

UCLA

UCLA Previously Published Works

Title

Strange Case of Dr. Bidder and Mr. Entrant: Consumer Preference Inconsistencies in Costly Price Offers

Permalink

<https://escholarship.org/uc/item/47q1f0g5>

Authors

Zeithammer, Robert

Stich, Lucas

Spann, Martin

et al.

Publication Date

2024-03-21

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Peer reviewed



Contents lists available at ScienceDirect

International Journal of Research in Marketing

journal homepage: www.elsevier.com/locate/ijresmar

Strange Case of Dr. Bidder and Mr. Entrant: Consumer Preference Inconsistencies in Costly Price Offers

Robert Zeithammer^{a,*}, Lucas Stich^b, Martin Spann^c, Gerald Häubl^d

^a Anderson School of Management, UCLA, 110 Westwood Plaza, Los Angeles, CA 90095, USA

^b Faculty of Business Management and Economics, Julius-Maximilians-Universität Würzburg, Sanderring 2, 97070 Würzburg, Germany

^c LMU Munich School of Management, Ludwig-Maximilians-Universität München, Geschwister-Scholl-Platz 1, 80539 Munich, Germany

^d Alberta School of Business, University of Alberta, 3-23 Business Building, Edmonton, Alberta T6G 2R6, Canada

ARTICLE INFO

Article history:

Received 5 June 2023

Available online xxxx

JEL Codes:

D12

D44

D81

C91

Keywords:

Pricing

Auctions

Entry costs

Behavioral economics

Experiments

ABSTRACT

Consumers make price offers to sellers in a variety of domains, such as when buying cars or houses or when bidding in auctions for airline upgrades, art, or collectibles. Submitting an offer typically entails 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 a series of incentive-compatible experiments. Our findings reveal a preference inconsistency, whereby the preferences implied by one of the decisions do not align with the preferences implied by the other. In particular, potential buyers enter more often than their offer amounts would predict based on standard economic models. This preference inconsistency is robust to two interventions designed to help consumers make offer-amount and entry decisions—(1) the provision of interactive-feedback decision aids and (2) the sequencing of the two sub-decisions in the normative order. Neither of these interventions resolves the inconsistency. Instead, the patterns of results suggest that consumers approach the offer-amount and entry decisions as if they were *unrelated*. We discuss the implications of our findings for the design of offer-submission interfaces, as well as for econometric attempts to infer consumer preferences from offer and bidding data.

© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

“I saw that, of the two natures that contended in the field of my consciousness, even if I could rightly be said to be either, it was only because I was radically both”

Robert Louis Stevenson, *Strange Case of Dr. Jekyll and Mr. Hyde*

Consumers make price offers to sellers in a variety of domains, including buying appliances, cars and houses, acquiring used household goods in garage sales and their online parallels (such as Craigslist), making offers to buy sneakers on StockX.com, and bidding in auctions for art, airline upgrades, or collectibles ([Bajari and Hortacsu 2005](#); [Yao and Mela 2008](#); [Haruvy and](#)

* Corresponding author.

E-mail addresses: rzeitam@ucla.edu (R. Zeithammer), lucas.stich@uni-wuerzburg.de (L. Stich), spann@imu.de (M. Spann), gerald.haeubl@ualberta.ca (G. Häubl).

<https://doi.org/10.1016/j.ijresmar.2024.08.006>

0167-8116/© 2024 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Popkowski Leszczyc 2009; Jindal and Newberry 2018). An important source of friction inherent in such participative pricing markets (Spann et al. 2018) is the *participation cost*, that is, the administrative cost of submitting a price offer,¹ as well as the time and opportunity costs associated with waiting for the outcome.² The auction literature often refers to these costs as “entry costs,” following Samuelson (1985).³ In all of the above examples, the cost of making an offer is substantial, so it is important for both researchers and firms to understand how consumers factor such costs into their decision-making.

We examine how participation costs influence consumer behavior through a series of incentive-compatible experiments that manipulate key design elements of the bidding interface. Compared to the related experimental auction literature (e.g., Palfrey and Pevnitskaya 2008; Ertaç et al. 2011), our study focuses on a simpler setting inspired by Priceline’s name-your-own-price mechanism (e.g., Amaldoss and Jain 2008; Fay 2004; Shapiro 2011). In our experiments, participants make binding price offers to a computer-simulated seller, whose random hidden reserve price determines offer acceptance. Our single-agent task is an ideally tractable setting for gaining insight into consumer offer-making with participation costs, without the strategic interactions present in multi-bidder auctions (e.g., Menezes and Monteiro 2000; Samuelson 1985). By focusing on single-agent decision making, our research thus complements the related experimental literature on excess entry into competitive situations (e.g., Laferrrière, Staubli, and Thöni 2023; Palfrey and Pevnitskaya 2008; Ertaç et al. 2011).

Our primary audience is researchers and analysts who estimate consumer preferences from bidding and offer data. Crucially, for estimating consumer preferences, the costly price-offer task involves two interrelated decisions: in addition to determining how much to offer, consumers facing participation costs must also decide whether to make an offer at all, i.e., whether to “enter” the uncertain situation at all. Both the offer-amount and the entry decisions can be interpreted separately as arising from underlying preferences, and we document a robust internal preference inconsistency, whereby the preferences implied by one of the decisions are inconsistent with the preferences implied by the other. Specifically, we find that potential buyers enter more often than their offer amounts would predict based on standard economic models.

Behavioral science has made significant progress in 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). In contrast, we document a preference inconsistency internal to a *single* task, where 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.

We report the results of four experiments. To test the decision difficulty hypothesis, we provide some participants with an interactive decision aid (e.g., Häubl and Trifts 2000) that provides real-time assistance in calculating the chances and payoffs given a potential offer amount (Experiment 1). Normative theories of behavior in our setting suggest a clear backward-solution order—a rational decision-maker should parse the problem into the offer-amount and the entry decision, and then solve the task “backward” starting with the offer amount. To help our participants follow the normative process, we both nudge them (Experiment 2) and force them (Experiment 4) to follow it. To rule out potential alternative explanations based on narrow bracketing (Read, Loewenstein, and Rabin 1999, Rabin and Weizsäcker 2009) and shifting reference points (Kahneman and Tversky 1979), we also reverse the normative order and solicit the two parts of the task starting with entry (Experiment 3).

Surprisingly, our manipulations—both the decision aid and various sequenced architectures—fail to resolve or reduce the preference inconsistency between the two decisions. While the manipulations do not bring the revealed preferences from entry closer to the revealed preferences from offers amounts, they do have two large effects on behavior: First, the decision aid improves consumers’ understanding of the probability of offer acceptance, influencing their entry behavior in line with their improved probability perception, but without a corresponding effect on offer amounts. Second, incentivizing participants to follow the normative process instead of merely nudging them towards it affects their offer amounts but does not affect their entry behavior. Specifically, participants offer more when their offers are submitted automatically compared to when they are only nudged to come up with offers first. But their behavior in the subsequent entry task remains the same, seemingly unrelated to the offer amount.

Taken together, our results suggest that the inconsistency is a robust feature of decision-making in costly participative pricing, not a bug that somehow needs to be fixed. One way to describe the situation we document is suggested by the title of this paper: within each participant, there resides a very risk-averse “Dr. Bidder” who sets the offer amount, and a nearly risk-neutral “Mr. Entrant” who decides whether to enter the game. Decision aids affect only Mr. Entrant, and the forcing of the normative architecture affects only Dr. Bidder. In other words, a typical participant in our experiments is not Dr. Bidder or Mr. Entrant, but “radically both”. We document an interesting inconsistency, propose several 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 consumer decision-making with entry costs—the normative assumption about one person with fixed preferences solving the two parts of the task using backward induction.

¹ Such costs include completing official paperwork for a house offer, traveling to a dealership to make an offer on a new car, joining and learning an online platform like [Greentoe.com](https://www.greentoe.com) to make offers on consumer electronics, or the nonpecuniary cost of initiating a negotiation as in Jindal and Newberry (2018).

² 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.

³ 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.

Our findings also hold immediate implications for decision architects (Johnson et al. 2012). Since our experiments involve induced valuations, we can directly measure the monetary surplus consumers earn under different architectures. Should interfaces guide consumers through the decision-making process, and to what extent should they enforce the normative process? We find that while gentle guidance can improve consumer earnings, forcing the backward-solution process can have adverse effects. Concerning decision aids, we observe that they can enhance consumer earnings, but only when combined with a sequential decision architecture. This suggests that consumers struggle to parse the task into the two decisions, and merely helping them “do the math” does not help them earn more. Only when combined with the correct parsing does the “what-if” calculator help them increase their payoffs.

2. Normative models and predictions for entry and bidding behavior

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”, following the auction literature (e.g., Samuelson 1985; Ertac 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:

$$\max_{\text{enter, not}} \left[\underbrace{\max_{b \geq 0} \Pr(\text{accept}|b)u_i(v - c - b) + (1 - \Pr(\text{accept}|b))u_i(-c), u_i(0)}_{V_i(v, c)} \right] \quad (1)$$

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

In words, Eq. (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 payoff (decreasing in the offer amount). Although submitting free ($c = 0$) offers is 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.

Note that Eq. (1) abstracts away from pre-decision costs such as cost of thinking analyzed in Hauser and Wernerfelt (1990). Given the simplicity of our decision interface, we do not expect such costs to be large in our experiments. To further minimize pre-decision costs, we provide training rounds that familiarize the participants with the task before the key experimental rounds begin. Our induced-value paradigm (Smith 1976) rules out the pre-decision cost of learning the private values analyzed in Levin and Smith (1994).

2.1. Specification-agnostic predictions: Backward-solution process

Equation (1) 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 first because it does not depend on any additional assumptions about the shape of the utility function.

Regardless of the shape of the utility function u_i , Eq. (1) constrains the rational decision-maker to first solve the inner maximization problem that finds the best offer given entry, and then compare the choice-specific value $V_i(v, c)$ with the utility of staying out of the market $u_i(0)$. Fig. 1 illustrates this normative process with preferences and probability beliefs as inputs. The nested structure of the two decisions of the problem shown in Eq. (1) implies a rational decision-maker needs to solve the two decisions in the order shown in Fig. 1. Because only one person with one utility function is making both decisions, they have to move together in a way consistent with the person’s preferences and beliefs. From this observation, we can make the following predictions:

- (a) A manipulation that systematically affects probability beliefs should have an effect on both offer amounts and entry decisions, and
- (b) A manipulation that systematically affects preferences should have an effect on both offer amounts and entry decisions.

As an example of (a), consider a manipulation that makes probability beliefs more optimistic. Such a manipulation should increase entry and reduce offer amounts. As an example of (b), consider a manipulation that makes preferences more risk-averse. Such a manipulation should increase offer amounts and reduce entry. By manipulating the bidding architecture of our interface, we can test the above specification-agnostic predictions by examining the predicted co-movement of the two parts of the decision.

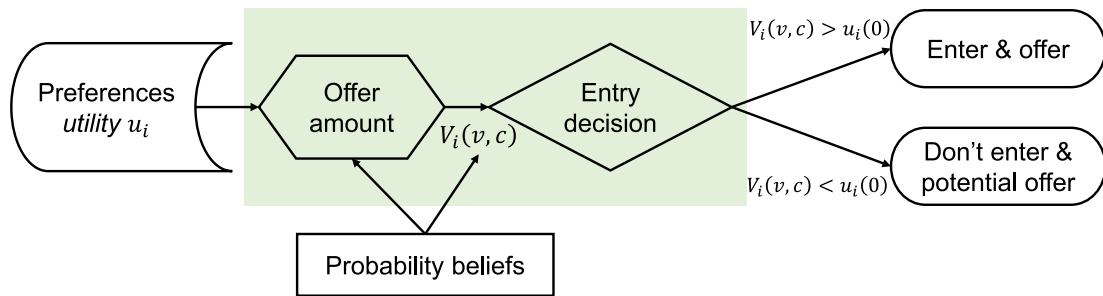


Fig. 1. Normative Backward-Solution Decision Process in Costly Price Offers.

2.2. Predictions from existing utility models and their behavioral enrichments

The general model in Eq. (1) includes several notable special cases that entail different assumptions about consumers' risk preferences. First, when $u_i(x) = x$, we obtain the risk-neutral model, which implies buyers should offer half their valuations whenever they enter (see [Web Appendix H](#) for a detailed derivation), and also implies that 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](#)).

Second, 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 $R \geq 0$ and 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 Hortacısu 2005](#); [Cox et al. 1988](#)), i.e. above $v/2$ here. It is also the winner of a modeling horse-race performed by [Bajari and Hortacısu \(2005\)](#), and hence a dominant candidate for capturing behavior in our experiments.

Third, 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](#)).⁴ For example, letting

$$u_i(x) = \begin{cases} x \geq 0 : \gamma x^r \\ x < 0 : -(-x)^s \end{cases} \quad (2)$$

for some $0 < r < 1$, $0 < s < 1$, and $0 < \gamma < 1$, captures the classic S-shaped value function with different curvatures in the gain and loss domains, as well as a concave kink at zero consistent with loss aversion.⁵ Freeing the r , s , and γ parameters to possibly exceed unity results in a flexible model that can accommodate a range of shapes, including standard globally concave risk-averse utility from expected utility theory when $r < 1 < s$.

In addition to specifying the curvature of the utility function, prior literature (e.g., [Ertaç et al. 2011](#)) has proposed enrichments of the standard model in Eq. (1), such as the idea of the *joy of winning*, where people derive additional utility from having their offer accepted. Such additional value, denoted by w , could take two forms: *additive*, where consumers with an actual valuation of v behave as if their valuation were $v + w$ for some $w > 0$, and *multiplicative*, where they behave as if their valuation were vw for some $w > 1$. Joy of winning differs from *joy of playing* proposed by [Palfrey and Pevnitskaya \(2008\)](#), which suggests that experimental subjects may simply receive an additional utility boost from entering regardless of the game outcome.

Given a utility function, the model in Eq. (1) can be used to predict how variations in valuation and submission costs affect buyers' entry and offer behavior. We designed the within-subject variation of our experiments to see which of the different utility models discussed above best fits the observed behavior (at the individual level) over a wide range of valuations and costs. We now introduce the common aspects of the paradigm used in all our experiments.

3. Experimental paradigm: Within-subject manipulation of valuation and offer submission cost

In this section, we discuss the key aspects of our experimental paradigm.⁶ We designed our experimental task to measure both decision elements of Eq. (1): offer amount and entry decision. To abstract from strategic considerations in auctions and the

⁴ Note that in our setting, there is no difference between the original Prospect Theory ([Kahneman and Tversky 1979](#)) and its cumulative version ([Tversky and Kahneman 1992](#)) because the task we study is a mixed gamble (some chance of loss, and the remaining chance of gain).

⁵ Note that Prospect Theory also includes probability weighting ([Gonzalez and Wu, 1999](#)); we abstract from it in this paper because probability weighting is not identified independently of r, s, γ in our data.

⁶ The complete instructions, stimuli, data, and code used in all our experiments are available in a project directory on the Open Science Framework: https://osf.io/q3ysf/?view_only=47df578c500b4ffc91aa06868e0e1e76.

associated explanations of the inconsistency, we used the single-agent setting motivated by name-your-own-price selling. Participants 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 participants) 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 randomly in each round, and with each value between 0 and 100 being equally likely. We explained to participants 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 they offered.⁷

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, participants 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) to provide incentives consistent with Eq. (1), 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 participants at the time of entry. After learning the current (v, c) pair, participants 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.

In each round, participants made money as follows. If sellers accepted their offer, they purchased the widget at 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 cost of submission c . Participants' 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. Based on prima facie irrational behavior by some participants in a pilot study (described in Web Appendix A), we only allowed offers below $v - c$, and therefore did not allow entry when $v < c$. In other words, we limited the bid amount to the $[0, v - c]$ interval.

Whenever a participant decided to submit an offer, they were asked on an immediately following 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 intermediate markers.

For all our experiments, we used CloudResearch's MTurk toolkit to recruit participants.^{8,9} To make sure each participant understood the game and payoff rules, we used five screening questions designed to test understanding of the rules. Only participants 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, participants completed nine training rounds (to familiarize participants with the task, but not used for analysis) that were indistinguishable to them 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, participants answered a brief demographic and risk-preference questionnaire, the details of which differed somewhat between experiments. In the concluding questionnaire, all experiments included a survey of attitudes to measure specific explanations of their behavior.¹⁰

4. Econometric model for estimating individual-level preferences in the proposed experimental paradigm

Our experimental paradigm varies the valuation v and cost c within subjects, and produces several dozen¹¹ observations of each participant's entry and offer decisions as a function of (v, c). These data allow us to estimate each participant's preferences using standard econometric techniques. The modeling section introduced two utility specifications prominent in the literature: expected utility theory with a constant relative risk-aversion (EU), and prospect-theoretic preferences with an S-shaped reference-dependent utility function and a kink at the origin (PT). In this section, we explain our estimation approach.

The EU model $u_i(x) = \frac{(W+x)^{1-R}}{(1-R)}$ has one free parameter, the constant risk-aversion R , and one parameter that is not identified in our data, the wealth parameter W . We fix $W=35$ throughout to ensure parameter comparability across experiments and conditions. The PT model $u_i(x) = \begin{cases} x \geq 0 : \gamma x^r \\ x < 0 : -(-x)^s \end{cases}$ has three free parameters: “gain liking” γ measuring the kink at zero (to capture loss aversion) and two curvatures in the loss and gain domain (s and r , respectively).

⁷ Given the inclusion of zero among the threshold prices and a tie-breaking rule in favor of acceptance, the true implemented probability was actually very slightly higher (in all experiments and treatments, so this cannot explain any effects we find).

⁸ Online panels are not generally representative of the U.S. population, so our findings do not represent the U.S. population either. However, a lack of representativeness cannot explain our results.

⁹ Our first experiment conceptually replicates the findings of our earlier study titled “Paying for a Chance to Save Money” using an experimental economics laboratory (available at https://osf.io/q3ysf/?view_only=47df578c500b4ffc91aa06868e0e1e76), indicating our documented empirical regularities extend beyond online participants..

¹⁰ On a 7-point scale (from 1 = not at all to 7 = very much), participants 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”).

¹¹ The exact number varies by experiment and condition, as explained below.

Let θ denote the vector of utility parameters we want to estimate, i.e., either $\theta = R$ in the EU model or $\theta = \{r, s, \gamma\}$ in the PT model. Our econometric exercise aims to compare the estimate $\hat{\theta}_{offers}$ based only on the observed offer data to $\hat{\theta}_{entry}$, based only on the observed entry behavior. Given the wide variation in individual risk preferences (e.g. documented by Holt and Laury 2002; Dohmen et al., 2011), we estimate both $\hat{\theta}_{data}$ separately for each participant, effectively allowing for arbitrary heterogeneity of preferences.

Before estimation, we normalize each utility function without loss of generality (WLOG) such that $u_i(1) - u_i(0) = 1$ and $u_i(0) = 0$, which in turn makes our estimation procedures more stable and the magnitude of the error terms more interpretable. We then adopt standard assumptions about the error terms, treating them as random differences between theoretical predictions and empirical observations: To model the entry decision, we use a logistic regression without an intercept of the choice-specific value function $V_i(v, c|\theta)$ from Equation 1:

$$Pr_{enter}(v, c|\theta, \tau) = \frac{\exp(\tau V_i(v, c|\theta))}{1 + \exp(\tau V_i(v, c|\theta))} \quad (3)$$

where the parameter $\tau > 0$ scales the choice-relevant value, and a larger τ correspond to a better fit of the model. The logistic model is a standard model of binary dependent variables, and it can be readily estimated by maximum likelihood.

To model the offer magnitude, we model the observed offer as $b^*(v, c|\theta) \exp(\varepsilon)$, where $b^*(v, c|\theta)$ is the optimal bid given θ , and ε is distributed normally with a mean of 0 and variance σ^2 , truncated such that $b^*(v, c|\theta) \exp(\varepsilon) < v - c$. In other words, we model the errors as lognormal, truncated to obey the basic rationality constraints enforced in our experimental design. The lognormal assumption also automatically ensures that the observed bids are positive. The lognormal error model is standard in the econometrics of bidding data and can be readily estimated by maximum likelihood. Please see Web Appendix B for the likelihood formula that accounts for the necessary truncation, as well as for additional technical details involved in our maximum likelihood estimation. We now turn to our first experiment.

5. Experiment 1: Effect of decision aids on offers, entry, and preference inconsistency

We designed our first experiment to investigate the preference inconsistency we found in our pilot study and to test whether decision difficulty could explain it. To manipulate decision difficulty, participants 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. Participants 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 Fig. 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 participants 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. To simplify the exposition, we do not report it in detail.

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: 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. A competing posted-price offer for the same good may in turn serve as the effective valuation bidders face. Finally, the auctioneer may also know the submission cost when it is primarily monetary, such as a fee the auctioneer is charging. The decision aid we implement can thus be a realistic option for real-world sellers, and so analyzing its effect is interesting in its own right.

5.1. Experimental design

All conditions used a “simultaneous” decision architecture that solicits offer and entry decisions on a single screen (see left side of Fig. 2). 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 participants 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 participants): 9 training rounds, 42 measurement rounds, and 4 retest rounds. The 9 training rounds exposed participants 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 participant to a $7(v) \times 6(c)$ full-factorial design involving all possible combinations of $v \in \{10, 25, 40, 55, 70, 85, 100\}$ and $c \in \{0, 1, 4, 8, 16, 32\}$.¹²

¹² The (v, c) levels were designed as follows: Maximum valuation was set to 100 cents to enable participants to easily understand acceptance probabilities. The additional valuation levels were set by subtracting a constant amount (15 to have some valuations that do not end in zero). Given these valuations, a cost of 32 is clearly too high for a risk-neutral participant (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. Offer cost of 1 is the smallest possible positive cost to test for knee-jerk aversion to paying for offer submission. 4 and 8 were included as non-negligible intermediate values.

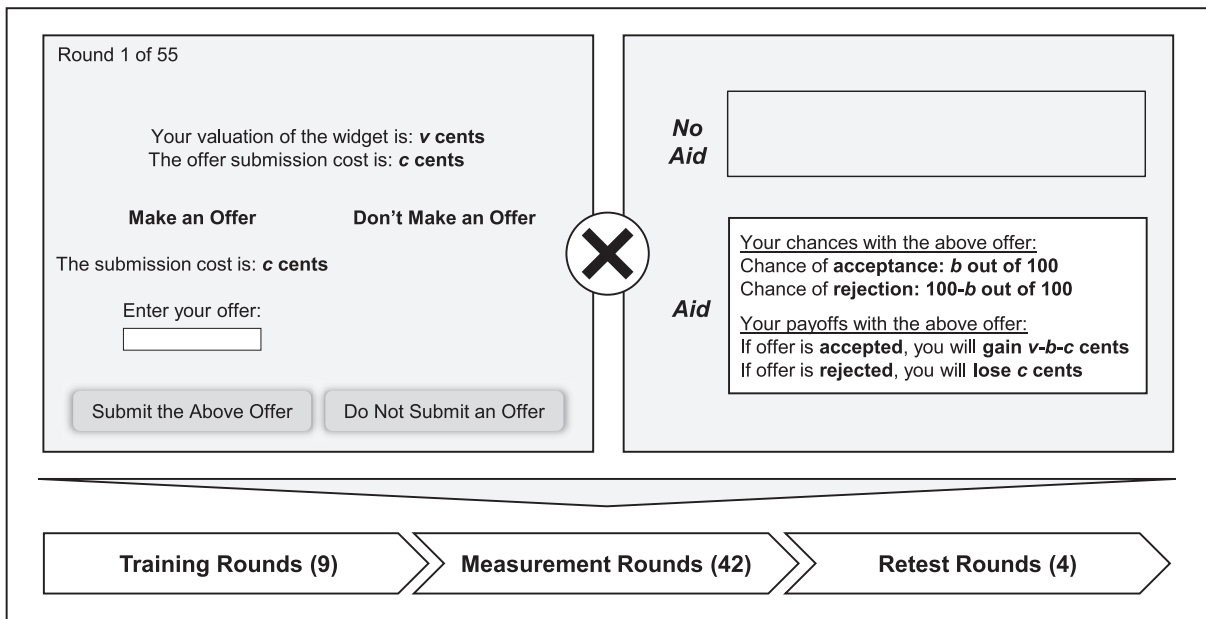


Fig. 2. Experiment 1: Overview of Experimental Design.

In the 3 (v, c) cells that involve $v < c$, participants were not allowed to make an offer because doing so would guarantee a loss. 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 participants with 4 (v, c) conditions¹³ 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 participants. After the 55 rounds, participants 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, participants received a base payment of \$2.50. In addition, they received their final account balance (with accounts initialized with \$1 to allow for small losses) 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 ($M_{\text{age}} = 38.59$, $SD_{\text{age}} = 11.84$; 38.71 % female) passed the screening questions and participated in the experiment, 100 of whom had access to the decision aid.¹⁴ The average participant earned about \$5.04 and took about 26 min (median 23 min) to complete the entire experiment. Only one participant ended up losing both the \$1 and their payoff from the Holt and Laury task; they received the base payment. All remaining participants received their total payoffs.

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 participants who entered in each (v, c) cell of the design and the average offers these entrants made. Our findings—summarized in Table 1—replicate the internal inconsistency found in the pilot study (described in Web Appendix A).

Consider the risk-neutral model $u_i(x) = x$ in Eq. (1) as the starting point of the analysis. The risk-neutral model implies buyers should offer half of their valuations regardless of the submission-cost amount whenever they enter (see section 2). 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. The risk-neutral model thus does not fit the offer data well,¹⁵ and risk-averse models (that generally predict less bid-shading) seem to fit better. By contrast, the risk-neutral model is good at predicting the average entry behavior if we allow for some noise: The entry predicted by the risk-neutral model is shown by the shading in Table 1: the model predicts entry in the unshaded cells (Web Appendix H shows that buyers with $v > 20\sqrt{c}$ enter). Note all but two (v, c) unshaded cells have actual entry probabilities above 50 %, and all shaded cells have actual entry probabilities below 50 %. In other words, Table 1 suggests that, on average, participants enter as if they were nearly risk-neutral, but bid as if they were risk-averse. This lack of overall fit of the risk-neutral model is quite consequential in that participants realize only 43 % of the payoff that risk-neutral agents are expected to earn.

¹³ The retest cells were (85,16), (55,4), (70,8), (40,1), with expected risk-neutral earnings between 2.1 and 4.2 cents.

¹⁴ The screening questions disqualified 45% of invited participants who consented to participate in the study.

¹⁵ Neither does any other model in which entry costs are effectively sunk, e.g., Smith and Levin (1996).

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

		Entry Probability	Submission Cost				
			0	1	4	8	16
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)					
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.

We now turn to our econometric method to document the internal inconsistency more precisely. Fig. 3 shows the distribution of “free” offers, i.e., offers where $c = 0$ and most participants enter, by decision aid condition. The left panel of Fig. 3 thus analyzes a subset of the offer data summarized in Table 1, and it is both clear that the risk-neutral model does not fit the subset, while the estimated homogeneous CRRA model¹⁶ (with $R=9.3$) fits the data well.

While clearly illustrating the difference in fit between the two models, the EU model used in Fig. 3 is homogeneous (same CRRA for all participants) and considers only a subset of the data (when $c = 0$). Table 2 presents results estimated on all data using a heterogeneous estimation approach that allows each participant their own CRRA coefficient.¹⁷ Justifying our individual-level approach, we find a high degree of population heterogeneity in that the individual-level CRRA coefficients vary widely across participants. Moreover, the population distribution is highly skewed (see Fig. 4), so the median is more relevant than the mean for capturing the central tendency of the data.

Looking at the medians in Table 2, we can sharpen our preliminary conclusion from Table 1 regarding the preference inconsistency central to this paper: participants set offer amounts as if they were extremely risk-averse with a median CRRA of about 9, but their entry behavior exhibits only very weak risk-aversion of about 1. The difference in the fit of these two estimates is profound: an analyst predicting offer amounts based on entry behavior would nearly double their error in terms of RMSE (Root Mean Squared Error between predictions and the actual offers) compared to the analyst analyzing observed offer amounts. Conversely, an analyst predicting entry behavior based on offer amounts would underestimate entry in the sense that their predicted entry probabilities among entrants would be only 60%—much less than the 82% of an analyst using entry data and the same underlying model of preferences.

One way to see why the CRRA expected-utility model calibrated on offer amounts underpredicts entry is to note the healthy entry in (v,c) cells of Table 1 just below the shaded region: a CRRA model with $R=9$ predicts no entry in cells with $c \geq 8$, as well as no entry in $(25,1)$, $(70,4)$, and $(55,4)$, but some of these cells have upwards of 80% of participants entering with the average entry probability across all eight cells of 61%.

Could the prospect-theoretic model resolve the preference inconsistency by allowing for more flexible and behaviorally realistic preferences? We cannot perform the same analysis as in Table 2 for the prospect-theoretic model because its parameters are not empirically identified from entry data alone (for example, both a low γ and a high s are consistent with non-entry). However, the model is identified based on offer data alone, so we can study the impact of adding in the entry data. Table A5 in Web Appendix C summarizes the estimation results. It is immediate from Table A5 that compared to using offer data alone, using both entry and offer data to estimate $\{r, s, \gamma\}$ results in a doubling of γ (to a population average of about 4) and halving of s (to a population average of about 1). The preference inconsistency thus persists even under the more flexible prospect-theoretic model.¹⁸

¹⁶ We pooled the data across participants, and excluded the few outliers marked by “+” in Fig. 3 from the maximum likelihood estimation for the purposes of Fig. 3.

¹⁷ To estimate the model on offer data, we need at least five offers submitted by each participant analyzed, and we exclude the few participants in each condition who submitted less often (note that this exclusion leaves out people who don’t enter at least twice even when entry is free). We also exclude these participants from the analysis of entry data to ensure that offer and entry decisions of exactly the same people are analyzed in Table 2.

¹⁸ We note that the $\{r, s, \gamma\}$ parameter values we estimate do not resonate with the classic prospect-theoretic S-shaped value function. Instead, the entry and offer data analyzed together implies the opposite of loss-aversion, and very little loss-side curvature. The average values also mask a lot of heterogeneity that is tangential to the main argument of this paper.

5.3. Results: Evidence against joy of winning, joy of playing, or participant fatigue explaining the inconsistency

The above apparent internal inconsistency of our participants 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) would suggest. But our data is not consistent with the explanations proposed in this literature. First, the distribution of offers is not consistent with participants deriving additional additive utility from winning (Ertaç et al. 2011): Given our rationality restriction that bids do not exceed $v-c$, additive utility from winning would manifest as a bunching of low-valuation free offers at their maximum allowed value of v . That is, the boxplot corresponding to $v = 10$ in Fig. 3 should be bunched up against its upper bound. Instead, the modal free offer at $v = 10$ is 5, and only one of the 107 participants in the condition without aid offered exactly 10. Our self-reported process data is not consistent with either (multiplicative or additive) form of additional utility from winning either: we regressed the free offers in Fig. 3 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 Web Appendix D). 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.

Joy of playing, namely the notion that participation in the market has entertainment value as suggested by Palfrey and Pevnitskaya (2008), is nearly impossible to rule out based on the offer and entry data alone because it simply implies that one should add a free intercept to $V_i(v, c|\theta)$ in Equation 1. Such a free intercept would tautologically “explain” excess entry, but one could not be sure about the underlying mechanism based on offer and entry data alone. One feature of Table 1 that is inconsistent with joy of playing is the fact that 9% of participants stay out in the (10,0) cell, which precludes the possibility of a loss and involves a positive expected gain, so everyone should enter in it—especially if participants experience even a small joy of playing. Similarly, almost a third of the participants stay out in the (25,1) cell, which also involves a positive expected gain only amplified by a joy of playing. If anything, the enrichment more consistent with the observed entry pattern is a small hassle cost proposed by Fay (2009), and later estimated to be quite large by Hann and Terwiesch (2003) and Jindal and Newberry (2018)—the opposite of the joy of playing. To address the possibility of joy of playing directly, we again turn to self-reports. We find that 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 ($p=.22$).

Finally, the inconsistency also cannot be fully explained by participants wearing out over the course of many rounds and gradually resorting to some sort of heuristic or random behavior that happens to be internally inconsistent. When we split the rounds in half and compare the “early” first 21 to the “late” second 21 rounds, we find that entry declines somewhat over time, but not enough to make late entry consistent with the risk-aversion implied by offer amounts. Specifically, Web Appendix G shows considerable entry in cells where a risk-averse person would not enter, both in the early and late rounds. Thus, the inconsistency we find is not somehow a mechanical product of a simple monotonic trend in our multiple-round design. This, of course, leaves the possibility of non-monotonic trends such as found in Li et al. (2022). But such analysis is beyond the scope of this paper.

With the existing 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.

5.4. Results: Decision aid impacts entry but leaves offer amounts unchanged

To what extent are the baseline findings outlined above driven by the computational difficulty of the task? Specifically, does the apparent preference inconsistency diminish when participants do not have to calculate probabilities and payoffs in their heads? We designed the interactive decision aid to answer these questions. Surprisingly, the aid does not resolve the inconsistency puzzle even directionally. However, the aid not reducing the inconsistency does not mean the aid has no effect—we find it impacts entry without an internally consistent impact on offer magnitudes. We explain these two findings in turn.

Fig. 3 is deliberately constructed to show the two conditions side by side on the same scale, and it shows that free offers are not affected by the aid at all. Table A1 in Web Appendix A demonstrates that this conclusion holds for all offers when we compare the average offer amounts in the same (v,c) cell between conditions. Table 2 then shows that everything we said above about the baseline condition also holds true for the decision-aid condition. Thus, computational difficulty is not the cause of the inconsistency we find.

While the offer amounts are unaffected by the decision aid treatment, the entry behavior exhibits a large and non-monotonic effect: it decreases entry when both valuations and costs are small (including zero cost) while increasing entry when valuations and costs are high. Fig. 5 illustrates the effect at the (v,c) cell level, and Table A1 in Web Appendix A shows the underlying averages. The effect is highly statistically significant: since we are performing 39 tests, we use a Bonferroni correction to set a familywise threshold for rejecting the null hypothesis, and even this conservative test rejects the null in six (v,c) cells.

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 where most participants enter from cells where most participants stay out. One way to visualize the effect of the decision aid on entry is to have the empirical entry-indifference




Fig. 3. Free Offers, Risk-Neutral and Risk-Averse Models (Experiment 1). *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 expected value of the estimated CRRA model (the estimated CRRA coefficients are 9.3 without decision aid and 9.2 with decision aid). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

curve rotate around its midpoint, falling from the upper-left corner while rising from the lower-right corner. Rotating the indifference curve in this way is inconsistent with entry becoming more risk averse as a result of the decision aid: if the manipulation made participants more risk averse, their indifference curve would pivot downward at positive cost levels but exhibit no change at “free entry” cells with $c = 0$ (where all rational agents should always enter).

What is the underlying mechanism of the surprising entry effect we observe? Part of the effect can be explained by the aid correcting a systematic misperception of probabilities. Fig. 6 takes the perceived probability of success of each submitted offer and plots it as a function of the true probability of success of the offer by decile bin. It shows that without the aid, the participants 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 closer to the truth in this way, the aid should increase entry when valuations are large and decrease entry when they are small—precisely the non-monotonic effect we find. While the belief correction shown in Fig. 6 explains *some* of the patterns in Fig. 5, 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 participants 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 to participants without the aid, significantly more participants with the aid report wanting to win more (No aid: $M=3.21$, $SD=1.72$; aid: $M=3.82$, $SD=1.93$; $p = 0.02$).

More importantly for our goal of explaining the apparent inconsistency between offer and entry behavior, offer magnitudes do not reflect the improvement in reported probability beliefs.¹⁹ Focusing on cells where a majority of participants enter

¹⁹ Note that there may be selection effects in the right panel of Fig. 5 because different proportions of participants enter with and without aid. Such a selection might be an alternative explanation for any differences in average offers submitted. The fact that we find no differences in any cell, including in high- v free-entry cells where almost all participants enter, suggests that such selection effects are not operating here.

Table 2
Individual-Level Estimates of the CRRA Model of Experiment 1, by Data Used.

Data used in estimation	Offer data		Entry data	
	No aid	Aid	No aid	Aid
Mean CRRA (R)	7.06	7.30	1.99	2.31
Median CRRA (R)	9.13	9.23	0.97	1.54
Population stdev. CRRA (R)	3.41	3.44	2.56	2.44
RMSE bid fit	9.4	11.4	17.1	15.8
Estimated Pr(entry no entry)	15%	18%	10%	13%
Estimated Pr(entry entry)	60%	60%	82%	83%

Note: Only participants with at least 5 submitted offers are analyzed. (112 in “No aid”, 92 in “Aid”).

(to reduce possible selection effects), we find no significant differences in the magnitude of the submitted offers (not even at the cell-wise 0.05 threshold; see Fig. 5 for details). If participants were following the normative process outlined in Fig. 1, the aid’s alignment of perceived probabilities with true probabilities should increase offers (relative to the baseline condition), regardless of the particular utility specification when the baseline beliefs are too optimistic and reduce offers when the baseline beliefs are too pessimistic. However, we find neither effect, suggesting that our participants behaved as if they were formulating offer amounts without considering acceptance probabilities.

In summary, the strong effect of our decision aid on entry via subjective probabilities without a corresponding internally consistent effect on offer amounts suggests that participants do not follow the normative process outlined in Fig. 1. The insensitivity of offer amounts to probability beliefs implies that the non-normative sequential order of the tasks—in which consumers first figure out whether to enter and then come up with the offer amount—does not fit the data either. We now turn to experiments that directly manipulate the decision order.

6. Experiments 2 and 3: Sequencing the decisions has almost no effect

6.1. Experiment 2 (nudge towards normative process): Design and results

The interface used in Experiment 1 may invite various suboptimal heuristics, so people may benefit from a nudge toward the normative process. If such a nudge made them parse the problem correctly, they should also appear less internally inconsistent, and the standard model would better explain their behavior. To test this candidate explanation, we modified the decision architecture as shown in Fig. 7 to nudge consumers toward the normative process.

Except for the different decision architecture, Experiment 2 had identical procedures and used the same measures as the “Baseline” (simultaneous) experiment described above.²⁰ The training rounds also followed the interface of Experiment 1, so Experiment 2 differed from Experiment 1 only after the training rounds. A total of 241 US residents ($M_{\text{age}} = 39.63$, $SD_{\text{age}} = 12.37$; 42.32 % female) passed the screening questions²¹ and participated in the experiment, taking 31 min (median 28 min) and earning on average about \$5.12. 108 of them had access to the decision aid. Only two participants ended up losing both the \$1 and the payoff from the Holt and Laury task; they received the base payment. All remaining participants received their total payoffs. Regardless of the decision aid, we found very few (v,c) cells with significant effects of the “backward” architecture on entry and offer magnitudes relative to the “simultaneous” architecture in Experiment 1 (details in Web Appendix A).

Table 3 is the analog of Table 2 applied to the data from Experiment 2, and it shows that the preference inconsistency “survives” the nudge. We conclude that whatever heuristics people used in the baseline condition do not appear to have been simple errors that can be corrected with a nudge toward the normative process.

6.2. Experiment 3 (forward decision architecture): Design and results

In Experiment 3, we maintained a sequential design but reversed the order of the two parts of the decision. Fig. 8 outlines the resulting “forward” decision architecture. A total of 243 US residents ($M_{\text{age}} = 39.89$, $SD_{\text{age}} = 12.27$; 44.44 % female) passed the screening questions²² and participated in the experiment, taking 25 min (median 24 min) and earning on average about \$5.00. 119 of them had access to the decision aid. Only one participant ended up losing both the \$1 and the payoff from the Holt and Laury task; they received the base payment. All remaining participants received their total payoffs.

²⁰ Analogously to Experiment 1, we also included the intermediate aid condition, not reported in detail.

²¹ The screening questions disqualified 53% of invited participants who consented to participate in the study.

²² The screening questions that check understanding disqualified 51% of invited participants who consented to participate in the study.




Fig. 4. Population Distribution of the CRRA Coefficient, by Data and Aid Condition (Experiment 1).

We consider the forward architecture to reinforce the natural process that someone not trained in economic reasoning might adopt—even in the simultaneous architecture of Experiment 1. Therefore, we expected the forward architecture to exacerbate the preference inconsistency. But much like the “backward nudge” sequencing of Experiment 2, the forward sequencing did not have a significant effect on behavior, so we do not report it in detail. Table 4 is the analog of Table 2 applied to the data from Experiment 3, and it shows that the preference inconsistency “survives” the forward sequencing.

6.3. The implications of experiments 2 and 3 for potential explanations of the mechanism behind the inconsistency puzzle

The fact that sequencing the two parts of the decision does not affect the preference inconsistency (Tables 2, 3, and 4 are nearly identical) allows us to rule out an entire class of possible explanations related to bracketing and shifting mindsets. One such class is the sunk cost fallacy, defined by Arkes and Blumer (1985) as “a greater tendency to continue an endeavor once an investment in money, effort, or time has been made.” (p. 124). Our design excludes one prominent type of sunk-cost-fallacy, namely the escalation of commitment to a failing course of action (e.g., Staw 1976, Gunia et al. 2009, Martens and Orzen 2021) because our participants do not receive any interim informative signals about their chances. However, it is also possible that the fallacy arises from a spontaneous shift in consumers’ minds between the time they make one of the component decisions and the time they consider the second one. For example, participants processing the two parts in the “forward” sequence may first decide whether or not to enter without processing the problem completely, and only then feel the need to justify their costly decision with a large bid. Or perhaps these participants are loss-averse, and first decide whether or not to enter without thinking of the entry cost as a loss or gain (perhaps as merely a neutral “price”), and then suddenly begin to worry about losing the submission cost when they consider the bid amount. Alternatively, participants processing the two parts in the “backward” sequence may feel invested in the offer amount they have just decided on (perhaps because they spent some effort on it) and enter the game too often to justify the investment to themselves.

More generally, the explanation may center on a theory that consumers do not bracket the two parts of the task together and process them sequentially or “narrowly” in the language of Rabin and Weizsäcker (2009). Note that in Experiment 1, we tried to minimize the potential for these explanations by asking for both the offer amount and the entry decision in a single task (see Fig. 2). However, it is still possible that our participants use a sequential process and that something shifts between the two parts of the task.

Instead of testing for specific bracketing or sunk-cost-fallacy mechanisms, we address the entire class by noting that if any such “Mr. Entrant first entering somewhat recklessly/heuristically, and Dr. Bidder later setting cautious/high offer amounts” (or the two alter egos operating in reverse order) explained the behavior in Experiment 1, then either Experiment 2 or Experiment 3 would have disrupted the process. For example, a “mere price becomes a potential loss after entry” mechanism cannot explain the inconsistency in the backward architecture, which solicits the offer amount first. On the flipside, a “feeling invested in one’s own offer amount” mechanism cannot explain the inconsistency in the forward architecture which solicits the entry decision first. In sum, the fact that the inconsistency is robust to sequencing means that the entire class of bracketing theories is not sufficient to account for it.

Unlike the forward architecture—which was incentive-compatible in the sense that decisions in the first (entry) stage were consequential, the backward nudge architecture allowed participants to escape to the outside option in the second stage. As a result, decisions in this first stage (offer amounts) were not necessarily consequential, and some participants made apparently non-serious offers: in about a third of all person-rounds, the offer field was left blank, and in another quarter of person-rounds, the offer entered was zero. 67 % of participants left the field blank or entered zero in at least one round.

The non-serious offers made by many participants in the offer stage of Experiment 2 suggest that the nudge may not have been strong enough to encourage solving the problem backwards. Would a stronger manipulation toward a backward-

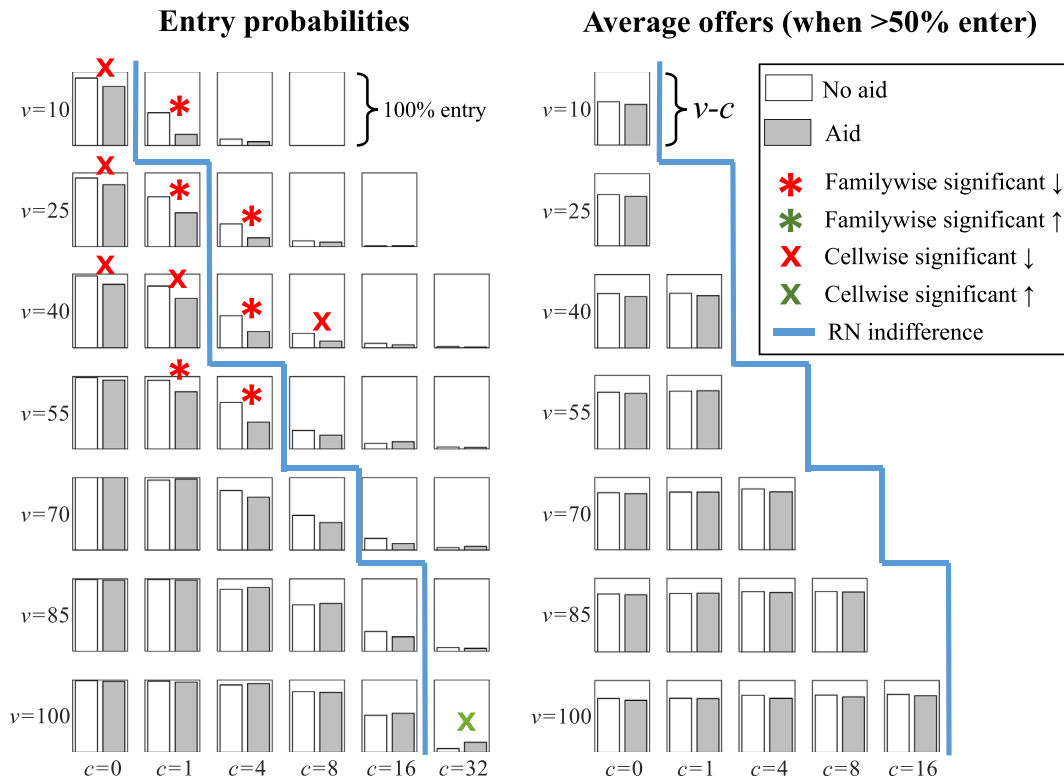


Fig. 5. Effect of Decision Aid on Entry and Average Offers Submitted (Experiment 1). *Note:* In the entry plot, the height of each box corresponds to 100 % of participants entering in that (v,c) cell. In the average offer plot, the height of each box is $v - c$, the upper bound on admissible offers. An "x" indicates a significant effect of aid at the 5 % level according to the z-test performed within the (v,c) cell only. A "*" indicates a significant effect of aid at the familywise 5 % level, adjusted with the Bonferroni correction for the 39 tests in each plot ($p < 0.0013 = 0.05/39$). The color denotes the direction of the change (red = decrease, green = increase). The thick (blue) line represents the indifference curve of the risk-neutral model. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 6. Effect of Decision Aid on Subjective Probability Beliefs (Experiment 1). *Note:* Error bars are 95% confidence intervals.

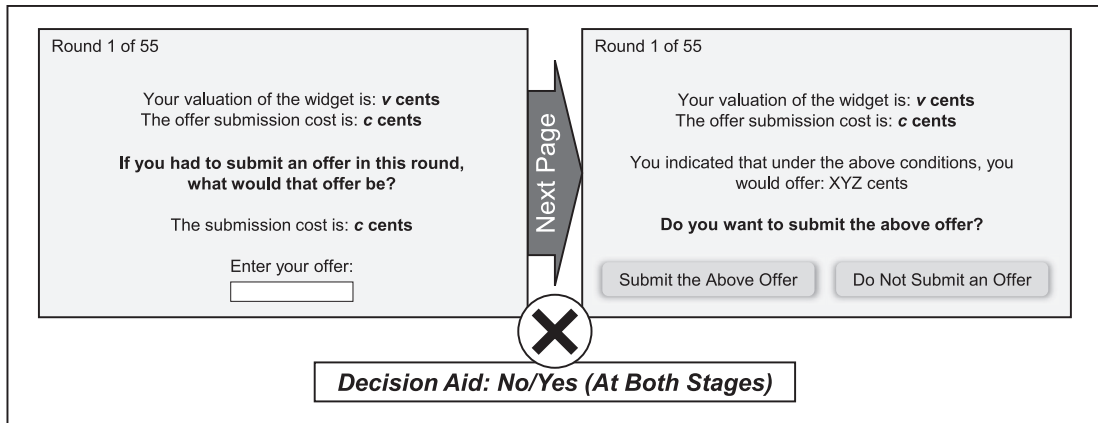


Fig. 7. Overview of Experiment 2 (“Backward”): Nudge Towards the Normative Process.

Table 3
Individual-Level Estimates of the CRRA Model of Experiment 2, by Data Used.

Data used in estimation	Offer data		Entry data	
	No aid	Aid	No aid	Aid
Decision aid condition	No aid	Aid	No aid	Aid
Mean CRRA (R)	6.51	6.63	1.81	2.13
Median CRRA (R)	9.03	8.95	0.81	1.10
Population stdev. CRRA (R)	3.74	3.63	2.55	2.68
RMSE bid fit	10.4	10.5	15.8	15.0
Estimated Pr(entry no entry)	14%	16%	11%	11%
Estimated Pr(entry entry)	59%	64%	78%	86%

Note: Only participants with at least 5 submitted offers are analyzed. (123 in “No aid”, 100 in “Aid”).

solution strategy eventually 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 4 (Incentivized normative decision architecture)

Instead of merely nudging people toward the normative process as we did in Experiment 2, we now incentivize the first stage by submitting all offers and paying the participants accordingly. After a filler task, we then incentivize the second stage (entry) given the response from the first stage (offer amount). The exact procedure was as follows (see Fig. 9): the experiment started with the same 9 training rounds that we used in the previous experiments. Participants then 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, which were chosen to be along the diagonal of the (v,c) space, as well as in the lower-left triangle of it—in order to empirically separate the effect of valuation from the effect of cost.²³ This block was followed by a filler task with demographic questions. In the next 21 rounds (fixed offers), we used participants’ offers from the forced-entry rounds, presented them in random order as exogenously given (one at a time and without informing the participants that they were, in fact, their own), and asked participants to decide whether they wanted to submit them. In the entry stage, we chose to pipe participants’ own offers to them to control for preference heterogeneity and to allow a direct comparison of the entry behavior with equivalent design cells in Experiment 1. To avoid any issues of sunk-cost fallacy or escalation of commitment, we did not inform participants of the exact source of the exogenous offers. Instead, we simply described the offer amounts as “predetermined” at the time of their presentation.

As before, whenever a participant 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, participants reported the extent of their agreement with the same statements we used in the previous experiments regarding what they were thinking during the task (see section 3). As with the other experiments, participants received a base payment of \$2.50, and they received their

²³ 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).

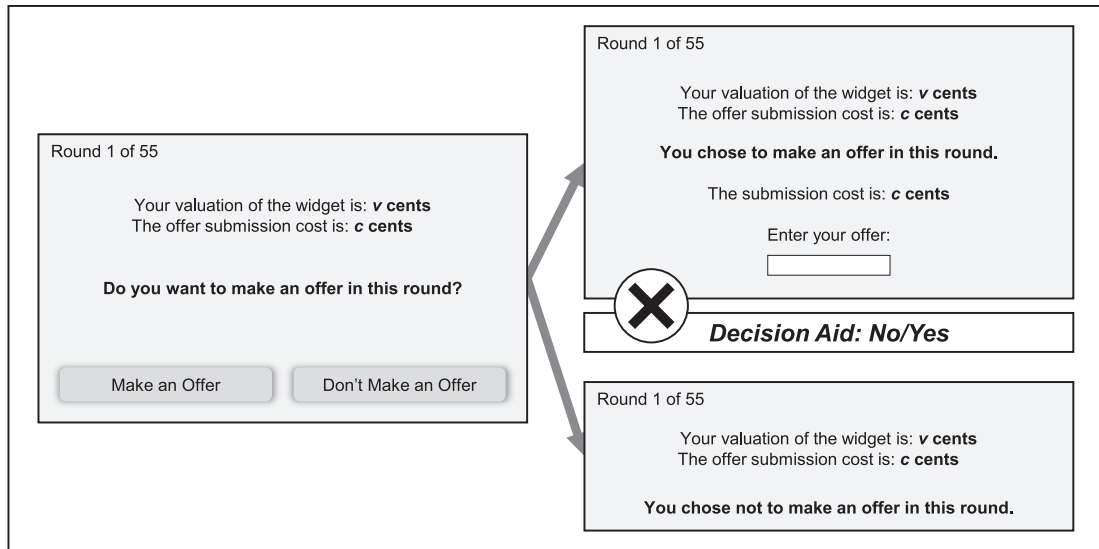


Fig. 8. Overview of Experiment 3: Forward Decision Architecture.

Table 4
Individual-Level Estimates of the CRRA Model of Experiment 3, by Data Used.

Data used in estimation	Offer data		Entry data	
	No aid	Aid	No aid	Aid
Decision aid condition	No aid	Aid	No aid	Aid
Mean CRRA (R)	7.43	6.97	2.36	2.08
Median CRRA (R)	9.13	9.13	1.14	0.84
Population stdev. CRRA (R)	3.14	3.49	3.22	2.91
RMSE bid fit	7.2	9.7	16.4	17.1
Estimated Pr(entry no entry)	17%	20%	13%	13%
Estimated Pr(entry entry)	58%	61%	80%	86%

Note: Only participants with at least 5 submitted offers are analyzed. (118 in “No aid”, 105 in “Aid”).

final account balance (including the \$1 starting balance) from all blocks as a bonus payment after completing the experiment.

A total of 242 US residents ($M_{\text{age}} = 40.48$, $SD_{\text{age}} = 12.26$; 53.72 % female) passed the screening questions²⁴ and participated in the study, 101 of whom had access to the decision aid. The average participant earned about \$4.78 and took about 31 min (median 29 min) to complete the entire experiment. Only one participant ended up losing both the \$1 and the payoff from the Holt and Laury task; they received the base payment. All remaining participants received their total payoffs.

7.1. Results: Preference inconsistency persists even in normative architecture

Table 5 is an analog of Tables 2, 3, and 4, and it shows the same preference inconsistency. Compared with the closest-related “backward nudge” design in Experiment 2 (Table 3), we see that the mean CRRA is substantially higher, and the entry model fits worse—especially when it is calibrated on the offer data. In the next subsection, we discuss what entry and offer behavior corresponds to these parametric patterns.

7.2. Results: Incentivizing the normative process instead of merely suggesting it increases offer magnitudes, but leaves entry unaffected

We start the discussion of the results in the condition without the decision aid and compare the observed behavior with that in Experiment 2 (backward nudge) to isolate the effect of incentivizing the normative process instead of merely suggesting it. Fig. 10 shows the difference in behavior due to the added offer incentives. Consistent with the higher mean CRRA dis-

²⁴ The screening questions that check understanding disqualified 54% of invited participants who consented to participate in the study.

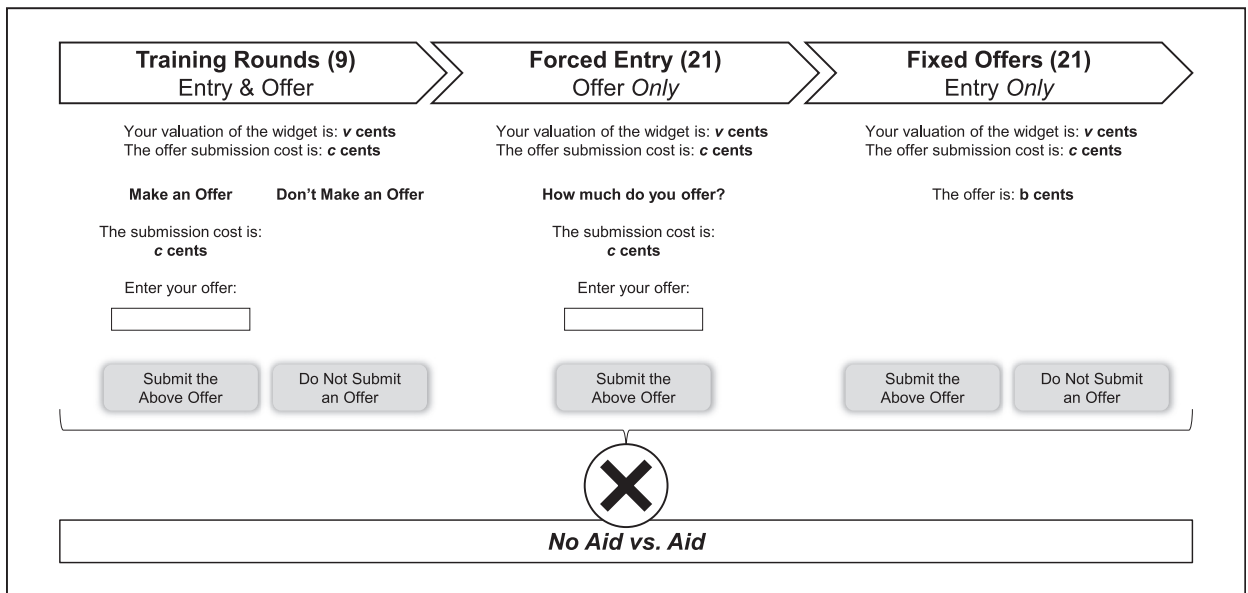


Fig. 9. Overview of Experiment 4: Incentivized Normative Decision Architecture.

Table 5
Individual-Level Estimates of the CRRA Model of Experiment 4, by Data Used.

Data used in estimation	Offer data		Entry data	
	No aid	Aid	No aid	Aid
Decision aid condition	No aid	Aid	No aid	Aid
Mean CRRA (<i>R</i>)	8.57	7.73	1.95	2.24
Median CRRA (<i>R</i>)	9.25	9.23	0.87	1.48
Population stdev. CRRA (<i>R</i>)	2.38	3.15	2.80	2.92
RMSE bid fit	9.9	10.2	17.1	15.1
Estimated Pr(entry no entry)	28%	31%	29%	29%
Estimated Pr(entry entry)	53%	57%	72%	78%

Note: Only participants with at least 5 submitted offers are analyzed. (139 in “No aid”, 101 in “Aid”).

cussed above, the submitted offers increase in all 21 (*v,c*) design cells, and the increase is highly statistically significant. On average, across all 21 cells, submitted offers increase by about 8 % relative to those in Experiment 2.

The dramatic increase in offers is not accompanied by a systematic change in entry (during the “Entry Only” block of the study flow; see Fig. 10): entry probabilities remain nearly unchanged in most cells, the observed changes involve a mix of signs, and we cannot reject the hypothesis that the entry pattern is identical to that in Experiment 2 (such a rejection would require familywise significance after the Bonferroni correction). The decision aid attenuates the effects described above but does not completely eliminate them.

The pattern of effects in Fig. 10 therefore implies that the entry decisions in Experiments 2–4 do not seem to take into account offer magnitudes, as Eq. (1) suggests they should. Under the expected-utility model, for example, an increase in offers would indicate an increase in risk aversion and thus should be accompanied by a decline in entry along and below 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 Eq. (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.

The higher offers in the incentivized treatment are potentially consistent with increased risk-aversion among the participants compared to Experiment 2. At the end of each experiment, we measured subjective risk-aversion using the classic “risk-taker” scale of Dohmen et al. (2012) to test whether Experiment 4 somehow made people more risk-averse than Experiment 2. We find no effect; please see Web Appendix F for details. Returning to the preference inconsistency findings in Table 5, the pattern documented in Fig. 10 shows how incentivizing the normative process exacerbates the inconsistency puzzle instead of resolving it: offers increase relative to Experiment 2, but entry remains about the same.

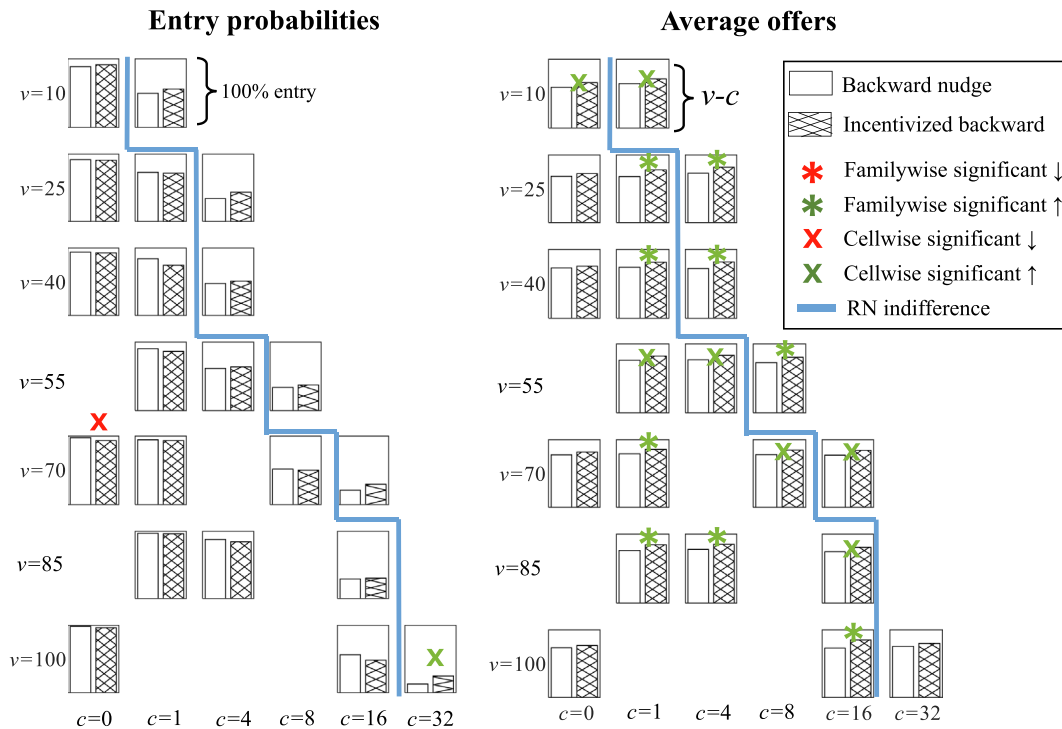


Fig. 10. Effect of Incentivizing Normative Process, Condition without Decision Aid. *Note:* In the entry plot, the height of each box corresponds to 100 % of participants entering in that (v,c) cell. In the average offer plot, the height of each box is $v - c$, the upper bound on admissible offers. An “x” indicates a significant effect of incentives at the 5 % level according to the z-test performed within the (v,c) cell only. A “*” indicates a significant effect of incentives at the familywise 5 % level, adjusted with the Bonferroni correction for the 39 tests in each plot ($p < 0.0024 = 0.05/21$). The color denotes the direction of the change (red = decrease, green = increase). The (blue) thick line is the indifference curve of the risk-neutral model. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

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, as if each participant turned into a very risk-averse “Dr. Bidder” when setting the offer amount, but at the same time entered the game as a nearly risk-neutral “Mr. Entrant”. This inconsistency is persistent in the sense that none of our decision-support manipulations reconciled the estimated preferences revealed from the two parts of the task. 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. The inconsistency also persists in a non-normative sequential architecture, which rules out potential alternative explanations based on bracketing (Read, Loewenstein, and Rabin 1999) or reference-point shifts (Kahneman and Tversky 1979), where Dr. Bidder has the same preferences as Mr. Entrant, but either applies a narrower decision focus or uses a different reference point when bidding.

Flexible specifications of the utility function, such as a generalization of prospect theory, do not resolve the inconsistency for two reasons: First, we estimated one such specification, and the inconsistency persisted (details in Web Appendix C). Second, and more generally, our manipulations do not affect the two parts of behavior in an internally consistent way: In Experiment 1, we found that reducing decision difficulty through a decision aid affects entry in a manner consistent with the effect on probability beliefs, but does not affect offer amounts, as (a) would predict. In Experiment 4, we found that incentivizing consumers to strictly follow the normative process affected 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 Eq. (1), not just with a particular popular assumption about the shape of the utility function (i.e., expected-utility theory or prospect theory).

Previous research assumed the analogous inconsistency in auctions 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 participants did not become more internally consistent when they received computational support or when they were guided through the normative process. Thus, the reference inconsistency appears to be a robust property of consumer decision-making under costly price offers, not a mistake to be “fixed”.

The inconsistency 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](#); [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 CRRA models on complete data is problematic because the models do not fit the complete behavior well. In other words, the assumption that a single set of preferences determines both bidding and entry appears too strong. Our findings thus present a challenge to modelers of offers and bidding with costly entry to develop new behaviorally realistic structural (i.e., based on some deep fixed preferences and involving some sort of consumer optimization) models not constrained by the normative structure of [Fig. 1](#).

In addition to informing future modeling and re-interpreting the preference inconsistency in terms of a non-normative decision process, our results also have immediate (as in not needing any modeling or econometrics) implications for market designers who need to design the interface they provide to consumers. Consider a policy maker who is interested in maximizing consumer earnings (i.e., consumer welfare under the risk-neutral model). In our baseline condition that mimics a simple real-world interface, our participants earn only about half the theoretical maximum monetary (43 % to be exact). With half of the potential consumer surplus being left on the table, there is substantial room for improvement, especially if the baseline behavior is due to a mistake or a bounded-rationality outcome. Theory and [Fig. 1](#) clearly predict that the incentivized backward architecture should be the best, and the decision aid should help. But what do our results recommend?

[Fig. 11](#) shows the median consumer earnings in Experiment 4 and the 21 (v,c) design cells that match the cells from Experiments 1, 2, and 3. Surprisingly, the normative architecture does not emerge as the winner: as long as the decision aid is not available, the incentivized backward treatment is actually the worst of all treatments we considered. Specifically, forcing participants to make a serious offer in the first stage reduces their ultimate consumer earnings by 45 % compared to participants who were allowed to make a casual 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 $p = .06$).

The pattern of earnings in [Fig. 11](#) is consistent with participants struggling to think through contingent payoffs in the first (offer) stage of the backward architecture, possibly focusing on other objectives at that moment—such as getting their offer accepted and forgetting about the cost of submission (in [Web Appendix E](#), we provide model-free evidence that both backward architectures indeed reduce the extent to which participants take submission cost into account when formulating offers). Another takeaway from [Fig. 11](#) is that the decision aid is not always helpful to consumer earnings: it does not help in the simultaneous architecture.

The preference inconsistency we find may exist in other market behaviors that have multiple component “sub-decisions”. For example, it may be fruitful to investigate the case of uncertain contingent rebates that involve a perceived cost of thinking akin to the cost of making an offer here ([Ailawadi et al. 2014](#)).

We now turn to the limitations. In any estimation of heterogeneous preferences, there is a tradeoff between flexibility in the model of heterogeneity and the model of individual behavior. Ideally, both models would be fully non-parametric, with




Fig. 11. Consumer Earnings, by Decision Architecture and Aid. *Note:* The expected earnings of a risk-neutral agent are 114 cents, literally off the chart. All bars represent the median consumer surplus in the 21 (v,c) design cells matched to the cells used in Experiment 4. Error bars represent the 95% confidence intervals.

preferences estimated at the individual level and each utility function allowed to be fully flexible. However, it is rare for a paper to have enough informative data per individual to accomplish this ideal—with one example being [Gonzalez and Wu \(1999\)](#), who accomplish full nonparametric estimation of a prospect-theoretic model at the individual level using a titration task. Our data is not as plentiful, so we resort to a parametric assumption about the utility shape while still estimating each participant's preferences at the individual level. It may be possible to solve the tradeoff in the opposite way, allowing for more flexibility at the individual level at the cost of partial pooling across participants via a Bayesian hierarchical model. Future research could explore this challenge.

Our tightly controlled experiments are well suited to identifying causal effects and understanding 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 (for more discussion, see [Charness et al. 2013](#)). 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.

CRediT authorship contribution statement

Robert Zeithammer: Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization. **Lucas Stich:** Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization. **Martin Spann:** Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization. **Gerald Häubl:** Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization.

Data availability

Data and code are available on OSF (https://osf.io/q3ysf/?view_only=47df578c500b4ffc91aa06868e0e1e76)

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary materials

Supplementary materials to this article can be found online at <https://doi.org/10.1016/j.ijresmar.2024.08.006>.

References

- Akerberg, 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.
- Ailawadi, K. L., Gedenk, K., Langer, T., Ma, Y., & Neslin, S. A. (2014). Consumer response to uncertain promotions: An empirical analysis of conditional rebates. *International Journal of Research in Marketing*, 31(1), 94–106.
- Amaldoss, W., & Jain, S. (2008). Joint Bidding in the Name-Your-Own-Price Channel: A Strategic Analysis. *Management Science*, 54(10), 1685–1699.
- Arkes, H. R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes*, 35(1), 124–140.
- Bajari, P., & Hortacsu, A. (2005). Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data. *Journal of Political Economy*, 113(4), 703–741.
- Bernhardt, M., & Spann, M. (2010). An Empirical Analysis of Bidding Fees in Name-Your-Own-Price Auctions. *Journal of Interactive Marketing*, 24(4), 283–296.
- Charness, G., Gneezy, U., & Kuhn, M. A. (2013). Experimental methods: Extra-laboratory experiments-extending the reach of experimental economics. *Journal of Economic Behavior & Organization*, 91, 93–100.
- Cox, J. C., Smith, V. L., & Walker, J. M. (1988). Theory and Individual Behavior of First-Price Auctions. *Journal of Risk and Uncertainty*, 1(1), 61–99.
- Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., & Wagner, G. G. (2011). Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences. *Journal of the European Economic Association*, 9(3), 522–550.
- Dohmen, T., Falk, A., Huffman, D., & Sunde, U. (2012). The Intergenerational Transmission of Risk and Trust Attitudes. *Review of Economic Studies*, 79(2), 645–677.
- Ertac, S., Hortacsu, A., & Roberts, J. W. (2011). Entry into Auctions: An Experimental Analysis. *International Journal of Industrial Organization*, 29(2), 168–178.
- Fay, S. (2009). Competitive Reasons for the Name-Your-Own-Price Channel. *Marketing Letters*, 20(3), 277–293.
- Fay, S. (2004). Partial-Repeat-Bidding in the Name-Your-Own-Price Channel. *Marketing Science*, 23(3), 407–418.
- Gonzalez, R., & George, Wu. (1999). On the Shape of the Probability Weighting Function. *Cognitive Psychology*, 38(1), 129–166.
- Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal Nonparametric Estimation of First-Price Auctions. *Econometrica*, 68(3), 525–574.
- Gunia, B. C., Sivanathan, N., & Galinsky, A. D. (2009). Vicarious Entrapment: Your Sunk Costs, My Escalation of Commitment. *Journal of Experimental Social Psychology*, 45(6), 1238–1244.
- Hann, I.-H., & Terwiesch, C. (2003). Measuring the Frictional Costs of Online Transactions: The Case of a Name-Your-Own-Price Channel. *Management Science*, 49(11), 1563–1579.
- Haruvy, E., Popkowski, P. T. L., & Leszczyc (2009). Bidder Motives in Cause-Related Auctions. *International Journal of Research in Marketing*, 26(4), 324–331.
- Hauser, J. R., & Wernerfelt, B. (1990). An Evaluation Cost Model of Consideration Sets. *Journal of Consumer Research*, 16(4), 393–408.
- Häubl, G., & Trifts, V. (2000). Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids. *Marketing Science*, 19(1), 4–21.
- Holt, C. A., & Laury, S. K. (2002). Risk Aversion and Incentive Effects. *American Economic Review*, 92(5), 1644–1655.
- Jindal, P., & Newberry, P. (2018). To Bargain or Not to Bargain: The Role of Fixed Costs in Price Negotiations. *Journal of Marketing Research*, 55(6), 832–851.
- Johnson, E. J., Shu, S. B., Dellaert, B. G. C., Fox, C., Goldstein, D. G., Häubl, G., Larrick, R. P., Payne, J. W., Peters, E., Schkade, D., Wansink, B., & Weber, E. U. (2012). Beyond Nudges: Tools of a Choice Architecture. *Marketing Letters*, 23(2), 487–504.

- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263–292.
- Krasnokutskaya, E., & Seim, K. (2011). Bid Preference Programs and Participation in Highway Procurement Auctions. *American Economic Review*, 101(6), 2653–2686.
- Laferrrière, V., Staubli, D., & Thöni, C. (2023). Explaining Excess Entry in Winner-Take-All Markets. *Management Science*, 69(2), 1050–1069.
- Levin, D., & Smith, J. L. (1994). Equilibrium in Auctions with Entry. *American Economic Review*, 84(3), 585–599.
- Li, Y., Krefeld-Schwalb, A., Wall, D. G., Johnson, E. J., Toubia, O., & Bartels, D. M. (2022). The More You Ask, the Less You Get: When Additional Questions Hurt External Validity. *Journal of Marketing Research*, 59(5), 963–982.
- Lichtenstein, S., & Slovic, P. (1973). Reversals of Preference Between Bids and Choices in Gambling Decisions. *Journal of Experimental Psychology*, 89(1), 46–55.
- Martens, N., & Orzen, H. (2021). Escalating commitment to a failing course of action – A re-examination. *European Economic Review*, 137, 10381.
- McAfee, R. P., & McMillan, J. (1987). Auctions with Entry. *Economic Letters*, 23(4), 343–337.
- Menezes, F. M., & Monteiro, P. K. (2000). Auctions with Endogenous Participation. *Review of Economic Design*, 5, 71–89.
- Moreno, D., & Wooders, J. (2011). Auctions with Heterogeneous Entry Costs. *RAND Journal of Economics*, 42(2), 313–336.
- Palfrey, T. R., & Pevnitskaya, S. (2008). Endogenous Entry and Self-Selection in Private Value Auctions: An Experimental Study. *Journal of Economic Behavior & Organization*, 66(3–4), 731–747.
- Rabin, M., & Weizsäcker, G. (2009). Narrow Bracketing and Dominated Choices. *American Economic Review*, 99(4), 1508–1543.
- Read, D., Loewenstein, G., & Rabin, M. (1999). Choice Bracketing. *Journal of Risk and Uncertainty*, 19(1/3), 171–197.
- Samuelson, W. F. (1985). “Competitive Bidding with Entry Costs.” *Economics Letters*, 17(1–2), 53–57.
- Shapiro, D. (2011). Profitability of the Name-Your-Own-Price Channel in the Case of Risk-Averse Buyers. *Marketing Science*, 30(2), 290–304.
- Smith, J. L., & Levin, D. (1996). Ranking Auctions with Risk Averse Bidders. *Journal of Economic Theory*, 68(2), 549–561.
- Smith, V. L. (1976). Experimental Economics: Induced Value Theory. *American Economic Review*, 66(2), 274–1229.
- Spann, M., Zeithammer, R., Bertini, M., Haruvy, E., Jap, S. D., Koenigsberg, O., Mak, V., Popkowski, P. T. L., Leszczyc, B. S., & Thomas, M. (2018). Beyond Posted Prices: The Past, Present, and Future of Participative Pricing Mechanisms. *Customer Needs and Solutions*, 5(1–2), 121–136.
- Spann, M., Zeithammer, R., & Häubl, G. (2010). Optimal Reverse-Pricing Mechanisms. *Marketing Science*, 29(6), 1058–1070.
- Staw, B. 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, A., & Kahneman, D. (1992). Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323.
- Tversky, A., Slovic, P., & Kahneman, D. (1990). The Causes of Preference Reversal. *American Economic Review*, 80(1), 204–217.
- Yao, S., & Mela, C. F. (2008). Online Auction Demand. *Marketing Science*, 27(5), 861–885.
- Zeithammer, R. (2015). Optimal Selling Strategies When Buyers Name Their Own Prices. *Quantitative Marketing and Economics*, 13(2), 135–171.