BOUNDED RATIONALITY AND LIMITED DATASETS*

Geoffroy de Clippel† and Kareen Rozen‡

May 2020

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 may lead to false conclusions that the data is consistent with the theory.

*This work greatly benefitted from suggestions by Eddie Dekel, Kfir Eliaz, Yusufcan Masatlioglu, Larry Samuelson, the editor Ran Spiegler, two anonymous referees, and 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 proposes insightful and plausible choice procedures in response to mounting evidence against rational choice. Great progress has been made to understand the choice functions these theories generate. Though theoretically insightful, such results do not apply to typical situations, in which only some choices are observed. In empirical settings, the modeler cannot control the choice problems individuals face. In experimental settings, generating a complete dataset can require an overwhelming number of decisions (e.g., 1,013 choice problems with 10 alternatives). A long literature studies Rationality under limited data (Samuelson, 1948; Houthaker, 1950; Richter, 1966; Afriat, 1967; Sen, 1971; Varian, 1982). We explore how these ideas can be brought into the discourse on bounded rationality.

A Decision Maker (DM)'s observed choices are consistent with a theory if they can be extended to a complete choice function that arises under the theory. The first works to study bounded rationality theories under limited data focus on explaining only observed choices, without considering out-of-sample implications (Manzini and Mariotti 2007, 2012; Tyson, 2013). This approach may seem natural, since we never worry about out-of-sample problems when testing Rationality in its standard description: if there is a preference ordering for which observed choices are maximal, then choices may be defined for out-of-sample problems by maximizing that same preference. Such extensibility need not hold in general, leading to an overfitting problem: choices could be attributed to a theory under which no extension of those choices can arise. This paper thus begins by clarifying 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 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.

For theories of bounded rationality, we observe that choices may reveal more elaborate information than the direct-preference comparisons revealed under Rationality. Moreover, that information may pertain to an acyclic relation that need not be complete, and need not be interpretable as a preference. Letting $c_{obs}(S)$ represent the observed choice of the DM at the choice set S, consider the choice pattern $c_{obs}(\{a,b,d,e\}) = a$ and $c_{obs}(\{a,b\}) = b$. Under Masatlioglu, Nakajima and Ozbay (2012)'s theory of Limited Attention, this choice pattern only reveals that the DM prefers a to at least one of d or e; it does not tell us whether d alone, e alone, or both options are worse than a. Under Rubinstein and Salant (2006a,b)'s theory of Triggered Rationality, which models reference dependence, the choice pattern above instead reveals information about the DM's salience ordering: namely, that the DM finds at least of one of d or e more salient than a (again, it does not tell us which).

This paper argues that the emergence of new forms of revealed information does not preclude us from pursuing the classic testing approach, at least for theories which have an acyclic relation in their description (or are observationally equivalent to one that does). We show that many prominent bounded-rationality theories are characterized by the ability to find an acyclic relation that simultaneously satisfies a collection of *restrictions* revealed by observed choices. This notion of *acyclic satisfiability* is the natural extension of SARP to accommodate revealed restrictions that may be more complex than those for Rationality.

Since Rationality only reveals direct comparisons like 'a is better than b', the data precisely identifies which relation to check for cycles. By contrast, checking acyclic satisfiability of more complex restrictions entails some guesswork: one must *find* an acyclic relation satisfying them. Nonetheless, we show that the same simple and elegant way we test Rationality generalizes to test theories that are characterized by acyclic satisfiability of *all upper-contour set (UCS) restrictions*, or *all lower-contour set (LCS) restrictions*. Intuitively, we test Rationality by constructing a preference ordering over the elements that respects the revealed preference: there must be a best element $x_1 \in X$ that is not revealed inferior to any alternative, a next-best element x_2 that is not revealed inferior to any alternative in $X \setminus \{x_1\}$, etc. A UCS restriction, which we formalize later, reveals that the upper-contour set of an element must contain some other element(s), but may not tell us which. Still, this information suffices to determine viable candidates for 'best' in a set. Hence the way we test Rationality applies more broadly than it seems.

Perhaps surprisingly, a variety of theories beyond Rationality can be tested in this manner. There are also several bounded-rationality theories which generate restrictions outside of the UCS (or LCS) class. Even then, applying the revealed-restrictions methodology results in insights and a deeper understanding of the theory. One may still wonder whether, by thinking some more, it is possible to find a 'better' characterization with UCS restrictions, or some more clever procedure, to make testing the theory as simple as testing Rationality. It turns out that even small, systematic departures from having all UCS (or all LCS) restrictions make testing significantly more challenging. We formalize the sense in which a theory can *generate rich sets of restrictions*, and show that testing such a theory is NP-hard. Given the current consensus that P≠NP, this means that there is no easy way to systematically check consistency, thereby addressing the questions above. But understanding the restrictions such theories generate can still help pinpoint datasets for which testing is tractable, because restrictions simplify to UCS (LCS).

We illustrate the methodology with four bounded-rationality theories: Choice Overload (Lleras, Masatlioglu, Nakajima and Ozbay, 2017), Triggered Rationality, Limited Attention and Shortlisting. We find that Choice Overload and Triggered Rationality boil down to the acyclic satisfiability of UCS restrictions, making them testable in a manner akin to Rationality. We show Limited Attention corresponds to acyclic satisfiability of more elaborate restrictions, and in fact generates rich sets of restrictions; however, these simplify to LCS if the dataset contains the intersection of any two choice problems that violate WARP. Similarly, we establish that Shortlisting generates rich sets of restrictions, but these simplify to UCS if the dataset contains all pairwise choices, and any triplets from which pairwise choices are cyclic.

The versatility of our approach does not mean there is no limitation. For each theory, seeing the empirical content through the lens of acyclic satisfiability takes a bit of thought. While many theories (explicitly or implicitly) involve a relation to examine, a theory may have a sufficiently different structure that trying to understand it through acyclic satisfiability is unfruitful. Still, the array of applications illustrates the scope of the approach we develop here. Some further applications appear in the earlier working paper or subsequent work, covering among other things choice correspondences (de Clippel and Rozen, 2019); various forms of satisficing (Barberà, de Clippel, Neme and Rozen, 2020); misperception in consumer theory (de Clippel and Rozen, 2018a); Manzini and Mariotti (2012)'s Categorization and Cherepanov, Feddersen and Sandroni (2013)'s Rationalization with limited data (de Clippel and Rozen, 2018b); and classic assignment methods in interactive decision-making with rational agents (de Clippel and Rozen, 2018c).

2. Testing the Consistency of Observed Choices with Theories

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 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* of choice describes the DM's choice procedure, and defines the set of choice functions that could possibly arise. For instance, Rationality posits that the DM applies a single preference ordering to select the best element from any choice problem: for some ordering P, $c(S) = \arg\max_P S$ for all $S \in \mathcal{P}(X)$. Manzini and Mariotti (2007)'s theory of *Shortlisting* posits that the DM makes a shortlist of undominated options using an asymmetric relation P_1, and chooses the undominated element in the shortlist according to an asymmetric preference relation P_2:

$$\{c(S)\} = \max(\max(S, P_1), P_2), \text{ for all } S \in \mathcal{P}(X).[1] \tag{1}$$

As another example, Masatlioglu, Nakajima and Ozbay (2012)'s theory of *Limited Attention* posits that the DM has some preference ordering P, and for each choice problem the DM maximizes P over his consideration set $\Gamma(S) \subseteq S$, with the restriction that consideration sets don't change when removing ignored alternatives:

$$c(S) = \arg\max_P \Gamma(S), \text{ for all } S \in \mathcal{P}(X), \text{ and} \tag{2a}$$

$$\Gamma(S) \subseteq T \subseteq S \Rightarrow \Gamma(T) = \Gamma(S), \text{ for all } S, T \in \mathcal{P}(X). \tag{2b}$$

To be clear, we use *relation* to mean a binary relation (possibly incomplete or cyclic); we use *ordering* to mean a complete, asymmetric (i.e., strict) and transitive relation; and for any relation P and $S \subseteq X$, denote $\arg\max_P S := \{x \in S \mid xPy, \ \forall y \in S \setminus \{x\}\}$.

Importantly, only the DM's choices, not her thought process or choice method, can be observed. 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. 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

[1] 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.

function that coincides with c_{obs} on $\mathcal{D}$, guaranteeing the ability to make coherent predictions for out-of-sample choice problems.

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

It is possible for two theories to describe quite different choice processes, and yet generate the same set of choice functions. While formally different, Definition 1 makes clear that such theories cannot be distinguished from each other using standard choice data alone. A theory's testable implications thus come only from the set of choice functions that it generates, not from the precise story or mathematical language the theory uses to define those choice functions.

We briefly expand on this point before moving to our main results in the next sections, as Definition 1 highlights an issue in the recent literature. To study Shortlisting with limited data, Manzini and Mariotti (Corollary 1, 2007) examine when there exist asymmetric relations P_1, P_2 such that (1) holds for $S \in \mathcal{D}$. Manzini and Mariotti (Definition 4, 2012) take a similar approach for Categorize-Then-Choose. To test Limited Attention with limited data, Tyson (2013) seeks conditions guaranteeing the existence of an ordering P and a consideration set mapping defined on $\mathcal{D}$ such that (2a) and (2b) hold for $S, T \in \mathcal{D}$. In other words, instead of using Definition 1, they check the conditions describing how choices emerge over observed choice problems.

The approach taken by this literature may seem natural at first, as it mirrors how we typically apply the classic definition of Rationality: we simply look for a preference ordering for which the observed choices are maximal, because it is trivial to extend them to a complete choice function under Rationality (just maximize that same preference). Such extensibility, however, does not hold for Shortlisting, Categorize-then-Choose, or Limited Attention. This means observed choices may be incorrectly attributed to a theory under which no extension of those choices can arise.

Before illustrating this overfitting problem for Limited Attention, it is helpful to illustrate the basic idea with a simpler example. Suppose $X = \{a, b, c\}$ and we observe $c_{obs}(\{a, b\}) = a$, $c_{obs}(\{b, c\}) = b$ and $c_{obs}(\{a, c\}) = c$. Imagine our theory says that the DM makes a choice in each set which is undominated according to an asymmetric binary relation. This theory generates the same choice functions as Rationality, because the DM's need to make a choice from $\{a, b, c\}$ forces the binary relation to be an ordering; Sen (1971) makes a related point for preference maximization with respect to general binary relations, observing the relation is necessarily a weak ordering when all finite subsets are in the feasible domain. It would be mistaken to deem c_{obs} consistent with this theory just because *observed* choices are undominated according to the relation $aPbPcPa$: there is no possible extension of c_{obs} to the out-of-sample problem $\{a, b, c\}$ that remains consistent with the theory.[2]

Going back to Limited Attention, suppose we observe the following choices:

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

The prior literature would deem c_{obs} consistent with Limited Attention, since (2a) and (2b) hold for $S, T \in \mathcal{D}$ using the ordering defined by $aPdPePbPf$, and with $\Gamma(S)$ given by the weak lower-contour set of $c_{obs}(S)$ for $S \in \mathcal{D}$. To the contrary, we claim there is no choice function under Limited Attention that extends c_{obs}. Suppose otherwise. Since removing d from $\{a, d, e\}$ changes the choice, (2b) implies d is considered in $\{a, d, e\}$. As a is chosen from $\{a, d, e\}$, and d is considered, we learn aPd. Similarly, $b \in \Gamma(\{b, e, f\})$ and ePb. Now consider the out-of-sample problem $\{b, d\}$. The ranking aPd implies $a \notin \Gamma(\{a, b, d\})$, since d is chosen from that set. Thus (2b) requires $\Gamma(\{b, d\}) = \Gamma(\{a, b, d\})$. Similarly, ePb implies $e \notin \Gamma(\{b, d, e\})$, thus (2b) also requires $\Gamma(\{b, d\}) = \Gamma(\{b, d, e\})$. This implies $\Gamma(\{a, b, d\}) = \Gamma(\{b, d, e\})$, which is impossible because the choices from $\{a, b, d\}$ and $\{b, d, e\}$ differ.

This caveat does not apply to all bounded-rationality theories. In some cases, though, minor rephrasings of a theory (that do not affect the set of choice functions generated) can switch extensibility on/off. The variant of Rationality above provides such an example: the issue would not occur when requiring the asymmetric relation to also be acyclic. Choice theorists could study how to phrase a theory in an extensible way, but we see no reason to be concerned with phrasing here. As Definition 1 makes clear, a correct test of consistency depends only on the set of choice functions arising under a theory. Thus our work in the next sections circumvents the matter of extensibility entirely.

[2]For another example, take a theory which says choices satisfy the IIA axiom: if $c(T) = x$ and $x \in S \subset T$ then $c(S) = x$. This also generates the same choice functions as Rationality. No extension of the above c_{obs} to $\{a, b, c\}$ satisfies IIA, and it would be mistaken to deem c_{obs} consistent with the theory merely because IIA is satisfied (vacuously) on $\mathcal{D} = \{\{a, b\}, \{b, c\}, \{a, c\}\}$.

3. Building Insights

With the proper definition of consistency, we can now ask how theories of bounded rationality can be tested. Though we don't have a universal answer to offer, we do show that the classic revealed-preference approach to testing rationality generalizes substantially. For this, we restrict attention to theories with at least one acyclic relation in their description. Of course, since testing is based only on the set of choice functions compatible with a theory, the methodology also applies to theories that don't include an acyclic relation in their description, but are observationally equivalent to one that does.

The common methodology goes as follows. Assuming consistency, start by identifying choice patterns revealing some restrictions on the acyclic relation. By construction, finding an acyclic relation satisfying these restrictions – what we will call *acyclic satisfiability* in the next section – is a necessary condition for consistency. To make sure the data cannot reveal any more critical information about this relation, one must prove sufficiency: the existence of an acyclic relation satisfying these revealed restrictions must guarantee consistency. Failing to do so (and finding a counter-example) means that the theory has more restrictions to reveal.

To illustrate, consider a DM who decides what to order from restaurant menus. A menu item $x \in X$ comprises both a description of the food and its price, which is denoted $p(x)$. The DM may be indecisive initially, as captured by an incomplete preference P_1 over menu items, but then makes a choice by selecting the cheapest item among those that are P_1-undominated. For expositional convenience, we assume the menu items in X have distinct prices.[3] The theory *Frugal when Undecided* posits that the DM picks:

$$\{c(S)\} = \arg \min_{x \in \max(S,P_1)} p(x), \text{ for all S} \in \mathcal{P}(X).$$

This theory is an instance of Manzini and Mariotti (2007)'s Shortlisting, with the second criterion P_2 known by the modeler (frugality in this case); we'll discuss the general case of unknown and possibly cyclic P_2 in Section 5.4. While Shortlisting presumes P_1 is asymmetric, one can equivalently require it to be acyclic: if there were a cycle, then the DM would have no shortlist for the menu of elements in that cycle.

[3]This ensures single-valued choice functions, as assumed throughout the paper. Correspondences can also be accommodated, see the working-paper version, de Clippel and Rozen (2019).

Consider what observed choices reveal about P_1. Suppose that in some observed menu S, there is a menu item x which is cheaper than $c_{obs}(S)$. For choices to still be consistent with the theory, there must be an item $y \in S$ which rules out x from consideration, that is, with yP_1x. 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 $S \setminus \mathcal{S}(x)$ that P_1-dominates x, where $\mathcal{S}(x)$ is the union of all observed choice problems from which x is chosen. Being able to find an acyclic relation satisfying these restrictions is necessary for consistency. The next proposition shows it is also sufficient.

Proposition 1 *Observed choices c_{obs} are consistent with Frugal when Undecided if and only if there is an acyclic relation P_1 satisfying the restrictions:*

(3) *For all $x \in S \in \mathcal{D}$ with $p(x) < p(c_{obs}(S))$, there is $y \in S \setminus \mathcal{S}(x)$ such that yP_1x*

Proof. Necessity follows from the above discussion. As for sufficiency, let P_1 be an acyclic relation satisfying (3). As P_1 may contain too many relationships, we construct a selection P_1^* in the following way: yP_1^*x if yP_1x and x belongs to some observed problem S for which $p(x) < p(c_{obs}(S))$ and $y \in S \setminus \mathcal{S}(x)$. Since P_1 is acyclic, so is P_1^*.

Because P_1^* is acyclic, the shortlist $\{x \in S \mid \nexists y \in S : yP_1^*x\}$ is nonempty for all choice problems S. We now check that for $S \in \mathcal{D}$, $c_{obs}(S)$ is the cheapest item in the shortlist. First, notice $c_{obs}(S)$ belongs to the shortlist: there is no $y \in S$ which is P_1^*-superior to $c_{obs}(S)$, because every $y \in S \setminus \{c_{obs}(S)\}$ belongs to $\mathcal{S}(c_{obs}(S))$. Finally, since P_1 satisfies (3), for any $x \in S$ that is cheaper than $c_{obs}(S)$ there is $y \in S \setminus \mathcal{S}(x)$ such that yP_1x. Thus yP_1^*x holds, and x does not belong to the shortlist. *Q.E.D.*

Determining whether or not there exists an acyclic relation satisfying restrictions like (3) may at first seem hard to check. As opposed to the classic SARP condition of Rationality, the restrictions in (3) don't immediately tell us which pairwise comparisons to examine for cycles. Consider the following example.

Example 1 *Suppose $X = \{a, b, d, e, f, g\}$ and we observe the choices:*

S	ag	bf	de	$abde$	$abdf$	$abdg$
$c_{obs}(S)$	a	b	d	e	f	g

with items coming earlier in the alphabet cheaper than items coming later. Notice that

$\mathcal{S}(a) = \{a, g\}$, $\mathcal{S}(b) = \{b, f\}$, and $\mathcal{S}(d) = \{d, e\}$. Since a, b and d are available when the more expensive item e is picked from $\{a, b, d, e\}$, the restrictions in (3) reveal:

$$\begin{aligned} bP_1a \textit{ or } dP_1a \textit{ or } eP_1a, \textit{ from } p(a) < p(e) \\ aP_1b \textit{ or } dP_1b \textit{ or } eP_1b, \textit{ from } p(b) < p(e) \\ aP_1d \textit{ or } bP_1d, \textit{ from } p(d) < p(e) \end{aligned} \tag{4}$$

Similarly, since f is chosen from $\{a, b, d, f\}$, we learn that:

$$\begin{aligned} bP_1a \textit{ or } dP_1a \textit{ or } fP_1a, \textit{ from } p(a) < p(f) \\ aP_1b \textit{ or } dP_1b, \textit{ from } p(b) < p(f) \\ aP_1d \textit{ or } bP_1d \textit{ or } fP_1d, \textit{ from } p(d) < p(f) \end{aligned} \tag{5}$$

Finally, since g is chosen from $\{a, b, d, g\}$, we learn that:

$$\begin{aligned} bP_1a \textit{ or } dP_1a, \textit{ from } p(a) < p(g) \\ aP_1b \textit{ or } dP_1b \textit{ or } gP_1b, \textit{ from } p(b) < p(g) \\ aP_1d \textit{ or } bP_1d \textit{ or } gP_1d, \textit{ from } p(d) < p(g) \end{aligned} \tag{6}$$

By Proposition 1, the data is consistent with Frugal when Undecided if, and only if, there exists an acyclic P_1 satisfying the above nine conditions.

At first, it may not be clear how to check whether there is an acyclic relation satisfying the restrictions in Example 1. The next section shows there is a simple, systematic method to do so (and that the above data is inconsistent with the theory). Perhaps surprisingly, this method is just the standard procedure for checking SARP, which generalizes to an extended family of restrictions that includes those in (3).

4. Acyclic Satisfiability, and How to Check It

Under Rationality, consistency reveals multiple direct preference comparisons: xPy whenever x is picked in the presence of y. For other theories, one may encounter restrictions that are conjunctions and disjunctions of direct comparisons. As is well known, any logical formula can be written in disjunctive normal form (i.e., as a disjunction of conjunctions). We restrict attention to this formulation.

Definition 2 *For each $i = 1, \ldots, n$, let $A_i \subseteq X \times X$ be a set of ordered pairs of options. The strict relation P* satisfies the restriction $r_{A_1 \vee \cdots \vee A_n}$ *if there exists i such that xPy for all $(x, y) \in A_i$.*

The restriction that P satisfies a direct comparison xPy is thus written r_{A_1}, where $A_1 = \{(x, y)\}$. For a more elaborate restriction, recall Example 1. The top restriction in (4) is written there as bP_1a or dP_1a or eP_1a. This corresponds to $r_{A_1 \vee A_2 \vee A_3}$, where $A_1 = \{(b, a)\}$, $A_2 = \{(d, a)\}$, and $A_3 = \{(e, a)\}$.

We now introduce the notion of acyclic satisfiability, a natural extension of SARP, to capture the empirical content of many theories of choice.

Definition 3 *A collection of restrictions is* acyclically satisfiable *if there exists a strict acyclic relation satisfying all the restrictions in the collection.*

In the absence of disjunctions (as in SARP), acyclic satisfiability simply boils down to there being no cycles in the (usually incomplete) relation that is pinned down by the restrictions. But in the presence of disjunctions, one does not know precisely which relation to test for acyclicity. In this section, we identify classes of restrictions for which acyclic satisfiability can be tested in the same way as one checks SARP, as well as cases for which testing is much more challenging.

Definition 4 *For $z \in X$, $r_{A_1 \vee \cdots \vee A_n}$* restricts z's upper-contour set *if $y = z$ for all $(x, y) \in \bigcup_i A_i$. In that case, we call $r_{A_1 \vee \cdots \vee A_n}$ an* upper-contour set (UCS) restriction*. Similarly, $r_{A_1 \vee \cdots \vee A_n}$* restricts z's lower-contour set *if $x = z$ for all $(x, y) \in \bigcup_i A_i$, and in that case, we call $r_{A_1 \vee \cdots \vee A_n}$ a* lower-contour set (LCS) restriction*.*

Stated differently, an UCS restriction $r_{A_1 \vee \cdots \vee A_n}$ on z's upper-contour set means that $\{x \in X | (x, z) \in A_i\}$ is included in z's upper-contour set, for some $i = 1, \ldots, n$. The restrictions in Example 1, and more generally (3), are a collection of UCS restrictions. For instance, the top restriction in (4) requires a's upper-contour set to contain at least one of b, d or e; while the second restriction in (4) requires b's upper-contour set to contain at least one of a, d or e.

Consider how to test SARP. In each step $k \geq 1$, one simply seeks and removes an option x_k that is not ranked below any remaining alternatives, according to Samuelson's revealed preference. Such an x_k is a candidate-maximal element among those remaining. It is possible to *enumerate* all of X in this manner if and only if SARP holds (that is, Samuelson's revealed preference is acyclic). Intuitively, this process constructs a possible preference ordering for the DM from the top down, by iteratively identifying a candidate for the best remaining option. Our first observation, formalized in Section 4.1, is that the acyclic satisfiability of all UCS restrictions (or all LCS restrictions) can be tested in the same manner as SARP. While UCS restrictions may not precisely identify *which* alternatives are ranked above an option, they do reveal that something is ranked above it, which is all we need to know to preclude it from being top ranked. By contrast, Section 4.2 shows that testing consistency for theories which systematically generate wider classes of restrictions is much harder.

4.1 Testing Comparable to SARP with UCS (LCS) Restrictions

The ability to apply the procedure for checking SARP to UCS 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 collection of UCS restrictions defined on X. If $\mathcal{R}$ is acyclically satisfiable, then there exists an option $x \in X$ such that no restriction in $\mathcal{R}$ restricts the upper-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 the maximal element. This provides a first, simple necessary condition for acyclic satisfiability: traverse elements of X to find one that does not appear at the bottom of a restriction.

Take an element x_1 with this property (if one exists), and make it the top element of the ordering. If a restriction $r_{A_1 \vee \cdots \vee A_n}$ has $A_i = \{(x_1, x)\}$, for some $x \in X$ and some $i = 1, \ldots, n$, then that restriction is satisfied, and can be safely eliminated from the collection $\mathcal{R}$. Otherwise, simplify the restriction to $r_{A'_1 \vee \cdots \vee A'_n}$, where $A'_i = \{(y, x) \in A_i | y \neq x_1\}$. Let $\mathcal{R}_1$ be this reduced set of UCS 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 top of any pairwise comparison.

Lemmas 1-2 hold independently of the set X and the set of UCS restrictions $\mathcal{R}$, so the reasoning may be iterated. The Lemmas thus provide a conceptual roadmap for defining the *enumeration procedure for UCS 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-best element thus far, then we treat $x_1, \ldots, x_{k-1}$ as if they are ranked above 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, that is, each restriction $r_{A_1 \vee \cdots \vee A_n} \in \mathcal{R}$ simplifies to $r_{A'_1 \vee \cdots \vee A'_n}$, where $A'_i = \{(y, x) \in A_i | y \notin \{x_1, \ldots, x_{k-1}\}\}$, and can be eliminated entirely if $A'_i = \emptyset$ for some i. 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 bottom of any UCS restriction in $\mathcal{R}_{k-1}$. Continue to the next step if and only if such an element is found.

Definition 5 *The enumeration procedure* fails *if in some step we cannot find a candidate x_k for the best 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 best 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. Vice versa, the ranking of options arising from a successful enumeration (that is, with x_k being the k-th best) will satisfy $\mathcal{R}$ by construction. We have thus shown the following.

Proposition 2 *Suppose that consistency of c_{obs} with a theory has been reduced to checking the acyclic satisfiability of a collection $\mathcal{R}$ of UCS restrictions. Then consistency holds if, and only if, the enumeration procedure using $\mathcal{R}$ succeeds.*

We have thus found a systematic way to process UCS restrictions when checking whether or not they are acyclically satisfiable. In particular, testing turns out to be tractable (in a number of steps that is at most polynomial in the size of the data) despite the possibility of disjunctions in the restrictions; this assumes, of course, that restrictions are tractably derived from observed choices, as is the case in all our applications of the enumeration procedure throughout the paper.

Example 1 (continued) *The nine restrictions (4-6) are UCS restrictions. To start, one must find options whose upper-contour set is left unrestricted. All three options e, f and g qualify; we can enumerate them first in any order. This allows us to cross out the six longer restrictions, which have two disjunctions each, in (4), (5) and (6). Three restrictions involving only a, b and d remain, and the enumeration procedure* *fails: for each of these remaining options, the upper-contour set is restricted. Thus the theory Frugal when Undecided is refuted by the example data.*

Clearly, the procedure can also be used to check consistency when it is reduced to checking the acyclic satisfiability of a collection of LCS restrictions. The only difference is that one seeks candidate *minimal* options in each step, i.e., options that do not appear at the top of any remaining LCS restrictions.

4.2 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.' The enumeration procedure shows how to easily work through UCS (LCS) restrictions to determine acyclic satisfiability. Does this mean that tedious guesswork is unavoidable for theories that generate more complex restrictions? Is there some other approach which makes testing consistency easy in those cases? Alternatively, could we have potentially come up with a 'better' characterization of the theory in terms of UCS (LCS) restrictions? Is it possible that some theories of bounded rationality are hard to test, whatever the methodology followed?

These are difficult questions, and we borrow techniques from computer science to provide some answers. The notions of P and NP are of interest to computer scientists for assessing running times of algorithms: 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 verified in polynomial time; and a problem is NP-hard if solving it is at least as complex as solving the most difficult problems in NP. We repurpose these ideas to instead get at the questions above: showing that a particular theory is NP-hard to test establishes there is neither a simple test nor a 'better' characterization that involves UCS (LCS) restrictions, at least given the widely-held belief that P$\neq$NP. In particular, if someone could find a simple way to check any observed choices for consistency with an NP-hard theory, then this would have the wide-reaching implication that P=NP.

Consider a collection of restrictions which is only a small departure from direct comparisons, but does not fall into the UCS (or LCS) class. As formalized below, imagine a collection where each restriction is a disjunction of two unrelated, direct comparisons ("z_1 must be ranked above z_2, or z_3 must be ranked above z_4").

Definition 6 *A* basic disjunctive restriction *on a set Z is a restriction of the form $r_{\{(z_1,z_2)\}\vee\{(z_3,z_4)\}}$, for some distinct $z_1, z_2, z_3, z_4 \in Z$.*

The departure from Section 4.1 may seem small (we'd have an LCS restriction if $z_1 = z_3$ and an UCS restrictions if $z_2 = z_4$), but has a significant impact.[4] To state the next result, we must first formalize what it means for a theory to generate 'rich sets' of restrictions such as these.

Definition 7 *A choice theory* generates rich sets of restrictions *if for any finite set Z and any collection $\mathcal{R}$ of basic disjunctive restrictions on Z, one can construct (in polynomial time) a set of options $X \supseteq Z$ and observed choices c_{obs} from some subsets of X, with the feature that c_{obs} is consistent with the theory if, and only if, $\mathcal{R}$ is acyclically satisfiable.*

For such theories, it turns out that there is no simple method to establish the consistency of any given dataset with the theory.

Proposition 3 *Testing a theory that generates rich sets of restrictions is NP-hard.*

It is important for Proposition 3 that basic disjunctions occur systematically under the theory, as formalized in Definition 7. Indeed, suppose we have a theory for which consistency amounts to acyclic satisfiability of UCS restrictions plus a single basic disjunction $r_{\{(z_1,z_2)\}\vee\{(z_3,z_4)\}}$. Testing remains easy: check by enumeration if the UCS restrictions are acyclically satisfiable while assuming z_1 is superior to z_2; and if that fails, then do the same while assuming z_3 is superior to z_4. The proof of Proposition 3, which appears in the Appendix, proceeds by showing that every instance of SAT3, a classic NP-hard problem, has a polynomial-time reduction to an equivalent problem of checking acyclic satisfiability of some basic disjunctive restrictions, which in turn has a polynomial-time reduction to testing consistency of some observed choices with the theory. Formally, the SAT3 problem takes a set of 'clauses' that are disjunctions of three 'literals' (variables or their negations), and asks whether there is a truth assignment for the variables that makes all clauses true. The idea is that if one could find a tractable way to test consistency with the theory, then one could leverage that

[4]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 is a special type of lower-contour set restriction. 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.

method to tractably solve SAT3, overturning the general consensus that P$\neq$NP.

Proposition 3 is broader than may seem at first glance. A restriction 'z_1 is superior to z_2, or z_3 is superior to z_4' could in general be

(i) an UCS (LCS) restriction when z_2 and z_4 (resp., z_1 and z_3) coincide, or

(ii) a basic disjunction when all z's are distinct, or

(iii) what one could call a *basic UCS-or-LCS restriction* 'a is superior to b, or b is superior to d' (the case where $z_2 = z_3$ or $z_1 = z_4$).

We have not yet discussed collections of basic UCS-or-LCS restrictions; nor have we discussed collections having *both* UCS and LCS restrictions. One could, for instance, imagine a collection of *mixed binary LCS and UCS* restrictions, where each restriction in the collection takes the form 'a is better than b or d' or takes the form 'a is worse than b or d'. What if a theory systematically generates one of these other types of collections of restrictions? It turns out that such a theory will also generate rich sets of restrictions, so focusing on basic disjunctions in Definition 7 is without loss of generality; see the Appendix.

Observation 1 *Suppose for each finite set Z and each collection of basic UCS-or-LCS restrictions (resp., mixed binary LCS and UCS restrictions) on Z, one can construct in polynomial time a set of options $X \supseteq Z$ and observed choices c_{obs} from some subsets of X, with the feature that c_{obs} is consistent with the theory if, and only if, $\mathcal{R}$ is acyclically satisfiable. Then the theory generates rich sets of restrictions.*

Thus we have shown that even small, systematic departures from all UCS, or all LCS, restrictions make theories much harder to test. Finding some choice patterns that reveals such restrictions is suggestive that testing is likely NP-hard. However, reaching that conclusion requires a formal argument to be sure the empirical content cannot generally be simplified further: the modeler must show the theory generates rich sets of restrictions. We demonstrate this in the next section, where we apply our methodology, and illustrate Propositions 2 and 3, for some prominent theories.

5. Applications

We now apply the methodology we developed to study some prominent theories of bounded rationality: Choice Overload, Triggered Rationality (a form of reference-dependent preferences), Limited Attention, and Shortlisting.

5.1 Choice Overload

Suppose the DM may become overwhelmed and unable to consider all alternatives in some choice problems, as in Lleras, Masatlioglu, Nakajima and Ozbay (2017)'s theory of *Choice Overload*. The DM has a preference ordering P and, for each choice problem S, maximizes P over his consideration set $\Gamma(S) \subseteq S$. Crucially, any option considered in a problem remains considered in smaller problems to which it belongs:[5]

$$S \subset T \Rightarrow \Gamma(T) \cap S \subseteq \Gamma(S), \text{ for all } S, T \in \mathcal{P}(X). \tag{7}$$

Notice from (7) that preference comparisons can be inferred from IIA violations. By (7), 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 $c_{obs}(T)$. These are direct preference comparisons, which qualify as both LCS and UCS restrictions, and are thus testable by enumeration.

The next result guarantees that there are no other, possibly more complex restrictions to consider. A nice implication is that the revealed preference identified by the above authors for full datasets happens to also capture the testable implications in limited datasets.

Proposition 4 *Observed choices c_{obs} are consistent with Choice Overload if and only if there is an acyclic relation P satisfying the UCS restrictions:*

$$c_{obs}(S) P c_{obs}(T), \text{ whenever } c_{obs}(S) \neq c_{obs}(T) \in S \subset T. \tag{8}$$

Proof. Suppose there is an acyclic relation satisfying (8), and let P be a transitive completion. Clearly P still satisfies (8). Define the consideration set mapping Γ by

$$\Gamma(S) = \{\arg\min_P S\} \cup \{c_{obs}(T) \mid S \subseteq T,\ T \in \mathcal{D},\ c_{obs}(T) \in S\},$$

for all $S \in \mathcal{P}(X)$. By construction, Γ satisfies (7). Let c be the choice function arising from maximizing P over $\Gamma(S)$ in each choice problem S.

To complete the proof, we show that c extends c_{obs}. Suppose, by contradiction,

[5] Property (7) also characterizes consideration sets in Manzini and Mariotti (2012) and Cherepanov, Feddersen and Sandroni (2013), who allow cyclic preferences. de Clippel and Rozen (2018c) show these theories are observationally equivalent to one with an acyclic relation in its description, and are testable by enumeration if the data includes all pairs, but NP-hard in general.

that $c(S) \neq c_{obs}(S)$ for some $S \in \mathcal{D}$. Then $\Gamma(S)$ contains at least two elements, and $c(S)$ must be the observed choice from some $T \in \mathcal{D}$ with $S \subset T$. Since P satisfies (8), this implies $c_{obs}(S) P c(S)$, contradicting P-maximality of $c(S)$ in $\Gamma(S)$. *Q.E.D.*

5.2 Reference Dependence

Suppose that the DM is subject to reference dependence. Rubinstein and Salant (2006a, 2006b)'s theory of *Triggered Rationality* posits the DM has a collection $\{P_x\}_{x \in X}$ of reference-dependent preference orderings and a salience ordering $\succ_\sigma$ over the alternatives in X. The most salient alternative in a set determines which preference P_x the DM maximizes.

Before attempting to infer the DM's preferences from his observed choices, we must figure out which preference is being maximized in each choice problem. This requires understanding which alternative is his reference point. Suppose we conjecture x is the most salient alternative in the choice problem S. Then x would also be the most salient alternative in all subsets of S in which it is contained, and any choices observed from those choice problems would all arise from the same preference ordering P_x. Let $P_{S,x}$ denote the Samuelson revealed-preference from those observations: $c_{obs}(R) P_{S,x} a$ if $c_{obs}(R) \neq a$ and $a, x \in R \subseteq S$. If $P_{S,x}$ has any cycles, then x could not possibly be the anchor for the preference used in S: that is, some other alternative in S must be more salient than x. Thus the data can reveal restrictions on relations that are not preferences (in this case, the salience ordering $\succ_\sigma$). These restrictions, which take the UCS form, are not only necessary but also sufficient for consistency with the theory.

Proposition 5 *Observed choices c_{obs} are consistent with Triggered Rationality if and only if there is an acyclic relation $\succ$ satisfying the UCS restrictions:*

$$(9) \qquad \exists y \in S \text{ with } y \succ x, \text{ whenever } P_{S,x} \text{ is cyclic.}$$

Proof. Suppose such an acyclic $\succ$ exists, and let $\succ_\sigma$ be a transitive completion; hence $\succ_\sigma$ still satisfies (9). Let x_i denote the i-th maximal element under $\succ_\sigma$, and let $X_i = \{x_i, x_{i+1}, \ldots, x_n\}$ be those elements weakly less salient than x_i. For each i, define the preference ordering P_{x_i} to be a transitive completion of the Samuelson-revealed preference P_{X_i,x_i}. A transitive completion exists, because x_i being $\succ_\sigma$-maximal in X_i implies P_{X_i,x_i} is acyclic. To complete the proof, we show the choice function generated by $\succ_\sigma$ and the reference-dependent preferences $(P_x)_{x \in X}$ extends c_{obs}. To see this, take any $S \in \mathcal{D}$ and let x_k be the $\succ_\sigma$-most salient (smallest-indexed) element in S. Then $S \subseteq X_k$. By construction, $c_{obs}(S) P_{x_k} y$ for all $y \in S \setminus \{c_{obs}(S)\}$. *Q.E.D.*

Testing a theory using data requires not just checking acyclic satisfiability of a given set of restrictions, but also constructing those restrictions from the data. The number of possible restrictions in (9) may at first seem worrying, since there could be restrictions for each $S \in \mathcal{P}(X)$. But, conveniently, applying the enumeration procedure does not require going through all S's, and testing Triggered Rationality remains tractable. Indeed, only restrictions affecting the viability of candidate-maximal elements for the $|X|-1$ non-singleton sets encountered along the way matters. Remember that the enumeration procedure starts by looking for $x_1 \in X$ whose upper-contour set is unrestricted. For Triggered Rationality, this simply means finding x_1 such that P_{X,x_1} is acyclic (if P_{X,x_1} is acyclic, then so is P_{S,x_1} for any subset S containing x_1, and hence focusing on restrictions arising from X is sufficient). Failing to find such an x_1 means observed choices are inconsistent with the theory. Otherwise, the enumeration procedure looks for $x_2 \in X \setminus \{x_1\}$ whose upper-contour set is unrestricted within $X \setminus \{x_1\}$. Next, one looks for x_2 such that $P_{X \setminus \{x_1\}, x_2}$ is acyclic; and so on. Thus one can simply construct restrictions 'online' for sets encountered as the procedure runs.

As an aside, recall our observation from the end of Section 2 that minor rephrasings of a theory can impact whether a theory is subject to the overfitting issue. Notice that Triggered Rationality does not have this issue as written, but it would if the salience ordering were replaced with a salience *function* satisfying IIA.

5.3 Limited Attention

Recall Masatlioglu, Nakajima and Ozbay (2012)'s theory of Limited Attention from Section 2. In this theory, an observed choice is only revealed preferred to alternatives in the DM's consideration set, which itself must be inferred from the choice data. What can we then learn about preferences from observed choices?

Masatlioglu et al. (2012) provides an answer for full data sets ($\mathcal{D} = \mathcal{P}(X)$). They show that 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 $S \setminus \{z\}$. Indeed, the DM must pay attention to z in S, else property (2b) requires her attention set (and thus her choice) to be the same in S and $S \setminus \{z\}$. Their result does not extend to limited data, as restriction can arise from choice problems that are not related by dropping one alternative. These are redundant with full data, but may be critical with limited data. Still, their underlying argument readily extends to any IIA violation: if the choice from T is available but not chosen from $S \subset T$, then the DM must consider at least one alternative in $T \setminus S$ when choosing from T. Otherwise, (2b) would require $\Gamma(T) = \Gamma(S)$, contradicting that the observed choices differ. The IIA violation thus reveals that $c_{obs}(T)Pz$ for some $z \in T \setminus S$.

More subtly, any violation of the Weak Axiom of Revealed Preference (WARP) reveals information about the DM's preference:

$$\text{(10)} \qquad \begin{array}{l} \text{For all } S, T \in \mathcal{D} \text{ with } c_{obs}(S) \neq c_{obs}(T) \text{ and } c_{obs}(S), c_{obs}(T) \in S \cap T: \\ c_{obs}(S)Pz \text{ for some } z \in S \setminus T \text{ or } c_{obs}(T)Pz' \text{ for some } z' \in T \setminus S. \end{array}$$

To see this, suppose the DM does not consider any option from $S \setminus T$ when choosing from S, and does not consider any option from $T \setminus S$ when choosing from T. Then (2b) would require $\Gamma(S \cap T)$ to equal both $\Gamma(S)$ and $\Gamma(T)$, so that $\Gamma(S) = \Gamma(T)$. This would be impossible to reconcile with observed choices from S and T being different. Consequently, it must be that the DM considers some option in $S \setminus T$ when choosing from S, or considers some option from $T \setminus S$ when choosing from T. Whatever that option is, he does not choose it. Hence $c_{obs}(S)Pz$ for some $z \in S \setminus T$ or $c_{obs}(T)Pz'$ for some $z' \in T \setminus S$. The restrictions in (10) encompass those revealed by IIA violations, which are a special case with S and T related by inclusion. Notice that when S and T are *not* related by inclusion, the restrictions in (10) are neither UCS nor LCS: since $c_{obs}(S) \neq c_{obs}(T)$ and since $S \setminus T$ and $T \setminus S$ are nonempty and disjoint, there is no single element whose upper-contour set (or lower-contour set) is restricted according to Definition 4.

In summary, if choices are consistent with Limited Attention, then there must be an acyclic relation (e.g., the DM's preference ordering) which satisfies all the restrictions that observed choices reveal in (10). The next result shows that the existence of an acyclic relation satisfying (10) guarantees consistency with the theory.

Proposition 6 *Observed choices $c_{obs} : \mathcal{D} \to X$ are consistent with Limited Attention if and only if there exists an acyclic relation P satisfying the restrictions (10).*

Proof. Only sufficiency remains. Suppose an acyclic relation satisfying (10) exists, and let P be a transitive completion (so P still satisfies (10)). Define $\Gamma : \mathcal{P}(X) \to \mathcal{P}(X)$ as follows. For $S \in \mathcal{D}$, $\Gamma(S) = \{c_{obs}(S)\} \cup \{x \in S | c_{obs}(S)Px\}$; and 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 $S \in \mathcal{D}$. It remains to show Γ is well-defined and satisfies (2b).

Suppose Γ is not well-defined. Then for some $S \notin \mathcal{D}$, there are $T, T' \in \mathcal{D}$ 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 $y \in T \setminus T'$. Since $S \subseteq T'$, $y \in T \setminus S$. Moreover, $\Gamma(T) \subseteq S$ implies $y \in T \setminus \Gamma(T)$. By definition of $\Gamma(T)$ for $T \in \mathcal{D}$, this means $yPc_{obs}(T)$. Similarly, if $y \in T' \setminus T$, then $yPc_{obs}(T')$, contradicting that P satisfies (10). Finally, for (2b), we show $\Gamma(S \setminus \{x\}) = \Gamma(S)$ for $S \in \mathcal{P}(X)$ and $x \in S \setminus \Gamma(S)$. There are four cases:

Case 1 ($S \setminus \{x\}, S \in \mathcal{D}$). As $S \in \mathcal{D}$, and $x \notin \Gamma(S)$, we know $xPc_{obs}(S)$. Suppose that $\Gamma(S \setminus \{x\}) \neq \Gamma(S)$. Then $c_{obs}(S) \neq c_{obs}(S \setminus \{x\})$. Applying (10) for choice problems S and $S \setminus \{x\}$, we conclude $c_{obs}(S)Px$, a contradiction.

Case 2 ($S \setminus \{x\} \in \mathcal{D}, S \notin \mathcal{D}$). As $S \setminus \{x\} \in \mathcal{D}$, we know $\Gamma(S \setminus \{x\}) = c_{obs}(S \setminus \{x\}) \cup \{y \in S | c_{obs}(S \setminus \{x\})Py\}$. Since $S \setminus \Gamma(S) \neq \emptyset$, there exists $T \in \mathcal{D}$ with $S \subseteq T$ and $\Gamma(T) \subseteq S$. Because $T \in \mathcal{D}$, $zPc_{obs}(T)$ for all $z \in T \setminus S$. Since $\Gamma(S) = \Gamma(T)$, we know $x \in T \setminus \Gamma(T)$. Hence $xPc_{obs}(T)$. If $\Gamma(S \setminus \{x\}) \neq \Gamma(S) = \Gamma(T)$, then $c_{obs}(S \setminus \{x\}) \neq c_{obs}(T)$ contradicting (10) for the pair of sets T and $S \setminus \{x\}$.

Case 3 ($S \setminus \{x\} \notin \mathcal{D}, S \in \mathcal{D}$). Since $S \in \mathcal{D}$, $\Gamma(S) = c_{obs}(S) \cup \{y \in S | c_{obs}(S)Py\}$. 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 ($S \setminus \{x\}, S \notin \mathcal{D}$). As $S \setminus \Gamma(S) \neq \emptyset$, there is $T \in \mathcal{D}$ with $S \subseteq T$ and $\Gamma(T) \subseteq S$. Since $x \in S \setminus \Gamma(S)$, we have $\Gamma(T) = \Gamma(S) \subseteq S \setminus \{x\}$. Then $\Gamma(S \setminus \{x\}) = \Gamma(T)$ by construction, and equals $\Gamma(S)$ by transitivity. *Q.E.D.*

The restrictions in (10) suggest that the theory might generate rich sets of restrictions. The next proposition confirms this intuition. Hence, while the above proposition provides an intuitive and helpful characterization of Limited Attention, there is neither a simple, systematic method to navigate the revealed-preference restrictions to check acyclic satisfiability, nor is there an entirely different method to make testing always simple. Checking acyclic satisfiability will require a lot of guesswork for some datasets.

Proposition 7 *Limited Attention generates rich sets of restrictions.*

Proof. Let Z be any finite set, and let $\mathcal{R}$ be any collection of basic disjunctive restrictions on Z. For each $r \in \mathcal{R}$, we use z_1^r, z_2^r, z_3^r, and z_4^r to denote the elements of Z so that r corresponds to "z_1^r must be ranked above z_2^r, or z_3^r must be ranked above z_4^r." Let X be derived from Z by adding a new option a_r for each restriction $r \in \mathcal{R}$. Consider then the following observed choices, for each $r \in \mathcal{R}$:

S	$a_r z_1^r z_2^r z_3^r$	$a_r z_1^r z_3^r z_4^r$
$c_{obs}(S)$	z_1^r	z_3^r

Clearly, X and c_{obs} are constructed from Z and $\mathcal{R}$ in polynomial time. Proposition 6 tells us c_{obs} is consistent with Limited Attention if, and only if, $\mathcal{R}$ is acyclically satisfiable. *Q.E.D.*

As should be clear from the definitions, generating rich sets of restrictions means that there exist datasets for which testing consistency is intractable. But there may be interesting classes of datasets for which testing can be done by enumeration. To illustrate this, notice that only LCS restrictions matter whenever $\mathcal{D}$ is closed under intersection (or at least contains the intersection of any two choice problems causing a WARP violation). Indeed, 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'. Thus, one can test consistency with Limited Attention in a manner similar to testing SARP for any dataset satisfying this intersection property.

5.4 Shortlisting

The theory Frugal when Undecided studied in Section 3 is an instance of Manzini and Mariotti (2007)'s theory of *Shortlisting* where P_2 is a known ordering. Shortlisting more generally allows P_2 to be any asymmetric relation. The observation that P_1 must be acyclic remains valid, and it remains the primitive of interest for our analysis.

We start by singling out restrictions on P_1 that specific choice patterns reveal. Suppose we observe $c_{obs}(\{a, x, y\}) = x$ and $c_{obs}(\{b, x, y\}) = y$. Since x and y are chosen in each other's presence, they must be incomparable under P_1, and thus comparable under P_2. To explain this data, x and y cannot *both* survive the shortlist in *both* $\{a, x, y\}$ and $\{b, x, y\}$. This data thus reveals a basic disjunctive restriction: aP_1y or bP_1x. This suggests Shortlisting might generate rich sets of restrictions, which the next result confirms. The proof is available in the Appendix.

Proposition 8 *Shortlisting generates rich sets of restrictions.*

Like for Limited Attention, we can identify a large class of datasets for which Shortlisting is easily testable. A first observation, going back to Manzini and Mariotti (2007, Remark 1), is that without loss of generality, P_2 can be taken to be the revealed preference arising from binary choice problems. Suppose then that we have all binary choice problems in our dataset, making P_2 known.

Here it is helpful to keep in mind the UCS characterization in Section 3 of the theory Frugal when Undecided, where the second rationale is also known, but is an exogenously given ordering. For that theory, when there is an option $x \in S$ which is cheaper than $c_{obs}(S)$, then we learn that some alternative $y \in S \setminus \mathcal{S}(x)$ eliminates x from the shortlist (recall that $\mathcal{S}(x)$ is the union of all observed choice problems where x is chosen). The analogy for Shortlisting is that when there is an option $x \in S$ which is preferred to $c_{obs}(S)$, based on observing $x = c_{obs}(\{x, c_{obs}(S)\})$, then there must be some $y \in S \setminus \mathcal{S}(x)$ that eliminates x from the shortlist. Formally,

$$\begin{aligned} &\text{For all } x \in S \in \mathcal{D} \text{ with } x \neq c_{obs}(S) \text{ and } c_{obs}(\{x, c_{obs}(S)\}) = x, \\ &\text{there is } y \in S \setminus \mathcal{S}(x) \text{ such that } yP_1x \end{aligned} \tag{11}$$

These are not quite all the testable implications, even with all binary choices. As P_2 may have cycles, P_1 must eliminate at least one element from any such cycle to ensure choices are well defined. Since the data provides us with a complete P_2, any cycle in P_2 includes a cycle of three options. Say that a triplet $\{a, b, c\}$ has *pairwise-cyclic choices* if a is chosen from $\{a, b\}$, b is chosen from $\{b, c\}$ and c is chosen from $\{a, c\}$. The next result shows that if our dataset also includes any triplet with pairwise-cyclic choices, then the above UCS restrictions do encapsulate consistency with the theory. Without knowing the choices from such triplets, more complex restrictions arise.

Proposition 9 *If $\mathcal{D}$ contains all pairs and all triplets with pairwise-cyclic choices, then c_{obs} is consistent with Shortlisting if, and only if, the UCS restrictions (11) are acyclically satisfiable.*

Proof. It remains to show sufficiency. Let P be an acyclic relation satisfying (11). We can assume without loss of generality that P is an ordering. For all $x, y \in X$, say yP_2x if $c_{obs}(\{x, y\}) = y$, and define P_1 by yP_1x if yPx and there is $S \in \mathcal{D}$ such that $y \in S \setminus \mathcal{S}(x)$ and $c_{obs}(\{x, c_{obs}(S)\}) = x$. We show $c_{obs}(S) = \max(\max(S, P_1), P_2)$, for all $S \in \mathcal{D}$. First, $c_{obs}(S) \in \max(S, P_1)$: there is no $x \in S$ with $xP_1c_{obs}(S)$, as $S \subseteq \mathcal{S}(c_{obs}(S))$. Second, $c_{obs}(S)$ is P_2-maximal in $\max(S, P_1)$: if $x \in S$ were P_2-superior to $c_{obs}(S)$, i.e. $c_{obs}(\{x, c_{obs}(S)\}) = x$, then there would be $y \in S$ with yP_1x.

Finally, we show $\max(\max(S, P_1), P_2)$ is single-valued for all $S \subseteq X$. Notice $\max(S, P_1)$ is nonempty since P_1 is acyclic. As P_2 is complete, $\max(\max(S, P_1), P_2)$ is singleton if, and only if, there is $x \in \max(S, P_1)$ such that xP_2y for all $y \in \max(S, P_1)$. So if it is not singleton, there is a P_2-cycle in $\max(S, P_1)$, and one can find $\{a, b, c\} \subseteq \max(S, P_1)$ with $aP_2bP_2cP_2a$. As it has pairwise-cyclic choices, $\{a, b, c\} \in \mathcal{D}$. To fix ideas, say $c_{obs}(\{a, b, c\}) = a$ (a similar argument applies in the other two cases). But $c_{obs}(\{a, c\}) = c$ since cP_2a, and hence bP_1c, which contradicts $c \in \max(S, P_1)$. *Q.E.D.*

What if the dataset is not as rich as Proposition 9 requires? An obvious corollary is that observed choices are consistent with Shortlisting if, and only if, they can be extended to some c_{obs} having all binary choice problems and any triplets with pairwise-cyclic choices, such that the resulting UCS restrictions (11) are acyclically satisfiable. Though this is a concise characterization for any dataset, there could still be many pairs (or corresponding triplets) missing if one is unlucky, or did not construct the dataset with Proposition 9 in mind. Proposition 8 tells us that there is no good way to get around the complexity: testing is necessarily hard for some datasets.

References

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

Barberà, Salvador, Geoffroy de Clippel, Alejandro Neme, and Kareen Rozen (2020), Order-k Rationality, *mimeo.*

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

de Clippel, Geoffroy, and Kareen Rozen (2018a), Consumer Theory with Misperceived Tastes, *SSRN Working Paper 3223654*.

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

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

de Clippel, Geoffroy, and Kareen Rozen (2019), Bounded Rationality and Limited Datasets, *mimeo*.

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

Lleras, Juan, Yusufcan Masatlioglu, Daisuke Nakajima, and Erkut Ozbay (2017), When More is Less: Choice by Limited Consideration, *Journal of Economic Theory*, 170, 70–85.

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

Manzini, Paola, and Marco Mariotti (2012), Categorize Then Choose: Boundedly Rational Choice and Welfare, *Journal of the European Economics Association*, 10(5), 1141–1165.

Masatlioglu, Yusufcan, Daisuke Nakajima and Erkut Ozbay (2012), Revealed Attention, *American Economic Review*, 102(5), 2183–2205.

Möhring, Rolf, Martin Skutella, and Frederik Stork (2004), Scheduling With And/Or Precedence Constraints, *Siam Journal on Computing*, 33(2), 393–415.

Richter, Marcel (1966), Revealed Preference Theory, *Econometrica*, 34, 635–645.

Rubinstein, Ariel and Yuval Salant (2006a), A Model of Choice from Lists, *Theoretical Economics*, 1, 3–17.

Rubinstein, Ariel and Yuval Salant (2006b), Two Comments on the Principle of Revealed Preference, *mimeo*.

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

Sen, Amartya (1971), Choice Functions and Revealed Preference, *The Review of Economic Studies*, 38, 307–317.

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

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

Appendix

Proof of Proposition 3. Fix an instance of SAT3 with a set V of variables and a set $\mathcal{C}$ of clauses. The 3 literals (a variable $v \in V$ or its negation $\bar{v}$) in a clause C are denoted ℓ_i^C for $i = 1, 2, 3$. C is true if at least one of ℓ_1^C, ℓ_2^C or ℓ_3^C is true. Let Z be the set with all variables v and negations $\bar{v}$; all clauses C; options x_C, x'_C, x''_C, y_C, z_C, per clause C; options a_v, b_v per variable v; and an option t. Let $\mathcal{R}$ be:

(i) $r_{\{(v,t)\}\vee\{(a_v,b_v)\}}$ and $r_{\{(\bar{v},t)\}\vee\{(b_v,a_v)\}}$,

(ii) $r_{\{(t,C)\}\vee\{(y_C,z_C)\}}$ and $r_{\{(t,C)\}\vee\{(z_C,y_C)\}}$, for each clause C,

(iii) $r_{\{(C,\ell_1^C)\}\vee\{(x_C,x'_C)\}}$, $r_{\{(C,\ell_2^C)\}\vee\{(x'_C,x''_C)\}}$, and $r_{\{(C,\ell_3^C)\}\vee\{(x''_C,x_C)\}}$, for each clause C.

Clearly, Z and $\mathcal{R}$ are derived in polynomial time. Since the theory generates rich sets of restriction, one can construct (in polynomial time as well) a superset X of Z and an observed choice function c_{obs} on X such that c_{obs} is consistent with the theory if, and only if, $\mathcal{R}$ is acyclically satisfiable. To conclude the proof, we show that 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 satisfies $\mathcal{R}$: place from worst to best, first all variables v that are true, then $\bar{v}$ for each false variable v, then all clauses C, then x_C, x'_C and x''_C in an order that respects surviving restrictions in (iii) above (which is possible since at most two of the three restrictions survive, given that at least one literal of C is ranked below C), then t, then b_v for true v's, then a_v for all v's, then b_v for false v's, then y_C and z_C for all clauses C (in any order), 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 an acyclic completion of P; this will still satisfy $\mathcal{R}$). Any variable v ranked below t is declared true, while all other variables v are declared false. Given that P is acyclic on $\{a_v, b_v\}$, we know from (i) that at most one of v or $\bar{v}$ can be ranked below t. For each clause C, given that P is acyclic on $\{x_C, x'_C, x''_C\}$, it must be by (iii) that at least one of its literal is ranked below C. Given that P is acyclic on $\{y_C, z_C\}$, it must be by (ii) that C is ranked below t. Thus we have a truthful assignment for the instance of SAT3. *Q.E.D.*

Proof of Observation 1. Take any finite Z and any collection $\mathcal{R}$ of basic disjunctions on Z. Consider the case of mixed binary LCS and UCS restrictions. Construct $Z' \supset Z$ by adding a_r and b_r for each $r \in \mathcal{R}$. We claim acyclic satisfiability of $\mathcal{R}$ is equivalent to acyclic satisfiability of the collection $\mathcal{R}'$ of binary LCS and UCS restrictions given by $r_{\{(z_1^r,a_r)\}\vee\{(b_r,a_r)\}}$, $r_{\{(a_r,z_4^r)\}\vee\{(a_r,b_r)\}}$, $r_{\{(b_r,z_2^r)\}\vee\{(b_r,a_r)\}}$, and

$r_{\{(z_3^r,b_r)\}\vee\{(a_r,b_r)\}}$, for each $r \in \mathcal{R}$. Suppose P (without loss, an ordering) satisfies $\mathcal{R}'$. If $a_r P b_r$, then $z_1^t P a^r$ and $b_r P z_2^r$, thus $z_1^r P z_2^r$ by transitivity; while if $a_r P b_r$, then similarly $z_3^r P z_4^r$. Conversely, if P satisfies $\mathcal{R}$, extend it by $z_1^r P a_r P_1 b_r P z_2^r$ if $z_1^r P z_2^r$ and by $z_3^r P b_r P_1 a_r P z_4^r$ otherwise. By the hypothesis, one can construct (in polynomial time) a set $X \supseteq Z' \supseteq Z$ and choices c_{obs} from some subsets of X, such that c_{obs} is consistent with the theory if, and only if, $\mathcal{R}'$ (equivalently, $\mathcal{R}$) are acyclically satisfiable.

Now consider the case of UCS-or-LCS restrictions. Note that acyclic satisfiability of $\mathcal{R}'$ (and thus of $\mathcal{R}$, by the above paragraph) amounts to acyclic satisfiability of the collection $\mathcal{R}''$ of UCS-or-LCS restrictions: $r_{\{(z_1^r,z_2^r)\}\vee\{(z_2^r,z_3^r)\}}$ and $r_{\{(z_2^r,z_1^r)\}\vee\{(z_1^r,z_3^r)\}}$, for each UCS restriction $r_{\{(z_1^r,z_3^r)\}\vee\{(z_2^r,z_3^r)\}} \in \mathcal{R}'$; and $r_{\{(z_1^r,z_2^r)\}\vee\{(z_2^r,z_3^r)\}}$ and $r_{\{(z_1^r,z_3^r)\}\vee\{(z_3^r,z_2^r)\}}$, for each LCS restriction $r_{\{(z_1^r,z_2^r)\}\vee\{(z_1^r,z_3^r)\}} \in \mathcal{R}'$. It is easy to see an acyclic relation P (without loss, an ordering) satisfies $\mathcal{R}$ if, and only if, it satisfies $\mathcal{R}''$. So by the hypothesis, one can construct (in polynomial time) a set $X \supseteq Z$ and choices c_{obs} from some subsets of X, such that c_{obs} is consistent with the theory if, and only if, $\mathcal{R}''$ (or equivalently $\mathcal{R}$) is acyclically satisfiable. *Q.E.D.*

PROOF OF PROPOSITION 8. Given Z and $\mathcal{R}$, derive X from Z by adding six options, a_r, b_r, c_r, d_r, e_r and f_r, for each $r \in \mathcal{R}$. We again use z_1^r, z_2^r, z_3^r, and z_4^r to denote the elements in a restriction r. Consider the following observed choices, for each $r \in \mathcal{R}$:

S	$a_r c_r$	$b_r d_r$	$e_r z_2^r$	$f_r z_4^r$	$a_r e_r z_2^r$	$b_r f_r z_4^r$	$a_r b_r c_r z_1^r$	$a_r b_r d_r z_3^r$
$c_{obs}(S)$	a_r	b_r	z_2^r	z_4^r	e_r	f_r	b_r	a_r

X and c_{obs} can be constructed from Z and $\mathcal{R}$ in polynomial time. We conclude by showing c_{obs} is consistent with Shortlisting if, and only if, $\mathcal{R}$ is acyclically satisfiable.

Necessity. Suppose (P_1, P_2) generate c_{obs} under Shortlisting. Since z_4^r and f_r are chosen in the other's presence, they are P_1-incomparable. So $z_4^r = c_{obs}(\{f_r, z_4^r\})$ means $z_4^r P_2 f_r$, and $z_4^r \neq c_{obs}(\{b_r, f_r, z_4^r\})$ means $b_r P_1 z_4^r$. Similarly, $a_r P_1 z_2^r$. From the last two choices, a_r and b_r are P_1-incomparable, and thus P_2-comparable. If $b_r P_2 a_r$, then the last choice requires $d_r P_1 b_r$ be P_1-inferior to d_r or z_3^r. As $b_r = c_{obs}(\{b_r, d_r\})$, $z_3^r P_1 b_r$. If $a_r P_2 b_r$, similar reasoning gives $z_1^r P_1 a_r$. Hence $z_1^r P_1 a_r$ or $z_3^r P_1 b_r$. Complete P_1 into an ordering P. Since $a_r P_1 z_2^r$ and $b_r P_1 z_4^r$, we have $z_1^r P z_2^r$ or $z_3^r P z_4^r$.

Sufficiency. Without loss, let P be an ordering satisfying $\mathcal{R}$. Define P_1 on X by $a_r P_1 z_2^r$, $b_r P_1 z_4^r$, $c_r P_1 z_1^r$, $d_r P_1 z_3^r$, $e_r P_1 a_r$, $f_r P_1 b_r$, $z_1^r P_1 a_r$ if $z_1^r P z_2^r$, and $z_3^r P_1 b_r$ if $z_3^r P z_4^r$, for each $r \in \mathcal{R}$. Notice P_1 is acyclic.[6] Let P_2 be an ordering on X with elements of Z at the top in any order; right below a_r's and b_r in any way such that $a_r P_2 b_r$ if $z_1^r P z_2^r$, and $b_r P_2 a_r$ if $z_3^r P z_4^r$ but not $z_1^r P z_2^r$; and below all remaining options, in any order. It is easy to see (P_1, P_2) generates c_{obs} under Shortlisting. *Q.E.D.*

[6]E.g., it extends to an ordering with $y_1 = \min_P Z$ at the bottom; above, elements of $\{a_r | z_2^r = y_1\} \cup \{b_r | z_4^r = y_1\}$ in any order; above, the analogous block of y_2, a_r's and b_r's for $y_2 = \min_P Z \setminus \{y_1\}$; and so on, till Z is exhausted; above c_r's and d_r's in any order; and then e_r's and f_r's in any order.