Testable Implications of Some Classic Assignment Methods

Geoffroy de Clippel* and Kareen Rozen†

July 2018

Abstract

We study the testable implications of serial dictatorship, stable many-to-one matchings, and the core of housing markets. We show that serial dictatorship is easy to test, and explain how elements of the power ranking between agents can be identified. Stability is also easy to test for an interesting class of many-to-one matching problems, and is tightly related to serial dictatorship. We provide an insightful characterization of the core of Shapley and Scarf (1974)'s housing markets using revealed top-trading cycles. This characterization proves useful in many examples, but is also used to prove that testing the core is generally NP-hard.

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

1 Introduction

Rationality, as captured by the maximization of a preference ordering, is a building block of most economic models. This creates a difficulty for testing: while outcomes may be observable, preferences are not. The testable implications of a model, or its empirical content, describe patterns of outcomes that can arise under the theory for *some* preference profile. For instance, the Strong Axiom of Revealed Preference (SARP) provides the empirical content of rational choice with a single decision maker. Building on this classic result, this paper studies simple models of *interactive* decision making: serial dictatorship, stable many-to-one matching, and the core/competitive equilibrium of housing markets.[1]

Under serial dictatorship, a collective bounty of indivisible objects have to be allocated to participants, each of whom needs at most one object. Who consumes what is determined through the combination of power and preferences: starting with the most powerful person and then moving down the ladder, each participant picks her most preferred object (or none if she prefers) from among the remaining ones. As coined by Piccione and Rubinstein (2007), this is the economics of the jungle.

Shapley and Scarf's (1974) housing markets also study the assignment of indivisible objects to participants with unit demand. A first difference is that housing markets come with property rights. Second, rather than objects being allocated through the exercise of power, houses are allocated through voluntary trades. Thus, no group of participants should be better off by trading among themselves. This corresponds to the classic core solution,

[1] Other works on testable implications in interactive settings include Brown and Matzkin (1996) on general equilibrium, Sprumont (2000) on Nash equilibrium and Pareto efficiency, Ray and Zhou (2001) on backwards induction, Echenique (2008) on stable matchings (which we expand upon later), Forges and Minelli (2009) on Nash equilibrium in market games, Carvajal, Deb, Fenske and Quah (2013) on Cournot oligopoly, and Chambers and Echenique (2014) on Nash bargaining.

which also happens to coincide with the Walrasian equilibrium.[2]

A third setting of interest is two-sided matching, in which participants are divided in two groups (e.g. firms and workers in a job market). Instead of assigning objects to participants, the problem is to match participants on one side of the market to participants on the other side. In many-to-one markets, participants on at least one side of the market do not want to be matched (or simply cannot be matched) to more than one participant on the other side. For instance, exclusivity clauses may prevent a full-time employee from working for another firm. While we focus on the job market as an application, there are many others: a student might be forbidden to enroll in multiple colleges at the same time, it may be illegal to be married to more than one person at the same time, etc. Once again, we will focus on the core, which also coincides with the set of stable matchings.

Our results characterizing the testable implications of these settings build on the general methodology of de Clippel and Rozen (2018). They introduce the concept of *acyclic satisfiability* as a natural extension of SARP. Quite simply, a collection of restrictions on a relation is acyclically satisfiable if there is an acyclic relation satisfying them. Acyclic satisfiability boils down to SARP when restrictions take the form of simple comparison (e.g. a must be ranked above b). For theories other than rational choice, the data often reveals more complex restrictions about an underlying ordering (e.g. a preference ordering that the decision maker imperfectly maximizes). For instance, a choice pattern may reveal that a is inferior to *either b or c*. This restricts the upper contour set of a, by requiring that it intersects $\{b, c\}$. Using a simple enumeration procedure, testing acyclic satisfiability is comparable to testing SARP when facing only upper-contour set restrictions. By contrast, testing acyclic satisfiability becomes NP-hard when allowing for *mixed* sets of restrictions, with some pertaining to upper-contour sets, and others to lower-contour sets.

[2]Preferences are assumed to be strict throughout the paper, hence the equivalence.

While de Clippel and Rozen (2018) apply these ideas to theories of bounded rationality for a single decision maker, this paper shows that they also prove useful for testing theories of interactive decision making with rational participants. Under serial dictatorship, we use observed assignments of objects to infer restrictions on the underlying power relation. To gain some intuition, suppose we conjecture that i is the most powerful individual. Then her assigned objects are systematically her most preferred. If the data reveals that this is impossible (as it would lead to a violation of SARP for i), then one must conclude that at least one other individual is more powerful than i. This corresponds to an upper-contour set restriction on the underlying power relation. We extend this reasoning in Section 2, and show that acyclic satisfiability of restrictions such as these captures the empirical content of serial dictatorship. With all restrictions pertaining to upper-contour sets, testing can be performed using a simple enumeration procedure.

Building on Echenique (2008), the work of Kalyanaraman and Umans (2008) establishes that testing stability in matching is NP-hard (even for the case of one-to-one matching). By contrast, when studying stable many-to-one matchings in Section 3, we use our serial-dictatorship result to provide an insightful and tractable testing procedure when all firms are conjectured to share the same ranking over workers.[3] Indeed, the set of stable matchings corresponds in that case to the outcome of serial dictatorship applied to workers, with firms as objects and the power relation given by the firms' common (unknown) ranking of workers.

In Section 4, we characterize the testable implications of core compatibility in housing markets by exploiting the well-known fact that the core coincides with the outcome of Gale's top-trading-cycle algorithm (Shapley

[3]Though restrictive, such an assumption is not uncommon. In fact, it is sometimes assumed that agents on *each* side of the market share a common preference, thereby having an objective ranking of the firms in addition to an objective ranking of the workers. In that case, a stable matching is *assortative*, with the best workers employed in the best firms.

and Scarf, 1974). In a nutshell, one must find a way to rank revealed trading cycles in each observation of the dataset. By consuming a house h, a participant not only reveals that she prefers h over all other houses in her trading cycle, but also over all houses appearing in later cycles. Core compatibility is equivalent to verifying SARP for the individual revealed preferences induced in this way. While simplifying matters greatly, one must still explore all the possible ways to rank revealed trading cycles, which can be demanding. Is there a much more efficient way to proceed? While the test can be refined marginally, we prove that there is no hope to do much better. Indeed, we establish that testing core compatibility is NP-hard, by showing that the problem is reducible in polynomial time to de Clippel and Rozen (2018)'s NP-hard problem of testing acyclic satisfiability of mixed sets of restrictions.

2 Serial Dictatorship

There is a group I of *individuals*, and a set X of conceivable *objects*, which are assumed to be distinct and indivisible. An *assignment problem* is described by the subset $S \subseteq X$ of objects to assign. A *feasible assignment* for S is a vector $(x_i)_{i \in I}$ whose components each belong to $S^* \equiv S \cup \{\emptyset\}$, and which does not assign the same object to two different individuals; that is, if $x_i = x_j$ for some $i \neq j$, then $x_i = x_j = \emptyset$. The symbol $\emptyset$ represent the absence of consumption. A *consumption function* c associates to each assignment problem S a feasible allocation $c(S)$ for S.

Under *Serial Dictatorship*, individuals are endowed with preference orderings $(\succ_i)_{i \in I}$ on $X \cup \{\emptyset\}$, and are themselves ranked according to a power ordering $\succ_p$. Consumption is determined by first letting the most powerful individual pick her preferred option in S^*; then letting the second-most powerful individual pick her preferred option among those that remain; and so on so forth. Formally, $c(S)$ is defined by induction using the following formula:

$$c_i(S) = \arg\max_{\succ_i} A_i(S) \cup \{\emptyset\},$$

where

$$A_i(S) = \{x \in S \mid \nexists j \text{ with } j \succ_p i \text{ and } x = c_j(S)\}$$

represents the set of objects available to individual i after the more powerful individuals have made their consumption decisions. An individual might not consume an object, either because she prefers not to consume, or because there are too few objects to go around.

A *dataset* $\mathcal{D}$ is a collection of assignment problems, and an *observed consumption function* c_{obs} associating to each allocation problem $S \in \mathcal{D}$ a feasible allocation $c_{obs}(S)$ for S. The observed consumption function is *consistent with serial dictatorship* if there exists a consumption function c under that theory such that $c_{obs}(S) = c(S)$ for all $S \in \mathcal{D}$.

We apply de Clippel and Rozen's (2018) testing methodology by figuring out what key information one may infer about the power ordering from observed consumption. Consider a group G of individuals, and i one of its members. As c_{obs} is vector-valued, let $c_{obs,i}(S)$ denote the object received by i under the observed assignment problem S. If i is most powerful within G, then her choice in each observed assignment problem $S \in \mathcal{D}$ must be consistent with maximizing a preference given the objects remaining once those allocated to individuals outside of G have been removed. Otherwise, she would not be the most powerful individual in G.

Formalizing this idea, for each group G and assignment problem S, define

$$S^*(G, S) = \{x \in S \mid \nexists j \in I \setminus G : c_{obs,j}(S) = x\} \cup \{\emptyset\}.$$

Then, let $\mathcal{R}_{SD}(c_{obs})$ denote the following collection of restrictions on the power relation: for each combination (G, i) with $G \subseteq I$ and $i \in G$, if there is a SARP violation[4] in the auxiliary data comprising the choice of $c_{obs,i}(S)$ from $S^*(G,S)$ for each $S \in \mathcal{D}$, then there exists $j \in G$ more powerful than i. Clearly, consistency with Serial Dictatorship means there must exist an acyclic relation that satisfies $\mathcal{R}_{SD}(c_{obs})$ or, using de Clippel and Rozen (2018)'s terminology, $\mathcal{R}_{SD}(c_{obs})$ must be *acyclically satisfiable.* Indeed, the restrictions are satisfied in that case by the underlying true power relation over individuals.

Before describing how acyclic satisfiability of these restrictions can be easily tested, we first establish that it encapsulates the empirical content of serial dictatorship.

Proposition 1. *The observed assignment function c_{obs} is consistent with serial dictatorship if and only if $\mathcal{R}_{SD}(c_{obs})$ is acyclically satisfiable.*

Proof. Necessity was established in the above paragraph. As for sufficiency, suppose that there is an acyclic relation $\succ_p$ that satisfies the restrictions in $\mathcal{R}_{SD}(c_{obs})$. We can assume without loss of generality that $\succ_p$ is an ordering, as satisfying restrictions is preserved under completion. For each individual i, let $G(i)$ be the set of individuals that are no more powerful than i:

$$G(i) = \{j \in I | j = i \text{ or } i \succ_p j\}.$$

Observe that the auxiliary data where i picks $c_{obs,i}(S)$ out of $S^*(G(i), S)$ for each $S \in \mathcal{D}$ satisfies SARP. Otherwise, $\mathcal{R}_{SD}(c_{obs})$ contains a restrictions of the form 'some $j \in G(i)$ is more powerful than i', which cannot be since $\succ_p$ violates that restriction. Let then $\succ_i$ be an ordering that is consistent with the revealed preferences arising from this auxiliary data, that is, such that $c_{obs,i}(S) \succ_i x$, for all $x \in S^*(G(i), S)$ and each $S \in \mathcal{D}$. To conclude the proof, we check that the assignment function generated by the power relation $\succ_p$

[4]SARP stands for the classic Strong Axiom of Revealed Preference. Violating it means that Samuelson's revealed preference associated to the auxiliary data is cyclic: there exists a sequence $(x_1, \ldots, x_K)$ of objects in S^* such that $x_k = c_{obs,i}(S^k)$ and $x_{(k+1) mod K} \in S^*(G, S^k)$ for each $k = 1, \ldots, K$.

and the preference profile $(\succ_i)_{i\in I}$ coincides with $c_{obs}(S)$, for each $S \in \mathcal{D}$. Indeed, if the object $x \in S$ has not been claimed by any more powerful individual (that is, $x \neq c_{obs,j}(S)$ for all $j \succ_p i$), then $x \in S^*(G(i), S)$ and $c_{obs,i}(S) \succ_i x$, by construction of $\succ_i$. $\square$

Proposition 1 is based on the idea that we must be able to construct a power ranking of individuals that leads to choices consistent with the data. There may be multiple such power rankings. When can we infer that an individual i is *revealed more powerful* than j, meaning that *every* power ranking which is consistent with the data ranks i above j? A small variation of Proposition 1 allows us to test this as well. Indeed, by appending the contrary restriction 'j is more powerful than i' to the original set of restrictions $\mathcal{R}_{SD}(c_{obs})$, it is easy to see that i is revealed more powerful than j if and only if the augmented restrictions *fail* to be acyclically satisfiable. Taking this a step further, if i is revealed more powerful than j, then i's assigned object is revealed preferred to j's according to i's preference ranking.

It turns out that testing acyclic satisfiability of $\mathcal{R}_{SD}(c_{obs})$ is quite simple, both in concept and computational complexity (the test is doable in polynomial time in the size of I and $\mathcal{D}$). Moreover, these features extend to testing revealed power. The simplicity of testing $\mathcal{R}_{SD}(c_{obs})$, from which the simplicity of testing revealed power easily follows, relies on two observations. First, all restrictions in $\mathcal{R}_{SD}(c_{obs})$ restrict the upper-contour set of the power ordering. Indeed, the restriction for a (G, i) combination, when one applies, requires some $j \in G \setminus \{i\}$ to be more powerful than i. Testing acyclic satisfiability can be done through de Clippel and Rozen (2018)'s enumeration procedure for upper-contour set (or lower contour-set restrictions), as we explain in the next paragraph. Second, while figuring out the set of all restrictions appearing in $\mathcal{R}_{SD}(c_{obs})$ would be demanding (as there are exponentially groups of individuals $G \subseteq I$ to consider), one need only identify the small subset of restrictions that matter for the implementation the enumeration procedure.

Here is how these broad ideas apply to test acyclic satisfiability in the present context. First, identify a candidate most-powerful individual in I, namely, an $i^1 \in I$ whose upper-contour set is not restricted in any way, that is, who is not involved at the 'bottom' of any restriction in $\mathcal{R}_{SD}(c_{obs})$. Notice that it suffices to check that (I, i^1) is not associated with a SARP violation, for if (G, i) led to a SARP violation for some $G \subset I$ and $i \in G$, then (I, i) would have a SARP violation too. Of course, if acyclic satisfiability holds, than such an individual i_1 can be found, while if every person in I must be less powerful than someone else, then acyclic satisfiability fails. Next, we identify a candidate most-powerful person in $I \setminus \{i^1\}$; namely, an $i^2 \in I \setminus \{i^1\}$ for whom there is no other person in $I \setminus \{i^1\}$ who is required to be more powerful according to $\mathcal{R}_{SD}(c_{obs})$. Again, this only requires finding a $i^2 \in I \setminus \{i^1\}$ such that $(I \setminus \{i^1\}, i^2)$ is not associated with a SARP violation. Continuing in this manner, one can enumerate all of I (with the interpretation that individuals enumerated earlier are more powerful) if and only if $\mathcal{R}_{SD}(c_{obs})$ is acyclically satisfiable. This test takes at most polynomial time, as the number of iterations of the procedure is bounded by the number of individuals in I, and only restrictions for those sets encountered along the path of the procedure need to be checked.

3 Stable Matchings

The technique developed in the previous section also proves useful when testing the core of some interesting matching problems. A *many-to-one matching problem* is characterized by two sets of agents. To fix ideas, say $\mathcal{F}$ is a set of firms and $\mathcal{W}$ is a set of workers. Each firm $f \in \mathcal{F}$ has a number $n_f \geq 0$ of positions to fill, where n_f is an integer. A *matching* is a function μ that associates to each worker $w \in \mathcal{W}$ an element of $\mathcal{F} \cup \{\emptyset\}$ such that $\mu^{-1}(f) = \{w \in \mathcal{W} \mid \mu(w) = f\}$ contains at most n_f elements for each firm f. Having $\mu(w) = \emptyset$ means that w is left unemployed.

Assume each firm f has a preference ordering $\succ_f$ over $\mathcal{W} \cup \{\emptyset\}$, and each worker w has a preference ordering $\succ_w$ over $\mathcal{F} \cup \{\emptyset\}$. This rules out the possibility of complementarity and substitutability between workers in a firm: if free to choose a subset of n_f workers from any set of prospective hires, a firm f systematically picks the top n_f options according to $\succ_f$. This is a classic assumption, which we maintain in this work, although we recognize it is restrictive beyond the special case of one-to-one matching. Workers ranked below $\emptyset$ according to $\succ_f$ are considered unworthy of hiring by firm f. Similarly, firms ranked below $\emptyset$ according to $\succ_w$ are considered by worker w to be unworthy of working for. While this model uses language relating to the job market, many other applications of such bilateral matching problems are discussed in the literature, including school assignment with priorities, resident-hospital matching, and marriage (oftentimes the special case of one-to-one matching).

A matching μ is *stable* if (a) no firm f hires an undesirable worker: $w \succ_f \emptyset$, for all $w \in \mu^{-1}(f)$, (b) no worker w works for an undesirable firm: $\mu(w) \succ_w \emptyset$ if $\mu(w) \neq \emptyset$, (c) there is no firm-worker pair (f, w) such that the worker prefers f over her current situation ($f \succ_w \mu(w)$) and the firm prefers w over leaving a vacant spot or over another putative hire ($w \succ_f \emptyset$ if $\mu^{-1}(f)$ is strictly inferior to n_f, or $w \succ_f w'$ for some $w' \in \mu^{-1}(f)$). As is well-known, the set of stable matching coincides with the core, that is, adding the possibility of deviations by larger coalitions of workers and firms does not further reduce the set of robust matchings.

The question we are interested in is the following. Suppose we observe different matchings crystallize for different configurations of labor supply (the subset of workers available) and demand (the number of positions open in each firm). How do we test whether there exist preference orderings for firms and workers such that these observed matchings are stable? Building on Echenique (2008), Kalyanaraman and Umans (2008) establish that this problem is NP-hard (even for the case of one-to-one matching). By contrast, we show that an insightful and tractable testing procedure is available when all firms are conjectured to share the same preference over workers.[5] Firms are thus assumed to have equal information about workers, and to share the same assessment function over worker characteristics. Under this assumption, professionals in the industry agree on an objective way to rank workers, but the modeler does not know this common ranking a priori (for instance, she may not know which characteristics of workers are most valued in the industry, or would need to interview the workers herself). Nonetheless, the modeler may be able to partially identify the industry ranking of workers by scrutinizing observed matchings. Workers may have different rankings over firms (for instance, due to locational preferences or firm culture), which the modeler similarly does not know but may be able to partially infer from the data.

An *initial market condition* is characterized by a collection of job openings $(n_f)_{f\in\mathcal{F}}$ and a set $\mathcal{W}'\subseteq\mathcal{W}$ of prosective employees. Note that n_f can be equal to zero for some f, meaning that such firms are not looking to hire under that initial condition. A *dataset* $\mathcal{D}$ is the collection of initial market conditions for which one has observed matchings occur. An *observed matching function* μ_{obs} associates a matching of prospective employees to open positions (or self-employment), for each initial market condition. It is *consistent with stability* if there exist a common preference ordering $\succ$ over workers, and a profile $(\succ_w)_{w\in\mathcal{W}}$ of preference orderings of the firms, such that the observed matching associated to each initial market condition in $\mathcal{D}$ is stable.

Testing for consistency with stability can be done using the techniques developed in the previous section. This follows from the close link between serial dictatorship and stability when firms share a common preference. In-

[5]Though restrictive, such an assumption is not uncommon. In fact, it is sometimes assumed that agents on *each* side of the market share a common preference, thereby having an objective ranking of the firms in addition to an objective ranking of the workers. In that case, a stable matching is *assortative*, with the best workers employed in the best firms.

deed, such problems admit a unique matching, which can be found as follows. For each firm $f \in F$ and each of the n_f open positions at this firm, define a new 'object', which corresponds to a position at that firm. The set X of all such objects contains $\sum_{f \in \mathcal{F}} n_f$ elements. As the n_f objects constructed for a firm f are all identical, each worker w is indifferent among them. Thus each worker's preference is only a weak ordering over $X \cup \{\emptyset\}$, with strict preferences only across those objects coming from different firms or $\emptyset$. Letting the set of individuals I be the set of workers $\mathcal{W}$, the definition of serial dictatorship and the analysis of the previous section easily generalize to this setting. Indeed, we can apply the serial dictatorship procedure using the firm's common ranking of the workers as their priority. It turns out that the unique stable matching coincides with the serial dictatorship allocation of this auxiliary problem.

Observation. *Matching problems with firms sharing a common preferences admit a unique stable matching, which coincides with the serial dictatorship allocation of the auxiliary allocation problem described above. The analysis of the previous section thus also provides a tractable test for consistency with this theory.*

In this setting, the implicit power relation that the test tries to infer is the firms' common preference, while information inferred about each individual's preferences over objects pertains to workers' preferences over firms. Revealed-power relationships thus correspond to identifying the industry preference over workers; and in turn, this can reveal information about a worker's preference over firms.

4 The Core of Housing Markets

As in Section 2, there is a group I of *individuals*, and a set X of *objects*, which are assumed to be distinct and indivisible. A first difference, though, is that housing markets come with property rights. We will assume that there is an equal number of objects and individuals, and each agent is endowed with one object. Second, rather than objects being allocated through the exercise of power, houses are allocated through voluntary trades. In addition to individual rationality, it will be required that no group of agents would be better off by trading their endowments among themselves. This is the standard definition of the core.

Formalizing these ideas, a *housing market* is described by an assignment ω of individuals to objects: their (observable) initial endowment. Each individual i is assumed to have a preference ordering $\succ_i$ on X. Individual preferences are unknown to the modeler. We write $h \succeq_i h'$ if $h \succ_i h'$ or $h = h'$. A *coalition* is any non-empty subset of X. Coalition S *blocks* an assignment α of individuals to objects if there exists an assignment α' of members of S to those objects they own ($\{\omega_i | i \in S\}$) such that $\alpha'_i(\omega) \succeq_i \alpha_i(\omega)$, for all $i \in S$, with strict preference for at least one $i \in S$. The assignment α belongs to the *core* given ω if it cannot be blocked by any coalition. As is well-known, the core of housing markets with preference orderings is single valued and coincides with the competitive equilibrium (Shapley and Scarf, 1974).

A *dataset* $\mathcal{D}$ is the collection of initial endowments for which one has observed trade occur. An *observed trading function* τ_{obs} associates to each endowment $\omega \in \mathcal{D}$ an assignment $\tau_{obs}(\omega)$ of individuals to objects. It is *core compatible* if there exist preference orderings $(\succ_i)_{i \in I}$ such that $\tau_{obs}(\omega)$ belongs to the core at ω, for all $\omega \in \mathcal{D}$.

The test we develop exploits the fact that the core of housing markets can be obtained by performing Gale's *top-trading cycle* algorithm (Shapley and Scarf, 1974). Each agent first points to their most preferred house. There must be at least one cycle of agents pointing to each others' houses (including trivial cycles in which an agent points to their own house). Any one such cycle is picked, with the corresponding agents and houses from this cycle removed (note that the final assignment will not depend on which cycle is chosen in each step, if multiple cycles exist). The process is then repeated with the remaining agents and their houses, until every house is assigned.

For a given endowment ω and observed trading assignment $\tau_{obs}(\omega)$, let $\mathcal{C}(\omega, \tau_{obs}(\omega))$ be the partition of X into *revealed trading cycles* when moving from ω to τ. Formally, $C \in \mathcal{C}(\omega, \tau_{obs}(\omega))$ if $C = \{\omega_i | \tau_{obs,i}(\omega) \in C\}$, and there is no $C' \subset C$ such that $C' = \{\omega_i | \tau_{obs,i}(\omega) \in C'\}$. Finding revealed trading cycles is easy. Start with any individual, say agent i. Her final consumption is $\tau_{obs,i}(\omega)$, which was the initial endowment of some agent j. If $j = i$ (that is, i does not trade), then we have already found a revealed trading cycle $C = \{\omega_i\}$. If $j \neq i$, then the object $\tau_{obs,j}(\omega)$ which j consumes is the initial endowment ω_k of some agent k. If $k = i$, then we have found a revealed cycle $C = \{\omega_i, \omega_j\}$; otherwise, the object $\tau_{obs,k}(\omega)$ which k consumes is itself the endowment of some other agent, etc. With a finite number of agents and both ω and $\tau_{obs}(\omega)$ one-to-one mappings of individuals to objects, we are bound to end up with an agent whose consumption is i's initial endowment. Thus one is sure to find a revealed trading cycle, and we can repeat the process with remaining individuals to find all the other revealed trading cycles as well.

Since the core can be derived by applying Gale's top-trading procedure, it must be that the object $\tau_{obs,i}(\omega)$ consumed by each agent i is part of the *top*-trading cycle for remaining individuals and objects when agent i 'exits' the procedure. As shown above, the modeler can infer the cycles from the data, but the order in which they exit in Gale's procedure isn't clear a priori.

Suppose the modeler conjectures that, for each $\omega \in \mathcal{D}$, the right order in which the cycles in $\mathcal{C}(\omega, \tau_{obs}(\omega))$ are exited is given by $\succ_\omega$; that is, if $C \succ_\omega C'$, then C exits before C'. Then the data at each $\omega \in \mathcal{D}$ would reveal, for each individual i, that she prefers $\tau_{obs,i}(\omega)$ over all other objects in her revealed trading cycle, as well as all objects in revealed trading cycles that are $\succ_\omega$-inferior. More succinctly, agent i prefers her assigned house $\tau_{obs,i}(\omega)$ to any object belonging to a $\succeq_\omega$-inferior trading cycle. This *revealed preference* must be acyclic for each individual i. Being able to find some profile $(\succ_\omega)_{\omega \in \mathcal{D}}$

of possible exit orderings of cycles which yields acyclic revealed preferences is thus a necessary condition for core compatibility. The next proposition shows that no additional restriction for core compatibility can be inferred from the data, as this condition is also sufficient.

Proposition 2. *τ_{obs} is core compatible if and only if there exist rankings $(\succ_\omega)_{\omega \in \mathcal{D}}$ of revealed trading cycles (elements of $\mathcal{C}(\omega, \tau_{obs}(\omega))$) at each $\omega \in \mathcal{D}$ such that the revealed preferences $\succ_1^*, \ldots, \succ_I^*$ are each acyclic, where $h \succ_i^* h'$ if there exists $\omega \in \mathcal{D}$ such that $\tau_i(\omega) = h$ and $h' \neq h$ belongs to a revealed trading cycle in $\mathcal{C}(\omega, \tau_{obs}(\omega))$ which is $\succeq_\omega$-inferior to the revealed trading cycle to which h belongs.*

Proof. Necessity follows from the discussion above. As for sufficiency, for each $i \in I$, let $\succ_i$ be a preference ordering that respects $\succ_i^*$ comparisons. By definition of $\succ_i^*$, for each $\omega \in \mathcal{D}$, the top-trading cycle procedure applied to $(\succ_i)_{i \in I}$ generates $\tau_{obs}(\omega)$, as desired. Indeed, the $\succ_\omega$-top revealed trading cycle is a top trading cycle according to Gale's procedure. After eliminating the goods involved in that cycle, the second-highest revealed trading cycle according to $\succ_\omega$ is the next top-trading cycle according to Gale's procedure, and so on so forth. $\square$

The next example and ensuing discussion illustrate the proposition.

Example 1. *Consider a problem with three individuals, $I = \{i, j, k\}$, three objects, $X = \{x, y, z\}$, and two observations, $\mathcal{D} = \{\omega, \omega'\}$. Initial endowments and observed trades are given in the following tables*

	i	j	k
ω	x	y	z
ω'	x	z	y

	i	j	k
$\tau_{obs}(\omega)$	y	x	z
$\tau_{obs}(\omega')$	x	y	z

The above proposition allows us to conclude that this data is not core compatible. To see this, first note that the revealed trading cycles are $C_1 = \{x, y\}$

and $C_2 = \{z\}$ at ω, and $C_1' = \{x\}$ and $C_2' = \{y, z\}$ at ω'. To see that the condition in the above proposition is violated, first note that because $\tau_{obs,j}(\omega) = x$ and x, y belong to the same revealed trading cycle at ω, it is revealed that j prefers x to y. Analogously, because $\tau_{obs,i}(\omega) = y$, it is revealed that i prefers y to x. Now consider the possible rankings of the ω'-revealed trading cycles. If C_1' dominates C_2', then x is revealed preferred to y for i, violating acyclicity of i's revealed preference. On the other hand, if C_2' dominates C_1', then y is revealed preferred to x for j, violating acyclicity of j's revealed preference.

A couple of points are worth emphasizing following this simple example. First, it illustrates how testing using the above proposition can be both more insightful and quicker than computing the assignment resulting from Gale's algorithm for all[6] 18 preference profiles and checking whether one matches the observed trading function. Second, it shows that taking into account the rankings of revealed preference cycles is critical for testing. Merely observing that an agent prefers her assigned object over others in the same trading cycle does not create any preference cycle in the above example.[7]

Though insightful and quick to check in many problems, testing using Proposition 2 can remain long and tedious in others (those with many observations that have many revealed cycles), as one must consider all possible rankings of revealed cycles at each observation. It may sometimes be possible to improve the result by finding tractable ways to narrow down the set of possible such rankings. In the example above, for instance, the data at ω reveals that i prefers y over x since it belongs to the same trading cycle. Knowing this, it must be that C_2' dominates C_1'. Hence, such straightforward revealed preference comparisons can reveal information about the ranking of revealed cycles. Unfortunately, one can only go so far. As we now prove, no

[6]More generally, this exhaustive approach would require computing the assignment for $n(n!)$ preference profiles, where n is the number of individuals or objects.

[7]This simpler observation leads to the following conclusion: i prefers y over x, j ranks the three objects alphabetically, and k ranks z above y. This information is compatible with preference orderings for all three agents. Yet the data is not core compatible.

test is tractable for all datasets, as testing core compatibility is NP-hard.

To prove this, we rely on Proposition 3 in de Clippel and Rozen (2018). There it is shown that the class of problems where one must figure out whether there is an acyclic relation satisfying a *mixed* set of binary restrictions, where some are lower-contour set (LCS) restrictions[8] and others are upper-contour set (UCS) restrictions,[9] is NP-hard. We show that every such acyclic-satisfiability problem is reducible in polynomial time to a problem of core compatibility for some observed trading function. While the complete proof of this result is available in the Appendix, a couple of examples can provide some intuition.

Example 2. *Consider a problem with four individuals, $\{i, j, k, \ell\}$, and two observations. The initial endowments and observed trades are as follows:*

	i	j	k	ℓ
ω	y	a	x	z
ω'	a	y	x	z

	i	j	k	ℓ
$\tau_{obs}(\omega)$	y	x	a	z
$\tau_{obs}(\omega')$	z	x	y	a

It is easy to check that τ_{obs} is core compatible by using Proposition 2.[10] *More importantly, this example illustrates how core compatibility can impose an UCS restriction on i's revealed preference: agent i ranks x below either y or z. Indeed, suppose that i ranks x above y. Then it must be that the revealed trading cycle $\{a, x\}$ in ω must exit before the revealed trading cycle $\{y\}$ associated to i. In that case, it must be that k ranks a above y. Hence the revealed trading cycle $\{a, z\}$ in ω' must exit before the revealed trading cycle $\{x, y\}$, in which case i prefers z over x.*

Example 3. *Consider a problem with six individuals, $I = \{i, j, k, \ell, m, n\}$, six objects, $X = \{a, b, c, x, y, z\}$, and two observations, $\mathcal{D} = \{\omega, \omega'\}$. Initial*

[8]Restrictions of the form, x is superior to either y or z, for some x, y, z.

[9]Restrictions of the form x is inferior to either y or z, for some x, y, z.

[10]E.g., the four individuals' revealed preferences are acyclic when ranking the revealed trading cycles as follows: $\{y\}$ above $\{a, x\}$ above $\{z\}$ in ω, and $\{x, y\}$ above $\{a, z\}$ in ω'.

endowments and observed trades are as follows:

	i	j	k	ℓ	m	n
ω	a	x	y	b	z	c
ω'	b	a	z	x	c	y

	i	j	k	ℓ	m	n
$\tau_{obs}(\omega)$	x	a	b	y	z	c
$\tau_{obs}(\omega')$	x	z	a	b	c	y

Using Proposition 2, it is easy to check that τ_{obs} is core compatible.[11] *More importantly, this example illustrates how core compatibility can impose an LCS restriction on i's revealed preference: agent i ranks x above either y or z. Indeed, suppose that i ranks y above x. Then it must be that the revealed trading cycle $\{b, y\}$ in ω exits before the revealed trading cycle $\{a, x\}$ associated to i. In that case, it must be that k ranks b above a. Hence the revealed trading cycle $\{x, b\}$ in ω' must dominate the revealed trading cycle $\{a, z\}$, in which case i prefers x over z.*

Examples 2 and 3 show how any given binary UCS (resp. LCS) restriction can be replicated as a revealed preference restriction on an individual i when a wisely-constructed auxiliary dataset is core compatible. Thus core compatibility of the auxiliary dataset implies acyclic satisfiability of the original problem (simply using the preference that arises for i from core compatibility). To establish that that testing core compatibility is NP-hard, it remains to show that acyclic satisfiability of the original problem implies core compatibility for the auxiliary dataset. The ordering satisfying the restrictions of the original problem is used as a preference for individual i, and one can construct preference orderings for the other individuals to check that the auxiliary dataset is core compatible. In particular, the constructed auxiliary dataset has the feature that core compatibility can be captured entirely as revealed preference restrictions for i. Details are available in the Appendix.

Proposition 3. *Testing core compatibility is NP-hard.*

[11]E.g., the six individuals' revealed preferences are acyclic when ranking the revealed trading cycles as follows: $\{z\}$ above $\{a, x\}$ above $\{b, y\}$ above $\{c\}$ in ω, and $\{y\}$ above $\{a, z\}$ above $\{b, x\}$ above $\{c\}$ in ω'.

References

Brown, D. and R. Matzkin (1996), Testable Restrictions on the Equilibrium Manifold, *Econometrica*, 64, 1249–1262.

Carvajal, A., R. Deb, J. Fenske, and J. Quah (2013), Revealed Preference Tests of the Cournot Model, *Econometrica*, 81(6), 2351–2379.

Chambers, C. and F. Echenique (2014), On the Consistency of Data with Bargaining Theories, *Theoretical Economics*, 9(1), 137–162.

de Clippel, G., and K. Rozen (2018), Bounded Rationality and Limited Datasets, *mimeo*.

Echenique, F. (2008), What Matchings Can Be Stable? The Testable Implications of Matching Theory, *Mathematics of Operations Research*, 33(3), 757–768.

Forges, F. and E. Minelli (2009), Afriat's Theorem for General Budget Sets, *Journal of Economic Theory*, 144, 135–145.

Kalyanaraman, S. and C. Umans (2008), The Complexity of Rationalizing Matchings, *ISAAC '08 Proceedings of the 19th International Symposium on Algorithms and Computation*, 171–182.

Piccione, M. and A. Rubinstein (2007), Equilibrium in the Jungle, *Economic Journal*, 117(522), 883–896.

Ray, I. and L. Zhou (2001), Game Theory via Revealed Preferences, *Games and Economic Behavior*, 37(2), 415–424.

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

Sprumont, Y. (2000), On the Testable Implications of Collective Choice Theories, *Journal of Economic Theory*, 93, 205–232.

Appendix: Proof of Proposition 3

Fix a mixed set $\mathcal{R}$ of binary restrictions defined on a set X, as in de Clippel and Rozen (2018, Proposition 3). For each restriction r, let x_r be the option whose contour set is being restricted, and let y_r, z_r be the two options potentially included in the upper (or lower) contour set of x_r if r is an UCS (or LCS) restriction. We assume wlog that $y_r \neq z_r$. Consider the following sets of agents and houses

$$A = X \cup \{i\} \cup \{j_r, k_r, \ell_r | r \in \mathcal{R}\}$$

$$H = X \cup \{d\} \cup \{a_r, b_r, c_r | r \in \mathcal{R}\}.$$

The dataset consists of two endowments – ω_r and ω'_r – for each restriction r. In all these endowments, each individual $x \in X \setminus \{x_r, y_r, z_r\}$ owns x, and for each $s \neq r$, j_s owns a_s, k_s owns b_s, and ℓ_s owns c_s. The remaining components of these endowments are defined in the two tables below:

	i	j_r	k_r	ℓ_r	x_r	y_r	z_r
ω_r	y_r	a_r	x_r	d	b_r	c_r	z_r
ω'_r	a_r	y_r	x_r	z_r	b_r	c_r	d

for each UCS restriction r, and

	i	j_r	k_r	ℓ_r	x_r	y_r	z_r
ω_r	a_r	x_r	y_r	b_r	d	c_r	z_r
ω'_r	b_r	a_r	z_r	x_r	c_r	y_r	d

for each LCS restriction r.

The observed trading function τ_{obs} is defined in the following two tables, where NT stands for 'no trade' (that is, consume one's own endowment)

	i	j_r	k_r	ℓ_r	other agents
$\tau_{obs}(\omega_r)$	NT	x_r	a_r	NT	NT
$\tau_{obs}(\omega_r')$	z_r	x_r	y_r	a_r	NT

for each UCS restriction r, and

	i	j_r	k_r	ℓ_r	other agents
$\tau_{obs}(\omega_r)$	x_r	a_r	b_r	y_r	NT
$\tau_{obs}(\omega_r')$	x_r	z_r	a_r	b_r	NT

for each LCS restriction r.

We conclude the proof, by showing that there exists an acyclic relation satisfying the restrictions listed in $\mathcal{R}$ if and only if τ_{obs} is core-compatible. If $\mathcal{R}$ is acyclically satisfiable, then let P be a strict acyclic relation on X satisfying the restrictions in $\mathcal{R}$. We can assume without loss of generality that P is complete, that is, an ordering. We now define rankings $(\succ_\omega)_{\omega\in\mathcal{D}}$ of the minimal trading cycles for each $\omega \in \mathcal{D}$. For each r, both $\succ_{\omega_r}$ and $\succ_{\omega_r'}$ places the singleton cycles $\{x\}$ on top (in any order), for each $x \in X \setminus \{x_r, y_r, z_r\}$. Right below these cycles, $\succ_{\omega_r}$ and $\succ_{\omega_r'}$ have:

(a) For an UCS restriction r with $y_r P z_r$: $\{y_r\} \succ_{\omega_r} \{a_r, x_r\} \succ_{\omega_r} \{z_r\}$ and $\{x_r, y_r\} \succ_{\omega_r'} \{a_r, z_r\}$.

(b) For an UCS restriction r with $z_r P y_r$: $\{z_r\} \succ_{\omega_r} \{a_r, x_r\} \succ_{\omega_r} \{y_r\}$ and $\{a_r, z_r\} \succ_{\omega_r'} \{x_r, y_r\}$.

(c) For an LCS restriction with $y_r P z_r$: $\{b_r, y_r\} \succ_{\omega_r} \{z_r\} \succ_{\omega_r} \{a_r, x_r\}$ and $\{y_r\} \succ_{\omega_r'} \{b_r, x_r\} \succ_{\omega_r'} \{a_r, z_r\}$.

(d) For an LCS restriction with $z_r P y_r$: $\{z_r\} \succ_{\omega_r} \{x_r, a_r\} \succ_{\omega_r} \{b_r, y_r\}$ and $\{y_r\} \succ_{\omega_r'} \{a_r, z_r\} \succ_{\omega_r'} \{b_r, x_r\}$.

Fix a ranking P' of H such that a's are above b's, b's are above c's, and c's are above d. The remaining cycles are singletons, and $\succ_{\omega_r}$ and $\succ_{\omega'_r}$ rank these at the bottom according to P'.

It remains to check that the individuals' revealed preferences induced by the observed trading function using these rankings of atoms are acyclic. We do this by providing such a preference ordering for each agent. Starting with agent i, let $\succ_i$ be any preference ordering that agrees with P on X and ranks all other houses below the elements of X.

- For an UCS restriction r such that $y_r P z_r$, the revealed preference is that y_r is superior to elements of $(H \setminus X) \cup \{x_r, z_r\}$ (from observed trade at ω_r), and that z_r is superior to elements of $H \setminus X$ (from observed trade at ω'_r). Given that $y_r P x_r$ or $z_r P x_r$ (since r is satisfied), it must be that $y_r P x_r$. Hence $\succ_i$ satisfies these revealed preferences.

- For an UCS restriction r such that $z_r P y_r$, the revealed preference is that y_r is superior to elements of $H \setminus (X \cup \{a_r\})$ (from observed trade at ω_r), and that z_r is superior to elements in $(H \setminus X) \cup \{x_r, y_r\}$ (from observed trade at ω'_r). Given that $y_r P x_r$ or $z_r P x_r$ (since r is satisfied), it must be that $z_r P x_r$. Hence $\succ_i$ satisfies these revealed preferences.

- For an LCS restriction r such that $y_r P z_r$, the revealed preference is that x_r is superior to elements of $(H \setminus X) \cup \{z_r\}$ (combining observed trades at ω_r and ω'_r). Given that $x_r P y_r$ or $x_r P z_r$ (since r is satisfied), it must be that $x_r P z_r$. Hence $\succ_i$ satisfies these revealed preferences.

- For an LCS restriction r such that $z_r P y_r$, the revealed preference is that x_r is superior to elements of $(H \setminus X) \cup \{y_r\}$ (from observed trade at both ω_r and ω'_r). Given that $x_r P y_r$ or $x_r P z_r$ (since r is satisfied), it must be that $x_r P y_r$. Hence $\succ_i$ satisfies these revealed preferences.

For j_r where r is an UCS restriction, any preference that ranks x_r at the top and a_r right below, satisfies the revealed preferences induced by observed trades. For j_r where r is an LCS restriction, any preference that ranks z_r at the top, following right after by a_r, satisfies the revealed preferences induced by observed trades.

For k_r where r is an UCS restriction with $y_r P z_r$, any preference that ranks y_r at the top, followed by a_r and then b_r, satisfies the revealed preferences induced by observed trades. For k_r where r is an UCS restriction with $z_r P y_r$, any preference that ranks a_r at the top, followed by y_r, and then b_r, satisfies the revealed preferences induced by observed trades. For k_r where r is an LCS restriction with $y_r P z_r$, any preference ranking b_r at the top, followed by a_r, satisfies the revealed preferences induced by observed trades. For k_r where r is an LCS restriction with $z_r P y_r$, any preference ranking a_r at the top, followed by b_r, satisfies the revealed preferences induced by observed trades.

For ℓ_r where r is an UCS restriction, any preference that ranks a_r at the top, and c_r right below it, satisfies the revealed preferences induced by observed trades. For ℓ_r where r is an LCS restriction, any preference that ranks y_r at the top, followed by b_r, and then c_r, satisfies the revealed preferences induced by observed trades.

Finally, for any agent w in X, any preference ordering with w at the top, following by the elements of H according to P', satisfies the revealed preferences induced by observed trades. This concludes the proof that τ_{obs} is core compatible if $\mathcal{R}$ is core compatible.

We now show the converse, namely that $\mathcal{R}$ is acyclically satisfiable if τ_{obs} is core-compatible. Thus there must exists a profile of preference orderings (one for each agent in A) such that the core at each endowment ω in $\mathcal{D}$ corresponds to $\tau_{obs}(\omega)$. Focus in particular on agent i's preference $\succ_i$. For any $x, y \in X$, say that xRy if i prefers x over y. Clearly R is acyclic, and it remains to prove that it satisfies the restrictions in $\mathcal{R}$. Consider first an UCS restriction r. Suppose that $x_r R y_r$, in which case i prefers x_r over y_r. Then, looking at trades for ω_r, the cycle $\{a_r, x_r\}$ must precede the cycle $\{y_r\}$, and hence k_r must prefer a_r over y_r. In that case, looking now at trades for ω'_r, it must be that the cycle $\{a_r, z_r\}$ precedes the cycle $\{x_r, y_r\}$, which reveals that i prefers z_r over x_r, or $z_r R x_r$. Thus $y_r R x_r$ or $z_r R x_r$, as desired. Consider next an LCS restriction r. Suppose, by way of contradiction, that both $y_r R x_r$ and $z_r R x_r$, in which case i prefers both y_r and z_r over x_r. Then, looking at trades for ω_r, the cycle $\{b_r, y_r\}$ must precede the cycle $\{a_r, x_r\}$, and hence k_r must prefer b_r over a_r. A similar reasoning applied to trades at ω'_r reveals that k_r prefers a_r over b_r. This contradiction confirms that the LCS restriction r holds. $\square$