The Inconsistency Rank*

Peter Caradonna†

October 20, 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.

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

*I would like to thank Axel Anderson, Chris Chambers, John Quah, Mauricio Ribeiro, Christopher Turansick, and Andrea Wilson for various helpful conversations over the course of this project, as well as seminar audiences at SAET 2019 and D-TEA 2020. 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.

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. As a consequence, in many practical settings, the data often fail to pass such an exact test.[1] Thus instead, what is needed are means of quantifying the severity, or magnitude, of observed deviations from rationality.

This paper provides a principled, transparent method of quantifying the degree of irrationality observed in any choice data set. Our measure is rooted in the basic observation that choice cycles rarely occur in isolation. Often, 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 are not indicative of 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 to cyclic choices over certain menus forcing subsequent choices to always 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 from $\{a, b\}$, b from $\{b, c\}$ 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

[1] For example, Harbaugh et al. (2001) find that 23 out of 31 second graders, 16 out of 42 sixth graders, and 19 out of 55 college undergraduates exhibited at least one cycle over choices between simple menus of juice and chips. In more a more complex setting, Echenique et al. (2011) find that 396 out of 494 individuals exhibit at least one choice cycle.

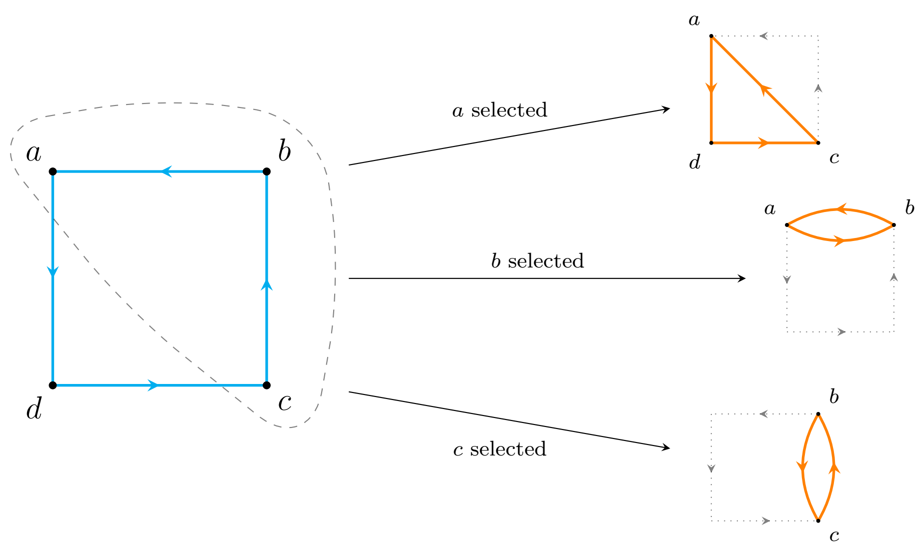

Figure 1: Given the subject's cyclic choices from the binary menus $\{a, b\}$, $\{b, c\}$, $\{c, d\}$, and $\{a, d\}$, every possible choice from $\{a, b, c\}$ creates at least one additional, knock-on cycle.

menus forces them to necessarily create another choice cycle now, regardless of their selection. If a is not chosen exclusively as the most-preferred alternative from $\{a, b, c\}$, a pairwise reversal obtains relative to their earlier choices.[2] Conversely, if a is indeed uniquely chosen from this new menu, then a is revealed preferable to c, creating a different, new cycle, this time between a, c, and d; see Figure 1. ■

In this example, the structure of the set of menus ensured that any cycle of choices over all four alternatives could never occur in isolation: any such cycle necessarily forced at least one other elsewhere, when choices from the tripleton menu are accounted for. In such cases, we interpret the forcing cycle as justifying, or *explaining*, the presence of the forced cycle.

To define our measure, we consider collections 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

[2]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.

the data. In general, many kernels will exist for a given data set. We show, however, that any two always contain precisely the same number of cycles. We define the our index, the inconsistency rank, to be the size of any such collection.

As illustrated in Example 1, the structure of the collection of menus may lead to non-trivial dependencies between cycles. Left unaccounted for, these dependencies lead many existing indices to 'double count' the effect of cycles, by mistakenly treating knock-on effects as evidence of further irrationality.[3]

In contrast, the inconsistency rank measures the number of distinct, *independent* choice cycles needed to fully account for the entirety of a subject's non-rationalizable behavior. Since no two elements of any kernel can be used to explain each other, our index does not 'double count' cycles. 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 provides a cardinal measure of irrationality: it is meaningful to say, for example, that the inconsistency in one subject's choices requires twice as many independent cycles to explain as another's. Moreover, it is valid to compare the inconsistency rank of two choice correspondences defined *across* domains. Normally, when making comparisons between choices over different sets of menus, 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 volumes of observed evidence of irrationality, in absolute terms.

In Section 3 we define our model; throughout, we consider the abstract

[3]For in-depth discussion of how the inconsistency rank relates to other well-known indices, see Section 5.2.

choice framework of Richter (1966).[4] 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 in the setting of price-consumption data.

This paper is most closely related to work seeking to quantify irrationality for general choice data environments. Most recently, Ribeiro (2024) has proposed a partial ordering over choice correspondences, where one data set is more rational than another if it is consistent on any sub-collection of menus on which the other is. Other classic contributions include Houtman and Maks (1985), who propose using as an index the size of any minimal set of observations which, when dropped, render the remaining observations consistent. In practice, the number of choice cycles, or number of observations belonging to some choice cycle, are also commonly employed measures (e.g. Famulari 1995; Harbaugh et al. 2001). Kalai et al. (2002) propose using the minimal number of preference relations needed to rationalize every observed choice; Apesteguia and Ballester (2015) propose a measure related to the minimal

[4]While we allow for set-valued choice, our primary results remain equally valid when only choice functions are observed.

number of binary swaps needed to transform the revealed preference relation into a preference. We postpone a more in-depth comparative discussion of these indices to Section 5.2.

There is also an extensive literature on measuring inconsistency in the specialized setting of price-consumption data. Afriat (1973) proposes the so-called 'critical cost efficiency' index (cf. Varian 1990).[5] More recent contributions include Echenique et al. (2011) (see also Lanier et al. 2024) who propose using the money pump as a means of quantifying inconsistency. Dean and Martin (2016) propose measuring the minimum cost needed to break all revealed preference cycles.

Finally, various notions of 'loss' relative to consistency have been implicitly used in recent work. Chambers et al. (2021) design a statistical estimator for a subject's (noisily observed) preferences minimizing the Kemeny distance to the observed data. Caradonna (2024) constructs a least-squares estimator for preferences from cardinal data on preference intensity.

3 Preliminaries

Let X be an arbitrary set of **alternatives** from which an agent chooses. A **preference relation** is a complete and transitive binary relation on X. Let $\Sigma \subseteq 2^X \setminus \{\varnothing\}$ be a collection of **budgets**, reflecting which menus an empiricist is able to observe choice from. 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.

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

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

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

and for a non-empty 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 correspondences for the environment (X, Σ), and $\mathcal{C}^f(X, \Sigma)$ the sub-collection of choice functions (i.e. singleton-valued choice correspondences). Given a choice correspondence $c \in \mathcal{C}(X, \Sigma)$, a preference relation $\succeq$ on X **rationalizes** c if:[6]

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

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:

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

where $x_0, \ldots, x_N$ are distinct. A choice correspondence is rationalizable if, and only if, it satisfies the generalized axiom, i.e. it contains no choice cycles; see Richter (1966).

3.1 The Budget Graph

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

[6] A choice function is rationalizable by a preference relation if and only if it is rationalizable by a strict preference, thus our definition of rationalizability coincides with the usual notion when c is a choice function.

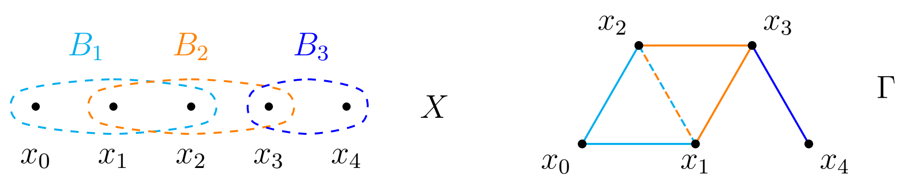

(a) A choice environment with five alternatives and three budget sets.

(b) The budget graph associated with this environment.

Figure 2: A 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.

two vertices belonging to some common budget:

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

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

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.[7] Let:

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

We refer to $\mathcal{Z}_c$ as the **cycle set** of c; in doing so, we are implicitly identifying cycles with their support. Let:

$$\mathcal{Z}_c^W = \left\{(\tilde{V}, \tilde{E}) \subseteq \mathcal{Z}_c : (\tilde{V}, \tilde{E}) \text{ is the support of a WARP violation}\right\}$$

denote the subset of $\mathcal{Z}_c$ that are (supports of) WARP violations.

[7] 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 $\tilde{E} = \{\{x_0, x_1\}\}$. This is not, however, a loop.

4 Dependencies Between Cycles

Informally, we say that a cycle forces others when, given some collection of choices yielding it, *every* possible combination of other choices from the remaining budgets in Σ necessarily creates other cycles.

This was precisely the case in Example 1: there, once the cycle of length four was chosen, every choice the subject could make from the remaining budget created (at least) one other cycle. The main result of this section is Theorem 1, which characterizes precisely how this behavior arises, in completely general environments.

4.1 Cyclic Collections and Covers

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

Given a cycle $z \in \mathcal{Z}_c \setminus \mathcal{Z}_c^W$ 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

(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.[8]

We refer to such a $\bar{B} \in \Sigma$ as a cover.[9] For a cycle $z \in \mathcal{Z}_c^W$, we define every cyclic collection to be uncovered.

Finally, we say a cyclic collection $\mathcal{G}_z$ is a **generator** for the cycle $z \in \mathcal{Z}_c$ if z is also a cycle of the choice correspondence restricted to $\mathcal{G}_z$, but not of

[8] Note that condition (i) implies (ii) if and only if $|V_\gamma| > 3$.

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

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

any proper sub-collection. Note that the same collection of menus may be a generator for multiple cycles, and may be covered when regarded as a generator for one cycle, but not another.[10]

4.2 A Characterization of Forcing

Our first result provides a complete characterization, in terms of the structure of Σ, of when a cycle z with generator $\mathcal{G}_z$ forces other cycles elsewhere in the data.

Theorem 1. *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; and

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

Moreover, this equivalence remains valid when c and c' are restricted to be choice functions.

Theorem 1 shows that when a cycle is generated by choices on some collection of budgets, it forces others if and only if the choice environment contains some budget covering the collection. This was precisely the stucture present in Example 1. There, the cycle of length four had, as a generator, the four binary menus. The tripleton budget covered this generator, and hence by Theorem 1, *every* choice from this menu necessarily created further cycles.

As a corollary, we also obtain the following 'ex-ante' characterization of forcing, describing which loops in the budget graph can support cycles which do not force others.

[10] For example, suppose $c(\{x_0, x_1\}) = \{x_0\}$ and $c(\{x_0, x_1, x_2\}) = \{x_0, x_1, x_2\}$. There are two cycles, a WARP violation z supported on $\{x_0, x_1\}$ and a cycle of length three, z', on $\{x_0, x_1, x_2\}$. The pair of these menus forms a generator for both z and z', but is uncovered as a generator for z, while covered as a generator for z'.

Corollary 1. *Let γ be a loop in $\Gamma(X, \Sigma)$. Then the following are equivalent:*

(i) Every $c \in \mathcal{C}(X, \Sigma)$ containing a cycle supported on γ contains at least one other cycle; and

(ii) Every cyclic collection for γ is covered.

Moreover, this equivalence remains valid when c is restricted to be a choice function.

5 Measurement of Inconsistency

In this section, we turn to the problem of quantifying the magnitude of any observed inconsistency in choice data. The premise of our approach is that, when choice cycles have forced others, the forced cycles should not be treated as evidence of any deeper degree of irrationality than what would implied by observation of only the forcing cycles themselves. Instead, we interpret forced cycles as artifacts arising from the structure of the experiment.

5.1 The Inconsistency Rank

Consider a choice correspondence $c \in \mathcal{C}(X, \Sigma)$, with cycle set $\mathcal{Z}_c$. Given $z, z' \in \mathcal{Z}_c$, we say that z **explains** z' (denoted $z \Longrightarrow z'$) if there exist generators $\mathcal{G}_{z'}$ for z' and $\mathcal{G}_z$ 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 were either (i) already involved in z, or (ii) 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 Theorem 1, had we observed only those choices directly involved in generating z, there are two possibilities. The first is that some collection of choices generating z also generate z'.[11] In this case, it 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' itself is not generated solely by some subcollection of $\mathcal{G}_z$, the choices making up z, but also involves choices made on budgets covering $\mathcal{G}_z$. By Theorem 1, given z, every choice on such a budget necessarily introduces further cycles. Thus while a priori the existence of the specific cycle z' may not have been a foregone conclusion, ex post we may still regard it as a particular realization of the extra cycle(s) forced by the presence of z, and the structure of the choice environment itself.

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 $\mathcal{I} \subseteq \mathcal{Z}_c$ 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.[12]

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

An irrationality kernel is simply a subset of $\mathcal{Z}_c$ with the property that no two cycles in it (even indirectly) explain each other, but which nonetheless together explain the entirety of the observed inconsistency, $\mathcal{Z}_c$.[13]

[11]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 $x_1 \succ_c x_0 \succsim_c x_1$, and z' the cycle $x_1 \succ_c x_0 \succsim_c x_2 \succsim_c x_1$. If one observes the choices generating z, here z' is also directly observed, as Σ itself is a generator for each cycle.

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

[13]Similar ideas have been studied in the context of cooperative game theory, e.g. the notion of a 'stable set' in Von Neumann and Morgenstern (1944).

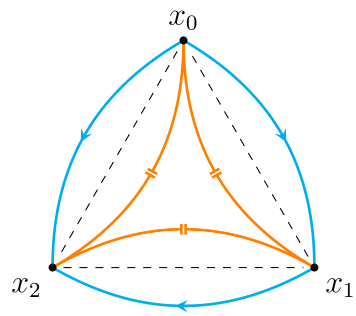

Figure 3: The revealed preference pair for Example 2. We plot $\succ_c$ in cyan, $\succsim_c \setminus \succ_c$ in orange. There are four choice cycles, three of which are WARP violations.

When $|\Sigma| < \infty$, irrationality kernels exist and are finite, for all choice correspondences in $\mathcal{C}(X, \Sigma)$. In general, there will be many irrationality kernels for a given data set; however, our next result shows that, for any correspondence, every irrationality kernel contains precisely the same number of elements.

Theorem 2. *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 2, 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 $\mathcal{Z}_c$ arising from the structure of Σ.

Example 2. Suppose $X = \{x_0, x_1, x_2\}$ and Σ is the complete domain, consisting of all subsets of X of cardinality two or more. Consider the correspondence:

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

There are four revealed preference cycles: three WARP violations, each supported on the sets $\{x_i, x_{i+1}\}$, where i is understood mod-three, and one cycle of length three; see Figure 3. Let z_i denote the WARP violation supported on $\{x_i, x_{i+1}\}$, and let z denote three-cycle. For each z_i, there is a unique generator $\mathcal{G}_i = \big\{\{x_i, x_{i+1}\}, X\big\}$. Moreover, the three generators $\mathcal{G}_i$ are precisely the possible generators for z.[14] Thus $z \Longleftrightarrow z_i$ for all i, but each pair z_i and z_j are $\Longrightarrow$-related if and only if $i = j$. But this means every pair of cycles are $\Longleftrightarrow^*$-related, hence any singleton set of cycles forms an irrationality kernel, and the inconsistency rank is one. ■

Irrationality kernels may be interpreted as a form of generalized 'basis' for the set of cycles, $\mathcal{Z}_c$. 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 $\mathcal{Z}_c$.[15] Theorem 2 then establishes that every such 'basis' for $\mathcal{Z}_c$ is of the same size.

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. 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]Note any generator for z must include at least one budget of the form $\{x_i, x_{i+1}\}$ and X itself. Given these, any other budgets would be redundant; choices on these two already yield the three-cycle hence any other budgets would violate the minimality requirement for generators.

[15]Note, however, that while every irrationality kernel is a maximal, independent set of cycles, not every such set is an irrationality kernel, nor are all such sets the same size. In Example 6, the choice correspondence c has a set of five WARP violations, all of which are pairwise independent, but an inconsistency rank of one.

[16]In practice, there may be reasons to assess certain cycles as being more or less severe than others, beyond their abstract choice-theoretic descriptions (e.g. Echenique et al. 2011; Lanier et al. 2024). In contexts where there exists some natural choice of scoring function quantifying the severity of individual cycles, our approach may be straightforwardly adapted by instead looking for a choice of irrationality kernel that minimizes the sum of the scores of the cycles it contains, and instead using this minimized sum as an index.

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

In this section, we show the inconsistency rank is not a monotone transformation of any existing index. We consider a number of well-known measures of inconsistency and show, by means of example, that in various circumstances, each yields the opposite ranking of the relative consistency of two choice correspondences than the inconsistency rank.

5.2.1 Counting Cycles

Perhaps the most straightforward inconsistency index is to simply count the number of choice cycles observed in the data (this is discussed, e.g., in Harbaugh et al. 2001). We define the cycle count index via:

$$i_{CC}(c) = |\mathcal{Z}_c|.$$

While natural, the cycle count treats all cycles independently; as a consequence, whenever there are non-trivial relations between cycles, i_{CC} 'double-counts' the forcing cycles.

Example 3. Suppose $X = \{x_0, \dots, x_4\}$ and Σ consists of the budgets:

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

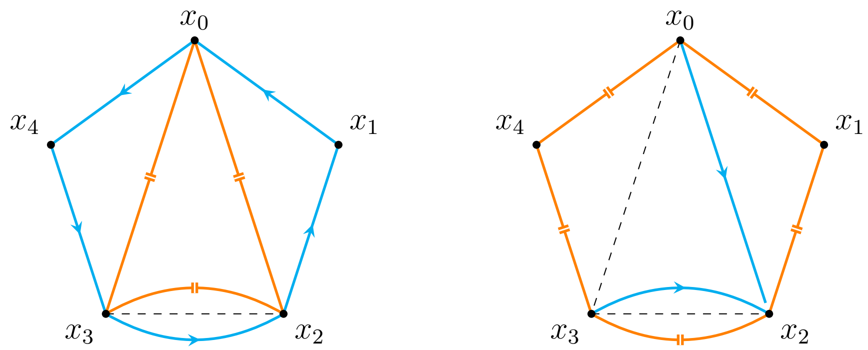

(a) The revealed preference pair for c. There are 7 cycles in $\mathcal{Z}_c$, one of which is a WARP violation.

(b) The revealed preference pair for c'. There are 4 cycles in $\mathcal{Z}_{c'}$, one of which is a WARP violation.

Figure 4: The revealed preference pairs for Example 3. The relations in $\succsim_c \setminus \succ_c$ (resp. $\succsim_{c'} \setminus \succ_{c'}$) are plotted in orange, and those in $\succ_c$ (resp. $\succ_{c'}$) are plotted in cyan.

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

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

where i is understood mod-five. The first correpsondence, c, has seven cycles: one supported on each loop of the budget graph, plus one WARP violation; see Figure 4. Hence $i_{CC}(c) = 7$. On the other hand, $i_{CC}(c') = 4$. Thus, the cycle count index ranks c as exhibiting a deeper degree of irrationality than c'.

However, note that for c, the cycle of length five explains every other, yielding an inconsistency rank $i_{IR}(c) = 1$, whereas for c', we have $i_{IR}(c') = 2$. Thus when we account for the effects of the structure of Σ on the set of choice cycles, we obtain a reversal relative to näively counting cycles. ■

5.2.2 Choices-In-Cycles

Another natural approach to quantifying inconsistency is to count the number (resp. fraction) 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. Define the number of choices-in-cycles index i_{NC} via:

$$i_{NC}(c) = \left|\{B \in \Sigma : \exists\, z \in \mathcal{Z}_c \text{ with generator } \mathcal{G}_z \text{ s.t. } B \in \mathcal{G}_z\}\right|.$$

Similarly, define the fraction of choices-in-cycles index, i_{FC}, as:

$$i_{FC}(c) = \frac{i_{NC}(c)}{|\Sigma|}.$$

As the following example shows, i_{NC} and i_{FC} penalize longer cycles more heavily than shorter, whereas the inconsistency rank simply treats all cycles as equal, leading to divergence in their assessments.[17]

Example 4. 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:

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

where i is understood mod-three, and k mod-$(K+1)$. Then $i_{NC}(c) = 6$, and its two cycles are independent, hence $i_{IR}(c) = 2$. Conversely, define:

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

[17] For example, this means that violations due to an accumulation of a large number of imperceptible differences (e.g. Armstrong 1950; Luce 1956) may be penalized much more harshly by these indices than outright incoherence of choices across a pair of visibly distinguishable alternatives. See also Dziewulski (2020).

Here, $i_{NC}(c') = K$, all of which appear in its single cycle, hence $i_{IR}(c') = 1$. Thus for K large enough, the inconsistency rank and choices-in-cycles (both i_{NC} and i_{FC}) 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. Define the Houtman-Maks index as:

$$i_{HM}(c) = \min_{\{\Sigma' \,:\, c|_{\Sigma \setminus \Sigma'} \text{ is rationalizable}\}} \left|\Sigma'\right|. \tag{2}$$

Example 5. Consider again the environment and choice correspondences from Example 3. There, $i_{HM}(c) = 2$ as one must, e.g., remove the budget $\{x_0, x_2, x_3\}$ and one binary budget to break every cycle. In contrast, $i_{HM}(c') = 1$, as it suffices to remove $\{x_0, x_2, x_3\}$ alone. However, $i_{IR}(c) = 1$ and $i_{IR}(c') = 2$, yielding the opposite ranking.

Here, the divergence is driven by the fact that the choice $c'\big(\{x_0, x_2, x_3\}\big)$ is crucial to two, independent cycles in $\mathcal{Z}_{c'}$. The Houtman-Maks index views the marginal contribution of this choice as simply 'creating inconsistency,' regardless the form it takes. In contrast, the inconsistency rank penalizes this choice more severely than it would, say, a choice that led to the formation of only a single cycle. ■

5.2.4 Multiple Rationales Index

Kalai et al. (2002) consider the minimal number of (strict) preferences needed to fully explain the entirety of a subject's choices. More formally, for any choice function $c \in \mathcal{C}^f(X, \Sigma)$, define the multiple rationales index as the minimal size of any such family:

$$i_{MR}(c) = \min \Big\{ \big|\{\succ_1, \ldots, \succ_K\}\big| : \forall B \in \Sigma,\ \exists k \text{ s.t. } c(B) \text{ is } \succ_k\text{-maximal in } B \Big\}.$$

The multiple rationales index does, in some instances, implicitly account for dependencies between cycles.[18] However, in general it yields different predictions than the inconsistency rank, as the next example shows.

Example 6. Let $X = \{x_0, \ldots, x_4\}$, and Σ consist of the sets $\{x_0, x_1\}$, $\{x_4, x_0\}$, $\{x_0, x_1, x_2\}$, $\{x_1, x_2, x_3\}$, $\{x_2, x_3, x_4\}$, $\{x_3, x_4, x_0\}$, and $\{x_4, x_0, x_1\}$. Define c via:

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

where i is understood mod-five. Let z_i denote the WARP violation supported on $\{x_i, x_{i+1}\}$, and consider the cycle z whose support is $\{x_0, x_1\}, \ldots, \{x_4, x_0\}$. For each WARP violation, $z \Longrightarrow z_i$, and these are the only non-trivial relations in $\Longrightarrow$. Thus $\Longrightarrow^*$ and $\Longrightarrow$ coincide, and we obtain that z is the unique irrationality kernel for the data. Hence $i_{IR}(c) = 1$.

However, $i_{MR}(c) > 2$. Clearly c cannot be rationalized by a single preference; suppose for sake of contradiction that it can be rationalized by two, $\succ_1$ and $\succ_2$. Call an alternative x_i a *local maximum* for a preference $\succ$ if $x_i \succ x_{i-1}, x_{i+1}$. Every preference must have at least one local maxima; suppose without loss that x_1 is a local maxima of $\succ_1$. Then $\succ_1$ rationalizes the choice on $\{x_0, x_1, x_2\}$, but cannot rationalize the choice on $\{x_1, x_2, x_3\}$, hence $\succ_2$ must. Thus x_2 is a local maxima for $\succ_2$, but this means that $\succ_2$ cannot rationalize the choice on $\{x_2, x_3, x_4\}$. Thus $\succ_1$ must, and hence x_3 is a local maxima of $\succ_1$, and by analogous reasoning, x_4 is a local maxima of $\succ_2$. However, this means $\succ_2$ cannot rationalize the choice of x_0 from $\{x_4, x_0, x_1\}$, nor can $\succ_1$, as x_1 was a local maxima of $\succ_1$ by hypothesis. Hence $i_{MR}(c) > 2$.

[18] For example, consider a simple binary-menu environment, a correspondence c possessing a single forcing and single forced cycle, and c' containing only one cycle. In this case, both i_{MR} and i_{IR} rank both c and c' as equally irrational. This is because on such a domain, the fact both cycles for c are $\Longrightarrow^*$-related implies both share a common edge, hence i_{MR} can rationalize all choices for c by introducing a single extra preference, the same as for c'. In contrast, i_{IR} simply says there is only one relevant cycle for each correspondence.

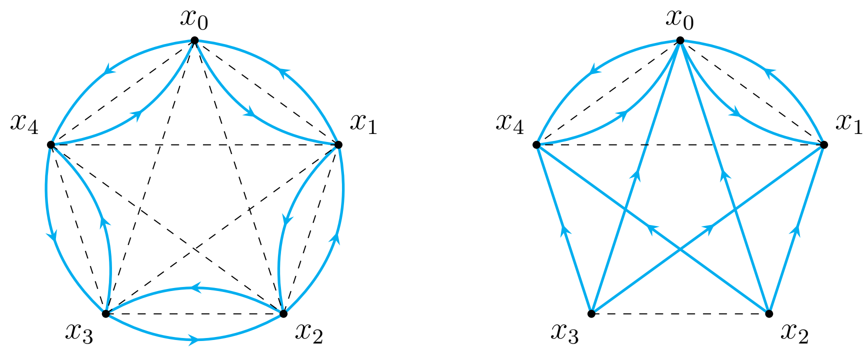

(a) The revealed preference pair for c. There are 6 cycles in $\mathcal{Z}_c$, five of which are WARP violations.

(b) The revealed preference pair for c'. The only cycles in $\mathcal{Z}_{c'}$ are the two WARP violations.

Figure 5: The revealed preference pairs for c and c' in Example 6. Since both c and c' are choice functions, all relations are strict.

Conversely, consider the choice function:

$$c'(B) = \begin{cases} B \setminus \{x_0\} & \text{if } B \text{ is doubleton} \\ \{x_0\} & \text{if } B \text{ is tripleton and contains } x_0 \\ B \setminus \{x_2, x_3\} & \text{otherwise.} \end{cases}$$

There are two choice cycles in c', the WARP violations supported on $\{x_4, x_0\}$ and $\{x_0, x_1\}$; see Figure 5. These are independent, hence $i_{IR}(c') = 2$. However, c' can be rationalized by a two preference family, e.g.:

$$x_0 \succ_1 x_4 \succ_1 x_1 \succ_1 x_3 \succ_1 x_2$$

and

$$x_4 \succ_2 x_1 \succ_2 x_0 \succ_2 x_3 \succ_2 x_2.$$

Thus i_{MR} and i_{IR} yield opposing conclusions about the relative inconsistency of c and c'. ■

5.2.5 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.[19] Formally, let X be finite and $\mathcal{L}_X$ denote the set of linear orders on X. Then for any choice function $c \in \mathcal{C}^f(X, \Sigma)$, the swaps index is defined as:

$$i_{SI}(c) = \min_{\rhd \in \mathcal{L}_X} \left\{ \sum_{B \in \Sigma} \left| \{x \in B : x \rhd c(B)\} \right| \right\}. \tag{3}$$

The swaps index seeks to find the (strict) preference that minimizes the empirical discrepancy with the data c, measured in terms of pairwise differences. In contrast, the inconsistency rank seeks to account for the internal structure of the non-rationalizable aspects of the data, and find a minimal justifying set of cycles. As the following example shows, these approaches do not always agree.

Example 7. Let $X = \{x_0, \ldots, x_5\}$ and Σ consist of the budgets:

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

Define the choice functions:

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

and

$$c'(B) = \begin{cases} \{x_{i+1}\} & \text{if } B = \{x_i, x_{i+1}\}, i \in \{1, 2, 3\} \\ \{x_i\} & \text{if } B = \{x_i, x_{i+1}\}, i \in \{0, 4, 5\} \\ \{x_0\} & \text{if } B = \{x_0, x_1, x_4\} \\ \{x_1\} & \text{if } B = \{x_1, x_3, x_4\}, \end{cases}$$

where i is understood mod-six.

Consider first the choice function c. There are six cycles in $\mathcal{Z}_c$: one of length six, two of length four, two of three, and one of length two, see Figure 6.

[19] Formally, Apesteguia and Ballester (2015) consider the problem of rationalizing a choice function by a linear order. This poses no particular difficulty to our framework.

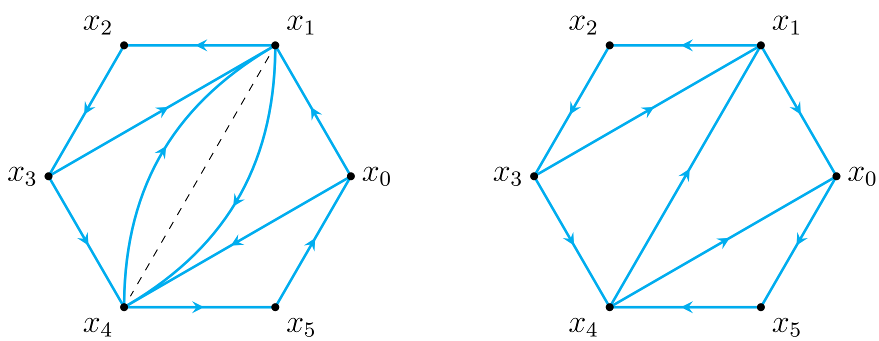

(a) The revealed preference for c. Here, there are 6 cycles, one of which is a WARP violation.

(b) The revealed preference for c'. Here, there are four cycles, none of which are WARP violations.

Figure 6: The revealed preference pairs for c and c' in Example 7. Since both c and c' are choice functions, all relations are strict.

However, the cycle of length six explains every other, hence $i_{IK}(c) = 1$. In contrast, there are four cycles in $\mathcal{Z}_{c'}$, supported on $\{x_1, x_2, x_3, x_4\}$, $\{x_0, x_1, x_4, x_5\}$, $\{x_1, x_2, x_3\}$, and $\{x_4, x_5, x_0\}$ respectively. The only non-trivial $\Longrightarrow$-relations are that the first of these cycles explains the the third, and the second explains the fourth. Thus $i_{IK}(c) = 2$, and hence the inconsistency rank deems c' as exhibiting a lesser degree of irrationality than c.

Conversely, the the linear order:

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

is a minimizer (3) for c.[20] Relative to this benchmark, c deviates by three swaps: two from choice on $\{x_0, x_1, x_4\}$ and one on $\{x_1, x_2\}$, hence $i_{SI}(c) = 3$. On the other hand, for c', the linear order:

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

[20]To see this, note that the revealed preference pair contains three cycles, supported on $\{x_1, x_2, x_3\}$, $\{x_1, x_4\}$, and $\{x_4, x_5, x_0\}$, which contain no edge in common. Any linear order must reverse at least one comparison in each of these cycles, hence $i_{SI}(c) \geq 3$; this bound is then attained by $\rhd$.

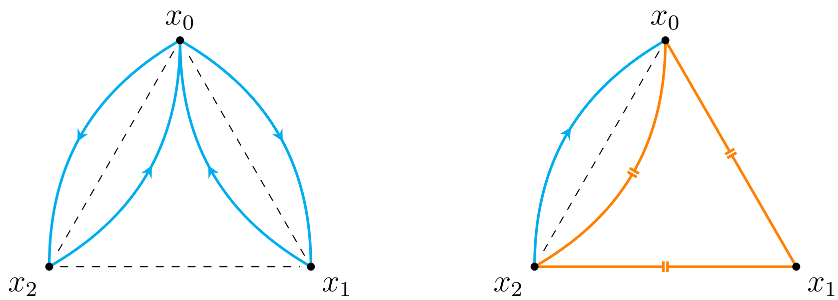

(a) The revealed preference pair for c. There are two choice cycles, both of which are WARP violations.

(b) The revealed preference pair for c'. There are two cycles, one of which is a WARP violation.

Figure 7: The revealed preference pairs for c and c' in Example 8.

is a minimizer.[21] Relative to this, c' requires only two swaps however: one $\{x_2, x_3\}$ and one on $\{x_4, x_5\}$, and hence $i_{SI}(c') = 2$, yielding a reversal. ■

5.2.6 Rationality Ordering

Ribeiro (2024) 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 8. Suppose $X = \{x_0, x_1, x_2\}$ and $\Sigma = \big\{\{x_0, x_1\}, \{x_0, x_2\}, X\big\}$. Define:

$$c(B) = \begin{cases} B \setminus \{x_0\} & \text{if } B \neq X \\ \{x_0\} & \text{if } B = X \end{cases}$$

and

$$c'(B) = \begin{cases} B & \text{if } B \neq \{x_0, x_2\} \\ \{x_0\} & \text{if } B = \{x_0, x_2\} \end{cases}$$

The correspondence c has two choice cycles, both of which are WARP violations, supported on $\{x_0, x_1\}$ and $\{x_0, x_2\}$. Denote these by z_1 and z_2 re-

[21] The revealed preference contains two cycles with disjoint generators (i.e. the cycles supported on $\{x_1, x_2, x_3\}$ and $\{x_4, x_5, x_0\}$) hence $i_{SI}(c') \geq 2$. This bound is attained by $\triangleright'$.

spectively. For each z_i, there is a unique generator consisting of $\{x_0, x_i\}$ and X. Since $\{x_0, x_{-i}\}$ does not belong to or cover this collection , z_1 and z_2 are independent and hence $i_{IK}(c) = 2$. Note, however, that the restriction of c to any sub-collection $\Sigma' \subseteq \Sigma$ is rationalizable if and only if Σ' does not contain both (i) one doubleton budget, and (ii) X itself.

Similarly, c' contains two cycles, one WARP violation supported on $\{x_0, x_2\}$ and one cycle of length three. However, the collection consisting of $\{x_0, x_2\}$ and X is a generator for both cycles, and hence $i_{IK}(c') = 1$. Despite this, the restriction of c' to some $\Sigma' \subseteq \Sigma$ is rationalizable if and only if Σ' does not contain both $\{x_0, x_2\}$ and X. Thus c' is not rationalizable whenever c is not, but not vice-versa, and we conclude c is more rational than c' in the rationality ordering. ■

6 Choice Environments with Special Structure

In this section, we provide descriptions two 'extremal' classes of choice environments. The first are domains on which no non-trivial dependencies can arise between cycles, for any choice correspondence; we show such environments are suitably degenerate, suggesting most experiments will feature the possibility of forced cycles emerging.

The second are domains where no loop in the budget graph is capable of supporting a cycle which does not force others. We show this richness property characterizes those domains on which satisfaction of the weak axiom is both necessary and sufficient for rationalizability.

6.1 Forcing-Free Design

A natural question is to what degree careful experimental design can ensure that there are no non-trivial forcing relations between cycles, no matter how choices are made. Such domains considerably simplify the problem of quantifying inconsistency. While environments with this property exist in trivial cases (e.g. when Σ is singleton) it is unclear how broad this class is.

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

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

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

Theorem 3. *Let (X, Σ) be a choice environment with finite budget graph $\Gamma(X, \Sigma)$. If Σ is forcing-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 3 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, on such domains, no budget of cardinality ≥ 3 can be involved in any choice cycle.

We interpret Theorem 3 as evidencing that most practical choice experiments will necessarily give rise to the possibility of non-trivial dependencies arising between cycles. In light of this, we expect the inconsistency rank to be broadly applicable and, in many cases, to yield different predictions than existing indices.

6.2 Well-Covered Environments

We now consider the opposite extreme: choice environments where no loop in the budget graph can support a non-forcing cycle. In light of Corollary 1, 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 (Arrow 1959; Sen 1971). This result is sometimes referred to as the 'fundamental theorem of revealed preference' (e.g. 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 to characterize which incomplete domains remain observationally rich enough for the weak axiom to remain characteristic.[22,23] This turns out to be intimately tied to our notion of forcing. Our next result shows that well-coveredness is precisely the minimal observability requirement needed for a 'strong' weak axiom.

Theorem 4. *Let (X, Σ) be an arbitrary choice environment. The following are equivalent:*

(i) The weak axiom of revealed preference is necessary and sufficient for strong rationalizability; and

(ii) (X, Σ) is well-covered.

Theorem 4 shows that, on any well-covered domain, testing rationalizability requires only checking the revealed preference for violations of the weak axiom. On such a domain, if a choice correspondence satisfies the weak axiom, any cycle of length greater than two always directly forces a strictly shorter cycle.[24] Since repeated application of this logic implies that any choice cycle entails a

[22]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).

[23]For a solution in the case of linear environments, see Cherchye et al. (2018).

[24]See Lemma 2 and Lemma 3 in the technical appendix.

WARP violation, we obtain that on well-covered domains, the satisfaction of the weak axiom is equivalent to the generalized.

This observation may be adapted straightforwardly to characterize those environments on which a correspondence is rationalizable if and only if it has no cycles of length less than or equal to some threshold K. Say (X, Σ) is **K-well covered** if, for every loop in $\Gamma(X, \Sigma)$ of length greater than K, every cyclic collection is covered.

Corollary 2. *Let (X, Σ) be an arbitrary choice environment. Then the following are equivalent:*

(i) For every $c \in \mathcal{C}(X, \Sigma)$, the correspondence c is rationalizable if and only if it exhibits no choice cycles of length $\leq K$; and

(ii) (X, Σ) is K-well covered.

7 Conclusions

Economic theory seeks to rationalize choice behavior by ascribing to individuals, a preference, whose maximization is consistent with their decisions. In practice, however, it is frequently the case that observed behavior is inconsistent with every possible preference. As such, it becomes crucial to be able to empirically assess the significance, or magnitude, of the observed violations.

We present a novel measure of irrationality for choice data, the inconsistency rank. Unlike existing measures, our index is based around the observation that the *structure* of a choice environment, or experiment, often leads to non-trivial dependencies between instances of irrational behavior. Our index quantifies the number of distinct instances of observed irrationality needed, in light of this structure, 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. We have shown that even in simple settings, choice cycles interact in subtle and non-trivial ways.[25] Understanding when and how these relations arise may additionally help inform the design of future experiments.

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.

Armstrong, Wallace Edwin, "A note on tue theory of consumer's behaviour," *Oxford Economic Papers*, 1950, *2* (1), 119–122.

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

Caradonna, Peter P, "A Least Squares Theory of Revealed Preferene," 2024.

Chambers, Christopher P, Federico Echenique, and Nicolas S Lambert, "Recovering preferences from finite data," *Econometrica*, 2021, *89* (4), 1633–1664.

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.

[25] See, e.g., Section 5.2.

Dziewulski, Paweł, "Just-noticeable difference as a behavioural foundation of the critical cost-efficiency index," *Journal of Economic Theory*, 2020, *188*, 105071.

_ , 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.

Luce, R Duncan, "Semiorders and a theory of utility discrimination," *Econometrica, Journal of the Econometric Society*, 1956, pp. 178–191.

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

Neumann, John Von and Oskar Morgenstern, "Theory of games and economic behavior," 1944.

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

Ribeiro, Mauricio, "Comparative rationality," *Mimeo*, 2024.

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.

Sen, Amartya K, "Choice functions and revealed preference," *The Review of Economic Studies*, 1971, *38* (3), 307–317.

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.

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

A Proof Appendix

A.1 Proof of Theorem 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 function $c' \in \mathcal{C}^f(X, \Sigma)$ which also contains z as a cycle, $\mathcal{G}_z$ 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 $N \geq 1$. Suppose first $N = 1$, i.e. z is a WARP violation. Then $\mathcal{G}_z$ consists of two budgets, B_1 and B_2, both of which contain $V_z = \{x_0, x_1\}$. Let $\trianglelefteq$ denote an arbitrary reverse well-ordering of $X \setminus \{x_0, x_1\}$.[26] Define:

$$c'(B) = \begin{cases} \{x_1\} & \text{if } B = B_1 \\ \{x_0\} & \text{if } x_0, x_1 \in B \text{ and } B \neq B_1 \\ B \cap \{x_0, x_1\} & \text{if } \big|B \cap \{x_0, x_1\}\big| = 1 \\ \max_{\trianglelefteq}(B) & \text{if } B \cap \{x_0, x_1\} = \varnothing. \end{cases}$$

As $\trianglelefteq$ is a reverse well-order, c' is well-defined and a choice function. By construction, z is a cycle for the restriction of c' to $\mathcal{G}_z$ and hence of the unrestricted choice function c' too. Moreover, it is the only cycle of c': to see this, note that (i) the only alternatives revealed preferred to x_0 (resp. x_1) are x_1 (resp. x_0), and (ii) the restriction of the revealed preference pair to $X \setminus \{x_0, x_1\}$ is rationalizable by $\trianglelefteq$. By (ii) any cycle $\tilde{z}$ must then include either x_0 or x_1, and by (i) it must be $\tilde{z} = z$.

Suppose now $N > 1$. We first define a choice function on $\Sigma|_{\mathcal{G}_z}$. Let $\trianglelefteq$

[26] That is, $\trianglelefteq$ is a linear order on $X \setminus \{x_0, x_1\}$ such that every non-empty subset has a (unique) $\trianglelefteq$-maximal element. Under the axiom of choice, such orderings always exist.

denote an arbitrary, reverse well-ordering of $\big(\cup_{\hat{B}\in\mathcal{G}_z} \hat{B}\big) \setminus V_z$. Now, define:

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

As $\mathcal{G}_z$ is a generator for z, every extension of $\tilde{c}'$ 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 $V_z = \{x_0, \dots, x_N\}$. Similarly, let $\trianglelefteq'$ denote a reverse well-ordering of $X \setminus \big(\cup_{\hat{B}\in\mathcal{G}_z} \hat{B}\big)$, and define an extension c' of $\tilde{c}'$ via:

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

As $\trianglelefteq$ and $\trianglelefteq'$ are reverse well-orders, c' is clearly well-defined and a choice function. Now, 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'.[27] First, note that for all $0 \le i \le M$, we must have $y_i \notin X \setminus \big(\cup_{\hat{B}\in\mathcal{G}_z} \hat{B}\big)$, as by construction, the only things that are c'-revealed weakly preferred to elements in this set are other elements in this set, and restricted to $X \setminus \big(\cup_{\hat{B}\in\mathcal{G}_z} \hat{B}\big)$, the revealed preference pair for c' is rationalizable by the linear order $\trianglelefteq'$. Thus for all i, we have $y_i \in \big(\cup_{\hat{B}\in\mathcal{G}_z} \hat{B}\big)$.

Suppose then that $y_i \in \big(\cup_{\hat{B}\in\mathcal{G}_z} \hat{B}\big) \setminus V_z$. By induction, every y_j must belong to $\big(\cup_{\hat{B}\in\mathcal{G}_z} \hat{B}\big) \setminus V_z$, as $y_i \succsim_{c'} y_{i+1}$ implies that $y_i, y_{i+1} \in B$ where $B \cap V_z = \varnothing$, and $y_i \trianglerighteq y_{i+1}$. But this is a contradiction, as $\trianglerighteq$ is also a linear order.

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, in fact, this cycle is precisely

[27] At least one such cycle must exist, as $z \in \mathcal{Z}_{c'}$ by construction.

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 $x^* \in \big(\cup_{\hat{B} \in \mathcal{G}_z} \hat{B}\big)$ and hence that $x^* \in B'$ for some $B' \in \mathcal{G}_z$. Since $c'(B') \cap V_z \neq \varnothing$ by minimality of the generator $\mathcal{G}_z$, 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 Corollary 1

Proof. Suppose γ has an uncovered cyclic collection, $\mathcal{B}_\gamma$. By the construction of c' in the proof (ii) implies (i) in Theorem 1, there exists a choice function which contains a single cycle, supported on γ, hence by contraposition, (i) implies (ii). Conversely, the other direction of the proof of Theorem 1 shows that given any set of choices generating a cycle on γ, no extension of these choices to Σ (indeed, to any covering budget) can avoid making another cycle. □

A.3 Proof of Theorem 2

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 $\hat{\mathcal{Z}} := \mathcal{Z}_c / \Longleftrightarrow^*$, i.e. the set of cycles of c modulo the equivalence relation $\Longleftrightarrow^*$, the symmetric component of the $\Longrightarrow^*$ relation.[28] 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) $\{[z_1], [z_2], \dots\} \subseteq \hat{\mathcal{Z}}$. 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 $\mathcal{I} \subseteq \mathcal{Z}_c$ by choosing one

[28]This is an equivalence relation. If $z_1 \Longleftrightarrow^* z_2$ and $z_2 \Longleftrightarrow^* z_3$, then $z_1 \Longleftrightarrow^* z_3$ by transitivity of $\Longrightarrow^*$.

cycle in $\mathcal{Z}_c$ 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 $\mathcal{I}, \mathcal{I}' \subseteq \mathcal{Z}_c$ 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.4 Proof of Theorem 3

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 Theorem 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^*.[29] 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 $\{b_0, \ldots, b_K\}$, define $E^o = \big\{\{b_0, b_1\}, \ldots, \{b_{K-1}, b_K\}\big\} \subseteq E_\Gamma$, and enumerate the path components of $\tilde{\gamma}$ as $\tilde{\gamma}_0, \ldots, \tilde{\gamma}_J$.[30] 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 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:

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

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

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

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 $B^{**} \in \mathcal{B}_\gamma \backslash \{B^*\}$ 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.5 Proof of Theorem 4

A.5.1 Notation

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

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 $(\succsim_c, \succ_c)\big|_e = \left(\succsim_c\big|_e, \succ_c\big|_e \right)$, where:

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

(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.5.2 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 $c \in \mathcal{W}(X,\Sigma)$. □

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 $x_{k'}$).[31] 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. □

[31] 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.5.3 Proof of Theorem 4

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.[32] Thus, letting (subscripts taken mod-J):

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

[32]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. □

A.6 Proof of Corollary 2

Proof. Suppose there exists some choice correspondence c that has no cycles of length less than or equal to K, but is not rationalizable. In particular, c obeys WARP hence the revealed preference pair reduces to a single relation $\succsim_c$. Since c is not rationalizable, there exists some loop γ in the budget graph of length $> K$ on which $\succsim_c |_{E_\gamma}$ is a cycle. Without loss, let γ be the shortest loop with this property. If every cyclic collection for γ is covered, by Lemma 3 there would exist a strictly shorter cycle, which cannot occur. Thus γ is uncovered, and hence (X, Σ) is not K-well covered. By contraposition, (ii) implies (i).

Conversely, suppose (X, Σ) is not K-well covered. It suffices to exhibit a choice correspondence with no cycles of length less than or equal to K, but which is nonetheless not rationalizable. In particular, it suffices to exhibit a choice correspondence with a cycle of length $K + 1$ or greater. Let γ be any loop in the budget graph which possesses an uncovered cyclic collection; by hypothesis at least one such γ exists, and has length $> K$. The construction for c in the second part of the proof of Theorem 4 supplies such a correspondence, hence, again by contraposition, we conclude (i) implies (ii). □