The Inconsistency Rank*

Peter Caradonna†

October 15, 2024

Abstract

Cycles in revealed preference data are often regarded as fundamental units of choice-theoretic inconsistency. Contrary to this, we show that in nearly any environment, cyclic choices over some menus necessarily force further cyclic choices elsewhere. In many cases, the entirety of a subject's inconsistency can be explained by only a handful of cycles. We characterize such dependencies, and show that every set of 'independent' cycles capable of explaining all others is necessarily of the same size. This quantity provides a simple, transparent measure of irrationality that accounts for the dependencies introduced by the structure of the choice environment or experiment.

*I would like to thank Axel Anderson, Chris Chambers, John Quah, Christopher Turansick, and Andrea Wilson for various helpful conversations over the course of this project, as well as seminar audiences at SAET 2019. This paper partially subsumes the earlier working paper 'How Strong is the Weak Axiom.'

†Division of Humanities and Social Sciences, Caltech. Email: ppc@caltech.edu.

1 Introduction

The hypothesis that agents are rational is perhaps the most ubiquitous and widely adopted assumption in all of economics. The testable implications of rationality have long since been characterized by the revealed preference literature (Samuelson 1938; Houthakker 1950; Richter 1966; Afriat 1967): a subject's choices or decisions are consistent with the maximization of a preference relation if, and only if, no *choice cycles* are observed in their behavior.

Despite the clarity and elegance of these results, they are binary in nature: behavior is either precisely consistent with the rational paradigm, or it is not. In practice, however, it is highly unlikely that any sufficiently rich data set would pass such an exact test (for example, due to measurement error or model misspecification). Instead, what is needed are means to quantify the severity, or magnitude, of any observed deviations from rationality.

This paper provides a principled, transparent method of quantifying the degree of irrationality observed in any choice data set. Our basic observation is that choice cycles rarely occur in isolation. Generally, once a subject has chosen cyclically from some collection of menus, there will be other menus on which *every* possible choice necessarily generates further cycles. We take the position that such 'forced' or 'knock-on' cycles should not be viewed as indicating any deeper degree of irrationality than what would be implied by observing only the initial, 'forcing' cycles alone.

The following example illustrates how the structure of a choice environment, or experiment, can lead cyclic choices over certain sub-collections to force subsequent choices to create additional cycles.

Example 1. Consider four alternatives $\{a, b, c, d\}$, and suppose an individual is presented with choices between $\{a, b\}$, $\{b, c\}$, $\{c, d\}$, and $\{d, a\}$. If this individual were to choose a in the presence of b, b in the presence of c and so forth cyclically, their choice behavior would be inconsistent with preference maximization, as it contains a revealed preference cycle.

Suppose now the agent is additionally presented with the opportunity to choose from the menu $\{a, b, c\}$. Their cyclic choices from the initial, binary menus forces them to necessarily make another choice cycle: if a is not chosen exclusively as the most-preferred alternative from this set, a pairwise reversal obtains, relative to their earlier choices.[1] On the other, if a is indeed chosen from this new menu, then a is revealed to be preferable to c, creating a different, new cycle, this time between a, c, and d. ■

In this example, the structure of the set of menus ensured that any cycle of choices over all four alternatives can never occur in isolation: any such cycle necessarily forces at least one other, elsewhere. In such cases, the forcing cycle may be used to justify, or *explain*, the existence of the forced cycle. We define the inconsistency rank of the data to be the size of any collection of mutually-independent cycles (i.e. which do not explain each other) but which nonetheless explain, in this manner, all others. We term such sets of cycles 'irrationality kernels' for the data. In general, many kernels may exist for a given data set; however, we show that any two necessarily contain the same number of elements.

The inconsistency rank may be interpreted as counting the observed number of distinct pieces of evidence of irrationality needed to justify the entirety of a subject's inconsistency. Since no two elements of any kernel can be used to explain each other, our index does not 'double count' cycles, as the majority of existing indices do.[2] Conversely, because every observed cycle can be explained by (at least) one cycle in any kernel, we ensure that our index reflects the entirety of a subject's inconsistency.

As a consequence, the inconsistency rank is a cardinal index: it is meaningful to say, for example, that the inconsistency in one choice data set requires twice as many pieces of evidence to justify as another. Moreover, it is similarly

[1]For example, if b belongs to the subject's choice set from this menu, then b is revealed weakly preferred to a, while a was earlier revealed to be strictly preferred to b, creating a cycle of length two.

[2]See Section 5.2 for an in-depth discussion of how the inconsistency rank relates to various other well-known measures of irrationality.

valid to compare the inconsistency rank of two choice correspondences defined over *distinct* collections of menus. Normally, when making cross-domain comparisons, there is the possibility that one domain may present a more exacting test. Frequently, this is because of the potential for more cycles to emerge from fewer inconsistent choices on one domain than another. The advantage of the inconsistency rank is that it normalizes precisely for this influence of the domain's structure on the set of observed cycles, allowing us to compare the volume of observed evidence of irrationality in absolute terms.

In Section 3 we define our model; throughout, we consider the abstract choice framework of Richter (1966). In Section 4 we characterize how the structure of the choice domain leads cyclic choices over some alternatives to sometimes force further cycles to emerge. We use this characterization in Section 5 to define the inconsistency rank. Finally, in Section 6 we examine extremal domains, on which either every (or no) choice cycle forces others. We show the former class characterizes those domains on which the 'fundamental theorem of revealed preference' (e.g. Ok et al. 2015) remains valid, while the latter class is suitably degenerate. We interpret this as providing evidence that in most practical experiments, forced cycles are likely to emerge.

2 Related Literature

There is an extensive literature on inconsistency measurement for revealed preference data; Dziewulski et al. (2024) is an excellent recent survey. Lanier and Quah (2024) study the incompatibility of several natural axioms such an index might obey. Mononen (2020) axiomatizes several classical measures for price-consumption data.

For general choice data sets, Houtman and Maks (1985) propose using as an index the size of any minimal set of observations which, when dropped, render the remaining observations consistent. Apesteguia and Ballester (2015) propose a measure related to the minimal number of binary swaps needed to transform the revealed preference relation into a preference. Ribeiro (2020) proposes a partial ordering, where one data set is more rational than another if it is consistent on any sub-collection of menus on which the other is. In practice, the number of choice cycles, or number of observations belonging to some choice cycle, are also commonly measures (e.g. Famulari 1995; Harbaugh et al. 2001).

For classical price-consumption data, Afriat (1973) proposes the so-called 'critical cost efficiency' index (cf. Varian 1990).[3] More recent contributions include Echenique et al. (2011) (see also Lanier et al. 2024) and Dean and Martin (2016).[4]

3 Preliminaries

Let X be an arbitrary set of **alternatives** from which an agent chooses. Let $\Sigma \subseteq 2^X \setminus \{\varnothing\}$ be a collection of **budgets** encoding the specifics of the collection of constraints under which choice occurs. When Σ contains all budgets of cardinality greater than one, we say that it is **complete**. We refer to the tuple (X, Σ) as a **choice environment**, and interpret any such environment as abstractly defining an experiment: it is simply the collection of problems over which an agent's choices are observed.

For any subset $A \subseteq X$, we define the restriction of Σ to A as those elements of Σ wholly contained in A:

$$\Sigma|_A = \big\{B \in \Sigma : B \subseteq A\big\},$$

and for a collection of subsets $\mathcal{A} \subseteq 2^X$, it will be convenient to define the shorthand:

$$\Sigma|_{\mathcal{A}} = \left\{B \in \Sigma : B \subseteq \bigcup_{A \in \mathcal{A}} A\right\}.$$

A mapping $c : \Sigma \to 2^X \setminus \{\varnothing\}$ is a **choice correspondence** if, for all $B \in \Sigma$, it satisfies $c(B) \subseteq B$. Let $\mathcal{C}(X, \Sigma)$ denote the collection of all choice

[3] See Echenique (2021) for a discussion of the interpretation of this index.

[4] See also Mononen (2023).

correspondences for the environment (X, Σ). Given a choice correspondence $c \in \mathcal{C}(X, \Sigma)$, a weak order $\succeq$ on X **strongly rationalizes** c if:

$$(\forall B \in \Sigma) \;\; c(B) = \big\{x \in B : \forall y \in B, x \succeq y\big\}.$$

Given a $c \in \mathcal{C}(X, \Sigma)$, its revealed preference pair $(\succsim_c, \succ_c)$ is defined via: $x \succsim_c y$ if there exists some $B \in \Sigma$ such that $x, y \in B$ and $x \in c(B)$, and $x \succ_c y$ if there exists some $B \in \Sigma$ such that $x, y \in B$, $x \in c(B)$ and $y \notin c(B)$.

A choice correspondence $c \in \mathcal{C}(X, \Sigma)$ satisfies the **weak axiom** of revealed preference if it contains no pairwise reversals: $x \succsim_c y$ implies $y \not\succ_c x$. We say c obeys the **generalized axiom** of revealed preference if $(\succsim_c, \succ_c)$ contains no cycles of the form:

$$x_0 \succsim_c x_1 \succsim_c \cdots \succsim_c x_N \succ_c x_0.$$

It is without loss to suppose that these alternatives are all distinct, as any cycle containing multiple appearances of the same alternative necessarily also contains a sub-cycle consisting only of distinct alternatives. In particular, it was shown by Richter (1966), making use of an extension theorem due to Szpilrajn (1930), that a choice correspondence is strongly rationalizable by a weak order if and only if it obeys the generalized axiom.[5] In light of this, we will interchangeably refer to the satisfaction of the generalized axiom as strong rationalizability.

3.1 The Budget Graph

It will be helpful to define an auxiliary structure that encodes, for a given choice environment, which pairs of alternatives it is possible for a preference to be revealed between. For any choice environment (X, Σ), let $\Gamma(X, \Sigma)$ denote the undirected graph whose vertex set $V_\Gamma = X$, and whose edge set E_Γ is given by the relation of two vertices belonging to some common budget:

$$\{x, y\} = e_{xy} \in E_\Gamma \iff \exists B \in \Sigma \text{ s.t. } \{x, y\} \subseteq B.$$

[5] We note, however, that Szpilrajn (1930) acknowledges the priority of Banach, Kuratowski, and Tarski in discovering, though not publishing, the result.

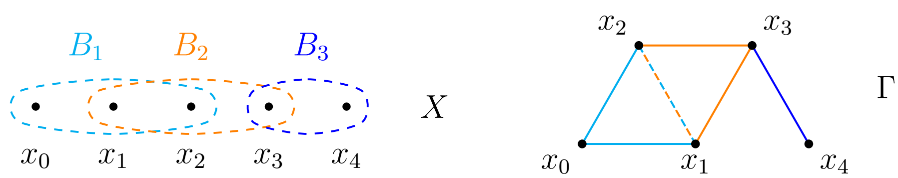

(a) A choice environment with five alternatives and three budget sets.

(b) The budget graph associated with this environment.

Figure 1: A simple choice environment and its corresponding budget graph. The coloring of the edges in the budget graph indicates which budgets are responsible for the edge's inclusion in the graph.

We term $\Gamma(X, \Sigma) = (V_\Gamma, E_\Gamma)$ the **budget graph**.

For a given $c \in \mathcal{C}(X, \Sigma)$ and any $e \in E_\Gamma$ there is a well-defined (possibly empty) restriction of the revealed preference pair $(\succsim_c, \succ_c)$ to the edge e, which we denote by $(\succsim_c, \succ_c)\big|_e = \left(\succsim_c\big|_e, \succ_c\big|_e \right)$, defined by:

$$\succsim_c\big|_e = \succsim_c \cap \{x, y\} \times \{x, y\},$$

(and respectively $\succ_c\big|_e$). Similarly, given a collection of edges $E' \subseteq E_\Gamma$, we define:

$$\succsim_c\big|_{E'} = \bigcup_{e \in E'} \succsim_c\big|_e.$$

A **loop** in Γ is a connected, finite subgraph $\gamma = (V_\gamma, E_\gamma)$ such that every vertex in V_γ belongs to precisely two edges in E_γ. Suppose that:

$$x_0 \succsim_c x_1 \succsim_c \cdots \succsim_c x_N \succ_c x_0 \tag{1}$$

is a cycle in $(\succsim_c, \succ_c)$. We refer to the subgraph with vertices $\{x_0, \ldots, x_N\}$ and edges $\{x_0, x_1\}, \ldots, \{x_N, x_0\}$ as the **support** of the cycle (1). The support of a cycle is a loop if and only if it does not correspond to a WARP violation.[6] Let:

$$\mathcal{Z} = \left\{ (\tilde{V}, \tilde{E}) \subseteq \Gamma(X, \Sigma) : (\tilde{V}, \tilde{E}) \text{ is the support of a cycle} \right\}.$$

[6] If we have a WARP violation $x_0 \succsim_c x_1 \succ_c x_0$, then this support is simply the subgraph $(\tilde{V}, \tilde{E})$ with $\tilde{V} = \{x_0, x_1\}$ and $\tilde{E} = \left\{ \{x_0, x_1\} \right\}$. This is not, however, a loop.

We refer to $\mathcal{Z}$ as the **cycle set** of c; by doing so, we are implicitly identifying cycles which have identical support.[7] Let:

$$\mathcal{Z}_W = \left\{ (\tilde{V}, \tilde{E}) \subseteq \Gamma(X, \Sigma) : (\tilde{V}, \tilde{E}) \text{ is the support of a WARP violation} \right\}$$

denote the set of (supports of) WARP violations. Note that by definition, we have $\mathcal{Z}_W \subseteq \mathcal{Z}$.

4 Propagation of Cycles

Informally, we say that a cycle propagates when, given some choices defining it, *every* possible combination of other choices from the remaining budgets in Σ necessarily lead to the formation of other choice cycles.

This was precisely what occurred in Example 1: once the cycle $x_0 \succ_c x_1 \succ_c x_2 \succ_c x_3 \succ_c x_4$ has been chosen, every choice the subject makes from the remaining budget creates (at least) one other cycle. We do not know which particular cycle will be realized until we know the subject's choice from the remaining budget, but even before this choice is made, we know that it will necessarily form a new choice cycle. In such cases, we say the cycle has the **propagation property**.

4.1 Cyclic Collections and Covers

Let $z \in \mathcal{Z} \setminus \mathcal{Z}_W$. A collection of budgets $\mathcal{B}_z \subseteq \Sigma$ is a **cyclic collection** for z if, for every $e \in E_z$, there exists a $B \in \mathcal{B}_z$ with $e \subseteq B$. Similarly, if $z \in \mathcal{Z}_W$, we say $\mathcal{B}_z \subseteq \Sigma$ is a cyclic collection for if it contains two distinct budgets B, B' that both contain the unique edge in E_z.

Given a cycle $z \in \mathcal{Z} \setminus \mathcal{Z}_W$ and cyclic collection $\mathcal{B}_z$, we say that $\mathcal{B}_z$ is **covered** if either:

(i) There exists a $\bar{B} \in \Sigma|_{\mathcal{B}_z}$ such that $V_\gamma \subseteq \bar{B}$; or

[7] For example, by doing so we are regarding the cycles $x_0 \succsim_c x_1 \succsim_c x_2 \succ_c x_0$ and $x_0 \succ_c x_1 \succ_c x_2 \succ_c x_0$ as equivalent.

(ii) There exists a $\bar{B} \in \Sigma|_{\mathcal{B}_z}$ such that $\bar{B}$ contains a pair of elements of V_γ that are not connected by any edge in E_z.

We refer to such a $\bar{B} \in \Sigma$ as a cover.[8] For a cycle $z \in \mathcal{Z}_W$, we define every cyclic collection to be uncovered, and note that condition (i) implies (ii) if and only if $|V_\gamma| > 3$. Conversely, if $z \in \mathcal{Z}_W$, we define any cyclic collection to be uncovered.

Finally, we say a cyclic collection $\mathcal{G}_z$ is a **generator** for the cycle $z \in \mathcal{Z}$ if z is also a cycle of the choice correspondence restricted to $\mathcal{G}_z$, and $\mathcal{G}_z$ is $\subseteq$-minimal with respect to this property.

4.2 Ex-Ante vs. Ex-Post Propagation

When considering the propagation of choice cycles, there are two natural questions we can ask:

- **Ex-Ante Propagation**: Given some loop γ in the budget graph, when does *every* choice correspondence which contains a cycle supported on γ necessarily also possess some other cycle?
- **Ex-Post Propagation**: When does every choice correspondence with cycle z and generator $\mathcal{G}_z$ also possess some other cycle?

A loop γ has the ex-ante propagation property if and only if there is no way of choosing cyclically around γ without creating some cycle elsewhere in the data. In particular, this notion imposes no restrictions on *which* budgets the choices generating the cycle on γ are made. Conversely, ex-post propagation requires that the choices generating the cycle supported on γ be made on a *specific* collection of budgets.

[8]Recall the restricted collection $\Sigma|_{\mathcal{B}_z}$ is defined as:

$$\Sigma|_{\mathcal{B}_z} = \left\{ B \in \Sigma : B \subseteq \bigcup_{\hat{B} \in \mathcal{B}_z} \hat{B} \right\},$$

We distinguish between these cases by regarding ex-ante propagation as a property of a loop on the budget graph alone, versus ex-post propagation as a property of the pair consisting of a loop (the support of the cycle) and cyclic collection for it (the generator for the cycle).

Our first results provide a characterization of both ex-ante and ex-post propagation in terms of the structure of the choice environment.

Proposition 1 (Characterization of Ex-Post Propagation)**.** *Let $c \in \mathcal{C}(X, \Sigma)$ be a choice correspondence with cycle z and associated generator $\mathcal{G}_z$. Then the following are equivalent:*

(i) $\mathcal{G}_z$ is covered.

(ii) Every choice correspondence which contains z as a cycle, and $\mathcal{G}_z$ as a generator for z, also contains at least one other cycle.

Proposition 1 shows that when a cycle is generate by choices on some collection of budgets, it forces others *if and only if* the choice environment contains some budget covering the collection. As a corollary of this, we obtain that a loop has the property that cycles supported on it always propagate, no matter how the choices generating this cycle were made, if and only if every cyclic collection for the loop is covered.

Corollary 1 (Characterization of Ex-Ante Propagation)**.** *Let γ be a loop in $\Gamma(X, \Sigma)$. Then every $c \in \mathcal{C}(X, \Sigma)$ which chooses cyclically around γ contains at least one other cycle if and only if every cyclic collection for γ is covered.*

5 Measurement of Inconsistency

Classical revealed preferences conditions such as GARP provide a black-and-white test of consistency: a subject satisfies the generalized axiom of revealed preference if and only if their choice behavior is consistent with the hypothesis of rational choice. However, such conditions are silent on the *severity* of any observed deviations from rational behavior.

In this section, we consider the problem of quantifying the magnitude of any observed inconsistency in an arbitrary choice data set. The premise of our approach is that, when choice cycles have forced others, any induced cycles should not be treated as evidencing a deeper degree of irrationality than what would have been surmised by observing the forcing cycles alone. Rather, such forced cycles are purely artifacts of the structure of the experimental structure itself. To the best of our knowledge, we are the first paper to consider the structure of the choice environment in this manner, and no other commonly used inconsistency index adequately differentiates between forcing and forced cycles in this manner.[9]

5.1 The Inconsistency Rank

Consider a choice correspondence $c \in \mathcal{C}(X, \Sigma)$, with cycle set $\mathcal{Z}$. Given $z, z' \in \mathcal{Z}$, we say that z **explains** z' (denoted $z \Longrightarrow z'$) if, for every generator $\mathcal{G}_{z'}$ for z', there exists a generator $\mathcal{G}_z$ for z such that, for each $B \in \mathcal{G}_{z'}$, either:

(i) B also belongs to $\mathcal{G}_z$; or

(ii) B covers $\mathcal{G}_z$.

A cycle z explains a cycle z' if, given some set of choices generating z, z' is generated by choices that either (i) were already involved in generating z, or (ii) were made on budgets on which *any* choice would have created new cycles, given those generating z.

To motivate this choice of terminology, suppose that $z \Longrightarrow z'$. In light of Proposition 1, had we observed only those choices directly involved in generating z, there are two possibilities. The first is that we also directly observe z', i.e. that some generator for z contains a generator for z'.[10] In this case, it

[9] See Section 5.2 for a discussion of this point.

[10] For example, if $X = \{x_0, x_1, x_2\}$ and $\Sigma = \big\{\{x_0, x_1\}, X\big\}$, define $c\big(\{x_0, x_1\}\big) = \{x_1\}$ and $c(X) = X$. Let z denote the cycle $x_1 \succ_c x_0 \precsim_c x_1$, and z' the cycle $x_1 \succ_c x_0 \precsim_c x_2 \precsim_c x_1$. If one observes the choices generating z, here z' is also directly observed, as Σ itself is a generator for each cycle.

is natural to conclude that z' is a direct consequence of the manner in which the choices making up z were made.

The second possibility is that z' is not direcly generated by the choices making up z alone but, by Proposition 1, we are able to nonetheless conclude that *every* possible combination of choices from the remaining budgets in Σ yield some other cycle(s). In this case, while the specific cycle z' need not have been directly caused by the choices making up z, ex-post we may still rationalize it as the particular realization of the additional inconsistency guaranteed by the choices yielding z and the structure of the choice environment.

Denote the transitive closure of $\Longrightarrow$ by $\Longrightarrow^*$. If $z \Longrightarrow^* z'$ then we say z **indirectly explains** z'; if two cycles are $\Longrightarrow^*$-unrelated, we call them **independent**. Define a subset of cycles $\mathcal{I} \subseteq \mathcal{Z}$ to be an **irrationality kernel** for the choice correspondence c if it satisfies:

(IK.1) **Independence**: For every pair of distinct cycles $z, z' \in \mathcal{I}$, z and z' are $\Longrightarrow^*$-unrelated.[11]

(IK.2) **Explanatory Power**: For every $z' \in \mathcal{Z}$, there exists some $z \in \mathcal{I}$ such that $z \Longrightarrow^* z'$.

An irrationality kernel is simply a subset of cycles with the property that no two cycles in it (even indirectly) explain each other, but nonetheless which together explain the entirety of the observed inconsistency, $\mathcal{Z}$.[12]

When $|\Sigma| < \infty$, irrationality kernels exist and are finite, for all choice correspondence in $\mathcal{C}(X, \Sigma)$. However, in general they will not be unique.[13] Nevertheless, as our next result shows, for any choice correspondence, *every* irrationality kernel has the same, finite, cardinality.

[11] Note that every $z \in \mathcal{Z}$ explains itself, i.e. the relation $\Longrightarrow$ is reflexive.

[12] Our definition of an irrationality kernel is also reminiscent of the notion of a 'stable set' in cooperative game theory; see Morgenstern and Von Neumann (1944).

[13] Note this does not require that X itself be finite.

Theorem 1. *Let* (X, Σ) *be a choice environment, with* $|\Sigma| < \infty$*. Then for every* $c \in \mathcal{C}(X, \Sigma)$ *there exists at least one irrationality kernel. Moreover, for any two kernels* $\mathcal{I}, \mathcal{I}' \subseteq \mathcal{Z}$*:*

$$|\mathcal{I}| = |\mathcal{I}'| < \infty.$$

In light of Theorem 1, we may associate to any such choice correspondence a well-defined number: the size of any its irrationality kernel(s); we term this quantity the **inconsistency rank** of the correspondence. It reflects the magnitude of the observed deviations from rationality, normalized for the dependencies between elements of $\mathcal{Z}$ introduced by the structure of Σ.

Irrationality kernels may equivalently be interpreted as a form of generalized 'basis' for the set of cycles, $\mathcal{Z}$. The requirement (IK.1) that all cycles in $\mathcal{I}$ be suitably unrelated is akin to requiring a set of vectors be linearly independent, while (IK.2) requires that, in this abstract sense, the cycles in $\mathcal{I}$ span all of $\mathcal{Z}$. From this perspective, Theorem 1 then establishes that every such 'basis' for $\mathcal{Z}$ is of the same size.[14]

The inconsistency rank admits a straightforward interpretation: it is the minimum number of choice cycles needed to fully justify, or explain, the entirety of a subject's inconsistency.[15] If each choice cycle is regarded as equally indicative of irrationality, the inconsistency rank simply tallies the minimum number of strikes against the hypothesis of rationality needed to justify their observed deviations.[16]

[14]Formally, the proof of Theorem 1 amounts to showing that set of cycles $\mathcal{Z}$ can be endowed with a *matroid* structure (see, e.g., Oxley 2006) in which the irrationality kernels $\mathcal{I}$ are precisely the bases.

[15]Kalai et al. (2002) consider a spiritually similar measure of inconsistency, the number of *preferences* needed to rationalize the data by 'multiple rationales.'

[16]Crucially, while every irrationality kernel is maximally independent, not every maximal independent set is an irrationality kernel. In general, such sets may be much larger than the inconsistency rank. For example, suppose a single cycle z directly explains cycles $z_1, \ldots, z_K$, but where where $z_1, \ldots, z_K$ are pairwise independent. Then the set $\{z_1, \ldots, z_K\}$ is (maximally) independent, but the unique irrationality kernel is simply $\{z\}$.

In light of this, the inconsistency rank may be regarded as a *cardinal* measure of inconsistency. It is meaningful to say, e.g., that given two choice correspondences, one requires twice as many cycles as another to explain away all the observed inconsistency. Indeed, such comparisons remain valid even across *differing* domains. Generally, different choice domains will affect the patterns of dependencies, and hence explanatory relations, between cycles differentially. However, the advantage of the inconsistency rank is that it normalizes for precisely these differences in explanatory relations, yielding a measurement in absolute terms.

5.2 Relation to Other Measures

The inconsistency rank is not a monotone transformation of any existing indices. In this section, we consider a number of well-known measures of inconsistency and show, by means of example, that each may yield the opposite ranking of the relative consistency of two choice correspondences when compared to the inconsistency rank. We interpret this as evidence that existing measures do not, in general, account for dependencies between cycles, and hence are prone to double counting.

5.2.1 Counting Cycles

The cycle count is of a choice correspondence c is simply $|\mathcal{Z}_c|$, i.e. the size of its set of revealed preference cycles (e.g. Harbaugh et al. 2001).

Example 2. Suppose $X = \{x_0, \ldots, x_4\}$ and Σ consists of the budgets:

$$\{x_0, x_1\}, \ldots, \{x_4, x_0\}, \text{ and } \{x_0, x_2, x_3\}.$$

Consider the following choice correspondences:

$$c(B) = \begin{cases} \{x_i\} & \text{if } B = \{x_i, x_{i+1}\} \\ \{x_0, x_2, x_3\} & \text{if } B = \{x_0, x_2, x_3\}, \end{cases}$$

and

$$c'(B) = \begin{cases} \{x_i, x_{i+1}\} & \text{if } B = \{x_i, x_{i+1}\} \\ \{x_2, x_3\} & \text{if } B = \{x_0, x_2, x_3\}. \end{cases}$$

The first correspondence has $|\mathcal{Z}_c| = 7$ (one cycle on every loop of the budget graph, plus one WARP violation), and the second $|\mathcal{Z}_{c'}| = 4$. Thus, the cycle count index ranks c as exhibiting a greater degree of irrationality than c'. However, for c, the cycle $x_0 \succ_c x_1 \succ_c \cdots \succ_c x_0$ explains *every* other cycle, and hence the inconsistency rank of c is one, whereas all of the irrational kernels for c' are of cardinality two, yielding the opposite conclusion when the structure of the domain is accounted for. ■

5.2.2 Choices-In-Cycles

Another natural approach to quantifying inconsistency is to count the number of choices involved in patterns inconsistent with rationality (e.g. Famulari 1995; Swofford and Whitney 1986). In our setting, this amounts to counting the number, or proportion, of choices which are involved in some revealed preference cycle.

Example 3. Let $X = \{x_0, x_1, x_2, y_0, y_1, y_2, z_0, \ldots, z_K\}$, and Σ consist of the sets:

$$\{x_0, x_1\}, \ldots, \{x_2, x_0\}, \{y_0, y_1\}, \ldots, \{y_2, y_0\}, \text{ and } \{z_0, z_1\}, \ldots, \{z_K, z_0\}.$$

Define:

$$c(B) = \begin{cases} \{x_i\} & \text{if } B = \{x_i, x_{i+1}\} \\ \{y_i\} & \text{if } B = \{y_i, y_{i+1}\} \\ \{z_i, z_{i+1}\} & \text{if } B = \{z_i, z_{i+1}\}. \end{cases}$$

Then c has 6 choices appearing in some choice cycle, and its two cycles are independent hence it has an inconsistency rank of 2. Conversely, define:

$$c'(B) = \begin{cases} \{x_i, x_{i+1}\} & \text{if } B = \{x_i, x_{i+1}\} \\ \{y_i, y_{i+1}\} & \text{if } B = \{y_i, y_{i+1}\} \\ \{z_i\} & \text{if } B = \{z_i, z_{i+1}\}. \end{cases}$$

Here, c' has K choices appearing in its one cycle, and has an inconsistency rank of 1. Thus here the inconsistency rank and choices-in-cycles (and the fraction of choices-in-cycles) index yield opposite comparisons across agents. ■

5.2.3 Houtman-Maks

Houtman and Maks (1985) propose using as an index the minimal number of choices which, when removed from the data, make the remaining observations rationalizable.

Example 4. Consider again the environment and choice correspondences from Example 2. The Houtman-Maks index for the first correspondence c is equal to two: one must, e.g., remove the budget $\{x_0, x_2, x_3\}$ and one binary budget to break every cycle. In contrast, for c' it suffices to remove $\{x_0, x_2, x_3\}$ alone, and hence its Houtman-Maks index is one. However, the inconsistency rank of c is one, versus two for c', again yielding a reversal. ■

5.2.4 Swaps Index

Apesteguia and Ballester (2015) propose measuring the inconsistency of a data set by counting the minimal number of binary 'swaps' needed to take the best-fitting linear order and render it consistent with the observed choices.[17]

Example 5. Let $X = \{x_0, \dots, x_5\}$ and Σ consist of the budgets:

$$\{x_0, x_1\}, \dots, \{x_4, x_5\}, \{x_5, x_0\}, \{x_0, x_1, x_4\}, \text{ and } \{x_1, x_3, x_4\}.$$

Define the choice functions:

$$c(B) = \begin{cases} \{x_{i+1}\} & \text{if } B = \{x_i, x_{i+1}\} \\ \{x_0\} & \text{if } B = \{x_0, x_1, x_4\} \\ \{x_1\} & \text{if } B = \{x_1, x_3, x_4\}, \end{cases}$$

[17] Formally, Apesteguia and Ballester (2015) consider rationalizing data by a linear order. This poses no particular difficulty to our framework.

and

$$c'(B) = \begin{cases} \{x_{i+1}\} & \text{if } B = \{x_i, x_{i+1}\}, i \in \{1,2\} \\ \{x_i\} & \text{if } B = \{x_i, x_{i+1}\}, i \in \{0,3,4\} \\ \{x_0\} & \text{if } B = \{x_0, x_1, x_4\} \\ \{x_1\} & \text{if } B = \{x_1, x_3, x_4\}. \end{cases}$$

Consider first the choice function c. There are five cycles in $\mathcal{Z}_c$, one each of length six, five, four, three, and two. However, the cycle of length six explains every other cycle, hence there is a unique irrationality kernel, consisting of just this cycle. In contrast, there are only three cycles in $\mathcal{Z}_{c'}$, one of length four and two of length three. Here, the cycles $x_0 \succ_{c'} x_1 \succ_{c'} x_4 \succ_{c'} x_5 \succ_{c'} x_0$ and $x_3 \succ_{c'} x_2 \succ_{c'} x_1 \succ_{c'} x_3$ together form an irrational kernel, and we conclude that c' exhibits a lesser volume of irrational behavior than does c, according to the inconsistency rank.

The opposite ordering, however, is obtained when considering the swaps index values for these choice functions. When computing the swaps index for c, the the linear order:

$$x_1 \rhd x_0 \rhd x_5 \rhd x_4 \rhd x_3 \rhd x_2$$

is most consistent with the data. This yields a swaps index of three (there are two inconsistencies from choice on $\{x_0, x_1, x_4\}$ and one on $\{x_1, x_2\}$ relative to this order).

On the other hand, for c', the linear order:

$$x_5 \rhd x_2 \rhd x_0 \rhd x_1 \rhd x_3 \rhd x_4$$

minimizes the swaps index, but requires only two swaps (one each on the budgets $\{x_2, x_3\}$ and $\{x_4, x_5\}$). This shows that the inconsistency rank is not a monotone transformation of the swaps index. ■

5.2.5 Rationality Ordering

Ribeiro (2020) proposes an ordinal ranking of choice data, called the 'rationality ordering.' A choice correspondence c dominates a correspondence c' in this ordering if, for every sub-collection $\Sigma' \subseteq \Sigma$ on which $c|_{\Sigma'}$ is not rationalizable, $c'|_{\Sigma'}$ is also not rationalizable.

Example 6. Suppose $X = \{x_0, x_1, x_2\}$ and Σ is the complete domain, consisting of all non-empty, non-singleton subsets of X. Define:

$$c(B) = \begin{cases} \{x_1\} & \text{if } B = \{x_0, x_1\} \\ B & \text{if } B = X \\ \{x_2\} & \text{else,} \end{cases}$$

and

$$c'(B) = \begin{cases} \{x_{i+1}\} & \text{if } B = \{x_i, x_{i+1}\} \\ B \text{ else.} \end{cases}$$

The correspondence c has four choice cycles, three of length two on $\{x_i, x_{i+1}\}$, label these z_i, $i = 0, 1, 2$, and one of length three over all of X, labelled z. Note that each two-cycle z_i has a unique generator, $\{x_i, x_{i+1}\}$ and X, and the three-cycle z has two possible choices of generator, $\{x_1, x_2\}$ and X, or $\{x_2, x_0\}$ and X. The set $\{z_0, z\}$ defines an irrationality kernel, as for each z_i, $i = 1, 2$, $z \implies z_i$, hence we obtain that the inconsistency rank of c equals 2.[18]

For correspondence c', there are again four cycles, with the same supports. However, note that now, the three-cycle has $\{x_0, x_1\}, \ldots, \{x_2, x_0\}$ as a generator, and hence it directly explains all the z_i, not just z_1 and z_2 as was the case for c. Thus the singleton set $\{z\}$ defines an irrationality kernel, and hence the inconsistency rank of c' is 1. However, by inspection, c is rationalizable on a strictly larger collection of subdomains of Σ than c. Thus in the rationality ordering, c is deemed strictly more rational than c', and hence we observe a reversal relative to the inconsistency index.[19] ■

[18] There are other irrationality kernels. For example, $\{z_0, z_1\}$ also defines a kernel because z_1 directly explains z, and z directly z_2, hence z_1 indirectly explains both z and z_2. By an analogous argument, $\{z_0, z_2\}$ defines yet another kernel.

[19] For an example of the strict part of this claim, note c' is not rationalizable restricted to the doubleton budgets, but c is.

6 Extensions

6.1 Propagation-Free Design

A natural question is to what degree careful experimental design can ensure that no choice cycles have the propagation property. Clearly such environments exist, e.g. when Σ is singleton, but a priori it is unclear how broad this class is.

Define a choice environment (X, Σ) to be **propagation-free** if, for all $c \in \mathcal{C}(X, \Sigma)$ and any cycles $z, z' \in \mathcal{Z}$,

$$z \Longrightarrow^* z' \quad \text{if and only if} \quad z = z'.$$

In other words, an environment is propagation-free if every cycle is independent of every other, for *every* choice correspondence. When this is the case, $\mathcal{Z}$ is the unique irrationality kernel for any data set, and hence the inconsistency rank always equals the cycle count, $|\mathcal{Z}|$. Our next result shows that, while propagation-free environments do exist, they are 'degenerate,' in a sense we make precise.

Theorem 2. *Let (X, Σ) be a choice environment with finite budget graph $\Gamma(X, \Sigma)$. If Σ is propagation-free then, for any loop γ in the budget graph, every cyclic collection $\mathcal{B}_\gamma$ consists of either:*

(i) A single budget containing the entire vertex set of γ; or

(ii) Consists exclusively of two-element budgets.

Theorem 2 shows that for a given loop in $\Gamma(X, \Sigma)$, there are two possibilities: either it is incapable of supporting a choice cycle from any $c \in \mathcal{C}(X, \Sigma)$, or it has a unique cyclic collection, its edge set. In particular, no budget of cardinality ≥ 3 can be a part of any choice cycle.

Given the particularly simple structure of such environments, we interpret Theorem 2 as evidencing that most practical choice experiments will necessarily give rise to the possibility of choice cycles propagation. This suggests that the inconsistency rank is broadly applicable and will, in most environments of interest, yield different predictions than existing indices.

6.2 Well-Covered Environments

In this section, we consider the opposite extreme: choice environments where *every* loop has the (ex-ante) propagation property. Call a choice environment (X, Σ) **well-covered** if, for every loop γ in the budget graph $\Gamma(X, \Sigma)$, every cyclic collection $\mathcal{B}_\gamma$ for γ is covered. Well-coveredness is a completeness, or observational richness, condition on the environment. In particular, the 'complete' domain consisting of X and all subsets of X of cardinality ≥ 2 is well-covered.

A well-known property of complete environments is that the weak axiom of revealed preference is equivalent to the generalized, and hence characterizes rationalizability (e.g. Arrow 1959). This result is sometimes referred to as the 'fundamental theorem of revealed preference' (Ok et al. 2015).

Dating back at least to the characterization of rationalizability for general environments by Richter (1966), it has been an open question in the revealed preference literature to characterize those 'observationally rich' domains on which the weak axiom remains characteristic.[20] Our next result builds on our theory of propagation to provide a complete solution: well-coveredness is precisely the minimal observability requirement needed for a 'strong' weak axiom.[21]

Theorem 3. *Let (X, Σ) be an arbitrary choice environment. The weak axiom of revealed preference is necessary and sufficient for strong rationalizability if and only if (X, Σ) is well-covered.*

[20]For an analogue of this result in the context of price-consumption data, see Cherchye et al. (2018).

[21]The problem of characterizing which domains the weak axiom implies the generalized may be viewed as an ordinal analogue of the problem in mechanism design of characterizing those type spaces on which weak and cyclic monotonicty coincide. See, e.g. Saks and Yu (2005); Kushnir and Lokutsievskiy (2019).

7 Conclusions

We present a novel measure of the degree of observed irrationality for choice data, the inconsistency rank. Our index is based around the observation that the structure of a choice environment, or experiment, can lead to non-trivial dependencies between choice cycles. Our index admits a simple interpretation as the number of distinct instances of observed irrationality needed, in order to explain the totality of a subject's inconsistency.

More generally, we clarify an important dimension of experiment design: how the structure of the selected menus or choice sets affects (and constrains) manner in which inconsistent choices can be made. By normalizing for this structure, our index provides a natural, cardinal measure of inconsistency that remains valid both within and across domains.

A Proof Appendix

A.1 Proof of Proposition 1

Proof. $(ii) \implies (i)$: Suppose z is a cycle of c, with uncovered generator $\mathcal{G}_z$. By contraposition, it suffices to exhibit a choice correspondence $c' \in \mathcal{C}(X, \Sigma)$ which also contains z as a cycle, and $\mathcal{G}_z$ as a generator for z, but which contains no other cycles.

Denote the revealed preference cycle z via:

$$x_0 \succsim_c x_1 \succsim_c \cdots \succsim_c x_N \succ_c x_0,$$

where $N \geq 1$. We first define a choice correspondence on $\Sigma|_{\mathcal{G}_z}$. Let:

$$\vec{c}'(B) = \begin{cases} x_i & \text{if } \{x_i, x_{i+1}\} \subseteq B \text{ for some } i, \\ B \cap V_z & \text{if } |B \cap V_z| = 1, \\ B & \text{if } B \cap V_z = \varnothing. \end{cases}$$

As $\mathcal{G}_z$ is a generator for z, every extension of $\vec{c}'$ to Σ must contain z as a cycle. Moreover, since $\mathcal{G}_z$ is uncovered, these three cases exhaust the possible ways a budget in $\Sigma|_{\mathcal{G}_z}$ can intersect $V_z = \{x_0, \ldots, x_N\}$. We now define an extension c' via:

$$c'(B) = \begin{cases} B \setminus \left(\cup_{B \in \mathcal{G}_z} B \right) & \text{if } B \notin \Sigma|_{\mathcal{G}_z}, \\ \tilde{c}'(B) & \text{else.} \end{cases}$$

Let:

$$y_0 \succsim_{c'} y_1 \succsim_{c'} \cdots \succsim_{c'} y_M \succ_{c'} y_0,$$

denote an arbitrary revealed preference cycle of c'. First, note that for all $0 \leq i \leq M$, we must have $y_i \notin X \setminus \left(\cup_{B \in \mathcal{G}_z} B \right)$, as by construction, the only things that are c'-revealed weakly preferred to elements in this set are other elements in this set, and no alternative is ever c'-strictly revealed preferred to any element of this set. Thus for all i, we have $y_i \in \left(\cup_{B \in \mathcal{G}_z} B \right)$.

Suppose then that $y_i \in \left(\cup_{B \in \mathcal{G}_z} B \right) \setminus V_z$. By induction, every y_j with $j \geq i$ must also belong to $\left(\cup_{B \in \mathcal{G}_z} B \right) \setminus V_z$, as $y_i \succsim_{c'} y_{i+1}$ implies that $y_i, y_{i+1} \in B$ where $B \cap V_z = \varnothing$. Thus, in particular $y_M \in \left(\cup_{B \in \mathcal{G}_z} B \right) \setminus V_z$, which is a contradiction as no elements in this set are ever c'-strictly revealed preferred to any other alternative.

Thus we conclude that every y_i must belong to V_z. Moreover, each $\{y_i, y_{i+1}\} \in E_z$, as by construction, if $y_i \succsim_{c'} y_{i+1}$, and $y_i, y_{i+1} \in V_z$, then $y_i, y_{i+1} \in B$ where $B \in \Sigma|_{\mathcal{G}_z}$. This means that if $\{y_i, y_{i+1}\} \notin E_z$, that B would cover $\mathcal{G}_z$, a contradiction. Thus every $\{y_i, y_{i+1}\} \in E_z$, and hence that, in fact, this cycle is precisely z. Since this cycle was arbitrary, we conclude that c' possesses precisely one cycle: z, as desired.

$(i) \implies (ii)$: Suppose now that $\mathcal{G}_z$ is covered, and that $c' \in \mathcal{C}(X, \Sigma)$ is arbitrary, other than (i) having cycle z and $\mathcal{G}_z$ generating z for c'. Let $B \in \Sigma$ cover $\mathcal{G}_z$; we consider two cases.

Case: $c'(B) \cap V_z = \varnothing$. Let $x^* \in c'(B)$; since $x^* \in B$ and B covers $\mathcal{G}_z$, we know $x^* \in \left(\cup_{\hat{B} \in \mathcal{G}_z} \hat{B} \right)$ and hence that $x^* \in B'$ for some $B' \in \mathcal{G}_z$. Since $c'(B') \cap V_z \neq \varnothing$, we know that, for some $0 \leq i, j \leq N$, we have:

$$x^* \succ_{c'} x_i \succsim_{c'} \cdots \succsim_{c'} x_j \succsim_{c'} x^*,$$

where $x_j \in c'(B')$, and $x^* \succ_{c'} x_i$ because $x^* \in c'(B)$ and $c'(B) \cap V_z = \varnothing$ by hypothesis. Since $x^* \notin V_z$, this cycle must necessarily be distinct from z.

Case: $c'(B) \cap V_z \neq \varnothing$. Let $x_i \in c'(B) \cap V_z$. If $x_j \in B \cap V_z$, where $\{x_i, x_j\} \notin E_z$ then, we obtain the cycle:

$$x_i \succsim_{c'} x_j \succsim_{c'} \cdots \succsim_{c'} x_N \succ_{c'} x_0 \succsim_{c'} \cdots \succsim_{c'} x_i.$$

Since x_i and x_j are non-adjacent in z, this means that at least x_{i+1} does not appear in the above cycle and hence it is distinct from z. If $c'(B)$ contains no other element of V_z there are two sub-cases: either $V_z \subseteq B$ or B contains two elements of V_z which do not form an edge in E_z. Consider first the former. If $V_z \subseteq B$, then we have $x_i \in c'(B)$ and $x_{i+1} \in B \setminus c'(B)$. Hence $x_i \succ_{c'} x_{i+1}$ and $x_{i+1} \succsim_{c'} x_i$, and this two-cycle is a distinct cycle from z, as z possesses an uncovered cyclic collection $\mathcal{G}_z$, and hence by definition must be supported on a loop (i.e. be of length greater than two). Consider then the latter sub-case. We have already shown that if $B \cap V_z$ contains any alternative non-adjacent in z to x_i then there is another cycle, hence suppose that every element of $B \cap V_z$ is adjacent to x_i. Since B contains a pair of alternatives non-adjacent in z, and every element of $B \cap V_z$ must be adjacent to x_i in z, this implies that $B \cap V_z = \{x_{i-1}, x_i, x_{i+1}\}$. Since we have shown already that if x_{i-1} or x_{i+1} belongs to $c'(B)$ there is another cycle, suppose that $x_i \in c'(B)$ and x_{i-1}, x_{i+1} are not. Then by an analogous argument to the prior sub-case we obtain a WARP violation. Thus we conclude that c' must contain some cycle other than z. □

A.2 Proof of Theorem 1

Proof. We first show that $|\Sigma| < \infty$ implies that, for any $c \in \mathcal{C}(X, \Sigma)$, at least one finite irrationality kernel must exist. By definition, $\Longrightarrow^*$ is a preorder. Consider $\hat{\mathcal{Z}} := \mathcal{Z}/\Longleftrightarrow^*$, i.e. the set of cycles of c modulo the equivalence relation $\Longleftrightarrow^*$. We claim $\hat{\mathcal{Z}}$ must be finite. To see this suppose, for sake of contradiction, that $\hat{\mathcal{Z}}$ is infinite. Then there exists (distinct) $\big\{[z_1], [z_2], \dots\big\} \subseteq \hat{\mathcal{Z}}$. Let $\{\mathcal{G}_{z_i}\}_{i=1}^{\infty}$ denote an arbitrary choice of generator for an arbitrary choice of cycle within each equivalence class. This defines a map $\{[z_1], [z_2], \dots\} \to 2^{\Sigma}$, via $[z_i] \mapsto \mathcal{G}_{z_i}$. Since $|\Sigma| < \infty$, by the pigeon-hole principle this map cannot be injective, and hence there exists $i \neq j$ such that $\mathcal{G}_{z_i} = \mathcal{G}_{z_j}$ and thus some $\tilde{z}_i \in [z_i]$ and $\tilde{z}_j \in [z_j]$ have a common generator, and therefore are $\Longleftrightarrow^*$-equivalent, implying $[z_i] = [z_j]$, a contradiction. We conclude $\hat{\mathcal{Z}}$ is finite, and hence $(\hat{\mathcal{Z}}, \Longrightarrow^*)$ is a finite partially ordered set. In particular, if non-empty, it must contain at least one undominated element. Forming $\mathcal{I} \subseteq \mathcal{Z}$ by choosing one cycle in $\mathcal{Z}$ from each $\Longrightarrow^*$-undominated equivalence class in $(\hat{\mathcal{Z}}, \Longrightarrow^*)$ then yields a finite irrationality kernel as desired.

We now show that any two irrationality kernels $\mathcal{I}$ and $\mathcal{I}'$ for a given $c \in \mathcal{C}(X, \Sigma)$ are equicardinal. Suppose then that $\mathcal{I}, \mathcal{I}' \subseteq \mathcal{Z}$ are distinct irrationality kernels for c. Then without loss of generality, there exists some $z \in \mathcal{I} \setminus \mathcal{I}'$. Since $\mathcal{I}'$ is an irrationality kernel, there exists some $z' \in \mathcal{I}'$ such that $z' \Longrightarrow^* z$. Since $\mathcal{I}$ is an irrationality kernel, by (IK.1) it must be the case that $z' \notin \mathcal{I}$ but that, by (IK.2), there exists some $z'' \in \mathcal{I}$ such that $z'' \Longrightarrow^* z'$. Since $\Longrightarrow^*$ is transitive, $z'' \Longrightarrow^* z$ as well, and hence $z'' = z$ and $z \Longleftrightarrow^* z'$. Moreover, by (IK.1) z (resp. z') must be the only element of $\mathcal{I}$ (resp. $\mathcal{I}'$) with this property.

Define $\phi : \mathcal{I} \to \mathcal{I}'$ via:

$$\phi(z) = \begin{cases} z & \text{if } z \in \mathcal{I} \cap \mathcal{I}' \\ z' & \text{if } z' \in \mathcal{I}' \text{ and } z' \Longleftrightarrow^* z. \end{cases}$$

In light of the above, we have shown this map is well-defined and a bijection between $\mathcal{I}$ and $\mathcal{I}'$; we conclude they are equicardinal, and, in particular, both finite. □

A.3 Proof of Theorem 2

Proof. Suppose, for purposes of contraposition, there exists a loop $\gamma = (V_\gamma, E_\gamma) \subseteq \Gamma(X, \Sigma)$ with a cyclic collection $\mathcal{B}_\gamma$ that satisfies neither (i) nor (ii). We will show that there exists a loop $\gamma' \subseteq \Gamma(X, \Sigma)$ that supports a cycle and has a covered cyclic collection, and hence by Proposition 1 that there exists a choice correspondence whose set of cycles admits a non-trivial relation.

To this end, let $B^* \in \mathcal{B}_\gamma$ denote a budget of cardinality > 2 (the existence of which is guaranteed by hypothesis). We consider two cases.

Case: $V_\gamma \nsubseteq B^*$

We suppose first that V_γ contains some point not in B^*. Let:

$$E^* = \{e \in E_\gamma : e \subseteq B^*\}$$

denote those edges in γ that are wholly contained in B^*. By hypothesis, both E^* and $E_\gamma \setminus E^*$ are non-empty. As γ is a loop, the edge set of the subgraph $\tilde{\gamma} = (V_\gamma, E_\gamma \setminus E^*)$ is a finite disjoint union of paths, the endpoints of which all lie in B^*.[22] Let $B^o = \{x \in B^* : \nexists e \in E_\gamma \setminus E^* \text{ s.t. } x \in e\}$ denote those elements of B^* that are not contained in any edge of $E_\gamma \setminus E^*$. Suppose this set is non-empty. Enumerate B^o as $\{b_0, \dots, b_K\}$, define $E^o = \big\{\{b_0, b_1\}, \dots, \{b_{K-1}, b_K\}\big\} \subseteq E_\Gamma$, and enumerate the path components of $\tilde{\gamma}$ as $\tilde{\gamma}_0, \dots, \tilde{\gamma}_J$.[23] For each $0 \leq j \leq J$, choose one of the two degree-one vertices of the path as the 'head,' which we will write as v_j^+, and the other as the 'tail,' denoted v_j^-. Define:

$$\hat{E} = E^o \cup E_{\tilde{\gamma}} \cup \big\{\{v_j^+, v_{j+1}^-\}\big\}_{j=0}^{J-1} \cup \{b_0, v_0^-\} \cup \{v_J^+, b_K\},$$

where $E_{\tilde{\gamma}} = E_\gamma \setminus E^*$. If, instead, B^o was empty, define $\hat{E}$ analogously, but replace $\{b_0, v_0^-\} \cup \{v_J^+, b_K\}$ with $\{v_0^-, v_J^+\}$ in the above expression. Now, every element of $\hat{E}$ is either an element of E_γ (and hence in E_Γ) or is a subset of B^*, and hence in E_Γ, thus $\hat{\gamma} = (V_\gamma \cup B^*, \hat{E}) \subseteq \Gamma(X, \Sigma)$. Moreover, by construction, $\hat{\gamma}$ is a loop whose every cyclic collection is covered: any cyclic collection for $\hat{\gamma}$ must contain $V_\gamma \cup B^*$ in the union of its elements. Since $|B^*| \geq 3$, this implies that B^* must cover the cyclic collection. It remains to show that $\hat{\gamma}$ supports a cycle for some $c \in \mathcal{C}(X, \Sigma)$. Since B^* does not contain V_γ we have that $\mathcal{B}_\gamma \setminus \{B^*\} \neq \varnothing$ and, in particular, there exists some $x^* \in B^*$ such that the two

[22] A path is a finite tree graph with two nodes of degree one, and all other nodes of degree 2.

[23] As $\Gamma(X, \Sigma)$ is finite by hypothesis, so too is every budget and hence $\hat{B}$.

edges of $\hat{\gamma}$ containing x^* are not both subsets of B^*. Then, define:

$$c(B) = \begin{cases} B^* \setminus \{x^*\} & \text{if } B = B^* \\ B & \text{else} \end{cases}$$

yields a revealed preference pair with a cycle supported on $\hat{\gamma}$.

Case: $V_\gamma \subseteq B^*$

Suppose first that $V_\gamma \subsetneq B^*$. Enumerate $B^* \setminus V_\gamma$ as $\{b_0, \ldots, b_K\}$, and suppose $e = \{x, y\} \in E_\gamma$ is an edge contained in B^*. Then let:

$$\hat{E} = \big(E_\gamma \setminus \{e\}\big) \cup \big\{\{b_k, b_{k+1}\}\big\}_{k=0}^{K} - 1 \cup \{x, b_0\} \cup \{b_K, y\},$$

and define $\hat{\gamma} = (V_\gamma, \hat{E})$. If, contrary to our initial assumption, $V_\gamma = B^*$, then simply define $\hat{\gamma} = \gamma$. Then every cyclic collection of $\hat{\gamma}$ is covered, as its vertex set is simply B^*, which is itself a budget. To show that $\hat{\gamma}$ supports a cycle, observe that by hypothesis, there is some budget $B^{**} \in \mathcal{B}_\gamma \setminus \{B^*\}$ that contains an edge e' of γ different from e. Denote such an $e' = \{a, b\}$. Then $e' \in \hat{E}$ and:

$$c(B) = \begin{cases} \{a\} & \text{if } B = B^{**} \\ B & \text{else} \end{cases}$$

yields a revealed preference cycle supported on $\hat{\gamma}$. □

A.4 Proof of Theorem 3

Let $\mathcal{W}(X, \Sigma)$ (resp. $\mathcal{G}(X, \Sigma)$) denote the set of choice correspondences satisfying the weak (resp. generalized) axioms.

A.4.1 Preliminary Lemmas

Lemma 1. *Let (X, Σ) be a choice environment and let γ be a loop in $\Gamma(X, \Sigma)$. Then there exists choice function $c \in \mathcal{W}(X, \Sigma)$ such that $\succsim_c |_{E_\gamma}$ is a cycle if and only if there exists a cyclic collection $\mathcal{B}_\gamma$ and choice function $\tilde{c} \in \mathcal{W}(X, \Sigma|_{\mathcal{B}_\gamma})$ such that $\succsim_{\tilde{c}} |_{E_\gamma}$ is a cycle.*

Proof. ($\Longrightarrow$): Suppose there exists a $c \in \mathcal{W}(X, \Sigma)$ such that $\succsim_c|_{E_\gamma}$ is a cycle. Then there exists some cyclic collection $\mathcal{B}_\gamma$ with the property that the choices inducing $\succsim_c|_{E_\gamma}$ are all made on elements of $\mathcal{B}_\gamma$. Then the restriction of c to $\Sigma|_{\mathcal{B}_\gamma}$ must still obey the weak axiom, and clearly satisfies the conclusion of the lemma.

($\Longleftarrow$): Suppose now there exists a cyclic collection $\mathcal{B}_\gamma$ and a $\tilde{c} \in \mathcal{W}(X, \Sigma|_{\mathcal{B}_\gamma})$ such that $\succsim_{\tilde{c}}|_{E_\gamma}$ is a cycle. Define an extension of $\tilde{c}$ to all of Σ as follows:

$$c(B) = \begin{cases} \tilde{c}(B) & \text{if } B \in \Sigma|_{\mathcal{B}_\gamma} \\ B \setminus \big(\cup_{\tilde{B} \in \mathcal{B}_\gamma} \tilde{B} \big) & \text{else.} \end{cases}$$

This defines a choice correspondence in $\mathcal{W}(X, \Sigma)$, for if $x \succsim_c y$ for distinct x, y, either $x, y \in \cup_{\tilde{B} \in \mathcal{B}_\gamma} \tilde{B}$, in which case there can be no violation of the weak axiom as $\tilde{c}$ is in $\mathcal{W}\big(X, \Sigma|_{\mathcal{B}_\gamma}\big)$, or $x \notin \cup_{\tilde{B} \in \mathcal{B}_\gamma} \tilde{B}$, in which case by construction $\neg\, y \succ_c x$, and thus $c \in \mathcal{W}(X, \Sigma)$. $\square$

Lemma 2. *Let (X, Σ) be a choice environment and let γ be a loop in $\Gamma(X, \Sigma)$ with $|V_\gamma| = 3$. Then there exists a choice correspondence $c \in \mathcal{W}(X, \Sigma)$ with $\succsim_c|_{E_\gamma}$ a cycle if and only if there exists a cyclic collection $\mathcal{B}_\gamma$ that is not covered.*

Proof. ($\Longleftarrow$): Suppose that $\mathcal{B}_\gamma$ is an uncovered cyclic collection for γ of minimal cardinality. Let us denote $E_\gamma = \{e_0, e_1, e_2\}$. Then, in particular, for every $e_j \in E_\gamma$, there is a unique $B_j \in \mathcal{B}_\gamma$ with $e_j \subseteq B_j$. Define $\tilde{c} \in \mathcal{C}(X, \Sigma|_{\mathcal{B}_\gamma})$ via:

$$\tilde{c}(B) = \begin{cases} e_j \cap e_{j+1} & \text{if } \exists\, e_j \in E_\gamma \text{ s.t. } B \cap V_\gamma = e_j \\ B \cap V_\gamma & \text{if } |B \cap V_\gamma| = 1 \\ B & \text{else.} \end{cases}$$

where all subscripts are taken mod-3. Note $\tilde{c}$ is well-defined, as $\mathcal{B}_\gamma$ is uncovered from which it follows the first two cases exhaust the possibilities for budgets in $\Sigma|_{\mathcal{B}_\gamma}$ that intersect V_γ. Moreover, $\tilde{c} \in \mathcal{W}(X, \Sigma|_{\mathcal{B}_\gamma})$. First, observe the restriction of the pair $(\succsim_{\tilde{c}}, \succ_{\tilde{c}})|_{E_\gamma}$ satisfies the weak axiom. But the only alternatives $\tilde{c}$ reveals strictly preferred to any others all lie in V_γ, and the only goods ever revealed preferred to elements of V_γ also lie in V_γ. Hence $\tilde{c} \in \mathcal{W}\big(X, B \in \Sigma|_{\mathcal{B}_\gamma}\big)$, and by Lemma 1 there exists a $c \in \mathcal{W}(X,\Sigma)$ such that $\succsim_c|_{E_\gamma}$ is cyclic.

($\Longrightarrow$): Let $c \in \mathcal{W}(X,\Sigma)$ be such that $\succsim_c|_{E_\gamma}$ is cyclic. Then there exists a cyclic collection $\mathcal{B}_\gamma$ on which choices generating the cycle $\succsim_c|_{E_\gamma}$ are made; fix such a collection. We now show that this cyclic collection must be uncovered, lest there exist some $B \in \Sigma|_{\mathcal{B}_\gamma}$ such that $V_\gamma \subseteq B$. Suppose, for sake of contradiction, that such a B exists.

Case 1: Suppose first that $c(B) \cap V_\gamma \neq \varnothing$. Then either $c(B)$ induces complete indifference across V_γ, or there exists some pair of elements of V_γ that is either strictly preferred to, or strictly dominated by the third element. Both possibilities preclude the existence of the cycle $\succsim_c|_{E_\gamma}$ for any $c \in \mathcal{W}(X,\Sigma)$.

Case 2: Suppose then that $c(B) \cap V_\gamma = \varnothing$: then for all $x \in V_\gamma$ and $y \in c(B)$ we have $y \succ_c x$. But $c(B) \subset B \subseteq \cup_{\tilde{B}\in\mathcal{B}_\gamma}\tilde{B}$, and since for all $x \in V_\gamma$ there exists some $\tilde{B}$ such that $x \in c(\tilde{B})$, there exists an $\tilde{x} \in V_\gamma$ and $\tilde{B} \in \mathcal{B}_\gamma$ such that $\tilde{x}, y \in \tilde{B}$ and $\tilde{x} \in c(\tilde{B})$. This contradicts our hypothesis that $c \in \mathcal{W}(X,\Sigma)$. $\square$

Lemma 3. *Let (X,Σ) be a choice environment and let γ be a loop in $\Gamma(X,\Sigma)$ with $|V_\gamma| > 3$. Suppose there exists a choice correspondence $c \in \mathcal{W}(X,\Sigma)$ with $\succsim_c|_{E_\gamma}$ a cycle. If every cyclic collection $\mathcal{B}_\gamma$ is covered, then there exists a loop γ' in $\Gamma(X,\Sigma)$ such that $|V'_\gamma| < |V_\gamma|$ and with $\succsim_c|_{E_{\gamma'}}$ a cycle.*

Proof. Let $\mathcal{B}_\gamma$ be a minimal cyclic collection on which choices inducing $\succsim_c|_{E_\gamma}$ are made, and suppose $\mathcal{B}_\gamma$ is covered. Then there exists some $B \in \Sigma|_{\mathcal{B}_\gamma}$ such that B contains a non-adjacent pair of vertices of γ. We proceed in two cases.

Case 1: Suppose first that $c(B)$ does not intersect V_γ. Let $x_k, x_{k'} \in B \cap V_\gamma$ be one such non-adjacent pair of vertices, and let $y \in c(B)$. As $c(B) \subseteq B \subseteq \cup_{\tilde{B}\in\mathcal{B}_\gamma}\tilde{B}$, and $\mathcal{B}_\gamma$ is a minimal cyclic collection on which choices inducing the cycle $\succsim_c|_{E_\gamma}$ are made, there is some $\tilde{B}_{k^*} \in \mathcal{B}_\gamma$ containing y, such that there is some $x_{k^*} \in c(\tilde{B}_{k^*}) \cap V_\gamma$. Without loss of generality, let $x_{k'} \succsim_c \cdots \succsim_c x_{k^*} \succsim_c$

$\cdots \succsim_c x_k$. In particular, by our hypothesis that c obeys the weak axiom, we cannot have $x_{k^*} = x_k$ (or $x_{k'}$).[24] As $c(B)$ does not contain any element of V_γ by hypothesis, but $x_{k'} \in B$, we have $y \succ_c x_{k'}$, and, as $x_{k^*}, y \in \tilde{B}_{k^*}$, it follows $x_{k^*} \succsim_c y$. Thus: $y \succ_c x_{k'} \succsim_c \cdots x_{k^*} \succsim_c y$. Define γ' to be the graph with $V_{\gamma'}$ given by the above collection of points, and $E_{\gamma'}$ consisting of those pairs related in the above cycle (clearly as there is a non-empty revealed preference for each pair this forms a loop in $\Gamma(X, \Sigma)$). By construction, $\succsim_c |_{E_{\gamma'}}$ is a cycle. Now, since $x_{k^*} \neq x_k$, $x_k \notin V_{\gamma'}$. Moreover, since x_k and $x_{k'}$ are non-adjacent in γ, under $\succsim_c |_{E_\gamma}$ we also have: $x_k \succsim_c \cdots \succsim_c \bar{x} \succsim_c \cdots \succsim_c x_{k'}$ along the 'other side' of the loop. Thus we also have that $\bar{x} \notin V_{\gamma'}$. So while we have added a point y not in V_γ to our $V_{\gamma'}$, we have omitted at least two others, x_k and $\bar{x}$, and we conclude: $|V_{\gamma'}| < |V_\gamma|$ as required.

Case 2: Suppose now that $c(B)$ intersects V_γ. As B contains the non-adjacent pair $x_k, x_{k'} \in V_\gamma$, the only way that $c(B)$ can avoid revealing a preference between x_k and $x_{k'}$ is if neither is in but both are adjacent in γ to $c(B)$. Moreover, this argument holds for every non-adjacent pair of vertices of γ contained in B. Now, if $c(B)$ induces a revealed preference $x_i \succsim_c x_j$ between any pair of non-adjacent vertices $x_i, x_j \in V_\gamma$ this partitions $\succsim_c |_{E_\gamma}$ into two sub-cycles, one of which must always contain a strict relation (either from $\succsim_c |_{E_\gamma}$ or resulting from a strict revealed preference between x_i and x_j). Letting γ' be defined by the vertices and pairs supporting any such sub-cycle suffices to prove the claim. Thus suppose that $c(B)$ does not induce any revealed preference between any non-adjacent pair (lest we be done). Thus $c(B)$ is adjacent to both x_k and $x_{k'}$ (and hence singleton) and $c(B) = \{x^*\}$ induces both $x_k \prec_c x^* \succ_c x_{k'}$. But these three points are all elements of V_γ, hence by virtue of $\succsim_c |_{E_\gamma}$ being a cycle we have either $x_k \succsim_c x^* \succsim_c x_{x'}$ or the reverse. But both of these yield contradiction via a violation of the weak axiom, and hence there exists a strictly shorter $\succsim_c$-cycle. □

[24]As $y \succ_c x_k$ and $y \succ_c x_{k'}$ by hypothesis, but $x_{k^*} \succsim_c y$ via choice on B_{k^*}.

A.4.2 Proof of Theorem 3

Theorem. *Let (X, Σ) be a choice environment. Then $\mathcal{W}(X, \Sigma) = \mathcal{G}(X, \Sigma)$ if and only if Σ is well-covered.*

Proof. ($\Longleftarrow$): For purposes of contraposition, suppose that $\mathcal{W}(X, \Sigma) \neq \mathcal{G}(X, \Sigma)$. Then there exists some loop γ in the budget graph $\Gamma(X, \Sigma)$ and some choice correspondence $c \in \mathcal{W}(X, \Sigma)$ such that $\succsim_c |_{E_\gamma}$ is a cycle. If $|V_\gamma| = 3$, then by Lemma 2, Σ is not well-covered and we are done. Hence suppose γ is of length strictly greater than three. Then there exists some cyclic collection $\mathcal{B}_\gamma$ on which choices generating the cycle $\succsim_c |_{E_\gamma}$ are made. If $\mathcal{B}_\gamma$ is not covered, we are done, hence suppose it is. Then by Lemma 3 there exists a loop γ' in the budget graph of strictly shorter length such that $\succsim_c |_{E_{\gamma'}}$ is also a cycle. As we have already concluded this process cannot repeat until it hits a three-cycle, we conclude that at some stage, there exists some loop $\gamma^{(n)}$ for which there exists a cyclic collection $\mathcal{B}_{\gamma^{(n)}}$ which is not covered and hence Σ is not well-covered.

($\Longrightarrow$): We again proceed by contraposition. If a cyclic collection for a budget graph loop of length 3 is uncovered, by Lemma 2, we immediately obtain $\mathcal{W}(X, \Sigma) \neq \mathcal{G}(X, \Sigma)$. Suppose then there exists some loop γ with $|V_\gamma| > 3$ with a cyclic collection $\mathcal{B}_\gamma$ that is uncovered (without loss of generality, let $\mathcal{B}_\gamma$ be a minimal such uncovered cyclic collection) In particular, let $E_\gamma = \{e_0, \ldots, e_{J-1}\}$. By virtue of γ being uncovered, for each $e_j \in E_\gamma$ there exists a $\tilde{B}_j \in \mathcal{B}_\gamma$ such that for all $j \in \{0, \ldots, J-1\}$ we have $e_j = \tilde{B}_j \cap V_\gamma$, and by the minimality of $\mathcal{B}_\gamma$, these $\{\tilde{B}_j\}$ are unique and completely exhaust $\mathcal{B}_\gamma$. Furthermore, for all $B \in \Sigma|_{\mathcal{B}_\gamma}$, $B \cap V_\gamma$ necessarily also either equals some e_j, is singleton, or is empty.[25] Thus, letting (subscripts taken mod-J):

$$\tilde{c}(B) = \begin{cases} e_j \cap e_{j+1} & \text{if } \exists\, e_j \text{ s.t. } e_j = B \cap V_\gamma \\ B \cap V_\gamma & \text{if } |B \cap V_\gamma| = 1 \\ B & \text{else,} \end{cases}$$

[25] The loop γ, viewed as a loop in the subgraph $\Gamma(X, \Sigma|_{\mathcal{B}_\gamma})$, is what is sometimes referred to as 'chordless' in graph theory.

we obtain a choice correspondence $c \in \mathcal{W}(X, \Sigma|_{\mathcal{B}_\gamma})$ by an argument identical to that in the proof of Lemma 2, only for a longer cycle. Clearly $\succsim_{\tilde{c}}|_{E_\gamma}$ is cyclic and by Lemma 1 this extends to a choice correspondence in $c \in \mathcal{W}(X, \Sigma)$ such that $\succsim_c|_{E_\gamma}$ is cyclic, and hence $\mathcal{W}(X, \Sigma) \neq \mathcal{G}(X, \Sigma)$. Thus, by contraposition, $\mathcal{W}(X, \Sigma) = \mathcal{G}(X, \Sigma)$ implies the well-coveredness of Σ. $\square$

References

Afriat, Sydney N, "The construction of utility functions from expenditure data," *International economic review*, 1967, *8* (1), 67–77.

_ , "On a system of inequalities in demand analysis: an extension of the classical method," *International economic review*, 1973, pp. 460–472.

Apesteguia, Jose and Miguel A Ballester, "A measure of rationality and welfare," *Journal of Political Economy*, 2015, *123* (6), 1278–1310.

Arrow, Kenneth J, "Rational choice functions and orderings," *Economica*, 1959, *26* (102), 121–127.

Cherchye, Laurens, Thomas Demuynck, and Bram De Rock, "Transitivity of preferences: when does it matter?," *Theoretical Economics*, 2018, *13* (3), 1043–1076.

Dean, Mark and Daniel Martin, "Measuring rationality with the minimum cost of revealed preference violations," *Review of Economics and Statistics*, 2016, *98* (3), 524–534.

Dziewulski, Paweł, Joshua Lanier, and John K-H Quah, "Revealed preference and revealed preference cycles: a survey," *Journal of Mathematical Economics*, 2024, p. 103016.

Echenique, F, "On the meaning of the critical cost efficiency index (August 2021)," Technical Report, Working Paper 2021.

Echenique, Federico, Sangmok Lee, and Matthew Shum, "The money pump as a measure of revealed preference violations," *Journal of Political Economy*, 2011, *119* (6), 1201–1223.

Famulari, Melissa, "A household-based, nonparametric test of demand theory," *The Review of Economics and Statistics*, 1995, pp. 372–382.

Harbaugh, William T, Kate Krause, and Timothy R Berry, "GARP for kids: On the development of rational choice behavior," *American Economic Review*, 2001, *91* (5), 1539–1545.

Houthakker, Hendrik S, "Revealed preference and the utility function," *Economica*, 1950, *17* (66), 159–174.

Houtman, Martijn and J Maks, "Determining all maximal data subsets consistent with revealed preference," *Kwantitatieve methoden*, 1985, *19* (1), 89–104.

Kalai, Gil, Ariel Rubinstein, and Ran Spiegler, "Rationalizing choice functions by multiple rationales," *Econometrica*, 2002, *70* (6), 2481–2488.

Kushnir, Alexey I and Lev Lokutsievskiy, "On the Equivalence of Weak- and Cyclic-Monotonicity," *Available at SSRN 3422846*, 2019.

Lanier, Joshua and John K-H Quah, "Goodness-of-fit and utility estimation: what's possible and what's not," *arXiv preprint arXiv:2405.08464*, 2024.

_ **, Matthew Polisson, and John K-H Quah**, "Money Pumps and Bounded Rationality," *arXiv preprint arXiv:2404.04843*, 2024.

Mononen, Lasse, "Foundations of measuring (cardinal) rationality," Technical Report, Mimeo, University of Zürich 2020.

_ , "Computing and comparing measures of rationality," Technical Report, Working Paper 2023.

Morgenstern, Oskar and John Von Neumann, *Theory of games and economic behavior*, Princeton university press, 1944.

Ok, Efe A, Pietro Ortoleva, and Gil Riella, "Revealed (p) reference theory," *American Economic Review*, 2015, *105* (1), 299–321.

Oxley, James G, *Matroid theory*, Vol. 3, Oxford University Press, USA, 2006.

Ribeiro, Mauricio, "Comparative rationality," *Unpublished Working Paper*, 2020.

Richter, Marcel K, "Revealed preference theory," *Econometrica: Journal of the Econometric Society*, 1966, pp. 635–645.

Saks, Michael and Lan Yu, "Weak monotonicity suffices for truthfulness on convex domains," in "Proceedings of the 6th ACM conference on Electronic commerce" ACM 2005, pp. 286–293.

Samuelson, Paul A, "A note on the pure theory of consumer's behaviour," *Economica*, 1938, *5* (17), 61–71.

Swofford, James L and Gerald A Whitney, "Flexible Functional Forms and the Utility Approach to the Demand for Money: A Nonparametric Analysis: Note," *Journal of Money, Credit and Banking*, 1986, *18* (3), 383–389.

Szpilrajn, Edward, "Sur l'extension de l'ordre partiel," *Fundamenta Mathematicae*, 1930, *16* (1), 386–389.

Varian, Hal R, "Goodness-of-fit in optimizing models," *Journal of Econometrics*, 1990, *46* (1-2), 125–140.