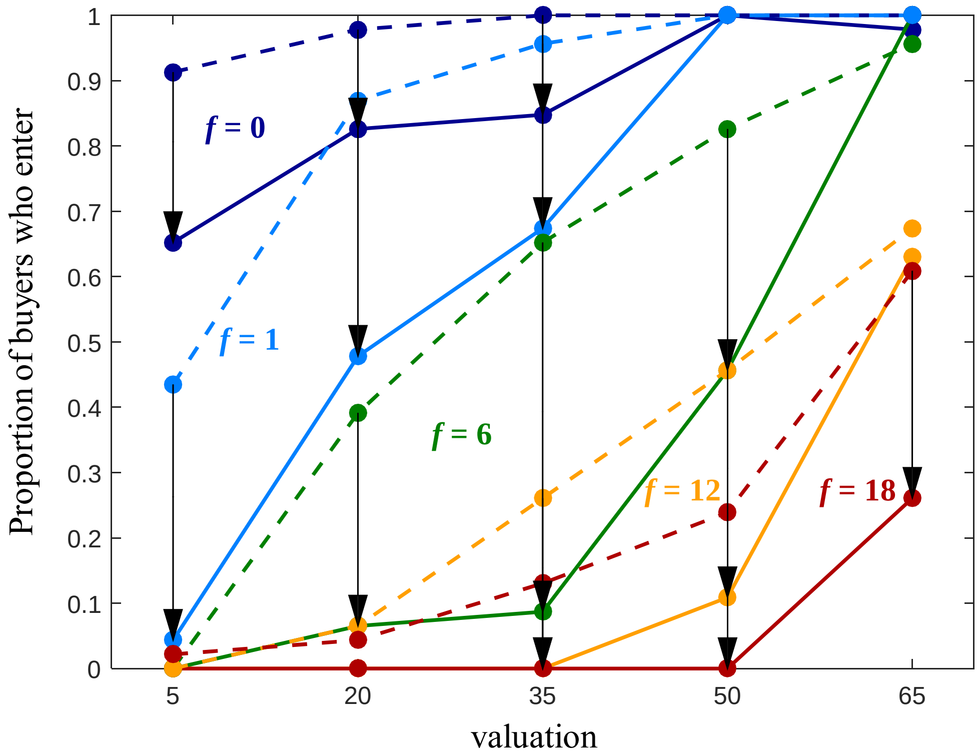

Note to figure: Only subjects who start without the decision aid (N=46) are considered. The solid lines connect the observed entry probabilities when the subjects do have the decision aid (second block); the dashed lines do the same when the subjects do not have the decision aid (first block). Arrows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

Table A1: Effect of Decision Aid on the Proportion of Subjects Who Enter

		Bidding fee				
	No aid (N=46)	0	1	6	12	18
Valuation	5	91%	43%	0%	0%	2%
	20	98%	87%	39%	7%	4%
	35	100%	96%	65%	26%	13%
	50	100%	100%	83%	46%	24%
	65	100%	100%	96%	67%	61%
	Decision aid (N=46)	0	1	6	12	18
Valuation	5	85%	9%	2%	0%	0%
	20	93%	54%	9%	7%	7%
	35	98%	72%	28%	13%	4%
	50	100%	100%	76%	24%	15%
	65	100%	100%	98%	72%	33%
	Difference (aid – no aid)	0	1	6	12	18
Valuation	5	-7%	**-35%**	2%	0%	-2%
	20	-4%	**-33%**	**-30%**	0%	2%
	35	-2%	**-24%**	**-37%**	-13%	-9%
	50	0%	0%	-7%	**-22%**	-9%
	65	0%	0%	2%	4%	**-28%**

Note to table: Proportion of subjects who enter in each (v,f) cell, by aid condition. The shaded area in the top two tables indicates when a risk-neutral buyer should not enter. The thick line in the bottom table shows the boundary of the shaded area. Bold differences in the bottom table are significant at the 5% level. All of them except (v=50, f=12) are also significant at the 1% level.

Appendix B: Individual Characteristics Related to Entry and Bidding Behavior

In both experiments analyzed in this paper, we found bids to be stable and entry to be declining with experience. Both entry and bids consistently exceed their risk-neutral benchmarks. The second experiment showed that helping the subjects with calculations reduces their excess entry and moderates the effect of experience on it. In this section, we pool the data across the two experiments and explore demographic and psychographic correlates of bid magnitude and excess entry. Table B1 describes the population distribution of demographics and psychographics across the two experiments. The population distributions of both the Holt-Laury (2002) and the Dohmen et al. (2012) measures are consistent with other findings in the literature. We omit age from the analyses in this section because the distribution of age in our population is highly concentrated around its median, with 91% of our subjects below the age of 30.

Table B1: Individual Characteristics across Both Experiments (N=183)

						Correlations			
	Mean	Median	Std. Dev.	Min	Max	Female	Age	Risk-taker	Fee refuser
Risk-averse	6.14	6	1.64	0	10	-0.01	0.10	-0.18	0.10
Female	0.55	1	0.50	0	1	1	0.09	-0.06	0.06
Age (years)	24.23	23	5.96	18	60		1	0.03	-0.08
Risk-taker	3.70	3	1.35	1	7			1	-0.07
Fee refuser	4.38	5	1.93	1	7				1

Note to table: Risk aversion is measured as in Holt and Laury (2002), and shown here as number of times the subject selects the safe option out of the 10 lottery-choice tasks. Subjective risk-taking is measured as in Dohmen et al. (2012), namely, as an answer to "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.

Table B2 reports a linear regression of each subject's observed probability of excess entry on the individual characteristics (each subject is the unit of observation, and the dependent variable is the proportion of the excess entry cells in which the subject entered), as well as on controls for the two different experiments and the presence of the decision aid. We expected more risk-averse subjects to enter less, and this expectation is indeed borne out by the negative sign on the Holt-Laury risk-aversion measure. The other significant correlate is gender. 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.

Table B2: Who Enters Too Much? Linear Regression of Individual Probability of Excess Entry

Variable	Estimate	*t*-stat	Estimate	*t*-stat
Intercept	**16.83%**	**7.92**	**16.46%**	**7.63**
Female	**6.52%**	**2.76**	**6.35%**	**2.63**
Decision aid	**-12.01%**	**-3.60**	**-13.60%**	**-4.01**
Experiment 2	3.66%	1.25	5.37%	1.83
Risk-averse	**-3.51%**	**-2.89**		
Risk-taker	1.85%	1.54		
Fee refuser	0.37%	0.31		
R^2	0.18		0.12	

Note to table: A unit of observation is one subject. Linear regression of the subject's 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. All scales are standardized within the entire population and hence are measured in population standard deviations. All other variables are dummies. Number of observations = 183 (91 subjects from Experiment 1, 92 from Experiment 2).

Table B3 reports five separate linear regressions of bid magnitude on the individual characteristics, holding constant the valuation and the fee. We expected risk aversion to correlate positively with bid magnitude, and this expectation is indeed borne out both by the positive signs of the Holt-Laury risk-aversion measure and by the negative sign of the Dohmen et al. subjective risk-taker score.

Table B3: Relationship between Bid Magnitude, Demographics, and Preference Survey Responses

	f=0 & *v*=35		*f*=0 & *v*=50		*f*=0 & *v*=65		*f*=1 & *v*=50		*f*=1 & *v*=65	
	Coeff	t-stat	Coeff	t-stat	Coeff	t-stat	Coeff	t-stat	Coeff	t-stat
Intercept	**20.50**	**23.45**	**29.48**	**26.25**	**40.55**	**29.33**	**32.19**	**31.18**	**39.83**	**32.55**
Female	1.18	1.21	1.68	1.36	-0.77	-0.50	0.15	0.13	**2.67**	**1.97**
Decision aid	**3.28**	**2.42**	-1.13	-0.65	-2.37	-1.10	-1.44	-0.90	-1.81	-0.94
Experiment 2	2.13	1.77	**8.17**	**5.32**	**7.97**	**4.17**	**6.93**	**4.90**	**7.27**	**4.31**
Risk-averse	0.61	1.21	**2.40**	**3.78**	**1.89**	**2.39**	**1.83**	**3.08**	1.16	1.66
Risk-taker	**-1.15**	**-2.34**	**-1.57**	**-2.50**	**-1.97**	**-2.53**	**-1.44**	**-2.49**	**-1.93**	**-2.79**
Fee refuser	0.01	0.02	0.34	0.55	0.59	0.76	0.10	0.16	0.62	0.91
N	179		182		182		181		183	
R^2	0.157		0.261		0.165		0.209		0.189	

Note to table: A unit of observation is one subject. Linear regression of observed and positive bids at a given (*v*,*f*) cell in Study 1 and the first block of Study 2. All scales are standardized within the entire population, and hence are measured in population standard deviations. All other variables are dummies.

Taken together, the analyses in Tables B2 and B3 show that risk preferences correlate with entry and bidding as expected in an obvious alternative model with heterogeneous risk aversion. By contrast, attitude toward the bidding fee does not have a systematic effect on behavior, and gender only affects entry. Much heterogeneity clearly remains after controlling for the demographics and stated preferences.