

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

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

November 4, 2021

Abstract

We provide a revealed preference characterization of expected utility maximization in binary lotteries with prize-probability trade-offs. We start by characterizing optimizing behavior when the empirical analyst exactly knows the utility function or the probability function of winning. Next, we consider the situation with both the probability function and the utility function unknown. In this case utility maximization has empirical content when imposing the mild shape restriction that at least one of these functions is log-concave.

*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.

Keywords: expected utility maximization, prize-probability trade-offs, revealed preference characterization, testable implications.

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, we take a framework where a lottery yields a reward $r - b$ with probability $P(b)$ and a payoff of zero with probability $1 - P(b)$. Here, 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 reward and the probability of winning.

This type of decision problem occurs frequently in economics. A notable example is the (independent private values, sealed-bid) 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 choose to increase the probability of winning the auction (in a monotone equilibrium) by increasing her bid b , but this implies that the final value of winning the auction, i.e. $(r - b)$, decreases. In what follows, we do not explicitly consider the strategic aspect of this game and concentrate mainly on the single-agent decision problem. Under the assumption that players play a Bayesian Nash equilibrium, the probability of winning, given the bid $P(b)$, captures all the relevant information for the the DM to choose her optimal bid. The first price auction is just one instance fitting in our general set-up. In Section 2 we will briefly discuss additional examples of often studied decision problems that are also characterized by prize-probability trade-offs in –admittedly– settings that are mostly more complex in reality.

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. Our set-up is clearly data restrictive since it assumes multiple observations of the same agent. This makes our approach more readily applicable to an experimental set-up that wants to investigate theoretical properties or identifications strategies for a given setting (see Capra, Croson, Rigdon, and Rosenblat, 2020, for a motivating review of different games and experiments). Coming back to our first price auction example, it is in particular interesting to note that there is a sizeable experimental literature that focuses on this specific decision situation (see, for example, Kagel and Levin (2016) for an overview).

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). A sizeable literature has emerged on testing decision theories under risk using this revealed preference approach. However, this literature has mainly focused on choices involving Arrow-Debreu securities from linear budgets (see, for example, Varian, 1983; Green and Osband, 1991; Kubler, Selden, and Wei, 2014; Echenique and Saito, 2015; Chambers, Echenique, and Saito, 2016; Polisson, Quah, and Renou, 2020), with a few papers focusing on the full mixture space (see, for example, Kim, 1996). We complement these earlier studies by considering expected utility maximization in a distinctively different decision setting.

Main theoretical results. To set the stage, 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 $r - b$).¹ 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

¹Admittedly, the assumption that P is perfectly observed is rather demanding. Therefore, in Appendix B we also present a statistical test derived from our testable conditions in Section 3 when P can (only) be estimated from a finite sample of observations.

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.

In most empirical settings, however, both P and U are unknown and the question arises whether we can obtain any testable implications in such a case. 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 . However, building on our first set of results, we show that this negative conclusion can be overcome by imposing minimalistic shape constraints that are often used in the literature.² 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. Log-concavity is a very weak assumption that is closely linked to monotonicity (see, for example, Cox et al., 1988; Bagnoli and Bergstrom, 2005). More specifically, log-concavity of U still allows the DM to be risk-loving but (only) excludes extremely risk-loving behavior. Intuitively, log-concavity of U imposes a single-crossing property of utility functions that is frequently used in game theory and mechanism design (see, for example, Maskin and Riley, 2000). Similarly, log-concavity of P is a minimal assumption that holds for most commonly used distributions in the literature, making it again a fairly weak restriction.

For each of these models, we derive necessary and sufficient testable conditions for expected utility maximization that are of the law-of-demand type. They require respectively that (1) higher rewards r must lead to higher payoffs $r - b$, (2) higher rewards r must lead to higher bids b , and (3) higher rewards r must lead to both higher payoffs $r - b$ and higher bids b . These results are in line with comparative static results that have been documented in the literature. A notable implication of our nonparametric characterizations is that the testable conditions are not only necessary but also sufficient for expected utility maximization.

Empirical implications. Our theoretical results can have alternative empirical applications. For compactness, we do not provide an empirical illustration of

²Dziewulski (2018) followed a related mathematical approach in a conceptually different setting. Particularly, this author developed revealed preference characterizations of rationalizable behavior for common specifications of the discounted utility model by referring to notions of stochastic dominance.

our theoretical characterization of expected utility maximization in the current paper.³ Next, our results can be employed to empirically test for equilibrium best-responding behavior of players in games with prize-probability trade-offs (such as auctions). If we assume that observed behavior is in equilibrium, then each player should maximize her expected utility with the prize-probability trade-off function ($P(b)$) defined by the equilibrium actions of all other players.

Furthermore, our characterizations entail two important conclusions with direct empirical relevance. 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, as we will discuss in Section 5, even if the rewards r are unobserved, the above comparative static results 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 distribution P and any data set with payoffs $r - b$ increasing in rewards r , we can find a utility function 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 function 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 more specific properties of these functions. These findings are similar in spirit to those of Manski (2002, 2004) on the impossibility to separately identify decision rules and beliefs.

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

³In an earlier version of the current paper (Cherchye, Demuynck, De Rock, and Freer, 2019) we used our revealed preference conditions to analyze Neugebauer and Perote (2008)’s experimental data on first-price auctions. For instance, there is a growing literature on the econometric analysis of auctions that focuses on identifying the distribution of values from the observed distribution of bids (see, for example, Guerre, Perrigne, and Vuong (2000) and Athey and Haile (2002, 2007)). Integrating our results in this econometric work may lead to verifying the underlying model assumptions and, consequently, to more robust conclusions (see also Appendix B for related results).

either the probability function P or the utility function U . Section 4 analyzes the setting with both P and U unknown. Section 5 discusses the usefulness of our theoretical results when the rewards r are unobserved. Section 6 presents our concluding discussion.

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 $r \leq \bar{r}$ for some exogenously given $\bar{r} \in \mathbb{R}$. The DM can choose a bid $b \in [0, \bar{r}]$. 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

$$U : [0, \bar{r}] \rightarrow \mathbb{R}_+,$$

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 $[0, \bar{r}]$, and that U is continuous and strictly increasing on \mathbb{R} .⁴ Observe that we can indeed restrict $b \leq r$ in this optimization problem, as any bid $b > r$ gives negative

⁴For our results, the monotonicity and continuity properties are inherited from P to U and vice versa. Particularly, we obtain readily similar revealed preference characterizations as in Theorems 1 and 2 and 3 when relaxing a property of P (e.g., assuming that it is just increasing instead of strictly increasing) and, simultaneously, the corresponding property of U (e.g., equally assuming that it is increasing instead of strictly increasing).

expected utility and is therefore dominated by a choice $b = r$, which gives zero expected utility. Next, we consider P to be independent of r mainly to ease our exposition. Our results in Section 3 (for known P or U) can be replicated for P dependent on r without extra assumptions.

In particular, our results would still hold under the assumption that the probability of winning is conditional on the reward, that is, $P(b|r)$. For instance, this would allow us to extend our results from auctions with independent private values to auctions with affiliated values. However, replicating our results in Section 4 (for unknown P and U) would require auxiliary assumptions when P can depend on r .

Our 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. Clearly our goal is not to claim that all these settings are special cases of our set-up. On the contrary, we only want to briefly show that the basic underlying game can be formulated to fit into our framework. As such, this paper presents at least the necessary testable implications that have to be satisfied in a more complex and general setting. Among other things, this also illustrates that our results can be used in other experimental settings than just first price auctions.

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. Thus, if we consider the Bayesian Nash equilibrium, the cdf P is generated as the distribution of highest bids given the equilibrium bidding of other players. As an implication, if we assume equilibrium play, the DM must optimize her expected utility as in (1).

Crowdfunding games. A crowdfunding game is an example of a mechanism to organize private provision of a public good.⁵ The participants in the game make

⁵Similar 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.

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 the public good being provided by:

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

In the Bayesian Nash equilibrium of this crowdfunding game, the cdf P equals the distribution of the sum of contributions of the other players as defined by their equilibrium bidding. Thus, if we assume equilibrium play, the DM has to maximize her expected utility as in (1).

Posted price problems. In a posted price problem, the DM (buyer) has a valuation r for the traded good. In order to obtain the good, the DM posts a price b at which she is willing to buy the good.⁶ 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} . In this case, the probability of the trade is given by:

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

which is determined by the distribution of the seller's reservation price. Thus, the Perfect Bayesian Equilibrium generates the DM's problem in which she maximizes

In this sense, we are closer to Tabarrok (1998). However, we do allow for differentiated (and not only binary) contributions, as in Zubrickas (2014).

⁶The 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.

her expected utility as in (1) for this specification of P .

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

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

The agent chooses the effort level that maximizes her utility while accounting for the cost of effort. At the same time, the probability $P(b)$ depends on b as the agent's utility is conditional on the bonus that is promised to her. Therefore, in a Subgame Perfect Nash equilibrium, the DM maximizes her expected utility as in (1), with the cdf P determined by the agent's optimal effort provision. While, admittedly, this constitutes a most basic version of a principal agent problem, the example illustrates once more that prize-probability trade-offs are relevant in many different settings.

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.⁷ As a first step of our analysis, we 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. A typical instance with observed U occurs when the empirical analyst assumes a risk neutral DM. Next, as indicated in the introductory section, a prime example of the case with observed P is the first price auction of which the participants play a symmetric equilibrium, in which case P equals

⁷We discuss the case of unobserved rewards in Section 5.

the cdf of the player types. In Appendix B, 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 can develop a statistical test of expected utility maximization by starting from our results in the current section. In Section 4, we will focus on the case where both P and U are unobserved.

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 a 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 $t = 1, \dots, T$,

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

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).⁸

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

⁸ Appendix A contains the proofs of our main theoretical results. We slightly abuse notation in Theorem 1 by assuming that $P(x) = 0$ if $x < 0$.

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

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

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 $t = 1, \dots, T$, $b^t < r^t$, and

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

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

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.

As mentioned above, the results of this section can be expanded to the case when the probability of winning depends on r . Evidently, this extension is trivial when $P(b|r)$ is observed. If $P(b|r)$ is unobserved, we need to slightly modify Theorem 1. Instead of using numbers P^t for every $t = 1, \dots, T$, we need to introduce numbers $P^{t,s}$ for $t, s = 1, \dots, T$, where index t corresponds to the potential value b^t and index s to the potential value r^s . In addition, we need to ensure that these numbers correspond to the same monotone function if $r^s = r^m$ for some $r, m = 1, \dots, T$.

Empirical content. We conclude this section 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

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

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

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

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 strictly 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).

4 When P and U are unknown

We next turn to the instance in which both the cdf P and utility function U are unknown to the empirical analyst. 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. As discussed in the Introduction, the assumption of log-concavity is a natural candidate to impose minimal structure on the decision problem.

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 $P(b) = e^{b-\bar{r}}$, which is a continuous and strictly increasing cdf on $[0, \bar{r}]$. 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 ,

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

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

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

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 $[0, \bar{r}]$. 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. $e^{r^t-\bar{r}}$).

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.

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 $t = 1, \dots, T$, $P(b^t) > 0$ and $b^t < r^t$, and
 - (b) for all observations $t, s = 1, \dots, T$, $r^t \geq r^s$ 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 $t = 1, \dots, T$, $b^t < r^t$, and
 - (b) for all observations $t, s = 1, \dots, T$, $r^t \geq r^s$ implies $b^t \geq b^s$.

The rationalizability conditions in Theorem 2 are of the law-of-demand type and 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 some strictly log-concave utility function U , it is rationalizable for any strictly log-concave utility function U .

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 $t, s = 1, \dots, T$,*

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

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 non-parametric 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 When rewards are unobserved

So far we have assumed that the rewards r are observed by the empirical analyst. This assumption holds well in experimental settings, 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 \bar{r} be strictly bigger than $\max_{t \in \{1, \dots, T\}} b^t$, and take any $\Delta > 0$ that satisfies

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

For every observation $t = 1, \dots, T$, we then consider the value $r^t = b^t + \Delta$, which is contained in $[0, \bar{r}]$. 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 (non-parametric) 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:

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

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.

6 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 $r - b$).

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 of the law-of-demand type 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.

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 empirical 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.

Follow-up research may fruitfully focus on extending our theoretical results towards a broader range of decision problems characterized by prize-probability trade-offs. A first avenue could focus on introducing heterogeneity to either P or U . Allowing for P to change across observations would allow for encompassing a broader set of applications such as settings where the DM is competing with different numbers of bidders. Our results can readily be adapted when P is observed, but will require extra identifying information if P is not observed. Similar conclusions hold for allowing U to change in order to capture for instance settings where one only has one observation from multiple participants (instead of multiple observations for one agent).

Next, 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

- Afriat, S. N., 1967. The construction of utility functions from expenditure data. *International Economic Review* 8, 67–77.
- Athey, S., Haile, P. A., 2002. Identification of standard auction models. *Econometrica* 70, 2107–2140.
- 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.
- Capra, C. M., Croson, R. T., Rigdon, M. L., Rosenblat, T. S., 2020. *Handbook of Experimental Game Theory*. Edward Elgar Publishing.
- Castillo, M., Freer, M., 2016. A revealed preference test of quasi-linear preferences. Tech. rep., Working Paper.
- Chambers, C. P., Echenique, F., Saito, K., 2016. Testing theories of financial decision making. *Proceedings of the National Academy of Sciences* 113 (15), 4003–4008.
- Cherchye, L., Demuynck, T., De Rock, B., Freer, M., 2019. Revealed preference analysis of expected utility maximization under prize-probability trade-offs. FEB Research Report Department of Economics DPS19. 13.
- 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.
- Cox, J., Smith, V. L., Walker, J. M., 1988. Theory and individual behavior of first-price auctions. *Journal of Risk and Uncertainty* 1, 61–99.

- Diewert, W. E., 1973. Afriat and revealed preference theory. *Review of Economic Studies* 40, 419–425.
- Dziewulski, P., 2018. Revealed time preference. *Games and Economic Behavior* 112, 67–77.
- 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.
- Guerre, E., Perrigne, I., Vuong, Q., 2000. Optimal nonparametric estimation of first-price auctions. *Econometrica* 68, 525–574.
- Kagel, J., Levin, D., 2016. Auctions: A survey of experimental research. In: Kagel, J. H., Roth, A. E. (Eds.), *The handbook of experimental economics*, volume 2. Princeton university press, Ch. 9, pp. 563–638.
- Kim, T., 1996. Revealed preference theory on the choice of lotteries. *Journal of Mathematical Economics* 26, 463–477.
- Kubler, F., Selden, L., Wei, X., 2014. Asset demand based tests of expected utility maximization. *American Economic Review* 104 (11), 3459–80.
- 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.
- Maskin, E., Riley, J., 2000. Equilibrium in sealed high bid auctions. *The Review of Economic Studies* 67 (3), 439–454.
- 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.

- Polisson, M., Quah, J. K.-H., Renou, L., 2020. Revealed preferences over risk and uncertainty. *American Economic Review* 110 (6), 1782–1820.
- Rochet, J.-C., 1987. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of mathematical Economics* 16 (2), 191–200.
- Rockafellar, R. T., 1970. *Convex analysis*. No. 28. Princeton university press.
- 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 dominant 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 $[0, \bar{r}]$, $P(b^t)$ 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 $U : \mathbb{R} \rightarrow \mathbb{R}$ that rationalizes the data set. Define:

$$U(x) = \min \left\{ \alpha x, \min_{t=1, \dots, T} \left\{ U^t \frac{P(b^t)}{P(r^t - x)} \text{ s.t. } P(r^t - x) > 0 \right\} \right\}, \quad (2)$$

where we choose:

$$\alpha > \max_t \frac{U^t}{r^t - b^t}. \quad (3)$$

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 (3), 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 (2). 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:

$$V(b) = \min \left\{ \alpha b, \min_{t=1, \dots, T} \left\{ P^t \frac{U(r^t - b^t)}{U(r^t - b)} \text{ s.t. } r^t > b \right\} \right\}, \quad (4)$$

where we choose:

$$\alpha > \max_t \frac{P^t}{b^t}. \quad (5)$$

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:

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

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

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 (5)) 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:

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

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:

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

a contradiction.

A.2 Proof of Theorem 2

In order to give the proof, we need to introduce some definitions and notation.

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 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 also denote the successor of n_i as n_i+ and its predecessor as n_i- .

To start, let us give some preliminary results.

Preliminary results

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

The following Lemma is close in spirit to the results of Rochet (1987) and Castillo and Freer (2016).

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 :

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

Summing the left and right hand sides over all nodes t of the cycle C gives:

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

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

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

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 condition in the lemma, an optimal solution to this problem will be path that does not have a cycle. Indeed, if a path includes a cycle, this 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

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:

$$\sum_{t \in C} p(r^t - r^{t+} + b^{t+}) - p(b^{t+}) \leq 0, \quad (6)$$

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 $P(r^s - r^t + b^t) = 0$, 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 (6)):

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

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 (6) 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:

$$\begin{aligned} \sum_{s \in C} p(r^s - r^{s+} + b^{s+}) - p(b^{s+}) &= \sum_{s \in C'} (p(r^s - r^{s+} + b^{s+}) - p(b^{s+})) \\ &\quad + p(r^{t-} - r^t + b^t) - p(b^t) + p(r^t - r^{t+} + b^{t+}) \\ &\quad - p(r^{t-} - r^{t+} + b^{t+}). \end{aligned} \quad (7)$$

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 (7) is negative by the induction hypothesis. As such, it suffices to show that:

$$p(r^{t-} - r^t + b^t) - p(b^t) \leq p(r^{t-} - r^{t+} + b^{t+}) - p(r^t - r^{t+} + b^{t+}). \quad (8)$$

Define $\Delta = r^{t-} - r^t \geq 0$ and set $r^{t-} - r^{t+} + b^{t+} = \tilde{b} \geq 0$. Then, substituting into (8) 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

We first state some preliminary results.

Preliminary results

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 $C = \{(t_1, t_2), (t_2, t_3), \dots, (t_n, t_1)\}$ contains two nodes t_i, t_j ($i < j$) with $y^{t_i} = y^{t_j}$, then we can break up C into two cycles of smaller length. In particular, we have the smaller cycles:

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

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,

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

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 until we 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 set $\mathcal{C} = \{\tilde{C}\}$. Else, let $\tilde{C} = \{(t_1, t_2), \dots, (t_n, t_1)\}$ be such that, for at least two nodes t_i, t_j ($i < j$) in \tilde{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:

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

Notice that \tilde{C}_1 and \tilde{C}_2 satisfy:

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

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 over and over until we 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:

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

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. \square

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:

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

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. \square

Main part of the 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 $0 < b^t < \bar{r}$.⁹ Then, strict concavity of u and p gives:

$$p(b^t) - p(b^s) \leq \partial p^s(b^t - b^s) = \partial u^s(b^t - b^s), \quad (9)$$

$$u(r^t - b^t) - u(r^s - b^s) \leq \partial u^s[(r^t - b^t) - (r^s - b^s)], \quad (10)$$

where the inequality (9) is strict if $b^s \neq b^t$ and the inequality (10) is strict if $r^t - b^t \neq r^s - b^s$. If we exchange t and s in conditions (9) and (10) and add them together, we obtain:

$$0 \leq (\partial u^s - \partial u^t)(b^t - b^s), \quad (11)$$

$$0 \leq (\partial u^s - \partial u^t)[(r^t - b^t) - (r^s - b^s)], \quad (12)$$

where (11) is strict if $b^t \neq b^s$ and (12) is strict if $r^t - b^t \neq r^s - b^s$. If $b^t > b^s$, then, for (11) 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 (12). 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 (12) 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 (11). As such we obtain $r^t - r^s > b^t - b^s \geq 0$ and thus also $r^t > r^s$. Again, by contraposition, we can conclude that $r^s \geq r^t$ implies $r^s - b^s \geq r^t - b^t$.

⁹ For the definition and basic properties of supergradients please see Rockafellar (1970).

(\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 :

$$u^t - u^s \leq r^s [(r^t - b^t) - (r^s - b^s)], \quad (13)$$

$$p^t - p^s \leq r^s (b^t - b^s), \quad (14)$$

where the inequality (13) is strict if $r^t - b^t \neq r^s - b^s$, and the inequality (14) 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 $\lim_{x \rightarrow 0} \tilde{u}(x) = -\infty$. Define:

$$U(x) = \exp(u(x)),$$

and:

$$P(b) = \exp(p(b) - p(\bar{r})).$$

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 $[0, \bar{r}]$.

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 b^t maximizes $p(b) + u(r - b)$. 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 :

$$p(b) + u(r^t - b) - (p(b^t) + u(r^t - b^t)) \leq r^t(b - b^t) + r^t[(r^t - b) - (r^t - b^t)] = 0,$$

as we needed to show.

B When U is unknown and P can be estimated

In this appendix, 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.

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 . The sample used for the cdf of G is a separate data set than the one used for the revealed preference test. We assume that the cdf G can be linked to the cdf P by a known function $\Gamma : [0, 1] \rightarrow [0, 1]$ such that, for all $b \in [0, \bar{r}]$,

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

This function Γ will generally depend on the specific setting at hand. For instance, in a first price auctions we can take G to represent the distribution of bids of a random participant in the auction, 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 $\Gamma(x) = x^k$.¹⁰ 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.

¹⁰The sample $(\hat{b}_j)_{j \leq m}$ of bids can then be obtained via m repetitions of the following procedure. Draw a random subject from the population, endow this subject with a random reward and ask her for her optimal bid.

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 $\mathbf{1}[\cdot]$ 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 $t, s \in \{1, \dots, T\}$. 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:¹¹

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

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

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

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 .¹²

Of course, in practice we do not observe the matrix Ω . We can approximate it

¹¹For simplicity, we assume that all values $r^t - r^s + b^s$ are distinct. Obviously, this does not affect the core of our argument.

¹²See, for example, Sepanski (1994).

using the finite sample analogue $\hat{\Omega}_m$, where:

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

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

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

We can use this last result as a basis for an asymptotic test of rationalizability. Specifically, consider the null hypothesis:

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

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

$$\begin{aligned} \text{O.P.I:} \quad Z_m &= \inf_{e_m, \hat{G}_{t,s} \in [0,1], U^t > 0} m e'_m(\hat{\Omega}_m)^{-1} e_m, \\ \text{s.t.} \quad \forall t, s : \quad e_{t,s} &= \mathbb{G}_m(r^t - r^s + b^s) - \hat{G}_{t,s}, & (15) \\ \Gamma(\hat{G}_{t,t}) U^t &\geq \Gamma(\hat{G}_{t,s}) U^s, & (16) \\ \Gamma(\hat{G}_{t,s}) &< \Gamma(\hat{G}_{t',s'}) \text{ for all } r^t - r^s + b^s < r^{t'} - r^{s'} + b^{s'}. & (17) \end{aligned}$$

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:

$$Z_m \leq m \varepsilon'_m(\hat{\Omega}_m)^{-1} \varepsilon_m.$$

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:

$$\lim_{m \rightarrow \infty} \Pr[Z_m > c_\alpha] \leq \lim_{m \rightarrow \infty} \Pr[m \varepsilon'_m(\hat{\Omega}_m)^{-1} \varepsilon_m > c_\alpha] = \alpha,$$

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

Two concluding 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 may be computationally difficult due to the constraints (16)-(17) that are nonlinear. For some particular instances of the function Γ , however, it may be possible to convert this problem into a problem that can be solved by standard algorithms. See, for example, the working paper version of this paper (Cherchye, Demuynck, De Rock, and Freer (2019)) for an application of this procedure to a first price auction setting.