BROWN
Orlando Bravo Center
for Economic Research

Bounded Rationality and Limited Datasets*

Bravo Working Paper # 2020-008

Geoffroy de Clippel† and Kareen Rozen‡

Abstract: Bounded rationality theories are typically characterized over exhaustive data sets. We develop a methodology to understand the empirical content of such theories with limited data, adapting the classic, revealed-preference approach to new forms of revealed information. We apply our approach to an array of theories, illustrating its versatility. We identify theories and datasets testable in the same elegant way as Rationality, and theories and datasets where testing is more challenging. We show that previous attempts to test consistency of limited data with bounded rationality theories are subject to a conceptual pitfall that can yield false positives and empty out-of-sample predictions.

---

*This work greatly benefitted from suggestions by Eddie Dekel, K_r Eliaz, Yusufcan Masatlioglu, Larry Samuelson, and several seminar audiences. First version 2012.

†Department of Economics. Email address: declippel@brown.edu.

‡Department of Economics. Email address: kareen rozen@brown.edu. Rozen thanks the NSF for financial support in early stages of this work (grant SES-0919955).

BOUNDED RATIONALITY AND LIMITED DATASETS*

Geoffroy de Clippel† and Kareen Rozen‡

December 2018

Abstract

Bounded rationality theories are typically characterized over exhaustive data sets. We develop a methodology to understand the empirical content of such theories with limited data, adapting the classic, revealed-preference approach to new forms of revealed information. We apply our approach to an array of theories, illustrating its versatility. We identify theories and datasets testable in the same elegant way as Rationality, and theories and datasets where testing is more challenging. We show that previous attempts to test consistency of limited data with bounded rationality theories are subject to a conceptual pitfall that can yield false positives and empty out-of-sample predictions.

*This work greatly benefitted from suggestions by Eddie Dekel, Kfir Eliaz, Yusufcan Masatlioglu, Larry Samuelson, and several seminar audiences. First version 2012.

†Department of Economics. Email address: declippel@brown.edu.

‡Department of Economics. Email address: kareen_rozen@brown.edu. Rozen thanks the NSF for financial support in early stages of this work (grant SES-0919955).

1. Introduction

The recent literature has proposed insightful and plausible choice procedures in response to mounting evidence against rational choice. Great progress has been made to understand the set of choice functions that these new theories generate, and which characteristics of the Decision Maker (DM) can be identified. Though theoretically insightful, such results do not apply to typical situations, in which only some choices are observed.[1] Indeed, one important purpose of a theory is the ability to make out-of-sample predictions in such cases. For the theory of Rationality, the problems of testable implications for limited data, out-of-sample prediction, and identification are well understood.[2] We explore how these ideas can be brought into the discourse on bounded rationality.

A DM's observed choices are consistent with a theory if they can be extended to a complete choice function arising under the theory. Previous attempts to study bounded rationality theories under limited data focus on explaining only observed choices, without considering out-of-sample implications.[3] Such an approach may at first seem natural, since there is no need to worry about out-of-sample problems when testing for Rationality in its standard description. Indeed, if one can find a preference ordering for which observed choices are maximal, then choices may be defined for out-of-sample problems simply by maximizing that same preference. We show that such extensibility need not hold in general, resulting in a problem of false positives: a DM's choices may be incorrectly attributed to a theory for which no extension of those choices can arise. A first contribution of this paper is thus to clarify what the right definition of consistency is, effectively moving the goalpost for consistency tests to the proper location.

We then explore how one might capture the empirical content of a wide range of choice theories. For this, we build on the classic, revealed-preference approach for testing Rationality. That insightful approach, culminating in the Strong Axiom of Revealed Preference (SARP), can be decomposed as follows. First, infer key preference comparisons under the presumption that the DM is rational. Next, observe that

[1] In empirical settings, the modeler cannot control the choice problems faced by individuals. In experimental settings, generating a complete dataset requires an overwhelming number of decisions by subjects: 26 choice problems when the space of alternatives contains 5 elements, 1,013 choice problems when it contains 10 elements, and 32,752 choice problems when it contains 15 elements.

[2] See Samuelson (1948), Houthaker (1950), Richter (1966), Afriat (1967), and Varian (1982).

[3] These include Manzini and Mariotti (2007), Manzini and Mariotti (2012) and Tyson (2013).

transitivity of the DM's preference requires these comparisons to be acyclic. Finally, prove that acyclicity is not just necessary, but also sufficient for consistency, to ensure that all key preference comparisons have been inferred from the data.

Moving on to recent preference-based theories of bounded rationality, a first observation we make through numerous illustrations is that choices can reveal more elaborate preference restrictions than the simple preference comparisons (like a is better than b) which are revealed under Rationality. Suppose the DM picks a from the choice problem $\{a, b, d, e\}$ and picks b from $\{a, b\}$. Under Masatlioglu, Nakajima and Ozbay (2012)'s theory of Limited Attention, for instance, such a choice pattern reveals that the DM prefers a to at least one of the options d or e, but does not tell us which one(s).

Importantly, we observe that the emergence of such new forms of revealed preference should not prevent us from pursuing the classic testing approach, which can be adapted as follows: first, infer key preference restrictions under the presumption that the theory holds; next note that consistency with the theory requires these restrictions to be *acyclically satisfiable*, meaning that there exists a strict acyclic relation satisfying them; and finally, check that all key restrictions have been identified, by proving that acyclic satisfiability is not only necessary but also sufficient. Acyclic satisfiability is the natural extension of SARP to accommodate restrictions that may be more complex than the simple comparisons revealed under Rationality.

Performing these steps is always insightful, and improves our understanding of the theory at hand. Yet one may wonder about how acyclic satisfiability itself can be checked in applications. This is straightforward when restrictions are simple comparisons (as in SARP), since such restrictions tell us precisely which incomplete relation must be acyclic. Checking acyclic satisfiability of more complex restrictions, however, entails some guesswork: one must *find* an acyclic relation satisfying them. Nonetheless, we show that one tests acyclic satisfiability in the same simple and elegant way as one tests SARP, when the restrictions all pertain to *lower-contour sets (LCS)*, as in the example above (note the lower-contour set of a is required to contain d or e), or all pertain to *upper-contour sets (UCS)*.

To understand why, remember how we test whether some (possibly incomplete) relation P is acyclic. Clearly, acyclicity implies that one can find at least one option, call it x_1, that is not P-superior to any alternative. Moreover, if there is no cycle among elements of $X \setminus \{x_1\}$, there must exist at least one option, call it x_2, that is not P-superior to any alternative in $X \setminus \{x_1\}$. Repeating this process, we see that acyclicity implies that one can enumerate all the options in such a way that each x_k is not P-superior to any x_j with $j > k$. Clearly, finding such an enumeration is also sufficient for acyclicity. The procedure is simple, and its outcome definitive, because success is independent of the way one makes a selection when multiple candidates for some x_k exist. We show in Proposition 2 that the *same process applies* when checking acyclic satisfiability for LCS restrictions on a relation P.[4] In a nutshell, the reason is this: the key property in each step is merely the ability to find an element whose lower contour set is not required to contain some remaining element(s); this does not require knowing precisely which elements belong to that lower-contour set.

Perhaps surprisingly, we illustrate that a variety of theories beyond Rationality can be tested in this manner, because their empirical content amounts to acyclic satisfiability of all LCS (or all UCS) restrictions. We also show that there are several bounded rationality theories which generate more complicated restrictions. How can we rule out the possibility that, by thinking some more, those restrictions would be simplifiable into LCS (UCS) ones? And, even if the restrictions cannot be simplified further, is there perhaps a cleverer procedure which makes testing simple?

Our work also sheds light on these questions. We show in Proposition 3 that, going beyond the LCS (or UCS) class, checking acyclic satisfiability is generally NP-hard; and we apply this result several times to identify theories that are NP-hard to test. Even in these cases, the characterization of the theory in terms of acyclic satisfiability yields insights and makes the theory much easier to test than through a brute force approach. Moreover, understanding the empirical content in terms of acyclic satisfiability can also help pinpoint classes of datasets that are easy to test, namely those where the corresponding restrictions fall into the LCS (or UCS) class. For instance, testing is easy for some theories if the dataset includes all binary choice sets. Indeed, being NP-hard to test means that there exist datasets where testing is quite difficult, but does not mean that testing is always difficult. It does imply, however, that if one could find a test of the theory that is easy for all datasets, then one would overturn the widely held belief, first conjectured by John Nash in the context of cryptography,[5] that no algorithm solves every instance of an NP-hard

[4] Acyclic satisfiability with UCS restrictions is checked analogously, looking in each step for options whose upper-contour sets are unrestricted.

[5] See Nash (1955), a recently declassified letter to the NSA.

problem in polynomial time.

While so far we have only discussed restrictions pertaining to preferences, we develop applications that illustrate the wide scope of our methodology. Restrictions may pertain to orderings other than preference (e.g. salience or priority), and to relations that may be incomplete (as in problems of just-noticeable differences). They can capture satisficing instead of maximizing behavior. The methodology applies to settings where X is not finite (e.g. consumer theory). It can apply when preferences are weak, when the choice data is stochastic, or when choices are observed only coarsely. Beyond bounded rationality, it even provides useful insight in some classic problems of interactive decision making with rational agents. Several of these applications are carried out in this paper and the Online Appendix; the rest are carried out in subsequent works, some authored by ourselves and some by others.

The versatility of our approach, however, does not mean that there is no limitation. For each application provided, seeing the theory's empirical content through the lens of acyclic satisfiability, and using Propositions 2 or 3 to see whether testing can be done tractably as in SARP, takes a bit of thought.[6] Of course, while many theories (explicitly or implicitly) involve a relation to examine, it is also possible for a theory to have a sufficiently different structure that capturing its empirical content through acyclic satisfiability is unfruitful.

This paper is organized as follows. We set up the framework in Section 2. We develop the appropriate notion of consistency for limited datasets in Section 3, and identify the issue with prior approaches. In Section 4, we build some insight for the methodology by using Limited Attention as a first example. In Section 5, we formalize restrictions and acyclic satisfiability, show that the same process to test SARP applies more broadly when all restrictions are LCS (or UCS), and that testing is generally NP-hard otherwise. We discuss a variety of applications in Section 6.

2. Framework

Consider a finite set X of alternatives. A *choice problem* is a nonempty subset of X and represents those alternatives that are feasible. The set of all conceivable

[6] This is not a feature unique to our work or to choice theory. For instance, neither the revelation principle in mechanism design, nor recursive techniques in repeated games, immediately yield answers in applications. In all these cases, the ideas require further work to bear fruit. See Myerson (1979) and Abreu, Pearce and Stacchetti (1990), respectively.

choice problems is denoted $\mathcal{P}(X)$. A *choice function* $c : \mathcal{P}(X) \to X$ associates an element $c(S) \in S$ to each choice problem S.

A *theory* $\mathcal{T}$ formally describes the DM's choice procedure. For instance, the classic theory of Rationality posits that the DM uses a single preference ordering P to select the best element from any choice problem: $c(S) = \arg\max_P S$ for all $S \in \mathcal{P}(X)$.[7] The literature has also proposed interesting and plausible choice procedures that depart from Rationality. As a start, consider the following two theories that will serve as illustrative examples for Section 3. In the theory of *Limited Attention* by Masatlioglu, Nakajima and Ozbay (2012), a DM facing a choice problem S maximizes a preference ordering P over a consideration set $\Gamma(S) \subseteq S$, with the restriction that consideration sets don't change when removing ignored alternatives:

$$\text{(1a)} \qquad c(S) = \arg\max_P \Gamma(S), \text{ for all } S \in \mathcal{P}(X), \text{ and}$$

$$\text{(1b)} \qquad \Gamma(S) \subseteq T \subseteq S \Rightarrow \Gamma(T) = \Gamma(S), \text{ for all } S, T \in \mathcal{P}(X).$$

In Manzini and Mariotti (2007)'s theory of *Shortlisting*, the DM makes a shortlist of undominated options using an asymmetric relation P_1. She then chooses the undominated element in the shortlist according to an asymmetric preference relation P_2:

$$\text{(2)} \qquad \{c(S)\} = \max(\max(S, P_1), P_2), \text{ for all } S \in \mathcal{P}(X).^8$$

This paper develops a methodology for characterizing the testable implications of choice theories. Importantly, only the DM's choices, not her thought process or choice method, are observable. Each theory $\mathcal{T}$ yields a collection $\mathcal{C}(\mathcal{T})$ of possible choice functions. In the presence of limited data, one observes the DM's choices only for problems in a *dataset* $\mathcal{D} \subseteq \mathcal{P}(X)$. An *observed choice function* $c_{obs} : \mathcal{D} \to X$ associates to each choice problem $S \in \mathcal{D}$ the alternative in S that the DM selected.

---

[7] We use the term *relation* to mean a (possibly incomplete or cyclic) binary relation, while we use the term *ordering* to mean a complete, asymmetric and transitive relation. For any relation P and any $S \subseteq X$, we denote $\arg\max_P S = \{x \in S \mid xPy, \ \forall y \in S \setminus \{x\}\}$.

[8] Following Manzini and Mariotti (2007)'s notation, $\max(S, R) = \{x \in S \mid \nexists y \in S \text{ s.t. } yRx\}$. Their notation $\{c(S)\}$ requires the undominated set to be a singleton.

3. Testing for Consistency Using Limited Data

We aim to understand when observed choices are consistent with (that is, do not refute) a given theory. This means that the theory must yield at least one choice function that coincides with c_{obs} on $\mathcal{D}$, guaranteeing the ability to make coherent predictions for unobserved choice problems.

Definition 1 (Consistency) *An observed choice function $c_{obs} : \mathcal{D} \to X$ is consistent with a theory $\mathcal{T}$ if there is $c \in \mathcal{C}(\mathcal{T})$ such that $c_{obs}(S) = c(S)$ for every $S \in \mathcal{D}$.*

Definition 2 (Prediction) *The set of predictions for $S \notin \mathcal{D}$ under a theory $\mathcal{T}$ is given by $\{c(S) \mid c \in \mathcal{C}(\mathcal{T})$ and c coincides with c_{obs} on $\mathcal{D}\}$.*

Addressing the question of out-of-sample predictions thus reduces to identifying the testable implications of a theory on limited data: x is a possible choice for $S \notin \mathcal{D}$ if and only if the expanded observed choice function, derived from c_{obs} by adding the counterfactual that the DM picks x from S, is consistent with $\mathcal{T}$.

The recent literature on bounded rationality has overlooked a potential pitfall one should keep in mind in the presence of limited data. To test Shortlisting with limited data, Manzini and Mariotti (Corollary 1, 2007) study when there exist asymmetric relations P_1, P_2 such that (2) holds for $S \in \mathcal{D}$. Manzini and Mariotti (Definition 4, 2012) take a similar approach for their theory of Categorize-Then-Choose. To test Limited Attention, Tyson (2013) seeks conditions guaranteeing the existence of an ordering P and a consideration set mapping defined on $\mathcal{D}$ such that (1a) and (1b) hold for $S, T \in \mathcal{D}$. In other words, the theory's conditions describing how choices emerge are checked only over *observed* problems. Such an approach may seem natural at first. Taking Rationality as a benchmark, if there is an ordering P such that $c_{obs}(S)$ is the P-maximal element for all $S \in \mathcal{D}$, then it is trivial to extend c_{obs} to a rational choice function c by letting $c(S)$ be the P-maximal element for all $S \in \mathcal{P}(X)$. In general, however, such an approach may yield 'false positives,' as it may be impossible to extend observed choices into a complete choice function under the theory. This extensibility issue affects prevalent theories of bounded rationality, leading to a potentially dangerous methodological pitfall.

The datasets below illustrate the possible issues that may arise. According to the prior literature's approach, c_{obs1} should be consistent with Limited Attention, and

c_{obs2} should be consistent with Shortlisting.[9] In fact, neither is true.

S	ae	ef	abd	ade	bde	bef
$c_{obs1}(S)$	e	f	d	a	b	e

S	ab	bd	ad	$abde$	$abdf$	$abdg$
$c_{obs2}(S)$	a	b	d	a	b	d

Consider Limited Attention. Suppose, by contradiction, that some Γ satisfying (1b) and some ordering P generate an extension of c_{obs1} under the theory. By (1b), d must be considered in $\{a, d, e\}$, since its removal changes the choice. As a is chosen from $\{a, d, e\}$, we learn aPd. Similarly, we conclude $b \in \Gamma(\{b, e, f\})$ and ePb. Now consider the out-of-sample problem $\{b, d\}$. The ranking aPd implies $a \notin \Gamma(\{a, b, d\})$, thus (1b) requires $\Gamma(\{b, d\}) = \Gamma(\{a, b, d\})$. At the same time, the ranking ePb implies $e \notin \Gamma(\{b, d, e\})$, thus (1b) also requires $\Gamma(\{b, d\}) = \Gamma(\{b, d, e\})$. This is impossible, as the choices from $\{a, b, d\}$ and $\{b, d, e\}$ differ. The problem here, quite simply, is that a mapping Γ satisfying the required property over $\mathcal{D}$ need not satisfy it elsewhere.

Next take Shortlisting. Suppose, by contradiction, that some asymmetric relations P_1 and P_2 yield an extension of c_{obs2} under the theory. As a is chosen from $\{a, b, d, e\}$, it must be P_1-undominated in $\{a, b, d, e\}$ and subsets thereof. Thus the choice of d from $\{a, d\}$ implies dP_2a. Symmetric reasoning for b and d yields the preference cycle $aP_2bP_2dP_2a$, and none of these elements can P_1-dominate one of the two others. But then choice cannot arise from $\max(\max(S, P_1), P_2)$ for the out-of-sample problem $S = \{a, b, d\}$. The data requires P_2 to be cyclic over $\{a, b, d\}$, and yet prevents P_1 from eliminating any from the shortlist. The problem here is that not all combinations of asymmetric P_1, P_2 are valid inputs to the theory, and checking validity requires thinking about both observed and unobserved choice problems.

4. Building Insights: Testing Limited Attention

Now that we have the proper notion of consistency, we can turn our attention to developing a methodology for testing. To build insight for the ideas to come, we begin by examining how the classic, revealed-preference approach can be adapted to characterize Limited Attention.

Following Samuelson (1948)'s study of Rationality, an option x is revealed preferred to an alternative y if there exists a choice problem where y is available but

[9]Indeed, (1a) and (1b) hold for $S, T \in \mathcal{D}$ using the ordering P defined by $aPdPePbPf$, with $\Gamma(S)$ given by $c_{obs1}(S)$ and its P-lower counter set for $S \in \mathcal{D}$; moreover, (2) holds for $S \in \mathcal{D}$ using P_1 given by eP_1d, fP_1a and gP_1b, and P_2 given by $aP_2bP_2dP_2a$ and xP_2y for $x \in \{a, b, d\}$, $y \in \{e, f, g\}$.

the DM picks x instead. Clearly, consistency with Rationality requires the revealed preference to be acyclic. Furthermore, the data cannot reveal any more critical information about the DM's preference, since acyclicity of the revealed preference is also sufficient for consistency.

Now consider Limited Attention. Observing the DM pick x from S does not imply that she prefers x over another option y in S, because she may have overlooked y. Option x is only revealed preferred to alternatives in her consideration set at S, which itself must be inferred from the choice data. What then is all the critical information about preferences that can be gleaned from observed choices?

Masatlioglu, Nakajima and Ozbay (2012) provide an answer for full data sets (that is, for $\mathcal{D} = \mathcal{P}(X)$). Consistency with Limited Attention is equivalent to acyclicity of the following revealed preference: the DM prefers x over $z \in S \setminus \{x\}$ if she picks x from S but not from $S \setminus \{z\}$.[10] This result does not extend to limited data. Revealed preference restrictions can also arise from observing choice problems that are not related by dropping a single alternative; and while these restrictions are redundant in case of full data, they may become critical with limited data.

However, the argument underlying Masatlioglu et al's (2012) revealed preference readily extends to more general IIA violations. Indeed, if the choice from T is available but not chosen from $S \subset T$, then the DM must have considered at least one alternative in $T \setminus S$ when choosing from T. Otherwise, (1b) would require $\Gamma(T) = \Gamma(S)$, contradicting that the observed choices differ. The IIA violation thus informs the modeler that there exists $z \in T \setminus S$ such that $c_{obs}(T)Pz$. More subtly, any violation of the Weak Axiom of Revealed Preference (WARP) reveals some critical information about the DM's preference:

$$\text{(3)} \qquad \begin{array}{l} \text{For all } S, T \in \mathcal{D} \text{ with } c_{obs}(S) \neq c_{obs}(T) \text{ and } c_{obs}(S), c_{obs}(T) \in S \cap T : \\ c_{obs}(S)Pz \text{ for some } z \in S \setminus T \text{ or } c_{obs}(T)Pz' \text{ for some } z' \in T \setminus S. \end{array}$$

Otherwise, the choice from $S \cap T$ would be ill-defined by a similar argument as above: $\Gamma(S \cap T)$ would have to be identical to both $\Gamma(S)$ and $\Gamma(T)$, contradicting that the observed choices differ. As IIA violations are special types of WARP violations where the sets are related by inclusion, the revealed preference restrictions inferred from IIA violations are encompassed by (3).

[10]The DM must pay attention to z in S, as otherwise, her attention set and thus her choice would be the same for S and $S \setminus \{z\}$.

To summarize the above discussion, we have learned that if choices are consistent with Limited Attention, then the DM's preference must satisfy all the restrictions that observed choices reveal in (3). Hence there must exist an acyclic relation satisfying (3). Could it be that we have missed other critical restrictions? The answer is no, as the existence of such an acyclic relation also guarantees consistency with the theory.

Proposition 1 *Observed choices $c_{obs} : \mathcal{D} \to X$ are consistent with Limited Attention if and only if there exists an acyclic relation P satisfying (3).*

Proposition 1 shows how the revealed preference approach to testing can be extended beyond Rationality, once we recognize that data can reveal restrictions about the DM's preference that are more complex than simple comparisons between two alternatives (as opposed to Rationality, or Limited Attention in the case of full data). We will see in Section 6 how this insight allows us to capture the empirical content of many other theories beyond Limited Attention.

Figuring out whether a set of restrictions can be met by some acyclic relation – what we will call *acyclic satisfiability* in the next section – seems harder than testing SARP in cases such as (3), where the restriction does not immediately tell you which comparison holds. The next section explores to what extent this intuition is correct. It turns out that it is sometimes, but not always true; and that we can cleanly distinguish which is which.

5. Acyclic Satisfiability, and How to Check It?

In this section, we introduce the notion of acyclic satisfiability, a natural extension of SARP that captures the empirical content of many theories of choice (see Sections 4 and 6). We identify cases for which acyclic satisfiability can be tested in the same way as SARP, as well as cases for which testing is much more challenging. The section includes a first application of these results to Limited Attention, as a follow-up to Section 4. Several other applications are provided in Section 6.

5.1 Formalizing Restrictions and Acyclic Satisfiability

For any $x, y \in X$, define the function $1_{(x,y)}$ that takes as input a strict relation P and tests whether the *simple comparison* xPy holds for that relation. In other words, the function outputs 'true' if x is ranked above y according to P, and 'false' otherwise. A simple comparison, if satisfied, thus provides definite information about how two elements compare. More generally, a *restriction* takes as input a strict relation P, and outputs 'true' or 'false' based on logical conjunctions ('and', $\wedge$) and/or disjunctions ('or', $\vee$) of simple comparisons under P. Thus a simple comparison is a special case, and may alternatively be called a simple restriction. There may be different ways a restriction with disjunctions can be satisfied, corresponding to the different ways of selecting which of the simple comparisons involved in the disjunctions are true.

Given an option x, we say that a restriction pertains to the *lower-contour set (LCS) of* x if all the comparisons in it take the form $1_{(x,\cdot)}$. As is well known, any logical formula can be written in disjunctive normal form, i.e., as a disjunction of conjunctions. Thus a restriction pertaining to the LCS of x can be described by a family of sets $\Sigma \subseteq \mathcal{P}(X)$ and denoted $1_{(x,\Sigma)}$. A relation P satisfies $1_{(x,\Sigma)}$ if and only if some set in Σ belongs to the P-lower contour set of x.[11] For notational convenience, we assume throughout that LCS restrictions are given using the notation $1_{(x,\Sigma)}$. Similar definitions apply for upper-contour set (UCS) restrictions, which we denote $1_{(\Sigma,x)}$.

Consider, for instance, the revealed preference restrictions arising from Limited Attention. The data $c_{obs}(\{a,b,d\}) = b$ and $c_{obs}(\{b,d,e\}) = d$ creates a revealed preference restriction $1_{(b,a)} \vee 1_{(d,e)}$, that is, "b is preferred to a, or d is preferred to e." More generally, any WARP violation generates a revealed preference restriction that is the disjunction of simple comparisons (see (3)). However, IIA violations exclusively generate LCS restrictions; for instance, $c_{obs}(\{w,x\}) = w$ and $c_{obs}(\{w,x,y,z\}) = x$ yields the restriction $1_{(x,\{\{y\},\{z\}\})}$, meaning "x is preferred to y or z." Many other examples of restrictions will be encountered in Section 6.

A collection of restrictions is *acyclically satisfiable* if there exists an acyclic relation satisfying them. Note that simple comparisons pin down a (typically incomplete) relation, and in this case acyclic satisfiability just boils down to this relation being acyclic. As illustrated in Proposition 1 and the results of Section 6, the empirical content of various theories is naturally captured through acyclic satisfiability of restrictions summarizing key information revealed by choices. Rationality being equivalent to SARP is now understood as just one instance of this broader approach.

[11]Slightly more generally, an LCS restriction may also *pertain to a set* $T \subseteq X$: a relation P satisfies $1_{(T,\Sigma)}$ if and only if some member of Σ is contained in the P-lower contour set of every $x \in T$ (and similarly for a UCS restriction pertaining to a set). One can immediately see that all our enumeration-related results (Lemmas 1-2, and Propositions 2 and 9) also hold for LCS (UCS) restrictions pertaining to sets. We avoid this more cumbersome notation in the text, since our only example of an application with LCS restrictions pertaining to sets appears in our study of 'good-enough' heuristics in Barberà et al (2018).

5.2 Testing Acyclic Satisfiability

Recall our discussion in Section 1 regarding how to test SARP. For each step $k \geq 1$, simply look for, and remove from the set, an option x_k that is not ranked above any remaining alternatives according to Samuelson's revealed preference. It is possible to *enumerate* all of X in this manner if and only if SARP holds (i.e., Samuelson's revealed preference is acyclic). Intuitively, one is trying to construct a possible preference ordering for the DM from the bottom up, by iteratively identifying a candidate for the worst remaining option. Our first observation, formalized in Section 5.2.1, is that acyclic satisfiability can be tested essentially in the same, tractable way when facing only LCS restrictions (or only UCS restrictions). Indeed, while lower-contour set restrictions may not precisely identify *which* alternative must be ranked below an option, they do reveal that *some* alternative must be ranked below it, which is all we need to know to rule it out as a candidate-worst option.

By contrast, Section 5.2.2 shows how testing acyclic satisfiability can become much harder to perform when considering wider classes of restrictions.[12] This will prove useful in applications, both to check that it is impossible to reduce the identified restrictions into all-LCS or all-UCS restrictions, and to ascertain that there isn't nonetheless an alternative, simple procedure to systematically test that theory.

5.2.1 Testing Comparable to SARP with LCS (UCS) Restrictions

The ability to apply the procedure used for checking SARP when considering a collection of LCS restrictions rests on two observations, presented as lemmas. The first lemma echoes the intuition above.

Lemma 1 *Let X be a set of options and $\mathcal{R}$ be a set of LCS restrictions defined on X. If $\mathcal{R}$ is acyclically satisfiable, then there exists an option $x \in X$ such that no restriction in $\mathcal{R}$ pertains to the lower-contour set of x.*

Indeed, any acyclic relation satisfying $\mathcal{R}$ can be completed into an ordering satisfying $\mathcal{R}$, and x may be taken to be its minimal element. This provides a first, simple necessary condition for acyclic satisfiability: traverse elements of X to find one that

[12]Testing acyclic satisfiability can be seen as an extension of the topological sort problem in computer science. Some extensions have been studied in problems of job-scheduling with waiting conditions; see Möhring et al. (2004) who provide a fast scheduling algorithm given conditions "job i comes before at least one job in a set J," which corresponds to a special case of $1_{(x,\Sigma)}$ with every $S \in \Sigma$ a singleton. They show scheduling is NP-hard for the generalization "some job in a set I comes before some job in a set J"; we show the problem is already NP-hard with simpler restrictions.

does not appear at the top of a restriction. Let x_1 be an element with this property (if one exists). Because x_1 will be treated as the bottom element of the ordering, any restriction $1_{(x,\Sigma)}$ such that $\{x_1\} \in \Sigma$ is now satisfied, and may be eliminated; all other restrictions $1_{(x,\Sigma)}$ simplify to $1_{(x,\Sigma')}$, where $\Sigma' = \{S \setminus \{x_1\} \mid S \in \Sigma\}$.[13] Let $\mathcal{R}_1$ be this reduced set of LCS restrictions over $X \setminus \{x_1\}$.

Lemma 2 *Let x_1 satisfy the property of Lemma 1. Then $\mathcal{R}$ (defined over X) is acyclically satisfiable if and only if $\mathcal{R}_1$ (defined over $X \setminus \{x_1\}$) is acyclically satisfiable.*

Necessity obtains by considering the restriction of the acyclic relation satisfying $\mathcal{R}$ to the set $X \setminus \{x_1\}$. Sufficiency obtains by augmenting the acyclic relation satisfying $\mathcal{R}_1$ by placing x_1 at the bottom of any pairwise comparison.

Lemmas 1-2 hold independently of the set X and the set of LCS restrictions $\mathcal{R}$, so the reasoning may be iterated. The Lemmas thus provide a conceptual roadmap for defining the *enumeration procedure for LCS restrictions*. The first step follows as in Lemma 1, while Lemma 2 shows how to iterate the procedure. In each step k, if there has been no failure to find a candidate-worst element thus far, then we treat $x_1, \ldots, x_{k-1}$ as if they are ranked below all remaining elements. Thus, we may restrict attention to a simplified set of restrictions $\mathcal{R}_{k-1}$, where $x_1, \ldots, x_{k-1}$ have been eliminated.[14] Writing $\mathcal{R}_0 = \mathcal{R}$, the enumeration procedure can be stated as follows.

> *Step k, for $k \geq 1$:* Look for an element $x_k \in X \setminus \{x_1, \ldots, x_{k-1}\}$ that does not appear at the top of any LCS restriction in $\mathcal{R}_{k-1}$. Continue to the next step if and only if such an element is found.

The enumeration procedure *fails* if in some step we cannot find a candidate x_k for the worst element; but if one can enumerate all of X in this way, then the enumeration procedure *succeeds*. Importantly, Lemma 2 ensures path independence: even if there are multiple candidates for the worst element in a step, a different selection among these would not convert failure of the procedure to success, or vice versa; that is, success and failure are definitive outcomes. The above reasoning shows that success of the enumeration procedure is a necessary condition for $\mathcal{R}$ to be acyclically satisfiable.

[13]Without using the notation $1_{(x,\Sigma)}$ for restrictions, this is equivalent to replacing functions of the form $1_{(y,x_1)}$ for any $y \in X$ by the logical value 'true'.

[14]That is, each $1_{(x,\Sigma)} \in \mathcal{R}$ simplifies to $1_{(x,\Sigma')}$, where $\Sigma' = \{S \setminus \{x_1, \ldots, x_{k-1}\} \mid S \in \Sigma\}$; and any restriction $1_{(x,\Sigma)}$ such that $S \subseteq \{x_1, \ldots, x_{k-1}\}$ for some $S \in \Sigma$ is eliminated entirely. Equivalently, note we can just ignore in step k all restrictions $1_{(x,\Sigma)}$ such that $S \subseteq \{x_1, \ldots, x_{k-1}\}$ for some $S \in \Sigma$.

Vice versa, ranking options in opposite order from a successful enumeration will satisfy $\mathcal{R}$ by construction. We have thus shown the following.

Proposition 2 *A set of LCS restrictions $\mathcal{R}$ is acyclically satisfiable if and only if the enumeration procedure succeeds.*

The procedure can also be used to check acyclic satisfiability of a set of UCS restrictions. The only difference is that one seeks candidate *maximal* options in each step, i.e., options that do not appear at the bottom of any remaining UCS restrictions.

Remark 1 *Proposition 2 is also helpful for identification (e.g., preference identification). If a set of LCS restrictions $\mathcal{R}$ is acyclically satisfiable and* every *acyclic relation P satisfying $\mathcal{R}$ has xPy, then we say that a ranking of x over y is* identified. *Clearly, one can test for such identification by adding the opposite restriction $1_{(y,x)}$ to the restrictions characterizing consistency. If the augmented restrictions are acyclically satisfiable, then it is possible to have the opposite ranking; but if acyclic satisfiability fails, then a ranking of x over y is identified. Conveniently, if the original restrictions are of the LCS (UCS) type, then adding a simple comparison does not change this. Thus identification is testable through enumeration in such cases.*

How easy is testing LCS (or UCS) restrictions through enumeration? The procedure requires at most $|X| - 1$ iterations, so testing is tractable if finding a candidate for minimal element is tractable in each iteration.[15] Testing a theory using data also requires constructing the restrictions in the first place, so one should confirm that doing so is easy. We check these steps for all theories identified as tractably testable in this paper. In those cases, testing can be carried out in a number of steps that is at most polynomial in the size of the dataset. As a benefit, computers can carry out the testing quickly for any dataset (and for relatively small datasets, answers can systematically be found by hand).

Consider, for instance, the theory of Limited Attention. The restrictions for Limited Attention, captured in (3), are in general neither LCS nor UCS. However, notice that only the LCS restrictions matter whenever $\mathcal{D}$ is closed under intersection (or at least contains the intersection of any two choice problems causing a WARP vio-

[15]In fact, for many theories (including Rationality), a restriction pertains to the lower-contour set of an option only if the DM chose it in some problem. In this case, options that are never chosen can be randomly enumerated in a preliminary step of the enumeration procedure, and thus the cardinality of the image of c_{obs} bounds the number of nontrivial steps in the procedure.

lation).[16] Thus, we may conclude from Proposition 2 that testing consistency with Limited Attention can be done in a way similar to checking SARP for datasets satisfying this intersection property. How easy is it to test Limited Attention for a general dataset? We defer this question until after we study non-LCS (and non-UCS) restrictions in the next subsection.

5.2.2 Hard to Test Otherwise

As we are expanding the realm of the classic, revealed-preference testing methodology beyond Rationality, it is natural to ask whether we have reached the frontier of 'tractable testing.' Suppose we are facing a set of restrictions that is more complex than those covered by Section 5.2.1. Is there some clever extension of the enumeration procedure, or an entirely different algorithm, that would make acyclic satisfiability easy to test for more general restrictions? The answer is essentially negative, in a sense we now make precise.

So far we analyzed cases where all restrictions pertain to lower contour sets, or all restrictions pertain to upper contour sets. To make the negative result most striking, say a collection of restrictions *is a mixed set of binary restrictions* if each restriction takes either the form $1_{(x,\{\{y\},\{z\}\})}$ ("x is better than y or z") or the form $1_{(\{\{y\},\{z\}\},x)}$ ("x is worse than y or z") for some $x, y, z \in X$. We view such collections as a minimal departure from those considered thus far, as each restriction must belong to either the LCS or UCS class, and can be more complicated than a simple comparison by at most one disjunction ("or") between options. We show that this small generalization of the problem becomes NP-hard; as a corollary, testing acyclic satisfiability with more permissive classes of restrictions is NP-hard as well.[17]

Proposition 3 *The problem of checking acyclic satisfiability for mixed sets of binary restrictions is NP-hard.*

If $P \neq NP$, as widely believed, no algorithm solves every instance of an NP-hard problem in polynomial time. To prove Proposition 3, we show that every instance of

[16]If S and S' cause a WARP violation, then $S \cap S'$ causes an IIA violation with S or S'. Suppose it occurs with S. Then $c_{obs}(S)$ must be preferred to some element of $S \setminus S'$, satisfying the 'or' condition from the WARP violation between S and S'.

[17]P is the set of problems solvable in polynomial time; NP is the set of problems that may or may not be solvable in polynomial time, but for which any conjectured solution can be checked in polynomial time. A problem is NP-hard if solving it is at least as complex as solving the most difficult problems in NP. Finding a polynomial-time solution for some NP-hard problem would have the important implication that $P = NP$. No one has found such a solution thus far.

SAT3 (a classic NP-hard problem[18]) has a polynomial-time reduction to an equivalent problem of checking acyclic satisfiability of mixed sets of binary restrictions.

Proposition 3 illustrates how small departures from only LCS or only UCS restrictions can make acyclic satisfiability much harder to test. Showing that the empirical content of a theory involves such restrictions is suggestive that testing is likely to be NP-hard, but reaching that conclusion requires a formal argument that the restrictions cannot be significantly simplified. Proposition 3 can be used to prove this formally: one must show that for any mixed set of binary restrictions $\mathcal{R}$, there is choice data (constructed in polynomial time given $\mathcal{R}$) such that the theory is consistent with these choices if and only if $\mathcal{R}$ is acyclically satisfiable.

Our first of several applications of Proposition 3 is to Limited Attention. Proposition 4 in the Appendix shows that testing consistency with this theory is NP-hard. Recall that Limited Attention generates LCS restrictions from IIA violations, but doesn't exactly generate UCS restrictions. For intuition on how Proposition 4 is proved, recall that WARP violations can give rise to restrictions of the form aPb or $a'Pb'$, with $a \neq a'$ and $b \neq b'$. Additional data, for example $c_{obs}(\{b,e\}) = e$ and $c_{obs}(\{b,d,e\}) = b$, could reveal that b is preferred to d. Similar data could reveal that b' is also preferred to d. Then, the restriction aPb or $a'Pb'$ implies the binary UCS restriction aPd or $a'Pd$. This lights a path towards an application of Proposition 3. The proof finds easy-to-construct classes of datasets for which testing consistency with Limited Attention is equivalent to testing acyclic satisfiability of mixed sets of binary restrictions.

Being NP-hard to test means that there exist datasets for which testing consistency is intractable. A first implication is that it may be possible to prove a theory is NP-hard even with a partial understanding of its testable implications. Indeed, it suffices to know how to test consistency within some class of datasets, and then show that testing acyclic satisfiability of mixed sets of binary restrictions reduces to testing consistency in those cases. A second implication is that there may be other classes of datasets for which testing remains tractable; and that these datasets can be identified by when the restrictions take only the LCS, or only the UCS, form. For instance, we argued in the previous subsection that the enumeration procedure applies to Limited Attention when the dataset has the intersection property. As noted later, a similar

[18] Given any set of 'clauses' that are disjunctions of three 'literals' (variables or their negations), the question is whether there is a truth assignment for the variables that makes all clauses true.

conclusion applies to some theories of categorization and rationalization, which are NP-hard to test in general, but easy to test by enumeration if the dataset includes, for instance, all pairs of observed choices.

6. Applications

Beyond the theory of Limited Attention that we have already covered, this section illustrates the surprising versatility of our approach. Restrictions can pertain to orderings other than preference (e.g. salience or priority), and to relations that may be incomplete (as in problems of just-noticeable differences). They can capture satisficing instead of maximizing behavior. The methodology applies to settings where X is not finite (e.g. consumer theory). It can apply when preferences are weak. It can apply when choices are observed only coarsely, as well as when the choice data is stochastic. It even provides useful insight in some classic problems of interactive decision making with rational agents. Details for the applications in Sections 6.1-6.4 are found in Appendix B, while details for Section 6.5 appear in the Online Appendix. Section 6.6 discusses applications carried out in subsequent works.

6.1 Choice Overload

A main alternative to condition (1b) when considering the maximization of a preference ordering over consideration sets is:

$$S \subset T \Rightarrow \Gamma(T) \cap S \subseteq \Gamma(S), \text{ for all } S, T \in \mathcal{P}(X). \tag{4}$$

In other words, if the DM considers an option in a choice problem, then she must consider it in smaller choice problems as well. This captures several heuristics a DM may use to deal with choice overload. Lleras, Masatlioglu, Nakajima and Ozbay (2017) impose the property directly, as it naturally encapsulates the idea that a DM is better at making choices from small sets than large ones. Cherepanov, Feddersen and Sandroni (2013) shows that property (4) characterizes the consideration sets arising from psychological rationalization, which comprise those elements that are top-ranked for some *rationale* (a binary relation).[19]

[19]Cherepanov et al (2013) also consider an extension permitting cyclic preferences; the resulting theory is observationally equivalent to Manzini and Mariotti (2012)'s categorization theory. While

Under (4), preference comparisons can be inferred from IIA violations. By definition, the DM pays attention to $c_{obs}(T)$ in every set $S \subseteq T$ such that $c_{obs}(T) \in S$. Hence an IIA violation, $c_{obs}(S) \neq c_{obs}(T)$, reveals that the DM prefers $c_{obs}(S)$ over $c_{obs}(T)$. This is a simple preference comparison. Are there other, perhaps more complex restrictions to consider? Not in this case. Proposition 5 in Appendix B shows that acyclic satisfiability of these restrictions guarantees consistency. Hence the theory is testable by enumeration. Interestingly, the same revealed preference identified by the above authors for full datasets happens to also capture the testable implications in limited datasets.

6.2 Reference Dependence

Reference effects provide an important context in which multiple preferences play a role. Without further restrictions, any choices can be explained by the maximization of a reference-dependent preference. Rubinstein and Salant (2006B)'s theory of *Triggered Rationality*[20] structures reference-point formation, by positing a salience ordering $\succ_\sigma$ over the alternatives and a collection of preference orderings $\{P_x\}_{x \in X}$. The most salient element in a set is the DM's reference point, anchoring the preference P_x maximized.

We apply our methodology to the salience ordering. Suppose we conjecture that an option $x \in S$ is the most salient alternative in S. This means x is the most salient alternative in all $R \subseteq S$ in which it is contained. In particular, the DM must apply the preference P_x in all those choice sets, and we can consider the resulting, Samuelson-revealed preference from those that are observed. If that revealed preference is cyclic, then our original conjecture must have been wrong: that is, we learn that some element in S is more salient than x. We show in Proposition 6 of Appendix B that acyclic satisfiability of these restrictions is both necessary and sufficient for consistency with Triggered Rationality. These restrictions are of the UCS type, and can be tractably tested by enumeration.

---

our methodology cannot be applied to the DM's preference in this case, de Clippel and Rozen (2019c) show that any consideration set mapping satisfying (4) generates a natural ranking of alternatives. They characterize the testable implications of these theories through acyclic satisfiability of restrictions on this ranking. These are UCS restrictions if the data is rich enough (e.g. includes all pairs), in which case consistency is testable by enumeration. Otherwise, they show testing is NP-hard as a corollary of Proposition 3.

[20] See also Rubinstein and Salant (2006A, Example 4).

6.3 *Shortlisting*

Noteworthily, our methodology applies to acyclic relations, not just orderings. This proves useful in studying Manzini and Mariotti (2007)'s theory of *Shortlisting*. Observe that the shortlisting relation P_1 must be acyclic for choices to be well defined, else the DM's shortlist would be empty in the choice problem corresponding to the cycle. However, P_1 may still be incomplete; indeed, if it were complete then the DM's observed choices would be consistent with rationality.

It is easy to see that UCS restrictions on the shortlisting relation P_1 are generated by certain choice configurations. For instance, suppose $c_{obs}(\{a, x\}) = x$ but $c_{obs}(\{a, x, y, z\}) = a$. Since a is not eliminated by P_1 in $\{a, x, y, z\}$, it is also not eliminated in $\{a, x\}$. Hence x, which is chosen from the latter set, must be preferred to a. As a is chosen from $\{a, x, y, z\}$ instead, y or z must eliminate x from consideration under P_1. However, non-UCS restrictions on P_1 can be generated as well. Suppose we also observe $c_{obs}(\{b, y, z\}) = y$ and $c_{obs}(\{d, y, z\}) = z$. Hence y and z are incomparable under P_1, and so they must be comparable under P_2. To explain this data, y and z cannot *both* survive the shortlist in *both* $\{b, y, z\}$ and $\{d, y, z\}$. Thus we learn the following restriction for P_1: P_1 ranks d above y, or P_1 ranks b above z. This is not quite a LCS restriction, since the elements at the top of the simple comparisons are not the same. Further data, though, may require a third alternative e to eliminate both b and d, lighting a path towards an application of Proposition 3.

Indeed, Proposition 7 in Appendix B shows that testing consistency with Shortlisting is NP-hard. To show this, a full characterization of Shortlisting's empirical content is not required. Indeed, it suffices to show there is subclass of datasets on which testing consistency of the theory amounts to testing acyclic satisfiability of a mixed set of binary restrictions. The result happens to also hold for *Order Shortlisting*, which requires the DM's preference to be an ordering.

6.4 *Undominated Alternatives*

Consider a DM who chooses the set of undominated alternatives according to an acyclic (possibly incomplete) relation. This choice procedure, which yields an observed choice correspondence, has several interesting interpretations. It has the same testable implications as Masatlioglu and Nakajima (2013)'s Markovian Choice by Iterative Search, when the starting point of the DM's search process is unobserved.

It corresponds to the shortlist in Manzini and Mariotti (2012)'s Shortlisting, which in some cases the modeler observes (e.g., alternatives considered by a committee, candidates interviewed for a position). It also relates to theories of just-noticeable differences. Following Luce (1956) and ensuing works, insights from psychology that people have difficulty discerning differences in stimuli (e.g., the Weber-Fechner law)[21] may apply to preference comparisons. Consider a DM with a utility function u and a threshold function $\tau : X \times X \to \mathbb{R}_+$. She discerns y is preferable to x if $u(y) > u(x) + \tau(x, y)$ (i.e., when the utility difference is big enough given the options); and selects alternatives for which she cannot discern anything better. This theory, too, is equivalent to choosing undominated elements.[22]

While the possibility of choice correspondences departs somewhat from our original framework, our methodology accommodates the testing of these theories. Given an acyclic relation $\succ$, let $C_\succ : \mathcal{P}(X) \to \mathcal{P}(X)$ associate with each choice problem the set of $\succ$-undominated elements. Say that the observed choice correspondence C_{obs} is consistent with the theory if there is an acyclic relation $\succ$ such that $C_{obs}(S) = C_\succ(S)$ for all $S \in \mathcal{D}$. Now consider what we learn about $\succ$ from observed choices. Suppose x is not among the observed choices for some $S \in \mathcal{D}$ with $x \in S$. Then, there must be $y \in S$ such that $y \succ x$. However, the data may disqualify some y's from playing this role. This would be the case, for instance, if x is ever picked in the presence of y. Thus, we may more accurately conclude that there is some y in $S \setminus \bigcup_{\{T \mid x \in C_{obs}(T)\}} T$ that $\succ$-dominates x. Proposition 8 in Appendix B establishes that this is all the information that can be gleaned: these UCS restrictions are both necessary and sufficient for consistency.[23] Hence the above theories can be tested by enumeration.

[21]See Fechner (1860)'s seminal work and the large literature that follows.

[22]A perceived-preference cycle $x_1 \succ x_2 \succ x_n \succ x_{n+1} = x_1$ implies $\sum_{i=1}^{n} \tau(x_{i+1}, x_i) < 0$, yet $\tau(\cdot, \cdot) \geq 0$. Conversely, for an acyclic $\succ$, take $u : X \to \mathbb{R}$ respecting $\succ$-comparisons, with $\tau(a, b) = 0$ if $a \succ b$ and $\tau(a, b) = \max_y u(y) - \min_x u(x)$ otherwise. Then $y \succ x$ iff $u(y) > u(x) + \tau(x, y)$. This theory includes Luce (1956)'s theory with constant thresholds as well as generalizations proposed in the large ensuing literature, which examines how the threshold varies with the alternatives compared, and the impact on perceived preferences.

[23]Notice that acyclic satisfiability does not require the relation satisfying the restrictions to be complete. Of course, the existence of an acyclic relation satisfying the restrictions, also implies the existence of an ordering satisfying the restrictions. However, that ordering need not yield the observed choices, generating for instance a single-valued choice function instead of a correspondence. Thus allowing the relation to be incomplete in the definition of acyclic satisfiability is useful indeed.

6.5 Indifferences

Most preference-based theories of bounded rationality assume strict preferences. As such, we have focused on restrictions pertaining to orderings, which are strict. Though it may be slightly precocious to advance our testing methodology to accommodate indifferences prior to further developments in bounded rationality, the Online Appendix nonetheless offers some steps in this direction. The methodology we advocate is the same: (a) identify necessary restrictions; (b) ensure one has gained a full understanding of the theory by checking that acyclic satisfiability of these restrictions is sufficient for consistency; (c) evaluate whether testing acyclic satisfiability can be done by enumeration like SARP, or is necessarily complex. However, interesting new features arise.

Data will typically reveal some restrictions that must hold strictly, while others may hold weakly. It may be tempting to apply the results of Section 5 by treating the weak ones as strict. This, however, can yield false negatives: $1_{(x,y)}$, $1_{(y,z)}$, and $1_{(z,x)}$ can be satisfied weakly (with x, y, z in one equivalence class) but not strictly. In settings with weak simple restrictions and strict LCS (UCS) restrictions, we show how testing can be performed tractably by enumerating *equivalence classes*.

As an illustration, we also show how this result applies to Choice Overload when the DM's preference may be a weak ordering. Characterizations in terms of acyclic satisfiability, and tractable tests by enumeration, are derived for consistency *à la* Afriat (observed choices belong to the choice set predicted under the theory), as well as consistency *à la* Richter (the set of observed choices coincides with the choice set predicted under the theory).[24] Interestingly, there is a new form of revealed (in)attention in theories where the DM maximizes a weak preference ordering over a consideration set. Witnessing the DM select both x and y in a choice problem reveals that she is indifferent between them. With Richter's consistency, seeing her select x but not y in a different problem reveals she did not consider y there. This may have important consequences for testing. For instance, adding the possibility of indifference to Choice Overload makes it subject to the extensibility issue highlighted in Section 3, and also generates revealed preference restrictions that are not simple. These added restrictions are nonetheless of the LCS type, so the theory remains testable by enumeration.

[24] See Richter (1966) and Afriat (1967) regarding these two notions for correspondences.

6.6 Further Applications (in Subsequent Works)

Our methodology is useful for other theories and contexts studied in separate works. We briefly discuss those applications, and which of Propositions 2 or 3 applies.

Barberà, de Clippel, Neme and Rozen (2018) study a variety of satisficing heuristics. For example, a DM may stop searching when her consideration set is large enough; she may be happy choosing a sufficiently well-ranked option (e.g., top two, or top quintile of a set); or she may first sample some alternatives to endogenously construct her satisficing threshold. They show the empirical content of these theories is captured by acyclic satisfiability of LCS restrictions on the DM's preference, and thus testable by enumeration. By contrast, they use Proposition 3 to show that it is NP-hard to test Sen (1993)'s theory of choosing the second best.

Applying our methodology to consumer demand data, de Clippel and Rozen (2018a) consider a consumer who fails to perfectly optimize due to misperception of prices or utility tradeoffs. They show that testing consistency of consumer-demand data with misperception amounts to acyclic satisfiability of LCS restrictions on the consumer's preference over bundles. Hence it is easy to test, by Proposition 2.

In both of the above papers, the DM (consumer) may face the same choice problem multiple times. Framing, perception and random events can lead to different choices for the same problem, so that the data may take the form of a correspondence. These two papers allow for the data to be limited in two senses: beyond our standard sense that not all choice problems are observed, there may also be too few repeated observations per choice problem to construct reliable probability distributions. The approach in both works is to consider the support of choices but ignore sampling frequencies. However, as discussed later below, our methodology can also be useful for theories that use stochastic choice data.

While motivated by bounded rationality, our methodology also applies to other theories. de Clippel and Rozen (2018b) consider rational agents implementing some classic assignment/matching methods. The unobserved constraints from others' choices means even rational agents appear to violate SARP. The modeler observes final allocations as a function of some (collective or individual) endowments, but not preferences or other primitives of the problem. Under serial dictatorship, for instance, agents are ranked according to a power relation, and pick their best option among those unchosen by more powerful agents. The data generates LCS restrictions on the power relation; thus testing is tractable by Proposition 2. They show stability is also easy to test for an interesting class of many-to-one matching problems. However, the core of Shapley and Scarf (1974)'s housing market is shown to be NP-hard to test using Proposition 3.

For another application, say an observed choice correspondence C_{obs} is *coarsely rationalizable* if there is a preference ordering $\succ$ such that the best element from each choice problem $S \in \mathcal{D}$ belongs to $C_{obs}(S)$. Fishburn (1976) first introduced this idea (under the name 'representability'), and proved it is equivalent to the following axiom: for every family of choice sets, there must exist at least one option x such that $x \in C_{obs}(S)$ for each choice set S in the family with $x \in S$. Let us instead study this problem with our methodology. Observe that for any set $S \in \mathcal{D}$ and $x \notin C_{obs}(S)$, x cannot be preference-maximal in S (that is, $1_{(\{y\}_{y \in S \setminus \{x\}}, x)}$). Acyclic satisfiability of these UCS restrictions is not only necessary for 'coarse rationalizability', but also sufficient, simply by taking the preference to be any transitive completion O of an acyclic relation satisfying them: for any $S \in \mathcal{D}$, $x = \arg\max_S O$ must belong to $C_{obs}(S)$, else there would have been a restriction saying x cannot be maximal in S. Thus our approach provides a revealed-preference test of coarse rationalizability that is also tractable. We are aware of this additional application thanks to Hu, Li, Quah and Tang (2018), who introduce the interpretation in terms of coarse data, and show that the testable implications of both minimax regret preferences and correspondences arising from frames amount to the coarse rationalizability of cleverly constructed, auxiliary datasets. For instance, consider correspondences arising from frames, where the DM has a fixed set of preferences and the observed choices are the maximizers for each preference. For each $S \in \mathcal{D}$ and each $x \in C_{obs}(S)$, consider the auxiliary choices $C_{S,x}$ which equals C_{obs} except that $C_{S,x}(S) = \{x\}$. Having a fixed set of preferences applied to all the problems means that the original data arises from frames if and only if each auxiliary $C_{S,x}$ is coarsely rationalizable. Since our approach applies to coarse rationalizability, it applies to those theories as well.

Finally, Cattaneo, Ma, Masatlioglu and Suleymanov (2018) study stochastic choice data arising from a DM who maximizes a preference ordering over stochastic consideration sets. Their theory requires an intuitive property of monotonic attention, which they verify is satisfied by many recent contributions in the literature on stochastic attention. Stochastic choice data is shown to be consistent with their theory if and only if a revealed preference is acyclic. Following up on the ideas presented here, they uncover restrictions on the revealed preference that are more complex than UCS (or

LCS), and show that acyclic satisfiability of those restrictions captures the empirical content with limited data. It is not difficult to see that these restrictions boil down, in the special case of deterministic data, to those we uncovered in Section 4 when studying Limited Attention. They can thus be seen as a (far from trivial) extension of (3) for stochastic data. Given Proposition 4 in the Appendix, testing is a fortiori NP-hard. One can check, however, that the restrictions they uncover greatly simplify when the dataset is closed under intersection. The relevant restrictions become: for all $S \subset T$, if the DM is strictly less likely to pick a from S than from T, then there exists $b \in T \setminus S$ such that aRb. This defines a collection of LCS restrictions, and we conclude that, in this case, testing can tractably be done by enumeration.

References

Abreu, D, D. Pearce and E. Stacchetti. (1990), Toward a Theory of Discounted Repeated Games with Imperfect Monitoring, *Econometrica*, 58, 1041–1063.

Afriat, S. (1967), The Construction of a Utility Function from Demand Data, *International Economic Review*, 8, 67–77.

Barberà, S., G. de Clippel, A. Neme, and K. Rozen (2018), Good Enough, *mimeo.*

Cattaneo, M., X. Ma, Y. Masatlioglu, and E. Suleymanov (2018), A Random Attention Model, *mimeo.*

Cherepanov, V., T. Feddersen and A. Sandroni (2013), Rationalization, *Theoretical Economics*, 8, 775–800.

de Clippel, G., and K. Rozen (2018a), Consumer Theory with Misperceived Tastes, *mimeo.*

de Clippel, G., and K. Rozen (2018b), Testable Implications of Some Classic Assignment Methods, *mimeo.*

de Clippel, G., and K. Rozen (2018c), A Note on Categorization and Rationalization: Testable Implications with Limited Data, *mimeo.*

Fechner, G.T. (1860), Elemente der Psychophysik, Leipzig: Breitkopf und Härtel.

Fishburn, P. (1976), Representable Choice Functions, *Econometrica*, 44, 1033-1043.

Houthaker, H. (1950), Revealed Preferences and the Utility Function, *Economica*, 17, 159–74.

Hu, G., J. Li, J. Quah and R. Tang (2018), Coarse Revealed Preference, *mimeo.*

Lleras, J., Y. Masatlioglu, D. Nakajima, and E. Ozbay (2017), When More is Less: Choice by Limited Consideration, *Journal of Economic Theory*, 170, 70–85.

Luce, R. D. (1956), Semiorders and a Theory of Utility Discrimination, *Econometrica*, 24, 178–191.

Manzini, P., and M. Mariotti (2007), Sequentially Rationalizable Choice, *American Economic Review*, 97(5), 1824–1839.

MANZINI, P., AND M. MARIOTTI (2012), Categorize Then Choose: Boundedly Rational Choice and Welfare, *Journal of the European Economics Association*, 10(5), 1141–1165.

MASATLIOGLU, Y. AND D. NAKAJIMA (2013), Choice by Iterative Search, *Theoretical Economics*, 8, 701–728.

MASATLIOGLU, Y., D. NAKAJIMA AND E. OZBAY (2012), Revealed Attention, *American Economic Review*, 102(5), 2183–2205.

MÖHRING, R., M. SKUTELLA, AND F. STORK (2004), Scheduling With And/Or Precedence Constraints, *Siam Journal on Computing*, 33(2), 393–415.

MYERSON, R. (1979), Incentive-Compatibility and the Bargaining Problem, *Econometrica*, 47, 61–73.

NASH, J. (1955), Letter to the United States National Security Agency, available at https://www.nsa.gov/news-features/declassified-documents/nash-letters/assets/files/nash_letters1.pdf.

RICHTER, M. (1966), Revealed Preference Theory, *Econometrica*, 34, 635–645.

RUBINSTEIN, A. AND Y. SALANT (2006A), A Model of Choice from Lists, *Theoretical Economics*, 1, 3–17.

RUBINSTEIN, A. AND Y. SALANT (2006B), Two Comments on the Principle of Revealed Preference, *mimeo*.

SALANT, Y. AND A. RUBINSTEIN (2008), (A, f): Choice with Frames, *Review of Economics Studies*, 75, 1287–1296.

SAMUELSON, P. (1948), Consumption Theory in Terms of Revealed Preference, *Economica*, 15, 243–253.

SHAPLEY, L. AND H. SCARF (1974), On Cores and Indivisibility, *Journal of Mathematical Economics*, 1, 23–27.

TYSON, C. (2013), Behavioral Implications of Shortlisting Procedures, *Social Choice and Welfare*, 41, 941–963.

VARIAN, H. (1982), The Nonparametric Approach to Demand Analysis, *Econometrica*, 50, 945–973.

APPENDIX A: Proofs for Sections 4-5

The proofs of Lemmas 1-2 and Proposition 2 are provided in the text.

Proof of Proposition 1: *Consistency for Limited Attention*

The argument for necessity appears in the main text. For sufficiency, suppose an acyclic relation satisfying (3) exists, and let P be a transitive completion (so P still satisfies (3)). We define $\Gamma : \mathcal{P}(X) \to \mathcal{P}(X)$ as follows. For $S \in \mathcal{D}$, $\Gamma(S) = \{c_{obs}(S)\} \cup \{x \in S | c_{obs}(S) P x\}$; for $S \notin \mathcal{D}$,

$$\Gamma(S) = \begin{cases} \Gamma(T) & \text{if } S \subseteq T,\ T \in \mathcal{D},\ \text{and } \Gamma(T) \subseteq S \\ S & \text{otherwise.} \end{cases}$$

Clearly $\Gamma(S) \neq \emptyset$ for any $S \in \mathcal{P}(X)$ and the P-maximal element in $\Gamma(S)$ is $c_{obs}(S)$ for any $S \in \mathcal{D}$. We show Γ is well-defined and satisfies (1b).

Suppose by contradiction that Γ is not well-defined for some S. This means that for some $S \notin \mathcal{D}$, there exist $T, T' \in \mathcal{D}$ such that $S \subseteq T \cap T'$ with $\Gamma(T) \cup \Gamma(T') \subseteq S$, but $\Gamma(T) \neq \Gamma(T')$. This implies $c_{obs}(T) \neq c_{obs}(T')$. Consider any $y \in T \setminus T'$. Then, since $S \subseteq T'$, $y \in T \setminus S$. Moreover, since $\Gamma(T) \subseteq S$, we know $y \in T \setminus \Gamma(T)$. By definition of $\Gamma(T)$ for $T \in \mathcal{D}$, this means $y P c_{obs}(T)$. Similarly, if $y \in T' \setminus T$, we conclude $y P c_{obs}(T')$, contradicting that P satisfies (3). To show Γ satisfies (1b), consider $S \in \mathcal{P}(X)$ and $x \in S \setminus \Gamma(S)$. We prove $\Gamma(S \setminus \{x\}) = \Gamma(S)$ in each of the four possible cases:

Case 1: $S \setminus \{x\}, S \in \mathcal{D}$. Since $S \in \mathcal{D}$, and $x \notin \Gamma(S)$, we know $x P c_{obs}(S)$. Suppose that $\Gamma(S \setminus \{x\}) \neq \Gamma(S)$. Then $c_{obs}(S) \neq c_{obs}(S \setminus \{x\})$. Applying (3) for choice problems S and $S \setminus \{x\}$, we conclude $c_{obs}(S) P x$, a contradiction.

Case 2: $S \setminus \{x\} \in \mathcal{D}, S \notin \mathcal{D}$. Since $S \setminus \{x\} \in \mathcal{D}$, we know $\Gamma(S \setminus \{x\}) = c_{obs}(S \setminus \{x\}) \cup \{y \in S | c_{obs}(S \setminus \{x\}) P y\}$. Since $S \setminus \Gamma(S) \neq \emptyset$, there exists $T \in \mathcal{D}$ with $S \subseteq T$ and $\Gamma(T) \subseteq S$. Because $T \in \mathcal{D}$, $z P c_{obs}(T)$ for all $z \in T \setminus S$. Since $\Gamma(S) = \Gamma(T)$, we know $x \in T \setminus \Gamma(T)$. Hence $x P c_{obs}(T)$. If $\Gamma(S \setminus \{x\}) \neq \Gamma(S) = \Gamma(T)$, then $c_{obs}(S \setminus \{x\}) \neq c_{obs}(T)$ contradicting (3) for the pair of sets T and $S \setminus \{x\}$.

Case 3: $S \setminus \{x\} \notin \mathcal{D}, S \in \mathcal{D}$. Since $S \in \mathcal{D}$, $\Gamma(S) = c_{obs}(S) \cup \{y \in S | c_{obs}(S) P y\}$. If $x \in S \setminus \Gamma(S)$ then $\Gamma(S) \subseteq S \setminus \{x\}$, so by construction $\Gamma(S \setminus \{x\}) = \Gamma(S)$.

Case 4: $S \setminus \{x\}, S \notin \mathcal{D}$. Since $S \setminus \Gamma(S) \neq \emptyset$, there exists $T \in \mathcal{D}$ with $S \subseteq T$ and $\Gamma(T) \subseteq S$. Since $x \in S \setminus \Gamma(S)$, then $\Gamma(T) = \Gamma(S) \subseteq S \setminus \{x\}$ and so $\Gamma(S \setminus \{x\}) = \Gamma(T)$ by construction, and equals $\Gamma(S)$ by transitivity. *Q.E.D.*

PROOF OF PROPOSITION 3: *Testing mixed binary restrictions is NP-hard*

Fix an instance of SAT3 with a set V of variables and a set $\mathcal{C}$ of clauses. The three literals (a variable or its negation) involved in a clause C are denoted ℓ_i^C for $i = 1, 2, 3$. Consider the set of options X that contains all variables v and their negations ($\bar{v}$), all clauses C, an option x_C for each clause C, and an option t. Let $\mathcal{R}$ be the following mixed set of binary restrictions: the restriction $1_{(\{v,\bar{v}\},t)}$ for each $v \in V$ and the restrictions $1_{(t,C)}$, $1_{(x_C,\{\ell_1^C,\ell_2^C\})}$ and $1_{(C,\{x_C,\ell_3^C\})}$ for each clause C.

We show the instance of SAT3 has a truthful assignment if and only if $\mathcal{R}$ is acyclically satisfiable. Given a truthful assignment for SAT3, an ordering constructed as follows will satisfy $\mathcal{R}$: place from worst to best, first all variables v that are true, then $\bar{v}$ for each false variable v, then x_C for all clause C such that ℓ_1^C or ℓ_2^C is true, then all clauses C, then t, then all remaining x_C's, then v for all false variables v, and finally $\bar{v}$ for all true variables v. Conversely, let P be an acyclic relation satisfying $\mathcal{R}$. We can assume without loss of generality that P is an ordering (otherwise take a completion of P; this will still satisfy $\mathcal{R}$). All variables ranked below t are declared true, while all others are declared false. It is easy to check that this defines a truthful assignment for the instance of SAT3. *Q.E.D.*

PROPOSITION 4 *Testing consistency with Limited Attention is NP-hard.*

Proof. Fix a mixed set $\mathcal{R}$ of binary restrictions defined on a set X. For each restriction r, let x_r be the option whose contour set is being restricted, and let y_r and z_r ($y_r = z_r$ is allowed) be the two options potentially included in the upper (or lower) contour set of x_r if r is an UCS (or LCS) restriction. Consider the set of options X' that contains all options in X, plus a new option t_r for each LCS restriction r and new options u_r, v_r, and w_r for each UCS restriction r, and the following observed choices:

S	$t_r x_r$	$u_r v_r$	$u_r w_r$	$u_r v_r x_r$	$u_r w_r x_r$	$v_r y_r z_r$	$w_r y_r z_r$	$t_r x_r y_r z_r$
$c_{obs}(S)$	t_r	u_r	u_r	v_r	w_r	y_r	z_r	x_r

for each restriction r. By Proposition 1, c_{obs} is consistent with Limited Attention if and only if there is an acyclic relation P satisfying the following restrictions $\mathcal{R}'$:

(i) $x_r P y_r$ or $x_r P z_r$, from $c_{obs}(\{t_r, x_r\}) = t_r$ and $c_{obs}(\{t_r, x_r, y_r, z_r\}) = x_r$,

for each LCS r, and the following three restrictions for each UCS r:

(ii) $y_r P v_r$ or $z_r P w_r$, from $c_{obs}(\{v_r, y_r, z_r\}) = y_r$ and $c_{obs}(\{w_r, y_r, z_r\}) = z_r$,

(iii) $v_r P x_r$, from $c_{obs}(\{u_r, v_r\}) = u_r$ and $c_{obs}(\{u_r, v_r, x_r\}) = v_r$,

(iv) $w_r P x_r$, from $c_{obs}(\{u_r, w_r\}) = u_r$ and $c_{obs}(\{u_r, w_r, x_r\}) = w_r$.

We show acyclic satisfiability holds for $\mathcal{R}$ if and only if it holds for $\mathcal{R}'$. Let P be an acyclic relation satisfying the restrictions in $\mathcal{R}$. We can assume without loss of generality that P is an ordering (otherwise take a completion of P that still satisfy restrictions in $\mathcal{R}$). Then extend P to any ordering on X' that ranks the new options v_r and w_r for each r directly above the corresponding x_r, and below any other $y \in X$ such that yPx_r. It is easy to check that any such extension satisfies $\mathcal{R}'$. Conversely, let P' be an acyclic relation satisfying the restrictions in $\mathcal{R}'$. We can assume without loss of generality that P' is an ordering (otherwise take a completion of P'; that will still satisfy $\mathcal{R}'$). Clearly, the restriction of P' to X satisfies $\mathcal{R}$. *Q.E.D.*

APPENDIX B: PROOFS FOR SECTIONS 6.1-6.4

B.1 Choice Overload

PROPOSITION 5 *Observed choices c_{obs} are consistent with the theory of maximizing a preference ordering over a consideration set satisfying (4) if and only if the following collection of restrictions is acyclically satisfiable:*

(5) *For all $S, T \in \mathcal{D}$ with $c_{obs}(S) \neq c_{obs}(T) \in S \subset T : c_{obs}(S) P c_{obs}(T)$.*

Proof. Necessity was given in the text. For sufficiency, suppose there is an acyclic relation satisfying (5), and let P be a transitive completion (hence P still satisfies (5)). Define Γ_P by $\Gamma_P(S) = \{\arg\min_P S\} \cup \{c_{obs}(T) \mid S \subseteq T,\ T \in \mathcal{D},\ c_{obs}(T) \in S\}$ for all $S \in \mathcal{P}(X)$. This Γ_P satisfies (4) and thus by CFS13 (Section 4.1), it is the set of rationalizable elements for some rationales $\{R_k\}_k$. Let c be the choice function arising from $(P, \{R_k\}_k)$ under the theory. For any $S \in \mathcal{D}$, we show $c(S) = c_{obs}(S)$. Suppose otherwise; then $\Gamma_P(S)$ contains at least two elements, and $c(S)$ must be the observed choice from some $T \in \mathcal{D}$ with $S \subset T$. This implies $c_{obs}(S)$ is revealed preferred to $c(S)$, contradicting P-maximality of $c(S)$ in $\Gamma_P(S)$ for a P satisfying (5). *Q.E.D.*

B.2 Reference Dependence

Formalizing the discussion in the text, the DM's revealed preference given the reference point x in S is $aP_{S,x}b$ if $c_{obs}(R) = a$ for some $R \subseteq S$ such that $b, x \in R$. Thus x cannot be most salient in S if $P_{S,x}$ is cyclic. Let $\mathcal{R}_{ref}$ be the following necessary UCS restrictions on $\succ_\sigma$: for each $S \in \mathcal{P}(X)$ and $x \in S$ such that $P_{S,x}$ is cyclic, there is $y \in S \setminus \{x\}$ with $y \succ_\sigma x$. These restrictions are also sufficient:

Proposition 6 *Observed choices c_{obs} are consistent with Triggered Rationality if and only if $\mathcal{R}_{ref}$ is acyclically satisfiable.*

Interestingly, the proof reveals that it is without loss of generality for the theory to require that if x is preferred to y when the reference point is y, then x is also preferred to y when the reference point is x, capturing a form of anchoring bias.

Proof. It remains to show sufficiency. Suppose an acyclic relation satisfying $\mathcal{R}_{ref}$ exists, and let $\succ_\sigma$ be a transitive completion (hence $\succ_\sigma$ still satisfies $\mathcal{R}_{ref}$). Let x_i denote the i-th maximal element according to $\succ_\sigma$. For each i, let P_{x_i} be a transitive completion of P_{X_i,x_i}. Such a completion exists, because x_i being $\succ_\sigma$-maximal in $X_i = \{x_i, x_{i+1}, \ldots, x_n\}$ implies P_{X_i,x_i} is acyclic. The choices generated by these primitives will now be shown to coincide with c_{obs} on $\mathcal{D}$. Take any $S \in \mathcal{D}$. Let k be the smallest index such that $x_k \in S$. Then $S \subseteq X_k$. By definition of P_{x_k}, $c_{obs}(S) \succ_{x_k} y$ for all $y \in S \setminus \{c_{obs}(S)\}$. *Q.E.D.*

While there could be many restrictions in $\mathcal{R}_{ref}$, the enumeration procedure always tests the theory tractably, since one need only consider the $|X| - 1$ nonsingleton sets $X,\ X \setminus \{x_1\},\ X \setminus \{x_1, x_2\}, \ldots$ encountered along the enumeration. If $P_{S,x}$ is acyclic, then so is $P_{S',x}$ for $x \in S' \subset S$. Moreover, within each such set, there is no need to consider all the restrictions: one may move to the next set upon finding any x for which $P_{S,x}$ is acyclic. Finally, a nested enumeration procedure tests if $P_{S,x}$ is acyclic.

B.3 Shortlisting

Proposition 7 *Testing consistency with Shortlisting is NP-hard. This remains true for Order Shortlisting (which requires the DM's preference to be an ordering).*

Proof. Fix a mixed set $\mathcal{R}$ of binary restrictions defined on a set X. For each restriction r, let x_r be the option whose contour set is being restricted; and without loss of generality assume $y_r \neq z_r$ are the two options potentially included in the upper (or lower) contour set of x_r if r is an UCS (or LCS) restriction. Consider the set of options X' that contains all options in X, plus the options a_r, b_r, c_r for each UCS restriction r and the options d_r, e_r, f_r, g_r for each LCS restriction r. Take the following data:

S	$a_r b_r$	$c_r x_r$	$b_r c_r x_r$	$a_r b_r y_r z_r$
$c_{obs}(S)$	b_r	x_r	c_r	a_r

for each UCS restriction r, and

S	$d_r f_r$	$e_r g_r$	$d_r f_r x_r$	$e_r g_r x_r$	$d_r y_r z_r$	$e_r y_r z_r$
$c_{obs}(S)$	d_r	e_r	f_r	g_r	z_r	y_r

for each LCS restriction r.

Define the following set of restrictions $\mathcal{R}'$ over X', which pertain to the shortlisting relation P_1 and are necessary conditions for consistency of c_{obs} with Shortlisting:

(i) $1_{(b_r, x_r)}$ from the two middle observed choices for UCS restrictions r, since c_r cannot eliminate x_r and x_r is revealed preferred to c_r.

(ii) $1_{(\{y_r, z_r\}, b_r)}$ from the first and last observed choices for UCS restrictions r, since a_r cannot eliminate b_r and b_r is revealed preferred to a_r.

(iii) $1_{(x_r, d_r)}$ and $1_{(x_r, e_r)}$ from the first four choices for LCS restrictions r, since by adding x_r, the DM chooses the revealed-worse options f_r and g_r, respectively.

(iv) The restriction $1_{(d_r, y_r)} \vee 1_{(e_r, z_r)}$ from the last two choices for the LCS restrictions r, since neither y_r nor z_r can eliminate each other, yet one is more preferred.

Acyclic satisfiability is also sufficient for consistency with Shortlisting. Let P' be an ordering on X' satisfying $\mathcal{R}'$, which exists by acyclic satisfiability of $\mathcal{R}'$. The restriction of P' to X, denoted P, satisfies the restrictions in $\mathcal{R}$. Let then P_1 be the relation on X' defined as follows: (a) for each LCS restriction r: $x_r P_1 d_r$, $x_r P_1 e_r$, $f_r P_1 x_r$, $g_r P_1 x_r$, $d_r P_1 y_r$ if $x_r P y_r$, and $e_r P_1 z_r$ if $x_r P z_r$, and (b) for each UCS restriction r: $a_r P_1 y_r$, $a_r P_1 z_r$, $b_r P_1 x_r$, $y_r P_1 b_r$ if $y_r P x_r$ and $z_r P_1 b_r$ if $z_r P x_r$. Notice that P_1 is acyclic.[25] Take P_2 as a preference ordering on X' such that $b_r P_2 a_r$ and $x_r P_2 c_r$, for

[25] Indeed, one can derive from P an ordering on X' satisfying these restrictions, simply by placing a_r above all options in X and b_r right above x_r, for each UCS restriction r, and by placing f_r and g_r above all options in X, and d_r and e_r below x_r but above any other element of X, for each LCS restriction r, and c_r at the bottom for each UCS restriction.

each UCS restriction r, and $y_r P_2 e_r$, $z_r P_2 d_r$, $d_r P_2 f_r$, $e_r P_2 g_r$, $y_r P_2 z_r$ if $z_r P y_r$, and $z_r P_2 y_r$ if $y_r P z_r$, for each LCS restriction r.[26] It is easy to check that the choice function associated to (P_1, P_2) coincides with c_{obs} on $\mathcal{D}$. Given that the constructed P_2 is an ordering, this proof will also establish that Order Shortlisting is NP-hard.

Finally, we claim $\mathcal{R}'$ is acyclically satisfiable if and only if $\mathcal{R}$ is acyclically satisfiable. Necessity follows from the argument at the beginning of the previous paragraph. For the converse, let P be an acyclic relation satisfying $\mathcal{R}$. Then, just as in Footnote 25, we may derive from P an acyclic relation P' over X' that satisfies $\mathcal{R}'$. *Q.E.D.*

B.4 Undominated Alternatives

Define $Z(x)$ to be the union of all the choice problems $T \in \mathcal{D}$ such that $x \in C_{obs}(T)$. As discussed in the text, if $x \notin C_{obs}(S)$ for some $S \in \mathcal{D}$, then there must exist some y in $S \setminus Z(x)$ that dominates x. Let $\mathcal{R}_{UD}(C_{obs})$ be the set of of strict UCS restrictions $1_{(\Sigma,x)}$, where $\Sigma = \{\{y\} | y \in S \setminus Z(x)\})$, derived by varying $S \in \mathcal{D}$ and $x \in S \setminus C_{obs}(S)$. If some $S \subseteq Z(x)$, then satisfying $\mathcal{R}_{UD}(C_{obs})$ is clearly impossible, and observed choices must be inconsistent with choosing the undominated elements for some acyclic relation. Thus, we henceforth assume that $S \setminus Z(x)$ is nonempty for all restrictions in $\mathcal{R}_{UD}(C_{obs})$.[27] These strict UCS restrictions capture the empirical content of the theory.

Proposition 8 *The observed choice correspondence C_{obs} is strongly consistent with choosing undominated elements according to an acyclic relation if and only if $\mathcal{R}_{UD}(C_{obs})$ is acyclically satisfiable.*

Proof. Necessity follows from the above discussion. As for sufficiency, let $\succ$ be acyclic relation satisfying the restrictions in $\mathcal{R}_{UD}(C_{obs})$. As $\succ$ may be overly complete, construct $\succ^*$ by $y \succ^* x$ if $y \succ x$ and there is $S \in \mathcal{D}$ such that $x \in S \setminus C_{obs}(S)$ and $y \in S \setminus Z(x)$. Then $\succ^*$ inherits the property of acyclicity from $\succ$. Let $T \in \mathcal{D}$. We have to check that $C_{obs}(T) = \{x \in T \mid \nexists y \in T : y \succ^* x\}$. Suppose that $x \in C_{obs}(T)$ and that $y \in T$. Then $y \in Z(x)$, and thus it is not the case that $y \succ^* x$. Conversely, let $x \in T \setminus C_{obs}(T)$. Then there is $y \in T \setminus Z(x)$ such that $y \succ x$ since $\succ$ satisfies the restrictions in $\mathcal{R}_{UD}(C_{obs})$, and moreover $y \succ^* x$ holds. *Q.E.D.*

[26] For instance, rank from bottom to top: a_r, b_r, c_r, for all UCS restriction r, then f_r, g_r, d_r, e_r for all LCS restriction r, then options in X in opposite order than P.

[27] This assumption is only to simplify notation; one can instead create an auxiliary option $\diamond$, augment $\mathcal{R}_{UD}(C_{obs})$ by including $\diamond$ in every $Z(x)$ and adding the restrictions $1_{(y,\diamond)}$ for all $y \in X$, and then test acyclic satisfiability of these restrictions over $X \cup \{\diamond\}$.