Essays in the Empirical Theory of Preferences

A Thesis
submitted to the Faculty of the
Graduate School of Arts and Sciences
of Georgetown University
in partial fulfillment of the requirements for the
degree of
Doctor of Philosophy
in Economics

By

Peter P. Caradonna, M.A.

Washington, DC
April 7, 2022

Essays in the Empirical Theory of Preferences

Peter P. Caradonna, M.A.

Thesis Advisor: Christopher P. Chambers, Ph.D.

Abstract

In this dissertation, I investigate various empirical aspects of the theory of preferences and decision. A classical result in revealed preference shows that when one observes a subject's choice on a rich enough collection of choice sets, the weak axiom of revealed preference is both necessary and sufficient for the choice data to be rationalized by a preference relation. In Chapter 2, I provide a complete characterization of how far these complete data assumptions may be relaxed, while still retaining a suitably powerful weak axiom. I then explore connections between these richness conditions and the classical literature on demand integrability. Relative to the existing literature which focused weakening the analytic regularity conditions under which the system of partial differential equations defining the integrability problem can be solved, in Chapter 3 prove a "nothing assumed" integrability theorem, that not only imposes no regularity conditions on model primitives, but also applies to arbitrary data sets, finite or infinite, in contrast with the traditional literature which requires an infinite set of observations. Finally, in Chapter 4, I develop a least squares regression theory for a novel form of choice data. I show that for a wide range of decision theoretic models, across a variety of domains, constructing statistical tests of consistency for these models may be reduced to a standard problem of testing multiple linear moment inequalities. Applications to trade, welfare, and the eliciting of subjective beliefs, are provided.

Index words: Microeconomics, Demand Theory, Decision Theory

Dedication

This paper is dedicated to my parents, without whose unwavering enthusiasm, encouragement, and support, I would not be where I am today.

ACKNOWLEDGMENTS

I would like to thank my stellar advisor Christopher Chambers, as well as my excellent committee members Roger Lagunoff and Nathan Miller for their guidance and encouragement over the years; you have each indelibly shaped the way I think about and do economics. It takes a village to raise a student, and I am immensely fortunate to have had the privilege to learn from each of you.

I would also like thank my fiancé Lauren Joseph for her kindness, patience, and bottomless support during my time as a doctoral student. It's not always easy to date a theorist. You have been there for me, through the highs and the lows, and gone out of your way to learn mathematical jargon along the way that no medical doctor should ever have to, all just to better understand whatever was exciting me. It has meant the world to me.

Finally, I would like to thank my stellar classmates, in particular Christopher Turanisck, JJ Naddeo, and Rachel Feldman, for always being willing to lend an ear and discuss my latest piece of progress or draft iterate. Your feedback has been of immense value to me, and I have learned endlessly from our discussions.

"It would be disingenuous of us to disguise the fact that the principal motive which prompted the work was the sheer fun of the thing."

-Alan Turing, on his contributions to game theory.

Table of Contents

List of Figures

Chapter 1

How Strong Is the Weak Axiom?

1.1 Introduction

The weak axiom of revealed preference, corresponding to the absence of pairwise reversals in observed choice behavior, is among the most elementary and normatively appealing consistency criteria. A particularly striking feature of the weak axiom is how dependent its implications are upon the structure of the domain of choice. When choice is observed on a complete collection of budgets, consistency with respect to the weak axiom is equivalent to rational behavior: the weak axiom completely characterizes the testable implications of rationality (see [11], [103]).[1] Conversely, when choice is observed only on an exceedingly sparse collection of budgets, the satisfaction of the weak axiom may become vacuous.

Complete domain hypotheses are commonplace in choice theory. In spite of this, the manner in which the structure of the domain of choice affects the implications of the weak axiom is generally very poorly understood. In the context of experiments, this implies a non-trivial interaction between the experimental design, that is the choice of which budgets to solicit subjects' choices from, and the interpretation of any potential inconsistency. For example, if the weak axiom of revealed preference is characteristic of rationality for a given experiment, then clearly no choice cycle of

[1] This is sometimes aptly referred to as the 'fundamental theorem of revealed preference.' See, for example, [90].

length three or more can occur in isolation: there must also be a choice reversal. For such experiments, the testable implications of the transitivity of preference are wholly subsumed by pairwise coherency of choices.

Reliance on such assumptions also limit our ability to test new models. It is common to characterize the testable implications of such theories under the assumption of a complete domain. Broadly speaking, the consequence of this assumption is that the empirical content of these models then tends to be characterized by an appropriate variant of the weak axiom (e.g. [81], [83], [43]) or at the least to rely heavily on the observation that on a complete domain, ‘all cycles imply two-cycles’ (e.g. [19]). Outside the realm of theory, however, full domain hypotheses are generally difficult to justify on either positive or normative grounds. [37] seek to understand what can be empirically tested under incomplete data; our work here may be seen as part of a dual approach of trying to better understand how robust such results are to the relaxation of these assumptions without fundamentally altering their testable implications. It seems likely that future results in this direction will require ideas formally extending those studied here in the ‘base case.’

We undertake the systematic study of how the power of the weak axiom varies with the richness of the collection of budgets choice is sampled on. In particular, we fully characterize those choice environments for which the weak axiom of revealed preference exhausts the testable implications of rational choice. We show that the class of environments includes not only complete collections of budgets, but also considerably smaller ones, and the property of having a strong weak axiom is not, in general, preserved under the addition of new budgets nor the restriction to sub-collections. We also consider the related problem, spiritually similar to the integrability theory of classical demand, of when the weak axiom, in conjunction with a 'local' no-cycles condition, characterizes rational choice. It turns out that in general, such a theories also require a suitably rich domain of budgets, though a weaker richness condition than is required for the weak axiom alone to suffice.

Example 1. Consider four alternatives $\{a, b, c, d\}$. 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, her choice behavior would be consistent with the weak axiom. This is because her choice behavior contains no preference *reversal*.[2] However, it would be inconsistent with preference maximization, as it would violate the generalized axiom: it contains a *cycle*.

Suppose now the agent were additionally presented with choices over $\{a, b, c\}$. The presence of the cycle from her other choices would necessarily force her to make another choice cycle over other alternatives: if she did not choose exclusively a as her most-preferred alternative, she would create a revealed preference reversal when this choice was considered alongside those preceding it. But, were she to choose exclusively a, then by revealing a to be preferable to c she would have chosen cyclically over a, d, and c. The structure of this collection of budgets ensures that any cycle of choices over all four alternatives necessarily induces other choice cycles in the data, though not necessarily a choice reversal.

Finally, suppose the agent is now presented with choices over the four binary budgets, $\{a, b, c\}$, and $\{c, d, a\}$. The presence of the cycle from her first four choices now necessarily forces her to make a preference reversal in her choices from the latter

[2]In fact, it would be impossible for the agent to *violate* the weak axiom as no budget in this environment contains a common pair of alternatives.

two budgets. If the agent were to choose anything but a from $\{a,b,c\}$, a reversal would be immediate. But then *any* choice from $\{a,c,d\}$ constitutes a choice reversal. Though this collection of budgets is far from complete, the budgets nevertheless intersect in such a manner as to force any revealed preference cycle to necessarily induce a concomitant revealed preference reversal. Were these choice sets selected by an experimenter to be presented to the individual, the experimental setup would preclude the existence of testable implications of preference transitivity beyond pairwise coherent choice. ∎

1.2 The Ex-Ante Power of the Weak Axiom

1.2.1 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** which we observe the agent choose from. We interpret the collection Σ as capturing the manner in which we are able to sample an agent's choices: we can observe an agent's choice on a set B if and only if it belongs to Σ. When Σ contains all non-empty, finite subsets of X, we will say that Σ is **complete**. We refer to the tuple (X,Σ) as a **choice environment**.

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). Given a choice correspondence $c \in \mathcal{C}(X,\Sigma)$, a preference relation $\succeq$ on X **strongly rationalizes** c if, for every budget we observe choice on, the chosen element(s) are precisely those $\succeq$-maximal alternatives:

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

Given a choice correspondence c, its revealed preference is a pair of relations $(\succsim_c, \succ_c)$ 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 satisfies the **weak axiom** of revealed preference if it contains no pairwise reversals: $x \succsim_c y$ implies $y \not\succ_c x$.[3] Notably, for choice correspondences satisfying the weak axiom, $\succ_c$ is indeed the asymmetric part of $\succsim_c$, allowing us to speak of a single revealed preference relation for such correspondences. We say c obeys the **generalized axiom** of revealed preference (sometimes referred to as 'congruence') if $(\succsim_c, \succ_c)$ contains no finite 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. We will denote the set of all choice correspondences for the environment (X, Σ) that satisfy the weak and generalized axioms, respectively, by $\mathcal{W}(X, \Sigma)$ and $\mathcal{G}(X, \Sigma)$. It was shown by [94], making use of an extension theorem due to [110], that a choice correspondence is strongly rationalizable by a preference relation if and only if it obeys the generalized axiom.[4] In light of this, we will interchangeably refer to the satisfaction of the generalized axiom as strong rationalizability.

[3] [82] provides a characterization of those choice correspondences that obey the weak axiom, for general environments, in terms of the the ability of the choices to be 'justified' by an asymmetric relation.

[4] We note, however, that [110] acknowledges the priority of Banach, Kuratowski, and Tarski in discovering, though not publishing, the result.

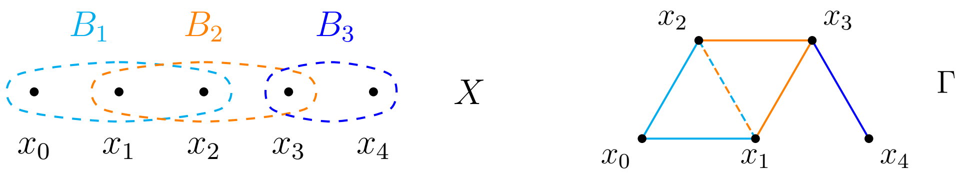

(a) A choice environment with five alternatives and three budget sets.

(b) The budget graph associated with this environment.

Figure 1.1: 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.

1.2.2 A Characterization

For purposes of combinatorial bookkeeping, it will be helpful to define an auxiliary structure that, for a given choice environment (X, Σ), encodes precisely which pairs of alternatives it is even possible for a preference to be revealed between. Let $\Gamma(X, \Sigma)$ be an undirected graph whose vertex set is X, and whose edge-set E_Γ is given by the (symmetric) relation of two vertices belonging to some common budget:

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

We term $\Gamma(X, \Sigma)$ the **budget graph**. Equivalently, the budget graph is the smallest undirected network with vertex set X for which the reflexive closure of the edge relation contains every revealed preference arising from a choice correspondence satisfying the weak axiom.[5]

For a given $c \in \mathcal{W}(X, \Sigma)$ and any $e \in E_\Gamma$ there is a well-defined (possibly empty) restriction of the revealed preference $\succsim_c$ to the edge $\succsim_c|_e$. This is because an edge

[5]The revealed preference arising from the 'complete indifference' choice correspondence, for example, obtains this bound.

$e = \{x, y\}$ is itself a two-element subset of the graph's vertex set, thus:

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

is well-defined. Similarly, given a collection of edges $E' \subseteq E_\Gamma$, we define:

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

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. Given a loop $\gamma \subseteq \Gamma(X, \Sigma)$, a collection of budgets $\mathcal{B}_\gamma \subseteq \Sigma$ is a **cyclic collection** for γ if, for every $e \in E_\gamma$ there exists a $B \in \mathcal{B}_\gamma$ with $e \subseteq B$. A cyclic collection for a loop γ is simply a collection of budgets for which there is some choice correspondence $\tilde{c} \in \mathcal{C}(X, \mathcal{B}_\gamma)$ that reveals a preference on every edge in the loop.[6] Our choice of terminology, however, betrays intent: we will be specifically interested in those collections that allow for cyclic choices around the loop and, in particular, those which admit extensions to all of Σ that obey the weak axiom.

Our ability use a particular cyclic collection to construct a choice correspondence that satisfies the weak axiom, but not the generalized, depends critically on how the collection intersects the remaining budgets in Σ. Given a loop γ and cyclic collection $\mathcal{B}_\gamma$, we say $\mathcal{B}_\gamma$ is **covered** if either:

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

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

where we define the restricted collection $\Sigma|_{\mathcal{B}_\gamma}$ via:

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

[6] Note that for every loop in $\Gamma(X, \Sigma)$, there exists at least one cyclic collection.

Note that condition (i) implies (ii) if and only if $|V_\gamma| > 3$. Practically speaking, covering budgets can be interpreted as choice sub-problems that are severely constrained by choices on a cyclic collection. If choices on some cyclic collection are constitute a GARP violation, it is easy to choose from budgets not contained within the cyclic collection without creating a WARP violation, by simply choosing from the (non-empty) subset of alternatives that do not lie within the cyclic collection. If a budget covers the cyclic collection, however, then the ability of a subject to make a pairwise consistent choice from the covering budget is constrained.

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, in essence, generalizes the classical argument that on a complete domain, every GARP violation implies a WARP violation (given a GARP violation, the complete domain forces the subject to choose from precisely the subset of alternatives making up the cycle). It requires instead only that the agent be forced to choose from some budget covering the collection on which the cycle is chosen. It turns out this is enough: the well-coveredness of Σ is both necessary and sufficient for the weak axiom of revealed preference to coincide with strong rationalizability for *any* choice correspondence.

Theorem 1. *Let (X, Σ) be a choice environment. The weak axiom of revealed preference is necessary and sufficient for strong rationalizability if and only if (X, Σ) is well-covered.*

Consider again the example from the introduction. In the case where the agent was presented with four budgets $\{a, b\}$, $\{b, c\}$, $\{c, d\}$, and $\{a, d\}$, the budget graph has a single loop, and the sole cyclic collection for this loop is uncovered. Thus this choice environment is not well-covered, and it is of course possible for the agent to choose cyclically in a manner violating the generalized axiom but consistent with the weak.

Now, consider the environment when the budget $\{a, b, c\}$ is added. This new budget serves to cover the loop of length four. However, it also adds two new loops of length three to the budget graph, formed by addition of the bisecting edge $\{a, c\}$. All of the cyclic collections for the loop with edges $\{a, b\}$, $\{b, c\}$, and $\{c, a\}$ are covered. However this is not true for the loop with edges $\{c, d\}$, $\{d, a\}$, $\{a, c\}$. Only by also adding yet another budget, $\{c, d, a\}$, is well-coveredness achieved. This last budget adds no new loops to the budget graph but, critically, serves to ensure that the cyclic collection for the loop $\{c, d\}$, $\{d, a\}$, $\{a, c\}$ becomes covered. It is this interlocking nature of the budget collections in the choice environment that well-coveredness characterizes.

1.2.3 Proof Sketch

The proof of the necessity of the well-coveredness of (X, Σ) for the weak axiom to coincide with the generalized proceeds by contraposition. We exhibit a means of constructing a choice correspondence, obeying the weak axiom but not the generalized, that relies only on the existence of a single loop with a single uncovered cyclic collection. The interested reader is referred to the Appendix. The proof of sufficiency is split over three lemmas. The first is a simple extension result which says, if we are given a loop γ and cyclic collection for it $\mathcal{B}_\gamma$, that if we can find a choice correspondence $\tilde{c}$ on the restricted domain $\Sigma|_{\mathcal{B}_\gamma}$ that cycles on γ and obeys the weak axiom, then there is no obstruction to extending $\tilde{c}$ to the full domain Σ.

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

The next lemma characterizes those minimal cycles that can arise from a choice correspondence that satisfies the weak axiom. It says that about any triangle in the budget graph, there is a choice correspondence that both (i) satisfies the weak axiom, and (ii) chooses cyclically around this triangle if and only if there exists an uncovered cyclic collection for the triangle.

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.*

Unfortunately, such a clean characterization does not obtain for longer loops. Lemma 3 however shows that, for loops of length four or more, if every cyclic collection for the loop is covered, then even if we cannot rule out the existence of a $c \in \mathcal{W}(X, \Sigma)$ that chooses cyclically around the loop, if such a c exists, it induces at least one other cycle elsewhere, around some strictly shorter loop.

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* $c \in \mathcal{W}(X, \Sigma)$ *where* $\succsim_c |_{E_\gamma}$ *contains 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* $\succsim_c |_{E_{\gamma'}}$ *contains a cycle.*

The 'sufficiency' direction of Theorem 1 then follows from a straightforward contraposition argument: suppose there exists some choice correspondence c which satisfies the weak, but not generalized, axiom. Then c contains some cycle of length three or more around some loop in the budget graph. If the loop was of length three, then by Lemma 2 the loop contains an uncovered cyclic collection and we conclude (X, Σ) is not well-covered. If the loop was of length four or greater and contains an uncovered cyclic collection, we again conclude (X, Σ) is uncovered, thus suppose that all of its cyclic collections are covered. Then by Lemma 3 there is a shorter cycle as well. Iterating this logic finitely many times, we obtain either a loop of length greater than three with an uncovered cyclic collection, or a cycle of length three, which by Lemma 2 implies an uncovered cyclic collection. In both these cases we conclude that (X, Σ) is not well-covered.

Examples of Well-covered Environments

Firstly, any complete environment is well-covered. Thus Theorem 1 extends the classical results of [11] and later [103].

Example 2 (Complete Abstract Environments)**.** Let X be a set, and suppose Σ contains all finite subsets of X. Then Σ is well-covered: letting γ be a loop, $V_\gamma \in \Sigma$. More generally, it is straightforward to show that if Σ either contains all cardinality two or all cardinality three budgets, it is well-covered.

More generally, if the budget collection is closed under finite unions, then it is well-covered. See, for example, Theorem 4 in [73].

Example 3 (Collections Closed Under Unions)**.** Let X be a set and suppose that, for all $B, B' \in \Sigma$, that $B \cup B' \in \Sigma$. Then Σ is well covered: for any loop γ, let $\mathcal{B}_\gamma$ denote an arbitrary cyclic collection. Since E_γ is finite, there exists a finite sub-collection of $\mathcal{B}_\gamma$ that is also a cyclic collection for γ. The union of this sub-collection is a budget by hypothesis, and covers $\mathcal{B}_\gamma$.

Another example of well-covered budget collections arise when there is some natural (weak) order on the space of alternatives, and budgets consist of intervals in this order. Such environments naturally arise when choice sets are defined simply by upper and lower bounds.

Example 4 (Interval Budgets)**.** Suppose $(X, \leq_X)$ is a weakly ordered set, and that Σ consists of order intervals, i.e. sets of the form $[x, y] = \{z \in X : x \leq_X z \leq_X y\}$, then it is well-covered. Letting γ be a loop in the budget graph, since V_γ is finite, it contains a $\leq_X$-minimal element, x_i. Without loss, suppose the adjacent vertices satisfy: $x_{i-1} \leq_X x_{i+1}$. Then, since budgets are intervals, every budget for the edge $\{x_i, x_{i+1}\}$ contains x_{i-1}, implying every cyclic collection for γ is necessarily covered and hence Σ is well-covered.

The argument showing any collection of interval budgets is well-covered relied critically on the 'intermediate value' property of order intervals. We may relax this requirement by substituting a suitable comparability criterion between budgets. Recall that if $(X, \leq_X)$ is a lattice, a subset B dominates a subset B' in the strong set order if, for all $x \in B$ and $x' \in B'$, $x \vee x' \in B$ and $x \wedge x' \in B'$.

Example 5 (Comparability of Budgets)**.** Suppose $(X, \leq_X)$ is a lattice, and that Σ consists of totally ordered subsets. If every pair of budgets in Σ is comparable in the strong set order, then Σ is well-covered. For a formal proof, see Appendix I.

Finally, it is easy to construct new well-covered collections from existing ones. In particular, well-coveredness is preserved by the taking of certain restrictions.

Example 6 (Restrictions of Well-covered Environments)**.** Suppose (X, Σ) is well-covered, and $A \subseteq X$. Then $(X, \Sigma|_A)$ is well-covered, where $\Sigma|_A = \{B \in \Sigma : B \subseteq A\}$. This follows straightforwardly by observing that, if $\Sigma|_A$ were not well-covered, then its uncovered cyclic collections could not become covered by passing to Σ, as all the added budgets must contain alternatives that do not belong to A.

1.2.4 Relation to the Classical Demand Framework

The question of when the weak axiom of demand theory is empirically distinguishable from the strong axiom in the classical demand framework has a long history. [97]

first proved that these axioms coincide in the case of two goods, though [50] soon after established that this result did not hold for the case of three or more goods. In a recent contribution, [34] characterized those linear budget collections for which several variants of the demand-theoretic weak and strong axioms coincide. Interestingly, they find that many widely used price-consumption datasets have large subsets exhibiting insufficient price variation to independently distinguish these axioms.[7] Given the apparent empirical shortcomings of field data for purposes of independently testing these phenomena, one is naturally led to consider how to construct simple, finite, laboratory experiments capable of rectifying this deficiency. Our Theorem 1 then provides a complete characterization of precisely which abstract choice experiments have testable implications of the generalized axiom in excess of the weak. Moreover, it is empirically and computationally desirable then to understand the problem for those environments in which one must take seriously indivisibilities, price nonlinearities, or other economic phenomena contrary to the linear budget paradigm, which our model speaks to.

While Theorem 1 holds equally well when $X = \mathbb{R}^n_+$ and elements of Σ are linear budgets, our results neither imply nor are implied by those of [34]. We assume no intrinsic order structure on the set of alternatives, thus make no requirement of a rationalizing preference being monotone. Particularly, we allow for choice correspondences that do not satisfy Walras' law. As such, our paper holds for more general data sets where the chosen commodity bundle does not lie on the budget frontier, but as a consequence we require a purely choice-based notion of revealed preference rather

[7] They find that roughly 70% of the Spanish survey ECPF (Encuesta Continua de Presupestos Familiares) panel dataset (see, for example [15]) satisfies their condition for when a WARP-based analysis is equally informative as SARP-based. Even more drastically, roughly 97% of price triples in the British FES (Family Expenditure Survey) cross-sectional data set (see, for example, [23], [21], [22]) satisfy their condition for WARP and SARP to coincide.

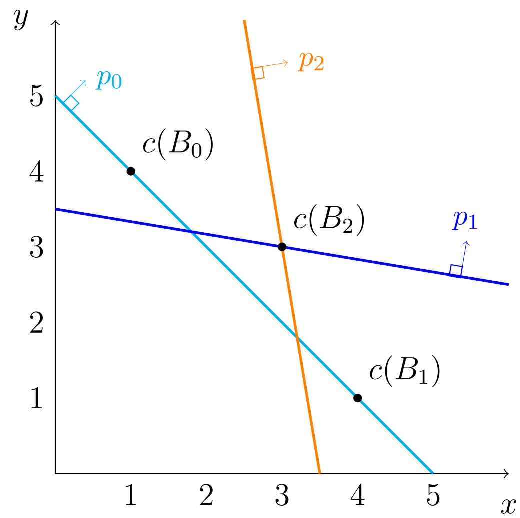

Figure 1.2: The frontiers of three linear budgets on $\mathbb{R}^2_+$. While the demand-theoretic weak and strong axioms coincide for any collection of linear budgets for two commodities, the above choices satisfy the choice-theoretic weak but not generalized axiom.

than one that makes use of the order structure of $\mathbb{R}^n$.[8] Additionally, we consider the solution concept of strong rationalizability, under which we require that the observed choices constitute the *entirety* of the agent's optimal choices from a given budget (classical references include [100], [70], [11]). This leads to a different notion of which 'cycles' constitute violations of our notion of rationalizability, and hence to differing characterizations of which environments lead to such cycles inducing reversals, even when the class of budgets considered is the same (see [84], [87]). In light of this, our results and those of [34] are best thought of as complementary, addressing different frameworks and valid in differing contexts.

[8] A consequence of this, however, is that our theory requires, as part of our definition of an observation, a complete description of the budget from which an agent chose. Walras' Law, on the other hand, provides an identifying assumption to pin down the budget set for an observation from only price and consumption data.

We conclude this section with an example of a collection of linear budgets in the two-commodity case (and hence for which the classical demand variants of the weak and strong axioms coincide) but which is not well-covered.

Example 7 (Non Well-covered Linear Budget Collection on $\mathbb{R}^2_+$)**.** Let $X = \mathbb{R}^2_+$, and consider three price tuples $p_0 = (\frac{1}{5}, \frac{1}{5})$, $p_1 = (\frac{1}{21}, \frac{1}{3.5}))$, and $p_2 = (\frac{1}{3.5}, \frac{1}{21})$. Let Σ consist of the three wealth-normalized linear budget sets formed by these price vectors: $B_i = B(p_i, 1)$ (see Fig. 1.2). Suppose an agent were to choose $c(B_0) = \{(1,4)\}$, $c(B_1) = \{(4,1)\}$, and $c(B_2) = \{(3,3)\}$. These are all alternatives belonging to each budget and, with the exception of the choice from B_1 all lie on the budget frontier (recall that we do not impose Walras' law). Moreover, c satisfies the weak axiom but exhibits a three-cycle:

$$(3,3) \succ_c (1,4) \succ_c (4,1) \succ_c (3,3),$$

and hence the collection cannot be well-covered. This stands in comparison to Rose's result that the classical demand version of the weak axiom coincides with (weak) rationalizability in the two-commodity case, no matter the budget collection.

Chapter 2

Abstract Choice and Integrability

2.1 Preliminaries

Well-coveredness of the budget collection is, in general, difficult to verify in practice, as it requires checking every cyclic collection for covering budgets. Without extra structure on the problem, this may become computationally difficult for larger experiments. Motivated by this difficulty, in this section we consider instead only those implications of well-coveredness that are reflected in the structure of the budget graph. If a budget collection is well-covered, clearly every loop in the budget graph of length four or more must possess a bisecting edge, that is an edge connecting vertices of the loop that does not belong to the loop's edge set. A graph with this property is said to be **chordal**. Critically, this property is efficiently verifiable: it is possible to determine whether a graph is chordal in linear time using standard methods (see, for example [96]).

In this section, we consider experiments with only a chordal budget graph, a necessary, though not sufficient, condition for the well-coveredness of the collection. We show that an experiment possesses a chordal budget graph if and only if strong rationalizability coincides with (i) the weak axiom, and (ii) a mild, discrete analogue of the Slutsky symmetry axiom of differential demand theory. This serves as a trade-off relative to Theorem 1: in exchange for requiring somewhat more structure than just the weak axiom on the part of the choice data, one obtains an efficiently verifiable minimal richness condition on the choice environment for no 'small' cycles to imply no cycles of any kind.

Such results appear also in the mechanism design literature, where it is of great interest to have criteria on type spaces that guarantee testing the global condition of cyclic monotonicity (a cardinal form of the generalized axiom) reduces to testing only pairwise comparisons (e.g. [99], [12], analogous to our Theorem 1) or pairwise comparisons plus only 'local' no cycle conditions (e.g. [10] and [76], which are analogous to our results in this section).

This result may be interpreted as an extension of integrability theory to the abstract choice framework. In particular, on rich enough domains our theory allows for incompletely observed, in particular potentially finite, data (corresponding to cases when Σ is far from complete) as opposed to the classical theory which takes as primitive a fully observed demand function. However, the relaxation to incomplete data can only go so far: our results also establish the chordality of the budget graph as the weakest possible richness condition on an environment under which strong rationalizability is equivalent to the classical integrability criteria of the weak axiom, plus a 'local no cycles' condition.

2.2 Abstract Analogues of the Integrability Conditions

Let (X, Σ) be a fixed choice environment, with $\Gamma(X, \Sigma)$ its budget graph. Let:

$$T_\Gamma = \big\{\{x, y, z\} \subseteq X : \{x, y\}, \{y, z\}, \{x, z\} \in E_\Gamma\big\}.$$

The combinatorial **domain** associated to the environment (X, Σ) is the triple $\mathcal{D}(X, \Sigma) = (X, E_\Gamma, T_\Gamma)$. The combinatorial domain essentially serves as a 'triangulation' of the set of alternatives using only the information encoded in the budget graph.

For a choice correspondence $c \in \mathcal{C}(X, \Sigma)$ with revealed preference pair $(\succsim_c, \succ_c)$, we say that c is **locally rationalizable** if we may extend $(\succsim_c, \succ_c)$ by a single relation $\succeq$ such that:[1]

$$(\forall \tau \in T_\Gamma) \quad \succeq|_\tau \text{ is complete and transitive.}$$

Local rationalizability is the ordinal analogue of the joint conditions of Slutsky negative semi-definiteness and symmetry. It says nothing more than we may strongly rationalize the revealed preference locally, about each triangle in the domain, much the same way as the usual properties on the Slutsky matrix guarantee an economically suitable local solution to the system of differential equations defining the integrability problem.

Local rationalizability is necessary, though not sufficient, for the strong rationalizability of c (as any strongly rationalizing preference relation is a local rationalization). The fact that one must potentially consider an extension of $\succsim_c$ is simply a consequence of allowing for the possibility that Σ is highly incomplete and c does not reveal any preference between some pairs in some triangles of T_Γ.[2]

2.2.1 A Decomposition of Local Rationalizability

We may decompose the property of local ratinonalizability into two properties: the weak axiom, and a property we term ordinal irrotationality, which serves as the analogue of Slutsky symmetry for the abstract choice model. For a choice correspondence $c \in \mathcal{C}(X, \Sigma)$ with revealed preference pair $(\succsim_c, \succ_c)$, we say that c is **ordinally irrotational** if the revealed preference pair $(\succsim_c, \succ_c)$ admits an order pair extension $(\succeq, \succ^*)$

[1]Formally, we mean that $(\succeq, \succ)$ is an order pair extension of $(\succsim_c, \succ_c)$, where $\succ$ is the asymmetric component of $\succeq$.

[2]If, for example, $E_\Gamma \subseteq \Sigma$, there would be no need to consider an extension. In particular, on a complete domain, no extension is ever necessary.

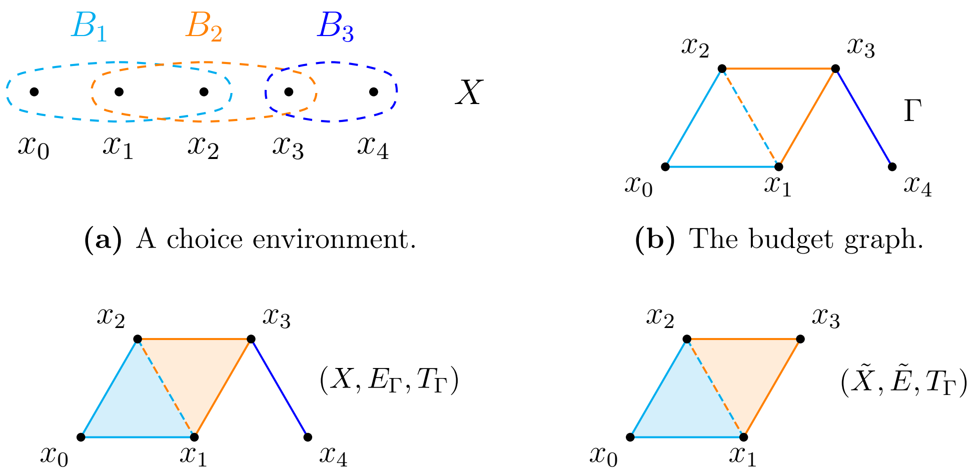

(a) A choice environment.

(b) The budget graph.

(c) The domain.

(d) The subdomain generated by $\tilde{T} = T_\Gamma$.

Figure 2.1: The various constructions associated with a choice environment. The shading of the triangles is intended to indicate their inclusion in T_Γ.

of its revealed preference such that:

$$(\forall \tau \in T_\Gamma) \quad (\succeq, \succ^*)|_\tau \text{ is complete and has no three-cycle,}$$

and for which $\succ^*$ is asymmetric and contains the asymmetric component of $\succeq$. Much the same as Slutsky symmetry, ordinal irrotationality is a 'no local cycles' condition, plus some minor regularity. Moreover, in conjunction with the weak axiom, it characterizes local rationalizability.

Proposition 1. *Let $c \in \mathcal{C}(X, \Sigma)$. Then c is locally rationalizable if and only if c both:*

(i) obeys the weak axiom; and

(ii) is ordinally irrotational.

The weak axiom and ordinal irrotationally are also logically independent, as the next example shows.

Example 8 (Independence of WARP and OI)**.** Let $X = \{x_0, x_1, x_2, x_3\}$, and suppose Σ consists of two budgets, $B_1 = \{x_0, x_1, x_2\}$, and $B_2 = \{x_1, x_2, x_3\}$. Let $c(B_1) = \{x_1, x_2\}$ and $c(B_2) = \{x_2\}$. Then c does not satisfy the weak axiom: $x_1 \succsim_c x_2$ but $x_2 \succ_c x_1$. However, c is ordinally irrotational: let $\succeq = \succsim_c \cup \{(x_1, x_3)\}$ and $\succ^* = \succ_c \cup \{(x_1, x_3)\}$. Then $\succ \subsetneq \succ^*$, $\succ^*$ remains asymmetric, and the restriction of $(\succeq, \succ^*)$ to each triangle in the domain (here, given by the two budgets) contains no three-cycle.[3]

2.3 Sampling and Integrability

The satisfaction of the generalized axiom of revealed preference by a choice correspondence implies its strong rationalizability and hence both the weak axiom and ordinal irrotationality, no matter the structure of the domain. However, the sufficiency of the weak axiom and ordinal irrotationality for the generalized axiom depends crucially on the structure of the domain. Intuitively, the denser the budget graph (that is, the greater the number of pairs of alternatives the experiment is capable of revealing a preference between), then the more triangles there will be and hence the more stringent the requirement of local rationalizability will be.

This is not the whole story, however. What turns out to be most important is the manner in which the triangles of the budget graph fit together. Certain collections of triangles may fit together in ways that permit even cyclic revealed preferences to be locally rationalizable (e.g. Figure 2.2). What is needed is that there be enough 'good' collections of triangles in the budget graph to cover every loop, allowing the condition of local rationalizability to rule out any possible cycles. This turns out to be possible, precisely when the budget graph of the environment is chordal. For any

[3]This establishes that OI does not imply WARP. To rule out the converse implication, it suffices to consider a three-element set of alternatives with three binary budget sets and any cyclic revealed preference.

environment with a chordal budget graph, if a choice correspondence (i) obeys the weak axiom, and (ii) is ordinally irrotational, then it is also strongly rationalizable. Moreover, possessing a chordal budget graph, it turns out, is the weakest possible richness condition on Σ under which *any* such traditional integrability result can possibly hold: for any environment without a chordal budget graph there always exist choice correspondences which obey the weak axiom and are ordinally irrotational (and hence locally rationalizable), yet nonetheless are not strongly rationalizable.

Theorem 2. *Let (X, Σ) be a choice environment with $\Gamma(X, \Sigma)$ chordal. Then a choice correspondence $c \in \mathcal{C}(X, \Sigma)$ is strongly rationalizable if and only if:*

(i) It obeys the weak axiom; and

(ii) It is ordinally irrotational.

Moreover, (i) and (ii) are jointly equivalent to the strong rationalizability of c if and only if $\Gamma(X, \Sigma)$ is chordal.

As local rationalizability is always necessary for strong rationalizability, regardless of the structure of the domain, we obtain the following corollary.

Corollary 1. *Let (X, Σ) be a choice environment, and suppose $\Gamma(X, \Sigma)$ is not chordal. Then there exists a choice correspondence which is locally, but not strongly, rationalizable.*

We emphasize that $\Gamma(X, \Sigma)$ being chordal is a requirement of sufficient ex-ante richness of the choice environment. A given pair of alternatives is connected by an edge in the budget graph if and only if there exists a choice correspondence on that environment capable of revealing a preference between them. Chordality specifically requires that, given any GARP violation in the data, that the experiment is rich enough that it is capable of revealing a preference between some non-adjacent pair of alternatives belonging to the cycle.[4]

In comparison with the classical integrability theory, we critically do not suppose a complete domain, or even an infinite data set. We allow for arbitrary environments, and our results fully characterize just what is needed observationally for the usual integrability conditions of the weak axiom and ordinal irrotationality to extend. In particular, Theorem 2 imposes no analytic, point-set, or order-theoretic assumptions on model primitives. This is particularly notable given the historical program of attempting to weaken the differentiability hypotheses of the classical integrability theory, e.g. [17], [60], and [18]. Nonetheless, there is a cost to this generality. In the classical theory, while one supposes a great deal more structure on the budgeter or demand function and its domain, one obtains a rationalizing utility with commensurately fine properties. Comparatively, all that is guaranteed by Theorem 2 is a rationalizing weak order. This is the price, it appears, of an integrability theorem that not only allows for finite data sets, but also imposes no hypotheses that are non-falsifiable by such data sets. Put another way, Theorem 2 operates wholly within the empirical content of abstract choice model, in the sense of [30].

[4]This contrasts with well-coveredness, which corresponds to the condition that, for any GARP violation, the experiment is rich enough to *guarantee* either a WARP violation, or a that some preference is revealed between a non-adjacent pair in the cycle (see Lemmas 2 and 3).

2.4 Proof Sketch

2.4.1 Simple Domains

While conceptually simple to state, the property of an environment having a chordal budget graph is a difficult global property to work with for purposes of proving Theorem 2. We first establish an equivalent characterization of chordality, in terms of collections of triangles in the associated domain, that is more suited to our purposes. Let $\tilde{T} \subseteq T_\Gamma$ be a collection of triangles in the budget graph. Let $\tilde{X}$ denote the points of X contained within some element of $\tilde{T}$, and $\tilde{E}$ the subset of edges in E_Γ contained in some element of $\tilde{T}$. Then the **subdomain** generated by $\tilde{T}$ is defined as the tuple:

$$\mathcal{D}(X, \Sigma)|_{\tilde{T}} = (\tilde{X}, \tilde{E}, \tilde{T}).$$

We will be particularly interested in subdomains possessing specific structure. For a finite $\tilde{T}$, we say the subdomain generated by $\tilde{T}$ is **combinatorially trivial** if, for every pair $\tau, \tau' \in \tilde{T}$, there is a unique sequence of of distinct elements of $\tilde{T}$:

$$\tau = \tau_1, \tau_2, \ldots, \tau_k = \tau'$$

where, for each $1 \leq j \leq k-1$, the triangles τ_j and τ_{j+1} share precisely a pair of common elements. If one imagines an undirected graph whose nodes are the elements of $\tilde{T}$ and whose edge relation is given by sharing a pair of common elements, then combinatorial triviality amounts to asking the graph associated with the collection $\tilde{T}$ be a tree.

Similarly, for a finite collection $\tilde{T}$ we say that the subdomain generated by $\tilde{T}$ is **topologically trivial** if it has no 'holes' in it in an appropriate sense.[5] If one imagines

[5] Formally, we say a subdomain $\mathcal{D}|_{\tilde{T}}$ is topologically trivial if it has a first Betti number of zero, i.e. $H_1(\mathcal{D}|_{\tilde{T}}; \mathbb{R}) = 0$, where H_* denotes simplicial homology with real coefficients. See [86] p. 34.

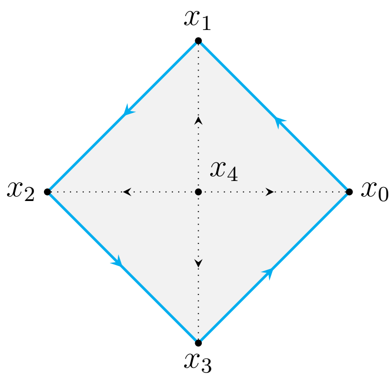

(a) A cyclic revealed preference (blue) with a local rationalization (black) on a topologically trivial subdomain that is not combinatorially trivial.

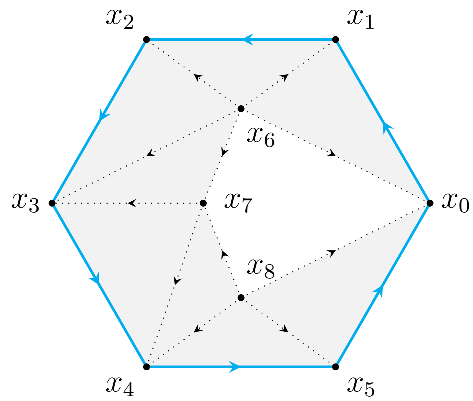

(b) A cyclic revealed preference (blue) and local rationalization (black) on a combinatorially trivial subdomain. The domain however retracts onto its outer boundary loop, hence it is topologically non-trivial.

Figure 2.2: Non-simple subdomains support locally rational cycles. When subdomains fail to be either combinatorially trivial or topologically trivial, the criterion of local rationalizability is too weak to guarantee the absence of cycles.

a subdomain as consisting of a graph whose triangles are 'filled in' forming a kind of triangulated surface, topological triviality roughly asks that this surface be simply connected or contractible.

We will say that a subdomain is **simple** if it is both combinatorially and topologically trivial. A loop $\gamma = (V_\gamma, E_\gamma)$ in the budget graph is contained in a subdomain $(\tilde{X}, \tilde{E}, \tilde{T})$ if $V_\gamma \subseteq \tilde{X}$ and $E_\gamma \subseteq \tilde{E}$. By abuse of notation, we will say the entire domain is simple if every loop in it is contained in a simple subdomain. This turns out to be equivalent to requiring that every loop in the budget graph possessing a bisecting edge.

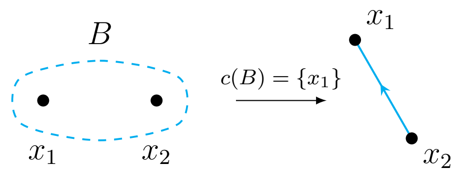

Figure 2.3: Choice correspondences and ordinal flows. Given a choice correspondence, we may view the relations of its revealed preference as specifying ordinal 'flows' along edges of the budget graph. The budget graph is the smallest network such that this interpretation remains valid for any choice correspondence.

Theorem 3. *Let* (X, Σ) *be a choice environment. Then the domain* $\mathcal{D}(X, \Sigma)$ *is simple if and only if the budget graph is chordal.*

2.4.2 Discrete Calculus

We now argue that, on a simple domain, every choice correspondence satisfying (i) the weak axiom, and (ii) ordinal irrotationality is strongly rationalizable. To do this, we will recast our ordinal problem into a cardinal form that shares strong formal similarities with the classical differential approach. In particular, we make use of the discrete exterior calculus, which is most suitable for our combinatorial structure.

A vector field, or 1-form, on the domain $\mathcal{D} = (X, E_\Gamma, T_\Gamma)$ is a map $F : \hat{E}_\Gamma \to \mathbb{R}$ such that $F(x, y) = -F(y, x)$, where $\hat{E}_\Gamma = \{(x, y) \in X \times X : \{x, y\} \in E_\Gamma\}$.[6] We interpret such a map as describing a magnitude of flow from x to y, and where a negative flow is simply interpreted as a flow in the opposite direction. Similarly, a 0-form on $\mathcal{D}$ is simply an element of $\mathbb{R}^X$, and a 2-form a map $\mathfrak{F} : \hat{T}_\Gamma \to \mathbb{R}$ such that $\mathfrak{F}(x_{\sigma(0)}, x_{\sigma(1)}, x_{\sigma(2)}) = \text{sign}(\sigma) \cdot \mathfrak{F}(x_0, x_1, x_2)$ for any permutation σ.

[6] We let $\hat{E}_\Gamma$ (resp. $\hat{T}_\Gamma$) denote the spaces of *oriented* edges (resp. triangles) of the budget graph.

There are natural operators between the spaces of forms on $\mathcal{D}$. For any 0-form f, we define the gradient of f to be the 1-form defined by:

$$\text{grad}(f)(x,y) = f(y) - f(x).$$

Similarly, for a 1-form F, its rotation, or curl, is defined pointwise as the 2-form:

$$\text{rot}(F)(x,y,z) = F(x,y) + F(y,z) + F(z,x).$$

It is straightforward to verify that both the gradient and curl are linear operators, and that the image of the gradient operator is vector subspace of the kernel of the curl operator. In particular, we term a 1-form F integrable if it belongs to the image of the gradient operator and irrotational if it belongs to the kernel of the rotation. One may think of the curl operator as measuring how far from (cardinally) transitive a given 1-form is about each triangle in the budget graph.

We will be interested in the existence of an integrable vector field that is consistent with the revealed preference pair (e.g. Figure 2.3). Let $(\succsim_c, \succ_c)$ denote the revealed preference of a choice correspondence c. We say that a 1-form F is a **cardinalization** of c if:

$$y \succsim_c x \implies F(x,y) \geq 0,$$

and

$$y \succ_c x \implies F(x,y) > 0.$$

For any choice correspondence c, let K_c denote the set of all 1-forms on $\mathcal{D}$ cardinalizing $(\succsim_c, \succ_c)$. For any c, K_c is a convex cone. Moreover, K_c is non-empty if and only if c obeys the weak axiom.[7]

[7] Cardinalizations as we have defiend them are intimately related to the idea of preference functions. See, for example [104], [93], or [4]. Unlike previous work, however, we do not assume that the cardinalization is defined for all pairs in X, but rather only those in $\hat{E}_\Gamma \subseteq X \times X$.

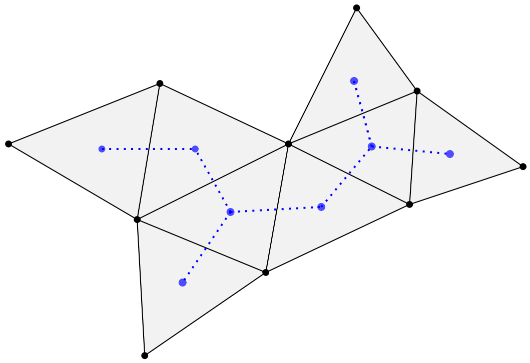

(a) A collection of triangles that is combinatorially trivial. The collection is combinatorially trivial becauses the 'sharing an edge' graph (dotted blue) is a tree.

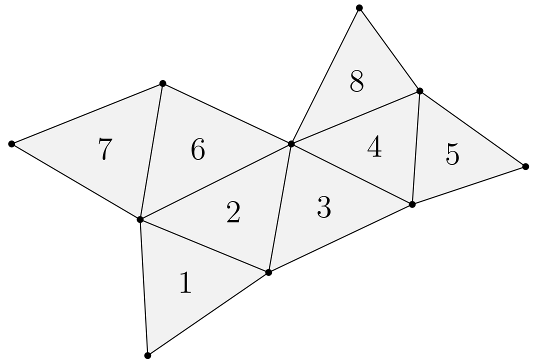

(b) A good ordering of the triangles, guaranteeing an irrotational cardinalization for any locally rational binary relation.

Figure 2.4: Good orderings and combinatorial triviality. On combinatorially trivial subdomains, one may always enumerate the triangles such that any triangle shares at most a single edge with the union of those triangles preceding it. The existence of such an enumeration guarantees a simple algorithm can construct irrotational cardinalizations for any locally rational binary relation.

2.4.3 Proof Sketch

Consider a subdomain $(\tilde{X}, \tilde{E}, \tilde{T})$ generated by some arbitrary finite collection $\tilde{T}$ of triangles, and consider some (for sake of exposition) asymmetric, locally rational $\succ_c$. To define a cardinalization F on this domain, it suffices to choose the values of F on those ordered pairs (x, y) where $x \succ_c y$, as this uniquely determines the value of F on all ordered pairs of the form (y, x). By this implicit choice of basis, we are identifying the cone K_c with the interior of the positive orthant of the space of 1-forms. We first consider the problem of whether there exists an irrotational cardinalization, a necessary though generally not sufficient condition for the existence of an integrable cardinalization:[8]

$$\operatorname{rot} F = 0$$
$$F \gg 0.$$

Combinatorial triviality ensures a solution to this problem for any locally rational binary relation. Consider the subdomain of Fig. 2.4. Call an enumeration of the triangles in the subdomain a 'good ordering' if it has the property that for any triangle in the enumeration, its intersection with the collection of all preceding triangles consists of at most a single edge. The combinatorial triviality of a subdomain guarantees good orderings exist. Such orderings provide a roadmap to constructing an irrotational cardinalization: restricting any one triangle, an irrotational cardinalization exists, simply by local rationality of $\succ_c$. Consider now any adjacent triangle. Again by local rationality we can find an irrotational cardinalization for just this triangle. However, we can form an irrotational cardinalization on the subdomain generated by both triangles together by ensuring the values of the flows on each triangle agree on the common edge the triangles share. This can always be attained by choosing irrotational but otherwise arbitrary flows on each triangle then multiplying the flows on one of the two by an appropriate positive scalar. More generally, in any good ordering no triangle shares more than a single edge with the collection of all preceding triangles, and thus an inductive application of this argument guarantees an irrotational cardinalization for any locally rational relation on a combinatorially trivial subdomain.[9]

[8]We employ the vector notation $x \geq y$ to denote that x is component-wise larger than y. We will write $x > y$ to denote that x is component-wise greater with at least one component strictly so, and $x \gg y$ to denote that every component of x strictly exceeds that of y.

[9]Without combinatorial triviality, irrotational cardinalizations are not guaranteeed exist, even for locally rational binary relations; see, for example, Figure 2.2.(a). The locally rational relation on Figure 2.2.(b) does admit an irrotational cardinalization (the subdomain is combinatorially trivial), albeit not an integrable one.

Thus we are guaranteed that for any locally rational $\succ_c$, on any combinatorially trivial subdomain, the cone of consistent cardinalizations K_c intersects the subspace of irrotational cardinalizations ker(rot). Indeed as K_c is the interior of the positive orthant of the space of flows, the intersection $\tilde{K}_c = K_c \cap \ker(\text{rot})$ will be of full dimension in ker(rot), and we may choose a basis for this subspace identifying $\tilde{K}_c$ with the interior of its positive orthant. Since the image of the gradient is a subspace of the kernel of the curl, we may view the gradient as a linear map taking a 0 form to an *irrotational* 1-form. Thus, given our choice of basis for ker(rot), the existence of an integrable cardinalization is equivalent existence of a $u \in \mathbb{R}^{\tilde{X}}$ such that:

$$\operatorname{grad} u \gg 0.$$

By Gordan's Alternative (see [54]), exactly one of the following holds: (i) the above problem admits a solution, or (ii) there exists a 1-form F such that:

$$\begin{aligned} \operatorname{grad}^\top F &= 0 \\ F &> 0 \\ \operatorname{rot} F &= 0. \end{aligned}$$

In particular, if no integrable cardinalization of $\succ_c$ exists, there is some non-zero $F \in \ker(\text{rot}) \cap \ker(\text{grad}^\top)$. However, this possibility is precisely ruled out by topological triviality, which ensures that this intersection is $\{0\}$. For any two matrices $A \in \mathbb{R}^{l\times m}$ and $B \in \mathbb{R}^{m\times n}$ such that $AB = 0$, there exist isomorphisms:

$$\ker(A) \cap \ker(B^\top) \cong \ker(A^\top A + BB^\top) \cong \ker(B^\top)/\operatorname{im}(A^\top).$$

Thus, in particular:

$$\ker(\text{rot}) \cap \ker(\text{grad}^\top) \cong \ker(\text{grad}^\top)/\operatorname{im}(\text{rot}^\top).$$

The reader familiar with simplicial homology will recognize that $\mathrm{grad}^\top$ is ∂_1 and $\mathrm{rot}^\top$ is ∂_2, that is they are the homological boundary operators between the appropriate spaces of real-valued chains on the simplicial complex $(\tilde{X}, \tilde{E}, \tilde{T})$. In particular, the homology group (here, vector space) of the subdomain in dimension 1 with real coefficients is precisely $\ker(\mathrm{grad}^\top)/\mathrm{im}(\mathrm{rot}^\top)$, and hence topological triviality implies $\ker(\mathrm{rot}) \cap \ker(\mathrm{grad}^\top) = \{0\}$. Thus for any locally rational relation on a combinatorially and topologically trivial subdomain, there exists an integrable cardinalization, precluding the existence of cycles supported on any loop contained in the subdomain. On a simple domain, every loop is contained in some such subdomain, allowing us to guarantee the generalized axiom holds.

2.5 Examples of Environments with Chordal Budget Graphs

Possessing a chordal budget graph is a much weaker notion than nearly any existing 'completeness' criterion for choice problems. Indeed, most notion of completeness actually yield a budget graph that is the complete graph on X, a significantly stronger condition.

Example 9 (Complete Abstract Environments)**.** Recall that for a general choice environment (X, Σ), the budget collection Σ is complete if it contains all finite subsets of X. If Σ is complete, then clearly its budget graph is complete as well.

Another example of a class of normatively complete environments that have a complete budget graph (and hence a simple domain) are complete collections of linear budgets.

Example 10 (Complete Collections of Linear Budgets)**.** Suppose $X = \mathbb{R}^L_+$ and Σ consists of all (income normalized) linear budgets $B(p, 1) = \{x \in \mathbb{R}^+_L : \langle p, x \rangle \leq 1\}$ for all $p \in \mathbb{R}^L_{++}$. In light of the argument in the complete abstract environments example, it suffices to show that every pair of distinct vectors of commodities forms an edge in the budget graph. Consider $x, y \in \mathbb{R}^L_+$, $x \neq y$ and let $x \vee y$ denote their component-wise supremum. Let $B(\tilde{p}, 1)$ denote any linear budget containing $x \vee y$; then $B(\tilde{p}, 1)$ contains both x and y and hence $\{x, y\} \in E_\Gamma$. Thus the budget graph is again complete, and hence chordal.

More generally, the linearity of the budgets played no role verifying the simplicity of the domain, yielding a natural generalization to the broader class of budget sets considered by [47].

Example 11 (Complete Forges-Minelli Environments)**.** Let $X = \mathbb{R}^L_+$. We will term Σ a complete Forges-Minelli budget collection if (i) every $B \in \Sigma$ is compact and there exists some increasing, continuous $g_B : \mathbb{R}^L_+ \to \mathbb{R}$ such that $B = \{x \in \mathbb{R}^L_+ : g_B(x) \leq 0\}$, and (ii) for every $x \in \mathbb{R}^L_+$, there exists some $B \in \Sigma$ such that $x \in B$. The budget graph associated with any complete Forges-Minelli environment is complete: for any distinct $x, y \in \mathbb{R}^L_+$, by (ii) there exists some $\tilde{B} \in \Sigma$ containing $x \vee y$. By (i), $\tilde{B}$ is downward closed, hence $x, y \in \tilde{B}$ too.

Our next example is of collections of budgets that satisfy a natural notion of completeness but nonetheless yield a budget graph that is generally less-than-complete.

Example 12 (Identifying Experiments)**.** Let X be a locally compact Polish space. We consider those budget collections studied by [32] in the context of the non-parametric identification of continuous preferences. These consist of those collections Σ that consist of (i) a countable collection of binary sets, (ii) such that $\cup_{B \in \Sigma} B$ is dense in X, and (ii) such that for all $x, y \in \cup_{B \in \Sigma} B$, $\{x, y\} \in \Sigma$. When X is uncountable, the budget graph will contain uncountably many isolated vertices. However, for any such experiment by (i) and (iii) the budget graph will be the complete graph on some countable subset of vertices (along with its isolated vertices) which nonetheless is chordal.

Of course, Theorems 1 and 2 are intimately related and, in particular, any well-covered environment has a simple domain.

Example 13 (Well-covered Abstract Environments)**.** If (X, Σ) is a general choice environment with Σ well-covered, then the budget graph is chordal

2.6 Applications

2.6.1 Cardinality-Constrained Choice

The complete cardinality-constrained problem, in essence considered by [11], is one of the most well-known examples of a domain on which the weak axiom is equivalent to strong rationalizability: if Σ contains every subset of X of cardinality at most three, then the weak axiom suffices for the generalized. While Theorem 1 of provides necessary and sufficient conditions with or without a cardinality constraint, for this special case Theorem 2 also provides an intuitive means of characterizing specifically which collections of small budgets are well-covered.

Suppose Σ contains no budget of cardinality greater than three, and denote the sub-collection of three-element budgets by $\Sigma_3 \subseteq \Sigma$. Our first observation is that, for the cardinality-constrained case, the weak axiom implies local rationalizability if and only if every triangle in the budget graph is itself a budget.

Proposition 2. *Let (X, Σ) be a cardinality-constrained choice environment. Then every choice correspondence c that obeys the weak axiom is locally rationalizable if and only if $T_\Gamma = \Sigma_3$.*

Notably, this holds independently of the structure of the budget graph. Making use of this, Theorem 2 immediately provides the following characterization of when the weak and generalized axioms coincide for domains consisting of small budgets.

Corollary 2. *Let (X, Σ) be a cardinality-constrained choice environment. Then the weak axiom characterizes strong rationalizability for any choice correspondence if and only (i) $T_\Gamma = \Sigma_3$, and (ii) the budget graph $\Gamma(X, \Sigma)$ is chordal.*

2.6.2 Deterministic Rationalizability of Stochastic Choice

There has been recent interest in deterministic notions of rationality that could be ascribed to models of stochastic choice. Let X be a finite set, and Σ a not-necessarily-complete collection of budgets. In this context, we take as primitive a collection of probability distributions $\mathbb{P}(\cdot, B)$, for each $B \in \Sigma$, corresponding to the observed frequency with which a given alternative is chosen when an agent is presented with choice set B.

[90] put forward two choice correspondences arising from such stochastic data. The 'upper' choice correspondence associated with $\mathbb{P}$ maps a budget to those alternatives in it that are observed to be chosen with positive probability:

$$C^{\mathbb{P}}(B) = \{x \in B : \mathbb{P}(x, B) > 0\}.$$

The 'lower' correspondence returns only those alternatives that are chosen with maximal frequency:

$$C_{\mathbb{P}}(B) = \{x \in B : \forall y \in B,\ \mathbb{P}(x, B) \geq \mathbb{P}(y, B)\}.$$

The authors term $\mathbb{P}$ completely upper (resp. lower) rational if $C^{\mathbb{P}}$ (resp. $C_{\mathbb{P}}$) is strongly rationalizable by a preference relation and obsesrve that, when Σ is complete, these rationality properties are characterized by $C^{\mathbb{P}}$ and $C_{\mathbb{P}}$ satisfying the weak axiom. Theorem 1 provides an immediate extension of these results to any well-covered budget collection, allowing considerable latitude by extending the results to experiments with less-than-complete domains. This is particularly valuable in the stochastic context, as to obtain reasonable empirical estimates of $\mathbb{P}$, one must sample each budget in Σ repeatedly. As such, reductions in the required breadth of Σ may lead to considerable savings in terms of the observational requirements of the theory.

2.6.3 Aggregation of Incomplete Preferences

Consider a set of national policies X, and let $\mathcal{I}$ be a set of agents. Given a subset of $A \subseteq X$, let $\mathcal{P}(A)$ denote the set of all preference relations on A. For each agent i, a regional preference is a relation $\succsim_i \in \mathcal{P}(S_i)$ for some fixed $S_i \subseteq X$. We interpret this as capturing that agents care only about policies affecting their particular region, or perhaps those neighboring regions. We term the tuple $(X, (S_i)_{i \in \mathcal{I}})$ a society. A social welfare function, for a given society, is simply a map $F : \prod_i \mathcal{P}(S_i) \to \mathcal{P}(X)$.

We say a social welfare function F satisfies the **Pareto** axiom if, whenever all agents who have preferences between two policies x and y agree on their relative ranking, this ranking is preserved by F. Notably, unlike in the case of complete preferences, the incompleteness of preferences may well lead to the set of Pareto social welfare functions being empty. When agents have preferences over only subset(s) of policies, it may be the case that every aggregation mechanism is forced to disregard the preferences of some populations, even when their views are unopposed. Moreover, this may be true even for profiles of (incomplete) preferences containing *no disagagreement of any kind.*

A social welfare function is said to satisfy the **strong unanimity** axiom if, whenever there is no disagreement over any pair of policies by any pair of agents, the social welfare function respects the preferences of the agents. Formally, F satisfies strong unanimity if, whenever for all $i, j \in \mathcal{I}$ it is the case that $\succsim_i |_{S_i \cap S_j} = \succsim_j |_{S_i \cap S_j}$, then $F(\succsim_1, \ldots, \succsim_N)$ is an order extension of $\cup_i \succsim_i$. Strong unanimity is far weaker condition on F than satisfying the Pareto axiom; indeed it requires the Pareto axiom to hold only for special case of unanimous profiles. Nonetheless, the regional nature of the preferences may lead to the set of social welfare functions that satisfy even just strong unanimity being empty. Theorem 2 provides a complete characterization of those societies for which the set of social welfare functions that satisfy strong unanimity is non-empty.

Let $\mathcal{S}$ be a society. In this setting, define the domain associated with $\mathcal{S}$ as the triple $\mathcal{D}(\mathcal{S}) = (X, E_S, T_S)$, where E_S (resp. T_S) are those pairs (resp. triples) of distinct policies belonging to some common S_i. We interpret this as follows: to each agent i, suppose we solicit their pairwise preferences on each binary choice set within S_i. Since we may restrict to profiles that are unanimous, there is agreement over $S_i \cap S_j$ for all pairs of agents, and since each agent i is rational over S_i, the union of the revealed preference relations is precisely a locally rational binary relation on our domain. Thus Theorem 2 implies that if (X, E_S) is chordal, there is always some weak order extending the union of these preferences, and we may always define the value of a social welfare function F at such a profile to be one of these extensions. Conversely, if (X, E_S) is not chordal, Theorem 2 guarantees a profile of unanimous regional preferences admitting no extending weak order, precluding the existence of any F satisfying the strong unanimity axiom for such a society.

Corollary 3. *Let $\mathcal{S} = (X, (S_i)_{i \in \mathcal{I}})$ be a society. The set of social welfare functions for $\mathcal{S}$ satisfying the strong unanimity axiom is non-empty if and only if $\mathcal{D}(\mathcal{S})$ is simple.*

Chapter 3

Preference Regression

3.1 Introduction

The predominant means of testing the predictive accuracy of models of individual decision making has been via the use of choice data. While natural, choice data is primarily ordinal in nature.[1] As a consequence, it is often difficult, both analytically and computationally, to obtain appropriate analogues of the standard econometric toolkit available in other settings. This paper considers instead a novel class of experiments capable of yielding cardinal measurements of the intensity of preference between pairs of alternatives. For a wide range of models, we show this data may be interpreted as observations of utility differences, under a canonical choice of representation. Proceeding from first principles, we develop a theory of mean squared error minimization for such data, and provide a unified framework capable of quantifying the accuracy of a model's predictions, obtaining point-estimates of underlying parameters, and, when data is measured with noise, providing non-parametric tests of rationalizability for individual models.

Any theory that places constraints upon observable behavior will be violated by some data. To quantify the predictive success of a theory then requires an appropriate measure of goodness of fit, or loss function. This is difficult for models of

[1] A noteworthy exception to this is the case of preferences over risk, where there is a large empirical literature based around the elicitation of certainty equivalents; see, for example, [25].

preference: for example, while it makes sense to speak of a firm falling 10% short of profit maximization, it is unclear what the appropriate ordinal analogue is. By considering a novel form of cardinal data reflecting the intensity of preference, we are able to directly measure consistency in 'utility space,' rather than the space of alternatives over which a subject chooses.[2] In this sense our approach is particularly natural: we propose a means of measuring the predictive success of models directly in terms of the primitives of the underlying theory, the preferences themselves.

More formally, many models of interest feature preferences which are preserved under particular collections of transformations of the consumption space. For example, quasilinearity and homotheticity in demand theory, stationarity axioms in dynamic settings, and various independence-type axioms for spaces of lotteries or Anscombe-Aumann acts may all be viewed as 'invariance' properties under appropriate families of transformations. We interpret such families as augmenting alternatives with varying quantities of an abstract, 'virtual' numeraire commodity. We prove a general representation theorem showing that, when a class of preferences satisfies such an invariance property, there is a canonical choice of utility representation with the property that, for any pair of alternatives, the difference in utilities across the pair equals the quantity of virtual numeraire needed to compensate the individual for receiving the less preferred alternative. This corresponds to an appropriate *equivariance* property of

[2] It is generally undesirable to test models of preference using loss functions defined over the consumption space. [112] argues such tests quantify the statistical, rather than economic significance of violations, which may be unrelated; see also [59] for discussion and extensions. [112] instead proposes using the minimal budgetary adjustments needed to remove all inconsistency from the data as a more economically reasonable metric. However, even this only imperfectly proxies for measurement in utility terms. Conversely, our approach may be viewed as precisely minimizing mean squared error over utility space, see subsection 3.5.2, and in particular (3.2).

the representation.[3] Critically, such compensation data is truthfully elicitable, providing an exact, observable proxy for utility differences for wide variety of models of preference.[4]

Relative to the existing literature, our approach provides a number of noteworthy advantages. Firstly, while many measures of consistency provide a numeric indication of the goodness of fit for a particular model, they do not speak to *which* particular preferences from the class of interest are the most consistent.[5] When models are parametric, our framework will generally identify a unique preference as the 'best fit' for a particular model, yielding point estimates of the underlying parameters. This makes it straightforward to estimate, for example, the parameters of a Cobb-Douglas or CES utility, or the prior of a subjective expected utility maximizer with given risk attitude.

When models are nonparametric, our framework instead identifies a set of such preferences. This may be interpreted as identifying the particular *structure* of those preferences which best reflect the (finite) data.[6] Often times the best-fit set is small enough to yield economically meaningful restrictions. For example, given data consistent with a maxmin expected utility preference, the set of priors will not generally be identified. It is, however, straightforward to obtain tight upper and lower bounds

[3]Theorem 4 provides an axiomatization of such utilities, for general families of transformations. We show our representation is unique up to an additive constant, hence its utility differences are well-defined. For details, see subsection 3.3.3.

[4]The elicitability of compensation data of this form is established in Theorem 5.

[5]For example, while the popular efficiency indices of [3] and [112] provide a means of quantifying the degree of failure of the strong axiom of revealed preference, alone they provide no means of ascertaining *which* rational preference(s) best approximate the data. [59] provides an extension to allow for calibrating various parametric models; see also [92].

[6]The set-valued identification is a consequence of our assumption of finite experiments; any two preferences in the best fit set will necessarily be indistinguishable by the experiment at hand. In this sense our identification result is the best possible for the class of problems considered.

on the true set of priors (see 24). As an application, we show how these bounds may be utilized to obtain predictions about speculative trade between agents, and the Pareto optimality of insurance in exchange economies with uncertainty. We provide an example in which our results are able to fully identify the Pareto frontier of an economy, even without identifying a single individual preference.

In many cases our framework is able to quantify not only the overall predictive accuracy of a model, but which specific axioms or constraints are most (or least) responsible. For example, in the context of maxmin expected utility preferences, we show not only how to evaluate the predictive success of the model as a whole, but show how to obtain straightforward estimates of the 'shadow price,' in model fit terms, of specifically relaxing the ambiguity aversion axiom of [52].[7] This granularity allows for the design of experiments which test models broadly, rather than focusing on specific aspects, but which nonetheless allow for fine-tuned analysis of the predictive success of individual components of the theory.

Finally, when data is observed with noise or is otherwise stochastic, we provide explicit statistical tests of rationalizability for a wide range of models and functional forms. For many models, the problem of mean squared error minimization reduces to computing the distance from the vector of observations to a polyhedral set of rationalizable vectors. We leverage this structure to construct explicit tests for rationalizability for a variety of models across multiple domains. As we consider measurements in utility space, this allows us to construct direct statistical tests of the *economic* significance of departures from rationalizability, in the sense of [112], for a large collection of models.

[7] See 21 for details.

When cross-sectional data is sampled from a population, this provides a means of testing hypotheses about whether a particular model or functional form reasonably represents a linear aggregate of a heterogeneous population.[8] For example, a social planner could use our methodology to test whether maximizing a particular choice of utility or objective function is a reasonable approximation to the social preferences of a population, in cases where there is insufficient information to conduct a traditional choice-based welfare analysis.

Example 14. Suppose one wishes to investigate how well-approximated a particular individual's preference over bundles of two commodities are by a continuous quasi-linear utility function of the form:[9]

$$U(x, y) = v(y) + x.$$

Preferences representable by such a utility are invariant under receiving more of the numeraire commodity, x. For example, $(3, 5) \succsim (4, 2)$ if and only if for all $\alpha \geq 0$, $(3 + \alpha, 5) \succsim (4 + \alpha, 2)$. In particular, this implies the indifference curves of any such preference are horizontal translates of one another.

Consider the three bundles $(x_1, y_1), (x_2, y_2)$, and (x_3, y_3) indicated in Figure 3.1. Suppose the subject is presented with all three possible pairs of bundles from this collection and one observes, for each pair of bundles, both (i) which bundle is preferred, and (ii) what quantity of numeraire, in conjunction with receiving the less-preferred

[8]For example, [25] finds strong evidence of heterogeneous risk attitudes in a large population of university students, with roughly 80% of subjects exhibiting departures from linear probability weighting consistent with prospect theory, while 20% broadly conform with the predictions of expected utility maximization. In the presence of individual-level heterogeneity, our results can be used as a robustness check for representative agent assumptions.

[9]Our use of the phrase 'the individual's preference' here is meant informally; in particular we do not assume, a priori, that the individual's choice behavior is consistent with the maximization of any complete and transitive binary relation.

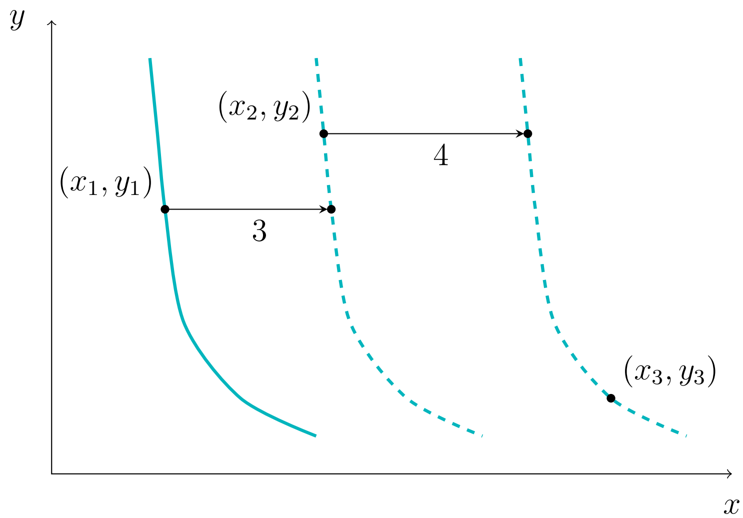

Figure 3.1: Quasilinear preferences and compensation. When preferences are representable by a quasilinear utility of the form $u(x, y) = v(y) + x$, indifference curves are horizontal translates of one another. Geometrically, translating any point on the (x_1, y_1) indifference curve 3 units to the right makes it fall precisely on the (x_2, y_2) indifference curve. If one translates the result again by 4, the point must then lie on the (x_3, y_3) curve.

bundle, would make the subject indifferent relative to receiving the more-preferred bundle. It is observed first that the subject is indifferent between $(x_1 + 3, y_1)$ and (x_2, y_2), and similarly between $(x_2 + 4, y_2)$ and (x_3, y_3). If the subject did possess a quasilinear preference, this would imply that translating the (x_1, y_1) indifference curve 3 units to the right makes it perfectly overlap the (x_2, y_2) indifference curve, and likewise, translating the (x_2, y_2) indifference curve 4 units to the right makes it overlap the (x_3, y_3) curve. Thus, for the data to be consistent with any quasilinear preference, it is necessary that, for the third pair, the subject be indifferent between $(x_1 + 7, y_1)$ and (x_3, y_3). In fact, this 'adding-up' condition is also sufficient for the existence of a continuous, quasilinear rationalizing utility.[10]

[10] This follows from Theorem 6.

More generally, the collection of 'quasilinear-rationalizable' data sets are precisely those which satisfy the above adding-up constraint. These vectors form a linear subspace (here, a plane) in the space of all possible data vectors that could result from the experiment. Given data *inconsistent* with these hypotheses, there is a unique mean squared error minimizing choice of consistent dataset, obtained by orthogonally projecting the data onto the subspace of rationalizable vectors. The squared distance between the original data and its 'best fit' serves as a natural measure of goodness of fit for the hypothesis of quasilinearity.

Finally, it is straightforward to further restrict to only those vectors rationalizable by quasilinear utilities with additional structure. For example, if one wished to additionally require $v(y)$ be increasing and concave, the set of rationalizable vectors would form a polyhedral subset of the quasilinear-rationalizable plane.[11] To compute the best fit, one would simply project the data onto this subset instead, and the distance between the data and its projection provides a quantification of the goodness of fit.

3.2 Related Literature

While our results hold across a variety of decision theoretic models, the literature on preferences under ambiguity is a rich source of applications for our methodology. We focus on preferences over monetary acts, as in, for example [20], [95], [24], [14], [5], or [31].[12] For any data set of pairwise compensation differences, we provide characterizations of the empirical content of a variety of models of preference under ambiguity,

[11] For an explicit description of this set for a general experiment, see 22.

[12] Preferences over monetary acts may alternatively be interpreted as 'risk-neutral' preferences over a richer domain that includes both subjective and objective uncertainty, but where each monetary lottery has been replaced with its certainty equivalent.

including subjective expected utility (Anscombe and Aumann 9), Choquet expected utility (Schmeidler 101), maxmin expected utility (Gilboa and Schmeidler 52), variational (Maccheroni et al. 80), as well as dual-self and dual-self variational preferences (Chandrasekher et al. 33).[13] Taken as a whole, our results yield not only a class of experiments capable of simultaneously differentiating between these models, but also revealed preference-like characterizations and explicit statistical tests for each.

Compensation differences involve two distinct pieces of information: which alternative in a particular pair is more preferred, and by how much. We introduce an extension of the Becker-DeGroot-Marschak mechanism capable of simultaneously and truthfully eliciting both these unknowns in experimental settings. The form of experiment considered here may then be viewed as a mixture of two standard methods of eliciting preferences over risky or uncertain prospects: having subjects make pairwise comparisons (e.g., Hey et al. 64, Abdellaoui et al. 1) and eliciting reservation prices (e.g., Becker et al. 16, Halevy 58).[14]

This paper contributes to a growing recent literature on the statistical testing of various decision- and demand-theoretic models. Much of this work has focused on constructing model-specific tests, for example [72], [38], [49], [28], and [106]. In constrast, we put forward a general methodology that applies to a wide range of different models, over a variety of domains. To obtain critical values for our test, we make use of an implementation of the non-differentiable delta method of [44] due to [69]. A notable benefit of this approach is its flexibility: we provide a means of testing both individual axioms and whole models, allowing far more granular insights into

[13]As noted in [33], dual-self expected utility is a particular choice of representation of the class of invariant biseparable preferences of [51]. As our results are ordinal in nature, all our results on empirical content and testing hold true for the underlying class of preferences, rather than resting upon a particular choice of representation.

[14]This stands in contrast to the ‘allocation’ approach of [79], [8], [35], [63], and [5].

not only which models may fail to be consistent with the data, but which aspects of the model are most responsible for the rejection.[15]

Finally, the use of least-squares techniques to aggregate incomplete and potentially inconsistent observations into a coherent ranking has a long history, e.g. [61], [109]. Recently, there has been renewed mathematical interest in the problem of establishing an optimal statistical ranking from an inconsistent, incomplete, and noisy dataset. We draw on a number of ideas from this literature, including the representation of data as a flow on a particular network whose structure reflects the incompleteness our observations, and the various associated regression theories for such problems, see [65], [71], [91] and references therein. These ideas have already found application elsewhere in economics, including game theory, social choice, and revealed preference (resp. Candogan et al. 26, Csató 36, Caradonna 27).

3.3 Invariant Preferences

3.3.1 Model

In this section, we specify a general, flexible, model of preferences possessing certain invariance properties. Let (X, d) be a metric space of **alternatives** that an agent has preferences over. A preference is a complete and transitive binary relation on X, which we will denote by $\succsim$. As is standard, we use $\succ$ and $\sim$ to denote the asymmetric (resp. symmetric) components. A preference is continuous if, for all $x \in X$, $\{x' \in X : x' \succsim x\}$ and $\{x' \in X : x \succsim x'\}$ are closed.

We say a jointly continuous function $\phi : \mathbb{R}_+ \times X \to X$ defines a **virtual numeraire** commodity if (i) for all $x \in X$, $\phi(0, x) = x$, and (ii) for all $\alpha, \beta \in \mathbb{R}_+$ and

[15] Our method of constructing test statistics also draws heavy from the statistical and econometric literature on shape-constrained regression, in particular [7], [75], and [102].

all $x \in X$, $\phi\big(\beta, \phi(\alpha, x)\big) = \phi(\alpha+\beta, x)$. Formally, such a map ϕ defines a continuous **action** of the monoid $\mathbb{R}_+$ on X.[16] For our purposes, ϕ corresponds to a collection of transforms $X \to X$, one for each $\alpha \geq 0$, which we interpret as augmenting an alternative with some quantity of virtual numeraire. Thus for any $\alpha \geq 0$ and any $x \in X$, we interpret $\phi(\alpha, x)$ as x plus α additional units of numeraire. Property (i) ensures that the transform corresponding to adding no units numeraire does not alter any alternative; property (ii) is a path-independence condition: adding β units of numeraire to the alternative consisting of x plus α units of numeraire is the same as simply adding $\alpha + \beta$ units of numeraire to x at once.

Given a space X equipped with a virtual numeraire ϕ, we will consider those continuous preferences which satisfy the following three axioms:

(N.1) **Invariance**: For all $\alpha \in \mathbb{R}_+$, $x, y \in X$:

$$x \succsim y \iff \phi(\alpha, x) \succsim \phi(\alpha, y).$$

(N.2) **Monotonicity**: For all $\alpha \in \mathbb{R}_+$, $x \in X$:

$$\phi(\alpha, x) \succsim x,$$

with indifference if and only if $\alpha = 0$.

(N.3) **Compensability**: For all $x, y \in X$,

$$x \succ y \implies \exists\, \alpha \in \mathbb{R}_+ \text{ s.t. } \phi(\alpha, y) \sim x.$$

Invariance says that adding the same quantity of numeraire to two alternatives does not affect the preference between them. It rules out numeraire-based 'wealth effects'

[16] A monoid is a semigroup with identity; see [48]. Formally, $\mathbb{R}_+$, equipped with the the usual notion of addition $+$, is a monoid.

where, when coupled with a high enough quantity of additional numeraire, an agent's preferences between two alternatives reverses. Monotonicity says the virtual numeraire commodity is a good. Compensability is a richness condition for ϕ. It ensures that any preference between two alternatives can always be offset by some quantity of numeraire. It rules out lexicographic behavior where no amount of numeraire could compensate an agent for receiving a less-preferred alternative.

If $y \succsim x$, then (N.3) guarantees there is some α such that $\phi(\alpha, x) \sim y$. We term this α the **compensation difference** from x to y. Note that by (N.2) this quantity is necessarily unique. Together, the axioms (N.1) - (N.3) will be shown to characterize those preferences which admit a representation under which the compensation difference, measured in numeraire units between any pair of alternatives, is precisely the utility difference of the alternatives.

3.3.2 Examples

Example 15. (Quasilinear Preferences): Suppose $X = \mathbb{R}^2_+$, and let $\phi\big(\alpha, (x, y)\big) = \big(\alpha + x, y\big)$. Any continuous preference $\succsim$ that admits a utility of the form $U(x, y) = v(y) + x$ clearly satisfies (N.1)-(N.3). Here, the compensation difference measures the quantity of the first commodity needed to offset a preference between two bundles.

Example 16. (Stationary Preferences for Dated Rewards): Suppose $X = \mathbb{R}_+ \times Z$ with $\phi\big(\alpha, (t, z)\big) = (\alpha + t, z)$ where, following [46], a pair $(t, z) \in X$ represents delivery of some alternative $z \in Z$ at t units of time in the future.[17] It is natural, when prizes are goods, to instead require (N.2) to hold with the opposite relation, to reflect

[17] [89] provide a complementary interpretation of preferences over X as the commitment preferences of an agent.

impatience on the part of the agent. Suppose $\succsim$ admits a utility of the form:

$$U(t, z) = \rho^t v(z),$$

where $0 < \rho < 1$ and $v : Z \to \mathbb{R}_{++}$. Then $\succsim$ satisfies (N.1) by an analogous argument to the preceding example.[18] Moreover, for any (t, z) and any $\alpha \geq 0$, $\rho^{t+\alpha} v(z) < \rho^t v(z)$. Hence the reverse analogue of (N.2) holds. Similarly, if $(t, z) \succ (t', z')$, then for some $\alpha > 0$, $U(\alpha + t, z) = U(t', z')$ hence the reverse analogue of (N.3) is satisfied. In this setting, the compensation difference measures the amount of time one must postpone the delivery of the more desirable dated reward to achieve indifference.

Example 17. (Homothetic Preferences): Let $X = \mathbb{R}^L_+ \setminus \{0\}$ represent the standard demand theoretic consumption space minus the origin, and define $\phi : \mathbb{R}_+ \times X \to X$ via:

$$\phi(\alpha, x) = e^{\alpha} x.$$

Suppose a preference $\succsim$ admits a utility U that is positive homogeneous.[19] The relation $\succsim$ satisfies (N.2) and (N.3) if, for example, U is strictly positive on X. It necessarily also satisfies (N.1). Indeed the preference is homothetic if and only if $\succsim$ satisfies (N.1). To see this, suppose $x \succsim y$. If $\lambda \geq 1$, $\lambda x \succsim \lambda y$ is equivalent to $e^{\ln \lambda} x \succsim e^{\ln \lambda} y$, hence the claim follows from (N.1) with $\alpha = \ln \lambda$. If $\lambda < 1$, by (N.1):

$$\begin{aligned} &\lambda x \succsim \lambda y \\ \iff &e^{\ln \frac{1}{\lambda}} \lambda x \succsim e^{\ln \frac{1}{\lambda}} \lambda y \\ \iff &x \succsim y, \end{aligned}$$

[18]In this setting, (N.1) corresponds to [46]'s (A.2), a monotonicity axiom, and (A.5), their stationarity axiom (cf. the stationarity axiom of Ok and Masatlioglu 89). Our axiom (N.2) is closely related to Fishburn and Rubinstein's axiom (A.3); they obtain (N.3) as a consequence of a continuity axiom and the particular topological structure of the consumption space.

[19]A function U is said to be positive homogeneous if, for all $\alpha > 0$, $U(\alpha x) = \alpha U(x)$.

and as $x \succsim y$ by hypothesis, it again follows that $\lambda x \succsim \lambda y$. Thus given ϕ, (N.1) corresponds to homotheticity.[20] Here compensation differences measure the amount, holding proportions fixed, one would need to scale up the less-preferred bundle of commodities to achieve indifference.

Example 18. (Mixture Independence): Let $\tilde{X}$ be a finite set, and let $\Delta(\tilde{X})$ denote the set of lotteries supported on $\tilde{X}$. Let $\bar{x} \in \tilde{X}$ be arbitrary. Define $\phi_{\bar{x}} : \mathbb{R}_+ \times \Delta(\tilde{X}) \to \Delta(\tilde{X})$ via:

$$\phi_{\bar{x}}(\alpha, p) = e^{-\alpha} p + (1 - e^{-\alpha})\delta_{\bar{x}},$$

where $\delta_{\bar{x}}$ denotes a dirac measure or point mass centered at $\bar{x}$. This defines a virtual numeraire, as:

$$\begin{aligned}
\phi_{\bar{x}}\big(\beta, \phi(\alpha, p)\big) &= e^{-\beta}\big(e^{-\alpha} p + (1 - e^{-\alpha})\delta_{\bar{x}}\big) + (1 - e^{-\beta})\delta_{\bar{x}} \\
&= e^{-(\beta+\alpha)} p + \big(1 - e^{-(\beta+\alpha)}\big)\delta_{\bar{x}} \\
&= \phi_{\bar{x}}(\beta + \alpha, p),
\end{aligned}$$

and $\phi(0, p) = p$. Now, let $\succsim$ be any von Neumann-Morgenstern preference on $\Delta(\tilde{X})$ that ranks $\delta_{\bar{x}}$ as the unique, most-preferred lottery. Then the restriction of $\succsim$ to $\Delta(\tilde{X}) \setminus \{\delta_{\bar{x}}\}$ satisfies (N.1) to (N.3).[21] Compensation differences here measures how much one would need to mix the less-desirable lottery with the sure-thing $\delta_{\bar{x}}$, before a subject deems the resulting lottery as good as the more-desirable lottery in a pair.

[20] More generally, there is no added implication from considering an invariance axiom for an action of the full group $(\mathbb{R}, +)$ (or, here, the isomorphic group $(\mathbb{R}_{++}, \cdot)$) when possible, because of the 'if and only if' in (N.1). In particular, even if ϕ extends from an action of $\mathbb{R}_+$ to an action of $\mathbb{R}$, the set of invariant preferences will be the same. This is true generally, see for example, [29], footnote 5.

[21] That $\delta_{\bar{x}}$ is the unique preference-maximal lottery is required to ensure (N.2) and (N.3). If there were some $p' \succsim \delta_{\bar{x}}$, then (N.2) would fail for $\phi(\alpha, p')$, as adding more numeraire corresponds to increasing the mixing coefficient on $\delta_{\bar{x}}$, which must weakly decrease utility. This is also motivates the restriction of $\succsim$ to $X \setminus \{\delta_{\bar{x}}\}$. Similarly, if there were p, p' such that $p \precsim \delta_{\bar{x}} \prec p'$, there would fail to be any compensation value for the pair $\{p, p'\}$ under $\phi_{\bar{x}}$.

Example 19. (Translation-Invariant Preferences): Let S be a finite set of states of the world, and $X = \mathbb{R}^S$ denote the space of all real-valued (monetary) acts.[22] Define $\phi : \mathbb{R}_+ \times X \to X$ via $\phi(\alpha, x) = x + \alpha \mathbb{1}_S$. A function $U : X \to \mathbb{R}$ is said to be translation-invariant if, for all $\alpha \in \mathbb{R}$,

$$U\big(x + \alpha \mathbb{1}_S\big) = U(x) + \alpha.$$

Then any preference $\succsim$ on X that admits a translation-invariant utility satisfies (N.1) to (N.3).[23] [56] interpret translation-invariance over utility acts as reflecting constant absolute ambiguity aversion. Interpreting acts as portfolios of Arrow securities, compensation differences in this setting correspond to the quantity of bonds (i.e. assets paying off the same across each possible state) an agent with risk-neutral preferences must additionally be awarded, in addition to a less preferred portfolio, to be indifferent with holding a more preferable one.

3.3.3 Representation

In this section, we establish a particular utility representation for any preferences satisfying (N.1) to (N.3). We say that a utility $U : X \to \mathbb{R}$ is **additive-equivariant** if, for all $\alpha \in \mathbb{R}_+$, and all $x \in X$:[24]

$$U\big(\phi(\alpha, x)\big) = U(x) + \alpha.$$

[22]We consider finite S to avoid measurability concerns. The example remains true if, for example, one has a measurable space (S, Σ) and X is a cone of real-valued measurable maps that contains the non-negative constant functions.

[23]By an argument analogous to that in the homothetic preferences example, even though translation invariance allows for negative α, it is still implied by (N.1).

[24]Given a monoid M, and actions ϕ_X and ϕ_Y of M on sets X and Y, a map $f : X \to Y$ is said to be **equivariant** if, for all $m \in M$ and all $x \in X$:

$$f\big(\phi_X(m, x)\big) = \phi_Y\big(m, f(x)\big).$$

Additive-equivariance corresponds to the special case where M is $\mathbb{R}_+$, Y is $\mathbb{R}$, and the action ϕ_Y is simply addition.

Roughly speaking, additive equivariance requires that the 'marginal utility' of additional numeraire (i) be constant across X, and (ii) normalized to unity. Thus for any additive equivariant representation, for every x, the addition of α units of numeraire yields precisely α extra utility on top of the utility from x. Crucially, if α is the compensation difference from x to y, then for any additive-equivariant utility:

$$\begin{aligned} U(y) &= U\big(\phi(\alpha, x)\big) \\ &= U(x) + \alpha, \end{aligned}$$

and hence $\alpha = U(y) - U(x)$.

Theorem 4. *Suppose that ϕ is a virtual numeraire for X. Then a continuous preference $\succsim$ on X satisfies (N.1) - (N.3) if and only if it admits a representation by a continuous, additive-equivariant utility. Additionally, if U and V are additive equivariant representations of the same preference, then there exists $\beta \in \mathbb{R}$ such that $U + \beta = V$.*

If a preference admits an additive-equivariant representation U, then the compensation difference between any pair of alternatives exists, and is equal to the utility difference under U. Theorem 4 establishes that the preferences admitting such a utility are precisely those which satisfy (N.1) - (N.3), and that the utility differences of every additive-equivariant representation for such a preference coincide. Thus additive-equivariant utilities form a canonical choice of representation for our purposes: if a preference satisfies (N.1) - (N.3) then its compensation differences are precisely the utility differences under some, and hence all, additive-equivariant representations.

Remark 1. Continuity of the virtual numeraire and preference plays no role in the proof of Theorem 4 other than in verifying the continuity of the representation. A non-topological variant of the result, where X is an arbitrary set equipped with an action ϕ, follows from an essentially identical argument. However, in this case nothing can be said about the continuity of U. This is notable as topological assumptions are often crucial in ensuring the existence of a utility representation. Here, (N.1) - (N.3) alone suffice without any further stipulations on X or $\succsim$.

Remark 2. The uniqueness of an additive-equivariant representation up to an additive constant, rather than increasing affine transform, follows from additive-equivariance implicitly normalizing utility units to numeraire units. If U is additive-equivariant for some numeraire ϕ, and $\gamma > 0$, then γU is not additive-equivariant for ϕ; it is, however, for the modified numeraire $\phi_\gamma(\alpha, x) = \phi(\alpha\gamma, x)$. Here, ϕ_γ can be interpretted as ϕ, but measured in different units: for example, if $\gamma = 2$, and ϕ measured numeraire in dollars, then ϕ_γ measures the same numeraire, denominated in half dollars.[25] One implication of this is the failure of ρ to be identified in 16, see [46].

3.3.4 Regularity Assumptions

We will henceforth impose the following three technical regularity conditions on model primitives, which, while not required for Theorem 4, are required in the sequel.

(A.1) **Injectivity**: For all $\alpha \in \mathbb{R}_+$, the map $\phi(\alpha, \cdot)$ is injective.[26]

[25] Formally, the group of order-preserving monoid isomorphisms $\mathbb{R}_+ \to \mathbb{R}_+$ (with composition as group operation) is isomorphic to the multiplicative group of positive reals; see [48]. Thus equivalence under linear change-of-numeraire-units of this form (and hence additional equivalence of additive-equivariant representations up to positive scalar multiples) is precisely the extra degree of freedom that would be obtained from requiring ϕ to be fixed only up to isomorphism (i.e. remaining agnostic of our measurement scale).

[26] If the set of preferences that satisfy (N.1) - (N.3) is non-empty, then for all $x \in X$, $\phi(\cdot, x)$ is necessarily injective: if for some $x \in X$ and $\alpha < \beta$

$$\phi(\alpha, x) = \phi(\beta, x),$$

then $\phi\big(\beta - \alpha, \phi(\alpha, x)\big) = \phi(\alpha, x)$, but as $\beta - \alpha > 0$, every reflexive relation on X violates (N.2).

We say that an alternative x is **reachable** from y, denoted $y \trianglelefteq x$ if there exists $\alpha \geq 0$ such that $x = \phi(\alpha, y)$. That is, if x is equal to y plus some additional numeraire. Let $\sim_{\trianglelefteq}$ denote the symmetric closure of this relation.[27] If ϕ satisfies (A.1), then $\sim_{\trianglelefteq}$ is an equivalence relation (see 4).

(A.2) **Cross Section**: There exists a continuous map $s : X/\sim_{\trianglelefteq} \to X$, such that, for all $y \in X/\sim_{\trianglelefteq}$,

$$(q \circ s)(y) = y,$$

where q is the quotient map taking X to $X/\sim_{\trianglelefteq}$, which carries its quotient topology.

(A.3) **No Loitering**: For all $x \in X$, there exists $\varepsilon > 0$ and $T > 0$ such that, for all $x' \in B_\varepsilon(x)$ and all $\alpha > T$:

$$\phi(\alpha, x') \notin B_\varepsilon(x),$$

where $B_\varepsilon(x)$ denotes the ε-ball about x.

Injectivity simply requires that there not be any pair of distinct alternatives x and y that become *equivalent* after being combined with a sufficient, common quantity of numeraire. Thus the process of adding numeraire is, in principle, reversible. Cross section is a weak technical assumption that ensures that at least some of the indifference curves of any preference satisfying (N.1) - (N.3) are sufficiently connected. Absent it, it could be the case that every indifference curve misses some equivalence

[27] Recall the symmetric closure of a relation R is the smallest symmetric relation containing R.

classes of $\sim_{\trianglelefteq}$.[28] Finally, no loitering ensures that no alternative can be regarded as the result of adding infinite numeraire to any other.

3.4 Data and Eliciitiation

An **experiment** is a finite collection $\mathcal{E}$ of pairs of alternatives such that no two alternatives (belonging even to differing pairs) are related under $\trianglelefteq$.[29] For a given agent and pair $\{x, y\} \in \mathcal{E}$, we will assume an observation of both (i) which alternative in $\{x, y\}$ is (weakly) more preferable than the other, and (ii) how much virtual numeraire is needed, in addition to receiving the less preferable alternative, to make the agent indifferent with receiving the more preferable alternative. That is, we assume we observe the compensation difference between the less and more favorable alternatives.[30]

We suppose a data set consisting of $N \geq 1$ repetitions of such an experiment. Formally, let $\vec{\mathcal{E}}$ denote the collection of all ordered pairs (x, y) such that $\{x, y\} \in \mathcal{E}$.

---

[28]For a drastic example, let $f : \mathbb{R} \to \mathbb{R}$ denote a discontinuous solution to Cauchy's functional equation $f(x + y) = f(x) + f(y)$. The graph of any such f is dense in the plane; see [2] Chapter 1 Theorem 3. Let $X = \{(x, y) \in \mathbb{R}^2 : y \geq f(x)\}$ be the epigraph of f, and let $\mathbb{R}_+$ act on X via addition along the second coordinate. Clearly (A.2) does not hold, and for any continuous preference on X obeying (N.1) - (N.3), *every* indifference curve is completely disconnected, and misses uncountably many equivalence classes in $X/\sim_{\trianglelefteq}$.

[29]The assumption that no pairs of alternatives $\{x, y\} \in \mathcal{E}$ are $\trianglelefteq$-related amounts to not inquiring how much numeraire would make a subject indifferent between receiving x versus x plus α units of numeraire. The stronger requirement that no two of alternatives belonging even to different pairs are $\trianglelefteq$-related is purely for convenience. It may be dropped, at the cost of requiring a slight modification to our our rationalizability condition. See subsection 3.5.1.

[30]Recall this is defined as the numerical quantity $\alpha \geq 0$ such that

$$\phi(\alpha, x) \sim y,$$

when $y \succsim x$.

A **data set** is a collection $\{Y^n\}_{n=1}^N$ of vectors in $\mathbb{R}^{\vec{\mathcal{E}}}$, where, for each $(x, y) \in \vec{\mathcal{E}}$:

$$Y_{xy}^n = \begin{cases} \alpha & \text{if } \phi(\alpha, x) \sim^n y \\ -\alpha & \text{if } \phi(\alpha, y) \sim^n x, \end{cases}$$

where $\phi(\alpha, x) \sim^n y$ denotes that α is the compensation difference between x and y in the n-th repetition. Since $Y_{xy}^n = -Y_{yx}^n$, we may identify the space of all possible data sets with $\mathbb{R}^{\mathcal{E}}$ by fixing a choice of ordering for each pair. When $N = 1$, a data set corresponds simply to observing a finite set of compensation differences for a single agent. Unless otherwise specified, we interpret the case of $N > 1$ as corresponding to observations of a sample of N agents from some fixed population.[31] In section 3.6 we will consider how to test hypotheses about the expected behavior of such a population.

3.4.1 Elicitation

In this section we will present a dominant-strategy incentive-compatible mechanism to truthfully elicit compensation differences data. Our approach may be seen as a generalization of [16]. We will consider the elicitation problem for a given observation, and extend to a full experiment via lottery.[32] [33] Let $\{x, y\} \in \mathcal{E}$ be an arbitrary pair of alternatives. We first define two intermediate mechanisms: in the x-mechanism, the agent is offered the opportunity to submit a non-negative 'sell price' in numeraire units for x, denoted s, to a computerized buyer. The buyer simultaneously and blindly selects a non-negative 'buy' price b. If $s > b$, no trade occurs and the agent is awarded

[31] We note such data may alternatively be interpreted as repeated observations of the noisy preference(s) of a single agent.

[32] Stemming back to [68] there has been concern that, in theory, random-lottery incentive systems rely on implicit assumptions about choice under uncertainty that may be problematic. However, there is a wide range of empirical evidence suggesting that these concerns do not bear out in practice, e.g. [108], [62], and [78].

[33] For a discussion of incentive compatibility of this mechanism in the case of experiments over ambiguous prospects, see [13].

x. If $b \geq s$, then a trade occurs, and instead of x, the agent receives $\phi(b, y)$. We analogously define the y-mechanism. Compensation differences may then be elicited by presenting the subject with a choice: they are invited to submit a sell price in either the x- or y-mechanism, but not both. However, in whichever mechanism they do not choose, a sell price of 0 will be submitted on their behalf. After the bids have been submitted, a coin is flipped to select either x or y, and the associated mechanism's reward is allocated to the agent, regardless of which intermediate mechanism they chose to manually submit a sell price for.

We model the agent's decision problem using the states of the world formalism. We do so to highlight that the incentive-compatibility of our mechanism does not depend on the manner in which the subject handles probabilities. Suppose that $\Omega = \mathbb{R}^2_+ \times \{x, y\}$ denotes the payoff-relevant states of the world; the tuple (b_x, b_y, z) denotes the state in which the computer selects bids b_x in the x-mechanism, b_y in the y-mechanism, and the payoff-determining mechanism is $z \in \{x, y\}$. A choice of action for the agent consists of a tuple in $\{x, y\} \times \mathbb{R}_+$, corresponding a choice of which intermediate mechanism to participate in, and what sell price to submit there. Let X^* denote the set maps from $\Omega \to X$ that are awarded by this mechanism. We assume the agent has preferences $\succsim^*$ over X^* and say these are **consistent** with their preference $\succsim$ over X if, for all $f, g \in X^*$, $f(\omega) \succsim g(\omega)$ for all $\omega \in \Omega$ implies $f \succsim^* g$.

Theorem 5. *Suppose an agent has preferences $\succsim$ on X that satisfy (N.2) and (N.3), and preferences $\succsim^*$ over X^* that are consistent with $\succsim$. Then choosing to submit a bid equal to their true compensation difference, in the mechanism corresponding to the more-preferred alternative, is $\succsim^*$-optimal.*

Remark 3. The assumption that the agent had a well-defined preference relation $\succsim^*$ over X^* is not required for the result. Even if $\succsim^*$ is a highly incomplete and non-transitive relation, Theorem 5 remains valid so long as $\succsim^*$ remains consistent in the above sense with $\succsim$. Consistency alone implies bidding equal to one's true compensation difference, in the more-preferred alternative's mechanism, yields an act that is at least as $\succsim^*$-preferable as (and hence comparable to) every other act in X^*.

3.5 Goodness of Fit

An experiment $\mathcal{E}$ may be associated with an undirected graph $(\mathcal{V}, \mathcal{E})$ whose vertices are those alternatives featuring in the experiment and whose edges are the pairs defining the experiment:

$$\mathcal{V} = \bigcup_{\{x,y\} \in \mathcal{E}} \{x, y\}.$$

We will assume henceforth that $(\mathcal{V}, \mathcal{E})$ is always a connected graph.[34] A **flow** on $(\mathcal{V}, \mathcal{E})$ is a skew-symmetric, real-valued function on $\vec{\mathcal{E}}$.[35] Given a data set $\{Y^n\}_{n=1}^N$, let $\bar{Y} = \frac{1}{N} \sum Y^n$. For a given agent, Y^n_{xy} is the compensation difference between x and y for agent n, hence $\bar{Y}_{xy}$ reflects the sample average compensation difference between x and y. Thus, for any experiment, $\bar{Y}$ defines a flow on $(\mathcal{V}, \mathcal{E})$. Conversely, every flow may be regarded as arising from a data set in this manner.

3.5.1 Rationalizability

We say that a data set $\{Y^n\}_{n=1}^N$ is **cardinally consistent** if, for every finite sequence (x_0, x_1), (x_1, x_2), $\ldots$, $(x_{L-1}, x_0) \in \vec{\mathcal{E}}$,

$$\sum_{l=0}^{L-1} \bar{Y}_{x_l x_{l+1}} = 0, \tag{3.1}$$

where subscripts are understood mod-L. By minor abuse of notation, we also refer to individual flows as being cardinally consistent if (3.1) holds. Cardinal consistency

[34] This is without loss of generality; if $(\mathcal{V}, \mathcal{E})$ is not connected, all results simply hold for each connected component.

[35] That is, $F : \vec{\mathcal{E}} \to \mathbb{R}$ is a flow if and only if, for all $(x, y) \in \vec{\mathcal{E}}$, $F_{xy} = -F_{yx}$.

requires a flow to belong to the kernels of a finite collection of linear functionals, thus the sub-collection of cardinally consistent flows forms a linear subspace of the space of all flows.

A data set is **rationalized** by a continuous, additive-equivariant preference if there exists a continuous preference relation $\succsim$ on X that satisfies (N.1) - (N.3), such that for all $(x, y) \in \vec{\mathcal{E}}$ with $\bar{Y}_{xy} \geq 0$:

$$\bar{Y}_{xy} = \alpha \iff \phi(\alpha, x) \sim y.$$

Similarly, we say $\{Y^n\}_{n=1}^N$ is rationalized by an additive-equivariant utility U if, for all $(x, y) \in \vec{\mathcal{E}}$:

$$\bar{Y}_{xy} = U(y) - U(x).$$

When the data contain observations of a sample population of agents, additive-equivariant rationalizability refers to whether the sample population is rationalizable in expectation. It is clear that if $\bar{Y}$ is rationalizable by an additive-equivariant utility, it will necessarily be cardinally consistent. Our next result establishes that cardinal consistency is also sufficient.

Theorem 6. *Let (X, ϕ) satisfy (A.1) - (A.3), and suppose the set of continuous preferences satisfying (N.1) - (N.3) is non-empty. Then for every experiment $\mathcal{E}$, for any dataset, the following are equivalent:*

(i) The data are cardinally consistent.

(ii) The data are rationalizable by a continuous preference satisfying (N.1) - (N.3).

(iii) The data are rationalized by a continuous, additive-equivariant utility.

Theorem 6 characterizes the testable implications of additive-equivariance for any experiment. It also highlights a benefit of additive-equivariant preferences: testing rationalizability amounts to investigating whether or not the data $\bar{Y}$ lies in a fixed, linear subspace that is explicitly determined by the structure of the experiment $\mathcal{E}$.

3.5.2 Preference 'Regression'

Least Squares Theory

In the preceding section we showed that the collection additive-equivariant rationalizable data vectors forms a linear subspace of the space of flows on $(\mathcal{V}, \mathcal{E})$. We now provide an alternative characterization. For a given experiment, let $\mathcal{F} \subsetneq \mathbb{R}^{\vec{\mathcal{E}}}$ denote the vector space of flows on $(\mathcal{V}, \mathcal{E})$, and let $\mathcal{U}$ denote the space of utility functions over the vertices $\mathcal{V}$:

$$\mathcal{U} = \{u : \mathcal{V} \to \mathbb{R}\}.$$

To any utility $u \in \mathcal{U}$ one may associate its **gradient**, a flow whose value on an oriented edge is given by the signed difference of the utility values at its endpoints:

$$(\operatorname{grad} u)_{xy} = u_y - u_x.$$

This defines a linear map grad : $\mathcal{U} \to \mathcal{F}$. The following two propositions are routine; for completeness, we provide proofs in section C.0.2.

Proposition 3. *A flow $F \in \mathcal{F}$ is cardinally consistent if and only if it belongs to the image of the gradient.*

In light of 3, the data $\bar{Y}$ are rationalizable by an additive-equivariant utility if and only if $\bar{Y}$ is the discrete gradient of a utility function $u \in \mathcal{U}$.[36] We postpone discussion of the economic content of 3 until section 3.5.2.

[36]This highlights a recurring parallel between the graph theoretic methods employed here and the differential calculus; cardinal consistency is a discrete analogue of the requirement that a gradient vector field integrate to zero around every closed curve in a domain.

Fix an ordering of $\mathcal{V} = \{v_1, \ldots, v_K\}$. By minor abuse of notation we will write i for v_i, F_{ij} for $F_{v_i v_j}$ and so forth when no confusion will result. A flow is uniquely determined by its values on oriented edges $(i,j) \in \vec{\mathcal{E}}$ with $i < j$. Thus we identify $\mathcal{F}$ with $\mathbb{R}^{\mathcal{E}}$, with basis $\left\{\mathbb{1}_{(i,j)}\right\}_{\{(i,j)\in\vec{\mathcal{E}}: i<j\}}$.[37] Using this basis, we endow $\mathcal{F}$ with an inner product via:

$$\langle F, F' \rangle = \sum_{\{(i,j)\in\vec{\mathcal{E}}\,:\, i<j\}} F_{ij} F'_{ij}.$$

The **divergence** of a flow is the real valued function on vertices defined component-wise by:

$$(\operatorname{div} F)_i = \sum_{j \in N(i)} F_{ij},$$

where $N(i) \subseteq \mathcal{V}$ denotes the set of neighbors of v_i. In other words, the divergence computes the vector of outflows net inflows at each vertex. This defines a linear map $\operatorname{div} : \mathcal{F} \to \mathcal{U}$. When $\mathcal{U}$ carries its standard inner product, $-\operatorname{div}$ is the adjoint of the gradient.

Proposition 4. *For any experiment, the space of flows on $(\mathcal{V}, \mathcal{E})$ admits an orthogonal direct-sum decomposition as:*[38]

$$\mathcal{F} = \operatorname{im}(\operatorname{grad}) \oplus \ker(\operatorname{div}).$$

By 3, the image of the gradient consists of the cardinally consistent and hence additive-equivariant rationalizable flows. Call a flow $C \in \mathcal{F}$ a **perfect cycle** if it has vanishing divergence, and is supported on a single loop in $(\mathcal{V}, \mathcal{E})$.[39] The kernel of the divergence

[37] For example, let $(\mathcal{V}, \mathcal{E})$ denote the complete graph on three vertices. We identify a flow F with the vector (F_{12}, F_{13}, F_{23}). If F is a cyclic flow of α units from v_1 to v_2, v_2 to v_3, and v_3 to v_1, then $F = (\alpha, -\alpha, \alpha)$.

[38] In differentiable terms, 4 is the graph theoretic analogue of the Helmholtz decomposition of vector calculus.

[39] A loop in $(\mathcal{V}, \mathcal{E})$ is a connected subgraph such that every vertex is contained in precisely two edges. A flow is supported on a loop if takes the value zero on every edge that does not belong to the edge set of the loop.

is precisely the span of the perfect cycles in $\mathcal{F}$.[40] By 4, every flow may be uniquely written as a sum of two orthogonal terms: a cardinally consistent flow, and a 'purely inconsistent' flow, expressible as a sum of perfect cycles.[41] For a given $\bar{Y}$, we define its **best fit** cardinally consistent approximation, denoted $\hat{Y}$, as its projection onto im(grad). This may be computed by solving the following least squares program:

$$\min_{u \in \mathcal{U}} \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2. \tag{3.2}$$

The value of this problem, $\|\bar{Y} - \hat{Y}\|_2^2$, reflects the mean squared error of imposing the hypothesis of additive-equivariance. By the Pythagorean theorem (suppressing subscripts):

$$\|\bar{Y}\|^2 = \|\hat{Y}\|^2 + \|\bar{Y} - \hat{Y}\|^2,$$

and thus we obtain an analogue of the coefficient of determination (i.e. R^2) for (3.2):

$$R^2 = 1 - \frac{\|\bar{Y} - \hat{Y}\|^2}{\|\bar{Y}\|^2} = \frac{\|\hat{Y}\|^2}{\|\bar{Y}\|^2}. \tag{3.3}$$

By definition, $R^2 \in [0,1]$, and $R^2 = 1$ if and only if $\bar{Y}$ is rationalizable by an additive-equivariant utility. Roughly, R^2 reflects the proportion of the variation in the (average) intensity of preference across the different pairs in $\mathcal{E}$ that is able to be explained by additive-equivariance. Thus R^2 provides an alternative measurement of goodness-of-fit that normalizes for the *structure* of the underlying experiment. This may be of use when comparing goodness of fit across experiments.[42]

A Social Choice Interpretation

The problem of choosing a best-fit, cardinally consistent approximation to a given flow is equivalent to choosing a social utility or *scoring function* for the alternatives

[40] See Godsil and Royle [53] Corollary 14.2.3.

[41] The decomposition of the inconsistent component into a sum of perfect cycles will not be unique, see, e.g., Figure 3.2.

[42] Similar ideas are leveraged in [40] to obtain a measure of the incentive alignment for players in normal form games.

in $\mathcal{V}$. By 3, a general rule for assigning a best-fit rationalizable flow to a data set may be identified with a function $b : \mathcal{F}^N \to \text{im}(\text{grad})$. Any such function factors as $b = \text{grad} \circ \tilde{b}$, where $\tilde{b} : \mathcal{F}^N \to \mathcal{U}$ is a social choice scoring function for incomplete and cardinal data.[43] Thus 3 establishes the problem of choosing a notion of rationalizable best-fit is isomorphic to the problem of choosing how to aggregate incomplete, cardinal preference information via scoring rule.

It is known that the mean squared error minimizing flow $\hat{Y}$ arises as the gradient of the score vector obtained by applying a natural generalization of the Borda count to the individual-level data $\{Y^n\}$.[44] Suppose, momentarily, that $(\mathcal{V}, \mathcal{E})$ is a complete graph. The *cardinal Borda score* is defined as:

$$s_{\text{CB}}(v_i) = \frac{1}{K} \sum_{j \neq i} \bar{Y}_{ji} = \frac{1}{NK} \sum_{j \neq i} \sum_{n=1}^{N} Y_{ji}^n,$$

where the leading coefficient is positive and hence does not affect the ranking. This terminology is justified, as the (ordinal) Borda score is:

$$s_{\text{B}}(v_i) = \sum_{j \neq i} \sum_{n=1}^{N} \text{sign}(Y_{ji}^n).$$

Note Y_{ji}^n is positive if $v_i \succ^n v_j$, and negative if $v_j \succ^n v_i$. Thus s_{B} counts 'net votes for v_i over v_j' across the sample population, summed over all $j \neq i$. The cardinal Borda score simply allows the intensity of preference, reflected in the magnitudes of the Y_{ji}^n, to factor in.

For a general experiment $(\mathcal{V}, \mathcal{E})$, let A denote its adjacency matrix, and D the diagonal matrix with $D_{ii} = \deg(v_i)$. Recall the Laplacian matrix $L = D - A$. Then

[43] See [114] for a characterization of such aggregation rules for the ordinal case.

[44] The observations of this subsection are implicit in [113], and noted explicitly in, e.g., [71] in the context of aggregating online rankings.

the normal equations for (3.2) may be compactly written as:[45]

$$L\,u^* = -\mathrm{div}\,\bar{Y}. \tag{3.4}$$

As $(\mathcal{V}, \mathcal{E})$ is assumed connected, the kernel of the Laplacian is spanned by the vector $(1, \ldots, 1)$, and thus every solution u^* to (3.4) will be unique up to addition of a constant vector. For sake of determinacy, we focus then on the minimum norm solution, or equivalently require $\sum_{i=1}^K u_i^* = 0$.[46] Equation (3.4) says that any solution u^* is determined by a strong 'averaging' property. In particular, for any v_i:

$$u_i^* = \frac{1}{\deg(i)}\left[\sum_{j \in N(i)} u_j^* + \bar{Y}_{ji}\right], \tag{3.5}$$

so u_i^* is an unweighted average of the utility values of each neighboring v_j, plus the observed flow from each neighboring v_j to v_i. If $(\mathcal{V}, \mathcal{E})$ is complete, then the unique (zero-sum) utility vector satisfying (3.5) is the cardinal Borda score, as (3.5) becomes:

$$u_i^* = \frac{1}{K-1}\left[\sum_{j \neq i} u_j^* + \bar{Y}_{ji}\right].$$

As $\sum_{j \neq i} u_j^* = -u_i^*$, this simplifies to the cardinal Borda score:

$$u_i^* = \frac{1}{K}\sum_{j \neq i} \bar{Y}_{ji}.$$

Even when $(\mathcal{V}, \mathcal{E})$ is less-than-complete, solutions to (3.2) remain wholly determined by the averaging property (3.5) which, for complete experiments, characterizes the cardinal Borda score up to a constant. Thus our choice of the mean squared error

[45] This makes use of the identity $-\mathrm{div} \circ \mathrm{grad} = L$, see Godsil and Royle [53] Lemma 8.3.2, and the fact that $-\mathrm{div}$ is the adjoint of grad (4).

[46] In vector notation, this solution is given by:

$$u^* = -L^\dagger \mathrm{div}\,\bar{Y},$$

where $L^\dagger$ denotes the Moore-Penrose pseudoinverse of L.

minimizer $\hat{Y}$ as the best-fit for $\bar{Y}$ may alternatively be regarded as the result of a two-step process. We first form a social ranking for our sample population using the natural generalization of the cardinal Borda score to (possibly incomplete) experiments, then define the best-fit flow as the gradient of this score vector.

A Money Pump Metric for the Residual

In section 3.5.2, we used the mean squared error $\|\bar{Y} - \hat{Y}\|_2^2$ to quantify the goodness of fit for the hypothesis of additive-equivariance. In this section, we consider an alternative criterion: the L^1 norm of the residual vector. In particular, we provide a novel interpretation of $\|\bar{Y} - \hat{Y}\|_1$ as the natural analogue of the money pump metric considered in Echenique et al. [42] for our setting.[47]

By 4, $R = \bar{Y} - \hat{Y}$ belongs to the kernel of the divergence, and thus consists of a sum of perfect cycles. Suppose first, that R is itself a perfect cycle. Then, for some finite sequence $(v^0, v^1), (v^1, v^2), \ldots, (v^{L-1}, v^0) \in \vec{\mathcal{E}}$, we have $R_{v^l v^{l+1}} = \bar{c} \geq 0$. Here, we have used superscript indices to distinguish them from the enumeration used to define the basis for $\mathcal{F}$, which we will always denote with a subscript.[48] If we interpret this residual as data arising from a single representative agent, then for this agent, for all l:

$$\phi(\bar{c}, v^l) \sim v^{l+1},$$

where the superscripts are understood mod-L. This implies that for every l, the agent would be willing to trade v^l, plus up to $\bar{c}$ units of numeraire, for v^{l+1}. A savvy arbitrageur could exploit such an agent as a 'numeraire pump,' and extract $\|R\|_1 = \bar{c}\, L$

[47] Roughly speaking, the money pump metric reflects the amount of money one could extract from a consumer who violates the generalized axiom of revealed preference. See Echenique et al. [42] for details.

[48] Thus, in particular, $R_{v^l v^{l+1}}$ here corresponds to the flow along a particular edge in this cycle.

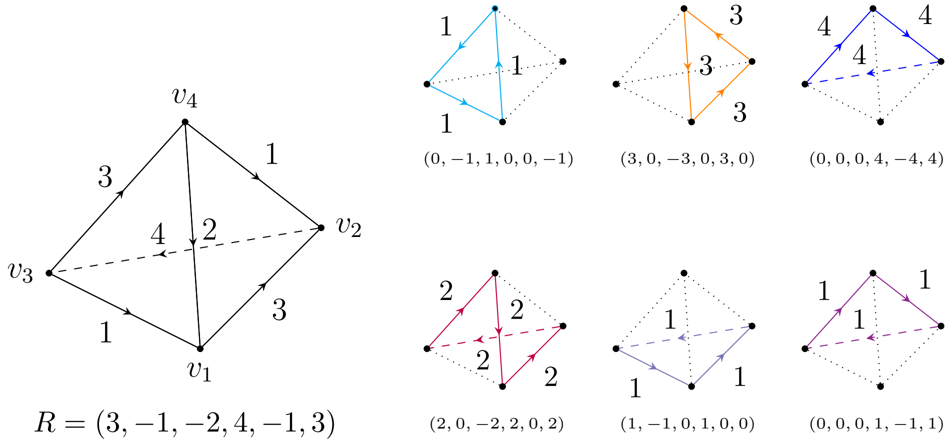

Figure 3.2: Residual flows. A residual flow $R = (R_{12}, R_{13}, R_{14}, R_{23}, R_{24}, R_{34})$, belonging to the kernel of the divergence, along with two decompositions into sums of perfect cycles. The lower bound of $\|R\|_1 = 14$ is attained by the sum of the money pump values of the bottom (though not the top) decomposition.

units of numeraire from the agent via a cyclic sequence of trades. Thus for residuals consisting of a single perfect cycle, the L^1 norm is precisely a numeraire-valued analogue of the money pump metric. In light of this, we formally define the **money pump** value of a pure cycle R as:

$$MP(R) = \|R\|_1 = \bar{c}\,L.$$

It is natural to seek to extend MP from single pure cycles to general residuals R linearly, by decomposing R into a sum of pure cycles and summing the money pump values of the cycles in this decomposition. However, for any such R, the decomposition into pure cycles will be non-unique; moreover the sum of the money pump values of different decompositions of the same residual will generally differ, see Figure 3.2. Instead, we consider the most conservative extension.

Let $\mathfrak{C} \subsetneq \mathcal{F}$ denote the set of pure cycles. For any $R \in \ker(\operatorname{div})$ let $\mathfrak{D}(R)$ denote the collection of all finite decompositions of R into pure cycles.[49] We extend $MP : \mathfrak{C} \to \mathbb{R}$ to a function $MP^* : \ker(\operatorname{div}) \to \mathbb{R}$ via:

$$MP^*(R) = \inf_{\{C_1,\dots,C_M\} \in \mathfrak{D}(R)} \sum_{m=1}^{M} MP(C_m).$$

In other words, MP^* attributes as little inconsistency to the agent as possible, by taking an infimum across all finite decompositions of the residual. In spite of its definition as a value function, our next result asserts that MP^* is in fact simply the L^1 norm on $\ker(\operatorname{div})$.

Proposition 5. *For all $R \in \ker(\operatorname{div})$, the money pump value of R is equal to its L^1 norm:*

$$MP^*(R) = \|R\|_1.$$

Moreover, the infimum over $\mathfrak{D}(R)$ is always attained.

Given its economic interpretation for elements of $\ker(\operatorname{div})$, it is tempting to consider the L^1 analogue of (3.2),

$$\min_{u \in \mathcal{U}} \left\| (\operatorname{grad} u) - \bar{Y} \right\|_1, \tag{3.2$'$}$$

and to interpret the value of this linear program as the magnitude of the inconsistency. However, (3.2′) lacks a great deal of the analytic structure of (3.2). Unlike (3.2), it is not a strictly convex program and hence can admit multiple distinct solutions. For example, if $\bar{Y}$ is a perfect cycle with unit flow over the circle graph on n vertices, the set of minimizing utilities for (3.2′) form an $(n-1)$-dimensional polytope, see [91].[50] Moreover, while the residual from (3.2) always belongs to $\ker(\operatorname{div})$, this need not be true for (3.2′). Thus the interpretation provided by 5 does not generally apply

[49] That is, those collections $\{C_1, \dots, C_M\} \subseteq \mathfrak{C}$ such that $\sum_m C_m = R$.
[50] In contrast, for any n, the minimizers for (3.2) are precisely the constant functions.

to residuals from (3.2′). As such, while the L^1 norm on ker(div) admits an economic interpretation, the L^2 theory appears better suited to quantifying inconsistency generally.

3.5.3 Shape Constraints

In most applications, it is of interest to test not only whether the data are rationalizable by an additive-equivariant utility, but also by one that possesses additional properties, such as quasi-concavity, monotonicity, homogeneity and so forth. This can be simply and tractably incorporated by considering constraint sets for (3.2).

Example 20 (CES Utility)**.** Let $X = \mathbb{R}^L_+ \setminus \{0\}$, and $\phi(\alpha, x) = e^{\alpha}\, x$. While CES utility functions are not additive-equivariant, it is straightforward to verify their natural logarithm is:

$$U(x) = \frac{1}{\rho} \ln\left(\sum_{l=1}^{L} x_l^{\rho}\right),$$

where $\rho \in (-\infty, 1]$. The following constrained analogue of (3.2) computes the MSE for CES preferences:

$$\begin{aligned} \min_{u,\rho} \quad & \left\|(\operatorname{grad} u) - \bar{Y}\right\|_2^2 \\ \text{subject to} \quad & u_i = \frac{1}{\rho} \ln\left(\sum_{l=1}^{L} v_{li}^{\rho}\right) \\ & \rho \leq 1, \end{aligned} \tag{3.6}$$

where $v_{i,l}$ denotes the l-th component of v_i.

When models are defined by a collection of axioms or properties, there is often a natural correspondence between the constraints constraints imposed and the restrictions characterizing the model. Examining which constraints bind at the solution to the MSE minimization problem then provides insight into the cost, in model fit terms, of imposing specific assumptions.

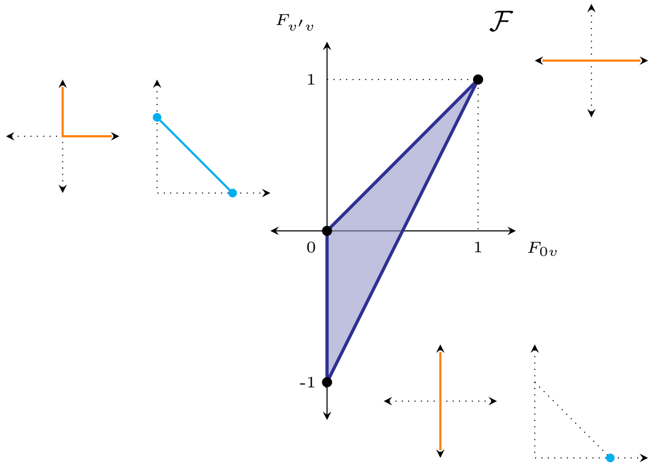

(a) The extremal MEU preferences.

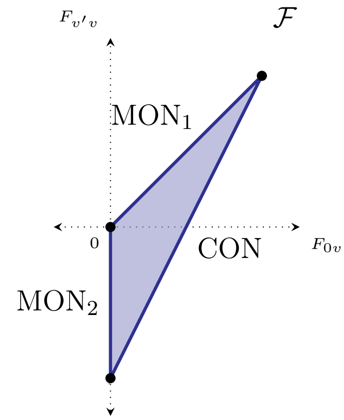

(b) The binding constraints at each face.

Figure 3.3: The MEU-rationationalizable flows (violet triangle) arising from the experiment $\mathcal{E} = \{\{0, v\}, \{v, v'\}\}$ where $v = (0, 1)$, $v' = (1, 0)$. For each vertex of the triangle, the corresponding MEU preference (indifference curve in orange) and set of priors (cyan) are shown. Each face of the triangle corresponds to a binding constraint: the top and left faces to monotonicity of consumption in state one (resp. two), and the bottom to convexity of the preference.

Example 21 (Risk-Neutral Maxmin Expected Utility)**.** Suppose we wish to test whether a subject's preferences over monetary acts (i.e. portfolios of Arrow securities) in a simple two-state model are consistent with a risk-neutral maxmin expected utility (MEU) function:

$$U(x) = \min_{\pi \in C} \mathbb{E}_\pi(x),$$

where C is a compact, convex set of priors over states of the world. Let $S = \{s_1, s_2\}$ denote the set of states, and $\mathbb{R}^S$ the consumption space. Any risk-neutral MEU preference is invariant under the addition of some quantity of bond, thus we let $\phi(\alpha, x) = x + \alpha \mathbb{1}_S$.

Let $v = (0,1)$ and $v' = (1,0)$ denote the portfolios consisting of a single unit of each Arrow security, and 0 the empty portfolio, and consider the experiment eliciting compensation differences over two pairs $\mathcal{E} = \big\{\{0,v\},\{v,v'\}\big\}$. If a utility $U : \mathbb{R}^S \to \mathbb{R}$ is increasing, concave, positively homogeneous, and additive equivariant, then there exists a compact, convex set $C \subseteq \Delta(S)$ such that $U = \min_{\pi \in C} \mathbb{E}_\pi(x)$.[51] Every data set arising from $\mathcal{E}$ is rationalizable by an additive-equivariant and positively homogeneous utility. As such, the only falsifiable properties for this experiment are monotonicity and convexity of the preference, i.e. ambiguity aversion.

Figure 3.3 plots those data vectors rationalizable by a risk-neutral MEU preference; each corner of the triangle corresponds to the data set that is (uniquely) rationalizable by the preference corresponding to one of the three extremal sets of priors.[52] To evaluate the goodness of fit for this model, we instead project the data vector onto this triangle. While the squared norm of the residual still reflects the goodness of fit, the image of the projection provides additional, granular insight into *which* properties of the model are the binding constraints to the fitting problem. The top (resp. left) face corresponds to those preferences that are only weakly increasing in consumption in the first (resp. second) state of the world. For these preferences, the monotonicity constraint binds. The bottom-right face corresponds to the set of subjective expected utility (SEU) preferences, which are ambiguity-neutral. These are the preferences for which ambiguity aversion is the binding constraint.

[51] See, e.g., [88] H.1.3 Lemma 2.

[52] The fact that each vector in the triangle is rationalizable by a *unique* MEU preference is a consequence of the fact that $|S| = 2$, and hence every closed convex set of priors is parametric. While the set of MEU rationalizable vectors remains polyhedral for larger experiments and with $|S| > 2$, the sets of priors associated with a given rationalizable vector will generally only be set-identified. See 24 for more details.

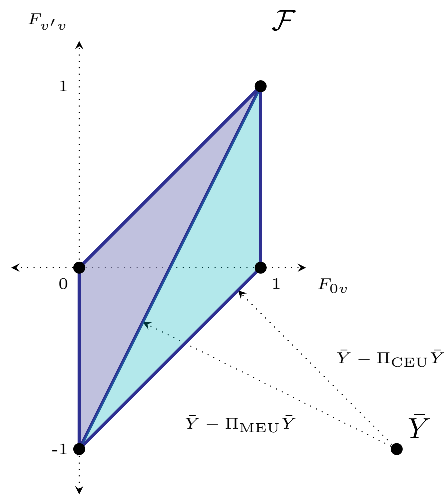

Figure 3.4: The set of MEU-rationalizable (violet) and CEU-rationalizable (violet or aquamarine) vectors for $\mathcal{E}$. Letting Π_{MEU} and Π_{CEU} denote the respective projections onto these sets, the quantity $\left\|\bar{Y} - \Pi_{\mathrm{MEU}}\bar{Y}\right\|_2^2 - \left\|\bar{Y} - \Pi_{\mathrm{CEU}}\bar{Y}\right\|_2^2$ reflects the *shadow price*, in mean squared error terms, of imposing ambiguity aversion, conditional upon requiring monotonicity, translation invariance, and homotheticity.

Suppose the data vector $\bar{Y}$ projects onto the relative interior of the lower-right face of the MEU-rationalizable triangle. As such, its best-fit is a risk-neutral SEU preference. At this projection the monotonicity constraints are slack; thus by considering the difference in mean squared error resulting from imposing both monotonicity and ambiguity aversion versus monotonicity alone, one obtains a measure of the *shadow price*, in model fit terms, of imposing ambiguity aversion. Figure 3.4 plots those vectors which admit additive-equivariant, positively homogeneous, and monotone rationalizations. Here, these are precisely the vectors rationalizable by a risk-neutral Choquet expected utility (CEU) preference.[53] This highlights an advan-

[53]More generally, this corresponds to the class of invariant biseparable preferences of [51]; see also [33] for an elegant representation theorem for such preferences. However, for this

tage of our theory: not only is it capable of measuring how well a particular models fit the data, it is capable of far more granular insights, including quantifying the severity of violations of individual axioms and properties.

Constraints for Non-parametric Models

Even when a class of preferences is non-parametric, it can often be tractably encoded into a constraint set:

$$\mathcal{K} = \{u \in \mathcal{U} : \exists U : X \to \mathbb{R} \text{ possessing the desired structure s.t. } U(v_i) = u_i\}.$$

Formally, we say a set $\mathcal{K} \subseteq \mathcal{U}$ defines a set of **shape constraints** if (i) $\mathcal{K}$ is convex, and (ii) $\mathcal{K} + \operatorname{span}\{(1, \ldots, 1)\}$ is closed.[54] Rather than solve an unconstrained least squares problem as in (3.2), one projects instead onto $\operatorname{grad}(\mathcal{K})$:

$$\min_{u \in \mathcal{K}} \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2, \tag{3.7}$$

which is closed and convex by our definition of shape constraints. In section C.0.5, we provide a number of explicit derivations of such constraint sets for a variety of ambiguity preferences, as well as formal proofs for some of the examples in this section.

By 4, the residual from (3.7) decomposes into two orthogonal components: one with vanishing divergence and one that is cardinally consistent. Let $\Pi_{\operatorname{grad}(\mathcal{K})}$ denote the projection onto $\operatorname{grad}(\mathcal{K})$, and recall that $\hat{Y}$ denotes the projection of $\bar{Y}$ onto the

simple experiment, the testable implications of this class of preferences coincide with those of the Choquet model of [101].

[54]This closure condition is innocuous and satisfied under all economically interesting cases we are aware of. It serves only to guarantee that the image $\operatorname{grad}(\mathcal{K})$ is a closed convex subset of $\mathcal{F}$ for any (connected) experiment. For such an experiment, $\ker(\operatorname{grad})$ is just the diagonal of $\mathcal{U}$ hence this condition is just a requirement $\mathcal{K} + \ker(\operatorname{grad})$ be closed, which is holds if and only if $\operatorname{grad}(\mathcal{K})$ is closed; see, e.g., [67], Lemma 17.H. It is necessarily satisfied if, for example, $\mathcal{K}$ is polyhedral, which is often the case.

cardinally consistent subspace. By the Pythagorean theorem:

$$\left\|\bar{Y} - \Pi_{\text{grad}(\mathcal{K})}\bar{Y}\right\|_2^2 = \left\|\bar{Y} - \hat{Y}\right\|_2^2 + \|\hat{Y} - \Pi_{\text{grad}(\mathcal{K})}\hat{Y}\|_2^2. \tag{3.8}$$

The magnitude of the first component of the right-hand side still reflects how well additive-equivariance is borne out in the data; the latter captures the shadow price, in mean squared error units, of imposing the shape constraints $\mathcal{K}$ beyond additive-equivariance.

When a family of models are defined on a common domain X and are additively-equivariant relative to the same virtual numeraire ϕ, equation (3.8) allows for a simple, tractable approach to model selection. Given a family of models, and hence shape constraints $\{\mathcal{K}_m\}_{m=1}^M$, computing $\|\hat{Y} - \Pi_{\text{grad}(\mathcal{K}_m)}\hat{Y}\|_2^2$ for each K_m yields an model-specific measure of how consistent are the data with respect to model m. By comparing these values, one obtains a ranking of these models for the data.

Example 22 (Quasilinearity Revisted)**.** Reconsider 14, but suppose now we wish to test whether the data are consistent with a utility:

$$U(x, y) = v(y) + x,$$

where v is, in addition, increasing and concave. Let $\mathcal{K}_{QIC}$ denote the set of vectors in $\mathcal{U}$ that are restrictions of quasilinear (in the first variable), increasing, and concave functions. For a general experiment $(\mathcal{V}, \mathcal{E})$, evaluating (3.7) with $\mathcal{K} = \mathcal{K}_{QIC}$ is equivalent to solving:

$$\begin{aligned} \min_{u \in \mathcal{U}} \quad & \left\|(\text{grad}\, u) - \bar{Y}\right\|_2^2 \\ \text{subject to} \quad & u_i = \langle \pi_i, v_i \rangle + \gamma_i && \forall i = 1, \ldots, K \\ & \langle \pi_i, v_i \rangle + \gamma_i \le \langle \pi_j, v_i \rangle + \gamma_j && \forall i, j = 1, \ldots, K \\ & \pi_i^1 = 1 && \forall i = 1, \ldots, K \\ & \pi_i \ge 0 && \forall i = 1, \ldots, K \end{aligned} \tag{3.9}$$

for $u \in \mathbb{R}^K$ and, for all $i = 1, \ldots, K$, $\pi_i \in \mathbb{R}^2$, $\gamma_i \in \mathbb{R}$ (where π_i^1 denotes the first component of π_i). Each feasible vector $(u, \pi_1, \ldots, \pi_K, \gamma)$ defines a quasilinear increasing, and concave function:

$$\tilde{U}(x) = \min_{i \in \{1,\ldots,K\}} \gamma_i + \langle x, \pi_i \rangle$$

whose restriction to $\mathcal{V}$ is u. The variables π_i act as supergradients at each v_i; the γ_i capture the vertical intercepts of the supporting hyperplanes defined by the π_i. Conversely, given a quasilinear, increasing, and concave function U, an arbitrary selection of supporting hyperplane for the subgraph of U at each v_i yields a choice of π_i and γ_i. These, along with the vector of utilities $u = U|_{\mathcal{V}}$, define a feasible element of the constraint set in (3.9).[55]

Calibration and Identification

In the experiment considered in 21, the set of rationalizable flows were in one-to-one correspondence with the set of risk-neutral MEU preferences. This was a consequence of considering a state space S with only two elements. Perfect identification will generally be obtainable, for a rich enough experiment, whenever the class of models is parametric. For such models, not only does (3.7) allow for a means of evaluating goodness of fit, it additionally allows for *calibration*. Whenever there is a one-to-one correspondence between rationalizable vectors and preferences, the best-fit exercise yields *point estimates* of the model's parameters.

Example 23. (Cobb-Douglas Preferences) Let $X = \mathbb{R}^L_{++}$ and $\phi(\alpha, x) = e^{\alpha} x$. Though Cobb Douglas preferences satisfy (N.1) - (N.3), the standard representation:

$$U(x) = \prod_{i=1}^{L} x_i^{\beta_i} \tag{3.10}$$

[55] The system of constraints in (3.9) may be viewed as essentially a system of Afriat inequalities, but where prices are unknown.

where $\beta \geq 0$ and $\sum_i \beta_i = 1$, is not not additive-equivariant. Define the homeomorphism $H : X \to \mathbb{R}^L$ via $H(x) = \big(\ln x_1, \ldots, \ln x_L\big)$. Then:

$$\ln U(x) = \langle \beta, H(x) \rangle,$$

which is additive-equivariant under the induced action under H on $\mathbb{R}_L$ (here, simply $\phi(\alpha, \tilde{x}) = \tilde{x} + \alpha(1, \ldots, 1)$).[56] Thus it suffices to instead treat $\mathbb{R}^L$ as the consumption space, rather than $\mathbb{R}^L_{++}$. Computing (3.7) for Cobb-Douglas preferences is then equivalent to evaluating the following quadratic program:

$$\begin{aligned} \min_{u \in \mathcal{U}} \quad & \big\|(\operatorname{grad} u) - \bar{Y}\big\|_2^2 \\ \text{subject to} \quad & u_i = \big\langle \beta, H(v_i) \big\rangle \quad \forall i = 1, \ldots, K \\ & \beta \geq 0, \end{aligned} \tag{3.11}$$

where $\beta \in \mathbb{R}^L$.[57] Clearly each feasible vector (u, β) for (3.11) corresponds to a unique Cobb-Douglas preference on $\mathbb{R}^L_{++}$ and vice versa. Moreover, when $\mathcal{V}$ is rich enough, there is a one-to-one correspondence between flows in $\operatorname{grad}(\mathcal{K}_{\text{C-D}})$ and feasible vectors (u, β) to (3.11).

When models are non-parametric, it will be impossible for a finite experiment to distinguish between some pairs of preferences. In such cases, each rationalizable flow will correspond to a *set* of preferences. In many cases, the set of preferences consistent with given flow will be small, in a formal sense, and will often encode additional further testable implications of behavior.

[56]Given an action ϕ of $\mathbb{R}_+$ on X and a homeomorphism $h : X \to Y$, the 'induced action' on Y refers to the map $\tilde{\phi} : \mathbb{R}_+ \times Y \to Y$ defined by:

$$\tilde{\phi}(\alpha, y) = h \circ \phi\big(\alpha, h^{-1}(y)\big).$$

[57]Note that additive-equivariance, in conjunction with the other constraints, imply the standard normalization $\langle \beta, \mathbb{1}_L \rangle = 1$.

Example 24 (Risk-neutral Maxmin Expected Utility Revisited)**.** Suppose we again wish to test whether a subject's preferences over monetary acts are consistent with a risk-neutral MEU preference (see also 21). However, we now consider a richer state space $S = \{s_1, s_2, s_3\}$. As a consequence, a general compact, convex set of priors is no longer describable by a finite set of parameters. For any experiment $(\mathcal{V}, \mathcal{E})$, evaluating (3.7) for risk-neutral MEU preferences is equivalent the following quadratic program:

$$\begin{aligned} \min_{u \in \mathcal{U}} \quad & \left\|(\operatorname{grad} u) - \bar{Y}\right\|_2^2 \\ \text{subject to} \quad & u_i = \langle \pi_i, v_i \rangle && \forall i = 1, \ldots, K \\ & \langle \pi_i, v_i \rangle \leq \langle \pi_j, v_i \rangle && \forall i, j = 1, \ldots, K \\ & \langle \pi_i, \mathbb{1}_S \rangle = 1 && \forall i = 1, \ldots, K \\ & \pi_i \geq 0 && \forall i = 1, \ldots, K, \end{aligned} \tag{3.12}$$

for $\pi_1 \ldots, \pi_K \in \mathbb{R}^S$.

Consider any feasible solution $(u, \pi_1, \ldots, \pi_K)$ to (3.12). The vector u, coupled with $\mathcal{V}$, defines a family of hyperplanes $H_{v_i,u_i} = \{x \in \mathbb{R}^S : \langle v_i, x \rangle = u_i\}$. Let $\tilde{H}_{v_i,u_i}$ denote the restrictions of these hyperplanes to the affine hull of $\Delta(S)$. The first and second constraints in (3.12) imply that each $\tilde{H}_{v_i,u_i}$ supports $\operatorname{co}\{\pi_1, \ldots, \pi_K\}$ at each π_i, and hence $\operatorname{co}\{\pi_1, \ldots, \pi_K\}$ is a set of priors consistent with the collection of utilities u. However, at any given solution to (3.12), the priors $(\pi_1, \ldots, \pi_K)$ will be non-unique, as there are will generally be infinitely many sets consistent with the $\tilde{H}_{v_i,u_i}$. Define:

$$\bar{C} = \left(\bigcap_{i=1}^{K} \tilde{H}^+_{v_i,u_i} \right) \bigcap \Delta(S),$$

where $\tilde{H}^+_{v_i,u_i}$ denotes the i-th upper half-space. A set of priors $C \subseteq \Delta(S)$ is consistent with u if and only if: (i) $C \subseteq \bar{C}$, and (ii) each facet of $\bar{C}$ contains some extremal point of C.[58] Given such a set, choosing an extremal point $\hat{\pi}_i \in \operatorname{ext}(C)$ from each $\tilde{H}_{v_i,u_i}$

[58] Recall that a facet of a polyhedron is a codimension-1 face.

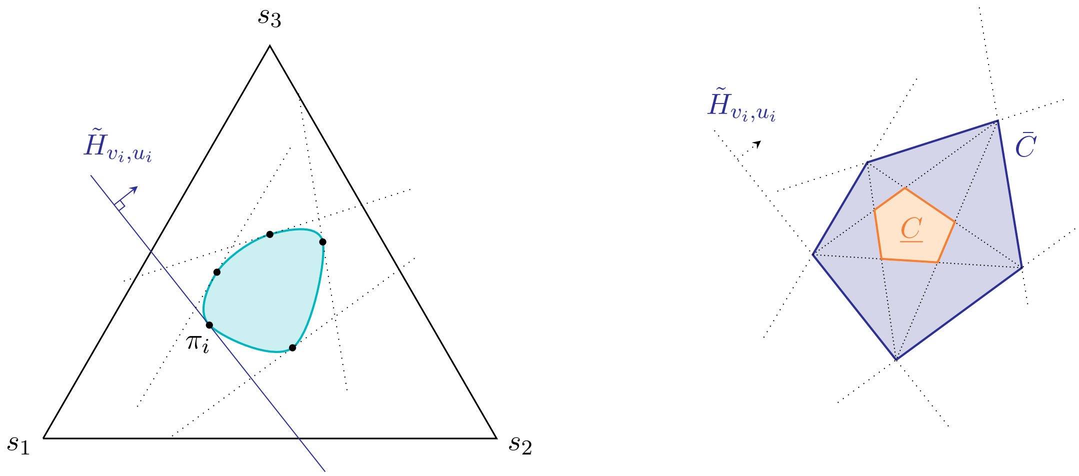

(a) The belief set C^* of a risk-neutral MEU preference. The vector $(u, \pi_1, \ldots, \pi_K)$ is a solution to (3.12), as for each $v_i \in \mathcal{V}$, the hyperplane $\tilde{H}_{v_i,u_i}$ supports C^* at π_i.

(b) Every MEU-rationalizable vector defines a belief polytope $\bar{C} = \cap_i \tilde{H}^+_{v_i,u_i}$. A belief set $C \subseteq \bar{C}$ rationalizes the data if and only if each facet of $\bar{C}$ contains some extremal point of C.

Figure 3.5: An experiment with $\mathcal{V} = \{v_1, \ldots, v_5\}$ and a rationalizing utility vector u define a system of hyperplanes on the belief simplex. From these hyperplanes, we obtain upper (purple) and lower (orange) envelope belief sets. Every belief set rationalizing the data is contained within the purple set, and contains the orange.

yields a tuple of priors $(\hat{\pi}_1, \ldots, \hat{\pi}_K)$ such that $(u, \hat{\pi}_1, \ldots, \hat{\pi}_K)$ is a feasible solution to (3.12). It follows that $\bar{C}$ is the largest set of priors consistent with u. This allows for bounding the subjective beliefs held by an individual, even absent full identification: if $\pi \notin \bar{C}$, it is not held by *any* risk-neutral MEU preference rationalizing the data.[59] Figure 3.5 shows the system of supporting hyperplanes for a belief set C^* arising from an experiment. Here, the upper envelope $\bar{C}$ corresponds to the intersection of the five

[59]Furthermore, for a fixed, feasible u, the collection of sets of priors $\mathcal{P}_u$ consistent with u is small. It is straightforward to show that $\mathcal{P}_U$ is nowhere dense in the space of all compact, convex subsets of $\Delta(S)$ endowed with the Hausdorff topology. Thus while a given solution to (3.12) will generally be consistent with many sets of priors, the collection of such sets remains topologically negligible.

half-spaces. It also depicts the set of priors, $\underline{C}$, which are contained in *every* set of priors consistent with u.[60] Thus, while C^* may be unknown, the vector u, along with $\mathcal{V}$, yield tight bounds.

These observations yield further economic predictions. For example, subjects with MEU preferences will engage in purely speculative trade if and only if they hold no common priors (Billot et al. 20, see also Rigotti et al. 95). Thus observing the $\bar{C}$ sets of two agents are disjoint implies predictions about their trade behavior. Similarly, in an economy of MEU agents without aggregate uncertainty, the Pareto frontier precisely corresponds to the set of full-insruance allocations if and only if the agents share at least one common prior.[61] Thus observing, for example, that the $\underline{C}$ sets of a population have non-empty intersection, not only yields welfare implications but in fact identifies the entire Pareto frontier, even when the individual preferences themselves may not be identified.

Remark 4. It is straightforward to extend 21 and 24 to more general risk attitudes. Suppose instead the experimenter wishes to test whether or not preferences over monetary acts are representable by a utility of the form:

$$U(x) = \min_{\pi \in C} \mathbb{E}_\pi \big[\tilde{u}(x)\big], \tag{3.13}$$

where $C \subseteq \Delta(S)$ and $\tilde{u}$ is a known continuous, increasing, and unbounded above Bernoulli utility $\tilde{u} : \mathbb{R} \to \mathbb{R}$.[62] Such a $\tilde{u}$ could be chosen, for example, on the basis of theoretical considerations, or first-stage non-parametric estimation (see 25 for such an estimator).

[60] Note, however, that generally $\underline{C}$ will not itself be consistent with u.
[61] See [20] Theorem 1.
[62] It is straightforward to adapt to the case where $\tilde{u} : \mathbb{R}_+ \to \mathbb{R}$.

Given such a $\tilde{u}$, let $\phi_{\tilde{u}}$ denote the action on $\mathbb{R}^S$ defined component-wise via $x_s \mapsto \tilde{u}^{-1}\big(\tilde{u}(x_s) + \alpha\big)$. For any x, $\phi_{\tilde{u}}(\alpha, x)$ is the monetary act which yields precisely α additional utility in each state. For compensation differences data measured using $\phi_{\tilde{u}}$, one can test the shape constraints corresponding to (3.13) exactly as in (3.12), but instead replacing each v_i with its vector of utilities under $\tilde{u}$. For simplicity, in section C.0.5 we present shape constraint characterizations of various utility functionals on $\mathbb{R}^S$ for the risk-neutral case (i.e. where $\tilde{u}$ is identity) with the understanding that all characterizations may be adapted in this manner to other choices of risk attitude.

3.5.4 Misspecification

In some cases, the specific invariance a model satisfies may vary across the preferences of the model itself. In this section we show that it is often possible to estimate the correct choice of numeraire under which a preference is invariant, even when data is elicited using a choice of numeraire for which the true preferences do *not* satisfy (N.1).

Example 25 (Non-parametric Estimation of EU Preferences)**.** Let S be a finite set, and let $\pi \in \Delta(S)$ denote some fixed, objectively known, and non-degenerate probability distribution over S. Let $X = \mathbb{R}^S_+$, identified with the set of monetary lotteries paying off x_s with probability π_s. Let $\phi(\alpha, x) = x + \alpha \mathbb{1}_S$, but suppose the subject actually has risk-averse expected utility preferences represented by:

$$U(x) = \sum_{s \in S} \pi_s \tilde{u}(x_s) = \mathbb{E}_\pi\big[\tilde{u}(x)\big],$$

for some increasing, unbounded, and concave $\tilde{u}$. While generally not ϕ-invariant, such a preference satisfies (N.1) for the action $\phi_{\tilde{u}}(\alpha, x)$, defined component-wise via $\tilde{u}^{-1}\big(\tilde{u}(x_s)+\alpha\big)$. Indeed, U is additive equivariant for $\phi_{\tilde{u}}$, hence the ij-th compensation difference, measured in $\phi_{\tilde{u}}$ numeraire, will equal $\mathbb{E}_\pi\big[\tilde{u}(v_j) - \tilde{u}(v_i)\big]$.

Despite the failure of (N.1) owing to misspecification (i.e. the use of ϕ rather than $\phi_{\tilde{u}}$ to elicit compensation differences), the subject's true preferences satisfy (N.2) and (N.3), hence Theorem 5 implies that for all observed α_{ij} (in ϕ numeraire):

$$v_i + \alpha_{ij}\mathbb{1}_S \sim v_j.$$

This implies the compensation difference between v_i and v_j under $\phi_{\tilde{u}}$ must also be equal to $\mathbb{E}_\pi\big[\tilde{u}(v_i + \alpha_{ij}\mathbb{1}_S) - \tilde{u}(v_i)\big]$, where α_{ij} is the observed compensation difference under ϕ.

In light of this, subject to monotonicity and concavity constraints, we choose an estimator $\hat{u}$ of $\tilde{u}$ that seeks to make $\mathbb{E}_\pi\big[\hat{u}(v_j) - \hat{u}(v_i)\big]$ as close as possible to $\mathbb{E}_\pi\big[\hat{u}(v_i + \alpha_{ij}\mathbb{1}_S) - \hat{u}(v_i)\big]$ for all (i, j).[63] This ensures that both the resulting 'transformed' compensation differences belong to the image of the gradient, and that they arise from a risk-averse expected utility preference.

Formally, let $\Theta = \big\{(s, i, j) \in S \times \{1, \ldots, K\}^2 : \{i, j\} \in \mathcal{E} \text{ or } i = j\big\}$, and for each $\{i, j\} \in \mathcal{E}$, define $v_{sij} = v_{si} + \alpha_{ij}$, and let $v_{sii} = v_{si}$. Let $\bar{\theta} = \arg\max_{\theta \in \Theta} v_\theta$, and consider the following quadratic program:

[63] Or, equivalently, $\mathbb{E}_\pi\big[\hat{u}(v_j)\big]$ as close as possible to $\mathbb{E}_\pi\big[\hat{u}(v_i + \alpha_{ij}\mathbb{1}_S)\big]$ for all (i, j).

$$
\begin{aligned}
\min_{u,\beta,\gamma} \quad & \left\|(\operatorname{grad} u) - \tilde{Y}\right\|_2^2 \\
\text{subject to} \quad & u_i = \sum_{s\in S} \pi_s\big(\gamma_{sii} + \beta_{sii} v_{sii}\big) && \forall i = 1,\dots,K \\
& \tilde{Y}_{ij} = \sum_{s\in S} \pi_s\big[(\gamma_{sij} + \beta_{sij} v_{sij}) - (\gamma_{sii} + \beta_{sii} v_{sii})\big] && \forall (s,i,j) \in \Theta \\
& \gamma_\theta + \beta_\theta v_\theta \le \gamma_{\theta'} + \beta_{\theta'} v_\theta && \forall \theta, \theta' \in \Theta \\
& \beta_\theta \ge 0 && \forall \theta \in \Theta \\
& \beta_{\bar{\theta}} = 1 \\
& \gamma_{\bar{\theta}} = 0.
\end{aligned}
\tag{3.14}
$$

where $u \in \mathbb{R}^K$, and $\beta, \gamma \in \mathbb{R}^\Theta$. A solution to this quadratic program yields a non-parametric estimator $\hat{u}$ of the true, unobserved $\tilde{u}$:

$$
\hat{u}(x) = \min_{\theta\in\Theta} \gamma_\theta + \beta_\theta x. \tag{3.15}
$$

For any increasing, concave, and unbounded above $\tilde{u}$, there is a $\hat{u}$ for which the value of (3.14) is zero.[64] Thus there exists some feasible solution to (3.14) which, for the experiment, is indistinguishable from the true $\tilde{u}$. Similarly, every feasible solution to (3.14) corresponds to an expected utility preference whose Bernoulli utility is increasing, concave, and unbounded above via (3.15).

Remark 5. This approach can be straightforwardly adapted to construct two-step experimental tests for a variety of other preferences, including non-expected utility theories such as cumulative prospect theory (Tversky and Kahneman 111) or popular parameterizations of disappointment aversion (Gul 57) such as in [98]. In the first stage, one estimates $\tilde{u}$ by varying payoffs under fixed, objectively known odds. In the

[64]This may be obtained by making a choice of tangent line to $\tilde{u}$ at each v_θ (hence obtaining γ_θ and β_θ) and defining $\hat{u}$ via (3.15).

second, one uses $\phi_{\hat{u}}$ while varying probabilities. Similar approaches allow for tests of additively separable models of dynamic preference à la [74] or quasi-hyperbolic preferences (e.g. Laibson 77). See also 4 for applications to translation-invariant ambiguity preferences.

3.6 Statistical Testing

In this section, we consider a stochastic analogue of our regression framework. We are interested in obtaining explicit hypothesis tests of rationalizability by various non-parametric models. Under our interpretation of a data set $\{Y^n\}$ as sampled from some population, we interpret such tests as of the hypothesis that a population of agents has preferences that are, in expectation, representable by some additive-equivariant utility U satisfying some collection of shape constraints. Formally, we assume an underlying linear model where for each $(i, j) \in \vec{\mathcal{E}}$, we observe the underlying 'true' population compensation difference Y^0_{ij}, polluted by a mean-zero, individual-specific shock.

> **Data Generating Process**: For all $\{x, y\} \in \mathcal{E}$ there is fixed, non-stochastic compensation difference $Y^0_{xy} = -Y^0_{yx}$. The data $\{Y^n\}_{n=1}^N$ is a random sample of N independent draws of the random flow Y, where for each $(i, j) \in \vec{\mathcal{E}}$ with $i < j$:
>
> $$Y_{ij} = Y^0_{ij} + \epsilon_{ij},$$
>
> where (i) $\mathbb{E}(\epsilon_{ij}) = 0$, and (ii) $\mathrm{Var}(\epsilon_{ij}) < +\infty$.[65]

In particular, we do not assume the ϵ shocks are uncorrelated or identically distributed across differing pairs in $\mathcal{E}$. This flexibility allows for a wide range of interpretations, ranging from models in which subjects compute compensation differences

[65] And hence for all $(i, j) \in \vec{\mathcal{E}}$ with $i > j$, $Y_{ij} = -Y_{ji}$.

from random utilities, to models in which shocks instead emerge due to idiosyncratic, pair-specific measurement or rounding errors.

We wish to test whether the vector of population compensation differences Y^0 arises from some some additive-equivariant function U, satisfying some collection of shape constraints $\mathcal{K}$. Phrased formally, we are interested in the hypothesis:

$$H_0 : Y^0 \in \operatorname{grad}(\mathcal{K}), \qquad H_1 : Y^0 \notin \operatorname{grad}(\mathcal{K}), \tag{3.16}$$

where $\mathcal{K} \subseteq \mathcal{U}$ is a set of shape constraints capturing the desired properties of U.[66] Following (3.7), the null and alternative hypotheses in (3.16) can be rephrased, up to a monotone transformation, as:

$$H_0 : \min_{u \in \mathcal{K}} \left\| (\operatorname{grad} u) - Y^0 \right\|_2 = 0, \qquad H_1 : \min_{u \in \mathcal{K}} \left\| (\operatorname{grad} u) - Y^0 \right\|_2 > 0. \tag{3.17}$$

A natural sample analogue of the objective function (3.17) is:

$$\psi(\bar{Y}) = \min_{u \in \mathcal{K}} \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2 = \min_{\hat{Y} \in \operatorname{grad}(\mathcal{K})} \|\hat{Y} - \bar{Y}\|_2 \tag{3.18}$$

where $\bar{Y}$ denotes the sample average $\frac{1}{N}\sum_n Y^n$. Under our assumptions on the data generating process, $\bar{Y}$ is a consistent estimator of Y^0, thus intuitively we should reject the null hypothesis when $\psi(\bar{Y})$ is large. However, ψ is not everywhere differentiable, which requires some care.

By an appropriate delta method due to [44], $\sqrt{N}\big(\psi(\bar{Y}) - \psi(Y^0)\big) \xrightarrow{L} \psi'_{Y^0}\big(N(0, \Sigma)\big)$, where $\psi'_{Y^0}(h)$ denotes the Hadamard directional derivative of ψ at Y^0 in the direction h, and Σ is the covariance matrix of ϵ.[67] Crucially, [44] show that ψ'_{Y^0} is linear if and only if ψ is differentiable at Y^0, and this is true if and only if the standard bootstrap $\sqrt{N}\big(\psi(\bar{Y}^*) - \psi(\bar{Y})\big)$ consistently estimates $\psi'_{Y^0}\big(N(0, \Sigma)\big)$.

[66] Recall from our definition of shape constraints we are guaranteed $\operatorname{grad}(\mathcal{K}) \subseteq \mathcal{F}$ is closed and convex.

[67] See also [105] and [41].

To obtain critical values for our statistic, first note that under H_0, $\psi(Y^0) = 0$, hence by [44], $\sqrt{N}\psi(\bar{Y}) \xrightarrow{L} \psi'_{Y^0}\big(N(0,\Sigma)\big)$. To simulate critical values for this distribution, the analogy principle suggests $\hat{\psi}'\big(\sqrt{N}(\bar{Y}^* - \bar{Y})\big)$, where $\hat{\psi}'$ is an appropriate nonparametric estimator of $\;'_{Y^0}$, and $\bar{Y}^*$ is a bootstrapped sample mean. The numerical derivative estimator of [69] provides a convenient method simulating this distribution without any further analytic calculations:

1. For $b = 1, \dots, B$, let $Z^{*(b)} = \sqrt{N}\big(\bar{Y}^{*(b)} - \bar{Y}\big)$, where $\bar{Y}^{*(b)}$ is a bootstrapped sample mean, drawn from the sample $\{Y^1, \dots, Y^N\}$.[68]

2. For all $b = 1, \dots, B$, compute:

$$\hat{\psi}'\big(Z^{*(b)}\big) = \frac{\big(\bar{Y} + \epsilon_N Z^{*(b)}\big) - \quad \big(\bar{Y}\big)}{\epsilon_N},$$

   for a choice of sequence of tuning parameters ϵ_N satisfying $\lim_N \epsilon_N = 0$, and $\lim_N \epsilon_N \sqrt{N} \to \infty$.

3. Use the empirical distribution of $\{\hat{\psi}'_N\big(Z^{*(b)}\big)\}_{b=1}^{B}$ to obtain critical values for (3.17).

When an explicit description of the tangent cones of $\text{grad}(\mathcal{K})$ is readily available, one can modify this procedure to make use of this extra analytic information [44]. Similarly, when $\text{grad}(K)$ is a closed, convex cone, [45] discuss a modification of this procedure with fine statistical properties.

[68]Note that the assumptions on our data generating process are sufficient to guarantee the consistency of the bootstrap.

Appendix A

Chapter Two Proofs

Proofs

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

Lemma. *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 $x_{k'}$).[1] 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,

[1] 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^*}.

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. □

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

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

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

[2] 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.

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. □

Corollary. *Let $(X, \leq_X)$ be a lattice, and suppose further that (i) Σ contains only totally ordered subsets of X, and (ii) every pair of elements of Σ is comparable in the strong set order. Then Σ is well-covered.*

Proof. Let γ denote an arbitrary loop in $\Gamma(X,\Sigma)$. By (i) we conclude that every edge pair of γ is related by $\leq_X$. As $\leq_X$ is a partial order and V_γ finite, $\leq_X |_{E_\gamma}$ admits a local minimum, in the sense of the existence of $x_{i-1}, x_i, x_{i+1} \in V_\gamma$ such that $x_i <_X x_{i-1}, x_{i+1}$, and $\{x_{i-1}, x_i\}, \{x_i, x_{i+1}\} \in E_\gamma$. Let $\mathcal{B}_\gamma$ be an arbitrary cyclic collection for γ. In light of (ii), without loss of generality suppose these two edges belong to different budgets in $\mathcal{B}_\gamma$, and let $B_{x_{i-1}x_i} \leq_{SSO} B_{x_i x_{i+1}}$ for two budgets in $\mathcal{B}_\gamma$ with $B_{x_{i-1}x_i}$ containing $\{x_{i-1}, x_i\}$ and $B_{x_i x_{i+1}}$ containing $\{x_i, x_{i+1}\}$. Then by the strong set order $x_{i-1} = x_i \vee x_{i-1} \in B_{x_{i+1}x_i}$, and hence $B_{x_{i+1}x_i}$ covers $\mathcal{B}_\gamma$. As γ and $\mathcal{B}_\gamma$ were arbitrary, we again conclude that Σ is well-covered. □

Appendix B

Chapter Three Proofs

Combinatorial Results on Simple Subdomains

We begin by recalling some definitions from the theory of simplicial complexes. A **simplicial complex** is a set of vertices $\{v_i\}_{i\in\mathcal{I}}$, and collection of non-empty finite subsets $\{s_j\}_{j\in\mathcal{J}}$ of $\{v_i\}$ called **simplices** such that:

1. Any set consisting of exactly one vertex is a simplex; and
2. Any non-empty subset of a simplex is a simplex.[1]

A simplex of cardinality $(n+1)$ is said to be of dimension n. Given a simplicial complex $\mathcal{D}_T$ generated by a collection of triangles T, the **boundary** $\mathcal{D}_T$, denoted $\dot{\mathcal{D}}_T$, is the sub-complex of generated by those 1-simplices which are the faces of exactly one triangle of T. The n-**skeleton** of a complex K, denoted $K^{(n)}$ is defined as the collection of all simplices of K of dimension n. We will commit the mild sin of occasionally using $K^{(n)}$ to denote both the set of n-simplices of K and also the subcomplex generated by these simplices where no confusion should result.

Theorem (Fundamental Theorem of Simple Sub-domains)**.** *Let $\mathcal{D}$ be an arbitrary domain, and $l \subseteq \mathcal{D}^{(1)}$ a loop. There exists a simple sub-domain for l if and only if there exists a simple sub-domain $\mathcal{D}|_{\bar{T}}$ for l that satisfies:*

[1] See [107] Section 3.1 (p. 108) for basic definitions.

(i) **Boundary***: $\dot{\mathcal{D}}|_{\bar{T}} = l$; and*

(ii) **Minimality***: The vertex set of $\mathcal{D}|_{\bar{T}}$ equals that of l.*

Proof. ($\Longleftarrow$) : Trivial.

($\Longrightarrow$) : Let $\tilde{T}$ generate a simple sub-domain for l, and consider a chain $\lambda \in C_1(\mathcal{D}|_{\tilde{T}}, \mathbb{R})$, given by:

$$\lambda = \sum_{\sigma \in l} n_\sigma \sigma,$$

with (i) zero coefficients on any $\sigma \in \mathcal{D}|_{\tilde{T}}$ that does not belong to l, and (ii) and such that, for all $\sigma \in l$, the coefficients satisfy $|n_\sigma| = 1$, where signs are chosen so $\lambda \in \ker \partial_1$.[2] As, by topological triviality, $H_1(\mathcal{D}|_{\tilde{T}}, \mathbb{R}) = 0$, $\operatorname{Im} \partial_2 = \ker \partial_1$ hence there exists some chain Λ in $C_2(\mathcal{D}|_{\tilde{T}}, \mathbb{R})$ solving:

$$\partial_2 \Bigg[\underbrace{\sum_{\tau \in \tilde{T}} n_\tau \tau}_{\Lambda} \Bigg] = \lambda$$

with some n_τ possibly equal to zero. Let $\bar{T} = \{\tau \in \tilde{T} : |n_\tau| \neq 0\}$ denote the support of Λ, and suppose for some 2-simplex $\tau \in \mathcal{D}|_{\bar{T}}$ there is a 1-face $\hat{\sigma}$ of τ such that $\hat{\sigma} \in \dot{\mathcal{D}}|_{\bar{T}}$ but $\hat{\sigma} \notin l$. Then we immediately obtain a contradiction: it must be the case actually $n_\tau = 0$, since $\partial_2 \Lambda$ would have coefficient equal in absolute value to $|n_\tau|$ on $\hat{\sigma}$, as by definition of a boundary, τ is the only 2-simplex in $\mathcal{D}|_{\bar{T}}$ containing $\hat{\sigma}$, and $n_{\hat{\sigma}} = 0$ as $\hat{\sigma} \notin l$. Therefore we conclude Λ is supported on a finite sub-collection $\bar{T}$ with the property that $\dot{\mathcal{D}}|_{\bar{T}} \subseteq l \subseteq \mathcal{D}|_{\bar{T}}$.

We claim first that $\mathcal{D}|_{\bar{T}}$ is combinatorially trivial. Suppose for sake of contradiction, that it fails to be so. Since $\tilde{T}$ was combinatorially trivial and $\bar{T} \subseteq \tilde{T}$, there exist a partition of $\bar{T}$ into maximal, non-empty collections of 2-faces $\bar{T}_1, \ldots, \bar{T}_K$, $K > 1$,

[2]The apparent indeterminacy of the signs of the coefficients in λ is simply a consequence of our being ambivalent about the choice basis for $C_1(\mathcal{D}|_{\tilde{T}}, \mathbb{R})$.

such that for each k, $\mathcal{D}|_{\bar{T}_k}$ is combinatorially trivial. This in turn implies that for all k, $\dot{\mathcal{D}}|_{\bar{T}_k} \neq \varnothing$, as any 'leaf' 2-face contains at least two 1-faces unique to it. Fix an arbitrary k and let $\hat{\Lambda}$ be a 2-chain in $C_2(\mathcal{D}|_{\bar{T}_k}, \mathbb{R})$ whose coefficients are are all 1 in absolute value, with signs chosen so that $\partial_2\hat{\Lambda}$ vanishes on any 1-face not in $\dot{\mathcal{D}}|_{\bar{T}_k}$. Then by construction,

$$\partial_2\hat{\Lambda} = \sum_{\sigma \in \dot{\mathcal{D}}|_{\bar{T}_k}} \hat{n}_\sigma \sigma$$

and for all such σ, $|\hat{n}_\sigma| = 1$. By identity, $(\partial_1 \circ \partial_2)(\hat{\Lambda}) = 0$, and thus for each vertex $x \in \dot{\mathcal{D}}|_{\bar{T}_k}^{(0)} \neq \varnothing$, x is contained in an even number of 1-faces in $\dot{\mathcal{D}}|_{\bar{T}_k}$, and hence $\dot{\mathcal{D}}|_{\bar{T}_k}$ consists of a union of loops. But as $\dot{\mathcal{D}}|_{\bar{T}_{\hat{k}}} \subsetneq \dot{\mathcal{D}}|_{\bar{T}} \subseteq l$, we obtain a contradiction, as no proper subcomplex of a loop may be a loop. Thus $\mathcal{D}|_{\bar{T}}$ is itself combinatorially trivial.

We now verify that $\dot{\mathcal{D}}|_{\bar{T}} = l$. Recall we have already obtained that $\dot{\mathcal{D}}|_{\bar{T}} \subseteq l \subseteq \mathcal{D}|_{\bar{T}}$. Suppose then, for sake of contradiction, that there exists some 1-face $\sigma \in l$, such that $\sigma \notin \dot{\mathcal{D}}|_{\bar{T}}$. Let $\tau \in \mathcal{D}|_{\bar{T}}$ denote one of the two 2-faces (combinatorial triviality) of $\mathcal{D}|_{\bar{T}}$ that contains σ, and let K denote the sub-complex of $\mathcal{D}|_{\bar{T}}$ generated by those 2-faces of $\mathcal{D}|_{\bar{T}}$ that may be reached from τ by a sequence of distinct 2-simplices with adjacent terms sharing a common face, but whose intersections do not contain σ. By construction K is combinatorially trivial; by an argument analogous to that of the preceding paragraph, $\dot{K}$ is a non-empty union of loops. But $\dot{K} \subsetneq \dot{\mathcal{D}}|_{\bar{T}} \cup \{\sigma\} \subseteq l$, where the first strict inclusion follows from the fact that the complement of K in $\mathcal{D}|_{\bar{T}}$ is a non-empty combinatorially trivial subcomplex too. Hence we obtain a contradiction, again because l cannot contain any proper sub-complex that is also a loop, and thus $\dot{\mathcal{D}}|_{\bar{T}} = l$ as claimed.

We turn to verifying our minimality claim, that the vertex sets of $\mathcal{D}|_{\bar{T}}$ and l coincide: $\mathcal{D}|_{\bar{T}}^{(0)} = l^{(0)}$. Let G denote the undirected graph whose vertex set is given by the 2-faces of $\mathcal{D}|_{\bar{T}}$ and whose edge set determined by the relation of having an intersection containing a 1-face. By combinatorial triviality of $\mathcal{D}|_{\bar{T}}$, G is a tree. Now, suppose toward a contradiction that the vertex sets of $\mathcal{D}|_{\bar{T}}$ and l do not coincide. Since $\dot{\mathcal{D}}|_{\bar{T}} = l$, this implies there is some vertex x of $\mathcal{D}|_{\bar{T}}$ not in l. Now, as $\mathcal{D}|_{\bar{T}}$ is combinatorially trivial, every 1-face σ of $\mathcal{D}|_{\bar{T}}$ that contains x is contained in precisely two 2-simplices. Let $\tilde{G}$ be the subgraph of G consisting of those 2-faces containing x as a vertex. Since each vertex τ of $\tilde{G}$ contains precisely two 1-faces that contain x, by finiteness $\tilde{G}$ is a cycle graph, contradicting the fact that G is a tree (i.e. that $\mathcal{D}|_{\bar{T}}$ is combinatorially trivial). Hence the vertex sets of $\mathcal{D}|_{\bar{T}}$ and l coincide.

Finally, we show the dimension-1 simplicial homology of $\mathcal{D}|_{\bar{T}}$ is zero in real coefficients, our last outstanding claim. As $\mathcal{D}|_{\bar{T}}$ is combinatorially trivial, its collection of 1-faces may be partitioned into two subsets: those faces in $\dot{\mathcal{D}}|_{\bar{T}}$ and those not. By definition, the edge-set of the graph G introduced in the preceding paragraph is in one-to-one correspondence with the the set of 1-faces of $\mathcal{D}|_{\bar{T}}$ not in $\dot{\mathcal{D}}|_{\bar{T}}$. By combinatorial triviality, G is a tree and hence has one more vertex (2-simplex of $\mathcal{D}|_{\bar{T}}$) than edge (1-face of $\mathcal{D}|_{\bar{T}}$ not in $\dot{\mathcal{D}}|_{\bar{T}}$). Similarly, $l = \dot{\mathcal{D}}|_{\bar{T}}$ is a loop, so the number of 1-faces must be the same as the number of vertices of l, which we have established is also the vertex set of $\mathcal{D}|_{\bar{T}}$. The Euler-Poincaré theorem ([86] Theorem 22.2) asserts the equivalence of the following two definitions of the Euler characteristic of $\mathcal{D}|_{\bar{T}}$:

$$\chi(\mathcal{D}|_{\bar{T}}) = \dim H_0(\mathcal{D}|_{\bar{T}}, \mathbb{R}) - \dim H_1(\mathcal{D}|_{\bar{T}}, \mathbb{R}) + \dim H_2(\mathcal{D}|_{\bar{T}}, \mathbb{R}) = V - E + F,$$

where V is the number of 0-simplices, E the number of 1-simplices, and F the number of 2-simplices in $\mathcal{D}|_{\bar{T}}$. By the above counting argument for the set of 1-faces of $\mathcal{D}|_{\bar{T}}$, we know:

$$E = \underbrace{V}_{\text{1-faces in } \dot{\mathcal{D}}|_{\bar{T}}.} + \underbrace{F - 1}_{\text{1-faces in } \mathcal{D}|_{\bar{T}} \setminus \dot{\mathcal{D}}|_{\bar{T}}.} \tag{B.1}$$

and hence $\chi(\mathcal{D}|_{\bar{T}}) = 1$. Now, since every 2-simplex in $\mathcal{D}|_{\bar{T}}$ intersects l, $\mathcal{D}|_{\bar{T}}$ is path-connected and hence $\dim H_0(\mathcal{D}|_{\bar{T}}, \mathbb{R}) = 1$. Moreover, by combinatorial triviality, $H_2(\mathcal{D}|_{\bar{T}}, \mathbb{R}) = 0$.[3] Then by Euler-Poincaré, $\dim H_1(\mathcal{D}|_{\bar{T}}, \mathbb{R}) = 0$, and thus $\mathcal{D}|_{\bar{T}}$ is a simple sub-domain as claimed. □

Lemma (Union Lemma)**.** *Let $\mathcal{D}|_T, \mathcal{D}|_{T'}$ be two simple sub-domains whose intersection consists of a single 1-face σ. Then $\mathcal{D}|_T \cup \mathcal{D}|_{T'}$ is a simple sub-domain.*

Proof. As $\mathcal{D}|_T$ and $\mathcal{D}|_{T'}$ are combinatorially trivial, it is immediate that so too is $\mathcal{D}|_T \cup \mathcal{D}|_{T'}$. Then, by the reduced simplicial Mayer-Vietoris theorem ([86] Theorem 25.1) there exists an exact sequence:

$$0 \to \tilde{H}_1(\mathcal{D}|_T \cap \mathcal{D}|_{T'}, \mathbb{R}) \to \tilde{H}_1(\mathcal{D}|_T, \mathbb{R}) \oplus \tilde{H}_1(\mathcal{D}|_{T'}, \mathbb{R}) \to \tilde{H}_1(\mathcal{D}|_T \cup \mathcal{D}|_{T'}, \mathbb{R}) \to \tilde{H}_0(\mathcal{D}|_T \cap \mathcal{D}|_{T'}, \mathbb{R})$$

which, making use of topological triviality of $\mathcal{D}|_T$ and $\mathcal{D}|_{T'}$ and the contractibility of $\mathcal{D}|_T \cap \mathcal{D}|_{T'}$ (i.e. $= \sigma$), reduces to:

$$0 \to (0) \to (0) \oplus (0) \to \tilde{H}_1(\mathcal{D}|_T \cup \mathcal{D}|_{T'}, \mathbb{R}) \to 0$$

and hence $\tilde{H}_1(\mathcal{D}|_T \cup \mathcal{D}|_{T'}, \mathbb{R}) = 0$, and equivalently $H_1(\mathcal{D}|_T \cup \mathcal{D}|_{T'}, \mathbb{R})$ by definition of reduced simplicial homology. □

[3] Since $\mathcal{D}|_{\bar{T}}$ is homengeously 2-dimensional it contains no simplices of dimension greater than two, hence $H_2(\mathcal{D}|_{\bar{T}}, \mathbb{R}) = 0$ if and only if the only solution to:

$$\partial_2 \left[\sum_{\tau \in \mathcal{D}|_{\bar{T}}^{(2)}} \tilde{n}_\tau \tau \right] = 0$$

is for $\tilde{n}_\tau = 0$ for all $\tau \in \mathcal{D}|_{\bar{T}}^{(2)}$. Clearly for any solution, any $\tau \in \mathcal{D}|_{\bar{T}}$ which contains a 1-face in $\dot{\mathcal{D}}|_{\bar{T}}$ must have have $\tilde{n}_\tau = 0$ by (PM.2). Hence in any non-zero solution to the above, the sub-collection of 2-simplices in $\mathcal{D}|_{\bar{T}}$ with non-zero coefficients must have the property that all of their 1-faces are contained also in some other (hence unique other) member of the sub-collection. But this sub-collection defines a subgraph of the graph G, and the above property implies again that this subgraph can have no leaves, contradicting the fact G is a tree, as $\mathcal{D}|_{\bar{T}}$ is combinatorially trivial. Hence $H_2(\mathcal{D}|_{\bar{T}}, \mathbb{R}) = 0$.

We now turn to the proof of Theorem 3 in the text.

Theorem. *Let (X, Σ) be a choice environment. Then the domain $\mathcal{D}(X, \Sigma)$ is simple if and only if the budget graph is chordal.*

Proof. ($\Longrightarrow$): Suppose the domain is simple, and let γ be a loop. By simplicity of the domain, there exists some simple subdomain containing γ. By the Fundamental Theorem of Simple Subdomains, there exists a collection of triangles $\tilde{T}$ such that the edge-set of the subdomain generated by $\tilde{T}$ consists either of edges of γ or bisections of γ. If $|E_\gamma| = 3$ then there is nothing to check thus suppose $|E_\gamma| > 3$. Then by combinatorial triviality, there exists at least one element of $\tilde{E}$ that does not belong to E_γ, and thus γ has a chord. Since γ was arbitrary, we conclude $\Gamma(X, \Sigma)$ is chordal.

($\Longleftarrow$): To prove the domain associated to any chordal budget graph is simple, we proceed by contraposition. Suppose, then, that $\mathcal{D}(X, \Sigma)$ is not simple. Then, as there exists a loop contained in no simple sub-domain, there exists a shortest such loop, which we will denote γ, with $E_\gamma = \big\{\{x_0, x_1\}, \ldots, \{x_{n-1}, x_n\}, \{x_n, x_0\}\big\}$. We know $|V_\gamma|$ must be strictly greater than three, lest V_γ be a triple in $\mathcal{D}$ and hence this triple serve trivially as a simple sub-domain containing γ. We now prove that for all $e \in E_\Gamma$, $e \subseteq V_\gamma$ if and only if $e \in E_\gamma$, that is, that γ is a chordless loop in $\Gamma(X, \Sigma)$. Clearly $E_\gamma \subseteq E_\Gamma$. Thus, for sake of contradiction, suppose there exists an $e \in E_\Gamma$ with $e \subseteq V_\gamma$ but $e \notin E_\gamma$. Then without loss of generality, $e = \{x_j, x_k\}$ with $k > j + 1$. Hence we obtain two loops, γ_1 and γ_2 via:

$$E_{\gamma_1} = \big\{\{x_0, x_1\}, \ldots, \{x_{j-1}, x_j\}, \{x_j, x_k\}, \{x_k, x_{k+1}\}, \ldots, \{x_n, x_0\}\big\}$$

and

$$E_{\gamma_2} = \big\{\{x_j, x_{j+1}\}, \ldots, \{x_{k-1}, x_k\}, \{x_k, x_j\}\big\},$$

both shorter than γ. By the minimality of γ, there exist simple sub-domains $\mathcal{D}|_{\tilde{T}_1}$ and $\mathcal{D}|_{\tilde{T}_2}$ of $\mathcal{D}(X, \Sigma)$ for γ_1 and γ_2 respectively, and by the fundamental theorem of simple sub-domains, these complexes may be taken to intersect only on the 1-face $\{x_j, x_k\}$. But by the union lemma, $\mathcal{D}|_{\tilde{T}_1} \cup \mathcal{D}|_{\tilde{T}_2}$ is a simple sub-domain for γ, a contradiction. Thus γ is chordless, and hence Γ is not chordal. □

Finally, we conclude with the (counter-)example mentioned in the text.

Example 26. Let $X = \{x_0, \ldots, x_4\}$, and let $\Sigma = E_\Gamma = \big\{\{x_0, x_1\}, \{x_1, x_2\}, \{x_2, x_3\}, \{x_3, x_4\}, \{x_4, x_0\}, \{x_1, x_3\}, \{x_3, x_0\}, \{x_0, x_2\}, \{x_2, x_4\}, \{x_4, x_1\}\big\}$. Note the domain associated to this environment corresponds to a triangulation of the Möbius strip. Let γ correspond to the boundary of the strip, i.e. the loop with edge set $E_\gamma = \big\{\{x_1, x_3\}, \{x_3, x_0\}, \{x_0, x_2\}, \{x_2, x_4\}, \{x_4, x_1$ As $X = V_\gamma$, clearly every edge in E_Γ consists of either an edge of E_γ or a bisecting edge, and it is simple to verify that every vertex belongs to some bisecting edge of γ. Finally, there is only one subdomain containing γ, the entire domain itself, and this is neither combinatorially trivial (its 'sharing a common face graph is a circle graph on five vertices) nor topologically trivial (the homology group of the Möbius strip in dimension one with $\mathbb{R}$-coefficients is $\mathbb{R}$).

Integrability Results

Given a simplicial complex K, a (discrete) n**-form** is a linear functional acting on oriented pieces of the n-dimensional skeleton $K^{(n)}$.[4] Define the space of all n-forms on K as:

$$C^n(K) = \big\{\phi : K^{(n)} \to \mathbb{R} : \phi([x_{\sigma(0)}, \ldots, x_{\sigma(n)}]) = \text{sign}(\sigma)\phi([x_0, \ldots, x_n])\big\},$$

[4]We intentionally adopt the analyst's terminology of 'forms' rather than the topologist's 'cochain' to further highlight the parallel to the exterior calculus arguments underpinning the solution to the classical integrability problem. See [71] and [55] for an in-depth discussion of the parallels between the smooth and discrete theories.

where $[x_0, \ldots, x_n]$ denotes an oriented n-simplex of K and σ is any permutation of $\{0, \ldots, n\}$. In particular, the vector space $C^0(K)$ consists precisely of all real valued functions on the vertices; $C^1(K)$ may be interpreted as the space of all real-valued flows on the 1-skeleton of K, where the permutation condition simply ensures that these flows are directed .

The gradient and curl operators are defined analogously as in the text. A 1-form F is said to be **exact** (or 'integrable') if there exists an $f \in C^0(K)$ such that $\operatorname{grad}(f) = F$. Similarly, if $\operatorname{rot}(F) = 0$, F is said to be **closed**. An exact 1-form is always closed; this may be succinctly stated as $\operatorname{Im}(\operatorname{grad}) \subseteq \operatorname{Ker}(\operatorname{rot})$. In particular, this implies the quotient vector space $\operatorname{Ker}(\operatorname{rot})/\operatorname{Im}(\operatorname{grad})$ is well-defined. This quotient is denoted $H^1(K, \mathbb{R})$ and is known as the first simplicial cohomology group of K (with $\mathbb{R}$-coefficients); its dimension may be interpreted as a measure of how far the closedness of a 1-form is from guaranteeing its exactness, or integrability.

Proposition. *Let $c \in \mathcal{C}(X, \Sigma)$. Then c is locally rationalizable if and only if c both:*

(i) obeys the weak axiom; and

(ii) is ordinally irrotational.

Proof. ($\Longrightarrow$): Suppose c is locally rationalizable. Then $\succeq$ and its asymmetric component $\succ$ form a suitable order pair extension to verify ordinal irrotationality. Moreover, c must obey the weak axiom for all those pairs of alternatives that are contained in some triangle of the budget graph. Thus we need only to verify that c does not violate the weak axiom for those pairs of distinct alternatives x, y which form edges not belonging to any triangle. But if $\{x, y\} \in E_\Gamma$ and $\{x, y\}$ is not contained in any triangle, then $\{x, y\} \in \Sigma$ and this must be the only budget containing both these alternatives, precluding any possible violation of the weak axiom over them.

($\Longleftarrow$): Suppose now c is ordinally irrotational and satisfies the weak axiom. Let $(\succeq, \succ^*)$ denote the order extension of the revealed preference guaranteed by ordinal irrotationality. In particular, the asymmetric component $\succ$ of $\succeq$ is a sub-relation of $\succ^*$, i.e. $\succ \subseteq \succ^*$. If $\succ = \succ^*$, then $\succsim$ is a local rationalization, thus it suffices to establish that one may always take $\succ^*$ to be the asymmetric component of $\succeq$.

Let $Z = \{(x, y) \in X \times X : x \succ^* y, y \succeq x, \text{ and } x \succeq y\}$ denote those pairs for which $x \succ^* y$ but $\neg x \succ y$ (recall that while $\succ^*$ contains $\succ$, by the definition of an order pair, $\succ^* \subseteq \succeq$). We partition Z into two subsets: $Z_0 = \{(x, y) \in Z : y \succsim_c x\}$ and $Z_1 = \{(x, y) \in Z : \neg y \succsim_c x\}$. Consider first those pairs in Z_0. If $(x, y) \in Z_0$ then $x \succ^* y$, $x \succeq y$, $y \succeq x$, and $y \succsim_c x$. As c satisfies the weak axiom, it cannot be the case that $x \succ_c y$. Moreover, it cannot be the case that $y \succ_c x$, as $\succ^*$ contains $\succ_c$ and $x \succ^* y$ already, and $\succ^*$ is asymmetric by hypothesis. Thus it must be that $y \succsim_c x$ and $x \succsim_c y$. In other words, both (x, y) and (y, x) are required to belong to $\succeq$, but we may omit (x, y) from $\succ^*$ without problem: $(\succeq, \succ^* \setminus \{(x, y)\})$ is still complete restricted to each triangle as $(x, y) \in \succeq$, and of course if $(\succeq, \succ^*)$ had no triangular cycles then neither could $(\succeq, \succ^* \setminus \{(x, y)\})$. More generally, $(\succeq, \succ^* \setminus Z_0)$ is complete restricted to each triangle in the budget graph, has no triangular cycles, and $\succ \subseteq \succ^* \setminus Z_0$.

Consider now Z_1. If $x \succ^* y$, $x \succeq y$, $y \succeq x$, but $\neg y \succsim_c x$, then there is no reason to include (y, x) in $\succeq$. As $\succ^*$ is asymmetric, we know (y, x) does not belong to $\succ^*$ (as (x, y) does by hypothesis). Moreover, its removal from $\succeq$ cannot affect the completeness of the order pair restricted to each triangle, nor can it create triangular cycles where none previously were. Hence starting from $(\succeq, \succ^*)$, we may instead consider the order pair $(\succeq \setminus Z_1 x, \succ^* \setminus Z_0)$, which has the property that the asymmetric part of $\succeq \setminus Z_1$ is $\succ^* Z_0$, and thus forms a local rationalization for c. $\square$

Henceforth we fix a choice environment (X, Σ), and will suppress the argument (X, Σ) appearing in domains. Let $\succeq$ be a binary relation on X that is locally rational.[5] Let γ be a loop in $\mathcal{D}$, and $\mathcal{D}|_{\tilde{T}}$ a simple sub-domain of $\mathcal{D}$ containing γ. A 1-form $F \in C^1(\mathcal{D}|_{\tilde{T}})$ is a **cardinalization** of $\succeq$ on $\mathcal{D}|_{\tilde{T}}$ if, for all 1-faces of $\mathcal{D}|_{\tilde{T}}$:

$$y \succeq x \implies F([x, y]) \geq 0,$$

and

$$y \succ x \implies F([x, y]) > 0.$$

Lemma. *(Closed Cardinalization Lemma) Let $\succeq$ be locally rational, and let $\mathcal{D}|_{\tilde{T}} \subseteq \mathcal{D}$ be a simple sub-domain. Then there exists a closed cardinalization of $\succeq$ on $\mathcal{D}|_{\tilde{T}}$.*

Proof. Let $\{\tilde{\tau}_1, \ldots, \tilde{\tau}_I\}$ be an enumeration of those 2-simplices of $\mathcal{D}|_{\tilde{T}}$ that each contain at least two distinct 1-faces in $\dot{\mathcal{D}}|_{\tilde{T}}$. Using this, we construct an enumeration of all 2-simplices of K as follows: between each $\tilde{\tau}_i, \tilde{\tau}_{i+1}$ insert the unique (via combinatorial triviality) ordered sequence of 2-simplices in $\mathcal{D}|_{\tilde{T}}$ connecting them, omitting any 2-simplices of $\mathcal{D}|_{\tilde{T}}$ that have appeared in the construction prior. Let $\{\tau_1, \ldots, \tau_J\}$ denote this enumeration, and let $\mathcal{D}|^j_{\tilde{T}}$ denote the sub-complex of $\mathcal{D}|_{\tilde{T}}$ generated by the first j elements of this enumeration.

We now inductively construct our closed cardinalization of $\succeq$. First, note that there is trivially a closed cardinalization of $\succeq$ on $\mathcal{D}|^1_{\tilde{T}}$: $\succeq$ restricted to the vertices of τ_1 is complete and transitive by local rationality, hence admits a utility function u_1 on these vertices. Let $F_1 \in C^1(\mathcal{D}|^1_{\tilde{T}})$ be defined as $\text{grad}(u_1)$. For our inductive step, suppose now that there is a closed cardinalization $F_j \in C^1(\mathcal{D}|^j_{\tilde{T}})$ of $\succeq$ on $\mathcal{D}|^j_{\tilde{T}}$, for some $j < J$. By analogous logic, there is a utility function u_{j+1} representing $\succeq$ restricted to the vertices of τ_{j+1}. Let $\tilde{F}_{j+1} = \text{grad}(u_{j+1})$ be the closed 1-form on the the complex

[5] That is, for all $T \in T_\Gamma$, $\succeq|_T$ is complete and transitive.

generated by τ_{j+1} alone. By virtue of the structure of the enumeration constructed above, τ_{j+1} and τ_j intersect on exactly a single 1-face, σ with vertex set $\{a, b\}$. There exists some $c \in \mathbb{R}_{++}$ such that $F_j([a, b]) = c\tilde{F}_{j+1}([a, b])$, with c unique if $\succeq$ is strict over this pair. Then define:

$$F_{j+1}([x, y]) = \begin{cases} F_j([x, y]) & \text{if } [x, y] \not\subset \tau_{j+1} \\ c\tilde{F}_{j+1}([x, y]) & \text{if } [x, y] \subset \tau_{j+1}, \end{cases}$$

completing the proof. □

Theorem (Ordinal Integrability Theorem)**.** *Let (X, Σ) be a choice environment with $\mathcal{D}(X, \Sigma)$ a simple domain. Then a choice correspondence $c \in \mathcal{C}(X, \Sigma)$ is strongly rationalizable if and only if:*

(i) It obeys the weak axiom; and

(ii) It is locally rationalizable.

Moreover, (i) and (ii) are jointly equivalent to the strong rationalizability of c if and only if $\mathcal{D}(X, \Sigma)$ is simple.

Proof. We begin first by verifying (i) and (ii) are equivalent to strong rationalizability for simple $\mathcal{D}$. Clearly, strong rationalizability always implies (i) and (ii), regardless of the structure of $\mathcal{D}$: any rationalizing weak order $\succeq_c$ of course is a local rationalization and implies $(\succsim_c, \succ_c)$ obeys the weak axiom.

Now, suppose $\mathcal{D}$ is simple, and let $c \in \mathcal{W}(X, \Sigma)$ be locally rationalizable. Let $\gamma \subseteq \mathcal{D}$ be an arbitrary loop. We will show that $\succsim_c |_{E_\gamma}$ cannot be cyclic. As γ is a loop, by simplicity of $\mathcal{D}$ there exists a simple sub-domain $\mathcal{D}|_{\tilde{T}} \subseteq \mathcal{D}$ containing γ, and $\succeq$ a local rationalization of $\succsim_c$ on $\mathcal{D}|_{\tilde{T}}$. By the preceding lemma there exists a closed cardinalization of $\succeq$ on $\mathcal{D}|_{\tilde{T}}$, which we will denote by $F \in C^1(\mathcal{D}|_{\tilde{T}})$. By the cohomology universal coefficient theorem (see [86] Theorem 53.1), there exists an isomorphism between $H_1(\mathcal{D}|_{\tilde{T}}, \mathbb{R})$ and $H^1(\mathcal{D}|_{\tilde{T}}, \mathbb{R})$ (see [86] Corollary 53.6 or [71] Theorem 4), and hence as $\mathcal{D}|_{\tilde{T}}$ is topologically trivial, $H^1(\mathcal{D}|_{\tilde{T}}, \mathbb{R}) = 0$, and therefore there exists an $f \in C^0(\mathcal{D}|_{\tilde{T}})$ such that $\text{grad}(f) = F$. Define the binary relation $\geq^*$ on the vertex set $\mathcal{D}|_{\tilde{T}}^{(0)}$ via $x_0 \geq^* x_1 \iff f(x_0) \geq f(x_1)$ (resp. strict). This is a weak order on the vertices of $\mathcal{D}|_{\tilde{T}}$ which, by consistency of F, is an extension of $\succeq$ on $\mathcal{D}|_{\tilde{T}}$.[6] Thus $\succsim_c|_{E_\gamma}$ is acyclic. As γ was arbitrary, and every potential cycle of $\succsim_c$ must be supported on some loop in $\mathcal{D}$, each contained in some simple sub-domain, we conclude $\succsim_c$ is acyclic. Thus, for all $c \in \mathcal{W}(X, \Sigma)$, if c is also locally rationalizable, it must satisfy the generalized axiom and hence is strongly rationalizable.

We now show that if $\mathcal{D}$ is not simple, (i) and (ii) do not imply strong rationalizability. Suppose, then, that $\mathcal{D}$ is not simple. By Theorem 3 there exists a chordless loop in the budget graph $\Gamma(X, \Sigma)$, which we will denote γ. Thus, there exists a cyclic collection for γ, denoted $\mathcal{B}_\gamma = \{B_1, \ldots, B_n\} \subseteq \Sigma$ such that for all $0 \leq j \leq n$ we have $\{x_j, x_{j+1}\} \subseteq B_j$, and that this collection is uncovered: as $|V_\gamma| > 3$ and γ is chordless, no budget in *all of* Σ contains any pair of non-adjacent points in γ. For all $B \in \Sigma|_{\mathcal{B}_\gamma}$, let:

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

and for all $B \in \Sigma$ define:

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

[6] It is generally an extension of $\succeq$ (which itself extends the revealed preference $\succsim_c$) as $\geq^*$ is complete and thus generally relates vertices not connected by any edge in $\mathcal{D}|_{\tilde{T}}^{(1)}$.

By an argument analogous to that in the proof of Theorem 1, $c \in \mathcal{W}(X, \Sigma)$ and not $\mathcal{G}(X, \Sigma)$.

We now verify that c is nonetheless locally rationalizable. To do this, we will explicitly construct a local rationalization $\succeq$. First, for all $e \in E_\gamma$, let $x_i \prec x_{i+1}$. Thus for all pairs $\{x, y\} \in E_\gamma$, $x \succ y$ if and only if $x \succ_c y$. For all $e \in E_\Gamma \setminus E_\gamma$ that intersect V_γ, we have shown this intersection must be singleton. For all such e, we know e is of the form $\{a, x_i\}$ for some $x_i \in V_\gamma$. For all pairs $\{a, x_i\}$ with $a \in (\cup_{\tilde{B} \in \mathcal{B}_\gamma} \tilde{B})$, define $a \prec x_i$, and if $a \notin (\cup_{\tilde{B} \in \mathcal{B}_\gamma} \tilde{B})$, let $a \succ x_i$. Finally, for those pairs $\{a, b\}$ that do not intersect V_γ, we consider two cases. If, either $\{a, b\} \subseteq (\cup_{\tilde{B} \in \mathcal{B}_\gamma} \tilde{B})$ or $\{a, b\} \subseteq X \setminus (\cup_{\tilde{B} \in \mathcal{B}_\gamma} \tilde{B})$, then let $a \succeq b$ and $b \succeq a$. If exactly exactly one element (without loss a) of $\{a, b\}$ is contained in $(\cup_{\tilde{B} \in \mathcal{B}_\gamma} \tilde{B})$, then let $b \succ a$. Finally, let $\succeq^*$ denote the reflexive closure of $\succeq$. Then $\succeq^* \supseteq \succsim_c$, and, by construction, $\succeq^*$ is locally rational. This follows from (i) for all $\{a, b\} \in E_\Gamma$, either $a \succeq^* b$ or $b \succeq^* a$, and (ii) for every $T \in T_\Gamma$, T only contains at most one pair in γ, as γ is chordless and of length greater than three. Thus, in particular, if T contains an edge of γ, denoted $\{x, y\}$, it contains some element z such that either $z \succ^* x, y$ or $x, y \succ^* z$. If T contains no edges of γ, then $\succeq^* |_T$ is clearly complete and transitive, and hence $\succeq^*$ is locally rational. $\square$

Application Theorems

Proposition. *Let (X, Σ) be a cardinality-constrained choice environment. Then every $c \in \mathcal{W}(X, \Sigma)$ is locally rationalizable if and only if $\mathcal{T}_\Gamma = \Sigma_3$.*

Proof. ($\Longleftarrow$): If $T_\Gamma = \Sigma_3$, consider the revealed preference of any choice function c obeying the weak axiom. If, for any pair $\{x, y\}$ contained in some $T \in \mathcal{T}_\Gamma$, we have neither $x \succsim_c y$ nor $y \succsim_c x$, it means that for every $T' \in \mathcal{T}_\Gamma = \Sigma_3$, it is the case that

that neither x nor y are chosen, hence $c(T') = T' \setminus \{x, y\}$. Note that, for any $T \in \Sigma_3$, at most one pair of elements may not have any preference revealed between them, as T is a budget itself so some choice must occur on it. Thus adding both (x, y) and (y, x) to $\succsim_c$ for every such $\succsim_c$-unrelated pair $\{x, y\}$ yields a locally rational extension.

($\Longrightarrow$): We proceed by contraposition. Suppose $\mathcal{T}_\Gamma \neq \Sigma_3$. Of course $\Sigma_3 \subseteq \mathcal{T}_\Gamma$, hence there exists some $\{x, y, z\} \in \mathcal{T}_\Gamma$ that is not a budget itself, but every pair of elements in it is contained in some budget. It is immediate then, due to the cardinality constraints on Σ, that $\{x, y\}, \{y, z\}, \{z, x\}$ is a loop in $\Gamma(X, \Sigma)$ that possesses an uncovered cyclic collection. By Lemma 2 we obtain the existence of a choice correspondence obeying the weak axiom whose revealed preference exhibits a three-cycle on this loop. Since the vertex set of this loop is in $\mathcal{T}_\Gamma$, this choice correspondence cannot be locally rationalizable. $\square$

Corollary. *Let (X, Σ) be a cardinality-constrained choice environment. Then the weak axiom characterizes strong rationalizability for any choice correspondence if and only (i) $T_\Gamma = \Sigma_3$, and (ii) the domain $\mathcal{D}(X, \Sigma)$ is simple.*

Proof. ($\Longrightarrow$): Suppose $\mathcal{W}(X, \Sigma) = \mathcal{G}(X, \Sigma)$. Then by Theorems 1 and 2, $\mathcal{D}(X, \Sigma)$ is simple and every choice correspondence in $\mathcal{W}(X, \Sigma)$ is locally rationalizable. By the preceding proposition, it follows then that $T_\Gamma = \Sigma_3$.

($\Longleftarrow$): Suppose $T_\Gamma = \Sigma_3$ and $\mathcal{D}(X, \Sigma)$ is simple, and let $\gamma \subseteq \Gamma(X, \Sigma)$ be an arbitrary loop. Then there exists some sub-collection $\tilde{\mathcal{T}}$ of three-good budgets that generate a simple sub-domain containing γ. By the fundamental theorem of simple subdomains, we may take this simple sub-domain's edge set to consist solely of edges of γ and bisections of γ. But, by combinatorial triviality, there exists a 'leaf' triangle in this sub-domain, hence for this triangle there exists a pair of edges $\{x, y\}, \{y, z\} \in E_\gamma$

such that $\{x, y, z\} \in \Sigma_3$. This implies that every cyclic collection for γ is covered, and by the arbitrariness of γ, Σ is well-covered. Theorem 1 then completes the proof. $\square$

Appendix C

Chapter Four Proofs

Proof of Theorem 4

Theorem 1. Suppose that ϕ is a continuous action of $\mathbb{R}_+$ on X. Then a continuous preference $\succsim$ on X satisfies (N.1) - (N.3) if and only if it admits a representation by a continuous, additive-equivariant utility.

Proof. It is immediate that if a preference relation admits a continuous additive-equivariant utility then it must satisfy (N.1) - (N.3), thus we focus on sufficiency.

Suppose then that $\succsim$ is a continuous weak order on X satisfying (N.1) - (N.3), and that there exists a $\succsim$ least-preferred alternative, $\underline{x}$. For all $x \in X$, define $c(x)$ as the (unique) solution to:

$$\phi\big(c(x), \underline{x}\big) \sim x.$$

For each x, existence of $c(x)$ follows from (N.3) and uniqueness from (N.2). Moreover, suppose $x \succsim y$. Then:

$$\phi\big(c(x), \underline{x}\big) \sim x \succsim y \sim \phi\big(c(y), \underline{x}\big),$$

hence by (N.3) there exists $\alpha \geq 0$ such that $\phi\big(\alpha, \phi(c(y), \underline{x})\big) = \phi\big(\alpha + c(y), \underline{x}\big) \sim \phi\big(c(x), \underline{x}\big)$. Thus by (N.2), $\alpha + c(y) = c(x)$, and hence $c(x) \geq c(y)$. Thus $c(\cdot)$ represents $\succsim$. As X is metric and $\succsim$ is continuous and admits the representation c, by [39] we conclude $\succsim$ admits a continuous utility representation $u : X \to \mathbb{R}$. Suppose $(x_n) \to x$. By continuity of u, $u(x_n) \to u(x)$. But $u(x_n) = u\big(\phi(c(x_n), \underline{x})\big)$ and $u(x) = u\big(\phi(c(x), \underline{x})\big)$.

As $\succsim$ satsifies (N.2), $\phi(\cdot,\underline{x})$ and $u|_{\phi(\mathbb{R}_+,\underline{x})}$ are injective, hence $\bar{u} = u|_{\phi(\mathbb{R}_+,\underline{x})} \circ \phi(\cdot,\underline{x})$ is injective and continuous. Thus as $\bar{u}\big(c(x_n)\big) \to \bar{u}\big(c(x)\big)$, $c(x_n) \to c(x)$, and as $x_n \to x$ was arbitrary, c is continuous.

To establish the additive-equivariance of c, note that by definition, for all x:

$$\phi\big(c(x),\underline{x}\big) \sim x. \tag{C.1}$$

Hence for all $x \in X$ and all $\alpha \geq 0$:

$$\phi\big(c(\phi(\alpha,x)),\underline{x}\big) \sim \phi(\alpha,x). \tag{C.2}$$

But by (C.1) and (N.1),

$$\phi\big(\alpha,\phi(c(x),\underline{x})\big) \sim \phi(\alpha,x), \tag{C.3}$$

and, as ϕ is an action:

$$\phi\big(\alpha,\phi(c(x),\underline{x})\big) = \phi(\alpha + c(x),\underline{x}). \tag{C.4}$$

Then by (C.2) - (C.4):

$$\phi(\alpha + c(x),\underline{x}) \sim \phi\big(c(\phi(\alpha,x)),\underline{x}\big),$$

and by (N.2) we conclude:

$$\alpha + c(x) = c(\phi(\alpha,x)). \tag{C.5}$$

Thus c is a continuous, additive-equivariant representation of $\succsim$.

Suppose now that $\succsim$ has no least-preferred alternative. Let $\underline{x} \in X$ be arbitrary, and define $c_{\underline{x}}(x)$ for all x in the upper contour set $\{x \in X : x \succsim \underline{x}\}$, as the unique solution to $\phi\big(c_{\underline{x}}(x),\underline{x}\big) \sim x$. By the preceding argument, $c_{\underline{x}}(\cdot)$ is continuous, additive-equivariant, and represents $\succsim$ on this subset of X. For any $x \in X$, define $c(x)$ as $c_{\underline{x}}(x)$ if $x \succsim \underline{x}$, and otherwise as $-d_x$, where d_x is the unique solution to:

$$\phi(d_x, x) \sim \underline{x}.$$

Note that such a d_x exists and is unique for each x by (N.3) and (N.2) respectively. Suppose $x \succsim y$. If $x \succsim \underline{x}$, then clearly $c(x) \geq c(y)$.[1] Consider then the case in which neither belongs to the $\underline{x}$ upper contour set. By (N.3) there exists $\alpha \geq 0$ such that:

$$\phi(\alpha, y) \sim x.$$

Then $\phi(d_x + \alpha, y) \sim \underline{x} \sim \phi(d_x, x)$, and by (N.2), $d_y = d_x + \alpha \geq d_x$, and therefore $c(x) \geq c(y)$. Thus c represents $\succsim$.

Let $\alpha \geq 0$. Since $c\big(\phi(\alpha, x)\big) = \alpha + c(x)$ if $x \succsim \underline{x}$, suppose instead $x \prec \underline{x}$. If $d_x \geq \alpha$, then:

$$\phi\big(d_x - \alpha, \phi(\alpha, x)\big) \sim \underline{x},$$

and hence $c\big(\phi(\alpha, x)\big) = -(d_x - \alpha) = c(x) + \alpha$. If, instead $\alpha > d_x$, then:

$$\phi(\alpha, x) = \phi\big(\alpha - d_x, \phi(d_x, x)\big) \sim \phi(\alpha - d_x, \underline{x}),$$

and thus $c\big(\phi(\alpha, x)\big) = \alpha - d_x + (0) = \alpha + c(x)$. Thus c is additive-equivariant.

Suppose now $x' \prec \underline{x}$. By hypothesis there is no $\succsim$-minimal element, hence there exists $y \in X$ such that $y \prec x' \prec \underline{x}$. Define $c_y(x)$ for all $x \succsim y$ as the unique solution to $\phi(c_y(x), y) \sim x$. By the preceding argument, c_y is continuous. Then for all $y \precsim x \prec \underline{x}$:

$$\phi\big(d_x, \phi(c_y(x), y)\big) \sim \underline{x}$$

by (N.1), thus

$$\phi\big(d_x + c_y(x), y\big) \sim \underline{x}.$$

By additive-equivariance of c_y:

$$d_x + c_y(x) + c_y(y) = c_y(\underline{x}),$$

[1] Either $y \succsim \underline{x}$ also and hence this follows from the preceding argument, or $y \prec \underline{x}$ in which case $c(x) \geq 0 > c(y)$.

and since $c_y(y) = 0$ by definition, re-arranging we obtain:

$$-d_x = c_y(x) - c_y(\underline{x}).$$

In particular, since $x \prec \underline{x}$, $-d_x = c(x)$. Thus:

$$c(x) = c_y(x) - c_y(\underline{x}).$$

Thus for any $y \prec \underline{x}$, the restriction of c to $\{x \in X : x \succsim y\}$ is continuous as it differs from the continuous function c_y by the constant, $c_y(\bar{x})$; hence c is continuous at x' in particular. Since every $x' \prec \underline{x}$ is contained within the upper contour set of some such y, we conclude that $c(x)$ is continuous.

Finally, suppose U and V are distinct, additive equivariant representations of $\succsim$. It suffices to show that the restrictions of U and V differ by a constant on any $\succsim$ upper contour set. Fix $\underline{x} \in X$ and let $x \succsim \underline{x}$, and suppose $\phi(\alpha, \underline{x}) \sim x$. Then: $U(x) - U(\underline{x}) = \alpha = V(x) - V(\underline{x})$, hence $U(x) = V(x) + \big[U(\underline{x}) - V(\underline{x})\big]$, implying the result. $\square$

Proof of Theorem 5

Theorem 2. Suppose an agent has preferences $\succsim$ on X that satisfy (N.1) - (N.3), and preferences $\succsim^*$ over X^* that are consistent with $\succsim$. Then choosing to submit a bid equal to their true compensation difference, in the mechanism corresponding to the more-preferred alternative, is $\succsim^*$-optimal.

Proof. Without loss of generality, let $x \succsim y$, with true compensation difference given by $\alpha \geq 0$, $\phi(\alpha, y) \sim x$. Since $\succsim$ satisfies (N.2) and (N.3), this α exists and is unique. Suppose first that the subject chooses to participate in the y-mechanism and submits a price of s. Then their state-dependent payoff is:

$$f_s(b_x, b_y, z) = \begin{cases} \phi(b_x, y) & \text{if } z = x \\ \phi(b_y, x) & \text{if } z = y,\ b_y \geq s \\ y & \text{if } z = y,\ s > b_y. \end{cases}$$

Similarly, if the agent instead submitted s in the x-mechanism, their reward would be:

$$g_s(b_x, b_y, z) = \begin{cases} \phi(b_y, x) & \text{if } z = y \\ \phi(b_x, y) & \text{if } z = x,\ b_x \geq s \\ x & \text{if } z = x,\ s > b_x \end{cases}$$

Suppose $s = \alpha$. By (N.2):

$$\phi(b_x, y) \succsim x \iff b_x \geq \alpha,$$

hence conditional upon $z = x$, the agent obtains $\max\{\phi(b_x, y), x\}$ from g_α.[2] Now, by (N.2), $\phi(b_y, x) \succsim y$ no matter the value of b_y, hence by consistency of $\succsim^*$ the most-preferred f act resulting from a bid in the y-mechanism is f_0.[3] Thus we wish to show $g_\alpha \succsim^* f_0$. But conditional upon $z = y$, both g_α and f_0 yield $\phi(b_y, x)$, and conditional upon $z = x$, g_α yields $\max\{\phi(b_x, y), x\}$ whereas f_0 yields $\phi(b_x, y)$. Thus by consistency, $g_\alpha \succsim^* f_0$. The final step is to show that $g_\alpha \succsim^* g_s$ for all other choices of s. This follows from the standard argument characterizing weak optimality of truthful bidding in Vickrey auctions, and we omit it. □

[2]The max here is understood in the preference sense.
[3]That is, it comes from setting $s = 0$.

Proof of Theorem 6

C.0.1 Overview

The proof of Theorem 6 proceeds in several steps. First, consider the case where X is homeomorphic to $X/{\sim_{\trianglelefteq}} \times \mathbb{R}_+$, and ϕ acts via addition along the second factor. Given a cardinally consistent data set, by 3 there exists some $u : \mathcal{V} \to \mathbb{R}$ such that $\operatorname{grad} u = \bar{Y}$. Without loss of generality, we may assume u is non-negative by adding a sufficiently large constant function. Define $\Delta : \mathcal{V} \to \mathbb{R}_+$ via $\Delta(v) = \|u\|_\infty - u_v$. If u is the restriction of any additive-equivariant utility U, then for all $v \in \mathcal{V}$:

$$U\big(\phi(\Delta(v), v)\big) = U(v) + \Delta(v) = U(v) + \big(\|u\|_\infty - U(v)\big) = \|u\|_\infty.$$

Thus by adding $\Delta(v)$ units of numeraire to each v, we obtain a collection of alternatives in X, $\{\phi\big(\Delta(v), v\big)\}_{v\in\mathcal{V}}$ that any additive-equivariant extension of u must be indifferent over. However, since $X = X/{\sim_{\trianglelefteq}} \times \mathbb{R}_+$ up to homeomorphism, we may view this set as the graph of a function $i : q(\mathcal{V}) \to \mathbb{R}_+$. Since $q(\mathcal{V})$ is a closed set, the Tietze extension theorem guarantees the existence of some continuous function $I : X/{\sim_{\trianglelefteq}} \to \mathbb{R}_+$ extending i. Whereas the graph of i was an 'incomplete' indifference curve, the graph of I supplies a 'complete' version. We then define a binary relation $\succsim$ on X by essentially translating the indifference curve given by the graph of I forward and backward using ϕ, and verify it possesses the desired structure.

However, X need not have such convenient structure. Hence the first section of the proof is dedicated to establishing that, even though X may itself not (up to homeomorphism) have any product structure, if (A.1) - (A.3) hold, then there is an equivariant embedding $\bar{s}$ of $X/{\sim_{\trianglelefteq}} \times \mathbb{R}_+$ into X such that for all $\alpha \in \mathbb{R}_+$ and all $y \in X/{\sim_{\trianglelefteq}}$, $q \circ \bar{s}(y, \alpha) = y$. Lemmas 1 to 5 verify the majority of the basic properties behind this construction. Lemmas 6-8 are of a technical nature and together establish the continuity of the inverse of $\bar{s}$, which proves it is an embedding, as claimed.

In the general case, we then work in $X/{\sim_{\trianglelefteq}} \times \mathbb{R}_+$ and proceed as before, obtaining a function I whose graph serves as an indifference curve for the preference we seek to construct. Here however, we then embed the graph of I into X using $\bar{s}$, and then once again define a relation $\succsim$ by 'translating' it using ϕ. We show that this still defines a continuous relation which both satisfies (N.1) - (N.3) and is consistent with the observed data.

C.0.2 Construction of Embedding

Lemma 4. *Let ϕ be a continuous action of $\mathbb{R}_+$ on X satisfying (A.1). Define the relation $x \sim_{\trianglelefteq} y$ if either:*

$$\exists\, \alpha \geq 0 \text{ s.t. } \phi(\alpha, x) = y,$$

or

$$\exists\, \beta \geq 0 \text{ s.t. } \phi(\beta, y) = x.$$

Then $\sim_{\trianglelefteq}$ is an equivalence relation.

Proof. Clearly $\sim_{\trianglelefteq}$ is reflexive and symmetric, hence all that remains is to verify transitivity. Suppose $x \sim_{\trianglelefteq} y$ and $y \sim_{\trianglelefteq} z$. We proceed in three cases: first suppose that only one of x and z is reachable from y; without loss $x \trianglelefteq y \trianglelefteq z$. Then there exists $\alpha_{xy}, \alpha_{yz} \geq 0$ such that $\phi(\alpha_{xy}, x) = y$ and $\phi(\alpha_{yz}, y) = z$ then clearly $\phi(\alpha_{xy} + \alpha_{yz}, x) = z$ and hence $x \trianglelefteq z$. Thus suppose $y \trianglelefteq x$ and $y \trianglelefteq z$. Then there exists $\alpha_{yz}, \alpha_{yz} \geq 0$ such that $\phi(\alpha_{yx}, y) = x$ and $\phi(\alpha_{yz}, y) = z$. Without loss of generality let $\alpha_{yx} \leq \alpha_{yz}$, so:

$$\phi\big(\alpha_{yz} - \alpha_{yx}, \phi(\alpha_{yx}.y)\big) = z,$$

and thus

$$\phi(\alpha_{yz} - \alpha_{yx}, x) = z,$$

and we obtain $x \sim_{\trianglelefteq} z$. Finally, suppose $x \trianglelefteq y$ and $z \trianglelefteq y$. Then there exists $\alpha_{xy}, \alpha_{zy} \geq 0$ such that $\phi(\alpha_{xy}, x) = y = \phi(\alpha_{zy}, z)$. Without loss, let $\alpha_{xy} \leq \alpha_{zy}$. Then:

$$\begin{aligned} y &= \phi(\alpha_{zy}, z) \\ &= \phi(\alpha_{xy} + (\alpha_{zy} - \alpha_{xy}), z) \\ &= \phi\big(\alpha_{xy}, \phi(\alpha_{zy} - \alpha_{xy}, z)\big). \end{aligned}$$

But, by (A.1), $\phi(\alpha_{xy}, \cdot)$ is injective hence, $\phi(\alpha_{zy}-\alpha_{xy}, z) = x$ and therefore $x \sim_{\trianglelefteq} z$. $\square$

In light of 4, there is a well-defined quotient space $X/{\sim_{\trianglelefteq}}$. In all that follows, we will consider $X/{\sim_{\trianglelefteq}}$ endowed with its quotient topology.

Corollary 4. *Let $q : X \to X/{\sim_{\trianglelefteq}}$ denote the canonical quotient map. Then for all $\alpha \geq 0$, for all $x \in X$,*

$$q(x) = (q \circ \phi)(\alpha, x).$$

Lemma 5. *Suppose (A.1) and (A.2). Then any continuous cross section s is an embedding of $X/{\sim_{\trianglelefteq}}$ into X.*

Proof. By hypothesis, s is continuous. Suppose then that $s(y') = s(y)$ for $y, y' \in X/{\sim_{\trianglelefteq}}$. Then:

$$(q \circ s)(y') = (q \circ s)(y)$$

and hence $y = y'$ as s is a cross section; thus s is injective. Moreover, by hypothesis, $q|_{\text{range}(s)} : \text{range}(s) \to X/{\sim_{\trianglelefteq}}$ is an inverse and continuous as $X/{\sim_{\trianglelefteq}}$ carries the quotient topology. Hence s is open. $\square$

For some fixed cross section s, define $\bar{s} : \mathbb{R}_+ \times X/{\sim_{\trianglelefteq}} \to X$ via:

$$\bar{s}(\alpha, y) = \phi\big(\alpha, s(y)\big),$$

and let $\bar{X} = \text{range}(\bar{s})$. We wish to show that $\bar{s}$ is an equivariant embedding, where the $\mathbb{R}_+$ acts on the domain by addition along the first factor. Clearly equivariance holds by construction:

$$\begin{aligned}\phi\big(\beta, \bar{s}(\alpha, y)\big) &= \phi\big(\beta, \phi(\alpha, s(y))\big)\\ &= \phi\big(\beta + \alpha, s(y)\big)\\ &= \bar{s}(\beta + \alpha, y).\end{aligned}$$

In all that follows we will assume (A.1) and (A.2), and a fixed s and hence fixed $\bar{s}$.

Lemma 6. *Let $\bar{q} : \bar{X} \to X/{\sim_{\trianglelefteq}}$ be the restriction of q to $\bar{X}$. Then $\bar{q}$ is an open map.*

Proof. Let $U \subset \bar{X}$ be open. Then:

$$\begin{aligned}\bar{q}(U) &= \big\{y \in X/{\sim_{\trianglelefteq}} : \exists\, \alpha \geq 0 \text{ s.t. } \phi\big(\alpha, s(y)\big) \in U\big\}\\ &= s^{-1}\big(\{x \in \text{range}(s) : \exists\, \alpha \geq 0 \text{ s.t. } \phi(\alpha, x) \in U\}\big)\\ &= s^{-1}\big(\text{range}(s) \cap \big[\cup_{\alpha \geq 0} f_\alpha^{-1}(U)\big]\big),\end{aligned}$$

where $f_\alpha = \phi(\alpha, \cdot)$. But, for all $\alpha \geq 0$, $\phi(\alpha, \cdot)$ is continuous hence $\text{range}(s) \cap \big[\cup_{\alpha \geq 0} f_\alpha^{-1}(U)\big]$ is a relatively open subset of range(s). Hence by 5, $\bar{q}(U)$ is open. □

Lemma 7. *Suppose that, for all $x \in X$, $\phi(\cdot, x)$ is injective. Then $\bar{s}$ is injective.*

Proof. Suppose $\bar{s}(\alpha, y) = \bar{s}(\alpha', y')$. Then:

$$\begin{aligned}\phi\big(\alpha, s(y)\big) &= \phi\big(\alpha', s(y')\big)\\ (q \circ \phi)\big(\alpha, s(y)\big) &= (q \circ \phi)\big(\alpha', s(y')\big)\\ s(y) &= s(y')\\ y &= y'\end{aligned}$$

where the second-to-last equality follows from 4, and the last from invoking 5. As $\phi\big(\cdot, s(y)\big)$ is injective, $\alpha = \alpha'$, and hence $\bar{s}$ is injective. □

For the remainder of this section, we will assume $\phi(\cdot, x)$ is injective for all x. Define $t : \bar{X} \to \mathbb{R}_+$ pointwise as the unique solution to:

$$\phi\big(t(x), (s \circ \bar{q})(x)\big) = x.$$

We will first show that the map $(t, \bar{q})$ is indeed the inverse of $\bar{s}$ (8). We then establish the regularity (i.e. continuity) of solutions to the above class of topological implicit function problems (9 - 11).

Lemma 8. *The map $(t, \bar{q}) : \bar{X} \to \mathbb{R}_+ \times X/{\sim_\trianglelefteq}$ is the inverse of $\bar{s}$.*

Proof. We will show $(t, \bar{q})$ is a left inverse. Thus let $(\alpha, y) \in \mathbb{R}_+ \times X/{\sim_\trianglelefteq}$. Then:

$$\begin{aligned}
\big((t, \bar{q}) \circ \bar{s}\big)(\alpha, y) &= \big((t \circ \bar{s})(\alpha, y), (\bar{q} \circ \bar{s})(\alpha, y)\big) \\
&= \big((t \circ \bar{s})(\alpha, y), (q \circ \phi)(\alpha, s(y))\big) \\
&= \big((t \circ \bar{s})(\alpha, y), y\big),
\end{aligned}$$

where the final equality follows from 4. Hence it remains to show $(t \circ \bar{s})(\alpha, y) = \alpha$. By definition of t,

$$\phi\big((t \circ \bar{s})(\alpha, y), (s \circ \bar{q} \circ \bar{s})(\alpha, y)\big) = \bar{s}(\alpha, y),$$

but by plugging in for $\bar{s}$ and appeal to 4, this simplifies to:

$$\phi\big((t \circ \bar{s})(\alpha, y), s(y)\big) = \phi\big(\alpha, s(y)\big).$$

Since $\phi\big(\cdot, s(y)\big)$ is injective, this implies $(t \circ \bar{s})(\alpha, y) = \alpha$ as desired. □

Lemma 9. *Suppose (A.1) - (A.3) and that ϕ is injective in its first factor. Then, for all $x \in \bar{X}$ there exists a finite open cover $\{N_{\alpha_i}\}_{i=1}^{K}$ of $\bar{s}\big([0, t(x)] \times \{\bar{q}(x)\}\big)$ with the following properties:*

1. *For all $i \in \{1, \ldots, K\}$, the set $\{\alpha : \bar{s}(\alpha, \bar{q}(x)) \in N_{\alpha_i}\}$ is a (relatively) open interval of $[0, \infty)$. For $i > 1$, denote this by $(\underline{\alpha}_i, \bar{\alpha}_i)$, and for $i = 1$, by $[0, \bar{\alpha}_1)$.*

2. *The indices* $\{\alpha_i\}_{i=1}^K$ *satisfy* $0 = \alpha_1 < \alpha_2 < \cdots < \alpha_K = t(x)$*, satisfy* $\alpha_i \in (\underline{\alpha}_i, \bar{\alpha}_i)$*, and, for all* $i, j = 1, \ldots, K$*,* $\alpha_i < \alpha_j$ *implies* $(\underline{\alpha}_i, \bar{\alpha}_i) \preceq_{SSO} (\underline{\alpha}_j, \bar{\alpha}_j)$*, where* $\prec_{SSO}$ *denotes the strong set order.*

3. *For all* i*,* N_{α_i} *satisfies the no loitering property of (A.3).*

Proof. Fix $x \in \bar{X}$. For all $\alpha \in [0, t(x)]$, define $x_\alpha = \bar{s}(\alpha, \bar{q}(x)) = \phi\big(\alpha, (s \circ \bar{q})(x)\big)$. By (A.3), for all $\alpha \in [0, t(x)]$, there exists $\varepsilon_\alpha, T_\alpha > 0$ such that, for all $x' \in B_{\varepsilon_\alpha}(x_\alpha)$, for all $\beta > T_\alpha$, $\phi(\beta, x') \notin B_{\varepsilon_\alpha}(x_\alpha)$. For each α, let U_α denote the connected component of $B_{\varepsilon_\alpha}(x_\alpha) \cap \bar{s}\big([0, t(x)] \times \{\bar{q}(x)\}\big)$ that contains x_α, and define $N_\alpha = B_{\varepsilon_\alpha}(x_\alpha) \setminus \big[\bar{s}\big([0, t(x)] \times \{\bar{q}(x)\}\big) \setminus U_\alpha\big]$. As $[0, t(x)] \times \{\bar{q}(x)\}$ is compact in $\mathbb{R}_+ \times X/\sim_{\trianglelefteq}$, by continuity $\bar{s}\big([0, t(x)] \times \{\bar{q}(x)\}\big)$ is a compact and hence closed subset of $\bar{X}$. U_α is a relatively open subset of $\bar{s}\big([0, t(x)] \times \{\bar{q}(x)\}\big)$, hence $\bar{s}\big([0, t(x)] \times \{\bar{q}(x)\}\big) \setminus U_\alpha$ is relatively closed in $\bar{s}\big([0, t(x)] \times \{\bar{q}(x)\}\big)$ and therefore also closed in $\bar{X}$. Then for all α, N_α is an open neighborhood of x_α. Moreover, by 7, $\bar{s}\big(\cdot, \bar{q}(x)\big)$ is injective (and continuous) hence for all α, $\{\alpha' : \bar{s}\big(\alpha', \bar{q}(x)\big) \in N_\alpha\}$ is an open interval in $[0, t(x)]$.

As $\bar{s}\big([0, t(x)] \times \{\bar{q}(x)\}\big)$ is compact and covered by $\{N_\alpha\}_{\alpha \in [0, t(x)]}$, there exists a finite set $0 = \alpha_1 < \cdots < \alpha_K = t(x)$ such that $\{N_{\alpha_i}\}_{i=1}^K$ form a finite subcover. By construction, for each i, $\alpha_i \in (\underline{\alpha}_i, \bar{\alpha}_i)$. Moreover, since properties (1.) and (3.) held for every element of $\{N_\alpha\}$ they hold for $\{N_{\alpha_i}\}$. Finally, it is without loss of generality to suppose that for all $i \neq j$, the intervals $(\underline{\alpha}_i, \bar{\alpha}_i) \not\subseteq (\underline{\alpha}_j, \bar{\alpha}_j)$, as if not, then some proper subcover does, and passing to this subcover preserves properties (1.) and (3.).

Then it remains only to verify $\{N_{\alpha_i}\}$ has the property that $\alpha_i < \alpha_j$ implies $(\underline{\alpha}_i, \bar{\alpha}_i) \preceq_{SSO} (\underline{\alpha}_j, \bar{\alpha}_j)$. Since neither interval contains the other, if $\underline{\alpha}_i < \underline{\alpha}_j$, then it must be that $\bar{\alpha}_i < \bar{\alpha}_j$, which implies $(\underline{\alpha}_i, \bar{\alpha}_i) \preceq_{SSO} (\underline{\alpha}_j, \bar{\alpha}_j)$ as desired.[4] If instead

[4]Note that as no interval in the collection is a subset of any other, it can never be the case that $\underline{\alpha}_i = \underline{\alpha}_j$ or $\bar{\alpha}_i = \bar{\alpha}_j$, thus considering only strict inequalities suffices.

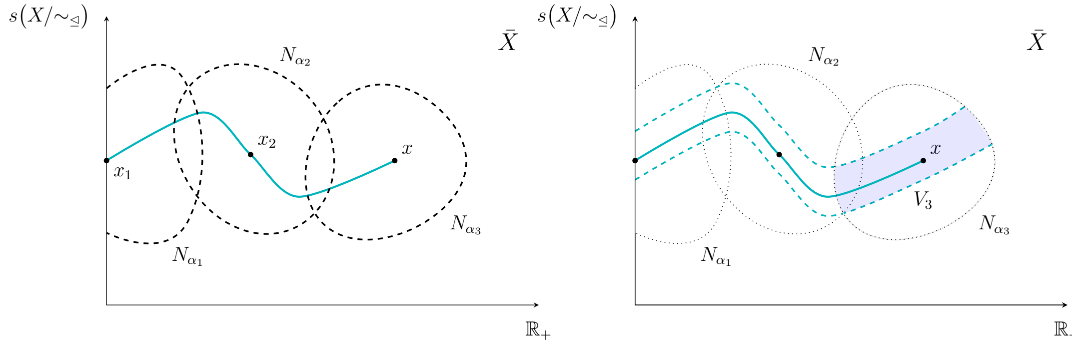

(a) An open cover of the path $\bar{s}\big([0, t(x)] \times \{\bar{q}(x)\}\big)$, here in aquamarine. This open cover satisfies all of the properties of 9.

(b) The construction of the neighborhood V_K (here, $K = 3$) for x on which t is bounded, from the open cover $\{N_{\alpha_i}\}_{i=1}^{3}$.

Figure C.1: An illustration of the construction underpinning Lemma 10. We have implicitly drawn the numeraire-paths of ϕ in $\bar{X}$ as vertical translates of one another.

$\underline{\alpha}_j < \underline{\alpha}_i$, then $\bar{\alpha}_j < \bar{\alpha}_i$, in which case $(\underline{\alpha}_j, \bar{\alpha}_j) \preceq_{SSO} (\underline{\alpha}_i, \bar{\alpha}_i)$, and hence $\alpha_i, \alpha_j \in (\underline{\alpha}_i, \bar{\alpha}_i) \cap (\underline{\alpha}_j, \bar{\alpha}_j)$. Thus swapping the labels of N_{α_i} and N_{α_j} preserves all salient properties but 'fixes' violations of property (2.). Repeating this process for each such pair cannot cycle (it simply sorts the indices via the $\{\underline{\alpha}_i\}$) and thus it terminates after some finite number of label swaps, resulting in a cover satisfying (2.). $\square$

Lemma 10. *Suppose (A.1) - (A.3) and that ϕ is injective in its first factor. Then for all $x \in \bar{X}$ there exists some open neighborhood of x on which t is bounded.*

Proof. Fix $x \in \bar{X}$, and let $\{N_{\alpha_i}\}_{i=1}^{K}$ denote an open cover of $\bar{s}\big([0, t(x)] \times \{\bar{q}(x)\}\big)$ of the form guaranteed by 9. Without loss of generality, suppose that N_{α_1} is the sole element to intersect $\bar{s}\big(\{0\} \times X/{\sim_{\trianglelefteq}}\big)$.[5] Define:

$$V_0 = \bar{s}\big(\{0\} \times X/{\sim_{\trianglelefteq}}\big)$$

and, for all $i = 1, \ldots, K$:

$$V_i = N_{\alpha_i} \bigcap \left[(\bar{q}^{-1} \circ \bar{q}) \left(\bigcup_{j<i} V_j \cap N_{\alpha_i} \right) \right],$$

see Figure C.1. We first verify, for all $i = 1, \ldots, K$, that V_i is open. Note that via 6 and our assumption that N_{α_1} is the only element of the open cover to intersect V_0, it suffices to show that V_1 is open. But

$$V_1 = N_{\alpha_1} \cap \big(\bar{q}^{-1} \circ \bar{q}\big)(V_0 \cap N_{\alpha_1}),$$

and $V_0 \cap N_{\alpha_1} = N_{\alpha_1} \cap \text{range}(s)$, and hence is relatively open in the range of s. As $\bar{q}$ is a left-inverse of s, $\bar{q}(N_{\alpha_1} \cap V_0)$ is open, and hence so too is V_1.

We now establish that, for all $i = 1, \ldots, K$,

$$\bar{s}\big([0, \bar{\alpha}_i) \times \{\bar{q}(x)\}\big) \subseteq \bigcup_{j \leq i} V_j,$$

where we recall that $(\underline{\alpha}_i, \bar{\alpha}_i) = \{\alpha \in [0, t(x)] : \bar{s}(\alpha, \bar{q}(x)) \in N_{\alpha_i}\}$ for $1 < i < K$, and $[0, \bar{\alpha}_i)$ is the analogue for $i = 1$.[6] For all $i = 1, \ldots, K$, let $x_{\alpha_i} = \bar{s}(\alpha_i, \bar{q}(x))$ and consider the case of $i = 1$. By hypothesis, $\alpha_1 = 0$, hence $x_{\alpha_1} = (s \circ \bar{q})(x) \in N_{\alpha_1} \cap V_0$. Then 4 implies $\bar{s}\big([0, t(x)] \times \{\bar{q}(x)\}\big) \subseteq (\bar{q}^{-1} \circ \bar{q})(N_{\alpha_1} \cap V_0)$, and thus $\bar{s}\big([0, \bar{\alpha}_1) \times \{\bar{q}(x)\}\big) \subseteq V_1$. Suppose now that, for all $1 \leq i \leq k$, that:

$$\bar{s}\big([0, \bar{\alpha}_i) \times \{\bar{q}(x)\}\big) \subseteq \bigcup_{j \leq i} V_j,$$

[5] For example, for all $i > 1$, redefine $N'_{\alpha_i} = N_{\alpha_i} \setminus \text{range}(s)$. N'_{α_i} is open as range(s) is closed: let $(x_n) \in \text{range}(s)$ and suppose $x_n \to x$. Then $q(x_n) \to q(x)$, and hence $(s \circ q)(x_n) \to (s \circ q)(x)$ by continuity. However, s is a cross-section thus, as $x_n \in \text{range}(s)$, x_n must be the value s takes at $q(x_n)$, hence $(s \circ q)(x_n) = x_n$ for all n. As X is metric and hence Hausdorff and as x_n converges to both x and $(s \circ q)(x)$, $(s \circ q)(x)$ must equal x, and thus $x \in \text{range}(s)$.

[6] This set is indeed an interval by 9.

but, for sake of contradiction, suppose that:

$$\bar{s}\big([0,\bar{\alpha}_{k+1})\times\{\bar{q}(x)\}\big)\not\subseteq\bigcup_{j\leq k+1}V_j.$$

As $(\underline{\alpha}_{k+1},\bar{\alpha}_{k+1})$ is an interval, if $\bar{\alpha}_k\in(\underline{\alpha}_{k+1},\bar{\alpha}_{k+1})$, the contradiction hypothesis would be false, thus it must be that $\bar{\alpha}_k\notin(\underline{\alpha}_{k+1},\bar{\alpha}_{k+1})$ and hence $\bar{s}(\bar{\alpha}_k,\bar{q}(x))\notin N_{\alpha_{k+1}}$. Then $(\underline{\alpha}_k,\bar{\alpha}_k)\cap(\underline{\alpha}_{k+1},\bar{\alpha}_{k+1})=\varnothing$. But 9 guarantees that, for all $l>k+1$, $\underline{\alpha}_l>\underline{\alpha}_{k+1}$, and for all $l<k$, $\bar{\alpha}_l<\bar{\alpha}_k$, hence:

$$\bar{s}(\bar{\alpha}_k,\bar{q}(x))\notin\bigcup_{i=1}^{K}N_{\alpha_i},$$

contradicting the fact that $\{N_{\alpha_i}\}_{i=1}^K$ is a cover for $\bar{s}\big([0,t(x)]\times\{\bar{q}(x)\}\big)$. Thus by induction $\bar{s}\big([0,\bar{\alpha}_K)\times\{\bar{q}(x)\}\big)\subseteq\bigcup_{j\leq K}V_j$, and in particular $x=x_{\alpha_K}\in V_K$.

We now verify that $t|_{V_i}$ is bounded for all $i=0,\dots,K$; since $x\in V_K$ and V_K is open, this suffices to establish the claim. For $i=0$ the claim is trivial as by definition, $t|_{V_0}$ is uniformly 0. Thus consider $i=1$, let $x'\in V_1$. Note that for any $\underline{x}'\in V_1$, if $\phi(\alpha,\underline{x}')=x'$, then $t(x')=\alpha+t(\underline{x}')$ by equivariance of $\bar{s}$.[7] But since N_{α_1} has a no-loitering bound of T_{α_1}, since both $x',\underline{x}'\in V_1\subseteq N_{\alpha_1}$,

$$t(x')<T_{\alpha_1}+t(\underline{x}').$$

However, if $x'\in V_1$, then $(s\circ\bar{q})(x')\in V_1$, and by definition $(t\circ s\circ\bar{q})(x')=0$. Thus for all $x'\in V_1$, $t(x')<T_{\alpha_1}$. Suppose now that, for all $i\leq k$, $t|_{V_i}$ is bounded, and let

[7] By definition,

$$\phi\big(t(\underline{x}'),(s\circ\bar{q})(\underline{x}')\big)=\underline{x}'$$

and, appealing to 4 to conclude $(s\circ\bar{q}\circ\phi)(\alpha,\underline{x}')=(s\circ\bar{q})(\underline{x}')$,

$$\phi\big((t\circ\phi)(\alpha,\underline{x}'),(s\circ\bar{q})(\underline{x}')\big)=\phi(\alpha,\underline{x}').$$

But as ϕ is an action:

$$\phi\big(\alpha+t(\underline{x}'),(s\circ\bar{q})(\underline{x}')\big)=\phi(\alpha,x')$$

too. By injectivity of ϕ in its first component, we conclude $t(\underline{x}')+\alpha=(t\circ\phi)(\alpha,\underline{x}')=t(x')$.

$x' \in V_{k+1}$. Then, $x' \in N_{\alpha_{k+1}}$ and there exists some $x'' \sim_{\trianglelefteq} x'$, where $x'' \in N_{\alpha_i} \cap V_j$ where $1 \leq j \leq k$. Suppose $x'' \trianglelefteq x'$. Then:

$$\begin{aligned} t(x') &< t(x'') + T_{\alpha_{k+1}} \\ &< \bar{T}_j + T_{\alpha_{k+1}} \\ &\leq \max_{i \leq k} \bar{T}_i + T_{\alpha_{k+1}}, \end{aligned}$$

where $T_{\alpha_{k+1}}$ is a no-loitering bound for $N_{\alpha_{k+1}}$, and $\bar{T}_j$ is any upper bound on $t|_{V_j}$ which exists by the induction hypothesis. Note that if $x' \trianglelefteq x''$, then $t(x')$ is bounded above by the same quantity. Thus for all $1 \leq i \leq K$, $t|_{V_i}$ is bounded; as $x \in V_K$ and V_K is open, this establishes the claim. □

Lemma 11. *Suppose (A.1) - (A.3) and that ϕ is injective in its first factor. Then t is continuous.*

Proof. Fix $x \in \bar{X}$. By 10, there exists $\varepsilon > 0$ such that $t|_{B_\varepsilon(x)}$ is bounded above by some constant K. Define $t^* : B_\varepsilon(x) \rightrightarrows \mathbb{R}_+$ via:

$$t^*(x') = \underset{\tilde{t} \in [0,K]}{\arg\min}\, d_X\big(\phi(\tilde{t}, (s \circ \bar{q})(x')), x'\big).$$

Since $t(x')$ is the unique unconstrained minimizer of this objective function, and $t(x') \in [0, K]$, it follows that $t^* = t|_{B_\varepsilon(x)}$ and hence t^* is a singleton-valued correspondence. But by the Theorem of the Maximum [6], t^* is upper hemicontinuous and hence continuous as a function. Thus for every $x \in X$ there is a restriction of t to some neighborhood of x on which it is continuous, hence it is continuous. □

Corollary 5. *Suppose (A.1)-(A.3), and that $\phi(\cdot, x)$ is injective for all $x \in X$. Then $\bar{s}$ is an equivariant embedding.*

Theorem 3. Let (X, ϕ) satisfy (A.1) - (A.3) and suppose Π_ϕ is non-empty. Then for every experiment $\mathcal{E}$, for any dataset, the following are equivalent:

(i) The data are cardinally consistent.

(ii) The data are rationalized by a continuous preference that satisfies (N.1) - (N.3).

(iii) The data are rationalized by a continuous, additive-equivariant utility function.

Proof. (i) $\implies$ (ii): Let $\mathcal{E}$ be an experiment, and $F \in \mathcal{F}$ a cardinally consistent flow on $(\mathcal{V}, \mathcal{E})$. By 3, there exists a utility $u \in \mathcal{U}$ such that $\operatorname{grad} u = F$. Without loss of generality we may suppose u is non-negative valued (by adding an appropriate constant that has no effect on its gradient). We may similarly without loss of generality assume that $\mathcal{V} \subsetneq \bar{X}$.[8] Then $i : \mathcal{V} \to \mathbb{R}_+$, where $i(v) = t(v) + (\|u\|_\infty - u_v)$ is well-defined. By definition of an experiment, $q(\mathcal{V})$ is in one-to-one correspondence with $\mathcal{V}$ hence we may instead view i as a map from $q(\mathcal{V}) \to \mathbb{R}_+$. By 5, $X/\sim_\trianglelefteq$ is homeomorphic to a subset of X and hence is metrizable and thus normal. By the Tietze extension theorem, e.g. [85], there exists a bounded, continuous function $I : X/\sim_\trianglelefteq \to \mathbb{R}_+$ such that $I|_{q(\mathcal{V})} = i$. The embedding of the graph of I under $\bar{s}$ will serve as a single 'full' indifference curve for the rationalizing preference we now construct.

We define a binary relation on X in three cases: first, if $x, y \in \bar{s}\big(\operatorname{epi}(I)\big) \subseteq \bar{X}$, then let $x \succsim y$ if and only if $t(x) - t(y) \geq (I \circ \bar{q})(x) - (I \circ \bar{q})(y)$.[9] If x but not y belong to $\bar{s}\big(\operatorname{epi}(I)\big)$, then let $x \succ y$. Finally, if neither x nor y belong to $\bar{s}\big(\operatorname{epi}(I)\big)$, then let:

$$x \succsim y \iff \min\{\alpha \in \mathbb{R}_+ : \phi(\alpha, y) \in \bar{s}\big(\operatorname{epi}(I)\big)\} \geq \min\{\alpha \in \mathbb{R}_+ : \phi(\alpha, x) \in \bar{s}\big(\operatorname{epi}(I)\big)\}.$$

Note that both minima are taken over closed sets that are bounded below and hence exist, thus the right-hand inequality is well-defined. As these cases are exhaustive, $\succsim$

[8] If it is not, there exists some $\bar{\alpha} > 0$ such that $\phi(\bar{\alpha}, \mathcal{V}) \subseteq \bar{X}$, and we may equivalently just work with this set of 'translates.'

[9] Note that here $t(x)$ and $t(y)$ are well defined because $x, y \in \bar{X}$.

is complete.[10] Now let $x \succsim y$ and $y \succsim z$, and suppose first that $x, y, z \in \bar{s}\big(\text{epi}(I)\big)$. Then

$$t(x) - t(y) \geq (I \circ \bar{q})(x) - (I \circ \bar{q})(y),$$

and

$$t(y) - t(z) \geq (I \circ \bar{q})(y) - (I \circ \bar{q})(z),$$

hence summing: $t(x) - t(z) \geq (I \circ \bar{q})(x) - (I \circ \bar{q})(z)$ and thus $x \succsim z$. By construction, if $x, y \in \bar{s}\big(\text{epi}(I)\big)$ but z is not, then $x \succsim z$, and by definition it is impossible that $y, z \in \bar{s}\big(\text{epi}(I)\big)$ but x is not, as $x \succsim y$ by hypothesis. Suppose finally that $x, y, z \notin \bar{s}\big(\text{epi}(I)\big)$. But then $x \succsim z$ by the transitivity of the usual order on $\mathbb{R}_+$. Thus $\succsim$ is transitive and hence a preference relation.

We now establish that $\succsim$ is continuous. First, let $x \in \bar{s}\big(\text{epi}(I)\big)$. Then, noting that $y \succsim x$ only if $y \in \bar{s}\big(\text{epi}(I)\big)$:

$$\begin{aligned}\{y \in X : y \succsim x\} &= \{y \in \bar{X} : t(y) - t(x) \geq (I \circ \bar{q})(y) - (I \circ \bar{q})(x)\}, \\ &= \{y \in \bar{X} : t(y) - (I \circ \bar{q})(y) \geq t(x) - (I \circ \bar{q})(x)\},\end{aligned}$$

where we define $\delta_x \equiv t(x) - (I \circ \bar{q})(x) \geq 0$. Consider the function $I_x : X/{\sim_{\trianglelefteq}} \to \mathbb{R}$ where $I_x(y) = I(y) + \delta_x$. This is continuous as I is, and by definition, $\{y \in X : y \succsim x\} = \bar{s}\big(\text{epi}(I_x)\big)$. By 5, this set is closed as $\text{epi}(I_x)$ is. Similarly, $\{y \in X : y \succ x\} = \bar{s}\big(\text{int}\,\text{epi}(I_x)\big)$, hence it is open; as $\succsim$ is complete, $\{y \in X : y \precsim x\} = \bar{s}\big(\text{int}\,\text{epi}(I_x)\big)^c$ is closed.

[10]From the definition of $\bar{s}$, if $x \in \bar{X}$, $x' \in X$ and $\phi(\alpha, x) = x'$, then $x' \in \bar{X}$ as well. Thus, in particular, the embedding under $\bar{s}$ of the graph of I partitions X into $\bar{s}\big(\text{epi}(I)\big)$ and $\{x \in X : \exists \alpha \in \mathbb{R}_+ \text{ s.t. } \phi(\alpha, x) \in \bar{s}\big(\text{epi}(I)\big)\}$, and thus the cases we consider are exhaustive.

Suppose now that $x \notin \bar{s}\big(\text{epi}(I)\big)$, and let $\hat{\alpha}_x = \min\{\alpha \in \mathbb{R}_+ : \phi(\alpha, x) \in \bar{s}\big(\text{epi}(I)\big)\}$. Then $\phi(\hat{\alpha}_x, x) = \bar{s}\big((I \circ q)(x), q(x)\big)$, hence:

$$
\begin{aligned}
\{y \in X : y \succsim x\} &= \{y \in \bar{s}\big(\text{epi}(I)\big)^c : y \succsim x\} \cup \bar{s}\big(\text{epi}(I)\big) \\
&= \{y \in \bar{s}\big(\text{epi}(I)\big)^c : \hat{\alpha}_x \geq \hat{\alpha}_y\} \cup \bar{s}\big(\text{epi}(I)\big) \\
&= \{y \in \bar{s}\big(\text{epi}(I)\big)^c : y \in f_{\hat{\alpha}_x}^{-1}\big(\bar{s}\big(\text{epi}(I)\big)\big)\} \cup \bar{s}\big(\text{epi}(I)\big) \\
&= f_{\hat{\alpha}_x}^{-1}\big(\bar{s}\big(\text{epi}(I)\big)\big).
\end{aligned}
$$

where $f_{\hat{\alpha}_x} = \phi(\hat{\alpha}_x, \cdot)$. As $f_{\hat{\alpha}_x}$ is continuous and $\bar{s}\big(\text{epi}(I)\big)$ is closed, we conclude the weak upper contour set at x is closed. Analogously, the strict upper contour set at x is open, and therefore the weak lower contour set at x is closed too. As these cases are exhaustive, $\succsim$ is continuous.

We now verify $\succsim$ obeys (N.1) - (N.3). Suppose then that $x \succsim y$, and let $\alpha \geq 0$. If $x, y \in \bar{s}\big(\text{epi}(I)\big)$, then, as $\phi(\alpha, x) = \phi\big(\alpha, \phi(t(x), (s \circ q)(x))\big)$ (and likewise y):

$$
\begin{aligned}
(t \circ \phi)(\alpha, x) - (t \circ \phi)(\alpha, y) &= (t \circ \phi)\big(\alpha, \phi(t(x), (s \circ q)(x))\big) - (t \circ \phi)\big(\alpha, \phi(t(x), (s \circ q)(x))\big) \\
&= (t \circ \phi)\big(\alpha + t(x), (s \circ q)(x)\big) - (t \circ \phi)\big(\alpha + t(y), (s \circ q)(y)\big) \\
&= (t \circ \bar{s})(\alpha + t(x), q(x)) - (t \circ \bar{s})(\alpha + t(y), q(y)) \\
&= (\alpha + t(x)) - (\alpha + t(y)) \\
&= t(x) - t(y) \\
&\geq (I \circ \bar{q})(x) - (I \circ \bar{q})(y) \\
&= (I \circ \bar{q} \circ \phi)(\alpha, x) - (I \circ \bar{q} \circ \phi)(\alpha, y)
\end{aligned}
$$

where the inequality follows from $x \succsim y$ and $x, y \in \bar{s}\big(\text{epi}(I)\big)$. Thus $\phi(\alpha, x) \succsim \phi(\alpha, y)$, as $\phi(\alpha, x), \phi(\alpha, y) \in \bar{s}\big(\text{epi}(I)\big)$. Suppose now x but not y belongs to $\bar{s}\big(\text{epi}(I)\big)$ (and thus that $x \succ y$. Then for all $0 \leq \alpha < \hat{\alpha}_y$, by definition $\phi(\alpha, x) \succ \phi(\alpha, y)$, hence suppose $\alpha \geq \hat{\alpha}_y$. Then as shown above, $(t \circ \phi)(\alpha, x) = t(x) + \alpha$, where $t(x) \geq (I \circ q)(x)$.

Similarly, since $y \notin \bar{s}\big(\text{epi}(I)\big)$, $(t \circ \phi)(\alpha, y) < (I \circ q)(y) + \alpha$. Hence:

$$\begin{aligned}(t \circ \phi)(\alpha, x) - (t \circ \phi)(\alpha, y) &= t(x) + \alpha - (t \circ \phi)(\alpha, y) \\ &\geq (I \circ q)(x) + \alpha - (t \circ \phi)(\alpha, y) \\ &> (I \circ q)(x) - (I \circ q)(y) \\ &= (I \circ q \circ \phi)(\alpha, x) - (I \circ q \circ \phi)(\alpha, y),\end{aligned}$$

hence $\phi(\alpha, x) \succ \phi(\alpha, y)$. Finally, suppose neither x nor y belong to $\bar{s}\big(\text{epi}(I)\big)$. Let $x \succsim y$ hence $\hat{\alpha}_y \geq \hat{\alpha}_x$. For all $\alpha < \hat{\alpha}_x$, $\hat{\alpha}_{\phi(\alpha,x)} = \hat{\alpha}_x - \alpha$, thus for all such α, $\phi(\alpha, x) \succsim \phi(\alpha, y)$. If $\alpha \geq \hat{\alpha}_x$, then $\phi(\alpha, x) \in \bar{s}\big(\text{epi}(I)\big)$; if $\phi(\alpha, y)$ is not then the preceding argument implies $\phi(\alpha, x) \succ \phi(\alpha, y)$. If $\phi(\alpha, y) \in \bar{s}\big(\text{epi}(I)\big)$, then:

$$\begin{aligned}(t \circ \phi)(\alpha, x) - (t \circ \phi)(\alpha, y) &= \hat{\alpha}_y - \hat{\alpha}_x \\ &\geq (I \circ q \circ \phi)(\alpha, x) - (I \circ q \circ \phi)(\alpha, y).\end{aligned}$$

Thus $\succsim$ satsifies (N.1). Property (N.2) holds by definition. Thus now suppose $y \succsim x$. Then $\phi(\hat{\alpha}_x, x), \phi(\hat{\alpha}_x, y) \in \bar{s}\big(\text{epi}(I)\big)$, thus, having verified (N.1) it suffices to find some α such that:

$$\phi(\alpha + \hat{\alpha}_x, x) \sim \phi(\hat{\alpha}_x, y).$$

Let:

$$\alpha = (t \circ \phi)(\hat{\alpha}_x, y) - (I \circ q)(y).$$

Note this is well-defined as $\phi(\hat{\alpha}_x, y) \in \bar{s}\big(\text{epi}(I)\big)$. But, since $(t \circ \phi)(\hat{\alpha}_x, x) = (I \circ q)(x)$,

$$\begin{aligned}(t \circ \phi)(\alpha + \hat{\alpha}_x, x) - (t \circ \phi)(\hat{\alpha}_x, y) &= \alpha + (t \circ \phi)(\hat{\alpha}_x, x) - (t \circ \phi)(\hat{\alpha}_x, y) \\ &= \alpha + (I \circ q)(x) - (t \circ \phi)(\hat{\alpha}_x, y) \\ &= (I \circ q)(x) - (I \circ q)(y).\end{aligned}$$

Thus $\succsim$ satisfies (N.3),

We now verify that the compensation differences under $\succsim$ for each pair in $\mathcal{E}$ precisely corresponds to the observed data, our last outstanding claim. Let $F_{yx} \geq 0$. Suppose first $x, y \in \bar{s}\big(\text{epi}(I)\big)$. Recall $(I \circ \bar{q})(x) = i(x) = t(x) + \big(\|u\|_\infty - u_x\big)$ (and likewise y) as $x, y \in \mathcal{V}$. Thus:

$$\begin{aligned} t(x) - t\big(\phi(F_{yx}, y)\big) &= t(x) - t(y) - F_{yx} \\ &= t(x) - t(y) - (\operatorname{grad} u)_{yx} \\ &= t(x) - t(y) - (u_x - u_y) \\ &= (I \circ \bar{q})(x) - (I \circ \bar{q})(y) \\ &= (I \circ \bar{q})(x) - (I \circ \bar{q} \circ \phi)(F_{yx}, y) \end{aligned}$$

where the first equality follows from an argument identical to 7, and the final from 4. By hypothesis $y \in \bar{s}\big(\text{epi}(I)\big)$ hence so too is $\phi(F_{yx}, y)$, and thus $\phi(F_{yx}, y) \sim x$ by definition of $\succsim$. If x or y do not belong to belong to $\bar{s}\big(\text{epi}(I)\big)$, then by invariance of $\succsim$ it suffices to verify that $\phi(F_{yx} + \max\{\hat{\alpha}_x, \hat{\alpha}_y\}, y) \sim \phi(\max\{\hat{\alpha}_x, \hat{\alpha}_y\}, x)$. But as $\phi(\max\{\hat{\alpha}_x, \hat{\alpha}_y\}, y), \phi(\max\{\hat{\alpha}_x, \hat{\alpha}_y\}, x) \in \bar{s}\big(\text{epi}(I)\big)$, this follows from the above case. Thus (i) $\implies$ (ii).

That (ii) $\implies$ (iii) follows from Theorem 4, so it remains only to prove (iii) $\implies$ (i). Let $x_0, \dots, x_L \in X$, and suppose that $U : X \to \mathbb{R}$ is additive-equivariant. Clearly:

$$\sum_{l=0}^{L} U(x_{l+1}) - U(x_l) = 0,$$

where subscripts are understood mod-L. Let:

$$\alpha_l = \begin{cases} \alpha_{l,l+1} & \text{if } x_{l+1} \sim \phi(\alpha_{l,l+1}, x_l) \\ -\alpha_{l,l+1} & \text{if } x_l \sim \phi(\alpha_{l,l+1}, x_{l+1}). \end{cases}$$

Then, for all $l = 0, \dots L$, if $U(x_{l+1}) \geq U(x_l)$:

$$U(x_{l+1}) - U(x_l) = U(x_l) + \alpha_{l,l+1} - U(x_l) = \alpha_l,$$

and if $U(x_l) \geq U(x_{l+1})$:

$$U(x_{l+1}) - U(x_l) = U(x_{l+1}) - \left[U(x_{l+1}) + \alpha_{l,l+1}\right] = \alpha_l.$$

Thus:

$$0 = \sum_{l=0}^{L} U(x_{l+1}) - U(x_l) = \sum_{l=0}^{L} \alpha_l.$$

Thus the compensation differences arising from any additive-equivariant utility will always be cardinally consistent. □

Remark 6. Conditions (A.2) - (A.3) are also necessary in the following sense. Suppose X is a metric space, ϕ a continuous action of $\mathbb{R}_+$ on X, and that (A.1) holds so that (A.2) is well-defined. Then if there exists an equivariant embedding $\hat{s} : \mathbb{R}_+ \times X/{\sim_{\trianglelefteq}} \to X$ (where the action of $\mathbb{R}_+$ on $\mathbb{R}_+ \times X/{\sim_{\trianglelefteq}}$ is simply addition along the first factor), then (A.2) and (A.3) must hold. This suggests that the technical conditions of Theorem 6 cannot be significantly relaxed without requiring a completely different proof approach.

Proposition Proofs

C.0.3 Proof of 3

Proof. Suppose first that $F \in \operatorname{im}(\operatorname{grad})$. Then there exists $u \in \mathcal{U}$ such that $\operatorname{grad} u = F$. Let $(v^0, v^1), (v^1, v^2), \ldots, (v^L, v^0) \in \vec{\mathcal{E}}$. Then:

$$\sum_{l=0}^{L} F_{v^l v^{l+1}} = \sum_{l=0}^{L} (\operatorname{grad} u)_{v^l v^{l+1}} = \sum_{l=0}^{L} (u_{v^{l+1}} - u_{v^l}) = 0.$$

Conversely, suppose F is cardinally consistent. Let $(\mathcal{V}, \mathcal{E}')$ denote a spanning tree for $(\mathcal{V}, \mathcal{E})$. Fix $\underline{v} \in \mathcal{V}$. Then for each $v \neq \underline{v}$, there is a unique sequence of edges in $\vec{\mathcal{E}}'$:

$$(\underline{v}, v^1), (v^1, v^2), \ldots, (v^L, v)$$

connecting $\underline{v}$ to v. Define $u(\underline{v}) = 0$ and:

$$u(v) = F_{\underline{v}v^1} + \sum_{l=1}^{L-1} F_{v^l v^{l+1}} + F_{v^L v}$$

The utility u is well-defined and does not depend on the choice of spanning tree: this follows from observing that if, for two different choices of spanning tree, the sums of F along two different paths from $\underline{v}$ to v differed, then by reversing one of the paths, one would obtain a violation of cardinal consistency. Finally, by construction, $\text{grad}\, u = F$, completing the proof. □

C.0.4 Proof of 4

Proof. By the fundamental theorem of linear algebra, it suffices to verify that, for all $u \in \mathcal{U}$, $F \in \mathcal{F}$:

$$\langle -\text{div}\, F, u \rangle = \langle F, \text{grad}\, u \rangle,$$

where $\mathcal{U}$ carries its standard Euclidean inner product. Then:

$$\begin{aligned}
\langle F, \text{grad}\, u \rangle &= \sum_{\{(i,j) \in \vec{\mathcal{E}}\, :\, i<j\}} F_{ij} \left[u_j - u_i \right] \\
&= \sum_{\{(i,j) \in \vec{\mathcal{E}}\, :\, i<j\}} F_{ij} u_j + \sum_{\{(i,j) \in \vec{\mathcal{E}}\, :\, i<j\}} F_{ji} u_i \\
&= \sum_{\{(i,j) \in \vec{\mathcal{E}}\, :\, j<i\}} F_{ji} u_i + \sum_{\{(i,j) \in \vec{\mathcal{E}}\, :\, i<j\}} F_{ji} u_i \\
&= \sum_{(i,j) \in \vec{\mathcal{E}}} F_{ji} u_i \\
&= \sum_{i \in \mathcal{V}} \left[\sum_{j \in N(i)} F_{ji} \right] u_i \\
&= \sum_{i \in \mathcal{V}} \left[- \sum_{j \in N(i)} F_{ij} \right] u_i \\
&= \langle -\text{div}\ F, u \rangle,
\end{aligned}$$

where the third to last line follows from the observation that summing over $\vec{\mathcal{E}}$ (i.e. each edge twice, once with each orientation) is equivalent to summing over, for each v_i, all of the edges connecting v_i to its neighbors, oriented away from v_i. Thus $\mathcal{F}$ admits an orthogonal decomposition as $\text{im}(\text{grad}) \oplus \ker(\text{div})$. $\square$

C.0.5 Proof of 5

Proof. As noted in the text, for a pure cycle C, $MP(C) = \|C\|_1$. Thus if $R = \sum_l C_l$ for some $\{C_1, \ldots, C_L\} \in \mathfrak{D}(R)$, then by the triangle inequality:

$$\|R\|_1 = \left\| \sum_{l=1}^{L} C_l \right\|_1 \leq \sum_{l=1}^{L} \left\| C_l \right\|_1 = \sum_{l=1}^{L} MP(C_l).$$

Taking infimums across all such decompositions of R we obtain $\|R\|_1 \leq MP^*(R)$. Thus it suffices to show that there always exists a decomposition in $\mathfrak{D}(R)$ attaining this lower bound.

Without loss of generality, suppose $R \geq 0$ componentwise.[11] If $R = 0$ then trivially $MP^*(R) = \|R\|_1 = 0$, hence suppose $R \neq 0$. Let $\mathcal{E}'$ denote the subset of edges on which $R \neq 0$, and let $\mathcal{V}' = \cup_{\{x,y\} \in \mathcal{E}'} \{x, y\}$ denote the associated vertex set. Choose $v^0 \in \mathcal{V}'$ arbitrarily. Since $v^0 \in \mathcal{V}'$ and $R \neq 0$, there exists some $v^1 \in N(v_0)$ such that $R_{v^0 v^1} \neq 0$. Since $R \in \ker(\text{div})$, v_1 may be chosen so that $R_{v^0 v^1} > 0$. Proceeding analogously we may construct a sequence of oriented edges in $\vec{\mathcal{E}'}$ such that $R_{v^j v^{j+1}} > 0$. We terminate this process when we choose a vertex that has appeared prior in the sequence.[12] Possibly by throwing out some initial segment of this sequence and relabelling indices, we obtain a sequence of oriented edges $(v^0, v^1), (v^1, v^2), \ldots, (v^{J_1}, v^0)$ such that $R_{v^j v^{j+1}} > 0$. Let $c_1 = \min_j R_{v^j v^{j+1}}$, and let $C_1 = \sum_{j=0}^{J_1} c_1 \mathbb{1}_{(v^j, v^j+1)}$. Then

[11] This simply amounts to a choice of orientation of each edge forming our basis for $\mathcal{F}$ in the same direction as the flow (if the flow is non-zero).

[12] This process necessarily termiantes as $(\mathcal{V}, \mathcal{E})$ is finite.

$0 \leq C_1 \leq R$ component-wise, and C_1 is equal to R on at least one component. Thus $R^1 = R - C_1$ also belongs to the positive cone of ker(div); however it is supported on a strict subgraph of $(\mathcal{V}', \mathcal{E}')$. Thus repeating this process, we obtain a finite decomposition $R = C_1 + \cdots + C_L$, where for all l, $C_l \geq 0$. Since every $C_l \geq 0$, however:

$$\|R\|_1 = \left\| \sum_{l}^{L} C_l \right\|_1 = \sum_{l=1}^{L} \|C_l\|_1 = \sum_{l=1}^{L} MP(C_l)$$

and hence the lower bound obtains. □

Shape Constraint 'Cookbook'

Quaslinear, Increasing, Concave Utility (Proof of 22)

Proof. Suppose first that U is a quasilinear, increasing, and concave utility. For all $i = 1, \ldots, K$, define $u_i = U(v_i)$ and let π_i denote an arbitrary choice of supergradient of U at each v_i. As U is increasing, it follows $\pi_i \geq 0$ for each i. Define $\gamma_i = u_i - \langle \pi_i, v_i \rangle$. Then for all $i = 1, \ldots, K$ and all $x \in X$:

$$U(x) \leq U(v_i) + \langle \pi_i, x - v_i \rangle.$$

Thus, in particular, $\langle \pi_i, v_i \rangle + \gamma_i \leq \langle \pi_j, v_i \rangle + \gamma_j$ for all i, j. Finally, as:

$$U\big(\phi(\alpha, v_i)\big) \leq U(v_i) + \langle \pi_i, (\alpha, 0) \rangle$$

it follows that:

$$\alpha \leq \pi_i^1 \alpha$$

hence $\pi^1 \geq 1$. If v_i is on the interior of $\mathbb{R}_+^2$ then there is some $\hat{v}$ such that, for some $\hat{\alpha} > 0$, $\phi(\hat{\alpha}, \hat{v}) = v_i$. Thus $U(\hat{v}) = U(v_i) - \alpha$, and:

$$U(\hat{v}) \leq U(v_i) + \langle \pi_i, (-\alpha, 0) \rangle,$$

which yields $-\alpha \leq -\alpha\pi_i^1$ and hence $\pi_i^1 \leq 1$. Thus for all v in the interior of X, their supergradients must have first component equal to 1. By the outer hemicontinuity of the supergradient correspondence ([66], Theorem 6.2.4) this remains true for those v on the boundary of X, and hence for all v_i, π_i is of the form $(1, \pi_i^2)$ as claimed.

Conversely, suppose $u, \{\pi_i\}_{i=1}^K, \{\gamma_i\}_{i=1}^K$ is a solution to (3.9). Define:

$$\tilde{U}(x) = \min_{i \in \{1,\ldots,K\}} \gamma_i + \langle x, \pi_i \rangle.$$

Then clearly $U(v_i) = u_i$, and $\tilde{U}$ is quasilinear, increasing, and concave. □

Cobb-Douglas Preferences (Proof of 23)

Proof. Consider $X = \mathbb{R}^L_{++}$, and $\phi(\alpha, x) = e^\alpha x$. Define $H : X \to \mathbb{R}^L$ via:

$$H(x) = \big(\ln x_1, \ldots, \ln x_L\big).$$

The transformation H induces an action of $\mathbb{R}_+$ on $\mathbb{R}^L$ via $\tilde{\phi}(\alpha, H(x)) = H(\phi(\alpha, x))$, here given by:

$$\tilde{\phi}(\alpha, H(x)) = H(x) + \alpha \mathbb{1}_L.$$

Critically, (X, ϕ) and $(\mathbb{R}^L, \tilde{\phi})$ are isomorphic in the above sense, and hence there is a one-to-one correspondence between observations of the form:

$$\phi(\alpha, x) \sim y$$

with:

$$\tilde{\phi}(\alpha, H(x)) \sim H(y).$$

A collection of observations of this latter form is rationalized by an affine utility on $H(X)$ (with gradient in $\Delta(L)$) if and only if the former form is rationalized by a

Cobb-Douglas utility, hence (3.7) under change of coordinates becomes:

$$\begin{aligned} \min_{u \in \mathcal{U}} \quad & \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2 \\ \text{subject to} \quad & u_i = \left\langle \beta, H(v_i) \right\rangle \quad \forall i = 1, \ldots, K \\ & \beta \geq 0 \end{aligned} \tag{C.6}$$

for $\beta \in \mathbb{R}^L$. Note that at any feasible solution (u, β), additive-equivariance implies $\langle \beta, \mathbb{1}_L \rangle = 1$ hence this condition would be redundant to include. □

Risk-Neutral Utility Functionals on $\mathbb{R}^S$

Let S be a finite set of states of the world and let $X = \mathbb{R}^S$ denote the space of monetary acts, along with $\phi(\alpha, x) = x + \alpha \mathbb{1}_S$. Let $(\mathcal{V}, \mathcal{E})$ denote an experiment; recall by definition, there does not exist any pair $v_i, v_j \in \mathcal{V}$ such for which there is some $\alpha \geq 0$ such that $\phi(\alpha, v_i) = v_j$. In light of 4, we will drop the 'risk-neutral' qualifier as it is understood that these characterizations may be straightforwardly extended to other Bernoulli utilities, and these may be estimated in advance via Equation 3.14.

Subjective Expected Utility

A map $U : X \to \mathbb{R}$ is said to be a subjective expected utility functional if it is of the form:

$$U(x) = \langle \pi, x \rangle,$$

for some $\pi \in \Delta(S)$. Define $\mathcal{K}_{\text{SEU}}$ as the collection of $u \in \mathcal{U}$ that are restrictions of subjective expected utility functionals. Then solving (3.7) with $\mathcal{K} = \mathcal{K}_{\text{SEU}}$ is equivalent to solving:

$$\begin{aligned}
\min_{u \in \mathcal{U}} \quad & \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2 \\
\text{subject to} \quad & u_i = \langle \pi, v_i \rangle \qquad \forall i = 1, \dots, K \\
& \langle \pi, \mathbb{1}_S \rangle = 1 \\
& \pi \geq 0.
\end{aligned} \tag{C.7}$$

Proof. Trivial. □

CHOQUET EXPECTED UTILITY

Recall that a function $\nu : 2^S \to \mathbb{R}$ is a capacity if (i) $\nu(\varnothing) = 0$, $\nu(S) = 1$, and (ii) for all $A \subseteq B$, $\nu(A) \leq \nu(B)$. By abuse of notation, let $S = \{1, \dots, S\}$, and let $\mathfrak{S}_S$ denote the set of permutations on $\{1, \dots, S\}$. For each $\sigma \in \mathfrak{S}_S$, define:

$$C_\sigma = \{x \in \mathbb{R}^S : x_{\sigma(1)} \geq x_{\sigma(2)} \geq \cdots \geq x_{\sigma(S)}\}. \tag{C.8}$$

The cones $\{C_\sigma\}_{\sigma \in \mathfrak{S}_S}$ cover $\mathbb{R}^S$. Note that if a functional $U : \mathbb{R}^S \to \mathbb{R}$ corresponds to Choquet integration with respect to ν, then for any σ, $U|_{C_\sigma}$ is linear, and indeed if $x \in C_\sigma$, then:

$$U(x) = \int_S x \, dP^\sigma,$$

where, for all $i = 1, \dots, S$, the probability measure P^σ is defined by:

$$P^\sigma(\sigma(i)) = \nu\big(\{\sigma(1), \sigma(2), \dots, \sigma(i)\}\big) - \nu\big(\{\sigma(1), \sigma(2), \dots, \sigma(i-1)\}\big). \tag{C.9}$$

See [51] for more discussion. Finally, for notational simplicity, define the shorthand A_i^σ for the set $\{\sigma(1), \sigma(2), \dots, \sigma(i)\}$.

We say that $U : X \to \mathbb{R}$ is said to be a Choquet expected utility (CEU) functional if:

$$U(x) = \int_S x \, d\nu,$$

where ν is a capacity the integral denotes Choquet integration. Define $\mathcal{K}_{CEU}$ as the collection of $u \in \mathcal{U}$ that are restrictions of CEU functionals. Then solving (3.7) with $\mathcal{K} = \mathcal{K}_{\text{CEU}}$ is equivalent to solving:

$$\begin{aligned}
\min_{u \in \mathcal{U}} \quad & \left\|(\operatorname{grad} u) - \bar{Y}\right\|_2^2 \\
\text{subject to} \quad & u_i = \langle P^\sigma, v_i \rangle && \forall \sigma \in \mathfrak{S}_S, \forall i = 1, \ldots, K \text{ s.t. } v_i \in C^\sigma \\
& P^\sigma_{\sigma(j)} = \nu_{A^\sigma_j} - \nu_{A^\sigma_{j-1}} && \forall \sigma \in \mathfrak{S}_S, \forall j = 1, \ldots, S \\
& \nu_A \leq \nu_B && \forall A, B \in 2^S \text{ s.t. } A \subseteq B \\
& \nu_\varnothing = 0 \\
& \nu_S = 1
\end{aligned} \tag{C.10}$$

Proof. Suppose U is a CEU functional. Then it corresponds to integration against some capacity ν which by definition then satisfies the last three constraints of (C.10). From the discussion, e.g., in [51] (see, in particular, Example 17), each v_i belongs to at least one C_σ cone, and restricted to each, U simply amounts to integration (i.e. a dot product) of v_i with the measure P^σ. Hence every CEU functional corresponds to a solution to (C.10). Conversely, it follows trivially that every solution to (C.10) defines a CEU functional. □

Convex Choquet Expected Utility

A capacity $\nu : 2^S \to \mathbb{R}$ is said to be a convex, if, for all $A, B \subseteq S$:

$$\nu(A) + \nu(B) \leq \nu(A \cap B) + \nu(A \cup B).$$

A map $U : X \to \mathbb{R}$ is said to be a convex Choquet expected utility (CCEU) functional if it is of the form:

$$U(x) = \int_S x \, d\nu,$$

for some convex capacity ν. Define $\mathcal{K}_{CCEU}$ as the collection of $u \in \mathcal{U}$ that are restrictions of CCEU functionals. Then, solving (3.7) with $\mathcal{K} = \mathcal{K}_{\text{CCEU}}$ is equivalent to solving:

$$\begin{aligned}
\min_{u \in \mathcal{U}} \quad & \left\|(\operatorname{grad} u) - \bar{Y}\right\|_2^2 \\
\text{subject to} \quad & u_i = \langle P^\sigma, v_i \rangle && \forall \sigma \in \mathfrak{S}_S, \forall i = 1, \dots, K \text{ s.t. } v_i \in C^\sigma \\
& P^\sigma_{\sigma(j)} = \nu_{A^\sigma_j} - \nu_{A^\sigma_{j-1}} && \forall \sigma \in \mathfrak{S}_S, \forall j = 1, \dots, S \\
& \nu_A \leq \nu_B && \forall A, B \in 2^S \text{ s.t. } A \subseteq B \\
& \nu_A + \nu_B \leq \nu_{A \cup B} + \nu_{A \cap B} && \forall A, B \in 2^S \\
& \nu_\varnothing = 0 \\
& \nu_S = 1
\end{aligned} \tag{C.11}$$

Proof. Follows from CEU case, where additionally the supermodularity of ν is enforced. □

Maxmin Expected Utility

A map $U : X \to \mathbb{R}$ is said to be a maxmin expected utility (MEU) functional if it is of the form:

$$U(x) = \min_{\pi \in P} \langle \pi, x \rangle,$$

for some compact, convex belief set $P \subseteq \Delta(S)$. Define $\mathcal{K}_{\text{MEU}}$ as the collection of $u \in \mathcal{U}$ that are restrictions of MEU functionals. Then solving (3.7) with $\mathcal{K} = \mathcal{K}_{\text{MEU}}$

is equivalent to solving:

$$\begin{aligned}
\min_{u \in \mathcal{U}} \quad & \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2 \\
\text{subject to} \quad & u_i = \langle \pi_i, v_i \rangle \qquad \forall i = 1, \dots, K \\
& \langle \pi_i, v_i \rangle \leq \langle \pi_j, v_i \rangle \quad \forall i, j = 1, \dots, K \\
& \langle \pi_i, \mathbb{1}_S \rangle = 1 \qquad \forall i = 1, \dots, K \\
& \pi_i \geq 0 \qquad \forall i = 1, \dots, K,
\end{aligned} \tag{C.12}$$

for $\pi_1 \dots, \pi_K \in \mathbb{R}^S$.

Proof. Suppose first that $u \in \mathcal{K}$ is the restriction to $\mathcal{V}$ of some MEU functional U. For $i = 1, \dots, K$, let $\pi_i \in \partial U(v_i)$ denote an arbitrarily selection of supergradients of U. As $U(0) = 0$, by homogeneity, $U(v_i) = \langle \pi_i, v_i \rangle$ for all $i = 1, \dots, K$. Furthermore, for all $x \in \mathbb{R}^S$ and all $v_i \in \mathcal{V}$:

$$\begin{aligned}
U(x) &\leq U(v_i) + \langle \pi_i, x - v_i \rangle \\
&= \langle \pi_i, v_i \rangle + \langle \pi_i, x - v_i \rangle \\
&= \langle \pi_i, x \rangle,
\end{aligned}$$

hence for all $v_j \in \mathcal{V}$, $\langle \pi_j, v_j \rangle \leq \langle \pi_i, v_j \rangle$. As U is increasing, for each i, $\pi_i \geq 0$. Let $\alpha \in \mathbb{R}$. Since U is translation-invariant, for all v_i:

$$U(v_i + \alpha \mathbb{1}_S) \leq U(v_i) + \langle \pi_i, \alpha \mathbb{1}_S \rangle$$

hence

$$U(v_i) + \alpha \leq U(v_i) + \langle \pi_i, \alpha \mathbb{1}_S \rangle$$

and

$$\alpha \leq \alpha \langle \pi_i, \mathbb{1}_S \rangle. \tag{C.13}$$

If $\alpha > 0$, $1 \leq \langle \pi, \mathbb{1}_S \rangle$, and if $\alpha < 0$, $1 \geq \langle \pi, \mathbb{1}_S \rangle$. Since (C.14) holds for all $\alpha \in \mathbb{R}$, we obtain $\langle \pi_i, \mathbb{1}_S \rangle = 1$.

Suppose now that for some collection $\pi_1, \ldots, \pi_K \in \Delta(S)$, we have a vector $u \in \mathcal{U}$ satisfying (i) $u_i = \langle \pi_i, v_i \rangle$ and (ii) $\langle \pi_i, v_i \rangle \leq \langle \pi_j, v_i \rangle$. Define

$$\hat{U}(x) = \min_{i \in \{1,\ldots,K\}} \langle \pi_i, x \rangle = \min_{\pi \in \text{co}\{\pi_1,\ldots,\pi_K\}} \langle \pi, x \rangle.$$

The latter equality follows from standard results on support functions see, e.g., [66] Theorem 3.3.2. By construction, $u_i = \hat{U}(v_i)$ and $\hat{U}$ is a risk-neutral MEU functional.

□

Variational Preferences

Proof. Suppose first that $u \in \mathcal{K}$ is the restriction to $\mathcal{V}$ of some risk-neutral variational utility functional U. For $i = 1, \ldots, K$, let $\pi_i \in \partial U(v_i)$ be an arbitrary selection of supergradients of U, one at each v_i. For all $i = 1, \ldots K$, let:

$$\gamma_i = u_i - \langle \pi_i, v_i \rangle.$$

Then, for all i, by construction $u_i = \gamma_i + \langle \pi_i, v_i \rangle$ and $\gamma_K = 0$. Moreover, for all $x \in \mathbb{R}^S$ and all $v_j \in \mathcal{V}$:

$$\begin{aligned} U(x) &\leq U(v_j) + \langle \pi_j, x - v_j \rangle \\ &= \gamma_j + \langle \pi_j, v_j \rangle + \langle \pi_j, x - v_j \rangle \\ &= \gamma_j + \langle \pi_j, x \rangle, \end{aligned}$$

hence in particular, for all $v_i \in \mathcal{V}$, $\gamma_i + \langle \pi_i, v_i \rangle \leq \gamma_j + \langle \pi_j, v_i \rangle$. As U is increasing, for each i, $\pi_i \geq 0$. Let $\alpha \in \mathbb{R}$. Since U is translation-invariant, for all v_i:

$$U(v_i + \alpha \mathbb{1}_S) \leq U(v_i) + \langle \pi_i, \alpha \mathbb{1}_S \rangle$$

hence

$$U(v_i) + \alpha \leq U(v_i) + \langle \pi_i, \alpha \mathbb{1}_S \rangle$$

and

$$\alpha \leq \alpha \langle \pi_i, \mathbb{1}_S \rangle. \tag{C.14}$$

If $\alpha > 0$, $1 \leq \langle \pi, \mathbb{1}_S \rangle$, and if $\alpha < 0$, $1 \geq \langle \pi, \mathbb{1}_S \rangle$. Since (C.14) holds for all $\alpha \in \mathbb{R}$, we obtain $\langle \pi_i, \mathbb{1}_S \rangle = 1$.

Suppose now that for some collection $\pi_1, \ldots, \pi_K \in \Delta(S)$ and $\gamma_1, \ldots, \gamma_K \in \mathbb{R}$ with $\gamma_K = 0$, we have a vector $u \in \mathcal{U}$ satisfying (i) $u_i = \gamma_i + \langle \pi_i, v_i \rangle$, and (ii) $\gamma_i + \langle \pi_i, v_i \rangle \leq \gamma_j + \langle \pi_j, v_i \rangle$. Define

$$\hat{U}(x) = \min_{i \in \{1,\ldots,K\}} \gamma_i + \langle \pi_i, x \rangle$$

By construction, $u_i = \hat{U}(v_i)$ and $\hat{U}$ is a (i) translation invariant, (ii) concave, (iii) increasing, (iv) normalized functional hence, by the results of [80], corresponds to a variational utility functional. □

Dual Self Expected Utility

A map $U : X \to \mathbb{R}$ is said to be a dual-self utility functional if it is of the form:

$$U(x) = \max_{P \in \mathbb{P}^*} \min_{\pi \in P} \langle \pi, x \rangle,$$

where $\mathbb{P}^*$ is a compact collection (in the Hausdorff topology) of compact, convex subsets of $\Delta(S)$.

Let $(\mathcal{V}, \mathcal{E})$ denote an experiment, where $v_K = 0$. Let $\mathcal{K}_{\text{DS}}$ denote the collection of $u \in \mathcal{U}$ that are restrictions of dual-self utility functionals. Then solving (3.7) with $\mathcal{K} = \mathcal{K}_{\text{DS}}$ is equivalent to solving:

$$\begin{aligned}
\min_{u \in \mathcal{U}} \quad & \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2 \\
\text{subject to} \quad & u_i = \langle \pi_{ii}, v_i \rangle && \forall i = 1, \ldots, K \\
& \langle \pi_{ii}, v_i \rangle \leq \langle \pi_{ij}, v_i \rangle && \forall i, j = 1, \ldots, K \\
& \langle \pi_{ji}, v_i \rangle \leq \langle \pi_{ii}, v_i \rangle && \forall i, j = 1, \ldots, K \\
& \langle \pi_{ij}, \mathbb{1}_S \rangle = 1 && \forall i, j = 1, \ldots, K \\
& \pi_{ij} \geq 0 && \forall i, j = 1, \ldots, K,
\end{aligned} \tag{C.15}$$

for $u \in \mathbb{R}^K$, $\{\pi_{ij}\}_{i,j=1}^K \in \mathbb{R}^S$.

Proof. Suppose, first, that $u, \{\pi_{ij}\}_{i,j=1}^K$ is a solution to (C.15). Define, for each $i = 1, \ldots, K$, the set $P_i = \text{co}\{\pi_{i,1}, \ldots, \pi_{i,K}\}$. Clearly $P_i \subseteq \Delta(S)$ for each i. Let $\mathbb{P}^* = \{P_i\}_{i=1}^K$. We claim that:

$$U(x) = \max_{P \in \mathbb{P}^*} \min_{\pi \in P} \langle \pi, x \rangle$$

defines a DSEU functional whose restriction to $\mathcal{V}$ is precisely u. Firstly, as $\langle \pi_{ii}, v_i \rangle \leq \langle \pi_{ij}, v_i \rangle$ for all $j = 1, \ldots, K$, it follows that:

$$u_i = \langle \pi_{ii}, v_i \rangle = \min_{\pi \in P_i} \langle \pi, v_i \rangle.$$

But, for all $j = 1, \ldots, K$ we have $\langle \pi_{ji}, v_i \rangle \leq u_i$, hence:

$$u_i \geq \langle \pi_{ji}, v_i \rangle \geq \min_{\pi \in P_j} \langle \pi, v_i \rangle,$$

as $\pi_{ji} \in P_j$. Thus:

$$\begin{aligned}
U(v_i) &= \max_{P \in \mathbb{P}^*} \min_{\pi \in P} \langle \pi, v_i \rangle \\
&= \min_{\pi \in P_i} \langle \pi, v_i \rangle \\
&= \langle \pi_{ii}, v_i \rangle \\
&= u_i.
\end{aligned}$$

Suppose now that $U(x) = \max_{P \in \mathbb{P}^*} \min_{\pi \in P} \langle \pi, x \rangle$ is a DSEU functional on $\mathbb{R}^S$. For $i = 1, \ldots, K$, let $P_i \in \mathbb{P}^*$ denote any belief set for which:

$$U(v_i) = \min_{\pi \in P_i} \langle \pi, v_i \rangle,$$

and let $\pi_{ii} \in P_i$ be any minimizer of the right-hand side.[13] Define, for each $i = 1, \ldots, K$, the utility value $u_i = \langle \pi_{ii}, v_i \rangle$. Since P_j is an 'active' belief set at v_j for each $j \neq i$, there exists, for each j, some $\pi_{ij} \in P_i$ such that $\langle \pi_{ij}, v_j \rangle \leq u_j$. Since each $\pi_{ij} \in P_i$, then $u_i \leq \langle \pi_{ij}, v_i \rangle$ for each i. Then, as clearly every $\pi_{ij} \in \Delta(S)$, the collection $u, \{\pi_{ij}\}_{i,j=1}^K$ is a solution to (C.15), as required.

□

Dual-Self Variational Utility

A map $U : X \to \mathbb{R}$ is said to be a dual-self variational utility functional if it is of the form:

$$U(x) = \max_{c \in \mathbb{C}} \min_{\pi \in \Delta(S)} \langle \pi, x \rangle + c(\pi),$$

where $\mathbb{C}$ is a collection of convex cost functions $c : \Delta(S) \to [0, \infty]$ such that $\max_{c \in \mathbb{C}} \min_{\pi \in \Delta(S)} c(\pi) = 0$. Such functionals are characterized by being (i) additive-equivariant, (ii) monotone, (iii) normalized, i.e. $U(\mathbb{1}_S) = 1$, see Supplementary Appendix to [33].

Let $(\mathcal{V}, \mathcal{E})$ denote an experiment, where $v_K = 0$. Let $\mathcal{K}_{\text{DSV}}$ denote the collection of $u \in \mathcal{U}$ that are restrictions of dual-self variational utility functionals. Then solving

[13] Such a belief set exists as $\mathbb{P}^*$ is compact (in the Hausdorff topology on the space of compact subsets of $\Delta(S)$), and $\min_{\pi \in P} \langle \pi, x \rangle$ is continuous in P for each x.

(3.7) with $\mathcal{K} = \mathcal{K}_{\text{DSV}}$ is equivalent to solving:

$$\begin{aligned} \min_{u \in \mathcal{U}} \quad & \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2 \\ \text{subject to} \quad & u_i \geq u_j \qquad\qquad \forall i, j \text{ s.t. } v_i \geq v_j \\ & u_K = 0, \end{aligned} \tag{C.16}$$

where $v_i \geq v_j$ is understood in the product order on $\mathbb{R}^S$.

Proof. Firstly, suppose U is a dual-self variational functional. Then it clearly is monotone, hence $v_i \geq v_j$ implies $U(v_i) \geq U(v_j)$. Moreover,

$$U(\mathbb{1}_S) = U\big(\phi(1, 0)\big) = U(0) + 1,$$

hence U is normalized if and only if $U(0) = 0$. Thus clearly letting $u_i = U(v_i)$ satisfies the constraints of (C.16).

Conversely, suppose u is a solution to (C.16). In light of the characterization provided in [33], it suffices to prove there exists an additive-equivariant and monotone extension from $\mathcal{V}$ to $\mathbb{R}^S$.[14] However, note that since $\mathcal{V}$ contains no pairs of $\sim_{\trianglelefteq}$-related elements, u is trivially additive-equivariant and by definition monotone on $\mathcal{V}$. Hence by Theorem 1 of [29],

$$U(x) = \sup\{u_{v_i} + b : v_i \in \mathcal{V}, b \in \mathbb{R}, \text{ and } v_i + b\mathbb{1}_S \leq x\}$$

defines an additive-equivariant, monotone, and normalized extension of u, and hence by [33] this corresponds to some dual-self variational utility functional. $\square$

[14]Normalization holds for any additive-equivariant extension, as $u_K = 0$.

Bibliography

[1] Mohammed Abdellaoui, Aurélien Baillon, Laetitia Placido, and Peter P Wakker. The rich domain of uncertainty: Source functions and their experimental implementation. *American Economic Review*, 101(2):695–723, 2011.

[2] János Aczél and J. Dhombres. *Functional equations in several variables.* Number 31. Cambridge university press, 1989.

[3] Sydney N Afriat. On a system of inequalities in demand analysis: an extension of the classical method. *International economic review*, pages 460–472, 1973.

[4] Victor Aguiar, Per Hjertstrand, and Roberto Serrano. A rationalization of the weak axiom of revealed preference. 2020.

[5] David Ahn, Syngjoo Choi, Douglas Gale, and Shachar Kariv. Estimating ambiguity aversion in a portfolio choice experiment. *Quantitative Economics*, 5(2): 195–223, 2014.

[6] Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: a Hitchhiker's Guide.* Springer, Berlin; London, 2006.

[7] Gad Allon, Michael Beenstock, Steven Hackman, Ury Passy, and Alexander Shapiro. Nonparametric estimation of concave production technologies by entropic methods. *Journal of Applied Econometrics*, 22(4):795–816, 2007.

[8] James Andreoni and John Miller. Giving according to garp: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70(2):737–753, 2002.

[9] Francis J Anscombe and Robert J Aumann. A definition of subjective probability. *Annals of mathematical statistics*, 34(1):199–205, 1963.

[10] Aaron Archer and Robert Kleinberg. Truthful germs are contagious: a local-to-global characterization of truthfulness. *Games and Economic Behavior*, 86: 340–366, 2014.

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

[12] Itai Ashlagi, Mark Braverman, Avinatan Hassidim, and Dov Monderer. Monotonicity and implementability. *Econometrica*, 78(5):1749–1772, 2010.

[13] Aurelien Baillon, Yoram Halevy, Chen Li, et al. Experimental elicitation of ambiguity attitude using the random incentive system. *University of British Columbia working paper*, 2021.

[14] Ralph C Bayer, Subir Bose, Matthew Polisson, and Ludovic Renou. Ambiguity revealed. *Available at SSRN 2046598*, 2013.

[15] Timothy KM Beatty and Ian A Crawford. How demanding is the revealed preference approach to demand? *American Economic Review*, 101(6):2782–95, 2011.

[16] Gordon M Becker, Morris H DeGroot, and Jacob Marschak. Measuring utility by a single-response sequential method. *Behavioral science*, 9(3):226–232, 1964.

[17] Melvyn S Berger and Norman G Meyers. On a system of nonlinear partial differential equations arising in mathematical economics. *Bulletin of the American Mathematical Society*, 72(6):954–958, 1966.

[18] Melvyn S Berger and NG Myers. Utility functions for finite dimensional commodity spaces", in: Js chipman, l. hurwicz, mk richter, and hf sonnenschein (eds.), preferences, utility, and demand, 1971.

[19] B Douglas Bernheim and Antonio Rangel. Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics. *The Quarterly Journal of Economics*, 124(1):51–104, 2009.

[20] Antoine Billot, Alain Chateauneuf, Itzhak Gilboa, and Jean-Marc Tallon. Sharing beliefs: between agreeing and disagreeing. *Econometrica*, 68(3):685–694, 2000.

[21] Richard Blundell, Martin Browning, and Ian Crawford. Best nonparametric bounds on demand responses. *Econometrica*, 76(6):1227–1262, 2008.

[22] Richard Blundell, Martin Browning, Laurens Cherchye, Ian Crawford, Bram De Rock, and Frederic Vermeulen. Sharp for sarp: nonparametric bounds on counterfactual demands. *American Economic Journal: Microeconomics*, 7(1): 43–60, 2015.

[23] Richard W Blundell, Martin Browning, and Ian A Crawford. Nonparametric engel curves and revealed preference. *Econometrica*, 71(1):205–240, 2003.

[24] Peter Bossaerts, Paolo Ghirardato, Serena Guarnaschelli, and William R Zame. Ambiguity in asset markets: Theory and experiment. *The Review of Financial Studies*, 23(4):1325–1359, 2010.

[25] Adrian Bruhin, Helga Fehr-Duda, and Thomas Epper. Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica*, 78(4):1375–1412, 2010.

[26] Ozan Candogan, Ishai Menache, Asuman Ozdaglar, and Pablo A Parrilo. Flows and decompositions of games: Harmonic and potential games. *Mathematics of Operations Research*, 36(3):474–503, 2011.

[27] Peter Caradonna. How strong is the weak axiom. 2020.

[28] Matias D Cattaneo, Xinwei Ma, Yusufcan Masatlioglu, and Elchin Suleymanov. A random attention model. *Journal of Political Economy*, 128(7):2796–2836, 2020.

[29] Simone Cerreia-Vioglio, Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini. Niveloids and their extensions: Risk measures on small domains. *Journal of Mathematical Analysis and Applications*, 413(1):343–360, 2014.

[30] Christopher P Chambers, Federico Echenique, and Eran Shmaya. The axiomatic structure of empirical content. *American Economic Review*, 104(8):2303–19, 2014.

[31] Christopher P Chambers, Federico Echenique, and Kota Saito. Testing theories of financial decision making. *Proceedings of the National Academy of Sciences*, 113(15):4003–4008, 2016.

[32] Christopher P Chambers, Federico Echenique, and Nicolas S Lambert. Preference identification. 2020.

[33] Madhav Chandrasekher, Mira Frick, Ryota Iijima, and Yves Le Yaouanq. Dual-self representations of ambiguity preferences. 2020.

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

[35] Syngjoo Choi, Raymond Fisman, Douglas Gale, and Shachar Kariv. Consistency and heterogeneity of individual behavior under uncertainty. *American economic review*, 97(5):1921–1938, 2007.

[36] László Csató. A graph interpretation of the least squares ranking method. *Social Choice and Welfare*, 44(1):51–69, 2015.

[37] Geoffroy De Clippel and Kareen Rozen. Bounded rationality and limited datasets. 2014.

[38] Rahul Deb, Yuichi Kitamura, John K-H Quah, and Jörg Stoye. Revealed price preference: theory and empirical analysis. *arXiv preprint arXiv:1801.02702*, 2018.

[39] Gerard Debreu. Continuity properties of paretian utility. *International Economic Review*, 5(3):285–293, 1964.

[40] Thomas Demuynck, Christian Seel, Giang Tran, et al. An index of competitiveness and cooperativeness for normal-form games. *American Economic Journal. Microeconomics*, 2020.

[41] Lutz Dümbgen. On nondifferentiable functions and the bootstrap. *Probability Theory and Related Fields*, 95(1):125–140, 1993.

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

[43] Ozgur Evren, Hiroki Nishimura, and Efe A Ok. Top" cycles and revealed preference structures. 2019.

[44] Zheng Fang and Andres Santos. Inference on directionally differentiable functions. *The Review of Economic Studies*, 86(1):377–412, 2019.

[45] Zheng Fang and Juwon Seo. A projection framework for testing shape restrictions that form convex cones. *arXiv preprint arXiv:1910.07689*, 2019.

[46] Peter C Fishburn and Ariel Rubinstein. Time preference. *International Economic Review*, pages 677–694, 1982.

[47] Francoise Forges and Enrico Minelli. Afriat's theorem for general budget sets. *Journal of Economic Theory*, 144(1):135–145, 2009.

[48] Laszlo Fuchs. *Partially ordered algebraic systems*, volume 28. Courier Corporation, 2011.

[49] Drew Fudenberg, Whitney Newey, Philipp Strack, and Tomasz Strzalecki. Testing the drift-diffusion model. *Proceedings of the National Academy of Sciences*, 117(52):33141–33148, 2020.

[50] David Gale. A note on revealed preference. *Economica*, 27(108):348–354, 1960.

[51] Paolo Ghirardato, Fabio Maccheroni, and Massimo Marinacci. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118(2):133–173, 2004.

[52] Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of mathematical economics*, 18(2):141–153, 1989.

[53] Chris Godsil and Gordon F Royle. *Algebraic graph theory*, volume 207. Springer Science & Business Media, 2001.

[54] Paul Gordan. Ueber die auflösung linearer gleichungen mit reellen coefficienten. *Mathematische Annalen*, 6(1):23–28, 1873.

[55] Leo J Grady and Jonathan R Polimeni. *Discrete calculus: Applied analysis on graphs for computational science*. Springer Science & Business Media, 2010.

[56] Simon Grant and Ben Polak. Mean-dispersion preferences and constant absolute uncertainty aversion. *Journal of Economic Theory*, 148(4):1361–1398, 2013.

[57] Faruk Gul. A theory of disappointment aversion. *Econometrica: Journal of the Econometric Society*, pages 667–686, 1991.

[58] Yoram Halevy. Ellsberg revisited: An experimental study. *Econometrica*, 75(2): 503–536, 2007.

[59] Yoram Halevy, Dotan Persitz, and Lanny Zrill. Parametric recoverability of preferences. *Journal of Political Economy*, 126(4):1558–1593, 2018.

[60] Philip Hartman. Frobenius theorem under carathéodory type conditions. *Journal of Differential Equations*, 7(2):307–333, 1970.

[61] David Harville. The use of linear-model methodology to rate high school or college football teams. *Journal of the American Statistical Association*, 72(358): 278–289, 1977.

[62] John D Hey and Jinkwon Lee. Do subjects separate (or are they sophisticated)? *Experimental Economics*, 8(3):233–265, 2005.

[63] John D Hey and Noemi Pace. The explanatory and predictive power of non two-stage-probability theories of decision making under ambiguity. *Journal of Risk and Uncertainty*, 49(1):1–29, 2014.

[64] John D Hey, Gianna Lotito, and Anna Maffioletti. The descriptive and predictive adequacy of theories of decision making under uncertainty/ambiguity. *Journal of risk and uncertainty*, 41(2):81–111, 2010.

[65] Anil N Hirani, Kaushik Kalyanaraman, and Seth Watts. Least squares ranking on graphs. *arXiv preprint arXiv:1011.1716*, 2010.

[66] Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Fundamentals of convex analysis*. Springer Science & Business Media, 2004.

[67] Richard B Holmes. *Geometric functional analysis and its applications*, volume 24. Springer Science & Business Media, 2012.

[68] Charles A Holt. Preference reversals and the independence axiom. *The American Economic Review*, 76(3):508–515, 1986.

[69] Han Hong and Jessie Li. The numerical bootstrap. *The Annals of Statistics*, 48(1):397–412, 2020.

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

[71] Xiaoye Jiang, Lek-Heng Lim, Yuan Yao, and Yinyu Ye. Statistical ranking and combinatorial hodge theory. *Mathematical Programming*, 127(1):203–244, 2011.

[72] Yuichi Kitamura and Jörg Stoye. Nonparametric analysis of random utility models. *Econometrica*, 86(6):1883–1909, 2018.

[73] Asen Kochov. The epistemic value of a menu and subjective states. *Available at SSRN 2715001*, 2010.

[74] Tjalling C Koopmans. Stationary ordinal utility and impatience. *Econometrica*, pages 287–309, 1960.

[75] Timo Kuosmanen. Representation theorem for convex nonparametric least squares. *The Econometrics Journal*, 11(2):308–325, 2008.

[76] Alexey I Kushnir and Lev Lokutsievskiy. On the equivalence of weak-and cyclic-monotonicity. *Available at SSRN 3422846*, 2019.

[77] David Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478, 1997.

[78] Jinkwon Lee. The effect of the background risk in a simple chance improving decision model. *Journal of risk and uncertainty*, 36(1):19–41, 2008.

[79] Graham Loomes. Evidence of a new violation of the independence axiom. *Journal of Risk and uncertainty*, 4(1):91–108, 1991.

[80] Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74(6):1447–1498, 2006.

[81] Paola Manzini and Marco Mariotti. Sequentially rationalizable choice. *American Economic Review*, 97(5):1824–1839, 2007.

[82] Marco Mariotti. What kind of preference maximization does the weak axiom of revealed preference characterize? *Economic Theory*, 35(2):403–406, 2008.

[83] Yusufcan Masatlioglu, Daisuke Nakajima, and Erkut Y Ozbay. Revealed attention. *American Economic Review*, 102(5):2183–2205, 2012.

[84] Rosa L Matzkin and Marcel K Richter. Testing strictly concave rationality. *Journal of Economic Theory*, 53(2):287–303, 1991.

[85] James R Munkres. *Topology; a First Course.* Prentice-Hall, 1974.

[86] James R Munkres. *Elements of algebraic topology.* Addison Wesley, Redwood City, California, 1984.

[87] Hiroki Nishimura, Efe A Ok, and John K-H Quah. A comprehensive approach to revealed preference theory. *American Economic Review*, 107(4):1239–63, 2017.

[88] Efe A Ok. *Real analysis with economic applications*. Princeton University Press, 2011.

[89] Efe A Ok and Yusufcan Masatlioglu. A theory of (relative) discounting. *Journal of Economic Theory*, 137(1):214–245, 2007.

[90] Efe A Ok and Gerelt Tserenjigmid. Deterministic rationality of stochastic choice behavior. 2019.

[91] Braxton Osting, Jérôme Darbon, and Stanley Osher. Statistical ranking using the l1-norm on graphs. *Inverse Problems & Imaging*, 7(3):907, 2013.

[92] Matthew Polisson, John K-H Quah, and Ludovic Renou. Revealed preferences over risk and uncertainty. *American Economic Review*, 110(6):1782–1820, 2020.

[93] John KH Quah. Weak axiomatic demand theory. *Economic Theory*, 29(3): 677–699, 2006.

[94] Marcel K Richter. Revealed preference theory. *Econometrica: Journal of the Econometric Society*, pages 635–645, 1966.

[95] Luca Rigotti, Chris Shannon, and Tomasz Strzalecki. Subjective beliefs and ex ante trade. *Econometrica*, 76(5):1167–1190, 2008.

[96] Donald J Rose, R Endre Tarjan, and George S Lueker. Algorithmic aspects of vertex elimination on graphs. *SIAM Journal on computing*, 5(2):266–283, 1976.

[97] Hugh Rose. Consistency of preference: the two-commodity case. *The Review of Economic Studies*, 25(2):124–125, 1958.

[98] Bryan R Routledge and Stanley E Zin. Generalized disappointment aversion and asset prices. *The Journal of Finance*, 65(4):1303–1332, 2010.

[99] Michael Saks and Lan Yu. Weak monotonicity suffices for truthfulness on convex domains. In *Proceedings of the 6th ACM conference on Electronic commerce*, pages 286–293. ACM, 2005.

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

[101] David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica: Journal of the Econometric Society*, pages 571–587, 1989.

[102] Emilio Seijo and Bodhisattva Sen. Nonparametric least squares estimation of a multivariate convex regression function. *The Annals of Statistics*, 39(3):1633–1657, 2011.

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

[104] Wayne J Shafer. The nontransitive consumer. *Econometrica (pre-1986)*, 42(5): 913, 1974.

[105] Alexander Shapiro. Asymptotic analysis of stochastic programs. *Annals of Operations Research*, 30(1):169–186, 1991.

[106] Bart Smeulders, Laurens Cherchye, and Bram De Rock. Nonparametric analysis of random utility models: computational tools for statistical testing. *Econometrica*, 89(1):437–455, 2021.

[107] Edwin H Spanier. *Algebraic topology*, volume 55. Springer Science & Business Media, 1989.

[108] Chris Starmer and Robert Sugden. Does the random-lottery incentive system elicit true preferences? an experimental investigation. *The American Economic Review*, 81(4):971–978, 1991.

[109] Raymond T Stefani. Football and basketball predictions using least squares. *IEEE Transactions on systems, man, and cybernetics*, 7(2):117–21, 1977.

[110] Edward Szpilrajn. Sur l'extension de l'ordre partiel. *Fundamenta Mathematicae*, 16(1):386–389, 1930. URL `http://eudml.org/doc/212499`.

[111] Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323, 1992.

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

[113] H Peyton Young. An axiomatization of borda's rule. *Journal of economic theory*, 9(1):43–52, 1974.

[114] H Peyton Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.