Consumer Preference Inconsistencies in Costly Price Offers

Robert Zeithammer Lucas Stich Martin Spann Gerald Häubl

December 7, 2022

ABSTRACT

Consumers submit price offers to sellers in a variety of domains. Submitting an offer often comes with administrative, waiting, and opportunity costs. Making such costly price offers involves two intertwined decisions: in addition to determining how much to offer, consumers must also decide whether to make an offer in the first place. We examine the impact of offer-submission costs on consumer behavior using three incentive-compatible experiments. Our findings reveal a preference inconsistency whereby the preferences implied by one of the decisions do not agree with preferences implied by the other: potential buyers enter more often than standard preference models predict they should. The inconsistency is robust to interventions designed to help consumers with their decisions: (1) the provision of interactive-feedback decision aids and (2) the sequencing of the two sub-decisions in the normative order. Our interventions have large effects on behavior, and jointly suggest that consumers approach the two decisions *as if* they were unrelated instead of following the normative backward solution process most existing models assume. We discuss the implications of our findings for the design of offer-submission interfaces, as well as for attempts to infer consumer preferences from offer and bidding data.

Keywords: Pricing, Auctions, Entry Costs, Behavioral Economics, Experiments

Contact: Robert Zeithammer, rzeitham@ucla.edu

1. Introduction

Consumers make price offers to sellers in a variety of domains, including buying cars and houses, acquiring used household goods in garage sales and their online parallels such as Craigslist, naming hotel-room prices on Priceline.com, making offers to buy trendy items on StockX.com, and bidding in auctions for art, airline upgrades, or collectibles (Bajari and Hortaçsu 2005; Haruvy and Popkowski Leszczyc 2018). An important source of friction inherent in such participative pricing markets (Spann et al. 2018) is the *participation cost* of making an offer, that is, the administrative cost of submitting a price offer (e.g., completing official paperwork for a house offer, traveling to a dealership to make an offer on a new car, or joining and learning an online platform like Greentoe.com to make offers on consumer electronics), as well as the time and opportunity cost associated with waiting for the outcome (e.g. waiting seven days for an eBay auction of a collectible to conclude, or losing out on a better travel option while waiting several days for a Lufthansa auction of a business-class seat to conclude).[1] In all of the above examples, the cost of making an offer is substantial, so it is important for both researchers and managers to understand how consumers account for such costs in their decision-making.

We examine the impact of participation costs on consumer behavior using three incentive-compatible experiments in which we manipulated both the cost of participation and several key design elements of the bidding interface. In the auction literature, the costs we study are often called "entry costs" following Samuelson (1985).[2] In contrast to prior experimental work on bidding in auctions with entry costs (summarized well in Palfrey and Pevnitskaya (2008)), our experiments use a simpler setting motivated by the name-your-own-price mechanism pioneered by Priceline (e.g., Amaldoss and Jain 2008; Fay 2004; Shapiro 2011). A subject in our experiments submits a binding price offer to a (computer-simulated) seller, and the seller's random hidden reserve price determines whether this offer is accepted. Our single-agent decision is an ideally tractable setting for gaining insight into consumer offer-making with participation costs because of

---

[1] Prior literature has examined these costs. Fay (2009) and Hann and Terwiesch (2003) study the hassle cost of submitting the offer and waiting for the outcome. Bernhardt and Spann (2010) study the monetary cost arising from various fees and commissions. Palfrey and Pevnitskaya (2008) consider the opportunity cost of participation.

[2] Additional types of costs associated with participation in auctions exist, for example, costs incurred to learn details about the good sold as in Levin and Smith (1994), but such costs are beyond the scope of this paper.

the absence of strategic interaction among multiple endogenously determined participants, which would be present in a multi-bidder auction with entry costs (e.g., Menezes and Monteiro 2000; Samuelson 1985).

Our main audience is researchers and analysts who estimate consumer preferences from bidding and offer data. Crucially for estimating consumer preferences from the behavior we observe in our studies, the costly price-offer task involves two interrelated decisions: in addition to determining how much to offer, consumers who face participation costs must also decide whether to make an offer at all–that is, whether to "enter" the uncertain situation in the first place. Both offer amounts and entry decisions can be separately interpreted as arising from underlying preferences, and we document a robust internal preference inconsistency whereby the preferences implied by one of the decisions do not agree with preferences implied by the other. Specifically, we find potential buyers enter more often than their offer amounts predict they should, based on standard economic models.

Behavioral science has made significant progress toward modeling human behavior under uncertainty by documenting various "preference reversals"—inconsistencies between underlying consumer preferences revealed by two seemingly unrelated yet structurally identical tasks (e.g., Lichtenstein and Slovic 1973; Tversky, Slovic, and Kahneman 1990; and many others). By contrast, we document a preference inconsistency internal to a *single* task, whereby a consumer's preferences revealed by one part of the task seem different from the same consumer's preferences revealed by another part of the same task.

The inconsistency we document is surprisingly resistant to reconciliation via enriched model specifications, both existing and newly proposed in this paper. We rule out two such existing enrichments: (1) Our single-agent setting rules out strategic explanations relevant in auction settings (e.g., Goeree et al. 2002; Battigalli and Siniscalchi 2003; Crawford and Iriberri 2007); and (2), our results are also not consistent with consumers getting additional utility from winning or playing as has been proposed in the related auction literature (e.g., Palfrey and Pevnitskaya 2008; Ertaç et al. 2011). We also test two new potential explanations of the inconsistency: (1) the costly offer decision may be so difficult for consumers that they resort to simplifying heuristics, which bias their behavior in one or both decisions and (2) consumers may not be processing the two decisions in the correct order. Normative theories of behavior in our setting have a clear backward solution structure: a rational decision-maker is supposed to parse the overall problem into the offer-amount and entry decision, and solve the decisions "backward" starting with the offer amount.

To test the decision difficulty hypothesis, we provide some of our subjects with an interactive decision aid (e.g., Häubl and Trifts 2000) that provides real-time assistance with calculating the chances and payoffs given a potential offer amount. To help our subjects follow the normative process, we use two manipulations of increasing forcefulness: (1) We merely guide them through the normative process by parsing the overall decision across two screens sequenced in the normative order; and (2) we force them to follow the normative process by first soliciting incentivized (automatically submitted) offer amounts and then piping them within subject as exogenously fixed-offer amounts in seemingly separate entry decisions.

Surprisingly, none of our manipulations—neither the decision aid nor either of the two levels of the process guidance—resolve or even reduce the inconsistency between the two decisions. Instead, both the decision aid and our guiding of respondents toward the normative process actually exacerbate the inconsistency, driving the preferences suggested by entry behavior further away from the preferences suggested by the offer amounts. The manipulations have two effects: First, we find that making the necessary computations easier via a decision aid improves consumers' understanding of the probability that their offer is accepted, affects their entry behavior consistently with their improved probability perception, but has no corresponding effect on offer amounts. Second, and in stark contrast to the decision aid's effect, forcing subjects through the normative process affects their offer amounts but not their subsequent entry behavior. Specifically, we find that subjects offer more when their offers are automatically submitted than when we merely guide them through the normative process. But their behavior in the subsequent entry task is the same—as if the offer magnitude did not matter much for the entry decision.

Taken together, our results imply the inconsistency is a robust feature of decision-making in costly participative pricing, not a mistake that somehow needs to be corrected. If the inconsistency is a feature and not a bug, what class of models are consistent with the inconsistency and the pattern of effects we document? In a post-hoc modeling synthesis, we propose that a behaviorally realistic model of consumer behavior with costly price offers needs to decouple the two parts of the decision. In such a model, consumers need to make the two decisions separately, for example in parallel, rather than sequentially as the normative process stipulates. The key evidence in favor of separate processes is the fact that our two experimental manipulations outlined above (decision aid and process guidance) affect only one of the decisions, with one manipulation affecting only entry (decision aid) and the other manipulation affecting only offers (process guidance). In other words, the two decisions do not co-move as they would if subjects were parsing the problem and solving it sequentially. We emphasize at the outset that this surprising implication for behaviorally realistic modeling of bidding with participation costs arises from a post-hoc synthesis of all the experiments we report, and we do not explicitly test any particular behavioral model. Instead, we document an interesting inconsistency, propose multiple plausible mechanisms, and rule them all out. In doing so, we find a pattern of results that violates a very fundamental standard assumption in modeling the phenomenon we study—the normative assumption about solving the two parts of the decision sequentially.

In addition to informing future modeling decisions and re-interpreting the preference inconsistency as a feature instead of a bug, our results also have immediate managerial implications for decision architects (Johnson et al. 2012). Our experiments involve induced valuations, so we can directly measure the amount of monetary surplus consumers earn under different architectures. We answer two questions: First, should interfaces parse the problem for consumers, and how strongly should they enforce the normative process? Second, should interfaces provide decision aids to help with computation? Regarding the sequential nature of the interface, our main finding is that whereas a gentle guiding help can help consumers make better decisions, forcing the backward solution process can backfire and make consumers much worse off than they would be in a simultaneous architecture. Regarding the decision aid, we find it can help consumers earn more surplus, but only when it is combined with a sequentially structured decision architecture. In other words, we find consumers struggle with finding the best amount to offer, and merely focusing on it more upfront does not help them; only when the focus comes with a “what-if” calculator of acceptance chances and contingent payoffs do their decisions improve. From the seller’s perspective, we find that regardless of the decision architecture, sellers make less profit when consumers have access to the interactive decision aid, and hence may be reluctant to offer it to consumers without government regulation or other incentives.

2. Model

The consumer behavior we study consists of two nested decisions: (1) how much to offer or "bid" and (2) whether to make an offer at all. We refer to the latter as the "entry decision" in accordance with the auction literature (e.g., Samuelson 1985; Ertaç et al. 2011; McAfee and McMillan 1987; Palfrey and Pevnitskaya 2008; Moreno and Wooders 2011). The normative economic model captures the behavior of consumer *i* with valuation *v* facing an offer-submission cost of *c* as solving the following problem:

$$\max_{enter,not}\left[\max_{b\geq 0} Pr(accept|b)u_i(v-c-b)+\big(1-Pr(accept|b)\big)u_i(-c),u_i(0)\right] \tag{1}$$

where u_i is the utility of consumer *i* as a function of change in monetary surplus measured in cents, the "*b*" notation for the offer magnitude highlights the fact that one can think of offers as bids in a single-bidder first-price auction with a hidden reserve, and where $Pr(accept|b)$ is the probability (potentially subjectively perceived by the consumer) that the seller accepts an offer of *b* cents. In all our experiments, we implement the true probability to be $\Pr(accept|b)=\left(\frac{b}{100}\right)$ to make the underlying computation as simple as possible, we train subjects to learn this formula, and we measure their probability beliefs to control for potentially biased misperceptions.

In words, equation (1) tells the decision-maker to first determine the optimal offer amount by solving the tradeoff between the probability of acceptance (increasing in the offer amount) and the utility of the monetary payoff (decreasing in the offer amount). Although submitting free (c=0) offers is then a "no brainer," deciding whether to submit an offer when *c*>0 involves comparing the expected utility from bidding with the utility of doing nothing and receiving no payoff.

The general model in equation (1) includes several notable special cases that entail different assumptions about consumers' risk preferences. When $u_i(x)=x$, we obtain the risk-neutral model, which implies buyers should offer half their valuations whenever they enter, and submission costs are effectively sunk (i.e. offer amounts do not depend on submission costs). This model is by far the most popular one in the econometrics of auctions (e.g., Ackerberg et al. 2007; Guerre et al. 2000) as well as the model used in the most closely related theory papers on name-your-own-price selling (Spann et al. 2010; Zeithammer 2015). A globally concave u_i gives rise to the expected utility model of a risk-averse consumer. For example, when $u_i(x) = \frac{(W+x)^{1-R}}{(1-R)}$ for some level of initial consumer wealth W, we obtain the expected utility model with constant relative risk aversion (CRRA) equal to R. This tractable CRRA model is often used to explain bidding above the risk-neutral level in first-price auctions (Bajari and Hortaçsu 2005; Cox et al. 1988). Finally, when the curvature and scale of the utility are allowed to depend on the sign of the argument, this model can accommodate reference-dependent "prospect-theoretic" preferences (Kahneman and Tversky 1979; Tversky and Kahneman 1992); see the Web Appendix for a specific example of such a model.

Prior literature (e.g., Ertaç et al. 2011) has proposed enrichments of the standard model in equation (1), for example the idea of the *joy of winning*, whereby people derive additional value from getting their offer accepted. Such additional utility, denoted as w, could take two forms: *additive*, whereby consumers with actual valuation of v behave as if their valuation were $v+w$ for some $w>0$; and *multiplicative*, whereby they behave as if their valuation were vw for some $w>1$.

Figure 1a: Normative Backward Solution Decision Process in Costly Price Offers

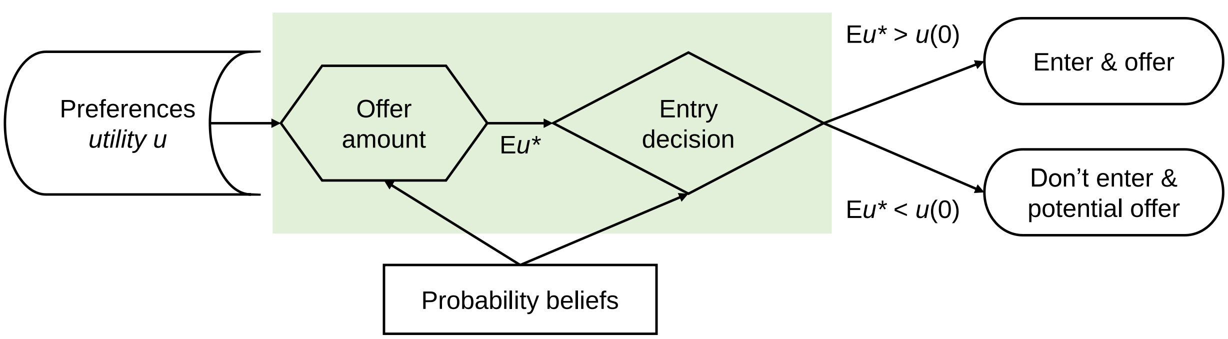

Given a utility function, the model in equation (1) can be used to predict how variations in valuation and submission cost influence the entry and offer behavior of buyers. We have designed the within-subject variation of our experiments to see which utility model fits the observed behavior best (at the individual level) across a wide range of valuations and costs. Equation (1) also implies that the two components of behavior should move together in response to experimental manipulations that affect preferences or probability beliefs. We explain this "co-movement prediction" of the model next.

Regardless of the shape of the utility function or its enrichment, equation (1) restricts the rational decision-maker to first solve the inner maximization problem that finds the best offer given entry, and then compare the expected utility from the best offer $Eu*$ with the utility of staying out of the market $u_i(0)$. Figure 1a illustrates this normative backward solution process with preferences and probability beliefs as inputs. The nested structure of the two decisions of the problem shown in equation (1) implies a rational decision-maker needs to solve the two decisions in the order shown in Figure 1a. Because only one person with one utility function is making the entire decision, the two sub-decisions have to move together in a way consistent with the person's preferences and beliefs. From this observation, we can make the following predictions about consumers who follow the normative process:

(a) A manipulation that systematically affects probability beliefs should have an internally consistent effect on both offer amounts and entry decisions, and

(b) A manipulation that systematically affects preferences should have an internally consistent effect on both offer amounts and entry decisions.

In running our experiments designed to improve the fit of the standard models by helping buyers in various ways, we identified manipulations that violate both of the above predicted patterns. In experiment 1, we found that reducing decision difficulty via a decision aid affects entry consistently with the effect on probability beliefs but does not affect offer amounts as (a) would predict. In experiment 3, we found that forcing (instead of merely nudging) consumers to strictly follow the normative process affects their offers as if buyers became more risk averse but has no associated effect on entry behavior as (b) would predict. Our results are thus inconsistent with the very structure of equation (1), and not only with a particular popular assumption about the shape of the utility function (i.e., expected-utility theory or prospect theory). Instead, our results point to an alternative separate structure illustrated in Figure 1b. Note that this surprising implication for behaviorally realistic modeling of bidding with participation costs arises from a post-hoc synthesis of all the experiments we report, it is not a hypothesis we originally set out to test. Nevertheless, the pattern of effects rules out the normative process in Figure 1a and restricts realistic models to ones that satisfy Figure 1b.

Figure 1b: Behaviorally Realistic Process in Costly Price Offers and Experimental Interventions

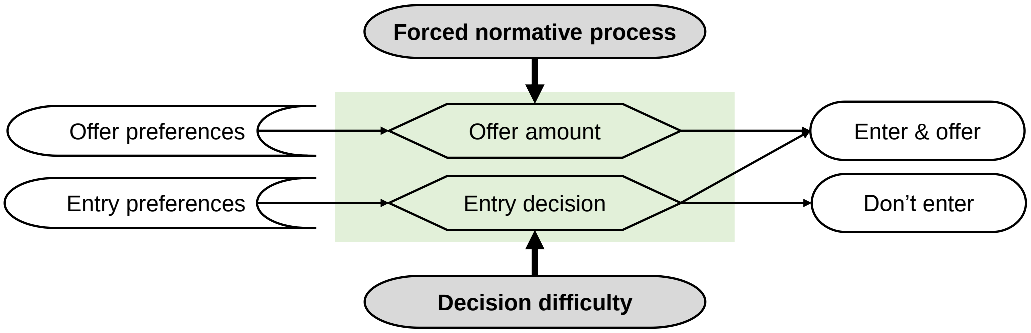

Having outlined the theory, model, and predictions, we now briefly introduce the common aspects of the paradigm used across all our experiments.

3. Experimental Paradigm: Within-Subject Manipulation of Valuation and Offer Submission Cost

We designed our experimental task to measure all the key elements of equation (1): offer amount (b) and entry decision. To provide incentives consistent with equation (1) and to measure earnings, we used the induced-value paradigm (Smith 1976). To abstract from strategic considerations in auctions and the associated explanations of the inconsistency,[3] we used the single-agent setting motivated by name-your-own-price selling. Subjects had the role of a buyer of "widgets" (small imaginary mechanical devices) in a market with only one seller. The widget seller was computerized and entertained a single binding price offer from the buyer. To decide whether an offer was accepted, the seller drew a secret (to subjects) threshold price and then accepted the offer if and only if it was equal to or greater than that threshold. The threshold price was a whole number of cents between 0 and 100 inclusive, chosen at random in each round, and with each value between 0 and 100 being equally likely. We explained to subjects that the probability of an offer being accepted (described in terms of the number of acceptances out of 100) was thus equal to the number of cents

[3] Specifically, the existing behavioral models that hinge on strategic considerations (e.g., strategic sophistication without equilibrium beliefs in Battigalli and Siniscalchi 2003, or k-level reasoning as in Crawford and Iriberri 2007) cannot explain our findings. Prior work has also used the quantal response equilibrium (QRE) model to provide explanations for both excessive entry (Palfrey and Pevnitskaya 2008) and overbidding (Goeree et al. 2002) in auctions. But in our single-agent decision context, the QRE model amounts to adding error terms to standard models of both the offer amount and the entry choice, so the QRE model also cannot explain the inconsistency we document.

they offered.

The key task of each experiment involved multiple rounds, each corresponding to an independent buying opportunity involving a different valuation and submission cost. In each of these rounds, subjects first learned their current valuation of the widget (v in equation (1)) and the current offer-submission cost (c in equation (1)). Following the standard induced-value paradigm (Smith 1976), the valuation was the value (in cents) that owning the widget had to them in that round. The widget had no other value to them. The submission cost was non-refundable and charged to subjects at the time of submitting an offer. After learning the current (v,c) pair, subjects were given the choice either to make the seller a binding offer (b in equation 1) for that round's widget or not to submit an offer to the seller and thus receive zero payoff that round. After running a pilot study described below, we limited the bid amount to the $[0, v - c]$ interval because bidding more than $v-c$ guarantees a loss, and we wanted to rule out irrationality of this type as a potential explanation of our results.

In each round, subjects made money as follows. If their offer was accepted by the seller, they purchased the widget for the price they offered, resulting in a payoff of $v-c-b$. If their offer was rejected by the seller, they did not purchase the widget and did not pay the price they offered, but still paid the current submission cost c. Subjects' payoff from each round was added to their account, which had a starting balance of $1 (to allow for moderate losses). If their payoff in a round was negative (if they lost money in a round), the lost amount was deducted from their account.

Whenever a subject decided to submit an offer, they were asked on an immediately subsequent screen to indicate their prediction of the likelihood of their offer being accepted, measured using a slider scale ranging from 0% (certainly rejected) to 100% (certainly accepted) without any discrete intermediate markers.

For all our experiments, we used CloudResearch's MTurk toolkit to recruit subjects.[4] To make sure each subject understood the game and payoff rules, we used five screening questions designed to test a

---

[4] Our first experiment conceptually replicates the findings of our earlier study using an experimental economics laboratory (details available from authors), so the empirical regularities we document are not limited to online respondents.

detailed understanding of the rules. Only subjects who answered all five questions correctly entered the experiment. We implemented this strict participation restriction to rule out misunderstanding and inattention as possible explanations for the behavior we observed.

After qualifying for the experiment, subjects completed nine training rounds indistinguishable to the subjects from the "measurement rounds" in the key experimental task. In the first two experiments, the measurement rounds were followed by retest rounds designed to assess learning from experience. Upon completion of all rounds, subjects answered a brief demographic and risk-preference questionnaire, the details of which differed somewhat between experiments (stimuli and instructions are available in a project directory on the Open Science Framework)[5].

In the concluding questionnaire, all experiments included a survey of attitudes to measure specific explanations of their behavior. On a 7-point scale (from 1 = *not at all* to 7 = *very much*), subjects were asked to indicate their agreement with statements characterizing what they were thinking while making or considering making an offer to the seller. We designed the statements to measure an a priori dislike of paying a submission cost ("I did not want to pay for submitting an offer"), loss aversion ("I was afraid of losing the submission cost" and "Losing the submission cost feels worse than gaining a payoff of the same amount feels good"), agreement with the objective of a risk-neutral buyer ("I wanted to maximize my potential payoff"), the joy of winning ("My offer being successful was more important to me than the potential payoff"), and the affective response to the situation as a source of potential joy of playing ("Submitting an offer was exciting" and "Submitting an offer was fun").

4. Pilot Study: Inconsistency in Preferences Revealed by Entry and by Offer Amounts

Our pilot study followed the above experimental paradigm but allowed any offer amount to be submitted. 105 subjects participated in the pilot study. We found some of them made an offer that exceeded $v-c$ at least once, guaranteeing a loss in that round. This behavior was highly concentrated in (v,c) cells of the pilot design with $v<c$, where entrants have no other choice but to lose money. To rule out such seemingly

[5] https://osf.io/q3ysf/?view_only=47df578c500b4ffc91aa06868e0e1e76

irrational behavior as a potential explanation of our findings, we decided to subsequently only allow offers that exceed $v-c$, and therefore to not allow entry when $v<c$.

More importantly, the pilot study found the preference inconsistency that motivates this paper: consistent with risk aversion, offers far exceeded the risk-neutral prediction of $v/2$, whereas the entry behavior we observed was broadly consistent with risk neutrality. In our first experiment, described next, we replicate this puzzling finding while explicitly disallowing offers above $v-c$ and testing for decision difficulty as a potential explanation.

5. Experiment 1: Effect of Decision Aids on Offers and Entry

We designed our first experiment to replicate the preference inconsistency we found in our pilot study and to test whether decision difficulty could explain it. To manipulate decision difficulty, subjects were randomly assigned to one of two conditions with different amounts of interactive feedback: the control condition (*no aid*) without interactive feedback, and the treatment condition (*aid*), in which the interface showed the probabilities of acceptance and rejection, as well as the associated contingent monetary payoffs, for any candidate offer amount typed into the “Enter your offer” box. Subjects in the aid condition were able to enter multiple amounts in the offer dialog box, consider the feedback, and decide which, if any, offer to make. See the right side of Figure 2 for the wording of the decision aid, and a temporal overview of the experimental design, which we explain next. A third condition provided the subjects with the probabilities of acceptance and rejection only. The effect of this intermediate aid condition was similar to the effect of the complete decision aid, and to simplify the exposition, we do not report it in detail (details available from authors upon request).

Note that our decision aid provides both interim computations needed to solve equation 1—the probability and the conditional surpluses. Depending on the real-world institution, the seller may be able to provide these computations, as well: regarding the probability computation, an auctioneer or a name-your-own-price seller often has historical data that can inform the probability of any given bid magnitude winning the auction. For example, Priceline can analyze a multitude of offers it receives on a given hotel stay, and estimate the chance of acceptance as a function of offer magnitude. Regarding the surplus computation, an auctioneer often knows the competing posted-price offer for the same good, and this alternative purchase opportunity becomes the effective valuation of all bidders who can afford the outside price. When the cost of participation is primarily monetary, such as the cost of transportation to the physical auction site or a fee the auctioneer is charging, the auctioneer knows the submission cost, as well. In summary, the decision aid we implement can be a realistic option for real-world sellers, and so analyzing its effect is interesting in its own right.

5.1. Experiment Design

All conditions used a "simultaneous" decision architecture that solicits offer and entry decisions on a single screen (see left side of Figure 2).

Figure 2: Experiment 1: Overview of Experimental Design

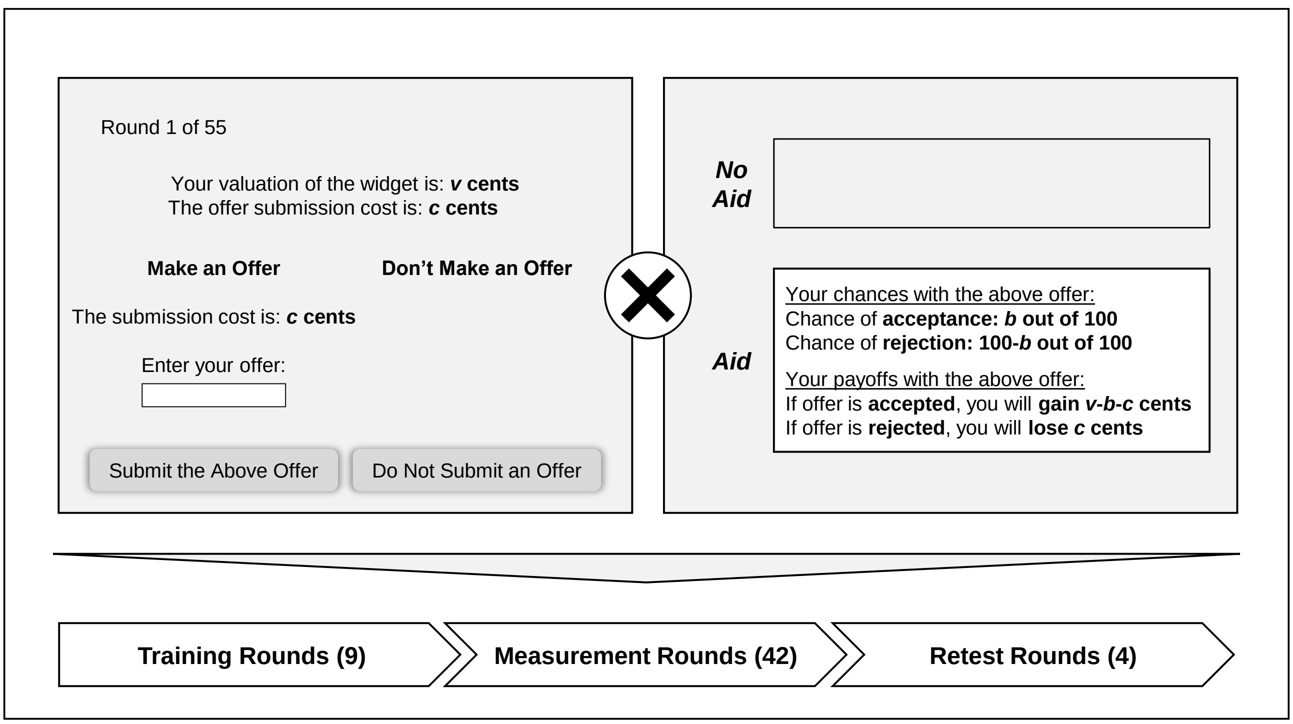

After qualifying for the experiment by correctly answering all five screening questions (see instructions on the Open Science Framework project directory for the questions), all subjects experienced 55 rounds of the costly offer task, followed by a short demographic and preference survey. The 55 rounds were divided into three blocks (not marked in any way to subjects): 9 training rounds, 42 measurement rounds, and

4 retest rounds.

The 9 training rounds exposed subjects to the following (v,c) treatments, selected to provide experience with the full range of possible values and tradeoffs: (10,0), (25,1), (40,1), (55,16), (70,32), (85,1), (100,4), (10,8), (55,4). The 42 measurement rounds treated each subject to a 7(v) × 6(c) full-factorial design involving all possible combinations of $v \in \{10, 25, 40, 55, 70, 85, 100\}$ and $c \in \{0, 1, 4, 8, 16, 32\}$.[6] In the three ($v$,$c$) cells that involve $v<c$, subjects were not allowed to make an offer, because doing so would guarantee a loss. In the rest of the design cells, entrants were not allowed to offer more than $v-c$ for the same reason. Hence, we had 39 active measurement rounds for analysis. Finally, the 4 retest rounds, which—to facilitate comparability— included no decision aids regardless of the experimental condition, presented subjects with four (v,c) conditions[7] in which a risk-neutral agent would enter but a sufficiently risk-averse agent would not.

Within each block, the presentation order was randomized across subjects. After the 55 rounds, subjects completed an incentivized paired-lottery task adapted from Holt and Laury (2002), responded to the subjective risk-taker scale by Dohmen et al. (2012), and indicated their agreement on 7-point scales (from 1 = *not at all* to 7 = *very much*) with statements about what they were thinking when they made or considered making an offer to the seller (see section 3). Finally, they provided their demographic information.

In return for completing the experiment, subjects received a base payment of $2.50. In addition, they received their final account balance (with accounts initialized with $1) and a payoff from a randomly selected lottery of the Holt and Laury (2002) task as a bonus payment after completing the experiment. A total of 217 US residents (M_{age} = 38.59, SD_{age} = 11.84; 38.71% females) passed the screening questions and

6 The (v,c) levels were designed as follows: We set the maximum valuation to 100 cents to enable subjects to easily understand acceptance probabilities (offer amount = probability in %). From there, we selected the additional valuation levels by subtracting a constant amount. To have some valuations that do not end in zero, we used increments of 15. Given these valuations, a cost of 32 is clearly too high (expected surplus of −7 cents with v=100), and thus, we used it as the highest cost. An offer cost of 16 is too high if people are very risk averse, and thus, we used it as one of the levels. We wanted to include an offer cost of 0 as a control, and 1 as the smallest possible positive cost to test for knee-jerk aversion to paying for offer submission. Finally, we also included 4 and 8 as non-negligible intermediate values to measure entry indifference with intermediate valuation levels relevant to risk-averse buyers.

7 The retest cells were (85,16), (55,4), (70,8), (40,1), with expected risk-neutral earnings between 2.1 and 4.2 cents.

participated in the experiment, 100 of whom had access to the decision aid.[8] The average subject earned about $5.04 and took about 26 minutes (median 23 minutes) to complete the entire experiment.

Table 1: Entry and Offers in Baseline Condition (No Decision Aid)

		Submission Cost					
	Entry Probability	**0**	**1**	**4**	**8**	**16**	**32**
Valuation	**10**	91%	44%	9%	0%		
	25	93%	68%	31%	8%	1%	
	40	97%	84%	44%	20%	6%	2%
	55	98%	95%	64%	26%	8%	3%
	70	100%	97%	82%	48%	16%	3%
	85	99%	99%	85%	64%	27%	5%
	100	99%	98%	93%	84%	51%	5%
	Offers Given Entry (Means)	**0**	**1**	**4**	**8**	**16**	**32**
Valuation	**10**	6	6	5	n/a		
	25	18	18	16	13	8	
	40	30	29	28	28	14	6
	55	42	42	41	37	32	13
	70	55	55	56	51	42	27
	85	67	67	66	63	55	37
	100	75	75	76	73	68	48

Note: Grey shading indicates the cells in which a risk-neutral person would not enter.

5.2. Results: Condition without Decision Aid ("Baseline")

We begin our discussion of the results with a detailed exposition of the "baseline" results, that is, results in the condition without a decision aid. Table 1 shows the proportion of subjects who entered in each (v,c) cell of the design and the average offers these entrants made. As we now explain, our findings, summarized in Table 1, replicate the internal inconsistency found in the pilot study.

Consider the risk-neutral model $u_i(x) = x$ in equation (1) as the starting point of analysis. The risk-neutral model implies buyers should offer half of their valuations regardless of the submission-cost amount whenever they enter. The average offers shown in Table 1 clearly exceed $v/2$, and clearly depend on the submission cost: for every valuation, the offers tend to decline with cost. Figure 3 shows a boxplot of "free" (zero-cost) offers and demonstrates that not only the population average but also almost all individual offers far exceed half of the valuations.

[8] The screening questions that check understanding disqualified 45% of invited respondents who consented to participate in the study.

Figure 3: Free Offers in Baseline Condition (Simultaneous Architecture, No Decision Aid)

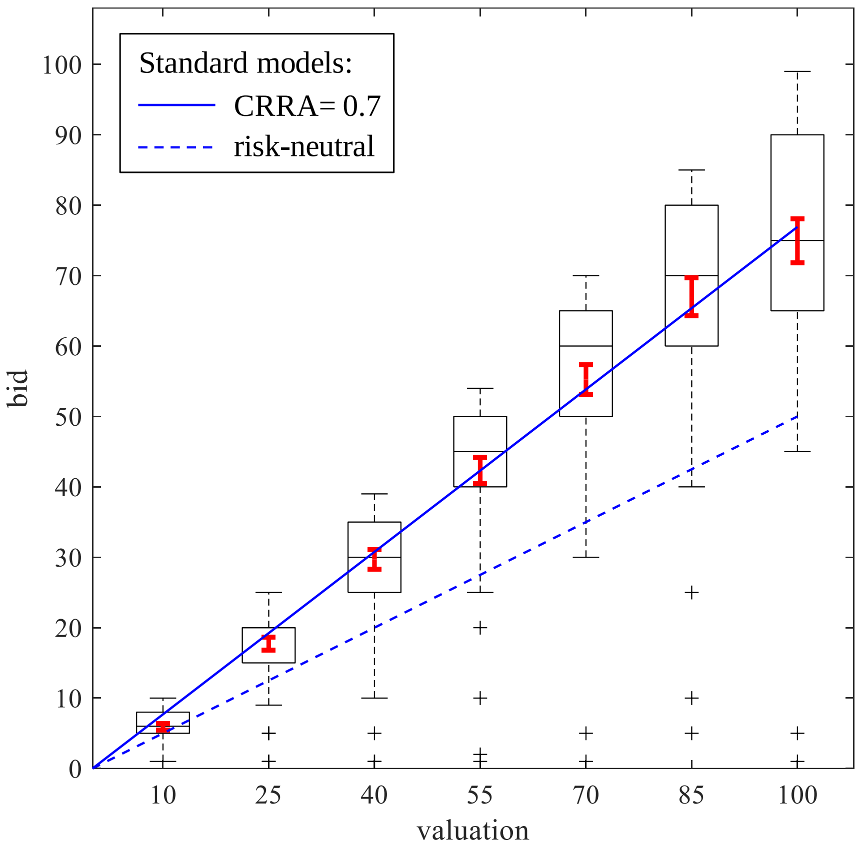

Note: The boxplots illustrate the distribution of submitted offers at each valuation level when *c*=0. The thicker (red) error bars represent 95% confidence intervals. The dashed (blue) line shows the optimal offer function by risk-neutral buyers. The solid (blue) line shows the optimal offer function by risk-averse buyers with CRRA=0.7 and *W*=0.

The risk-neutral model thus does not fit the offer data well (and neither does any other model in which entry costs are effectively sunk, e.g., the additively separable risk-averse model of Smith and Levin (1996)). By contrast, the standard CRRA risk-averse expected-utility model fits at least the average free (*c*=0) offers well: when offers are free and *W*=0, the model has a closed-form linear solution for optimal offers of $v/(2-R)$, and Figure 3 shows the average free offers are indeed approximately linear in valuations, and the model with *R*=0.7 fits these averages well. In the Web Appendix, we describe our maximum likelihood estimation of risk preferences performed at the individual level to account for preference heterogeneity. This rigorous estimation confirms the evidence from Table 1 and Figure 3 that the offer amounts we observe are consistent with risk aversion among most of our subjects.

Although the risk-neutral model does not fit the offer data, it is good at predicting the average entry behavior if we allow for some noise, following the standard random utility framework. The entry predicted by the risk-neutral model is shown by the shading in Table 1: the model predicts entry in the unshaded cells. Note all but two (v,c) unshaded cells have actual entry probabilities above 50%, and all shaded cells have actual entry probabilities below 50%. By contrast, the CRRA expected-utility model calibrated on free offers underpredicts entry: it contrasts with the risk-neutral model by predicting no entry in four additional cells of our design (specifically, 55,4; 70,8; 85,16; and 100,16)—cells with healthy participation in the experiment. In the Web Appendix, we again solidify this simple visual analysis with a statistically rigorous model. Specifically, we use a logistic regression of entry on an intercept and the risk-neutral expected surplus (the objective of a risk-neutral buyer), and we find that the estimated intercept is almost zero while the coefficient on the expected surplus is positive and large.

The above seeming internal inconsistency of our subjects is reminiscent of the "excess entry" inconsistency in auction settings (Palfrey and Pevnitskaya 2008; Ertaç et al. 2011) in the sense that consumers enter more often than their preferences calibrated on bids ("offers" here) suggest. But our data is not consistent with the explanations proposed in that literature: First, the offer data are not consistent with subjects deriving additional utility from winning (Ertaç et al. 2011): Additive utility of winning would manifest as a positive intercept of the observed offers in Figure 3 and a bunching of low-valuation free offers at their maximum allowed value of v. We observe neither pattern.

Second, the offer data is not consistent with a multiplicative utility from winning: Ertaç et al. (2011) show a multiplicative utility of winning would manifest as free offers following $wv/(2-R)$ in our paradigm. We tested this prediction by regressing free offers on valuation, interacted with the self-reported focus on winning, namely, agreement with "My offer being successful was more important to me than the potential payoff." The coefficient on the key interaction has the "wrong" (negative) sign and is not significant (details in Appendix). In summary, winning does not seem to be a significant incremental motivation in our paradigm, possibly because there is no other bidder to beat, as in auctions.

Third, we also do not find any evidence for joy of playing, namely, the notion that participation in the market has entertainment value, as suggested by Palfrey and Pevnitskaya (2008). The correlation between agreement with the statement “Submitting an offer was fun” and the individual-level probability of entry over the measurement rounds is 0.11, which is not statistically different from zero (p=.22).

Fourth, the inconsistency we find cannot be explained by some sort of a mental shift that occurs after entry but before bidding. In a different experimental paradigm that forced the entry decision to occur before the offer amount decision, one might suspect that participants who have entered might feel that they have already incurred an irreversible sunk cost and hence be more prone to “bidding” higher due to this sunk cost shifting their risk preferences, reference points, or other underlying preference constructs. However, our experimental design is deliberately simultaneous, asking for both the offer amount and the entry decision in a single task (see Figure 2). Therefore, a “Dr. Jekyll first entering and Mr. Hyde then setting offers” explanation is not useful in understanding the results of Experiment 1.

Finally, the inconsistency also cannot be explained by subjects wearing out during the course of many rounds, and gradually resorting to some sort of heuristic or random behavior that happens to be internally inconsistent. When we consider only the first five rounds of the study, both the entry decisions and the offer amounts remain approximately the same as in Table 1, and there is no systematic direction in which they differ (for example, the offers are slightly larger in 55 percent of the cells and slightly lower in the remaining cells). When we split the rounds in half and compare the “early” first 21 to the “late” second 21 rounds, the difference in the average offers across all (v,c) cells is about −1 cent and the difference in the average entry probability across all (v,c) cells is about −3%. There is no systematic shift, and the (v,c) cell-specific differences again having mixed signs without any discernible patterns (detailed data available from authors). More importantly, the evidence for the inconsistency shown above across all 42 rounds is just as strong in the early rounds and in the late rounds. Therefore, the inconsistency we find is not somehow a mechanical product of our design with multiple rounds.

With the existing and mechanical explanations of similar inconsistencies not working in our context, we set out to test new explanations, the first of which is decision difficulty discussed next.

5.3. Results: Decision Aid Impacts Entry but Leaves Offer Amounts Unchanged

To what extent are the baseline findings outlined above driven by computational difficulty of the task? Specifically, does the seeming inconsistency between entry and offers diminish when subjects do not have to calculate probabilities and payoffs in their heads? We designed the interactive decision aid to answer these questions. Surprisingly, we found our decision aid accentuated the puzzle instead of alleviating it.

The table titled "Experiment 1" in the Appendix shows the entry and offer data in both aid conditions, as well as the difference in both sub-decisions between the conditions. We find no significant effects on offer amounts. Specifically, the 14 (v,c) cells that involve at least 75% of the subjects entering in both conditions, and hence are not subject to much sample selection driven by any difference in entry behavior, show no statistically significant differences between offer amounts with and without the aid.[9] In 10 of the relevant (v,c) cells, the average offer amounts are within a cent of those in Table 1, and the remaining averages are within a few cents. The baseline analysis of offer amounts being consistent with risk aversion thus applies with the decision aid, as well.

In contrast to the insensitivity of offer amounts to our decision aid, we find it has a large effect on entry. If computationally constrained *risk-averse* subjects in the baseline condition were simply making a mistake at the entry stage of the normative process and our decision aid corrected this mistake, we should have found a *reduction* in entry (specifically, along and under the risk-neutral indifference line in the (v,c) space). Instead, the aid affects entry non-monotonically: it decreases entry when both valuations and costs are small (including zero cost), while increasing entry when valuations and costs are high. Figure 4 illustrates the effect at the (v,c) cell level. To interpret the effect of the decision aid on our understanding of the inconsistency puzzle, let the empirical entry-indifference curve be the curve in the (v,c) space that separates cells with most subjects entering from cells with most subjects staying out. One way to visualize the effect of the decision aid on entry is the empirical entry-indifference curve rotating around its midpoint: dropping from the upper-left corner while rising from the bottom-right corner. The indifference curve rotating in this

[9] The relevant cells are c=0, c=1 & v≥55, c=4 & v ≥85, and c=8 & v=100. The conclusion is not sensitive to the 75% threshold, because the only significant differences between average *submitted* offer amounts with and without the aid occur either in cells with very little entry or in cells with a big difference in entry between the two conditions.

way is inconsistent with entry becoming more risk averse as a result of the decision aid: if the manipulation made subjects more risk averse, their indifference curve would pivot downward at positive cost levels but exhibit no change at "free entry" cells with c=0 (When entry is costless, all rational agents should enter regardless of their risk aversion).

Figure 4: Effect of Decision Aid on Entry

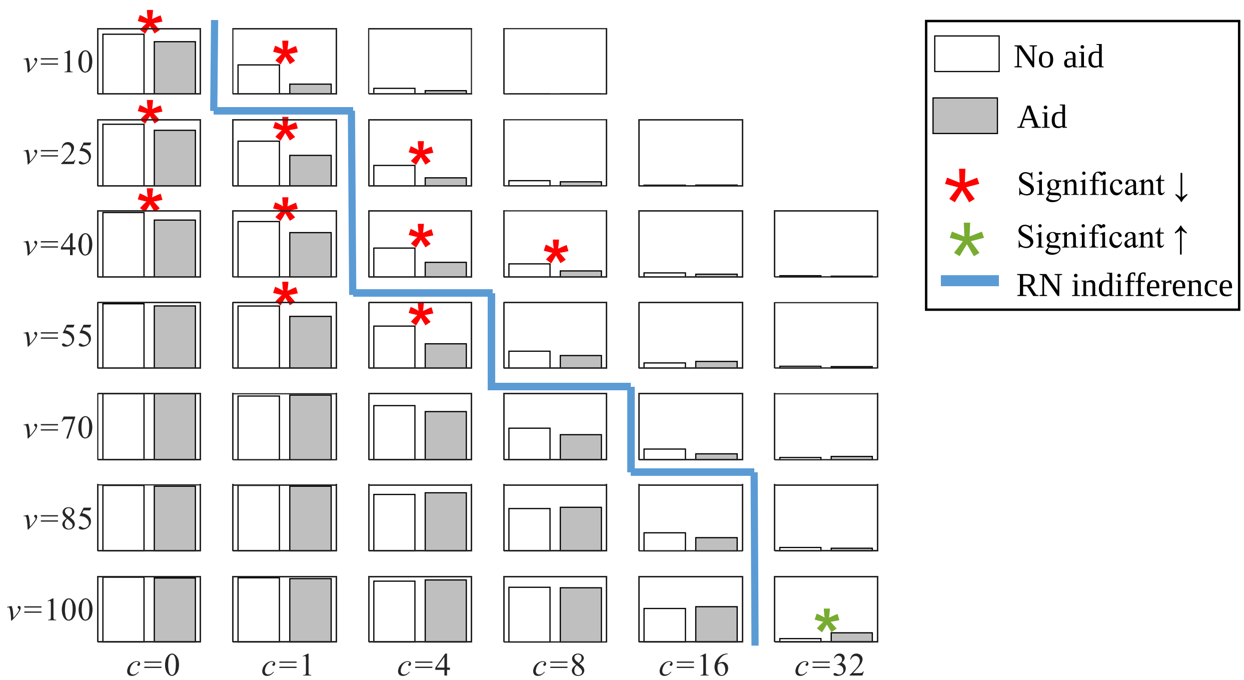

Note: The height of each box corresponds to 100% of subjects entering in the particular (v,c) cell. The * indicates a significant change between the proportions at the 5% level, and the color indicates the direction of the change (red = decrease, green = increase). The thick line is the indifference curve of the risk-neutral model.

Because the seemingly risk-averse offer amounts are unaffected by aid and the aid reduces the fit of the risk-averse entry model, the decision aid accentuates the inconsistency puzzle instead of resolving it. What is the underlying mechanism behind the surprising entry effect we observe? Some of the effect can be explained by the aid correcting a systematic misperception of probabilities. Figure 5 takes each submitted offer's perceived probability of success and plots it as a function of the true probability of the offer's success by decile bin. It shows that without the aid, the subjects are optimistic when their valuations (and hence their offers, bounded above by $v-c$) are small and pessimistic when their valuations (and hence most offers) are large. By bringing subjective beliefs more in line with truth in this way, the aid should increase entry when valuations are large and reduce entry when they are small—precisely the non-monotonic effect we find.

Figure 5: Effect of Decision Aid on Subjective Probability Beliefs

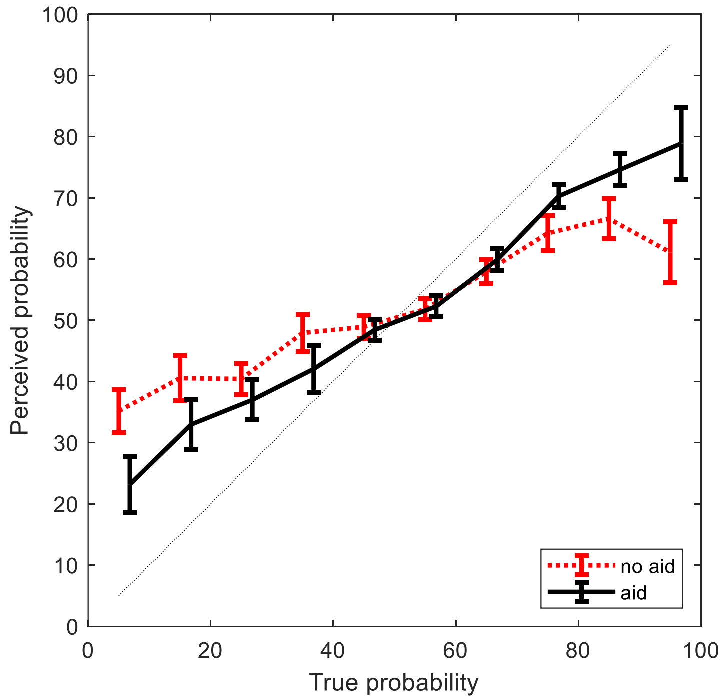

Note: Error bars are 95% confidence intervals.

While the belief correction shown in Figure 5 does explain some patterns in Figure 4, it does not explain the reduced entry when costs are zero and does not justify more entry when the expected consumer surplus is negative (e.g., in the (100,32) cell). Part of these deleterious effects may be due to the aid guiding subjects to enter more when winning probabilities are high and enter less when they are low, irrespective of cost. The attitude questions at the end of the experiment corroborate this interpretation of the aid putting too much emphasis on getting the offer accepted: compared with subjects with no aid, significantly more subjects with the aid report wanting to win more.[10] More importantly for our goal of developing a behaviorally realistic model of the overall behavior and explaining the seeming inconsistency between offer and entry behavior, offer magnitudes do not reflect the large improvement in reported probability beliefs.

If subjects followed the normative process outlined in Figure 1a, the aid's alignment of perceived probabilities with true probabilities should increase offers (relative to the baseline condition) when the

[10] No aid (M = 3.21, SD = 1.72), aid (M = 3.82, SD = 1.93). The p-value of the relevant t-statistic is .02.

baseline beliefs are too optimistic and reduce offers when the baseline beliefs are too pessimistic. Yet, we find neither effect, indicating that our subjects behaved as if they formulated offer amounts without considering acceptance probabilities. Importantly for consumer welfare, the decision aid also did not actually help subjects earn more money in the experiment, as we discuss next.

5.4. Results: Decision Aid Has No Impact on Consumer Earnings

We now briefly discuss the expected earnings of our subjects and the profit of the seller our study simulates. To average over the actual realizations of seller price thresholds during the experiment, we take the submitted offers as given and compute the *expected* earnings they imply as $\left(\frac{offer}{100}\right)(valuation - offer) - cost$. A risk-neutral subject would expect to make about 220 cents in the 39 measurement rounds of our experiment. Of the 117 subjects in the no-aid condition, none made more, and most subjects made a lot less: the average earnings were only 96 cents—43% of the theoretical maximum—and the median earnings were only a bit higher at 107 cents. Because our buyers enter on average as if they were risk neutral, the main culprit for their low consumer surplus is clearly their tendency to make high offers upon entry.

The expected earnings vary dramatically across subjects: the standard deviation of total individual earnings within the measurement rounds is 63 cents, and the range of earnings is from −70 cents to +212 cents.[11] Note from the definition of expected earnings, we are averaging over seller cost realizations, so this variation of earnings is solely attributable to differences between strategies of individual participants. Given the effective show-up fee of 350 cents, the highest-earning subjects were thus able to double their earnings compared to the lowest-earning subjects, and earn almost 3 dollars more from participating in our experiment. We propose these are steep incentives for careful bidding, and our design thus does not suffer from the standard Harrison "flat maximum" critique in experimental economics research (Harrison 1989).

We expected the decision aid to help consumers earn more by reducing their offer amounts to be more in line with their risk-neutral entry behavior. Instead, their offers remained the same on average, and the beneficial effects of the aid on entry (e.g., reduced costly entry when valuations are small and expected

[11] Note that the subject losing 70 cents is an outlier – only 6 percent of subjects lose money in the measurement rounds.

returns negative) were offset by the deleterious effects (e.g., reduced entry when offers are free and increased entry when costs are too high for positive earnings). Table 2 shows the increase in average expected earnings is small and statistically insignificant. Given a few very low-earning outliers scattered across the conditions, the median captures the central tendency of the surplus data better than the average, and the median surplus does not show even a directional effect. Note that not finding a better (in terms of earnings) decision aid is not somehow a failure of our design—our goal was to see if helping consumers better understand the decision they were making would resolve the preference inconsistency, not necessarily to develop the best possible decision aid in terms of earnings (risk-averse agents do not maximize earnings, see equation 1).

Table 2: Consumer Expected Surplus, Seller Profit, and Gains from Trade, by Decision Aid

	Consumer Earnings				**Seller Profits**		**Gains from Trade**	
	Median	*SE*	*Mean*	*SE*	*Mean*	*SE*	*Mean*	*SE*
No Aid	107.1	7.3	95.6	5.8	300.7	6.7	396.3	8.8
Aid	105.3	6.7	103.5	5.3	276.4	6.9	379.9	8.7

Table 2 also shows the expected profit of a seller simulated by the computer in our experiments, under the assumption that the submission costs are a pure market friction. Seller profits thus defined decline with the aid faster than consumer earnings rise, and the aid thus does not have a significant effect on market efficiency.[12] Our results thus suggest sellers may not want to offer decision aids, and regulators who care about overall market efficiency may not want to require such aids either.

5.5. Discussion of First Experiment Results

To summarize our findings so far, we find a preference inconsistency reminiscent of the "excess entry" effect in the experimental auction literature in that our subjects submit seemingly risk-averse offers while entering as if they were risk neutral. Our task is a single-agent decision, so existing explanations around strategic competition between bidders and additional utility of winning do not apply here. We also ruled out a new possible explanation that the inconsistency arises from the computational difficulty of the task: contrary to this plausible hypothesis, our decision aid did not resolve the inconsistency puzzle, but

[12] The 24.4 cent decline in seller profits is statistically significant, with a standard error of 9.6 cents ($p<.02$).

rather exacerbated it. A more fundamental generalization of the standard model *is* consistent with our data: the strong effect of our decision aid on entry and subjective probabilities without any matching effect on offer amounts suggests subjects do not follow the normative process outlined in Figure 1a. The same insensitivity of offer amounts to the probability beliefs implies the non-normative sequential ordering of tasks, whereby consumers first figure out whether to enter and then come up with the offer amount, does not fit the data either. Therefore, the explanation of the puzzle may lie in an alternative, *seemingly* separate process (see Figure 1b) instead of in fine-tuning the details of the utility function as proposed in much of the extant literature.

The interface used in Experiment 1 may invite various suboptimal heuristic processes, so people may benefit from a nudge toward the normative process. If such a nudge made them parse the problem correctly, perhaps they would also seem less internally inconsistent, and the entire standard model would explain their behavior better. To examine this candidate explanation, we changed the decision architecture of our interface to walk the consumers through the correct process, as we discuss next.

6. Experiment 2: Nudge Toward Normative Decision Architecture

Basic economics teaches us to solve the problem in equation (1) "backward," starting with determining the offer amount given participation (the subproblem in the square brackets in equation (1)), and then calculating whether the associated expected utility of participation exceeds the utility of staying out. Should the market designer help guide consumers through this "backward solution" logic instead of simply soliciting offers simultaneously with entry the way our first experiment did? Would this decision architecture resolve the inconsistency puzzle? In Experiment 2, we implement a "backward" decision architecture to guide subjects through the two sub-decisions in the correct order as illustrated in Figure 1a.[13] Except for the different decision architecture, Experiment 2 had identical procedures and used the same measures as the

[13] We also experimented with a "forward" sequencing of the two tasks (first entry, then offer). Much like the "backward" sequencing analyzed here, the forward sequencing did not have a significant effect on behavior, so we do not report it in detail. Analogously to Experiment 1, we also included the intermediate aid condition, in which behavior closely mimicked that with the complete decision aid, so we again do not report the results in detail. All results not reported in detail are available upon request from the authors.

"Baseline" (simultaneous) experiment described above. The training rounds also followed the interface of Experiment 1, so Experiment 2 only diverged from Experiment 1 after the training rounds. See Figure 6 for an overview of the experimental design of Experiment 2.

Figure 6: Overview of Experiment 2: Offer, Then Entry (Backward)

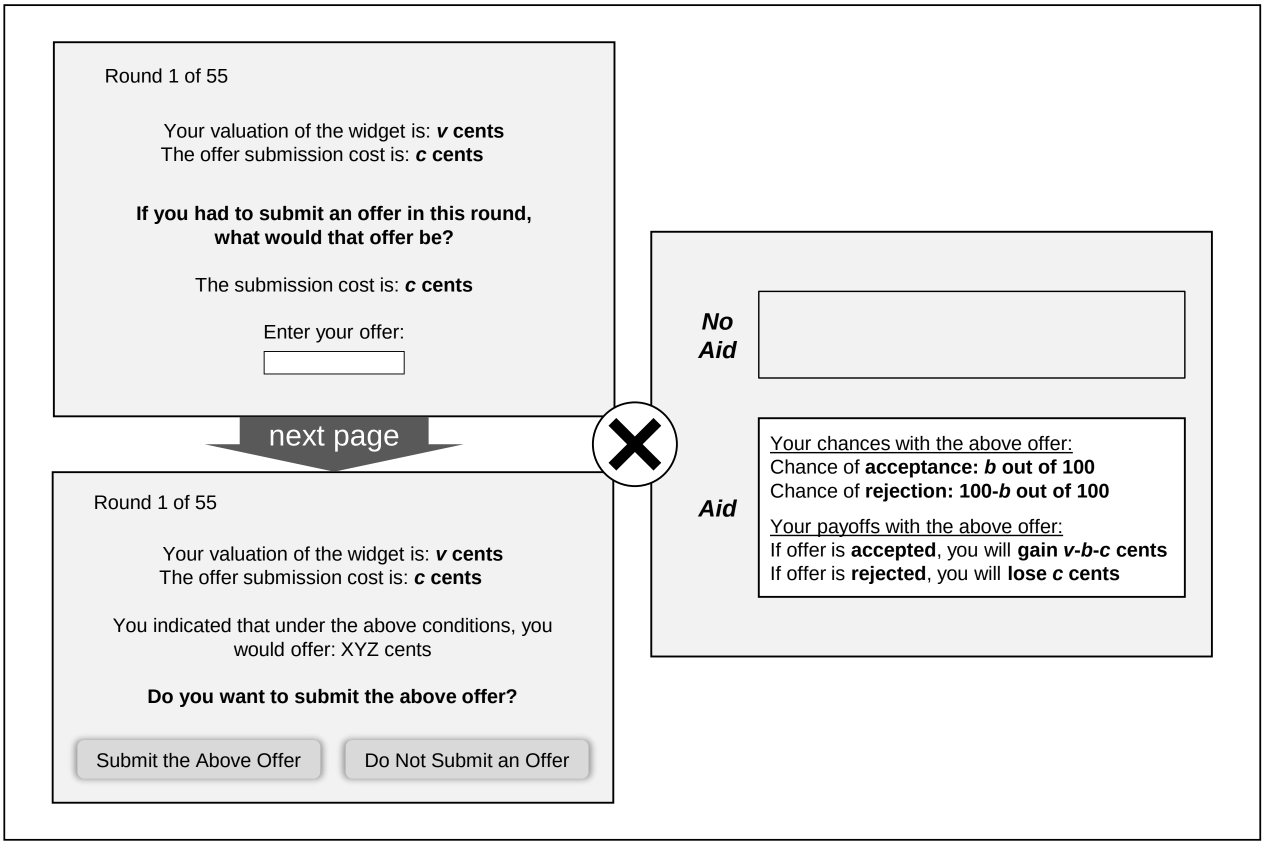

A total of 241 US residents (M_{age} = 39.63, SD_{age} = 12.37; 42.32% females) passed the screening questions[14] and participated in the experiment, taking 31 minutes and earning approximately $5.12, on average. 108 of them had access to the decision aid. The table titled "Experiment 2" in the Appendix shows the difference between Experiment 2 and Experiment 1 for both sub-decisions and both decision-aid conditions.

6.1. Results: Without Decision Aid, the Backward Decision Architecture Has No Effect on Earnings and Almost No Effect on Entry and Offer Magnitude

In the condition without the aid, we found very few (v,c) cells with significant effects of the

[14] The screening questions that check understanding disqualified 53% of invited respondents who consented to participate in the study.

backward architecture on entry and offer magnitudes. Submitted offers generally decline slightly in most (v,c) cells, significantly so in (70,4) by 3.5 cents and in (100,8) by 4.4 cents. These modest declines still result in offers far above the risk-neutral level. The CRRA approximation based on free offers shown in Figure 3 yields essentially the same slope of the average free offers in valuation (not reported in detail). In the Web Appendix, we document that controlling for heterogeneity by individual-level estimation does not change the general conclusion that revealed preferences are very similar between Experiments 1 and 2. The lack of systematic effects on entry or offers means the inconsistency survives the backward-architecture manipulation. Thus, whatever heuristics people employed in the baseline condition do not seem to have been a simple mistake fixable with a nudge toward the normative process. While not resolving the inconsistency puzzle, average high-valuation offers declining by a few cents does imply potential for higher earnings and less risk aversion. Alas, the subjects did not end up earning more than in Experiment 1, because the backward architecture worsens their entry decisions, as we discuss next.

Figure 7 shows the effect of the backward architecture on entry, which increased significantly in three money-losing cells and did not change anywhere else. Our data do not isolate the underlying mechanism by which the backward direction without a decision aid worsens (in terms of earnings) the entry decision. One psychological mechanism consistent with the effect is escalation of commitment (e.g., Staw 1976): *Once I enter a binding offer into the system, I feel that I should see the process through.* Or: *Having already invested effort thinking about the offer, I now want to know the outcome of submitting it.* The overall effect on expected earnings is a statistically insignificant decrease of 2 cents. The slightly worse entry decisions thus offset the beneficial slightly reduced offer amounts.

The backward nudge seems to be unhelpful in one additional respect: many subjects in the backward architecture seem to anticipate their entry decision already at the offer-amount stage, and they enter evidently non-serious offers: in about one third of the overall person-rounds, the offer field was left empty, and in an additional quarter person-rounds, the entered offer was zero. 67% of the subjects left the field empty or entered zero in at least one round. Contrary to the instructions, our subjects thus do not focus only on entry-contingent offer magnitudes at the first stage, and they continue solving the two decisions of the overall problem separately. The manipulation is not only nudging customers toward the normative process, but also seems to be interfering with their natural decision process. Additionally, the overall decisions may be more—rather than less—difficult for subjects in our backward architecture. The longer average completion time certainly suggests the task challenged at least some subjects.

Figure 7: Effect of Backward Architecture on Entry, Condition without Decision Aid

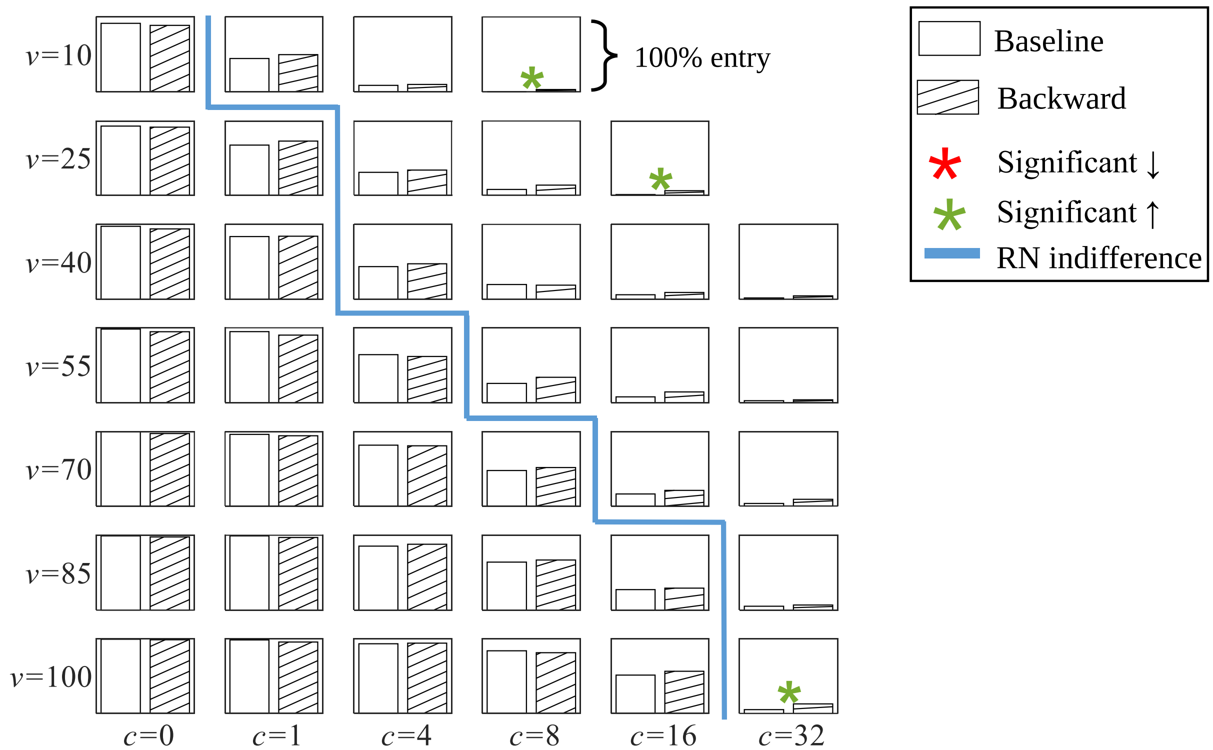

Note: The height of each box corresponds to 100% of subjects entering in that (v,c) cell.

6.2. Results: With Decision Aid, the Backward Decision Architecture Improves Earnings but Preserves the Internal-Inconsistency Puzzle

Once the backward architecture is combined with the decision aid, the subjects finally benefit from the nudge by earning more. Given a few very low-earning outliers scattered across conditions, the median captures the central tendency of earnings better than the average, and the median earnings increased 22% to \$1.31 from \$1.07 (Wilcoxon test p<.01). The behavioral effects behind this large effect on earnings are subtle: no (v,c) cells show a significant entry effect relative to entry in Experiment 1 with the aid, but we see a systematic pattern of small insignificant increases just under the risk-neutral indifference curve. Coupled with the small but systematic reduction in offer magnitude analogous to the reduction observed without decision aid, the overall behavior yields higher earnings. In other words, subjects with decision aids offered less and entered more in the backward architecture, in line with the sign of expected surplus, relative to subjects with the same decision aids in the simultaneous architecture used in Experiment 1.

One way to interpret the implications of Experiment 2 for earnings is that consumers struggle with finding the best amount to offer, and merely focusing on it more does not help them earn more. Only when the focus comes with a "what-if" calculator of acceptance chances and contingent payoffs do their decisions improve. Although the combination of decision aid and backward nudge increases consumer earnings, the manipulations analyzed in Experiments 1 and 2 fail to explain the inconsistency puzzle: even in the backward-architecture condition with decision aid, offer magnitudes still far exceed $v/2$, and entry behavior is still approximately risk neutral.

The non-serious offers many subjects submitted in the offer stage of Experiment 2 suggest the nudge may not have been strong enough to encourage solving the problem backwards. Would a stronger manipulation toward a backward-solution strategy finally resolve the inconsistency puzzle by lowering offers and/or reducing entry along the risk-neutral indifference curve? We address this possibility in our last experiment.

7. Experiment 3: Forced Normative Decision Architecture

Experiment 2's nudge toward the normative process did not result in most subjects following the backward process. Instead of solving the inside optimization problem of equation (1) before moving on to the entry decision, many subjects at least occasionally decided not to enter without solving for the optimal contingent offer (i.e., they did not input a serious offer when asked). This behavior is consistent with (a small) cost of thinking involved in deriving the offer. Might this cost of thinking be the source of the internal preference inconsistency we found in Experiments 1 and 2? If yes, forcing the subjects to follow the normative process should resolve the puzzle. If the associated realignment of offer amounts with entry decisions occurs primarily via a reduction of offer amounts, forcing the subjects to follow the normative process should also increase subjects' earnings by better aligning their behavior with the risk-neutral model. In Experiment 3, we test these possibilities by first eliciting separately incentivized, automatically submitted offers (i.e., offers that solve the inside optimization of equation (1)) and then observing how each subject makes the entry decision when given a seemingly exogenous (though actually their own) offer amount as fixed and asked a simple yes/no entry choice. To contrast the suggested "backward" treatment from Experiment 2 with the forced "backward" treatment of Experiment 3, we hereafter refer to the Experiment 2 architecture as "suggested backward." See Figure 8 for the illustration of the experimental design.[15]

A total of 242 US residents (M_{age} = 40.48, SD_{age} = 12.26; 53.72% females) passed the screening questions[16] and participated in the study. They were randomly assigned to the same two decision-aid conditions as in Experiments 1 and 2. 101 of them had access to the decision aid. The exact procedure was as follows: the experiment started with the same 9 training rounds we used in the previous two experiments. After that, subjects experienced 21 rounds in which they had to submit an offer to the seller ("forced-entry" rounds). The 21 rounds corresponded to a subset of the 42 full-factorial (v,c) cells selected to be along the diagonal of the (v,c) space, as well as in the lower-left triangle of it, to separate the effect of valuation from the effect of cost.[17] This block was followed by a filler task with demographic questions. In the next 21 rounds (fixed offers), we used subjects' offers from the forced-entry rounds, presented them in random order as exogenously given (one at a time and without informing the subjects that they were, in fact, their own), and asked subjects to decide whether they wanted to submit them. We chose to pipe subjects' own offers to them to control for preference heterogeneity and allow a direct comparison of the entry behavior with equivalent design cells in Experiments 1 and 2. To avoid any issues of sunk-cost fallacy or escalation of

[15] After completing the experimental procedure of Experiment 3, each subject also experienced the same problem framed as a lottery. We do not report the results of that framing in this paper, because it is not related to our thesis. Analogously with Experiments 1 and 2, Experiment 3 also had an intermediate decision-aid treatment with only the probability of acceptance shown. All results not reported in detail are available upon request from the authors.

[16] The screening questions that check understanding disqualified 54% of invited respondents who consented to participate in the study.

[17] Specifically, participants were presented with the following (v,c) pairs: (10,0), (25,0), (40,0), (70,0), (100,0), (10,1), (25,1), (40,1), (55,1), (70,1), (85,1), (25,4), (40,4), (55,4), (85,4), (55,8), (70,8), (70,16), (85,16), (100,16), (100,32).

commitment, we did not inform subjects about the precise source of the exogenous offers. Instead, we merely described the offer amounts as "predetermined" at the time of their presentation.

Figure 8: Overview of Experiment 3: Offer Only, Then Entry Only (Forced Backward)

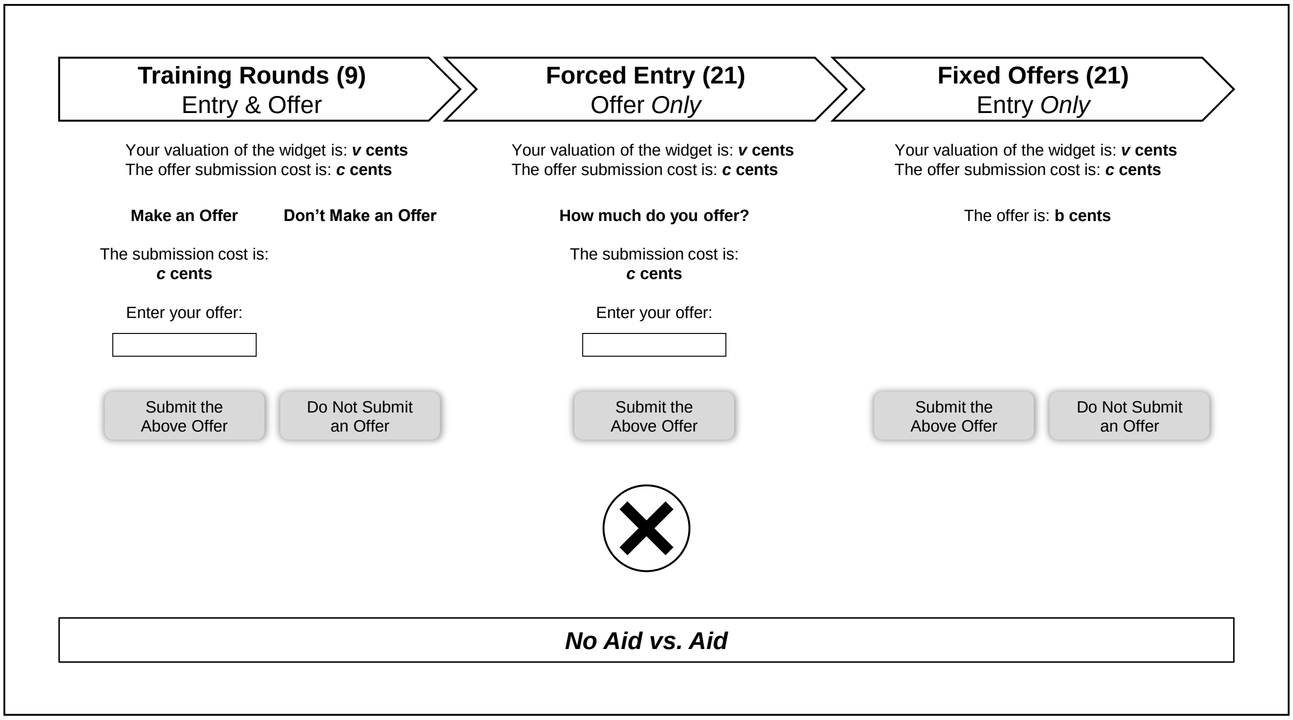

As in Experiments 1 and 2, whenever a subject decided to submit an offer, they were also asked to indicate their prediction of the likelihood that their offer would be accepted. Following the fixed-offer rounds, subjects reported the extent of their agreement with the same statements we used in the previous experiments regarding what they were thinking when they made or considered making an offer to the seller (see section 3). As with the other experiments, subjects received a base payment of $2.50, and they received their final account balance (including the $1 starting balance) from all blocks as a bonus payment after completing the experiment.

7.1. Results: Forcing the Normative Process Instead of Merely Suggesting It Increases Offer Magnitudes, but Leaves Entry Unaffected

We start the discussion of the experiment results in the condition without the decision aid and compare the observed behavior with that in Experiment 2 to isolate the effect of forcing the normative process instead of merely suggesting it. The table titled "Experiment 3" in the Appendix shows the difference between Experiment 3 and Experiment 2 for both sub-decisions and both decision-aid conditions. The submitted offers increase relative to Experiment 2 levels in all 21 (v,c) design cells, and the increase is statistically significant in 16 of the cells. On average, across all 21 cells, submitted offers increase by about 8% relative to those in Experiment 2. Figure 9 shows the pattern of offer increases among all entrants.

Figure 9: Effect of Forcing Backward Process on Offers, Condition without Decision Aid

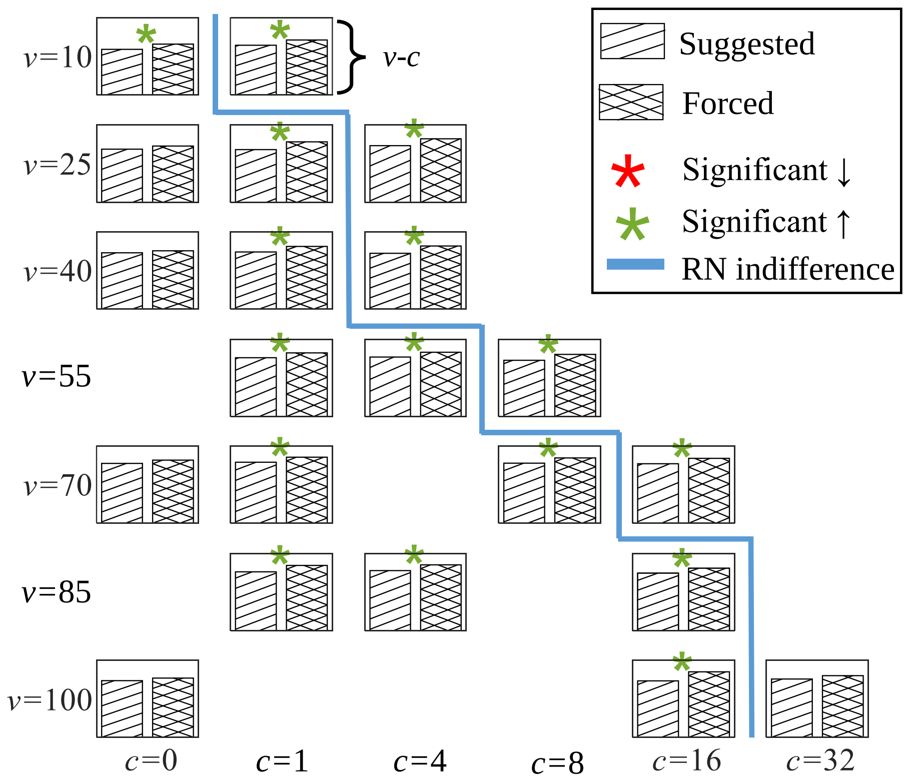

Note: The height of each box corresponds to $v-c$, i.e., the upper bound on admissible offers.

The dramatic increase in offers is not accompanied by a systematic change in entry (during the "Entry *Only*" block of the study flow; see Figure 8): entry probabilities remain nearly unchanged in most cells, and observed changes involve a mix of signs, with only one significant decrease in the (40,1) cell and only one significant increase in the (100,32) cell (details in the Appendix).

Providing the subjects with the decision aid attenuates the effects described above but does not cancel them out completely. Specifically, submitted offers still generally increase by about 8%, on average, across all 21 cells, but the increase is only significant in eight cells of the design—cells approximately along the risk-neutral indifference curve (see "Experiment 3" in the Appendix for details). Entry-probability changes continue to have mixed signs, but the magnitudes of the changes are smaller, and none are statistically significant. The differences between Experiment 2 and 3 are less pronounced than without the aid, but offers still rise relative to those in the suggested backward condition (see Appendix for details).

Having described the effects of the key manipulation of Experiment 3, we now turn to the implications for the inconsistency puzzle. Because offers increase relative to Experiment 2 but entry remains about the same, forcing the normative process exacerbates the inconsistency puzzle instead of resolving it. The same pattern of effects also implies the entry decisions in Experiments 2 and 3 do not seem to take into account offer magnitudes as equation (1) suggests they should. Under the expected-utility model, for example, an increase in offers would indicate an increase in risk aversion, and so should be accompanied by a decline in entry along and under the risk-neutral indifference line (see prediction (b) in section 2). Other utility specifications, for example, a prospect-theoretic value function in place of u in equation (1) may have more subtle predictions, but all would imply a systematic co-movement of offers and entry behavior, contrary to an effect on only one of the decisions we find here.

7.2. Results: Forced-Backward Architecture Reduces Consumer Earnings Relative to the Suggested Backward Architecture, Especially When Subjects Have No Decision Aid

Figure 10 shows the median consumer earnings in Experiment 3 and the matched cells from Experiments 1 and 2. Specifically, the two leftmost groups of bars represent Experiment 1 and 2 consumer earnings in the 21 (v,c) design cells matched to the cells used in Experiment 3. Note that comparing the two leftmost groups of bars with each other replicates the same pattern of earnings as in all 42 cells of those studies. At least in terms of this pattern, the 21 (v,c) design cells used in Experiment 3 thus constitute a representative subset of all cells.

Figure 10 shows the forced-backward treatment did not help subjects earn more money than the suggested-backward treatment. Rather, we see the opposite: forcing the subjects to give a serious offer in the first stage reduces their ultimate consumer earnings by 45% compared with subjects who were only casually to enter their offer in the first stage (a highly significant effect, Wilcoxon test $p<.01$). Even with the decision aid, forcing the normative process on consumers is slightly worse for them than merely suggesting it (the 7% decrease is marginally significant, Wilcoxon test $p=.06$). This pattern is consistent with subjects struggling to think through contingent payoffs in the first (offer) stage, possibly focusing on other objectives in that moment, such as getting their offer accepted, and forgetting about the cost of submission. We document this mechanism in our last result next.

Figure 10: Consumer Earnings in 21 (*v*,*c*) Design Cells Comparable across All Studies

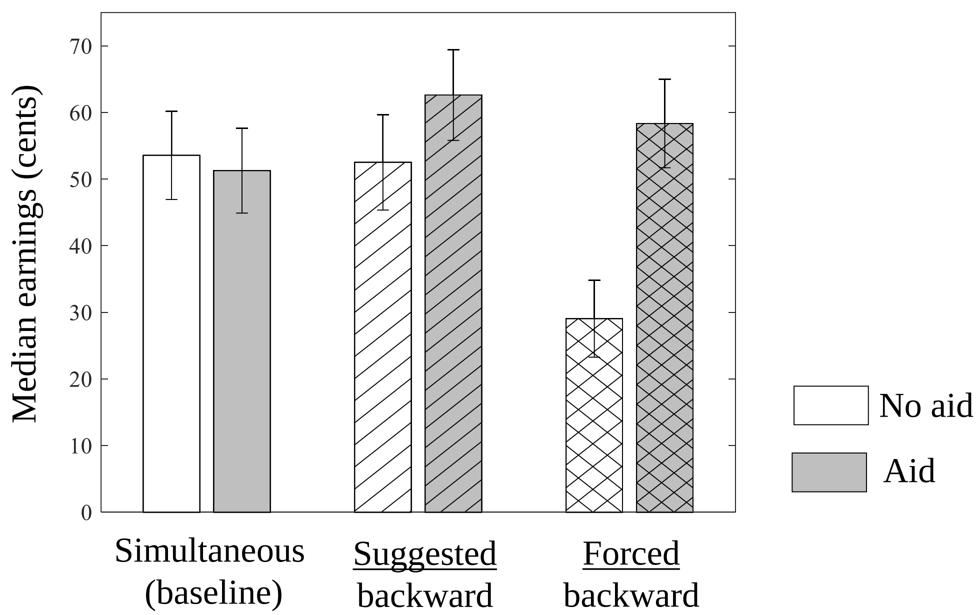

Note: The two leftmost groups of bars represent the Experiment 1 and 2 consumer surplus in the 21 (*v*,*c*) design cells matched to the cells used in Experiment 3. Error bars represent the 95% confidence intervals.

7.3. Results: Backward Architectures Reduce the Extent to Which Subjects Take into Account Submission Cost When Formulating Offers

Table 3 zooms in on two specific cells while conditioning on buyers who enter in both of them. The two cells highlighted in Table 3 are instructive because they allow us to test whether buyers are taking into account the submission cost when formulating their offers above and beyond the crude upper bound $v-c$ we impose on offers. Because $100-16=85-1$, the buyer's problem involves the same upside, and offers should thus be higher in (100,16) than in (85,1) as long as buyers take into account the submission cost when formulating their offers. To see why, let x=v−c be the net valuation of receiving the widget. Then, the offer problem can be reformulated as

$$b(x,c) = \underset{b\geq 0}{argmax} \left(\frac{b}{100}\right) \underbrace{u(x-b)}_{u_G>0} + \left(1-\frac{b}{100}\right) \underbrace{u(-c)}_{u_L(-c)<0}, \qquad (2)$$

where u_G does not depend on c. The first-order condition is $b = \frac{u_G - u_L(-c)}{u_G^{'}}$, so $\frac{db}{dc} \propto u_L^{\prime}(-c) > 0$. In other words, in any set of (v,c) cells such that v−c is constant and c varies, the optimal offers rise as c rises, as long as losing higher c hurts more (irrespective of risk preferences or exact shape of u).

To apply the above idea, and to avoid selection out of the higher-cost cell based on risk preferences, considering only buyers who make offers in both cells is important. Unfortunately, this conditioning makes the test underpowered in conditions other than forced backward, with only about 30 subjects entering in both cells in each sequencing-aid combination. However, the forced-backward condition solicits offers from over 100 subjects per aid condition, and Table 3 shows their offers do not depend on the submission cost. By contrast, the offers clearly do depend on costs in the predicted direction in Experiment 1 despite the test being underpowered.

Table 3 also clearly shows how much the submitted offers in Experiment 3 exceed those in Experiment 1. Among entrants without decision aids, Experiment 3 involves the highest offer amounts among all architectures considered in this paper in both cells. The same is true in the most lucrative (100,0) cell (not shown in table). Together, these increases in offer amounts explain the dramatic decrease in earnings documented in the previous subsection of the paper.

Table 3: Average Offer among Subjects Who Enter in Both (100,16) and (85,1) Cells

	Decision Aid	**Simultaneous (Experiment 1)**	**Suggested backward (Experiment 2)**	**Forced backward (Experiment 3)**
(85,1)	no	61.3	59.7	71.0
	yes	65.3	62.6	65.9
(100,16)	no	67.9	62.3	71.2
	yes	67.4	63.9	65.5
Diff	no	6.6*	2.6	0.2
	yes	2.1	1.3	−0.4

Note: * indicates significant difference at the 5% level.

8. General Discussion

Using a series of incentivized experiments, we examine how consumers make costly price offers to sellers. We document a persistent within-task preference inconsistency whereby subjects enter as if they were approximately risk neutral but place offers consistent with substantial risk aversion. This inconsistency is persistent in that neither of our manipulations aimed at helping consumers to make decisions reconciled the two sub-decisions of the overall decision with each other. Specifically, we tested two types of such manipulations: an interactive decision aid (e.g., Häubl and Trifts 2000) that concretizes the consequences of a potential offer amount, and a sequential-decision architecture (e.g., Johnson et al. 2012) that guides decision-makers through the normative decision process. We find the decision aid affects entry consistently with the improved probability beliefs it causes, but it leaves the offer amounts unaffected despite a clear prediction that they too should change in response to the changing probability beliefs. But offer amounts are not set in stone—the more forceful version of our sequential-decision architecture increases them while leaving the entry behavior unaffected. This insensitivity of entry is again internally inconsistent with the associated increase in offer amounts.

Taken together, our experiments reveal that, instead of following the normative backward-solution process, people treat the two decisions of the costly offer task *as if* they were unrelated decisions and process them separately. Figure 1b illustrates the separate processes along with the two moderators we document (decision difficulty and forced sequential architecture).

The seemingly separate nature of the entry and offer processes that our findings suggest also recasts the preference inconsistency as more of a feature and less of a bug. Whereas previous research assumed the inconsistency arises from consumers making a mistake in one of the decisions (see, e.g., the "excess entry" nomenclature for a similar inconsistency in Palfrey and Pevnitskaya, 2008), our subjects did not become more internally consistent when they received computational support or were guided through the normative process. Instead, both of these interventions designed to help consumers actually exacerbated the inconsistency. Thus, this form of preference inconsistency appears to be a robust property of consumer decision-making and should therefore be incorporated into behaviorally realistic models.

The separate nature of the entry and offer processes is a challenge to the structural program that attempts to simulate counterfactual behavior using estimates of underlying preferences (reviewed, e.g., in Ackerberg et al., 2007, and Krasnokutskaya and Seim, 2011). The key idea of the structural program is that one can estimate a model of stable underlying preferences, and simulate counterfactual behavior using the estimate. Our results call into question any attempt to estimate consumer preferences from either entry or bidding behavior alone, and then making predictions about consumer behavior in the more complete setting of bidding with costly entry. Even estimation of standard models on complete data is problematic because the models do not fit the complete behavior well. Our results imply that counterfactual entry needs to have its own *separately estimated* model, not a mathematical construction based on a model estimated from bids alone. More generally, the assumption that a single set of preferences determines both bidding and entry is too strong. Instead, behaviorally realistic models need to allow for separate preferences and estimates. All this places a higher burden on data than previously thought.

On the left side of Figure 1b, we highlight that no single set of underlying preferences may be giving rise to mutually consistent behaviors in both decisions of the task. But we may have not yet understood how such theoretically appealing single underlying preferences get modulated and expressed into the two seemingly unrelated behaviors. Our findings thus present a challenge to modelers of offers and bidding with costly entry to develop new behaviorally realistic structural (i.e., based in some deep fixed preferences and involving some sort of consumer optimization) models not constrained by the normative structure of Figure 1a.

The above conclusion regarding the entry decision and the offer formulation operating as separate processes fits all our data and rationalizes the pattern of moderation we document. In a post-hoc exercise reported in the Web Appendix, we have also explored the potential of a more flexible form of utility fitting the data while remaining within the normative structure of Figure 1a. We found that preferences following a prospect-theoretic model fit the data better than a more restrictive expected-utility formulation with CRRA (with both specifications estimated at the individual level to allow for preference heterogeneity). The more flexible model connects the two decisions of costly offer behavior by allowing S-shaped utility functions with downside neglect—a tendency to care more about gains than about equivalent losses (i.e., a convex kink at zero). Although S-shaped preferences with convexity in the loss domain resemble prospect theory, downside neglect is the opposite of prospect theory's loss aversion. We therefore find mixed evidence of prospect-theoretic preferences fitting our data within each experiment. Across experiments, subjects with decision aids have more markedly S-shaped preferences, providing further evidence that decision aids do not improve the fit of the standard expected-utility model to the data. Importantly, no amount of flexibility of the utility function in Figure 1a can explain the qualitative patterns of moderation we document across experiments, namely, the fact that each effective moderator affects only one of the two sub-decisions of the overall decision.

In addition to informing future modeling decisions and re-interpreting the preference inconsistency in terms of a non-normative parallel process, our results also have immediate managerial implications for market designers who need to design the interface they provide to consumers. In our baseline condition that mimics a simple real-world interface, our subjects earn only about half the theoretical maximum monetary payoff. The half of potential consumer surplus left on the table leaves ample room for improvement, and both of our moderator treatments should help consumers earn more if their behavior in the baseline condition is some sort of a mistake or a bounded-rationality outcome. We find our decision aid that concretizes the probability and payoff consequences of candidate offer amounts can increase consumer earnings by more than 25%, but only when the decision architecture nudges them toward the normative process. The sellers we simulate do not benefit from providing decision aids in terms of profit, suggesting real-world sellers would not adopt decision aids or normative-process nudges voluntarily.

Without the decision aid, the backward nudge does not help consumers earn more, and sharpening the incentive to follow the normative process surprisingly backfires, reducing earnings by 45% relative to the nudge. Thus, when consumers are thinking about how much to offer, they are better off knowing they will have a choice of whether to submit the offer than knowing the offer will be submitted no matter what—they end up offering less and ultimately earning more when they anticipate being able to decide eventual entry. Although using backward solution seems simple to us as researchers, the incentivized backward architecture is difficult for subjects to adopt. When faced with formulating a forced offer, subjects may resort to either simplifying heuristics (Tversky and Kahneman 1974) or adopting a more simplistic objective function that treats the submission cost as sunk.

The managerial and policy recommendations stemming from our results are straightforward. Sellers and regulators who wish to help consumers get more surplus in participative pricing markets with entry costs should both provide instant feedback about chances and potential savings relative to market or other reference prices, and also sequence their interface to separate the two sub-decisions, ideally letting consumers formulate offers first before deciding whether to submit them. Because sellers may not benefit from providing decision aids, regulators advocating for consumer welfare may want to mandate them. Once decision aids are in place, we find sellers are approximately indifferent between the different sequencings of the interface, so they may be more easily persuaded to use the backward decision architecture.

When the seller's profit includes the submission cost, as in the case of a two-part tariff involving a "submission fee" studied theoretically by Spann et al. (2010), we find that the seller does not benefit from offering a decision aid either (details in Web Appendix). A consultant to such a seller would make a grave error by following the risk-neutral model proposed in Spann et al. (2010): such a model would set the fee too high and lose over a third of potential expected revenue. Our results demonstrate that such a consultant should not rely on a risk-averse model either: such a model would dramatically under-predict potential revenue by under-predicting buyer entry. Please see the last figure in the Web Appendix for details of this additional analysis.

Our tightly controlled experiments are well suited to identify causal effects and to understand when and why anomalies in consumers' entry and bidding behavior occur. However, the analyses do not reveal much about the magnitude of the observed effects in the wild. Such quantification would require field experiments, which could also further investigate the design and implementation of decision aids and other elements of consumer decision architecture.

References

Ackerberg, Daniel, C. Lanier Benkard, Steven Berry, and Ariel Pakes (2007), "Econometric Tools for Analyzing Market Outcomes," in *Handbook of Econometrics*, J. J. Heckman and E. E. Leamer, eds. Vol. 6A: Elsevier.

Amaldoss, Wilfred and Sanjay Jain (2008), "Joint Bidding in the Name-Your-Own-Price Channel: A Strategic Analysis," *Management Science*, 54 (10), 1685-99.

Bajari, Patrick and Ali Hortaçsu (2005), "Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data," *Journal of Political Economy*, 113 (4), 703-41.

Battigalli, Pierpaolo and Marciano Siniscalchi (2003), "Rationalizable Bidding in First-Price Auctions," *Games and Economic Behavior*, 45, 38-72.

Bernhardt, Martin and Martin Spann (2010), "An Empirical Analysis of Bidding Fees in Name-Your-Own-Price Auctions," *Journal of Interactive Marketing*, 24 (4), 283-96.

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.

Crawford, Vince P. and Nagore Iriberri (2007), "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?", *Econometrica*, 75 (6), 1721-70.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde (2012), "The Intergenerational Transmission of Risk and Trust Attitudes," *Review of Economic Studies*, 79 (2), 645-77.

Ertaç, Seda, Ali Hortacsu, and James W. Roberts (2011), "Entry into Auctions: An Experimental Analysis," *International Journal of Industrial Organization*, 29 (2), 168-78.

Fay, Scott (2009), "Competitive Reasons for the Name-Your-Own-Price Channel," *Marketing Letters*, 20 (3), 277-93.

---- (2004), "Partial-Repeat-Bidding in the Name-Your-Own-Price Channel," *Marketing Science*, 23 (3), 407-18.

Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (2002), "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, 104 (1), 247-72.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong (2000), "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 68 (3), 525-74.

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), 1563-79.

Harrison, Glenn W. (1989), “Theory and Misbehavior of First-Price Auctions,” *American Economic Review*, 79 (4), 749-762.

Haruvy, Ernan and Peter T. L. Popkowski Leszczyc (2018), "A Study of Bidding Behavior in Voluntary-Pay Philanthropic Auctions," *Journal of Marketing*, 82 (3), 124-41.

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.

Holt, Charles A. and Susan K. Laury (2002), "Risk Aversion and Incentive Effects," *American Economic Review*, 92 (5), 1644-55.

Johnson, Eric J., Suzanne B. Shu, Benedict G. C. Dellaert, Craig Fox, Daniel G. Goldstein, Gerald Häubl, Richard P. Larrick, John W. Payne, Ellen Peters, David Schkade, Brian Wansink, and Elke U. Weber (2012), "Beyond Nudges: Tools of a Choice Architecture," *Marketing Letters*, 23 (2), 487-504.

Kahneman, Daniel and Amos Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47 (2), 263-92.

Krasnokutskaya, Elena and Katja Seim (2011), "Bid Preference Programs and Participation in Highway Procurement Auctions," *American Economic Review*, 101 (6), 2653-86.

Levin, Dan and James L. Smith (1994), "Equilibrium in Auctions with Entry," *American Economic Review*, 84 (3), 585-99.

Lichtenstein, Sarah and Paul Slovic (1973). "Reversals of Preference Between Bids and Choices in Gambling Decisions," *Journal of Experimental Psychology*, 89(1), 46-55.

McAfee, R. Preston and John McMillan (1987), "Auctions with Entry," *Economic Letters*, 23 (4), 343-47.

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

Moreno, Diego and John Wooders (2011), "Auctions with Heterogeneous Entry Costs," *RAND Journal of Economics*, 42 (2), 313-36.

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), 731-47.

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.

Smith, James L. and Dan Levin (1996), "Ranking Auctions with Risk Averse Bidders," *Journal of Economic Theory*, 68 (2), 549-61.

Smith, Vernon L. (1976), "Experimental Economics: Induced Value Theory," *American Economic Review*, 66 (2), 274-79.

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), 121-36.

Spann, Martin, Robert Zeithammer, and Gerald Häubl (2010), "Optimal Reverse-Pricing Mechanisms," *Marketing Science*, 29 (6), 1058-70.

Staw, Barry M. (1976). "Knee-Deep in the Big Muddy: A Study of Escalating Commitment to a Chosen Course of Action," *Organizational Behavior & Human Performance*. 16 (1), 27-44.

Tversky, Amos and Daniel Kahneman (1992), "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5 (4), 297-323.

---- (1974), "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185 (4157), 1124-31.

Tversky, Amos, Paul Slovic, and Daniel Kahneman (1990), "The Causes of Preference Reversal," *American Economic Review*, 80 (1), 204-17.

Zeithammer, Robert (2015), "Optimal Selling Strategies When Buyers Name Their Own Prices," *Quantitative Marketing and Economics*, 13 (2), 135-71.

Appendix: Tables and Additional Analyses

Experiment 1: Entry Probabilities and Offer Amounts, by Decision-Aid Condition

		Entry probability						Average offer amount given entry					
	v\c	0	1	4	8	16	32	0	1	4	8	16	32
No aid	10	91%	44%	9%	0%			5.9	6.0	4.6			
	25	93%	68%	31%	8%	1%		17.7	17.8	16.4	13.2	8.0	
	40	97%	84%	44%	20%	6%	2%	29.7	29.0	28.0	27.7	13.6	6.0
	55	98%	95%	64%	26%	8%	3%	42.3	42.4	40.9	37.4	31.7	13.3
	70	100%	97%	82%	48%	16%	3%	55.2	55.2	55.6	51.2	42.0	27.3
	85	99%	99%	85%	64%	27%	5%	67.0	66.8	66.3	63.2	55.1	37.0
	100	99%	98%	93%	84%	51%	5%	75.0	74.7	76.3	73.3	67.9	48.0
Aid	10	80%	15%	5%	0%			6	5	3			
	25	84%	46%	12%	6%	1%		17	18	15	9	8	
	40	86%	67%	22%	9%	4%	1%	28	28	27	24	16	2
	55	95%	79%	37%	19%	10%	2%	42	43	38	38	29	7
	70	100%	98%	73%	38%	9%	5%	54	55	53	46	35	28
	85	98%	98%	88%	66%	20%	4%	66	67	65	63	49	40
	100	98%	97%	95%	83%	54%	14%	72	74	72	71	66	44
Difference	10	**-11.5%**	**-29.4%**	-3.5%	0.0%			-0.4	-1.2	-2.0			
	25	**-9.2%**	**-21.5%**	**-18.8%**	-1.7%	0.1%		-0.6	0.4	-1.7	**-4.4**	0.0	
	40	**-11.4%**	**-16.8%**	**-21.6%**	**-10.7%**	-2.0%	-0.7%	-1.7	-1.2	-1.2	**-3.8**	1.9	-4.0
	55	-3.3%	**-15.9%**	**-27.1%**	-6.6%	2.3%	-0.6%	-0.8	0.3	**-3.3**	0.2	-2.5	**-6.3**
	70	0.0%	1.4%	-9.1%	-9.9%	-7.2%	1.6%	-0.8	0.0	-2.6	**-5.3**	**-6.7**	0.4
	85	-1.1%	-1.1%	2.5%	1.9%	-7.4%	-1.1%	-1.0	0.3	-1.0	-0.4	**-6.3**	3.0
	100	-1.1%	-1.3%	1.8%	-0.8%	2.7%	**8.9%**	-2.5	-0.7	-4.0	-2.3	-1.6	-3.6

Note: Color heatmap indicates high (green) and low (red) values within every block. **Bold** indicates a statistically significant difference at 5% level. The thick border shows the risk-neutral entry-indifference line.

Experiment 2 vs. Experiment 1: Effect of Nudging Toward the Backward Process

		Entry probability						Average offer amount given entry					
	v\c	0	1	4	8	16	32	0	1	4	8	16	32
No aid	10	-3%	5%	1%	**3%***			0.0	-0.2	-0.9			
	25	-1%	5%	3%	6%	**5%***		-0.7	-1.5	-1.1	-1.2	**-3.8***	
	40	-3%	0%	4%	-1%	3%	3%	-0.5	-0.1	-1.9	**-3.4***	2.4	-0.2
	55	-4%	-5%	-2%	8%	7%	1%	-1.5	-1.3	-1.5	-3.2	**-7.0***	-3.7
	70	-2%	-2%	-1%	4%	5%	6%	-0.9	-0.6	**-3.5***	-2.9	-0.4	-5.0
	85	-1%	-2%	2%	3%	2%	2%	-1.6	-2.8	-3.3	-2.7	-3.7	-4.9
	100	-1%	-3%	1%	-3%	5%	**8%***	-1.5	0.6	-3.0	**-4.4***	**-6.5***	3.4
Aid	10	0%	9%	-1%	0%			-0.1	0.5	1.9			
	25	3%	8%	-1%	-2%	1%		-1.3	**-2.9***	-1.3	0.7	-1.0	
	40	5%	4%	3%	-1%	-1%	0%	-1.3	-0.7	-3.0	-1.0	2.8	**6.0***
	55	-1%	6%	6%	-1%	-4%	0%	-1.6	-2.5	-2.2	-1.7	-6.4	0.5
	70	-3%	-3%	-1%	6%	2%	-2%	-1.1	-2.4	-1.0	-0.2	0.4	-4.3
	85	2%	-1%	0%	3%	4%	2%	-2.8	-1.5	**-3.7***	-1.9	1.3	-4.8
	100	1%	2%	1%	3%	3%	-5%	1.0	-3.0	-1.9	-1.3	-2.1	3.8

Note: Color heatmap indicates positive (blue) and negative (red) differences between suggested backward and baseline, within every dependent variable (separate heatmaps for entry probabilities and offer amounts) and across both decision-aid conditions. **Bold** followed by **asterisk (*)** indicates statistically significant difference at 5% level. The thick border shows the risk-neutral entry-indifference line.

Experiment 3 vs. Experiment 2: Effect of Forcing the Backward Process

		Entry probability						Average offer amount given entry					
	v\c	0	1	4	8	16	32	0	1	4	8	16	32
No aid	10	3%	6%					**0.7***	**0.6***				
	25	-1%	-1%	9%				1.0	**2.4***	**1.9***			
	40	-1%	**-10%***	4%				1.0	**2.8***	**3.4***			
	55		-4%	2%	4%				**3.4***	**3.2***	**3.6***		
	70	-4%	-1%		-2%	9%		3.0	4.5		**4.2***	**3.8***	
	85		-1%	-4%		1%			7.0	6.0		**4.4***	
	100	-2%				-8%	**12%***	**3.3**				**9.9***	3.0
Aid	10	4%	8%					0.5	**0.7***				
	25	-1%	-5%	8%				0.9	**1.7***	**2.8***			
	40	-2%	-2%	0%				**2.6***	1.5	**4.2***			
	55		-3%	1%	-1%				2.3	2.8	1.3		
	70	2%	0%		1%	4%		-0.2	1.9		**4.1***	**9.9***	
	85		1%	-1%		9%			0.4	3.3		**4.8***	
	100	-2%				-4%	5%	-1.2				-1.2	**10.7***

Note: Color heatmap indicates positive (blue) and negative (red) differences between forced and suggested backward treatments, within every dependent variable (separate heatmaps for entry probabilities and offer amounts) and across both decision-aid conditions. **Bold** followed by **asterisk (*)** indicates statistically significant difference at 5% level. The thick border shows the risk-neutral entry-indifference line.

Multiplicative Joy of Winning Does Not Explain High Offers

Does a multiplicative joy of wining explain the high offer magnitudes we observe in Experiment 1? Ertaç et al. (2011) show that a joy of winning that makes buyer *i* behave as if all her valuations were inflated by the factor $J_i > 1$ would manifest as $offer_i(v, c = 0) = \frac{vJ_i}{2-R_i}$ in our paradigm. The following regression uses our "want to win" statement ("My offer being successful was more important to me than the potential payoff.") to show the multiplicative joy of winning does not explain the cross-sectional variation in free offers. To maximize power, the regression combines all three decision-aid conditions in Experiment 1.

Regression of Offers on "Want to Win"

	Estimate	**t-stat**	**t-stat (clustered SE)**
Constant	−1.84	−1.55	−2.31 *
Probability-only aid	−2.05	−3.37 *	−1.79
Decision aid	−1.09	−1.71	−0.96
Valuation	0.79	43.72 *	39.48 *
Want To Win scale	0.39	1.37	2.51 *
Valuation X Want2Win	−0.01	−1.43	−1.14
Number of observations	2205		
Number of subjects	336		
R^2	0.79		

Note: * indicates significance at the 5% level.

Web Appendix

<u>Table of Content:</u>

Evidence for Constant Relative Risk Aversion Preferences Based on Offer Data

Heterogeneity in risk aversion implies heterogeneity in the overall level of offers: more risk-averse people offer more, ceteris paribus. To control for this heterogeneity when analyzing offers, we construct a “relative offer”—the offer divided by the average free offer of the same subject. Table WA1 reports the results of a linear regression of the relative offer on valuation, cost, and decision-aid dummies. By definition of the relative-offer construct, all effects are in percentages of the average free offer the same consumer makes when offers are free to submit.

Table WA1: Regressions of Relative Offers

	Simultaneous		**Backward**	
	Estimate	t-stat	Estimate	t-stat
constant	14.6% *	11.7	15.7% *	12.5
v_25	28.1% *	42.9	26.0% *	48.6
v_40	55.0% *	84.3	54.2% *	57.2
v_55	85.4% *	89.3	82.7% *	70.6
v_70	115.8% *	108.6	115.5% *	85.0
v_85	142.6% *	100.7	140.1% *	82.5
v_100	161.4% *	108.5	160.7% *	81.6
prob_aid	−2.9%	−1.3	−4.0%	−1.8
full_aid	−1.9%	−0.6	−2.5%	−1.1
cost_1	2.2% *	2.3	1.4%	1.5
cost_4	2.5%	1.6	−1.2%	−1.0
cost_8	−2.3%	−1.1	−3.8% *	−2.5
cost_16	−13.8% *	-4.2	−17.8% *	−7.4
cost_32	−55.2% *	−10.8	−51.6% *	−11.5
N	6,578		6,839	
R^2	0.69		0.74	

Note: * indicates significance at 5% level.

Note the data used in the regression involves selection based on individual risk aversion, with fewer (and presumably less risk averse) subjects submitting offers when costs are higher. If offers were declining monotonically with costs, such a selection might explain the relationship between submitted offers and costs. However, the relationship we find is non-monotonic, and hence harder to explain by selection alone.

What does the non-monotonic relationship between offers and costs teach us about the utility function? Consider the following notation for the gain-side and loss-side utilities u_G and u_L:

$$b(v,c) = \arg\max_{b\geq 0} \left(\frac{b}{p}\right) \underbrace{u(v-c-b)}_{u_G>0} + \left(1-\frac{b}{p}\right) \underbrace{u(-c)}_{u_L<0} . \tag{A1}$$

Then, the first-order condition is $bu_G' = u_G - u_L$, equating the marginal utility loss from lower surplus upon winning on the LHS to the net marginal gain from winning more often on the RHS.

Differentiating both sides of the FOC in terms of *c* yields:

$$\frac{\partial b}{\partial c}u_G' + bu_G''\left(-1-\frac{\partial b}{\partial c}\right) = u_G'\left(-1-\frac{\partial b}{\partial c}\right) + u_L' \Leftrightarrow \frac{\partial b}{\partial c} = \frac{u_L' - \left(u_G' - bu_G''\right)}{2u_G' - bu_G''} \tag{A2}$$

where the $u_G' - bu_G''$ term in the parentheses is obviously positive as long as u_G is concave.

It is clear from equation (A2) that the slope of offer in cost is ambiguous. To demonstrate that risk-aversion alone is not sufficient for the pattern observed in the data, consider two examples while focusing on the sign of $\left.\frac{\partial b}{\partial c}\right|_{c=0}$. Assume that $u_L'(0) = 1$, i.e., there is no kink at zero (no loss or gain aversion). Then:

<u>*Example* 1: CARA in gains</u>

FOC is: $bu_G' = u_G - u_L : Rb = \exp\left(R\left(v-b\right)\right) - 1$, and the numerator of the crucial derivative is:

$u_L' - \left(u_G' - bu_G''\right) = -R\exp\left(-R\left(W+v-c-b\right)\right)\left[\exp\left(Rb\right) - 2\exp\left(Rv\right) + \exp\left(R\left(2v+b\right)\right)\right] < 0$.

Therefore, CARA utility has $\frac{\partial b}{\partial c} < 0$ everywhere.

<u>*Example* 2: CRRA in gains:</u>

The numerator of the crucial derivative is $u_L' - \left(u_G' - bu_G''\right) = 1 + \frac{cR}{v-c-b} - \frac{1}{\left(1-R\right)\left(v-c-b\right)^R}$.

Since $b\left(v,0\right) = \frac{v}{2-R}$ is in closed form, we can simply plug it in, and obtain:

$u_L' - \left(u_G' - bu_G''\right) = 1 - \frac{1}{\left(1-R\right)\left(v\frac{1-R}{2-R}\right)^R} > 0$ when *v* is large enough.

Therefore, CRRA utility in gains implies $\left.\frac{\partial b}{\partial c}\right|_{c=0} > 0$ for large-enough valuation.

To summarize, the slope of offers in cost at zero is informative about the shape of the gain-side utility function. A positive slope is consistent with CRRA but not CARA.

Comparing Fit of Standard Entry Models in Experiment 1

Using a parametric random-utility entry model can sharpen the assessment of fit between the risk-neutral benchmark and the entry data. We ran a logistic regression of the entry decisions on the theoretical expected surplus of a risk-neutral agent, rescaled so that the standard deviation across the 39 (v,c) design cells is 1 (but not de-meaned, so the sign of the rescaled surplus is the same as the sign of expected surplus). The re-scaling measures the expected surplus in standard deviations across the (v,c) cells, facilitating a comparison with analogously scaled expected utility—another objective function the buyers may follow. We found an intercept of −0.21 and slope of 2.58, both highly significant. The negative intercept sign implies that our subjects enter a bit less often than 50-50 when their expected consumer surplus (the objective of a risk-neutral buyer) would be zero (the model predicts they enter with probability 0.45), but the large positive coefficient on the Scaled Expected Surplus indicates its strong predictive ability. Comparing the two coefficients to each other shows that our average subjects' indifference curve follows the expected surplus of about 0.1 standard deviations—less than one cent in monetary terms. To illustrate this indifference curve at our valuation points, note that the risk-neutral expected surplus of a free offer in our setting is $v^2/400$, and so the average subject would face an expected surplus of 1 cent at the following valuation-cost pairs: (25,0.5), (40,3), (55,6.5), (70,11), (85,17), and (100,24).

An analogous logistic regression on equally-scaled expected utility from the EU CRRA model with wealth of 35 and a CRRA coefficient of 4 (a calibration derived from matching the model to free offers) finds an intercept of 0.39 and slope of 1.50. The positive intercept indicates over-entry relative to this alternative model, and the smaller coefficient on the rescaled objective function indicates that the model is actually less predictive of entry than the risk-neutral model. We conclude that the risk-neutral benchmark captures the entry behavior remarkably well.

Table WA2 shows the logistic regression of entry on the scaled objective as described above, for the two standard entry models. The behavior of subjects with decision aids is less consistent with the two standard models than behavior without aids. The fit with the risk-neutral model worsens via a reduced intercept—a sign of reduced entry along the diagonal of the (*v*,*c*) space clearly visible in Figure 4. The fit with the risk-averse EU model also worsens in terms of the slope, significantly in the probability-aid condition.

Table WA2: Comparing the Fit Of Two Standard Entry Models: Risk-Neutral and Expected Utility With Constant Relative Risk-Aversion Calibrated on Free Offers

	Model: Risk-Neutral		**Model: Risk-Averse EU**	
	Estimate	t-stat	Estimate	t-stat
intercept	−0.21 *	−4.97	0.39 *	12.05
prob_aid dummy	−0.11	−1.85	−0.10 *	−2.26
full_aid dummy	−0.45 *	−7.16	−0.28 *	−5.93
scaled objective	2.58 *	32.34	1.50 *	14.29
scaled objectives X prob_aid	−0.05	−0.47	−0.47 *	−3.82
scaled objective X full_aid	0.06	0.49	−0.24	−1.68

Note: N = 13104, * indicates significance at 5% level.

Estimation of Expected-Utility and Prospect-Theoretic Models

How well do the canonical expected utility theory and the more flexible prospect theory fit the entry and offer data we observe? Can a prospect-theoretic model resolve the preference inconsistency we document? We answer these questions by estimating a parametric version of the model in equation 1 using both the entry and the offer-amount data, thus examining both key components of consumer behavior jointly. We consider two versions of the model: CRRA expected utility theory with $u_i(x) = x^{R_i}$ and prospect theory with $u_i(x) = \begin{cases} x \geq 0 : \gamma_i x^{r_i} \\ x < 0 : -(-x)^{s_i} \end{cases}$.

While the risk-neutral model discussed extensively in the paper makes the same prediction for all subjects, the risk-averse model has a free parameter—the constant risk-aversion *R*. A prospect-theoretic model we specify above has three free parameters: "gain liking" γ measuring the kink at zero (opposite of loss aversion) and two curvatures in the loss and gain domain (*s* and *r*, respectively). Regardless of the preference model, these free parameters need to be estimated at the individual level to allow for the preference heterogeneity suggested by the *i* subscripts in the equations above.

We model the behavior of consumer *i* with valuation *v* facing an offer-submission cost of *c* as following equation 1 with the appropriate parametric utility. Our goal is to estimate the structural parameters θ, i.e., either *R* in the CRRA expected utility model or $\{r, s, \gamma\}$ in the prospect-theoretic model, for each subject. We model the observed offer as $offer = b^*(v,c\,|\,\theta) + \varepsilon$, where $b^*(v,c\,|\,\theta)$ is the argmax of the inner (inside square brackets) maximization in equation 1, and ε is distributed normally with a mean of 0 and variance σ^2, truncated such that $0 < offer < v - c$. The truncation restricts the observed offer to the basic rationality constraints enforced in our experiment design. Given the normality assumption, it is straightforward to define the likelihood of the structural parameters θ, as well as the variance term σ^2, as follows:

$$\log \mathrm{L}_b\left(\theta,\sigma \middle| v,c,offer\right) = -\log(\sigma) - \frac{\left(offer\text{-}b^*(v,c\,|\,\theta)\right)^2}{2\sigma^2} - \log\left[\Phi\left(\frac{v-c-b^*(v,c\,|\,\theta)}{\sigma}\right) - \Phi\left(\frac{-b^*(v,c\,|\,\theta)}{\sigma}\right)\right]$$

where the last term accounts for the truncation.

To model the entry decision, let $V_{enter}\left(v,c \middle| \theta\right)$ be the choice-relevant value of entering, namely

$$V_{enter}\left(v,c \middle| \theta\right) = \max_{b\geq 0} Pr\left(accept\,|\,b\right)u_i\left(v-c-b,\theta\right) + \left(1 - Pr\left(accept\,|\,b\right)\right)u_i\left(-c,\theta\right).$$

To allow an econometric error, we model each entry choice as a logistic transformation of this choice-relevant value:

$$\mathrm{Pr}_{enter}\left(v,c \middle| \theta,\tau\right) = \frac{\exp\left(\tau V_{enter}\left(v,c \middle| \theta\right)\right)}{1+\exp\left(\tau V_{enter}\left(v,c \middle| \theta\right)\right)}.$$

Overall, the likelihood of the data combines the entry information with the observed offer information across observations indexed by *n* as follows:

$$\log \mathrm{L}\left(\theta,\sigma,\tau\right) = \sum_{n \text{ where } entry} \log\left[\mathrm{Pr}_{enter}\left(v_n,c_n \middle| \theta,\tau\right)\right] + \log \mathrm{L}_b\left(\theta,\sigma \middle| v_n,c_n,offer_n\right) + \sum_{n \text{ where } no\ entry} \log\left[1-\mathrm{Pr}_{enter}\left(v_n,c_n \middle| \theta,\tau\right)\right].$$

This likelihood function combines probabilities and densities in the same way censored models, such as the Tobit, do.

Combining the likelihood from the entry model with the likelihood of the offer model discussed above into a single overall likelihood function, it is straightforward to obtain individual-level estimates of $\{\theta,\sigma,\tau\}$ using maximum likelihood**.** The likelihood maximization is straightforward except for the need to solve the bidding problem in equation 1 at every step, which we accomplish by searching for the optimal bid on a fine grid. To keep the estimates in a realistic range, we conduct a constrained search for the maximum-likelihood estimates, allowing both risk-averse and risk-seeking preferences in gains and losses. Specifically, we constrain $r, s \in [0.05, 4]$ $\sigma, \tau \in [0.1, 20]$, and $\gamma \in [0.1, 10]$. To avoid local maxima, we build up from constrained specifications (i.e., starting with offers only with $\gamma = 1$), and start each optimization with an informed guess of the parameter value using a combination of constrained results and grid search of the newly freed parameter.

We begin the discussion of our estimates with an estimation based on observed offer data alone. This mimics most auction econometrics that focuses on observed bids. Table WA3 shows the population moments of the parameter estimates for both preference models. The estimated σ clearly shows that the more flexible prospect-theoretic model fits better than the standard CRRA model. The globally concave CRRA model struggles to fit costly offers, while the more flexible prospect-theoretic model suggests that about half of the subjects have an S-shaped value function (s<1). The γ parameter further suggests that half the subjects care more about a small gain than about a small loss. Our evidence for loss aversion is thus mixed in that half the subjects seem to have preferences that are the opposite of loss aversion. Between the γ exceeding unity and the S-shaped utility function, many of our subjects thus seem to neglect the downside of making a costly offer in formulating their offer amounts. This downside neglect is one way to explain why offers do not initially rise as much as a function of cost as the simpler CRRA model predicts (see Web Appendix titled "Evidence for Constant Relative Risk Aversion Preferences Based on Offer Data" for details of this observation).

Table WA3: Models Estimated on Offer Data Only, Parameter Estimates

	Expected Utility Model with CRRA			**Prospect-Theoretic Model**			
Parameter	Mean	Median	SD	Mean	Median	SD	> 1
R	9.01	6.19	7.42				
r				0.39	0.30	0.45	6%
s				1.52	0.84	1.56	42%
σ	7.01	5.86	4.47	6.04	4.62	4.50	100%
γ				2.77	1.00	2.78	50%

We now add the entry behavior to the data used in estimation. Table WA4 shows the population moments of the parameter estimates for the prospect-theoretic model while demonstrating how the estimated value function varies by the maximum submission cost actually paid by a buyer during the experiment.

Table WA4: Prospect-Theoretic Model Estimated on Both Offer and Entry Data, Parameter Estimates

					Median by Max Cost Paid			
Parameter	Mean	Median	SD	> 1	4	8	16	32
r	0.36	0.28	0.33	4%	0.11	0.22	0.38	0.26
s	1.15	1.07	0.75	54%	1.70	1.41	0.88	0.53
σ	6.23	5.24	4.52	100%	2.83	4.84	5.99	8.63
τ	5.37	1.42	7.28	56%	1.03	1.22	2.26	0.63
γ	4.96	3.97	3.81	84%	3.07	6.03	3.42	3.57
Subjects	117				10	38	52	13

Table WA4 and Figure WA1 show how the heterogeneous prospect-theoretic model addresses our main modeling challenge through flexibility on the loss side: the model connects the two seemingly mutually incompatible decision components through increasing convexity in the loss domain among buyers who pay higher costs. We clearly find a lot of preference heterogeneity: the population average of approximate risk neutrality in the losses shown in WA4 is actually an aggregate of rather concave (and hence compatible with risk-averse expected utility theory) preferences among the buyers who never paid more than 4 cents, and rather convex preferences among the buyers who paid 16 or 32 cents at least once.

Figure WA1: Estimated Value Function, Median by Maximum Submission-Cost Paid

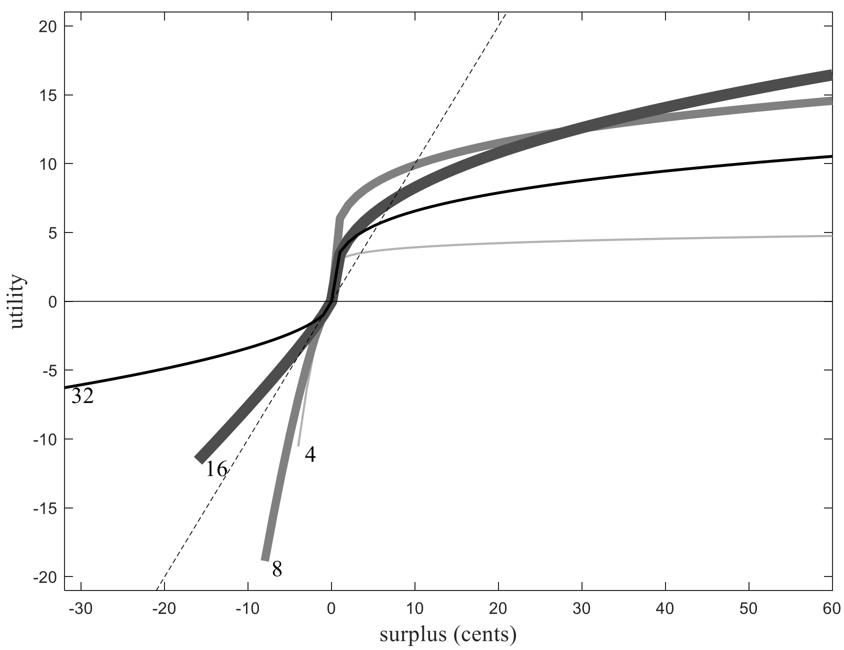

Note: Each utility function starts at the level of surplus equal to –c. Thickness corresponds to the number of subjects for whom that was the highest cost paid throughout the experiment.

Another notable feature of the parameter estimates shown in Table WA4 is the γ that far exceeds unity—the opposite of loss aversion traditionally defined as a concave kink at zero. The γ higher than 1 indicates our subjects are gain-loving "in the small," meaning they like a small gain much more than they dislike the monetarily equivalent small loss. Subjects who have paid 16 or 32 cents at least once also continue to care more about gains than losses "in the large"; that is, they exhibit a consistent downside neglect at all cost levels. By contrast, more conservative buyers who never paid more than 8 exhibit loss aversion "in the large" in the sense that they dislike a large loss more than they like an equivalent monetary gain. Our model thus captures a range of possible behaviors from staying out of high-cost cost rounds and offering a lot upon entry to entering even high-cost rounds and giving low-ball offers upon entry.

The model predicts entry well: the average predicted probability when a subject did not enter is 0.074, whereas the average predicted probability when entry occurred is 0.860. Compared with the same model estimated on offer data only (Table WA3), the model estimated on all the data involves only a marginally worse fit to the observed offer magnitude. The main difference in terms of model parameters is the larger γ in Table WA4 than in Table WA3 to explain the large amount of entry.

Decision Aid and Learning

The four retest rounds allow us to assess learning from experience during the experiment. We find that learning generally involves a reduction in entry, whereby subjects enter somewhat less in the retest rounds than in the identical test (measurement) rounds. Table WA5 shows the average entry probability across the four retest rounds compared with the average entry probability in the same design cells during measurement rounds. that the table shows the magnitude of the entry reduction declines with the decision aid, as well as with the backward direction of the sequencing. Decision aids and backward sequencing are thus substitutes for learning from experience regarding entry. However, the lower entry is not necessarily beneficial in terms of earnings: subjects generally earn slightly less in the retest tasks than in the matched test tasks (the differences are not significant; results not reported in detail).

Table WA5: Learning to Enter Less from Experience

	No Aid			Decision Aid		
	test	retest	diff	test	retest	diff
Baseline	56%	45%	−11% *	41%	34%	−6%
Backward	57%	47%	−10% *	45%	41%	−4%

Note: * indicates significance at the 5% level, with the four observations per subject in both the test and retest interpreted as independent observations.

Additional Figures

Effect of Decision Aids on Subjective Probability Beliefs, Suggested Backward

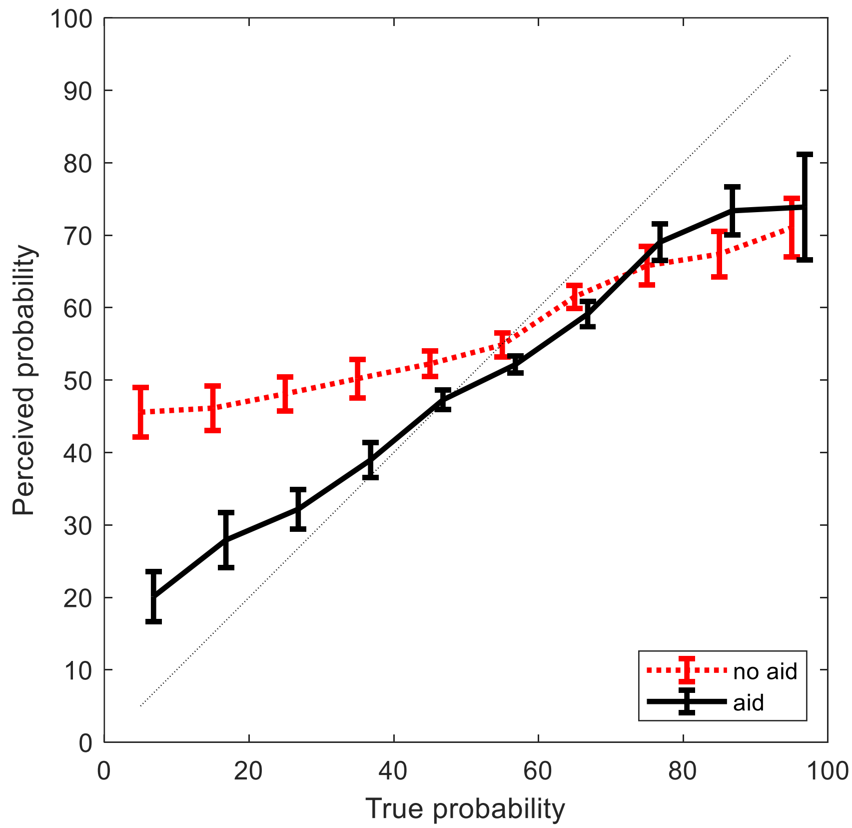

Effect of Forcing Backward Process on Offers, Condition with Decision Aid

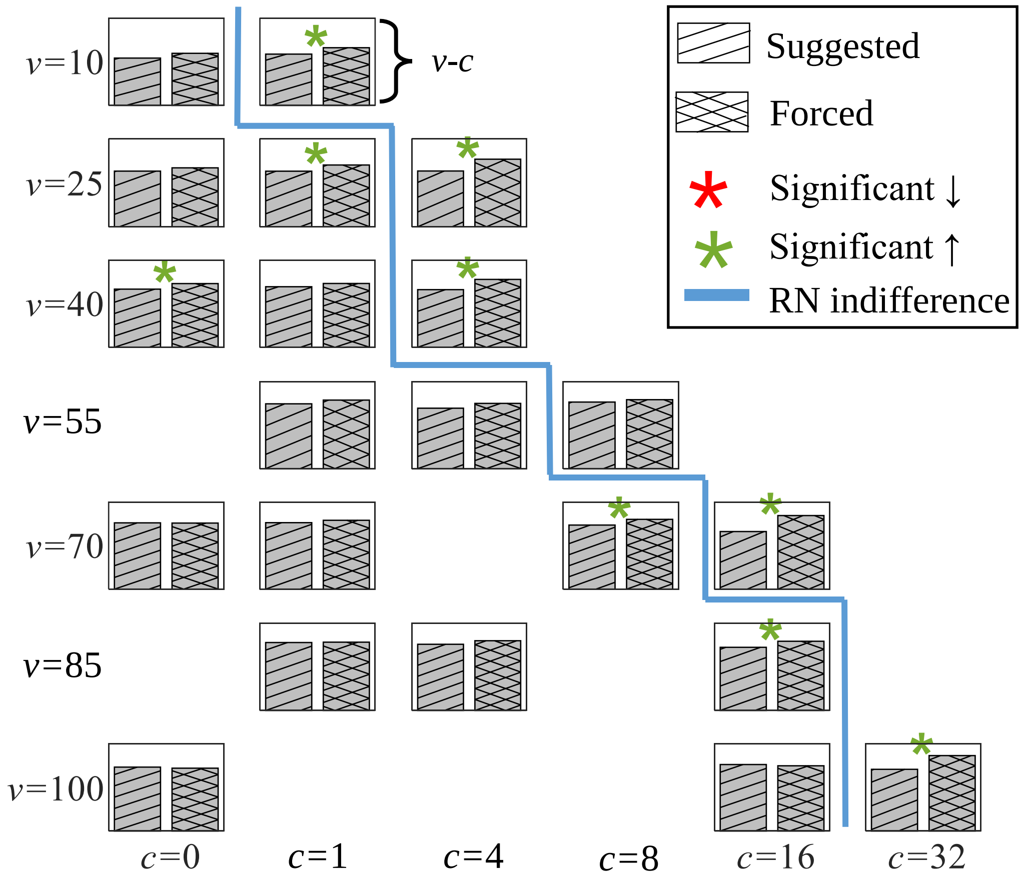

Effect of Decision Aid on Seller Capturing the Submission Cost as a "Fee"

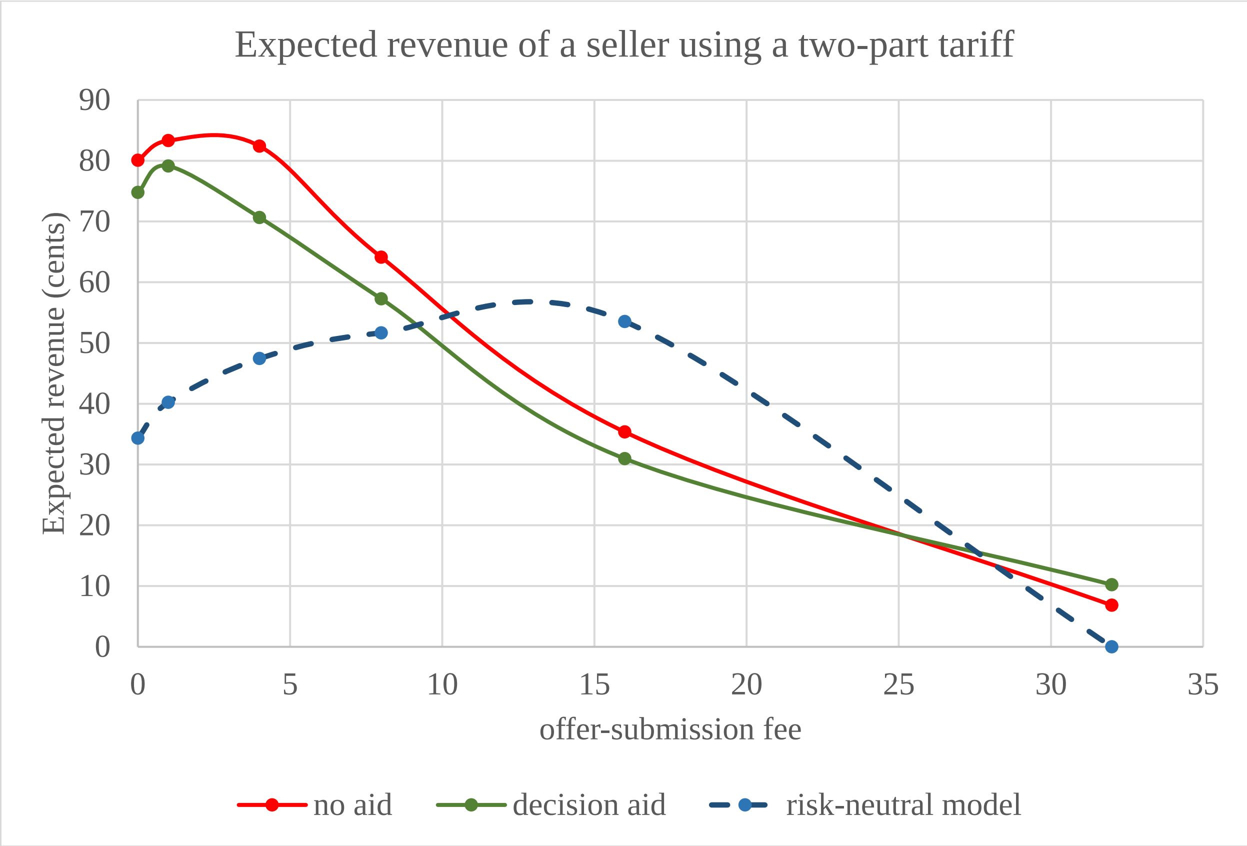

Note: Empirical revenues calculated in the simultaneous decision architecture.