Revealed Preference Analysis of Expected Utility Maximization under Prize-Probability Trade-Offs

Laurens Cherchye
KULeuven

Thomas Demuynck
SBS-EM, ECARES, Université libre de Bruxelles

Bram De Rock
SBS-EM, ECARES, Université libre de Bruxelles

Mikhail Freer
University of Essex

December 2019

ECARES working paper 2019-27

ECARES
ULB - CP 114/04
50, F.D. Roosevelt Ave., B-1050 Brussels BELGIUM
www.ecares.org

Revealed Preference Analysis of Expected Utility Maximization under Prize-Probability Trade-Offs*

Laurens Cherchye† Thomas Demuynck‡ Bram De Rock§ Mikhail Freer¶

November 27, 2019

Abstract

We provide a revealed preference characterization of expected utility maximization in binary lotteries with prize-probability trade-offs. This characterization applies to a wide variety of decision problems, including first price auctions, crowdfunding games, posted price mechanisms and principal-agent problems. We start by characterizing optimizing behavior when the empirical analyst exactly knows either the probability function of winning or the decision maker's utility function. Subsequently, we provide a statistical test for the case where the utility function is unknown and the probability function has to be estimated. Finally, we consider the situation with both the probability function and utility function unknown. We show that expected utility maximization has empirical content when these functions satisfy log-concavity assumptions. We demonstrate the empirical usefulness

*This paper replaces and extends our previous working paper "Equilibrium play in first price auctions: Revealed preference analysis". We also thank the reading group at the University of Sussex for useful discussion.

†Department of economics, University of Leuven. E. Sabbelaan 53, B-8500 Kortrijk, Belgium. E-mail: laurens.cherchye@kuleuven.be. Laurens Cherchye gratefully acknowledges the European Research Council (ERC) for his Consolidator Grant 614221. Part of this research is financed by the Fund of Scientific Research Flanders (FWO-Vlaanderen).

‡ECARES, Université Libre de Bruxelles. Avenue F. D. Roosevelt 50, CP 114, B-1050 Brussels, Belgium. E-mail: thomas.demuynck@ulb.be. Thomas Demuynck acknowledges financial support by the Fonds de la Recherche Scientifique-FNRS under grant nr F.4516.18

§ECARES, Université Libre de Bruxelles, and Department of Economics, University of Leuven (KU Leuven). Avenue F. D. Roosevelt 50, CP 114, B-1050 Brussels, Belgium. E-mail: bram.de.rock@ulb.be. Bram De Rock gratefully acknowledges FWO and FNRS for their financial support.

¶Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, United Kingdom. E-mail: m.freer@essex.ac.uk.

of our theoretical findings through an application to an experimental data set.
Keywords: expected utility maximization, prize-probability trade-offs, revealed preference characterization, testable implications, experimental data.

1 Introduction

We analyze models of expected utility maximization in which the decision maker (DM) faces a binary lottery that is characterized by a prize-probability trade-off. In particular, the lottery yields a reward $r-b$ with probability $P(b)$ and a payoff of zero with probability $(1-P(b))$, where the value of r is exogenously given and P is a cumulative distribution function. The DM's problem is to choose the optimal value of b. In other words, she faces a trade-off between the value of the prize and the probability of winning the prize.

This type of decision problem occurs frequently in economics. A prime example is a first price auction where the DM is one of the participants. In this case the prize of the lottery is given by the value r of the object for the DM minus the DM's bid b to win the auction. The DM can increase the probability of winning the auction by increasing her bid b, but this implies that the final value of winning the auction, i.e. $(r-b)$, decreases. We will discuss additional examples of often studied decision problems characterized by prize-probability trade-offs further on.

Our main contribution is that we develop a revealed preference approach to characterize behavior that is expected utility maximizing under price-probability trade-offs. A distinguishing and attractive feature of our revealed preference characterizations is that they do not require a (non-verifiable) functional specification of the optimization problem. They define testable conditions for optimizing behavior that are intrinsically nonparametric and, therefore, robust to specification bias. To define these testable conditions, we will assume that the empirical analyst can use, for a given DM, a sequence of observations on rewards r (received when winning the lottery) and on money amounts b (called "bids" in what follows) that the DM is willing to forego in order to increase her probability of winning.

Testing consistency of behavior with expected utility maximization constitutes a prominent problem in the economics literature. It is particularly relevant in the context of market organization and mechanism design. Apart from auctions, another notable example of an experimentally studied game with a prize-probability trade-off is the bilateral trade problem (see, for example, Abrams, Sefton, and Yavas (2000); Cason, Friedman, and Milam (2003)). One of the best-studied trade mechanisms is the posted price mechanism, in which the buyer or seller posts a price and the other party can choose to either accept or reject this offer. This mechanism is considered to be an important alternative to the standard double auction. Recent research has drawn considerable attention to this mechanism by arguing that it is more robust to limited strategic sophistication of players (Borgers and Li, forthcoming). Still, the relevance of the mechanism crucially depends on the rationality of the participants' behavior. In this respect, the existing empirical literature on checking rationality usually proceeds by imposing parametric structure on the participants' preferences. This approach is clearly vulnerable in that refutations of rational behavior might simply be due to parametric misspecification. As we will show in Section 2, posted price problems also fit in the general set-up that we analyze nonparametrically in the current paper.

As a preliminary remark, the nonparametric revealed preference approach that we present in this paper follows the tradition of Afriat (1967), Diewert (1973) and Varian (1982). From this perspective, our paper also fits into the literature that developed revealed preference tests for decision models with uncertainty. Varian (1983) was the first to derive a revealed preference tests for the expected utility maximization model. In follow-up work, Green and Osband (1991) introduced alternative revealed preference tests of the expected utility maximization model, and Echenique and Saito (2015) focused on subjective expected utility maximization. A main distinguishing feature of our study pertains to the fact that these other studies all assume that the objective or subjective probabilities characterizing the DM's uncertainty are exogenous for the DM. By contrast, in our set-up the DM can impact the probability of winning by varying her bids.[1]

Overview of results. We start by assuming that the analyst perfectly knows either the probability of winning P (as a function of b) or the DM's utility function U (as a function of $r-b$). For this set-up, we show that the assumption of expected utility maximization generates strong testable implications. Particularly, we derive a revealed preference characterization of optimizing behavior that takes the form of a set of inequalities that are linear in unknowns. The characterization defines necessary and sufficient conditions for the existence of a utility function (when P is known) or a probability function (when U is known) such that the DM's observed decisions on b are consistent with expected utility maximization.

Admittedly, the assumption that either U or P is exactly known is too demanding in many practical settings. Therefore, we proceed our analysis by dropping this assumption. As a first case, we consider that U is unknown, but the empirical analyst can use a finite number of i.i.d. observations to estimate P. For this situation, we construct a statistical test for consistency of observed bidding behavior with expected utility maximization. This test has multiple practical applications. For

[1] Kim (1996) also allows for endogenous probabilities in a setting that is formally close to ours. However, he considers finite choice sets, while our focus is on continuous choice sets. This allows for covering a wide range of economic decision problems.

instance, in our own empirical exercise (in Section 5) we will apply the test to an experimental data set on first price auctions. An alternative interesting application pertains to crowd-funding games with refunds if the project is not successful. In this case, the reward is the valuation of the project by a certain participant (DM) and the bid is the amount of money the DM promises to the project. In practice, the probability P can then be estimated by using observed data on biddings from similar crowd-funding games in the past.

In various other settings, however, the empirical analyst cannot even construct a reliable estimate of P, which makes that both U and P are unknown. For example, this observability restriction applies in a natural way to dynamic choice problems of the principal-agent type. In a most simple specification, the principal is the DM who can obtain a reward r of which the probability depends on the agent's effort level. The DM can choose to increase the effort level of the agent by increasing the wage b offered to the agent, with this wage being paid only if the principal receives the reward. The lottery's value then equals the reward r minus the wage b offered to the agent, and the probability of winning the prize increases in the wage rate b offered to the agent. In such a setting, the probability of success P is usually determined in the second stage of the game, which is often not observed by the empirical analyst. Typically, the analyst only observes the wages b but not the effort levels that govern the probability of success.

Not very surprisingly, we find that the assumption of optimizing behavior does not generate any testable restrictions for observed behavior when not imposing any structure on U and P. Interestingly, however, we also show that this negative conclusion can be overcome when putting shape constraints on U or P that are often used in the relevant literature. Specifically, we focus on the following three cases: (1) P is strictly log-concave; (2) U is strictly log-concave; and (3) both P and U are strictly log-concave.[2] For each of these models, we derive necessary and sufficient testable conditions for expected utility maximization: (1) requires that higher rewards r must lead to higher payoffs $r - b$; (2) requires that higher rewards r must lead to higher bids b; and (3) requires that higher rewards r must lead to both higher payoffs $r - b$ and higher bids b. These results are in agreement with comparative static results that have been documented in the literature (see, for example, Cox and Oaxaca (1996) for the first price auction model). A notable implication of our nonparametric characterizations is that the testable conditions are not only necessary but also sufficient for expected utility maximization.

Importantly, these characterizations entail two additional conclusions. First, they show that the assumption of expected utility maximization does have empirical content even under minimalistic shape restrictions for P and/or U. Moreover,

[2] See, for example, Bagnoli and Bergstrom (2005) for an overview of alternative applications that make use of such log-concave specifications.

as we will discuss in Section 6, even if the rewards r are unobserved, the above comparative static results (for the scenarios (1), (2) and (3)) still enable partial identification of the reward structure when (only) using information on the observed bids. Second, our result for scenario (1) shows that, for any log concave P and any data set with payoffs $r - b$ increasing in rewards r, we can find utility functions U such that the combination (P, U) generates this observed data set. Similarly, it follows from our result for scenario (2) that, for any log concave utility functions U and any data set with bids b increasing in rewards r, we can construct a probability distributions P such that (P, U) generates the data set. In other words, even if we assume that either P or U is log-concave, it turns out to be empirically impossible to (partially) identify the functions. These findings are similar in spirit to those of Manski (2002, 2004) on the impossibility to separately identify decision rules and beliefs.

We also illustrate the empirical usefulness of our theoretical results through an application to Neugebauer and Perote (2008)'s experimental data set on first price auctions. We consider (1) the setting with P estimated and (2) the setting with both P and U unknown but at least one of them assumed to be log-concave. We will conclude that the hypothesis of expected utility maximization cannot be rejected for a large fraction of the subjects in the sample. Further, by comparing our results for the settings (1) and (2) we show that some observed DMs behave consistently with expected utility maximization for the actual distribution of bids in the population, while other DMs are better modeled as expected utility maximizers with respect to a subjective (log-concave) belief of this distribution (which may be very different form the true distribution).

Outline. The remainder of this paper is structured as follows. Section 2 introduces our theoretical set-up and notation. It also provides a more formal description of the above cited examples of decision problems that fit in our general framework. Section 3 considers the case in which the empirical analyst knows either the probability function P or the utility function U. Section 4 analyzes the setting with both P and U unknown. Section 5 presents an illustrative application to experimental data. Section 6 discusses the usefulness of our theoretical results when the rewards r are unobserved (which is often relevant in non-experimental settings). Section 7 presents our concluding discussion. The Appendix contains the proofs of our main theoretical results as well as some robustness exercises related to our empirical application.

2 Set-up and notation

As explained in the introductory section, we consider a setting where the DM can win a reward r with a certain probability. We assume that $r > 0$ and $r \leq \overline{r}$ for some exogenously given $\overline{r} \in \mathbb{R}$. The DM can choose a bid $b \in [0, \overline{r}]$. Choosing a higher value of b increases the probability of winning the reward. We model this through a latent random variable $\tilde{b}$ (unobserved by the DM) with cumulative distribution function (cdf) P such that the award is won whenever $b \geq \tilde{b}$. In other words, the probability of winning is equal to $P(b) = \Pr(\tilde{b} \leq b)$. The downside of increasing b is that the value of winning is decreasing with the bid. As such, the DM obtains $r - b$ if the reward is won (with probability $P(b)$), while the DM's payoff is zero if the reward is not won (with probability $1 - P(b)$).

The standard expected utility model assumes that the DM has a Bernoulli utility function

$$U : [0, \overline{r}] \to \mathbb{R}_+,$$

such that b solves

$$\max_{b \in [0,r]} P(b)\ U(r - b), \tag{1}$$

where we normalize the utility associated with zero payoff to zero, i.e. $U(0) = 0$. We will assume throughout that P is continuous and strictly increasing on $[0, \overline{r}]$, and that U is continuous and strictly increasing on $\mathbb{R}$. Observe that we can indeed restrict $b \leq r$ in this optimization problem, as any bid $b > r$ gives negative expected utility and is therefore dominated by a choice $b = r$, which gives zero expected utility.

This general set-up applies to a wide variety of decision problems that are frequently encountered in economics. We illustrate this by discussing in turn first price auctions, crowdfunding games, posted price problems and principal-agent problems.

First price auctions. In a first price auction, the DM (bidder) has a value r for the object. Placing a bid of b decreases the value of winning the auction to $r - b$, while it increases the probability of winning. In this case, the random variable $\tilde{b}$ is the value of the highest bid of all other participants, and $P(b) = \Pr(\tilde{b} \leq b)$ is the probability that the DM wins the auction. As such, the expected utility of participating in the auction is given by (1).[3]

[3]We refer to Kagel and Levin (2011) for an overview of the experimental literature on auctions.

Crowdfunding games. A crowdfunding game is an example of a mechanism to organize private provision of a public good.[4] The participants in the game make bids for the public good. If the sum of these bids is above a certain threshold, then the public good is provided. Otherwise the payoff to all participants is zero. This fits in our general set-up for the DM being a participant of the crowdfunding game and r being the DM's value of the public good. Placing a bid lowers the value of the public good to $r - b$ when the public good is provided. Let $\tilde{t}$ be the random variable capturing the sum of the bids of all other participants, and let t be the threshold above which the public good is provided. When using $\tilde{b} = t - \tilde{t}$, we can define the probability of providing the good by

$$Pr(b + \tilde{t} \geq t) = \Pr(t - \tilde{t} \leq b) = \Pr(\tilde{b} \leq b) = P(b),$$

and the DM's expected utility of participating in the crowdfunding game is given by (1).

Posted price problems. In a posted price problem, the DM (buyer) has a valuation r for the traded good.[5] In order to obtain the good, the DM posts a price b at which she is willing to buy the good. The seller (second-mover) then decides whether or not to accept this offer. The DM receives a reward of $r - b$ if the seller accepts, and a payoff of zero if the seller rejects. As such, the seller's decision is based on her (unobserved) value $\tilde{b}$ for the good, which we can assume to be random from the buyer's point of view. The seller will accept the offer if and only if the posted price is at least as large as her reservation price $\tilde{b}$. Thus, the probability of the trade is given by

$$\Pr(\tilde{b} \leq b) = P(b),$$

which makes that the DM's expected utility equals (1).

Principal-agent problems. In a principal-agent model, the DM (as principal) can receive a reward of size r with a probability that depends on the effort e of the agent. In order to stimulate the agent to exert effort, the principal can promise a conditional bonus of b to the agent, which the agent only gets if the principal receives the prize. Thus, the DM's payoff in case the effort is high enough equals $r - b$. It is also natural to assume that e is an increasing function of b, say $e(b)$, and

[4]Similar games are discussed by Tabarrok (1998) and Zubrickas (2014). We here consider a simplified version of the game in which there is no lottery reward and only refund of contributions. In this sense, we are closer to Tabarrok (1998). However, we do allow for differentiated (and not only binary) contributions, as in Zubrickas (2014).

[5]The literature also frequently considers the alternative version with the seller posting the price. It is easily verified that this seller-posted price problem equally fits in general set-up.

that the reward is received only if the value of e is above some random variable $\tilde{e}$. When using the random variable $\tilde{b} = e^{-1}(\tilde{e})$, we can define

$$\Pr(\tilde{e} \leq e(b)) = \Pr(\tilde{b} \leq b) = P(b),$$

such that the DM's expected utility is equal to (1).

3 When P or U is known

We assume that the empirical analyst observes a finite number of rewards and bids for a given DM.[6] We first consider a setting where the researcher either knows the cdf P or the utility function U. For these cases, we derive the nonparametric revealed preference conditions for consistency with expected utility maximization. Next, we relax the assumption that P is fully observable and (only) assume that the analyst can estimate the empirical distribution of P by using a finite sample of observed winning probabilities. Under this assumption we develop a statistical test of expected utility maximization. In Section 4, we will focus on the case where both P and U are unobserved.

3.1 Rationalizability

We assume that the empirical analyst observes a DM who decides T times on the value of the bid b for various values of the reward r. This defines the data set

$$D = (r^t, b^t)_{t=1}^T,$$

which contains a return $r^t > 0$ and corresponding bid $b^t \in [0, r^t]$ for each observation $t \leq T$.

For a given cdf P and a utility function U, we say that the data set D is (P, U)-rationalizable if the observed bids b^t maximize the expected utility of the DM given the primitives P and U. This yields the next definition.

Definition 1. *For a given cdf P and utility function U, a data set $D = (r^t, b^t)_{t=1}^T$ is (P, U)-**rationalizable** if $U(0) = 0$ and, for all observations $t = 1, \ldots, T$,*

$$b^t \in \underset{b \in [0, r^t]}{\operatorname{argmax}}\, P(b)U(r^t - b).$$

The following theorem provides the revealed preference conditions for a data set D to be rationalizable if the researcher knows either P (but not U) or U (but not P).[7]

[6]We discuss the case of unobserved rewards in Section 6.

[7]We slightly abuse notation in Theorem 1 by assuming that $P(x) = 0$ if $x < 0$

Theorem 1. *Let $D = (r^t, b^t)_{t=1}^T$ be a data set.*

1. *Let P be a cdf. Then, there exists a utility function U such that the data set D is (P, U)-rationalizable if and only if,*

   (a) *for all observations $t = 1, \ldots, T$, $P(b^t) > 0$ and $b^t < r^t$, and*

   (b) *there exist numbers $U^t > 0$ such that, for all observations $t, s = 1, \ldots, T$,*

   $$P(b^t)U^t \geq P(r^t - r^s + b^s)U^s.$$

2. *Let U be a utility function. Then, there exists a cdf P such that the data set $D = (r^t, b^t)_{t=1}^T$ is (P, U)-rationalizable if and only if,*

   (a) *for all observations $t = 1, \ldots T$, $b^t < r^t$, and*

   (b) *there exists numbers $P^t > 0$ such that, for all observations $t, s = 1, \ldots, T$,*

   $$P^t U(r^t - b^t) \geq P^s U(r^t - b^s).$$

Conditions 1.a and 1.b of Theorem 1 present a set of inequalities that give necessary and sufficient conditions for rationalizability when the cdf P is given. The inequalities in 1.b are linear in the unknown numbers U^t, which makes them easy to verify. Intuitively, every number U^t represents the utility of winning the auction in period t, i.e. $U^t = U(r^t - b^t)$. Further, condition 1.b corresponds to the individual's maximization problem in Definition 1. In particular, the expected utility of choosing the observed bid b^t should be at least as high as the expected utility of making any other bid, including the bid $r^t - r^s + b^s$. This yields the condition

$$\begin{aligned} P(b^t)U^t &= P(b^t)U(r^t - b^t) \\ &\geq P(r^t - r^s + b^s)U(r^t - r^t + r^s - b^s) \\ &= P(r^t - r^s + b^s)U^s. \end{aligned}$$

Next, conditions 2.a and 2.b present a set of inequalities that give necessary and sufficient conditions for rationalizability when the utility function U is given. In this setting, the numbers P^t can be interpreted as the probabilities of winning if the bid equals b^t, i.e. $P^t = P(b^t)$. It is required that the expected utility of choosing the bid b^t is at least as high as the expected utility of choosing another bid b^s, which yields

$$\begin{aligned} P^t U(r^t - b^t) &= P(b^t)U(r^t - b^t), \\ &\geq P(b^s)U(r^t - b^s) = P^s U(r^t - b^s). \end{aligned}$$

This shows that necessity of the conditions 1.a-1.b and 2.a-2.b in Theorem 1 is relatively straightforward and may seem a rather weak implication. Interestingly, however, Theorem 1 states that data consistency with these condition is not only necessary but also sufficient for rationalizability. Particularly, in Appendix A.1 we provide a constructive proof that specifies a data rationalizing utility function U and a data rationalizing cdf P based on the conditions in statements 1 and 2 of Theorem 1.

We conclude this subsection by illustrating the empirical content of the rationalizability conditions in Theorem 1. Particularly, we show that the conditions can be rejected as soon as the data set D contains (only) two observations. First, for conditions 1.a-1.b we assume a data set D with the observations t, s such that $r^s - b^s$, $r^t - b^t$, $r^t - r^s + b^s$, and $r^s - r^t + b^t$ are all strictly positive and

$$P(b^t) = \frac{1}{10} \qquad P(r^t - r^s + b^s) = \frac{1}{4},$$
$$P(b^s) = \frac{1}{3} \qquad P(r^s - r^t + b^t) = \frac{1}{2}.$$

Then, condition 1.b in Theorem 1 requires that there exists strictly positive U^t and U^s such that

$$\frac{1}{10}U^t \geq \frac{1}{4}U^s \Leftrightarrow \frac{U^t}{U^s} \geq 2.5, \text{ and}$$
$$\frac{1}{2}U^s \geq \frac{1}{3}U^t \Leftrightarrow \frac{U^t}{U^s} \leq 1.5,$$

which is impossible. We conclude that the data set is not rationalizable.

Next, for conditions 2.a-2.b we assume that $U(x) = x$, which means that utility is linear, and that both $r_t - b_t$ and $r_s - b_s$ are strictly positive. Then, we must have

$$\frac{P^t}{P^s} \geq \frac{r^t - b^s}{r^t - b^t} \text{ and } \frac{P^s}{P^t} \geq \frac{r^s - b^t}{r^s - b^s},$$

for any two observations t and s. Since at least one of the two right hand sides must be stirctly positive, it must hold that

$$\begin{aligned}
1 \geq& \frac{r^t - b^s}{r^t - b^t}\frac{r^s - b^t}{r^s - b^s} \\
\Leftrightarrow& (r^t - b^t)(r^s - b^s) \geq (r^t - b^s)(r^s - b^t) \\
\Leftrightarrow& -b^t r^s - r^t b^s \geq -b^s r^s - b^t r^t \\
\Leftrightarrow& (r^s - r^t)(b^s - b^t) \geq 0.
\end{aligned}$$

This is violated as soon as $r^t > r^s$ and $b^s > b^t$ (or vice versa).

3.2 Estimating P

So far we have assumed that either the cdf P or the utility function U is fully observed by the empirical analyst. In this section, we show how to use the characterization in statement 1 of Theorem 1 to derive a statistical test of rationalizability when U is unknown, but the empirical analyst can construct an estimate of the cdf P from a finite sample of observations. We will use this statistical test in our empirical application in Section 5.

More formally, let us assume that we have a random sample of m values $(\hat{b}_j)_{j\leq m}$, drawn i.i.d. from a cdf G. We assume that the cdf G can be linked to the cdf P by a function $\Gamma : [0,1] \to [0,1]$ such that, for all $b \in [0, \overline{r}]$,

$$P(b) = \Gamma\left(G(b)\right).$$

This function Γ will generally depend on the specific setting at hand. For instance, as we will explain in more detail in Section 5, our own empirical application to first price auctions will take G to represent the distribution of bids for a random participant, while P equals the distribution of the highest bid among all participants different from the DM. Then, for an auction with $k+1$ randomly drawn participants in total (i.e. k participants different from the DM) and independent bids, we get

$$P(b) = (G(b))^k,$$

which yields the function $\Gamma(x) = x^k$.

In what follows, we assume that the function $\Gamma : [0,1] \to [0,1]$ is known by the empirical analyst (or is part of the underlying decision making model). Of course, if it is possible to directly obtain i.i.d. draws from the distribution P, we can set Γ equal to the identity function.

Given the finite sample $(\hat{b}_j)_{j\leq m}$, it is possible to construct an estimator of the cdf G by using the empirical distribution function

$$\mathbb{G}_m(b) = \frac{1}{m}\sum_{j=1}^{m} \mathbf{1}[\hat{b}_j \leq b],$$

where $\mathbf{1}[.]$ is the indicator function that equals 1 if the premise is true and zero otherwise. This estimator has a small sample bias equal to

$$\varepsilon_m(b) = \mathbb{G}_m(b) - G(b).$$

Next, we recall that our characterization in statement 1 of Theorem 1 only requires us to evaluate the distribution P (and hence G) at a finite number of values $r^t - r^s + b^s$, where $P(r^t - r^s + b^s) > 0$ for $t, s \in \{1, \ldots, T\}$. From now on, we will assume that $G(r^t - r^s + b^s) > 0$ for all such t, s. Correspondingly, we construct a finite vector of errors ε_m, with entries[8]

$$(\varepsilon_m)_{t,s} = \mathbb{G}_m(r^t - r^s + b^s) - G(r^t - r^s + b^s).$$

The vector $\sqrt{m}\varepsilon_m$ has an asymptotic distribution that is multivariate normal with mean zero and variance-covariance matrix Ω, where

$$\Omega_{(t',s'),(t,s)} = \begin{cases} G(r^t - r^s + b^s)(1 - G(r^{t'} - r^{s'} + b^{s'})) \text{ if } r^t - r^s + b^s < r^{t'} - r^{s'} + b^{s'} \\ G(r^{t'} - r^{s'} + b^{s'})(1 - G(r^t - r^s + b^s)) \text{ if } r^{t'} - r^{s'} + b^{s'} < r^t - r^s + b^s \end{cases}.$$

Standard results yield

$$m\,\varepsilon_m'(\Omega)^{-1}\varepsilon_m \sim^a \chi^2(K)),$$

where $\sim^a$ denotes convergence in distribution and K is the size of the vector ε.[9]

Of course, in practice we do not observe the matrix Ω. We can approximate it using the finite sample analogue $\widehat{\Omega}_m$, where

$$(\widehat{\Omega}_m)_{(t',s'),(t,s)} = \begin{cases} \mathbb{G}_m(r^t - r^s + b^s)\left(1 - \mathbb{G}_m(r^{t'} - r^{s'} + b^{s'})\right) \text{ if } r^t - r^s + b^s < r^{t'} - r^{s'} + b^{s'} \\ \mathbb{G}_m(r^{t'} - r^{s'} + b^{s'})\left(1 - \mathbb{G}_m(r^t - r^s + b^s)\right) \text{ if } r^{t'} - r^{s'} + b^{s'} < r^t - r^s + b^s \end{cases}.$$

Because $\widehat{\Omega}_m$ is a consistent estimate of Ω, it follows that

$$m\,\varepsilon'(\widehat{\Omega}_m)^{-1}\varepsilon \sim^a \chi^2(K),$$

We use this last result as the basis for our asymptotic test of rationalizability. Specifically, consider the null hypothesis

$$H_0 : \left\{ \begin{array}{c} \text{there is a utility function } U \text{ such that the data set} \\ D = (r^t, b^t)_{t=1}^T \text{ is } (P, U)\text{-rationalizable.} \end{array} \right\}.$$

To empirically check this hypothesis, we can solve the following minimization problem.

OP.I: $$Z_m = \min_{e_m, \widehat{G}_{t,s} \in [0,1], U^t > 0} m\, e_m'(\widehat{\Omega}_m)^{-1} e_m,$$

s.t.
$$\forall t, s: \; e_{t,s} = \mathbb{G}_m(r^t - r^s + b^s) - \hat{G}_{t,s}, \tag{2}$$
$$\Gamma(\hat{G}_{t,t})U^t \geq \Gamma(\hat{G}_{t,s})U^s, \tag{3}$$
$$\Gamma(\hat{G}_{t,s}) < \Gamma(\hat{G}_{t',s'}) \text{ for all } r^t - r^s + b^s < r^{t'} - r^{s'} + b^{s'}. \tag{4}$$

[8] For simplicity, we assume that all values $r^t - r^s + b^s$ are distinct. Obviously, this does not affect the core of our argument.

[9] See, for example, Sepanski (1994).

If the hypothesis H_0 holds true, the above problem has a feasible solution with

$$\hat{G}_{t,s} = G(r^t - r^s + b^s).$$

As such, we must have

$$Z_m \leq \sum_{t \in T} m\, \varepsilon'_m (\widehat{\Omega}_m)^{-1} \varepsilon_m.$$

Let us denote by c_α the $(1-\alpha) \times 100$th percentile of the $\chi^2(K)$ distribution. Then, if H_0 holds, we obtain

$$\lim_{m \to \infty} \Pr[Z_m > c_\alpha] \leq \lim_{m \to \infty} \Pr\left[\sum_{t \in T} m\, \varepsilon'_m (\widehat{\Omega}_m)^{-1} \varepsilon_m > c_\alpha\right] = \alpha,$$

which implies that we can construct an asymptotic test of H_0 by solving problem **OP.I** for the given data set, to subsequently verify whether its solution value exceeds c_α.

In view of our following discussion, two remarks are in order. First, our empirical hypothesis test is conservative in nature when compared to the theoretical test (based on Theorem 1) that uses the true distributions P and G. Second, implementing our hypothesis test in principle requires solving the minimization problem **OP.I**, which is computationally difficult because the constraints (3)-(4) are highly nonlinear in the unknowns $\hat{G}_{t,s}$ and U^t. For some particular instances of the function Γ, however, it is possible to convert this problem into a problem that can be solved by standard algorithms. We will consider such an instance in our empirical application in Section 5.

4 Unobserved P and U

We next turn to the instance in which both the cdf P and utility function U are unknown to the empirical analyst. For example, this applies when the procedure outlined above cannot use a reliable estimate of P, or when the researcher simply wants to avoid using this statistical approach. We start by a negative result: if no structure is imposed on P and U, then any data set D is rationalizable (i.e. expected utility maximization has no empirical content). Subsequently, we show that this negative conclusion can be overcome by imposing a (strict) log-concavity condition on P or U or on both. The assumption of log-concavity is a natural candidate to impose minimal structure on the decision problem. Log-concavity of U still allows for various risk attitudes, as it only requires that subjects are not too risk-loving. Log-concavity of P also allows the model to encompass most of the commonly used distribution functions.

4.1 A negative result

A natural first question is whether the assumption of expected utility maximization generates testable implications if we do not impose any structure on P or U. The following corollary shows that the answer is negative.

Corollary 1. *Let $D = (r^t, b^t)_{t=1}^T$ be a data set. If $b^t < r^t$ for all observations t, then there always exists a cdf P and a utility function U such that D is (P, U)-rationalizable.*

We can show this negative conclusion by using the cdf $P(b) = e^{b-\overline{r}}$, which is a continuous and strictly increasing cdf on $[0, \overline{r}]$. This function satisfies $P(b^t) > 0$ for all t, which makes that condition 1.a of Theorem 1 is satisfied. Thus, to conclude rationalizability of D we only need to verify condition 1.b in Theorem 1. Specifically, it suffices to construct numbers $U^t > 0$ such that, for all t, s,

$$P(b^t)U^t = e^{b^t-\overline{r}}U^t \geq P(r^t - r^s + b^s)U^s = e^{r^t-r^s+b^s-\overline{r}}U^s,$$

We meet this last inequality requirement when specifying $U^t = e^{r^t-b^t} > 0$ for all observations t, as this gives

$$P(b^t)U^t = e^{b^t-\overline{r}}e^{r^t-b^t} = e^{r^t-\overline{r}} = e^{r^t-r^s+b^s-\overline{r}}e^{r^s-b^s} = P(r^t - r^s + b^s)U^s.$$

A crucial aspect of this rationalizability argument is that we have used a cdf P that is log-linear. In such a case, we can always set the utility function U to be equally log-linear on a suitable interval of $[0, \overline{r}]$. Such a combination of P and U rationalizes any data set D, as any choice of $b < r$ gives the same level of expected utility (i.e. $e^{r^t-\overline{r}}$).

In what follows, we will show that we can overcome the negative result in Corollary 1 when imposing strict log-concavity on P or U, thereby also excluding the log-linear specifications. As we will argue, this minimal structure suffices to give specific empirical content to the hypothesis of expected utility maximization.

4.2 Log-concave P or U

We first consider the case with P strictly log-concave. Take any two observations t and s from a data set D. When assuming that the cdf P is known but not the utility function U, condition 1.b of Theorem 1 requires

$$\begin{aligned} &P(b^t)U^t \geq P(r^t - r^s + b^s)U^s, \text{ and} \\ &P(b^s)U^s \geq P(r^s - r^t + b^t)U^t. \end{aligned}$$

For $P(r^t - r^s + b^s) > 0$ and $P(r^s - r^t + b^t) > 0$, we can take the log of both sides to obtain

$$\begin{aligned} &p(r^t - r^s + b^s) - p(b^t) \leq u^t - u^s, \text{ and} \\ &p(r^s - r^t + b^t) - p(b^s) \leq u^s - u^t, \end{aligned}$$

where $p = \ln P$ and $u = \ln U$. Adding up these two conditions gives,

$$p(r^t - r^s + b^s) - p(b^s) \leq p(b^t) - p(r^s - r^t + b^t).$$

Without loss of generality, we can assume $r^t \geq r^s$. Using $\Delta = r^t - r^s \geq 0$, we get

$$p(\Delta + b^s) - p(b^s) \leq p(b^t) - p(b^t - \Delta).$$

Because the cdf P is strictly log-concave, the function p is strictly concave. Then, the above inequality will be satisfied if and only if $\Delta + b^s \geq b^t$ or, equivalently,

$$r^t - b^t \geq r^s - b^s.$$

Thus, strict log-concavity of P requires that, if the rewards r weakly increase (i.e. $r^t \geq r^s$), then the prizes $r - b$ must also weakly increase (i.e. $r^t - b^t \geq r^s - b^s$). In Appendix A.2, we show that this testable implication is not only necessary but also sufficient for rationalizability of the data set D.

We can develop an analogous argument when U is strictly log-concave. In this case, we obtain that a weak increase in the rewards r (i.e. $r^t \geq r^s$) must imply a weak increase in the bids b (i.e. $b^t \geq b^s$). Again, this requirement is both necessary and sufficient for rationalizability. The following theorem summarizes our conclusions.

Theorem 2. *Let $D = (r^t, b^t)_{t=1}^T$ be a data set.*

1. *Let P be a strictly log-concave cdf. Then, there exists a utility function U such that the data set D is (P, U)-rationalizable if and only if,*
   (a) *for all observations $t = 1, \ldots, T$, $P(b^t) > 0$ and $b^t < r^t$, and*
   (b) *for all observations $t, s = 1, \ldots, T$, $r^t \geq r^s$ implies $r^t - r^s \geq b^t - b^s$.*
2. *Let U be a strictly log-concave utility function. Then, there exists a cdf P such that the data set D is (P, U)-rationalizable if and only if,*
   (a) *for all observations $t = 1, \ldots T$, $b^t < r^t$, and*
   (b) *for all observations $t, s = 1, \ldots, T$, $r^t \geq r^s$ implies $b^t \geq b^s$.*

The rationalizability conditions in Theorem 2 have a clear economic interpretation. If P is strictly log-concave, then any increase in the reward r must lead to an increase in the prize $r - b$ that is obtained when winning the lottery. Analogously, if U is strictly log-concave, then any increase in the reward r must lead to an increase in the optimal bid b. More surprisingly, these are the *only* testable implications for (P, U)-rationalizability. They fully exhaust the empirical content of expected utility maximization under the stated observability conditions.

Importantly, the conditions in statement 1 of Theorem 2 are independent of a particular form for the cdf P. In other words, as soon as the data set D is (P, U)-rationalizable by some utility function U for a strictly log-concave cdf P, it is rationalizable for any strictly log-concave P that satisfies $P(b^t) > 0$. This is a clear non-identification result. Apart from the property of strict log-concavity and the fact that the observed bids must lead to strictly positive probabilities, we will never be able to recover any additional property of the function P.

The same non-identification conclusion holds for the rationalizability conditions in statement 2 of Theorem 2. As soon as the data set D is (P, U)-rationalizable for a strictly log-concave utility function U, it is rationalizable for any strictly log-concave utility function U.

4.3 Log-concave P and U

We conclude this section by considering the case where both P and U are assumed to be strictly log-concave. In such a situation, rationalizability requires that the data set D satisfies simultaneously the conditions in statements 1 and 2 of Theorem 2. As we state in the following theorem, this requirement is both necessary and sufficient for (P, U)-rationalizability.

Theorem 3. *Let $D = (r^t, b^t)_{t=1}^T$ be a data set. Let P be a strictly log-concave cdf and let U be a strictly log-concave utility function. Then, the data set D is (P, U)-rationalizable if and only if,*

(a) for all observations $t = 1, \dots, T$, $P(b^t) > 0$ and $b^t < r^t$, and

(b) for all observations $t, s = 1, \dots, T$,

$$r^t \geq r^s \text{ implies } \left(b^t \geq b^s \text{ and } r^t - b^t \geq r^s - b^s\right).$$

Interestingly, this (nonparametric) characterization naturally complies with existing theoretical findings in the (parametric) literature on auctions. In that literature, it is well-established that, when both P and U are strictly log-concave (and satisfy some additional smoothness conditions), the DM's (unique) optimal bid b is increasing in r with a slope less than one (see, for example, Cox and Oaxaca

(1996)). We equally obtain that $r^t \geq r^s$ requires $b^t \geq b^s$. In addition, in our nonparametric setting the slope condition corresponds to $r^t - r^s \geq b^t - b^s$ for $r^t \geq r^s$. From Theorem 3, we conclude that these conditions are not only necessary but also sufficient for rationalizability by a strictly log-concave cdf and strictly log-concave utility function.

5 Illustrative application

We demonstrate the empirical usefulness of our theoretical results through an application to the experimental data set of Neugebauer and Perote (2008) on first-price auctions. This experiment contained 28 subjects. Every subject participated in 50 rounds of 7-player first price auctions. At the beginning of every round, every subject received an i.i.d. value drawn from the uniform distribution on $[0, 100]$, i.e. $\overline{r} = 100$. Subjects received no feedback between the different rounds. The data set contains the observed bidding behavior for each of the 28 subjects.[10] Interestingly, the absence of feedback between the subsequent auctions allows us to exclude various behavioral reasons for overbidding. In particular, there is no winner or loser regret. Moreover, there is no basis for updating subjects' beliefs about the distribution of bids in the population.

5.1 Methodology

Our following analysis will empirically test the alternative model specifications that we discussed above. A main focus will be on the standard expected utility maximization model, which we will refer to as our baseline model. Next, we will also consider three models that impose additional log-concavity assumptions on the cdf P and/or the utility function U. Before entering our empirical analysis, we present a computationally tractable statistical framework to test our baseline model (starting from our general set-up in Section 3). Subsequently, we introduce the techniques that we will use to evaluate our different log-concave models. Finally, we discuss the pass/fail benchmarks that we will use to compare the empirical performance of the various models that are subject to testing.

[10]We actually only consider a subset of Neugebauer and Perote (2008)'s original data set. In their original experiment, the authors also conducted a second session with every subject participating in 50 additional auctions. However, in this session the subjects did receive feedback after every auction round. As the absence of feedback is important to motivate our following tests, we have excluded these data from our application. Further, Neugebauer and Perote's original study contained 28 subjects who played 50 rounds with feedback in a first session, and 50 periods without feedback in a second session. For the same reason as before, we chose not to include these subjects in our empirical application.

Baseline model. For our baseline model, we assume that each participant maximizes her expected utility given the distribution of bids of the other participants. We remark that this implicitly assumes that every subject knows this distribution or, alternatively, has a belief of the distribution that matches the true distribution. This will be relevant when interpreting our empirical results.

The observed bids of the other players can be assumed to be independent, as there is no communication between the subjects and the groups are randomly formed. This makes that the observed distribution of bids of the other participants can be used to estimate the true distribution of these bids. Notice that the observed bids do contain multiple observations from the same individuals. However, this does not violate the i.i.d. assumption as long as we assume the population of players to equal the set of participants in the experiment.[11]

Let G be the cumulative distribution function of the bids of the players (excluding the bids of the DM herself). Then, if the bids are i.i.d., it follows that the probability of winning the auction when offering a bid b equals

$$P(b) = (G(b))^k,$$

where $k + 1$ is the total number of participants in the auction, i.e. there are k other participants in addition to the DM. Using our notation of Section 3, this corresponds to $\Gamma(x) = (x)^k$. Then, **OP.I** becomes

$$Z_m = \min_{e_m, \widehat{G}_{t,s} \in [0,1], U^t > 0} m\, e'_m (\widehat{\Omega})^{-1} e_m,$$

$$\text{s.t. } \forall t, s: \quad e_{t,s} = \mathbb{G}_m(r^t - r^s + b^s) - \hat{G}_{t,s}, \tag{5}$$

$$\left(\hat{G}_{t,t}\right)^k U^t \geq \left(\hat{G}_{t,s}\right)^k U^s, \tag{6}$$

$$\hat{G}_{t,s} < \hat{G}_{t',s'} \text{ for all } r^t - r^s + b^s < r^{t'} - r^{s'} + b^{s'}. \tag{7}$$

Interestingly, we can convert this problem into a minimization problem with linear constraints and a quasi-convex objective function, which is easily solvable through standard algorithms for finding global minima.[12] To see this, we define $\hat{g}_{t,s} =$

[11] If the of players population equals all possible potential participants in the experiment (i.e. also the individuals who did not participate), then the i.i.d. assumption might be violated. In this case, a solution would be to include only one (random) bid for every participant in the auction, which would make the data generating process closer to the i.i.d. assumption. However, for our exercise this would mean that we can use only 27 observations to estimate the distribution of bids, which is too low to rely on asymptotic convergence results. Since all subjects participated to the same session and every participant has seen all possible opponents, we believe that equating the total population with the population of participants is not overly restrictive for our data set.

[12] In particular the minimum can be found by using the sub-gradient projection method (see Kiwiel, 2001). Moreover, it can be well approximated using standard techniques for local minimization (e.g. trust-region methods).

$\ln(\hat{G}_{t,s})$ and $u^t = \ln(U^t)$. Then, we can linearize the constraints (6) as

$$k\,\hat{g}_{t,t} + u^t \geq k\,\hat{g}_{t,s} + u^s,$$

while the constraints (5) yield

$$e_{t,s} = \mathbb{G}_m(x^t - x^s + b^s) - \exp(\hat{g}_{t,s}).$$

Taken together, this obtains the following problem with only linear constraints and a quasi-convex objective function.

$$\begin{aligned} &\textbf{OP.II:} && \min_{g_{t,s}\leq 0} m\, e_m'(\widehat{\Omega})^{-1} e_m, \\ &\text{s.t.} && k\,\hat{g}_{t,t} + u^t \geq k\,\hat{g}_{t,s} + u^s, \\ &&& \hat{g}_{t,s} > \hat{g}_{t',s'} \text{ if } r^t - r^s + b^s > r^{t'} - r^{s'} + b^{s'}, \end{aligned}$$

where we defined $e_{t,s} = \mathbb{G}_m(x^t - x^s + b^s) - \exp(\hat{g}_{t,s})$.

A few remarks are in order. First, instead of using G, an alternative approach could use a direct estimator $\mathbb{P}$ of P which is based on the observed values of highest bids among the other (non-DM) participants. This corresponds to setting Γ equal to the identity function. However, the estimator $\mathbb{P}$ would be based on only a fraction $1/k$ of the number of observations used to estimate $\mathbb{G}$, which makes it less efficient.

Next, in our empirical application, we will split the data per subject into five blocks of 10 periods each. Our motivation to do so is threefold. A first reason relates to the power of our statistical test. The critical values of our test statistic are based on the percentiles of the χ^2 distribution with degrees of freedom proportional to the number of observations squared. As an implication, increasing the number of observations does not necessarily improve the power of our test. For example, by going from 10 to 20 observations we scale up the degrees of freedom from 100 to 400, which leads to a substantial increase in our critical values. Including 10 observations per person appears to define a good compromise between fit and power. Second, and more importantly, restricting the number of observations used in the statistical test is also useful from a computational point of view. Although the optimization problem **OP.II** is well-behaved, the number of constraints is proportional to the square of the number of observations. Finally, an attractive feature of splitting up the sample into blocks of 10 rounds is that it allows us to investigate dynamic patterns in the data. Particularly, as we will explain below, it allows us to study learning behavior versus fatigue.

Summarizing, in our statistical test for each DM we compute $\mathbb{G}_m$ as the empirical distribution function of all bids except from the bids of the DM herself. This gives a total of 27×10 observed bids when using five blocks of 10 rounds per subject. The estimator $\mathbb{G}_m$ is used to solve **OP.II** for each DM individually.

Log-concave models. Our log-concave models assume that the DM maximizes her expected utility for a log-concave cdf P or log-concave utility function U. Testing these models involves two important differences with the baseline model. First, we no longer use the observed bidding behavior of the other (non-DM) participants to reconstruct P. Therefore, we do not need to assume that the DM's belief about the cdf P coincides with the true probability of winning the auction. Next, while our test of the baseline model is statistical in nature, our tests for the log-concave models are exact. These tests do not involve randomness and produce answers that equal either yes or no.

In our empirical application, we will evaluate the data fit of the various log-concave models by using a nonparametric goodness-of-fit measure that is often used for exact tests of the revealed preference type. Specifically, we will make use of the Houtman-Maks index (HMI) (Houtman and Maks, 1985), which measures the relative size of the maximum subset of observations that are consistent with the (exact) revealed preference conditions that are subject to testing. Putting it differently, the HMI identifies the smallest subset of observations that we have to eliminate from the original data set so that the remaining data set satisfies the nonparametric testable implications of the log-concave model that is under evaluation. Intuitively, one may interpret the HMI as an upper bound on the probability that DMs make model-consistent decisions.

Pass/fail benchmark. We recall from our discussion in Section 3 that our statistical test of the baseline model is conservative in nature. This means that it likely under-rejects the null hypothesis. Moreover, the tests of log-concave models are of an exact and non-stochastic nature. This raises the question of which pass/fail benchmark we can use to meaningfully compare the empirical performance of the various models. In our following application, we will construct a benchmark that is based on a calibration exercise. Particularly, it is based on the model of random/irrational behavior that underlies a power measurement procedure proposed by Bronars (1987), which is commonly used in empirical revealed preference analysis.[13]

Following Bronars' procedure, we simulate irrational behavior by generating 500 random data sets $(\tilde{b}^t, r^t)_{t=1}^T$ per DM under evaluation. We do so by randomly drawing (using a uniform distribution) bids $\tilde{b}^t$ between zero and the observed value r^t.[14] We set the benchmark cut-off level for pass/fail as the significance level (for the baseline model) or HMI value (for the log-concave models) such

[13]In Appendix B, we also report results for more standard statistical significance levels (for the baseline model) and alternative HMI pass/fail benchmarks (for the log-concave models).

[14]Because different DMs face different levels of rewards r^t, the 500 random data sets are DM-specific.

that 95% of these random data sets fail the corresponding revealed preference test. By using this procedure for the pass/fail conclusion of the different tests under consideration, we guarantee a test of which the power equals 95% (i.e. the probability that randomly generated data fail the test is 95%).[15]

More specifically, for the baseline model we first solve the optimization problem **OP.II** for each of the 500 random data sets. This gives us 500 objective function values. Then, we rank these values from smallest to largest, and we define as a cut-off value the 10th smallest value (i.e. the 5th percentile) of the resulting distribution. We say that the evaluated DM fails the test of our baseline model if the optimal value of **OP.II** for this subject (based on her observed choices) exceeds this cut-off value. We proceed similarly for the log-concave models. In this case, we compute the HMI for the 500 random data sets. Subsequently, we rank these resulting HMI values from largest to smallest, and we take the 10th largest value as the relevant cut-off value. Analogous to before, we say that the DM under evaluation fails the test of the given log-concave model if the HMI value for her actual choices is below this cut-off value.

5.2 Test Results

We will start by discussing the pass/fail test results for the four model specifications under consideration. Subsequently, we will illustrate that differences between our test results for the baseline model and the log-concave models may be interpreted in terms of expected utility maximization with respect to objective probabilities versus subjective beliefs of winning the lottery. Finally, we will use our test results to study the behavioral phenomena of learning and fatigue at the level of individual DMs.

Pass/fail results. Table 1 presents the pass/fail rates of our four models for each block of 10 periods together with the 95% confidence intervals. The first row in every block presents the percentage of individuals for which we do not reject null hypothesis of rationalizable behavior (using the pass/fail benchmarks that we defined above). The number of DMs that pass the test is given in parenthesis. In the second row of each block, we show 95% confidence intervals for the percentage of individuals passing our tests. These confidence intervals are computed by using the Clopper–Pearson procedure.

Some interesting observations can be drawn from Table 1. First, the pass rates of all models are significantly above 5% for all time periods. Given our specific pass/fail benchmarking procedure (outlined above), 5% would be the expected pass rate if actual behavior were uniformly random. Thus, we may safely conclude that

[15]As a robustness check, Appendix B also reports results for alternative power levels.

the four models subject to testing describe actual behavior better than random behavior. This suggests that the real DMs in our sample behave substantially more rational than the randomly generated subjects. Interestingly, when considering the 95% confidence intervals for our pass rates, we conclude that this difference is also statistically significant in all cases. This shows that our tests are sufficiently powerful to effectively discriminate between rational and random behavior.

Table 1: Pass rates for various models ($N = 28$).

rounds		Baseline	Log-conc P	Log-conc U	Log-conc P and U
$1 \leq t \leq 10$	Pass Rate	32% (9)	29% (8)	89% (25)	68% (19)
	95% CI	(16% – 52%)	(13% – 49%)	(72% – 98%)	(48% – 84%)
$11 \leq t \leq 20$	Pass Rate	50% (14)	39% (11)	89% (25)	82% (23)
	95% CI	(31% – 69%)	(22% – 59%)	(72% – 98%)	(63% – 94%)
$21 \leq t \leq 30$	Pass Rate	71% (20)	50% (14)	96% (27)	89% (25)
	95% CI	(51% – 87%)	(31% – 69%)	(82% – 99%)	()72% – 98%)
$31 \leq t \leq 40$	Pass Rate	39% (11)	57% (16)	93% (26)	82% (23)
	95% CI	(22% – 59%)	(37% – 76%)	(77% – 99%)	(63% – 94%)
$41 \leq t \leq 50$	Pass Rate	61% (17)	64% (18)	100% (28)	82% (23)
	95% CI	(41% – 79%)	(44% – 81%)	–	(63% – 94%)

The cut-off values for pass/fail are computed such that the null hypothesis of rationalizability is rejected for 95% of 500 randomly generated data sets.

Further, we learn from Table 1 that the numbers of subjects consistent with the different model specifications are lowest in the first rounds. The maximum share of subjects for which rationalizability is not rejected equals 71% for the baseline model, 64% for the model with log-concave P model, 100% for the model with log-concave U, and 89% for the model with log-concave P and U. In the absense of feedback between different auction rounds, the observed increase in pass rates over time may suggest a passive learning effect. We will discuss the presence of learning dynamics in more detail below.

Objective versus subjective probabilities. When interpreting the pass rates in Table 1, it is worth remarking that the baseline model and the models with log-concave P correspond to two distinct cases of rational behavior. The baseline model assumes that DMs optimize their expected utility subject to a belief function that equals an objective probability distribution (i.e. the cdf P is the distribution of highest bids submitted by all other players). By contrast, the model with log-concave P in principle allows the probability function P to be very different from the actual distribution of highest bids, as long as it is log-concave. In this case, the distribution P need not correspond to the actual objective probability of winning

Table 2: Consistency between baseline and log-concave models

rounds	Log-conc P	Log-conc U	Log-conc P and U
$1 \leq t \leq 10$	3/5/6	8/17/1	7/10/2
$11 \leq t \leq 20$	5/6/9	13/12/1	12/11/2
$21 \leq t \leq 30$	11/3/9	19/8/1	19/6/1
$31 \leq t \leq 40$	7/9/4	10/16/1	10/13/1
$41 \leq t \leq 50$	11/7/6	17/11/0	14/9/3

In each cell $x/y/z$, x equals the number of subjects that pass the tests of both the log-concave and baseline models, y the number of subjects that pass the test of the log-concave model but not of the baseline model, and z the number of subjects that pass the test of the baseline model but not of the log-concave model.

the lottery and, therefore, can also be interpreted as a subjective belief of winning the lottery. From this perspective, it is interesting to check which subjects are consistent with every model separately. Is it the case that the same DMs pass both the baseline model and the model with log-concave P, or do the models pick up different patterns of behavior (i.e. based on objective versus subjective beliefs)?

Table 2 allows us to address this question: in each cell $x/y/z$, x equals the number of subjects that pass the tests of both the log-concave and baseline models, y the number of subjects that pass the test of the log-concave model but not the baseline model, and z the number of subjects that pass the test of the baseline model but not the log-concave model. From the second column in this table, we can infer that DMs who are consistent with the baseline model are not always consistent with the model with log-concave P, and vice versa. A substantial fraction of subjects are consistent with only one of the two models.

The results for the log-concave U model are less informative as the behavior of most subjects is rationalizable by this model and, therefore, nearly all subjects who are consistent with the baseline model are also consistent with this log-concave model. A similar conclusion applies to the model with log-concave P and U, albeit to a somewhat lesser extent. Nevertheless, we can still make the claim that there are multiple subjects who are consistent with log-concave U or log-concave P and U, but not with our baseline specification. In other words, a considerable fraction of DMs act as expected utility maximizers, while their beliefs about the actions of the other players are far from true. These findings fall in line with the results of Kirchkamp and Reiß (2011), who show that part of the observed deviations from equilibrium play in first price auction is caused by a failure to form correct beliefs.

Table 3: Dynamic patterns at the individual level

Patterns	Baseline	Log-conc P	Log-conc U	Log-conc P and U
Learning	8 (28%)	14 (43%)	23 (82%)	17 (61%)
Learning and fatigue	10 (36%)	2 (7%)	0 (0%)	5 (18%)
Other patterns	10 (36%)	14 (50%)	6 (18%)	6 (21%)

Learning and fatigue. As a final exercise, we consider the dynamics of the pass rates for the alternative model specifications. Particularly, we can study the behavioral phenomena of learning and fatigue. The underlying intuition is that subjects may need some time to learn their optimal bidding behavior but, after a significant amount of repetitions, they may also get tired (or bored) and start to act less rational. Hence, DMs can be inconsistent with rationalizability in the initial blocks, become consistent later (learning), and finally again exhibit inconsistent behavior (fatigue). In our following evaluation, we will exploit the fact that our rationalizability test can be performed at the level of the individual DMs. We can evaluate how these behavioral dynamics affect the individual consistency patterns.

For a given behavioral model, we say that an individual exhibits "learning" behavior if she fails the test in the earlier blocks but passes the test in the later blocks, while she exhibits "learning and fatigue" behavior if she fails the test in the first blocks, passes the test in the middle blocks, and again fails the test in the last blocks. Table 3 shows the number of subjects that conform to each of these patterns. Between 50 and 85 per cent of subjects show either learning or learning and fatigue effects. This evidence of for learning behavior falls in line with the findings reported by Neugebauer and Perote (2008). However, the revealed patterns for the baseline model are quite different from these for the log-concave models. For the baseline model, the fraction of DMs displaying learning behavior is (only) marginally smaller than the faction of DMs displaying learning and fatigue behavior. By contrast, for the log-concave models a large majority of DMs fits the learning pattern, while there is only a very small share of subjects that exhibits learning and fatigue behavior.

6 When rewards are unobserved

So far we have assumed that the rewards r are observed by the empirical analyst. This assumption holds well in an experimental setting, where the exogenous variables are usually under the control of the experimental designer. However, in a real life setting this type of data set is often not available. From this perspective, it is interesting to investigate the usefulness of our above theoretical results in settings where the rewards r are unobserved.

In what follows, we start by showing that the model of expected utility maximization no longer has testable implications in such a case. This conclusion holds even when either the cdf P or the utility function U is perfectly observable. For compactness, we will only provide the argument for P observed and U unobserved, but the reasoning for U observed and P unobserved proceeds analogously. Importantly, however, this non-testability result does not imply that it is impossible to identify bounds on the rewards that are consistent with the observed bids under the assumption of rationalizability. We will show this by discussing the (partially) identifying structure that rationalizable behavior imposes on the unobserved rewards.

A non-testability result. We consider a setting where the empirical analyst only observes a finite number of bids $(b^t)_{t=1}^T$. Further, we assume that the empirical analyst knows the true cdf P but not the utility function U. For simplicity, we assume that $P(b^t) > 0$ for all observations t. If this last condition were violated, the bids would violate condition 1.a in Theorem 1 and, thus, the observed behavior would not be (P, U)-rationalizable. To address the issue of testability, we must characterize a finite collection of rewards $(r^t)_{t=1}^T$ such that the data set $D = (r^t, b^t)_{t=1}^T$ together with P satisfies the rationalizability conditions 1.a and 1.b in Theorem 1.

More formally, we must define $(r^t)_{t=1}^T$ such that $b^t < r^t$ for all t and there exist numbers $U^t > 0$ such that, for all observations s, t,

$$P(b^t)U^t \geq P(r^t - r^s + b^s)U^s.$$

We will show that, for any $(b^t)_{t=1}^T$ and cdf P, we can always specify such a set $(r^t)_{t=1}^T$, which effectively implies non-testability of expected utility maximization. Let $\overline{r}$ be strictly bigger than $\max_{t \in \{1,\ldots,T\}} b^t$, and take any $\Delta > 0$ that satisfies

$$\Delta \in \left]0, \bar{r} - \max_{t \in \{1,\ldots,T\}} b^t\right[.$$

For every observation $t = 1, \ldots, T$, we then consider the value $r^t = b^t + \Delta$, which is contained in $[0, \overline{r}[$. This specification of the rewards ensures $r^t - b^t = \Delta$, i.e. the payoff when winning is the same for each observation t. Furthermore, for all t, s, we let $U^t = U^s = 1$. It then follows that

$$P(b^t)U^t = P(b^t) \text{ and } P(r^t - r^s + b^s)U^s = P(b^t),$$

which implies that the rationalizability condition 1.b in Theorem 1 is trivially satisfied. We thus obtain the following non-testability result.

Corollary 2. *For every data set $D = (b^t)_{t=1}^T$ and cdf P such that $P(b^t) > 0$ for all observations t, there exist values $(r^t)_{t=1}^T$ and a utility function U such that the data set $D' = (r^t, b^t)_{t=1}^T$ is (P, U)-rationalizable.*

Partial identification of rewards. Importantly, the negative conclusion in Corollary 2 does not imply that it is impossible to identify the underlying values r^t that (P, U)-rationalize the observed behavior. Since our characterizations in Theorems 1, 2 and 3 define necessary and sufficient conditions for (P, U)-rationalizability, they can still be used to partially identify the distribution of rewards. This (partially) identifying structure defines the strongest possible (nonparametric) restrictions on the unobserved rewards for the given assumptions regarding U and P.

Let us first consider identification on the basis of Theorem 1. Assuming $P(b^t) > 0$ for all observations, we have for any two observations t and s that the values r^t and r^s providing a (P, U)-rationalization for some U must satisfy the inequality:

$$\frac{P(r^t - r^s + b^s)}{P(b^t)} \frac{P(r^s - r^t + b^t)}{P(b^s)} \leq \frac{U_t}{U_s} \frac{U_s}{U_t} = 1,$$

which puts restrictions on the reward differences $r^t - r^s$. In general, these restrictions will depend on the shape of the cdf P.

This illustrates that, generically, the rewards r^t can only be partially identified, meaning that there are multiple values of $(r^t)_{t=1}^T$ that satisfy the rationalizability restrictions. As an implication, the distribution of rewards cannot be uniquely recovered when only using information on P. This may seem to contradict the vast literature on auction theory that focuses on identifying the distribution of rewards from the distribution of bids (see, for example, Athey and Haile, 2007). However, these existing identification results all rely on additional functional structure that is imposed on the utility function U. By contrast, our result in Theorem 1 is nonparametric in nature, and only assumes that U is strictly increasing.

Next, if the empirical researcher does not know P but assumes that it is strictly log-concave, then we can use statement 1 of Theorem 2 to partially identify the rewards. Specifically, these rewards must satisfy $b^t < r^t$ and, in addition,

$$r^t - r^s \geq 0 \text{ implies } r^t - r^s \geq b^t - b^s.$$

This last statement is equivalent to

$$b^t > b^s \text{ implies } (r^t - r^s \geq b^t - b^s \text{ or } r^t - r^s < 0),$$

which again puts bounds on the reward differences $r^t - r^s$.

Similarly, if U is assumed to be strictly log concave but P is unconstrained, then statement 2 of Theorem 2 imposes $b^t < r^t$ and

$$r^t \geq r^s \text{ implies } b^t \geq b^s.$$

This condition can be rephrased as

$$b^t > b^s \text{ implies } r^t - r^s > 0,$$

which defines restrictions on the sign of $r^t - r^s$.

Finally, if we assume that both P and U are strictly log-concave, then Theorem 3 requires $b^t < r^t$ and

$$r^t \geq r^s \text{ implies } (b^t \geq b^s \text{ and } r^t - b^t \geq r^s - b^s),$$

This is equivalent to

$$b^t > b^s \text{ implies } r^t - r^s \geq b^t - b^s,$$

which once more specifies restrictions on $r^t - r^s$.

We conclude with a simple example that illustrates the application of these identification constraints to retrieve information on latent rewards. Specifically, we assume a data set with four observations (i.e. $T = 4$) containing the bids $b^1 = 1, b^2 = 4, b^3 = 8$ and $b^4 = 10$. Then, if we assume that both P and U are strictly log-concave, (P, U)-rationalizability imposes the restrictions

$$\begin{aligned} &r^1 > 1, \\ &r^2 \geq r^1 + 3, \\ &r^3 \geq r^2 + 4, \\ &r^4 \geq r^3 + 2. \end{aligned}$$

It follows from our argument that any rewards r^1, r^2, r^3 and r^4 satisfying these constraints will provide a (P, U)-rationalization of the observed behavior. This clearly shows the partially informative nature of our nonparametric identification results.

7 Concluding discussion

We provided a nonparametric revealed preference characterization of expected utility maximization in binary lotteries with trade-offs between the final value of the prize and the probability of winning the prize. We have assumed an empirical analyst who observes a finite set of rewards r and bids b for the DM under study.

We started by characterizing optimizing behavior when the empirical analyst also perfectly knows either the probability distribution of winning P (as a function of b) or the DM's utility function U (as a function of $r - b$). Building on this characterization, we next developed a statistical test for the empirical setting where the function U is unknown and the probability P has to be estimated.

In a following step, we considered the case where both functions U and P are fully unknown. For this setting, we first showed that any observed bidding behavior is consistent with expected utility maximization if no further structure is imposed on these unknown functions. However, we also established that imposing log-concavity restrictions does give empirical bite to the hypothesis of expected utility maximization. Specifically, we derived testable implications when either the probability distribution P or the utility function U is assumed to be log-concave. Log-concavity of P imposes that rewards and final prizes should go in the same direction, and log-concavity of U requires that rewards and bids must be co-monotone. Interestingly, these co-monotonicity properties fully exhaust the empirical content of expected utility maximization under the stated log-concavity assumptions.

We demonstrated the practical usefulness of our theoretical results through an application to an experimental data set on first price auctions. We found that the hypothesis of expected utility maximization cannot be rejected for a large fraction of the DMs in our sample. At a more general level, our results suggest that observed bidding behavior is largely consistent with expected utility maximization as long as we control for possible learning effects and do not impose (non-verifiable and potentially implausible) parametric assumptions on DM preferences. Moreover, some DMs appear to be successfully best responding to the actual distribution of bids in the population, while other DMs can be modeled as expected utility maximizers with respect to a subjective belief of this distribution.

Finally, while our main focus was on testing expected utility maximization when both rewards r and bids b are observed, we have also considered the use of our results in the case where the rewards are no longer observed (which is often relevant in non-experimental settings). On the negative side, we have shown that expected utility maximization is no longer testable in such a case, even if P or U is fully known. On the positive side, we have demonstrated that our characterizations do impose partially identifying structure on the rewards r that can rationalize the observed behavior in terms of expected utility maximization.

An interesting avenue for future research consists of extending our characterizations by relaxing or strengthening our assumptions regarding individual preferences. For example, we may consider the weaker assumption that the DM's utility function U is continuous and satisfies first-order stochastic dominance. More formally, this means that the individual's utility is increasing in both the final prize

$r - b$ and the probability of winning $P(b)$. This extension would lead to a test of expected utility maximization for preferences that (only) satisfy first order stochastic dominance.

Next, follow-up research may fruitfully focus on extending our results towards a broader range of decision problems characterized by prize-probability trade-offs. For instance, an interesting alternative application concerns contest or all-pay auctions. The key difference between this setting and our current set-up is that the DM has to pay the bid even if she loses the auction. Thus, increasing the probability of winning decreases not only the DM's potential prize but also her payoff when she does not get the prize. Another possible application pertains to the double-auction bilateral trade mechanism. This mechanism differs from the posted price model presented in Section 2 in that the seller and the buyer simultaneously post a price. Trade occurs at the average of these two prices if the seller's price does not exceed the buyer's price, while there is no trade otherwise. Once more, the DMs face a clear prize-probability trade-off as posting a higher/lower price increases the probability of trade for the buyer/seller. However, a main difference with our set-up is that the potential prize becomes stochastic, as it depends on the (randomly) posted price of the other party.

References

Abrams, E., Sefton, M., Yavas, A., 2000. An experimental comparison of two search models. Economic Theory 16 (3), 735–749.

Afriat, S. N., 1967. The construction of utility functions from expenditure data. International Economic Review 8, 67–77.

Athey, S., Haile, P. A., 2007. Handbook of Econometrics. Vol. 6. Elsevier, Ch. Nonparametric approaches to auctions, pp. 3847–3965.

Bagnoli, M., Bergstrom, T., 2005. Log-concave probability and its applications. Economic Theory 26, 445–569.

Borgers, T., Li, J., forthcoming. Strategically simple mechanisms. Econometrica.

Bronars, S. G., 1987. The power of nonparametric tests of preference maximization. Econometrica 55, 693–698.

Cason, T. N., Friedman, D., Milam, G. H., 2003. Bargaining versus posted price competition in customer markets. International Journal of Industrial Organization 21 (2), 223–251.

Cox, J., Oaxaca, R. L., 1996. Is bidding behavior consistent with bidding theory for private value auctions? In: Research in Experimental Economics. Vol. 6. JAI Press, pp. 131–148.

Diewert, W. E., 1973. Afriat and revealed preference theory. Review of Economic Studies 40, 419–425.

Echenique, F., Saito, K., 2015. Savage in the market. Econometrica 83, 1467–1495.

Green, E. J., Osband, K., 1991. A revealed preference theory for expected utility. Review of Economic Studies 58, 677–695.

Houtman, M., Maks, J. A. H., 1985. Determining all maximal data subsets consistent with revealed preference. Kwantitatieve Methoden 19, 89–104.

Kagel, J. H., Levin, D., 2011. Auctions: a survey of experimental research, 1995–2010. Handbook of experimental economics 2.

Kim, T., 1996. Revealed preference theory on the choice of lotteries. Journal of Mathematical Economics 26, 463–477.

Kirchkamp, O., Reiß, J. P., 2011. Out of equilibrium bids in first price auctions: wrong expectaions of wrong bids. Economic Journal 121, 1361–1397.

Kiwiel, K. C., 2001. Convergence and efficiency of subgradient methods for quasiconvex minimization. Mathematical Programming 90, 1–25.

Manski, C. F., 2002. Identification of decision rules in experiments on simple games of proposal and response. European Economic Review 46 (4-5), 880–891.

Manski, C. F., 2004. Measuring expectations. Econometrica 72 (5), 1329–1376.

Matzkin, R. L., Richter, M. K., 1991. Testing strictly concave rationality. Journal of Economic Theory 53, 287–303.

Neugebauer, T., Perote, J., 2008. Bidding 'as if' risk neutral in experimental first price auctions without information feedback. Experimental Economics 11, 190–202.

Sepanski, S. J., 1994. Asymptotics for multivariate t-statistics and hotellings t^2-statistic under infinite second moments via bootstrapping. Journal of Multivariate Analysis 49, 41–54.

Tabarrok, A., 1998. The private provision of public goods via diminant assurance contracts. Public Choice 96, 345–362.

Varian, H. R., 1982. The nonparametric approach to demand analysis. Econometrica 50, 945–974.

Varian, H. R., 1983. Non-parametric tests of consumer behavior. The Review of Economic Studies 50, 99–110.

Zubrickas, R., 2014. The provision point mechanism with refund bonuses. Journal of Public Economics 120, 231–234.

A Proofs

A.1 Proof of Theorem 1

Statement 1: P is known but U is not.
($\Rightarrow$) Let $D = (r^t, b^t)_{t=1}^T$ be (P,U)-rationalizable. Let us first derive condition 1.a. Given that P is strictly increasing on $[0, \overline{r}]$, $P(b^t)$ can only be zero if $b^t = 0$. Then, the expected utility of choosing $b^t = 0$ is given by

$$P(0)U(r^t) = 0.$$

Notice that, as U is strictly increasing and $U(0) = 0$, we have that $U(r^t) > 0$. Given continuity of U and the fact that P is strictly increasing, there must exist a $\varepsilon > 0$ such that $P(\varepsilon) > 0$ and $U(r^t - \varepsilon) > 0$. As such,

$$P(\varepsilon)U(r^t - \varepsilon) > 0,$$

which means that $b^t = 0$ can never be an optimal choice.

Next, if $b^t = r^t$, and consequentially $U(r^t - b^t) = U(0) = 0$, we have that

$$P(r^t)U(0) = 0.$$

Notice that $P(r^t) > 0$ as $r^t > 0$. Given continuity of P and the fact that U is strictly increasing, there must exist a ε such that,

$$P(r^t - \varepsilon)U(\varepsilon) > 0.$$

Again this implies that $b^t = r^t$ can never be an optimal bid.

Finally, to derive condition 1.b, let $U^t = U(r^t - b^t) > 0$. Then, by optimality of b^t, we have that

$$\begin{aligned} P(b^t)U^t &= P(b^t)U(r^t - b^t), \\ &\geq P(r^t - (r^s - b^s))U(r^t - (r^t - r^s + b^s)), \\ &= P(r^t - r^s + b^s)U^s, \end{aligned}$$

which is exactly condition 1.b.

($\Leftarrow$) To prove sufficiency, we construct a regular Bernoulli utility function $U : \mathbb{R} \to \mathbb{R}$ that rationalizes the data set. Define

$$U(x) = \min\left\{\alpha x, \min_{t=1,\dots,T}\left\{U^t \frac{P(b^t)}{P(r^t - x)} \text{ s.t. } P(r^t - x) > 0\right\}\right\}, \tag{8}$$

where we choose

$$\alpha > \max_t \frac{U^t}{r^t - b^t}. \tag{9}$$

Notice that $U(x)$ is well-defined (i.e. finite valued), continuous and strictly increasing as it is the minimum of a finite number of strictly increasing, continuous functions. Also, for all observations t,

$$0 < U^t \frac{P(b^t)}{P(r^t)},$$

which follows from the fact that $P(b^t) > 0$, strict monotonicity of P and $U^t > 0$. As such, we have $U(0) = \alpha 0 = 0$.

Next, for all t we have $U(r^t - b^t) = U^t$. Indeed, from the definition, we immediately obtain the inequality $U(r^t - b^t) \leq U^t$ and, by assumption (9), we have $U^t < \alpha(r^t - b^t)$. If the inequality would be strict, i.e. $U(r^t - b^t) < U^t$, then there must be an observation s such that

$$U^s \frac{P(b^s)}{P(r^s - r^t + b^t)} < U^t.$$

This, however, contradicts condition 1.b.

Finally, let us show that the data set $D = (r^t, b^t)_{t=1}^T$ is (P, U)-rationalizable by the function $U(x)$ defined in (8). Consider any $b \in [0, r^t]$, then we have

$$P(b)U(r^t - b) \leq P(b)U^t \frac{P(b^t)}{P(b)} = P(b^t)U^t.$$

Statement 2: U is known but P is not.

($\Rightarrow$) Let $D = (r^t, b^t)_{t=1}^T$ be (P, U)-rationalizable. As in our proof of statement 1, we can show that $b^t < r^t$ for all t, which obtains condition 2.a. To derive condition 2.b, let us set $P^t = P(b^t)$. As in our proof of statement 1, we can show that $P^t > 0$. Then, choosing b^t should provide at least as much utility as choosing b^s. As such,

$$P^t U(r^t - b^t) = P(b^t)U(r^t - b^t) \geq P(b^s)U(r^t - b^s) = P^s U(r^t - b^s),$$

which obtains condition 2.b.

($\Leftarrow$) To prove sufficiency, we need to construct a cdf P. Define the function

$$V(b) = \min\left\{\alpha b, \min_{t=1,\ldots,T}\left\{P^t \frac{U(r^t - b^t)}{U(r^t - b)} \text{ s.t. } r^t > b\right\}\right\}, \tag{10}$$

where we choose

$$\alpha > \max_t \frac{P^t}{b}. \tag{11}$$

Notice that $V(b)$ is well-defined (i.e. finite valued), non-negative, continuous, and strictly increasing as it is the minimum of a finite number of strictly increasing, continuous functions. Given this, define

$$P(b) = \frac{V(b)}{V(\overline{r})},$$

which obtains that P is a cdf on $[0, \overline{r}]$.

Next, for all t we have $V(b^t) = P^t$. Indeed, as $r^t > b^t$, we have that $V(b^t) \leq P^t$. If the inequality is strict, then $P^t < \alpha b^t$ (by condition (11)) implies that there is an observation s such that

$$V(b^t) = P^s \frac{U(r^s - b^s)}{U(r^s - b^t)} < P^t.$$

This, however, contradicts condition 2.b.

Let us finish the proof by showing that the data set $D = (r^t, b^t)_{t=1}^T$ is (P, U)-rationalizable. If not, then there is a $b \in [0, r^t]$ such that

$$P(b)U(r^t - b) > P(b^t)U(r^t - b^t) = \frac{P^t}{V(\overline{r})}U(r^t - b^t).$$

This inequality requires that $U(r^t - b) > 0$, which implies that $b < r^t$. As such, $V(b) \leq P^t \frac{U(r^t - b^t)}{U(r^t - b)}$. Given this,

$$P(b)U(r^t - b) \leq \frac{P^t}{V(\overline{r})} \frac{U(r^t - b^t)}{U(r^t - b)} U(r^t - b) = \frac{P^t}{V(\overline{r})} U(r^t - b^t),$$

a contradiction.

A.2 Proof of Theorem 2

Please see Appendix A.4 for some network concepts and Appendix A.5 for the proofs of the lemmata.

Statement 1: P is strictly log-concave.

Lemma 2 shows that there exists a utility function U such that the data set D is (P, U)-rationalizable if and only if, for all cycles C on $G = (T, T \times T)$, which satisfy $P(r^t - r^{t+} + b^{t+}) > 0$ for all for all nodes t, we have

$$\sum_{t \in C} p(r^t - r^{t+} + b^{t+}) - p(b^{t+}) \leq 0, \tag{12}$$

with $p(x) = \ln(P(x))$. We will show that this condition is satisfied if and only if for all observations t, s, $r^s \geq r^t$ implies $r^s - b^s \geq r^t - b^t$.

($\Rightarrow$) Consider two observations t and s. If $P(r^t - r^s + b^s) = 0$, then it must be that $r^t - r^s + b^s \leq 0$, since P is strictly increasing. In particular,

$$r^t \leq r^s - b^s.$$

As $b^s \geq 0$, this implies $r^t \leq r^s$ and also $r^t - b^t \leq r^s - b^s$. Similarly, if $P(r^s - r^t + b^s) = 0$, we obtain $r^s \leq r^t$ and $r^s - b^s \leq r^t - b^t$. So the result holds for both these cases.

Next, consider the case where both $P(r^t - r^s + b^s) > 0$ and $P(r^s - r^t + b^t) > 0$. Without loss of generality, assume that $r^s \geq r^t$. Then, given the cycle $C = \{(t, s), (s, t)\}$, we must have (by (12))

$$\begin{aligned} &p(r^t - r^s + b^s) - p(b^s) + p(r^s - r^t + b^t) - p(b^t) \leq 0 \\ \Leftrightarrow &p(r^s - r^t + b^t) - p(b^t) \leq p(b^s) - p(b^s - (r^s - r^t)). \end{aligned}$$

Given strict concavity of p, this can only hold if $r^s - r^t + b^t \geq b^s$ or, equivalently, $r^s - b^s \geq r^t - b^t$, as we needed to show.

($\Leftarrow$) We work by induction on the length of the cycle C in order to show that condition (12) is satisfied. If C has length 2, the proof is similar to the necessity part above. Let us assume that the condition holds for all cycles up to length $n-1$ and consider a cycle of length n. Let t be the node of the cycle with the lowest value of r^t. Denote by C' the cycle where the edges (t^-, t) and (t, t^+) are removed and the edge (t^-, t^+) is added. Using this notation we have,

$$\begin{aligned} \sum_{s \in C} p(r^s - r^{s+} + b^{s+}) - p(b^{s+}) = &\sum_{s \in C'} \big(p(r^s - r^{s+} + b^{s+}) - p(b^{s+})\big) \\ &+ p(r^{t-} - r^t + b^t) - p(b^t) + p(r^t - r^{t+} + b^{t+}) \\ &- p(r^{t-} - r^{t+} + b^{t+}). \end{aligned} \tag{13}$$

Notice that $P(r^t - r^{t+} + b^{t+})$ being strictly positive implies also that $P(r^{t-} - r^{t+} + b^{t+}) > 0$ since $r^{t-} \geq r^t$. As such we can indeed take the logarithm.

The first term on the right hand side of (13) is negative by the induction hypothesis. As such, it suffices to show that,

$$p(r^{t-} - r^t + b^t) - p(b^t) \leq p(r^{t-} - r^{t+} + b^{t+}) - p(r^t - r^{t+} + b^{t+}). \tag{14}$$

Define $\Delta = r^{t-} - r^t \geq 0$ and set $r^{t-} - r^{t+} + b^{t+} = \tilde{b} \geq 0$. Then, substituting into (14) gives,

$$p(\Delta + b^t) - p(b^t) \leq p(\tilde{b}) - p(\tilde{b} - \Delta).$$

As p is strictly concave and strictly increasing, this holds whenever

$$\begin{aligned} &\Delta + b^t \geq \tilde{b} \\ \Leftrightarrow &r^{t-} - r^t + b^t \geq r^{t-} - r^{t+} + b^{t+} \\ \Leftrightarrow &r^{t+} - b^{t+} \geq r^t - b^t. \end{aligned}$$

This is indeed the case, as $r^{t+} \geq r^t$.

Statement 2: U is strictly log-concave.
This proof is readily analogous to the proof of statement 1.

A.3 Proof of Theorem 3

($\Rightarrow$) First, notice that, by continuity and monotonicity of P and U, we have that $P(b^t) > 0$ and $U(r^t - b^t) > 0$. As such, the choice b^t also optimizes the log of $P(b)U(r^t - b)$, denoted by $p(b) + u(r^t - b)$. This objective function is strictly concave, so a solution has to satisfy the first order condition,

$$\partial p^t - \partial u^t = 0,$$

where ∂p^t is a suitable supergradient of $p(b^t)$ and ∂u^t is a suitable supergradient of $u(r^t - b^t)$, and where we use that $0 < b^t < \bar{r}$. Then, strict concavity of u and p gives,

$$p(b^t) - p(b^s) \leq \partial p^s(b^t - b^s) = \partial u^s(b^t - b^s), \tag{15}$$
$$u(r^t - b^t) - u(r^s - b^s) \leq \partial u^s(r^t - b^t - (r^s - b^s)), \tag{16}$$

where the inequality (15) is strict if $b^s \neq b^t$ and the inequality (16) is strict if $r^t - b^t \neq r^s - b^s$. If we exchange t and s in conditions (15) and (16) and add them together, we obtain,

$$0 \leq (\partial u^s - \partial u^t)(b^t - b^s), \tag{17}$$
$$0 \leq (\partial u^s - \partial u^t)(r^t - b^t - (r^s - b^s)), \tag{18}$$

where (17) is strict if $b^t \neq b^s$ and (18) is strict if $r^t - b^t \neq r^s - b^s$. If $b^t > b^s$, then, for (17) to hold, we must have that $\partial u^t < \partial u^s$, which implies we need in turn that $r^t - r^s \geq b^t - b^s$ to satisfy (18). As such, we obtain that $r^s \geq r^t$ implies $b^s \geq b^t$.

Next, if $r^t - b^t > r^s - b^s$, then for (18) to hold, we must have that $\partial u^t < \partial u^s$, which implies we need in turn that $b^t \geq b^s$ to satisfy (17). As such we obtain $r^t - r^s > b^t - b^s \geq 0$ and thus also $r^t > r^s$. Again we can therefore conclude that $r^s \geq r^t$ implies $r^s - b^s \geq r^t - b^t$.

($\Leftarrow$) Taking the contraposition, we have that $b^t > b^s$ implies $r^t > r^s$ and $r^t - b^t > r^s - b^s$ implies $r^t > r^s$. Then, by combining Lemmata 3 and 5 we have that there are numbers u^t and p^t such that, for all observations t, s,

$$u^t - u^s \leq r^s(r^t - b^t - (r^s - b^s)), \tag{19}$$
$$p^t - p^s \leq r^s(b^t - b^s), \tag{20}$$

where the inequality (19) is strict if $r^t - b^t \neq r^s - b^s$, and the inequality (20) is strict if $b^t \neq b^s$. As shown in Matzkin and Richter (1991), these inequalities imply the existence of continuous, strictly increasing and strictly concave functions $\tilde{u}$ and p such that, for all t,

$$\tilde{u}(r^t - b^t) = u^t, \text{ and } p(b^t) = p^t.$$

and r^t is a supergradient of $u(r^t - b^t)$ and $p(b^t)$. Define the function

$$u(x) = \min\{\ln(\alpha x), \tilde{u}(x)\},$$

where we choose $\alpha > 0$ such that, for all t,

$$\ln(\alpha(r^t - b^t)) > u^t.$$

The function $u(x)$ is still strictly concave, strictly monotone and continuous. In addition, for all t we have that $u(r^t - b^t) = u^t$ and r^t is a supergradient of $u(r^t - b^t)$, but now we also have that $\lim_{x\to 0} \tilde{u}(x) = -\infty$. Define

$$U(x) = \exp(\tilde{u}(x)),$$

and

$$P(b) = \exp\left(p(b) - p(\overline{r})\right).$$

Then, U is strictly increasing, strictly log-concave and $U(0) = 0$ and P is between 0 and 1, strictly increasing and strictly log-concave on $[0, \overline{r}]$.

For these definitions of U and P, let us show that the data set $D = (r^t, b^t)_{t=1}^T$ is (P, U)-rationalizable. That is, that t b^t maximizes $p(b) + \tilde{u}(r - b)$. We know that $P(b^t)U(r^t - b^t) > 0$, so we only need to consider values $b < r^t$ with $P(b) > 0$. By concavity of p and u we have, for all such b,

$$p(b) + \tilde{u}(r^t - b) - \big(p(b^t) - \tilde{u}(r^t - b^t)\big) \leq r^t(b - b^t) + r^t(r^t - b - (r^t - b^t)) = 0,$$

as we needed to show.

A.4 Some notation and definitions on networks

A directed network $G = (T, E)$ consists of a finite set of nodes T and edges $E \subseteq T \times T$. An edge $e \in E$ is called an incoming edge for the node t if $e = (s, t)$ for some $s \in T$ and it is called an outgoing edge if $e = (t, s)$ for some $s \in T$. Two nodes t, s are connected if there is a sequence of edges

$$e_1 = (t, n_1), e_2 = (n_1, n_2), \dots e_k = (n_{k-1}, s),$$

connecting t to s. We also call $e_1, \dots, e_k$ a path from t to s.

A cycle $C = (e_1, \dots, e_k)$ on the network G consists of a collection of edges such that

$$e_1 = (n_1, n_2), e_2 = (n_2, n_3), \dots e_k = (n_k, n_1).$$

We call $\{n_1, \dots, n_k\}$ the nodes of the cycle and k the length of the cycle. For a node n_i in the cycle , n_{i+1} is called the successor of n_i if $i < k$ and n_1 if $i = k$. Similarly, n_{i-1} is called the predecessor of n_i if $i > 1$ and n_k if $i = 1$. We denote the successor as n_i+ and the predecessor as n_i-.

A.5 Lemmata

Lemma 1. *Let P be a cdf and let $D = (r^t, b^t)_{t=1}^T$ be a data set such that $P(b^t) > 0$ and $b^t < r^t$ for all t. Then, there exists a utility function U such that D is (P, U)-rationalizable if and only if, for all t, there exists numbers u_t such that, for all t, s with $P(r^t - r^s + b^s) > 0$,*

$$p(r^t - r^s + b^s) - p(b^t) \leq u^t - u^s,$$

where $p(x) = \ln(P(x))$.

Proof. $(\Rightarrow)$ Let D be (P, U)-rationalizable. Then, from condition 1.b in Theorem 1 we know there exist number $U^t > 0$ such that, for all t, s,

$$P(b^t)U^t \geq P(r^t - r^s + b^s)U^s.$$

If $P(r^t - r^s + b^s) > 0$ we can take logs on both sides, which gives

$$p(r^t - r^s + b^s) - p(b^t) \leq u^t - u^s,$$

as we wanted to show.
$(\Leftarrow)$Assume that there are numbers u_t such that, for all t, s with $P(r^t - r^s + b^s) > 0$,

$$p(r^t - r^s + b^s) - p(b^t) \leq u^t - u^s.$$

Taking exponents on both sides gives $P(r^t - r^s + b^s)U^s \leq P(b^t)U^t$ shows that condition 1.b of Theorem 1 holds in the case where $P(r^t - r^s + b^s) > 0$. For the case where $P(r^t - r^s + b^s) = 0$ then condition 1.b is always satisfied as the left hand side is then equal to zero. Applying Theorem 1 shows that there exists a utility function U such that D is (P, U)-rationalizable. $\square$

Lemma 2. *Let P be a cdf and let $D = (r^t, b^t)_{t=1}^T$ be a data set such that $P(b^t) > 0$ and $b^t < r^t$ for all t. Then, there exists a utility function U such that D is (P, U)-rationalizable if and only if, for all cycles C on the network $G = (T, T \times T)$, which satisfy $P(r^t - r^{t+} + b^{t+}) > 0$ for all nodes t, we have*

$$\sum_{t \in C} p(r^t - r^{t+} + b^{t+}) - p(b^{t+}) \leq 0.$$

Proof. $(\Rightarrow)$ From Lemma 1 we have that there are numbers u^t such that, for all nodes t of C,

$$p(r^t - r^{t+} + b^{t+}) - p(b^t) \leq u^t - u^{t+}.$$

Now, summing the left and right hand sides over all nodes t of the cycle C gives

$$0 \geq \sum_{t \in C} \left(p(r^t - r^{t+} + b^{t+}) - p(b^t)\right) = \sum_{t \in C} \left(p(r^t - r^{t+} + b^{t+}) - p(b^{t+})\right).$$

$(\Leftarrow)$ Assume m is the node with the highest value r^m. It follows that, for all nodes t,

$$r^m - r^t + b^t > 0,$$

so by strict monotonicity of P, $P(r^m - r^t + b^t) > 0$. Let E be the set of edges (t, s) such that $p(r^t - r^s + b^s) > 0$. Let $\mathcal{P}_t$ be the set of all paths on the graph $G'(N, E)$ that start at m and end at t. Notice that $\mathcal{P}_m$ includes the path (m, m). Given that $p(r^m - r^t + b^t) > 0$ exists for all nodes t, the set $\mathcal{P}_t$ is non-empty. Now define, for all t,

$$u^t = \min_{\pi \in \mathcal{P}_t} \sum_{(s, s^+) \in \pi} p(b^s) - p(r^s - r^{s+} + b^{s+}).$$

Because of the cyclical monotonicity condition, the solution set to the minimization problem will always contain a path that does not contain any cycles, as including a cycle makes the right hand side only larger. This shows that the minimum is bounded from below and, therefore, the value u^t is well-defined.

Also, if $P(r^t - r^s + b^s) > 0$ then, for any path in $\mathcal{P}_t$, we can create a path in $\mathcal{P}_s$ by adding the edge (t, s). Therefore, for all s, t,

$$u^s \leq u^t + p(b^t) - p(r^t - r^s + b^s).$$

Using Lemma 1, we can conclude that the data set D is (P, U)-rationalizable for some utility function U. $\square$

Lemma 3. *Let $(z^t, y^t)_{t=1}^T$ be a collection of numbers $z^t, y^t \in \mathbb{R}$. Then, the following statements are equivalent:*

1. *For all cycles C in $G = (T, T \times T)$ where the values y^t are not equal over all nodes t in C, we have that*
$$\sum_{t \in C} z^t (y^{t+} - y^t) > 0.$$
2. *For all t, s we have that*
$$y^t > y^s \Rightarrow z^t < z^s.$$

Proof. ($1 \Rightarrow 2$) Suppose the condition in statement 1 holds. Then, given a cycle $C = \{(t,s),(s,t)\}$ we have that, if $y^t \neq y^s$,

$$\begin{aligned} &z^t(y^s - y^t) + z^s(y^t - y^s) > 0, \\ \Leftrightarrow &(z^s - z^t)(y^t - y^s) > 0. \end{aligned}$$

As such, $y^t > y^s$ implies $z^t < z^s$, as we wanted to show.
($2 \Rightarrow 1$) We use induction on the length of the cycle C. For a cycle of length 2 the proof is similar to the first part of the proof. Assume that the equivalence holds for all cycles up to length $n-1$ and consider a cycle C of length n. If the cycle $C = \{(t_1,t_2),(t_2,t_3),\ldots(t_n,t_1)\}$ contains two nodes t_i, t_j ($i < j$) where $y^{t_i} = y^{t_j}$, then we can break up C into two cycles of smaller length. In particular, we have the cycles,

$$\begin{aligned} &C_1 = \{(t_1,t_2),\ldots,(t_{i-2},t_{i-1}),(t_{i-1},t_j),(t_j,t_{j+1}),\ldots,(t_n,t_1)\}, \text{ and} \\ &C_2 = \{(t_i,t_{i+1}),(t_{i+1},t_{i+2}),\ldots(t_{j-2},t_{j-1}),(t_{j-1},t_i)\}. \end{aligned}$$

Also, as $y^{t_i} = y^{t_j}$ we have,

$$\sum_{t \in C} z^t(y^{t+} - y^t) = \sum_{t \in C_1} z^t(y^{t+} - y^t) + \sum_{t \in C_2} z^t(y^{t+} - y^t).$$

By the induction hypothesis, the sum on the right hand side is greater than 0, so the sum on the left is then also greater than 0.

Next, we consider the case where there is a cycle C of length n and where, for all nodes $t, s \in C$, $y^t \neq y^s$. Let t be the node in C with the smallest value of y^t, and let C' be the cycle obtained from C by removing the edges (t^-, t), (t, t^+) and adding the edge (t^-, t^+). Then,

$$\begin{aligned} \sum_{s \in C} z^s(y^{s+} - y^s) = &\sum_{s \in C'} z^s(y^{s+} - y^s), \\ &+ z^{t-}(y^t - y^{t-}) + z^t(y^{t+} - y^t) - z^{t-}(y^{t+} - y^{t-}). \end{aligned}$$

The first expression on the right hand side is strictly greater than zero by the induction hypothesis. As such, it suffices to show that,

$$\begin{aligned}
&z^{t-}(y^t - y^{t-}) + z^t(y^{t+} - y^t) - z^{t-}(y^{t+} - y^{t-}) \geq 0 \\
\Leftrightarrow &z^{t-}(y^t - y^{t-}) + z^t(y^{t+} - y^t) - z^{t-}(y^{t+} - y^t) - z^{t-}(y^t - y^{t-}) \geq 0 \\
\Leftrightarrow &(z^t - z^{t-})(y^{t+} - y^t) \geq 0.
\end{aligned}$$

By assumption, we have $y^{t+} > y^t$, so the second part of the product is strictly positive. In addition, we have $y^{t-} > y^t$ so $z^{t-} < z^t$ by statement 2 of the lemma, which shows that the first part of the product is also strictly positive. $\square$

Lemma 4. *Let $(z^t, y^t)_{t=1}^T$ be a collection of numbers $z^t, y^t \in \mathbb{R}$ and let C be a cycle in $G = (T, T \times T)$. Then, there exists a collection of cycles $\mathcal{C}$ such that,*

1. *For all $\tilde{C} \in \mathcal{C}$ and all nodes $t, s \in \tilde{C}$ we have $y^t \neq y^s$,*
2. *$\sum_{s \in C} z^s(y^{s+} - y^s) = \sum_{\tilde{C} \in \mathcal{C}} \sum_{s \in \tilde{C}} z^s(y^{s+} - y^s)$,*
3. *$\sum_{s \in C} 1[y^s \neq y^{s+}] = \sum_{\tilde{C} \in \mathcal{C}} \sum_{s \in \tilde{C}} 1[y^s \neq y^{s+}]$.*

Proof. Consider a cycle C in $G = (T, T \times T)$. We will build the collection $\mathcal{C}$ in two steps. First, we remove from C all edges (t, s) where $y^t = y^s$. In order to do this, if C contains an edge (t, s) where $y^t = y^s$ we construct a new cycle C by deleting the edges (t^-, t) and (t, s) and adding the edge (t^-, s). The resulting cycle C' has the feature that,

$$\sum_{s \in C} z^s(y^{s+} - y^s) = \sum_{s \in C'} z^s(y^{s+} - y^s),$$

and,

$$\sum_{s \in C} 1[y^{s+} \neq y^s] = \sum_{s \in C'} 1[y^{s+} \neq y^s].$$

This process can be repeated to finally arrive at a cycle $\tilde{C}$ such that, for any edge (t, s) we have $y^t \neq y^s$ together with,

$$\sum_{s \in C} z^s(y^{s+} - y^s) = \sum_{s \in \tilde{C}} z^s(y^{s+} - y^s),$$

and,

$$\sum_{s \in C} 1[y^{s+} \neq y^s] = \sum_{s \in \tilde{C}} 1[y^{s+} \neq y^s].$$

We take $\tilde{C}$ as a starting point of the second step. If $\tilde{C}$ contains no two nodes t and s (not connected by an edge) such that $y^t = y^s$, then we have constructed $\mathcal{C} = \{\tilde{C}\}$. Else, let $\tilde{C} = \{(t_1, t_2), \ldots, (t_n, t_1)\}$ be such that, for at least two nodes

t_i, t_j ($i < j$) in C, we have $y^{t_i} = y^{t_j}$. We decompose $\tilde{C}$ into two new cycles $\tilde{C}_1$ and $\tilde{C}_2$, in the following way:

$$\tilde{C}_1 = \{(t_1, t_2), \dots (t_{i-2}, t_{i-1}), (t_{i-1}, t_j), (t_j, t_{j+1}), \dots, (t_n, t_1)\} \text{ and}$$
$$\tilde{C}_2 = \{(t_i, t_{i+1}), \dots (t_{j-1}, t_j)\}.$$

Notice that $\tilde{C}_1$ and $\tilde{C}_2$ inherit from $\tilde{C}$ the property that $y^{t_k} = y^{t_l}$ for all edges (t_k, t_l) together with

$$\sum_{s \in \tilde{C}} z^s(y^{s+} - y^s) = \sum_{s \in \tilde{C}_1} z^s(y^{s+} - y^s) + \sum_{s \in \tilde{C}_2} z^s(y^{s+} - y^s),$$

and

$$\sum_{s \in \tilde{C}} 1[y^{s+} \neq y^s] = \sum_{s \in \tilde{C}_1} 1[y^{s+} \neq y^s] + \sum_{s \in \tilde{C}_2} 1[y^{s+} \neq y^s].$$

Again we can repeat this process to obtain a collection $\mathcal{C}$ of cycles such that, for all nodes $t_i, t_j \in \tilde{C} \in \mathcal{C}$, we have $y^{t_i} \neq y^{t_j}$. Moreover,

$$\sum_{(t,t+) \in C} z^s(y^{s+} - y^s) = \sum_{\tilde{C} \in \mathcal{C}} \sum_{(t,t+) \in C} z^s(y^{s+} - y^s),$$

and,

$$\sum_{(t,t+) \in C} 1[y^t \neq y^{t+}] = \sum_{\tilde{C} \in \mathcal{C}} \sum_{(t,t+) \in \tilde{C}} 1[y^t \neq y^{t+}],$$

which we wanted to show. □

Lemma 5. *Let $(z^t, y^t)_{t=1}^T$ be a collection of numbers such that $z^t, y^t \in \mathbb{R}$. Then, the following statements are equivalent.*

1. *For all cycles C in $G = (T, T \times T)$ where the values y^t are not all equal over the nodes t of C, we have that*

$$\sum_{t \in C} z^t(y^{t+} - y^t) > 0.$$

2. *There exist numbers u^t such that, for all t, s,*

$$u^t - u^s \leq z^s(y^t - y^s),$$

with a strict inequality if $y^t \neq y^s$.

Proof. ($2 \Rightarrow 1$) This is easily obtained by summing the inequality in statement 2 over all edges (t, t^+) of the cycle C.
($1 \Rightarrow 2$) Let $\mathcal{M}$ be the collection of all cycles in $G = (T, T \times T)$ such that, for all $M \in \mathcal{M}$ and all nodes t, s in M, $y^t \neq y^s$. Notice that any cycle in $\mathcal{M}$ can have at most $|T|$ nodes, so the number of elements in $\mathcal{M}$ is finite.

Given that there are only finitely many cycles in $\mathcal{M}$, there should exist an ε such that, for all $M \in \mathcal{M}$,

$$\sum_{s \in M} z^s(y^{s+} - y^s) > \varepsilon |M|,$$

where $|M|$ is the number of nodes in M.

Now, fix a node m and let $\mathcal{P}_t$ denote the collection of all finite paths in $G = (T, T \times T)$ from m to node t. Define,

$$u^t = \min_{\pi \in \mathcal{P}_t} \sum_{s \in \pi} z^s(y^{s+} - y^s) - \varepsilon 1[y^{s+} \neq y^s].$$

In order to show that this is well-defined, we need to show that there are no cycles C in $G = (T, T \times T)$ such that,

$$\sum_{s \in C} z^s(y^{s+} - y^s) - \varepsilon 1[y^{s+} \neq y^s] < 0.$$

If $y^{s+} = y^s$ for all $s \in C$, then this is obviously satisfied. Else we have, by Lemma 4, a collection of cycles in $\mathcal{M}$ such that,

$$\sum_{s \in C} z^s(y^{s+} - y^s) = \sum_{M \in \mathcal{C}} \sum_{s \in M} z^s(y^{s+} - y^s),$$

and,

$$\sum_{s \in C} 1[y^{s+} \neq y^s] = \sum_{M \in \mathcal{C}} \sum_{s \in M} 1[y^{s+} \neq y^s].$$

Then,

$$\begin{aligned}
&\sum_{s \in C} z^s(y^{s+} - y^s) - \varepsilon 1[y^{s+} - y^s], \\
&= \sum_{M \in \mathcal{C}} \sum_{s \in M} z^s(y^{s+} - y^s) - \varepsilon \sum_{M \in \mathcal{C}} \sum_{s \in M} 1[y^{s+} \neq y^s], \\
&= \sum_{M \in \mathcal{C}} \left(\sum_{s \in M} z^s(y^{s+} - y^s) - \varepsilon 1[y^{s+} \neq y^s] \right) > 0,
\end{aligned}$$

by assumption on the value of ε. As such, we can restrict the minimization over the set of all paths without cycles, which shows that u^t is bounded from below and therefore well-defined. Now, for all paths from m to t we can define a path from m to s by adding the edge (t, s). This means that,

$$u^s \leq u^t + z^t(y^s - y^t) - \varepsilon 1[y^s \neq y^t],$$

so $u^s \leq u^t + z^t(y^s - y^t)$ and $u^s < u^t + z^t(y^s - y^t)$ if $y^s \neq y^t$ as we wanted to show. □

B Robustness checks

To check robustness of our empirical results, we conduct two additional analyses. First, we consider a number of alternative pass/fail benchmarks to decide whether or not the observed behavior passes a particular test. Second, we present results for different lengths of the time blocks used in our rationalizability tests.

B.1 Alternative pass/fail benchmarks

As an introduction, we first provide some additional results for our statistical rationalizability test of the baseline model and our exact rationalizability tests of the log-concave models. Column 3 of Table 4 presents the means and standard deviations of the p-values for our statistical test, both for the actual data and the 500 random data sets (each one obtained by uniformly drawing bids between 0 and the reward r^t, as we explained in the main text). For all but the first time block, we find that the median of the p-values for "real" DMs is significantly above that for the "randomly generated" DMs (i.e. the H_0 of equal medians has a probability $p < .001$). We obtain this conclusion on the basis of Wilcoxon rank-sum tests for differences in medians.[16]

A similar conclusion holds for the log-concave models, of which the results are reported in columns 4-6 of Table 4. Particularly, we conclude that both the mean (using t-tests) and median (using Wilcoxon rank-sum tests) of the HMI values for the real DMs significantly exceed those of the random subjects (i.e. the H_0 of equal medians/means has a probability $p < .001$) in all instances, with the sole exception of the log-concave P model in the first time block. From all this, we may safely conclude that our four models of expected utility maximization provide a better description of the observed decision behavior than the model that predicts random behavior, which further motivates our conclusion in the main text.

[16] We focus on medians for the baseline model as the p-values seem to follow a bimodal distribution (which makes a t-test for differences in means less applicable). Still, it is worth adding that t-tests give the same qualitative conclusions as our Wilcoxon rank-sum tests.

Table 4: Mean and standard deviation of p-values (for the baseline model) and HMI values (for the log-concave models)

rounds		Baseline	Log-conc P	Log-conc U	Log-conc P and U
$1 \leq t \leq 10$	Real Subj	.65 (.47)	.66 (.16)	.86 (.10)	.63 (.14)
	Random Subj	.29 (.42)	.62 (.11)	.62 (.11)	.47 (.10)
$11 \leq t \leq 20$	Real Subj	.76 (.41)	.72 (.14)	.93 (.10)	.70 (.15)
	Random Subj	.34 (.43)	.61 (.11)	.61 (.11)	.44 (.09)
$21 \leq t \leq 30$	Real Subj	.83 (.38)	.74 (.16)	.92 (.10)	.72 (.16)
	Random Subj	.34 (.44)	.61 (.11)	.62 (.11)	.46 (.10)
$31 \leq t \leq 40$	Real Subj	.71 (.43)	.78 (.19)	.94 (.10)	.75 (.19)
	Random Subj	.46 (.47)	.62 (.10)	62 (.11)	.46 (.09)
$41 \leq t \leq 50$	Real Subjects	.86 (.34)	.78 (.17)	.94 (.08)	.74 (.19)
	Random Subjects	.35 (.44)	.62 (.10)	.61 (.10)	.45 (.09)

Baseline model. Next, Table 5 presents the pass rates for our rationalizability test of the baseline model for four different pass/fail benchmarks. Particularly, we consider the nominal sizes $\alpha = .05$ and $\alpha = .10$ (columns 3 and 4) as well as cut-off values that are based on the 90% and 95% rejection rates for the random data sets (columns 5 and 6). For each time block, we report the pass rate for our sample of DMs, the 95% confidence interval for the percentage of DMs passing the test, and the share of random subjects for which we do not reject rationalizability.

A first observation from Table 5 is that the pass rate for the real DMs is everywhere above the one for the random subjects. Next, for more than half of the random subjects we do not reject rationalizability when using $\alpha = .05$ and $\alpha = .10$. This shows that the test is rather conservative and justifies our use of a power-corrected threshold in our main analysis in Section 5. Further, when considering the last two columns of Table 5, we observe that the number of subjects passing our rationalizability test is gradually increasing over time, especially in the earlier time blocks. This is consistent with the patterns of learning and fatigue that we document in the main text.

Log-concave models. Table 6 presents the pass rates of our rationalizability tests of the log-concave models when using alternative pass/fail benchmarks for the HMI. The structure of the table is similar to the one of Table 5. In addition to power-corrected thresholds (columns 5 and 6), we show the pass rates when using HMI = 1 and HMI = .9 as cut-offs (column 3 and 4). The table makes clear that these alternative cut-offs imply stricter rationalizability tests than the

Table 5: Pass rates for the baseline model for different levels of type-one errors

rounds		$\alpha \geq .05$	$\alpha \geq .10$	Power$= .9$	Power$= .95$
$1 \leq t \leq 10$	Pass Rate	19 (68%)	19 (68%)	12 (43%)	9 (32%)
	95% CI	(48% – 84%)	(48% – 84%)	(25% – 63%)	(16% – 52%)
	Random subj	40%	37%	10%	5%
$11 \leq t \leq 20$	Pass Rate	23 (82%)	23 (82%)	17 (61%)	14 (50%)
	95% CI	(63% – 94%)	(63% – 94%)	(41% – 79%)	(13% – 49%)
	Random subj	45%	42%	10%	5%
$21 \leq t \leq 30$	Pass Rate	24 (86%)	24 (86%)	20 (71%)	20 (71%)
	95% CI	(67% – 96%)	(67% – 96%)	(51% – 87%)	(51% – 87%)
	Random subj	44%	42%	10%	5%
$31 \leq t \leq 40$	Pass Rate	22 (79%)	22 (79%)	12 (43%)	11 (39%)
	95% CI	(59% – 92%)	(59% – 92%)	(25% – 63%)	(22% – 59%)
	Random subj	55%	53%	10%	5%
$41 \leq t \leq 50$	Pass Rate	25 (89%)	25 (89%)	20 (71%)	17 (61%)
	95% CI	(72% – 98%)	(72% – 98%)	(51% – 87%)	(41% – 79%)
	Random subj	46%	43%	10%	5%

power-corrected cut-offs.

Let us start with the log-concave P model. The exact rationalizability test (with HMI=1 as the pass/fail benchmark) appears to be quite restrictive: at most 7 subjects pass this test. Once we allow subjects to make at most one mistake (using HMI=.9 as the pass/fail benchmark), the maximum pass rate increases up to 13 subjects. In this case, the probability that a random subject passes the rationalizability test still does not exceed 1%. Further, the results in Table 6 reveal dynamic patterns that are similar to the ones documented in the main text. Particularly, pass rates are increasing for later time blocks, which reflects a learning effect.

We next turn to the model with log-concave U. Even though the associated rationalizability test is really restrictive when using HMI=.9 and HMI=1 to define the pass/fail benchmarks (with 0% and 1% of the random subjects passing the respective tests), we do identify a significant number of observed DMs that behave consistent with the model. In this respect, we recall that the model basically imposes the restriction that bids are increasing if rewards go up, which makes that these high pass rates should not be too surprising. Moreover, even if we do not use power-corrected pass/fail benchmarks, we observe that the pass rate increases over time (especially for HMI= .9). Further, we note that the results for the two

Table 6: Pass rates for the log-concave models for different HMI cut-off levels

		Log-concave P model			
rounds		HMI= 1	HMI= .9	Power= .9	Power= .95
$1 \leq t \leq 10$	Pass Rate	1 (4%)	5 (18%)	13 (46%)	8 (29%)
	95% CI	(0% – 18%)	(6% – 37%)	(28% – 66%)	(13% – 49%)
	Random subj	0%	1%	10%	5%
$11 \leq t \leq 20$	Pass Rate	1 (4%)	6 (21%)	11 (39%)	11 (39%)
	95% CI	–	(8% – 41%)	(22% – 59%)	(22% – 59%)
	Random subj	0%	1%	10%	5%
$21 \leq t \leq 30$	Pass Rate	2 (7%)	7 (25%)	14 (50%)	14 (50%)
	95% CI	(1% – 24%)	(11% – 45%)	(31% – 69%)	(31% – 69%)
	Random subj	0%	2%	10%	5%
$31 \leq t \leq 40$	Pass Rate	7 (25%)	12 (43%)	19 (68%)	16 (57%)
	95% CI	(11% – 45%)	(25% – 63%)	(48% – 84%)	(37% – 76%)
	Random subj	0%	1%	10%	5%
$41 \leq t \leq 50$	Pass Rate	3 (11%)	13 (46%)	21 (75%)	18 (64%)
	95% CI	(2% – 28%)	(28% – 66%)	(55% – 89%)	(44% – 81%)
	Random subj	0%	1%	10%	5%
		Log-concave U model			
rounds		HMI= 1	HMI= .9	Power= .9	Power= .95
$1 \leq t \leq 10$	Pass Rate	6 (21%)	15 (54%)	25 (89%)	25 (89%)
	95% CI	(8% – 41%)	(34% – 73%)	(72% – 98%)	(72% – 98%)
	Random subj	0%	1%	10%	5%
$11 \leq t \leq 20$	Pass Rate	16 (57%)	22 (79%)	28 (100%)	25 (89%)
	95% CI	(37% – 76%)	(59% – 92%)	–	(72% – 98%)
	Random subj	0%	2%	10%	5%
$21 \leq t \leq 30$	Pass Rate	15 (54%)	19 (68%)	27 (96%)	27 (96%)
	95% CI	(34% – 73%)	(48% – 84%)	(82% – 99%)	(82% – 99%)
	Random subj	0%	1%	10%	5%
$31 \leq t \leq 40$	Pass Rate	16 (57%)	25 (89%)	26 (93%)	26 (93%)
	95% CI	(37% – 76%)	(72% – 98%)	(77% – 99%)	(77% – 99%)
	Random subj	0%	1%	10%	5%
$41 \leq t \leq 50$	Pass Rate	16 (57%)	23 (82%)	28 (100%)	28 (100%)
	95% CI	(37% – 76%)	(63% – 94%)	–	–
	Random subj	0%	1%	10%	5%
		Log-concave P and U model			
rounds		HMI= 1	HMI= .9	Power= .9	Power= .95
$1 \leq t \leq 10$	Pass Rate	1 (4%)	3 (11%)	19 (68%)	19 (68%)
	95% CI	(0% – 18%)	(2% – 28%)	(48% – 84%)	(48% – 84%)
	Random subj 0%	0%	10%	5%	
$11 \leq t \leq 20$	Pass Rate	1 (4%)	6 (21%)	23 (82%)	23 (82%)
	95% CI	(0% – 18%)	(8% – 41%)	(63% – 94%)	(63% – 94%)
	Random subj	0%	0%	10%	5%
$21 \leq t \leq 30$	Pass Rate	2 (7%)	6 (21%)	25 (89%)	25 (89%)
	95% CI	(1% – 24%)	(8% – 41%)	(72% – 98%)	(72% – 98%)
	Random subj	0%	0%	10%	5%
$31 \leq t \leq 40$	Pass Rate	6 (21%)	9 (32%)	23 (82%)	23 (82%)
	95% CI	(8% – 41%)	(16% – 52%)	(63% – 94%)	(63% – 94%)
	Random subj	0%	0%	10%	5%
$41 \leq t \leq 50$	Pass Rate	2 (7%)	12 (43%)	23 (82%)	23 (82%)
	95% CI	(1% – 24%)	(25% – 63%)	(63% – 94%)	(63% – 94%)
	Random subj	0%	0%	10%	5%

power-corrected pass/fail benchmarks are exactly the same for all but the second time block.

Finally, we consider the model with log-concave P and U. By construction, the test of this model is more strict than the tests of the models with log-concave P or log-concave U. As an implication, its pass rates for the absolute thresholds HMI= 1 and HMI= .9 cannot exceed those of these other models. Following the same argument, it is not surprising that almost none of the random subjects passes the rationalizability tests for these absolute pass/fail benchmarks. For the two power-corrected tests, we get the same results for all time blocks. This observation motivates the robustness of our results.

B.2 Varying the length of the time blocks

As a further robustness check, we vary the length of our time blocks. Particularly, in our main analysis we have split the data per DM into five blocks of 10 rounds each. In what follows, we redo our rationalizability tests for blocks consisting of 8 and 12 rounds. This will show that our main conclusions are not overly reliant on our particular choice of block length.

8 round blocks. Table 7 presents the results for our statistical test of the baseline model when using 8 round blocks. Since we shorten the time period, our test uses less rationalizability constraints per DM, which makes the test less powerful. This results in larger differences between the pass rates for generic pass/fail benchmarks α (columns 3 and 4) and power-adjusted pass/fail benchmarks (columns 5 and 6). Importantly, however, we find again that all pass rates are higher for the "real" DMs than for the "randomly generated" subjects. Next, the increasing trend in pass rates (suggesting a learning effect) is somewhat less pronounced in this case.

Table 8 presents the results for the log-concave models when using 8 round blocks. The various panels present results for the models with log-concave P, log-concave U, and log-concave P and U. Similar to before, we present results for HMI = 1, HMI = 7/8 (which means that we allow DMs to make one mistake) and power-corrected cut-offs. Once more, the rationalizability tests are more permissive since the number of rationalizability constraints decreases per DM. This makes that a larger number of subjects pass the test for the cut-offs HMI= 1 and HMI= 7/8. Moreover, the lower number of observations implies that our power-correction is generally less efficient. Nonetheless, the power-corrected pass rates are comparable to those that we report for 10 round blocks in the main text. In addition, in almost all cases we observe increasing pass rates over time.

Table 7: Pass rates for the baseline model with 8 round blocks

Rounds		$\alpha \geq .05$	$\alpha \geq .10$	Power= .9	Power= .95
	Pass Rate	22 (79%)	20 (71%)	15 (54%)	14 (50%)
$1 \leq t \leq 8$	95% CI	(59% – 92%)	(51% – 87%)	(34% – 72%)	(31% – 69%)
	Random subj	46%	44%	10%	5%
	Pass Rate	22 (79%)	21 (75%)	11 (39%)	9 (32%)
$9 \leq t \leq 16$	95% CI	(59% – 92%)	(55% – 89%)	(22% – 59%)	(16% – 52%)
	Random subj	57%	55%	10%	5%
	Pass Rate	18 (64%)	18 (64%)	14 (50%)	12 (43%)
$17 \leq t \leq 24$	95% CI	(44% – 81%)	(44% – 81%)	(31% – 69%)	(25% – 63%)
	Random subjects	40%	40%	10%	5%
	Pass Rate	27 (96%)	27 (96%)	23 (82%)	15 (54%)
$25 \leq t \leq 32$	95% CI	(82% – 100%)	(82% – 100%)	(63% – 94%)	(34% – 73%)
	Random subjects	47%	44%	10%	5%
	Pass Rate	24 (86%)	24 (86%)	13 (46%)	12 (43%)
$33 \leq t \leq 40$	95% CI	(67% – 96%)	(67% – 96%)	(28% – 66%)	(25% – 63%)
	Random subjects	56%	55%	10%	5%
	Pass Rate	24 (86%)	24 (86%)	19 (68%)	17 (61%)
$41 \leq t \leq 48$	95% CI	(67% – 96%)	(67% – 96%)	(48% – 84%)	(41% – 79%)
	Random subjects	46%	44%	10%	5%

Table 8: Pass rates for the log-concave models with 8 round blocks

Log-concave P model					
Rounds		HMI= 1	HMI= 7/8	Power= .9	Power= .95
$1 \leq t \leq 8$	Pass Rate	5 (18%)	7 (25%)	7 (25%)	7 (25%)
	95% CI	(6% – 37%)	(11% – 45%)	(11% – 45%)	(11% – 45%)
	Random subj	1%	11%	10%	5%
$9 \leq t \leq 16$	Pass Rate	4 (14%)	10 (36%)	10 (36%)	10 (36%)
	95% CI	(4% – 33%)	(19% – 56%)	(19% – 56%)	(19% – 56%)
	Random subj	1%	5%	10%	5%
$17 \leq t \leq 24$	Pass Rate	4 (14%)	8 (29%)	21 (75%)	8 (29%)
	95% CI	(4% – 33%)	(13% – 49%)	(55% – 89%)	(13% – 49%)
	Random subj	1%	9%	10%	5%
$25 \leq t \leq 32$	Pass Rate	5 (18%)	16 (57%)	21 (75%)	16 (57%)
	95% CI	(6% – 37%)	(37% – 76%)	(55% – 89%)	(37% – 76%)
	Random subj	1%	8%	10%	5%
$33 \leq t \leq 40$	Pass Rate	7 (25%)	15 (54%)	15 (54%)	15 (54%)
	95% CI	(11% – 45%)	(34% – 73%)	(34% – 73%)	(34% – 73%)
	Random subj	1%	10%	10%	5%
$41 \leq t \leq 48$	Pass Rate	5 (18%)	15 (54%)	21 (75%)	15 (54%)
	95% CI	(6% – 37%)	(34% – 73%)	(55% – 89%)	(34% – 73%)
	Random subj	0%	8%	10%	5%
Log-concave U model					
Rounds		HMI= 1	HMI= 7/8	Power= .9	Power= .95
$1 \leq t \leq 8$	Pass Rate	11 (39%)	22 (79%)	27 (96%)	22 (79%)
	95% CI	(22% – 59%)	(59% – 92%)	(82% – 100%)	(59% – 92%)
	Random subj	0%	9%	10%	5%
$9 \leq t \leq 16$	Pass Rate	13 (46%)	25 (89%)	27 (96%)	27 (96%)
	95% CI	(28% – 66%)	(72% – 98%)	(82% – 100%)	(82% – 100%)
	Random subj	0%	5%	10%	5%
$17 \leq t \leq 24$	Pass Rate	18 (64%)	22 (79%)	22 (79%)	22 (79%)
	95% CI	(44% – 81%)	(59% – 92%)	(59% – 92%)	(59% – 92%)
	Random subj	1%	11%	10%	5%
$25 \leq t \leq 32$	Pass Rate	17 (61%)	23 (82%)	27 (96%)	23 (82%)
	95% CI	(41% – 79%)	(63% – 94%)	(82% – 100%)	(63% – 94%)
	Random subj	1%	11%	10%	5%
$33 \leq t \leq 40$	Pass Rate	21 (75%)	26 (93%)	28 (100%)	26 (93%)
	95% CI	(55% – 89%)	(77% – 99%)	–	(77% – 99%)
	Random subj	1%	9%	10%	5%
$41 \leq t \leq 48$	Pass Rate	19 (68%)	26 (93%)	28 (100%)	26 (93%)
	95% CI	(48% – 84%)	(77% – 99%)	–	(77% – 99%)
	Random subj	1%	7%	10%	5%
Log-concave P and U model					
Rounds		HMI= 1	HMI= 7/8	Power= .9	Power= .95
$1 \leq t \leq 8$	Pass Rate	3 (11%)	5 (18%)	27 (96%)	20 (71%)
	95% CI	(2% – 28%)	(6% – 37%)	(82% – 100%)	(51% – 87%)
	Random subj	0%	0%	10%	5%
$9 \leq t \leq 16$	Pass Rate	3 (11%)	7 (25%)	27 (96%)	22 (79%)
	95% CI	(2% – 28%)	(11% – 45%)	(82% – 100%)	(59% – 92%)
	Random subj	0%	0%	10%	5%
$17 \leq t \leq 24$	Pass Rate	3 (11%)	8 (29%)	22 (79%)	25 (89%)
	95% CI	(2% – 28%)	(13% – 49%)	(59% – 92%)	(72% – 98%)
	Random subj	1%	0%	10%	5%
$25 \leq t \leq 32$	Pass Rate	4 (14%)	11 (39%)	27 (96%)	24 (86%)
	95% CI	(4% – 33%)	(22% – 59%)	(82% – 100%)	(67% – 96%)
	Random subj	0%	0%	10%	5%
$33 \leq t \leq 40$	Pass Rate	6 (21%)	14 (50%)	28 (100%)	25 (89%)
	95% CI	(8% – 41%)	(31% – 69%)	–	(72% – 98%)
	Random subj	0%	1%	10%	5%
$41 \leq t \leq 48$	Pass Rate	5 (18%)	13 (46%)	28 (100%)	24 (86%)
	95% CI	(6% – 37%)	(28% – 66%)	–	(67% – 96%)
	Random subj	0%	1%	10%	5%

Table 9: Pass rates for the baseline model with 12 round blocks

Rounds		$\alpha \geq .05$	$\alpha \geq .10$	Power= .9	Power= .95
$1 \leq t \leq 12$	Pass Rate	21 (75%)	21 (75%)	19 (68%)	18 (57%)
	95% CI	(55% – 89%)	(55% – 89%)	(48% – 84%)	(44% – 81%)
	Random subj	35%	34%	10%	5%
$13 \leq t \leq 24$	Pass Rate	24 (86%)	24 (86%)	18 (57%)	15 (54%)
	95% CI	(67% – 96%)	(67% – 96%)	(44% – 81%)	(34% – 73%)
	Random subj	44%	38%	10%	5%
$25 \leq t \leq 36$	Pass Rate	25 (89%)	25 (89%)	24 (86%)	22 (79%)
	95% CI	(72% – 98%)	(72% – 98%)	(67% – 96%)	(59% – 92%)
	Random subj	49%	48%	10%	5%
$37 \leq t \leq 48$	Pass Rate	23 (82%)	23 (82%)	24 (86%)	21 (75%)
	95% CI	(63% – 94%)	(63% – 94%)	(67% – 96%)	(55% – 89%)
	Random subj	41%	40%	10%	5%

12 round blocks. Table 9 presents the results for our test of the baseline model when using blocks of 12 rounds. In this case, the test has become stricter, as the number of rationalizability restrictions increases. Interestingly, we again find that the pass rates are consistent with those presented in the main text. In particular, the pass rates for the real DMs are significantly above those for the random subjects in all cases. In addition, the pass rates exhibit an increasing trend over time, especially in the earlier time periods (albeit somewhat less pronounced than in Table 1 in the main text).

Finally, Table 10 presents the results for the log-concave models when using 12 round blocks. Because our rationalizability tests become stricter, fewer subjects pass the tests for the cut-offs HMI= 1 and HMI= 11/12. However, the power-corrected pass rates in the table are comparable to those for 10 round blocks that we report in the main text. Moreover, we systematically observe an increase in the pass rates over time.

Table 10: Pass rates for the log-concave models with 12 round blocks

		Log-concave P model			
Rounds		HMI= 1	HMI= 11/12	Power= .9	Power= .95
$1 \leq t \leq 12$	Pass Rate	1 (4%)	3 (11%)	9 (32%)	9 (32%)
	95% CI	(0% – 18%)	(2% – 28%)	(16% – 52%)	(16% – 52%)
	Random subj	0%	0%	10%	5%
$13 \leq t \leq 24$	Pass Rate	1 (4%)	3 (11%)	20 (71%)	11 (39%)
	95% CI	(0% – 18%)	(2% – 28%)	(51% – 87%)	(22% – 59%)
	Random subj	0%	0%	10%	5%
$25 \leq t \leq 36$	Pass Rate	3 (11%)	9 (32%)	15 (54%)	15 (54%)
	95% CI (2% – 28%)	(16% – 52%)	(34% – 73%)	(34% – 73%)	
	Random subj	0%	0%	10%	5%
$37 \leq t \leq 48$	Pass Rate	1 (4%)	6 (21%)	21 (75%)	18 (64%)
	95% CI	(0% – 18%)	(8% – 41%)	(55% – 89%)	(44% – 81%)
	Random subj	0%	0%	10%	5%
		Log-concave U model			
Rounds		HMI= 1	HMI= 11/12	Power= .9	Power= .95
$1 \leq t \leq 12$	Pass Rate	5 (18%)	12 (43%)	25 (89%)	25 (89%)
	95% CI	(6% – 37%)	(25% – 63%)	(72% – 97%)	(72% – 97%)
	Random subj	0%	0%	10%	5%
$13 \leq t \leq 24$	Pass Rate	12 (43%)	20 (71%)	28 (100%)	26 (93%)
	95% CI	(25% – 63%)	(51% – 87%)	–	(77% – 99%)
	Random subj	0%	0%	10%	5%
$25 \leq t \leq 36$	Pass Rate	14 (50%)	21 (75%)	26 (93%)	26 (93%)
	95% CI	(31% – 69%)	(55% – 89%)	(77% – 99%)	(77% – 99%)
	Random subj	0%	0%	10%	5%
	Pass Rate	11 (39%)	18 (64%)	28 (100%)	28 (100%)
$37 \leq t \leq 48$	95% CI	(22% – 59%)	(44% – 81%)	–	–
	Random subj	0%	0%	10%	5%
		Log-concave P and U model			
Rounds		HMI= 1	HMI= 11/12	Power= .9	Power= .95
$1 \leq t \leq 12$	Pass Rate	1 (4%)	2 (7%)	22 (79%)	18 (64%)
	95% CI	(0% – 18%)	(1% – 24%)	(59% – 92%)	(44% – 81%)
	Random subj	0%	0%	10%	5%
$13 \leq t \leq 24$	Pass Rate	1 (4%)	3 (11%)	28 (100%)	23 (82%)
	95% CI	(0% – 18%)	(2% – 28%)	–	(63% – 94%)
	Random subj	0%	0%	10%	5%
$25 \leq t \leq 36$	Pass Rate	3 (11%)	5 (21%)	21 (75%)	21 (75%)
	95% CI (2% – 28%)	(6% – 37%)	(55% – 89%)	(55% – 89%)	
	Random subj	0%	0%	10%	5%
$37 \leq t \leq 48$	Pass Rate	1 (4%)	6 (21%)	26 (93%)	23 (82%)
	95% CI	(0% – 18%)	(8% – 41%)	(77% – 99%)	(63% – 94%)
	Random subj	0%	0%	10%	5%