Preference Regression[*]

Peter P. Caradonna[†]

October 22, 2021

Abstract

This paper investigates the problem of testing and calibrating models of individual decision making. We consider a consumption space equipped with an endogenous notion of abstract 'numeraire,' and characterize those preferences for which the quantity of numeraire needed to compensate an agent between a pair of alternatives provides a consistent, cardinal measure of the intensity of preference. This framework includes many well-known preferences over classical commodity spaces, finite or infinite horizon consumption streams, and a wide range of models of preference over uncertainty and risk as special cases. For data consisting of observed or experimentally elicited compensation differences, we develop a least squares theory for quantifying a model's predictive accuracy and estimating underlying parameters. We additionally provide a general class of explicit, non-parametric statistical tests of rationalizability by particular models for stochastic data. Applications to model selection, welfare analysis and elicitation of subjective beliefs are given.

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

[*]I am indebted to my advisor Christopher Chambers for his encouragement and innumerable helpful discussions over the course of this project. I would also like to thank Axel Anderson, Thomas Demuynck, Ivana Komunjer, Yusufcan Masatlioglu, Alexandre Poirier, Margit Reischer, Mauricio Ribeiro, Christopher Turansick, as well as seminar audiences at Georgetown, SAET and D-TEA for numerous helpful comments.

[†]Department of Economics, Georgetown University. Email: ppc14@georgetown.edu. Comments are welcome.

[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, Bruhin et al. (2010).

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 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

[2] It is generally undesirable to test models of preference using loss functions defined over the consumption space. Varian (1990) argues such tests quantify the statistical, rather than economic significance of violations, which may be unrelated; see also Halevy et al. (2018) for discussion and extensions. Varian (1990) 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 Section 5.2, and in particular (2).

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 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 on the true set of priors (see Example 11). 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.

[3] Theorem 1 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 Section 3.3.

[4] The elicitability of compensation data of this form is established in Theorem 2.

[5] For example, while the popular efficiency indices of Afriat (1973) and Varian (1990) 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. Halevy et al. (2018) provides an extension to allow for calibrating various parametric models; see also Polisson et al. (2020).

[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.

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 Gilboa and Schmeidler (1989).[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 Varian (1990), for a large collection of models.

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

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

[7]See Example 8 for details.

[8]For example, Bruhin et al. (2010) 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

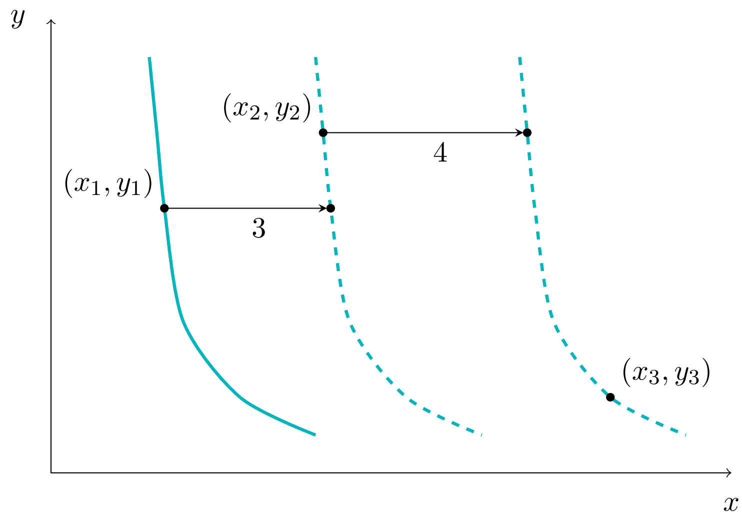

Figure 1: When preferences are representable by a quasilinear utility of the form $u(x, y) = v(y) + x$, indifference curves are horizontal translates of one another. Geometrically, translating any point on the (x_1, y_1) indifference curve 3 units to the right makes it fall precisely on the (x_2, y_2) indifference curve. If one translates the result again by 4, the point must then lie on the (x_3, y_3) curve.

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 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 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

priori, that the individual's choice behavior is consistent with the maximization of any complete and transitive binary relation.

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]

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.

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 Billot et al. (2000), Rigotti et al. (2008), Bossaerts et al. (2010), Bayer et al. (2013), Ahn et al. (2014), or Chambers et al. (2016).[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, including subjective expected utility (Anscombe and Aumann 1963), Choquet expected utility (Schmeidler 1989), maxmin expected utility (Gilboa and Schmeidler 1989), variational (Maccheroni et al. 2006), as well as dual-self and dual-self variational

[10] This follows from Theorem 3.

[11] For an explicit description of this set for a general experiment, see Example 9.

[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.

preferences (Chandrasekher et al. 2020).[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. 2010, Abdellaoui et al. 2011) and eliciting reservation prices (e.g., Becker et al. 1964, Halevy 2007).[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 Kitamura and Stoye (2018), Deb et al. (2018), Fudenberg et al. (2020), Cattaneo et al. (2020), and Smeulders et al. (2021). 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 Fang and Santos (2019) due to Hong and Li (2020). 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 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. Harville (1977), Stefani (1977). 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

[13] As noted in Chandrasekher et al. (2020), dual-self expected utility is a particular choice of representation of the class of invariant biseparable preferences of Ghirardato et al. (2004). 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 Loomes (1991), Andreoni and Miller (2002), Choi et al. (2007), Hey and Pace (2014), and Ahn et al. (2014).

[15] Our method of constructing test statistics also draws heavy from the statistical and econometric literature on shape-constrained regression, in particular Allon et al. (2007), Kuosmanen (2008), and Seijo and Sen (2011).

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 Hirani et al. (2010), Jiang et al. (2011), Osting et al. (2013) and references therein. These ideas have already found application elsewhere in economics, including game theory, social choice, and revealed preference (resp. Candogan et al. 2011, Csató 2015, Caradonna 2020).

3 Invariant Preferences

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 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 $\alpha \in \mathbb{R}_+$, $x, y \in X$:

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

[16] A monoid is a semigroup with identity; see Fuchs (2011). Formally, $\mathbb{R}_+$, equipped with the the usual notion of addition $+$, is a monoid.

(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' 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.

3.2 Examples

Example 2. (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 3. (Stationary Preferences for Dated Rewards): Suppose $X = \mathbb{R}_+ \times Z$ with $\phi\big(\alpha, (t, z)\big) = (\alpha + t, z)$ where, following Fishburn and Rubinstein (1982), 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

[17] Ok and Masatlioglu (2007) provide a complementary interpretation of preferences over X as the commitment preferences of an agent.

instead require (N.2) to hold with the opposite relation, to reflect 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 4. (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):

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

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.

[18] In this setting, (N.1) corresponds to Fishburn and Rubinstein (1982)'s (A.2), a monotonicity axiom, and (A.5), their stationarity axiom (cf. the stationarity axiom of Ok and Masatlioglu 2007). 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 $\alpha > 0$, $U(\alpha x) = \alpha U(x)$.

[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, Cerreia-Vioglio et al. (2014), footnote 5.

Example 5. (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.

Example 6. (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] Grant and Polak (2013) 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.

[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 $\phi_{\bar{x}}$.

[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).

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]

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

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 1. *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 1 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 1 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

[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.

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 Example 3, see Fishburn and Rubinstein (1982).

3.4 Regularity Assumptions

We will henceforth impose the following three technical regularity conditions on model primitives, which, while not required for Theorem 1, are required in the sequel.

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

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 Lemma 1).

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

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

[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 Fuchs (2011). 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).

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

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 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.

4 Data & Elicitiation

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]

[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 Aczél and Dhombres (1989) 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 Section 5.1.

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

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

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 $\{x, y\} \in \mathcal{E}$. 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 6 we will consider how to test hypotheses about the expected behavior of such a population.

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 Becker et al. (1964). 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 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

when $y \succsim x$.

[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 Holt (1986) 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. Starmer and Sugden (1991), Hey and Lee (2005), and Lee (2008).

[33]For a discussion of incentive compatibility of this mechanism in the case of experiments over ambiguous prospects, see Baillon et al. (2021).

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 2. *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 2 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^*.

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.

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), $\ldots$, $(x_{L-1}, x_0) \in \vec{\mathcal{E}}$,

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

where subscripts are understood mod-L. By minor abuse of notation, we also refer to individual flows as being cardinally consistent if (1) holds. Cardinal consistency 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 3. *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:*

[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 $(x, y) \in \vec{\mathcal{E}}$, $F_{xy} = -F_{yx}$.

(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 3 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 $\mathcal{E}$.

5.2 Preference 'Regression'

5.2.1 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 $\operatorname{grad} : \mathcal{U} \to \mathcal{F}$. The following two propositions are routine; for completeness, we provide proofs in Appendix D.

Proposition 1. *A flow $F \in \mathcal{F}$ is cardinally consistent if and only if it belongs to the image of the gradient.*

In light of Proposition 1, 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 Proposition 1 until Section 5.2.2.

[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 div : $\mathcal{F} \to \mathcal{U}$. When $\mathcal{U}$ carries its standard inner product, $-$div is the adjoint of the gradient.

Proposition 2. *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 Proposition 1, 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 is precisely the span of the perfect cycles in $\mathcal{F}$.[40] By Proposition 2, 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:

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

[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, Proposition 2 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.

[40] See Godsil and Royle (2001) 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 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 (2):

$$R^2 = 1 - \frac{\|\bar{Y} - \hat{Y}\|^2}{\|\bar{Y}\|^2} = \frac{\|\hat{Y}\|^2}{\|\bar{Y}\|^2}. \tag{3}$$

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]

5.2.2 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 in $\mathcal{V}$. By Proposition 1, 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 Proposition 1 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:

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

[42] Similar ideas are leveraged in Demuynck et al. (2020) to obtain a measure of the incentive alignment for players in normal form games.

[43] See Young (1975) for a characterization of such aggregation rules for the ordinal case.

[44] The observations of this subsection are implicit in Young (1974), and noted explicitly in, e.g., Jiang et al. (2011) in the context of aggregating online rankings.

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

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

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 the normal equations for (2) may be compactly written as:[45]

$$L\, u^* = -\operatorname{div} \bar{Y}. \tag{4}$$

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 (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 (4) says that any solution u^* is determined by a strong 'averaging' property. In particular, for any v_i:

$$u_i^* = \frac{1}{\deg(i)} \Bigg[\sum_{j \in N(i)} u_j^* + \bar{Y}_{ji} \Bigg], \tag{5}$$

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 (5) is the cardinal Borda score, as (5) becomes:

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

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

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

[45] This makes use of the identity $-\operatorname{div} \circ \operatorname{grad} = L$, see Godsil and Royle (2001) Lemma 8.3.2, and the fact that $-\operatorname{div}$ is the adjoint of grad (Proposition 2).

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

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

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

Even when $(\mathcal{V}, \mathcal{E})$ is less-than-complete, solutions to (2) remain wholly determined by the averaging property (5) which, for complete experiments, characterizes the cardinal Borda score up to a constant. Thus our choice of the mean squared error 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.

5.2.3 A Money Pump Metric for the Residual

In Section 5.2.1, 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. (2011) for our setting.[47]

By Proposition 2, $R = \bar{Y} - \hat{Y}$ 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 $\|R\|_1 = \bar{c}\,L$ 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

[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. (2011) for details.

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

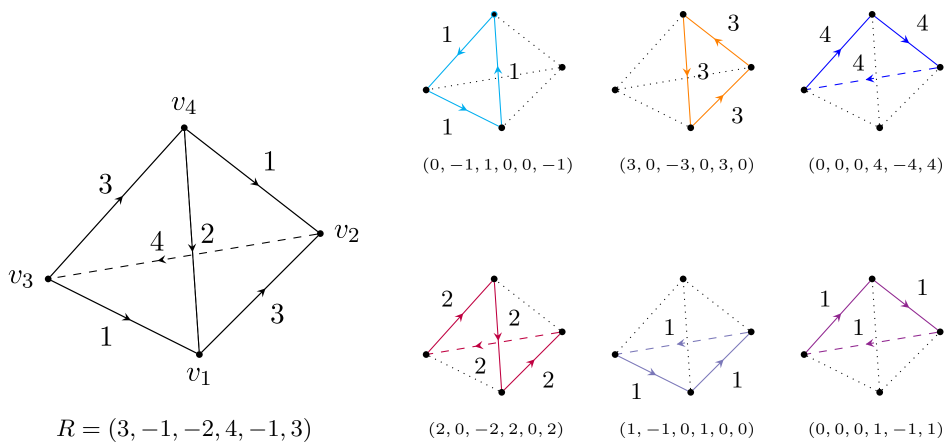

Figure 2: A residual flow $R = (R_{12}, R_{13}, R_{14}, R_{23}, R_{24}, R_{34})$, belonging to the kernel of the divergence, along with two decompositions into sums of perfect cycles. The lower bound of $\|R\|_1 = 14$ is attained by the sum of the money pump values of the bottom (though not the top) decomposition.

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 2. Instead, we consider the most conservative extension.

Let $\mathfrak{C} \subsetneq \mathcal{F}$ denote the set of pure cycles. For any $R \in \ker(\text{div})$ 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 $MP^* : \ker(\text{div}) \to \mathbb{R}$ via:

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

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 $\ker(\text{div})$.

Proposition 3. *For all $R \in \ker(\text{div})$, 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.

[49] That is, those collections $\{C_1, \ldots, C_M\} \subseteq \mathfrak{C}$ such that $\sum_m C_m = R$.

Given its economic interpretation for elements of ker(div), it is tempting to consider the L^1 analogue of (2),

$$\min_{u \in \mathcal{U}} \left\| (\operatorname{grad} u) - \bar{Y} \right\|_1, \tag{2'}$$

and to interpret the value of this linear program as the magnitude of the inconsistency. However, (2′) lacks a great deal of the analytic structure of (2). Unlike (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 (2′) form an $(n-1)$-dimensional polytope, see Osting et al. (2013).[50] Moreover, while the residual from (2) always belongs to ker(div), this need not be true for (2′). Thus the interpretation provided by Proposition 3 does not generally apply to residuals from (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.

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 (2).

Example 7 (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 (2) computes the MSE for CES preferences:

$$\begin{aligned} \min_{u,\rho} \quad & \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2 \\ \text{subject to} \quad & u_i = \frac{1}{\rho} \ln \left(\sum_{l=1}^{L} v_{li}^\rho \right) \\ & \rho \leq 1, \end{aligned} \tag{6}$$

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

[50] In contrast, for any n, the minimizers for (2) are precisely the constant functions.

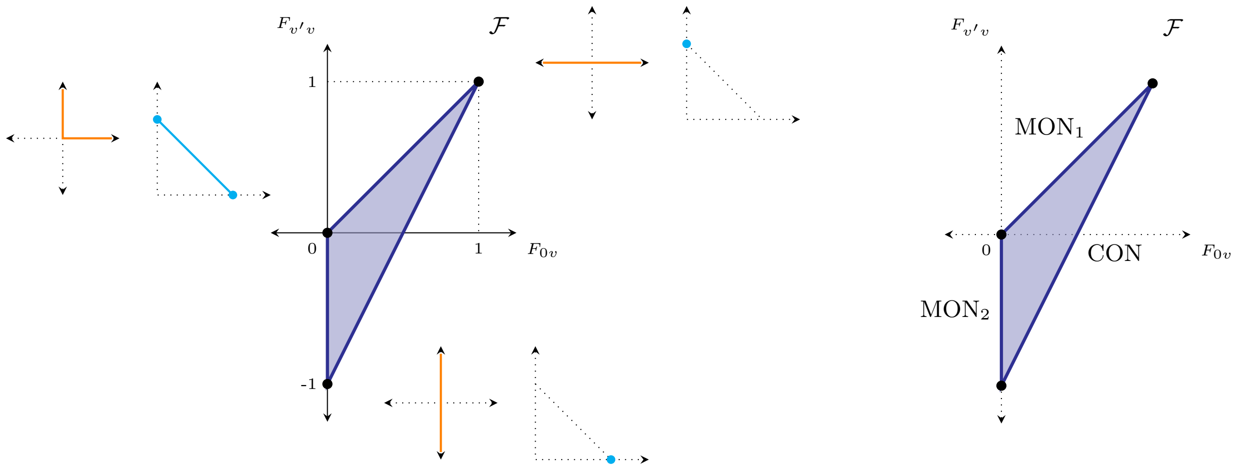

(a) The extremal MEU preferences.

(b) The binding constraints at each face.

Figure 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.

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.

Example 8 (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 $\mathcal{E} = \{\{0, v\}, \{v, v'\}\}$. 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 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.

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 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 advantage 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.

[51] See, e.g., Ok (2011) 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 Example 11 for more details.

[53] More generally, this corresponds to the class of invariant biseparable preferences of Ghirardato et al. (2004); see also Chandrasekher et al. (2020) for an elegant representation theorem for such preferences. However, for this simple experiment, the testable implications of this class of preferences coincide with those of the Choquet model of Schmeidler (1989).

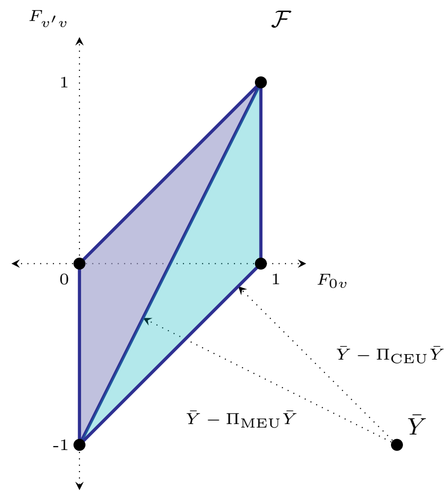

Figure 4: The set of MEU-rationalizable (violet) and CEU-rationalizable (violet or aquamarine) vectors for $\mathcal{E}$. Letting Π_{MEU} and Π_{CEU} denote the respective projections onto these sets, the quantity $\|\bar{Y} - \Pi_{\text{MEU}}\bar{Y}\|_2^2 - \|\bar{Y} - \Pi_{\text{CEU}}\bar{Y}\|_2^2$ reflects the *shadow price*, in mean squared error terms, of imposing ambiguity aversion, conditional upon requiring monotonicity, translation invariance, and homotheticity.

5.3.1 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 (2), one projects instead onto $\operatorname{grad}(\mathcal{K})$:

$$\min_{u \in \mathcal{K}} \|(\operatorname{grad} u) - \bar{Y}\|_2^2, \tag{7}$$

which is closed and convex by our definition of shape constraints. In Appendix E, we provide a number of explicit derivations of such constraint sets for a variety of ambiguity preferences, as well

[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., Holmes (2012), Lemma 17.H. It is necessarily satisfied if, for example, $\mathcal{K}$ is polyhedral, which is often the case.

as formal proofs for some of the examples in this section.

By Proposition 2, the residual from (7) decomposes into two orthogonal components: one with vanishing divergence and one that is cardinally consistent. Let $\Pi_{\mathrm{grad}(\mathcal{K})}$ denote the projection onto $\mathrm{grad}(\mathcal{K})$, and recall that $\hat{Y}$ denotes the projection of $\bar{Y}$ onto the cardinally consistent subspace. By the Pythagorean theorem:

$$\left\|\bar{Y} - \Pi_{\mathrm{grad}(\mathcal{K})}\bar{Y}\right\|_2^2 = \left\|\bar{Y} - \hat{Y}\right\|_2^2 + \left\|\hat{Y} - \Pi_{\mathrm{grad}(\mathcal{K})}\hat{Y}\right\|_2^2. \tag{8}$$

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 (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 $\|\hat{Y} - \Pi_{\mathrm{grad}(\mathcal{K}_m)}\hat{Y}\|_2^2$ 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 9 (Quasilinearity Revisted)**.** Reconsider Example 1, 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 (7) with $\mathcal{K} = \mathcal{K}_{QIC}$ is equivalent to solving:

$$\begin{aligned}
\min_{u \in \mathcal{U}} \quad & \left\|(\mathrm{grad}\, u) - \bar{Y}\right\|_2^2 \\
\text{subject to} \quad & u_i = \langle \pi_i, v_i \rangle + \gamma_i & \forall i = 1, \ldots, K \\
& \langle \pi_i, v_i \rangle + \gamma_i \le \langle \pi_j, v_i \rangle + \gamma_j & \forall i, j = 1, \ldots, K \\
& \pi_i^1 = 1 & \forall i = 1, \ldots, K \\
& \pi_i \ge 0 & \forall i = 1, \ldots, K
\end{aligned} \tag{9}$$

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 (9).[55]

5.3.2 Calibration and Identification

In the experiment considered in Example 8, 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 (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 10. (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:

$$U(x) = \prod_{i=1}^{L} x_i^{\beta_i} \tag{10}$$

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 $\phi(\alpha, \tilde{x}) =$

[55] The system of constraints in (9) may be viewed as essentially a system of Afriat inequalities, but where prices are unknown.

$\tilde{x} + \alpha(1, \ldots, 1))$.[56] Thus it suffices to instead treat $\mathbb{R}^L$ as the consumption space, rather than $\mathbb{R}^L_{++}$. Computing (7) for Cobb-Douglas preferences is then equivalent to evaluating the following quadratic program:

$$\begin{aligned} \min_{u \in \mathcal{U}} \quad & \left\|(\operatorname{grad} u) - \bar{Y}\right\|_2^2 \\ \text{subject to} \quad & u_i = \big\langle \beta, H(v_i) \big\rangle \quad \forall i = 1, \ldots, K \\ & \beta \geq 0, \end{aligned} \tag{11}$$

where $\beta \in \mathbb{R}^L$.[57] Clearly each feasible vector (u, β) for (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 (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.

Example 11 (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 Example 8). 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 (7) for risk-neutral MEU preferences is equivalent the following quadratic program:

$$\begin{aligned} \min_{u \in \mathcal{U}} \quad & \left\|(\operatorname{grad} u) - \bar{Y}\right\|_2^2 \\ \text{subject to} \quad & u_i = \langle \pi_i, v_i \rangle & \forall i = 1, \ldots, K \\ & \langle \pi_i, v_i \rangle \leq \langle \pi_j, v_i \rangle & \forall i, j = 1, \ldots, K \\ & \langle \pi_i, \mathbb{1}_S \rangle = 1 & \forall i = 1, \ldots, K \\ & \pi_i \geq 0 & \forall i = 1, \ldots, K, \end{aligned} \tag{12}$$

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

[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 $\langle \beta, \mathbb{1}_L \rangle = 1$.

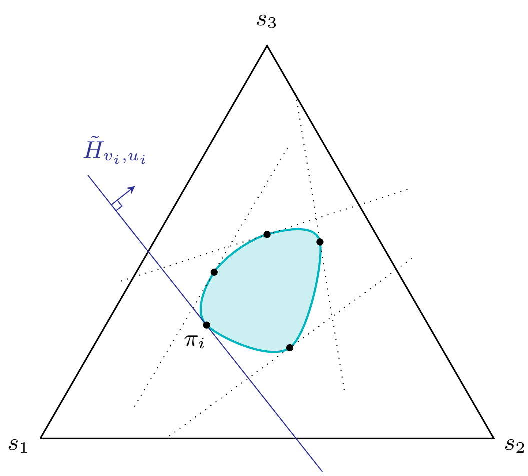

(a) The belief set C^* of a risk-neutral MEU preference. The vector $(u, \pi_1, \ldots, \pi_K)$ is a solution to (12), as for each $v_i \in \mathcal{V}$, the hyperplane $\tilde{H}_{v_i,u_i}$ supports C^* at π_i.

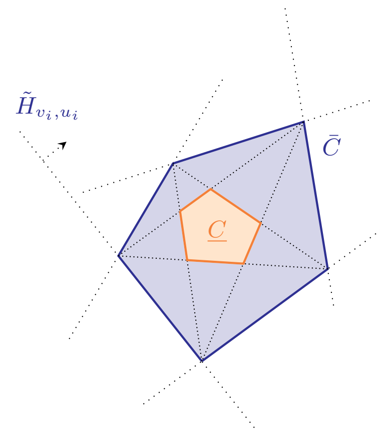

(b) Every MEU-rationalizable vector defines a belief polytope $\bar{C} = \cap_i \tilde{H}^+_{v_i,u_i}$. A belief set $C \subseteq \bar{C}$ rationalizes the data if and only if each facet of $\bar{C}$ contains some extremal point of C.

Figure 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 and lower envelope belief sets $\bar{C}$ and $\underline{C}$. Every belief set rationalizing the data is contained within $\bar{C}$ and contains $\underline{C}$.

Consider any feasible solution $(u, \pi_1, \ldots, \pi_K)$ to (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 (12) imply that each $\tilde{H}_{v_i,u_i}$ supports $\text{co}\{\pi_1, \ldots, \pi_K\}$ at each π_i, and hence $\text{co}\{\pi_1, \ldots, \pi_K\}$ is a set of priors consistent with the collection of utilities u. However, at any given solution to (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 $\hat{\pi}_i \in \text{ext}(C)$ from each $\tilde{H}_{v_i,u_i}$ 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 (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

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

full identification: if $\pi \notin \bar{C}$, it is not held by *any* risk-neutral MEU preference rationalizing the data.[59] Figure 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 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. 2000, see also Rigotti et al. 2008). 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 Example 8 and Example 11 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:

$$U(x) = \min_{\pi \in C} \mathbb{E}_\pi \big[\tilde{u}(x)\big], \tag{13}$$

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 Example 12 for such an estimator).

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

[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 (12) will generally be consistent with many sets of priors, the collection of such sets remains topologically negligible.

[60] Note, however, that generally $\underline{C}$ will not itself be consistent with u.

[61] See Billot et al. (2000) Theorem 1.

[62] It is straightforward to adapt to the case where $\tilde{u} : \mathbb{R}_+ \to \mathbb{R}$.

corresponding to (13) exactly as in (12), but instead replacing each v_i with its vector of utilities under $\tilde{u}$. For simplicity, in Appendix E 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.

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 12 (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 2 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 $\Theta = \{(s,i,j) \in S \times \{1,\dots,K\}^2 : \{i,j\} \in \mathcal{E} \text{ or } i = j\}$, 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:

$$\begin{aligned}
\min_{u,\beta,\gamma} \quad & \left\| (\operatorname{grad} u) - \tilde{Y} \right\|_2^2 & \\
\text{subject to} \quad & u_i = \sum_{s \in S} \pi_s \big(\gamma_{sii} + \beta_{sii} v_{sii}\big) & \forall i = 1,\dots,K \\
& \tilde{Y}_{ij} = \sum_{s \in S} \pi_s \big[(\gamma_{sij} + \beta_{sij} v_{sij}) - (\gamma_{sii} + \beta_{sii} v_{sii})\big] & \forall (s,i,j) \in \Theta \\
& \gamma_\theta + \beta_\theta v_\theta \le \gamma_{\theta'} + \beta_{\theta'} v_\theta & \forall \theta, \theta' \in \Theta \\
& \beta_\theta \ge 0 & \forall \theta \in \Theta \\
& \beta_{\bar{\theta}} = 1 & \\
& \gamma_{\bar{\theta}} = 0. &
\end{aligned} \tag{14}$$

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 $\tilde{u}$:

$$\hat{u}(x) = \min_{\theta \in \Theta} \gamma_\theta + \beta_\theta x. \tag{15}$$

For any increasing, concave, and unbounded above $\tilde{u}$, there is a $\hat{u}$ for which the value of (14) is zero.[64] Thus there exists some feasible solution to (14) which, for the experiment, is indistinguishable from the true $\tilde{u}$. Similarly, every feasible solution to (14) corresponds to an expected utility preference whose Bernoulli utility is increasing, concave, and unbounded above via (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 1992) or popular parameterizations of disappointment aversion (Gul 1991) such as in Routledge and Zin (2010). In the first stage, one estimates $\tilde{u}$ by varying payoffs under fixed, objectively known odds. In the second, one uses $\phi_{\hat{u}}$ while varying probabilities. Similar approaches allow for tests of additively separable models of dynamic preference

[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).

[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 (15).

à la Koopmans (1960) or quasi-hyperbolic preferences (e.g. Laibson 1997). See also Remark 4 for applications to translation-invariant ambiguity preferences.

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 nonparametric 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 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:

$$H_0 : Y^0 \in \mathrm{grad}(\mathcal{K}), \qquad H_1 : Y^0 \notin \mathrm{grad}(\mathcal{K}), \tag{16}$$

where $\mathcal{K} \subseteq \mathcal{U}$ is a set of shape constraints capturing the desired properties of U.[66] Following (7), the null and alternative hypotheses in (16) can be rephrased, up to a monotone transformation, as:

$$H_0 : \min_{u \in \mathcal{K}} \left\|(\mathrm{grad}\, u) - Y^0\right\|_2 = 0, \qquad H_1 : \min_{u \in \mathcal{K}} \left\|(\mathrm{grad}\, u) - Y^0\right\|_2 > 0. \tag{17}$$

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

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

A natural sample analogue of the objective function (17) is:

$$\psi(\bar{Y}) = \min_{u \in \mathcal{K}} \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2 = \min_{\hat{Y} \in \operatorname{grad}(\mathcal{K})} \|\hat{Y} - \bar{Y}\|_2 \tag{18}$$

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 Fang and Santos (2019), $\sqrt{N}\big(\psi(\bar{Y}) - \psi(Y^0)\big) \xrightarrow{L} \psi'_{Y^0}\big(N(0,\Sigma)\big)$, 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, Fang and Santos (2019) 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 $\psi'_{Y^0}\big(N(0,\Sigma)\big)$.

To obtain critical values for our statistic, first note that under H_0, $\psi(Y^0) = 0$, hence by Fang and Santos (2019), $\sqrt{N}\psi(\bar{Y}) \xrightarrow{L} \psi'_{Y^0}\big(N(0,\Sigma)\big)$. 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 ψ'_{Y^0}, and $\bar{Y}^*$ is a bootstrapped sample mean. The numerical derivative estimator of Hong and Li (2020) provides a convenient method simulating this distribution without any further analytic calculations:

1. For $b = 1, \ldots, B$, 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 $\{Y^1, \ldots, Y^N\}$.[68]

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

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

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 (17).

[67] See also Shapiro (1991) and Dümbgen (1993).

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

When an explicit description of the tangent cones of grad($\mathcal{K}$) is readily available, one can modify this procedure to make use of this extra analytic information (Fang and Santos, 2019). Similarly, when grad(K) is a closed, convex cone, Fang and Seo (2019) discuss a modification of this procedure with fine statistical properties.

Appendix A Proof of Theorem 1

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 Debreu (1964) 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:

$$\phi\big(c(x), \underline{x}\big) \sim x. \tag{19}$$

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

$$\phi\big(c(\phi(\alpha, x)), \underline{x}\big) \sim \phi(\alpha, x). \tag{20}$$

But by (19) and (N.1),

$$\phi\big(\alpha, \phi(c(x), \underline{x})\big) \sim \phi(\alpha, x), \tag{21}$$

and, as ϕ is an action:

$$\phi\big(\alpha, \phi(c(x), \underline{x})\big) = \phi\big(\alpha + c(x), \underline{x}\big). \tag{22}$$

Then by (20) - (22):

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

and by (N.2) we conclude:

$$\alpha + c(x) = c(\phi(\alpha, x)). \tag{23}$$

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 $c(x) \geq c(y)$.[69] 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.

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

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:

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

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. □

Appendix B Proof of Theorem 2

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.[70] 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.[71] 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. □

Appendix C Proof of Theorem 3

C.1 Overview

The proof of Theorem 3 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 Proposition 1 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 $v \in \mathcal{V}$:

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

[70] The max here is understood in the preference sense.

[71] That is, it comes from setting $s = 0$.

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.2 Construction of Embedding

Lemma 1. *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 $x \sim_{\trianglelefteq} z$. □

In light of Lemma 1, 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 1. *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 2. *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. □

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 3. *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 Lemma 2, $\bar{q}(U)$ is open. □

Lemma 4. *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 Corollary 1, and the last from invoking Lemma 2. 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 $\bar{s}$ (Lemma 5). We then establish the regularity (i.e. continuity) of solutions to the above class of topological implicit function problems (Lemma 6 - Lemma 8).

Lemma 5. *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 Corollary 1. 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 Corollary 1, 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. $\square$

Lemma 6. *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 Lemma 4, $\bar{s}\big(\cdot, \bar{q}(x)\big)$ 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 $[0, t(x)]$.

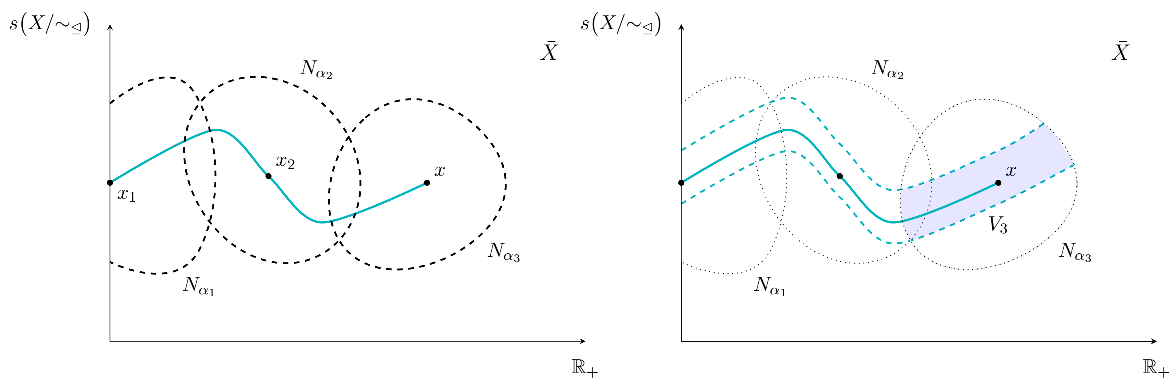

(a) An open cover of the path $\bar{s}([0, t(x)] \times \{\bar{q}(x)\})$, here in aquamarine. This open cover satisfies all of the properties of Lemma 6.

(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 6: An illustration of the construction underpinning Lemma 7. We have implicitly drawn the numeraire-paths of ϕ in $\bar{X}$ as vertical translates of one another.

As $\bar{s}([0, t(x)] \times \{\bar{q}(x)\})$ 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.[72] If instead $\underline{\alpha}_j < \underline{\alpha}_i$, 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.). □

Lemma 7. *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.*

[72] 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.

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 Lemma 6. Without loss of generality, suppose that N_{α_1} is the sole element to intersect $\bar{s}\big(\{0\} \times X/\sim_{\trianglelefteq}\big)$.[73] Define:

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

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

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

see Figure 6. We first verify, for all $i = 1, \dots, K$, that V_i is open. Note that via Lemma 3 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 $i = 1, \dots, K$,

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

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 $i = 1$.[74] For all $i = 1, \dots, K$, 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 Corollary 1 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 $1 \le i \le k$, that:

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

but, for sake of contradiction, suppose that:

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

[73] 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 $x \in \text{range}(s)$.

[74] This set is indeed an interval by Lemma 6.

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 Lemma 6 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 $\bar{s}\big([0, \bar{\alpha}_K) \times \{\bar{q}(x)\}\big) \subseteq \bigcup_{j \le K} V_j$, 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 $\bar{s}$.[75] 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 $i \le k$, $t|_{V_i}$ is bounded, and let $x' \in V_{k+1}$. 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 $1 \le j \le k$. Suppose $x'' \trianglelefteq x'$. Then:

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

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 $1 \le i \le K$, $t|_{V_i}$ is bounded; as $x \in V_K$ and V_K is open, this establishes the claim. □

[75] By definition,

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

and, appealing to Corollary 1 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 $t(\underline{x}') + \alpha = (t \circ \phi)(\alpha, \underline{x}') = t(x')$.

Lemma 8. *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 Lemma 7, 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 (Aliprantis and Border, 2006), 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 2. *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) $\Longrightarrow$ (ii): Let $\mathcal{E}$ be an experiment, and $F \in \mathcal{F}$ a cardinally consistent flow on $(\mathcal{V}, \mathcal{E})$. By Proposition 1, 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 $\mathcal{V} \subsetneq \bar{X}$.[76] 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 Lemma 2, $X/\!\sim_{\trianglelefteq}$ is homeomorphic to a subset of X and hence is metrizable and thus normal. By the Tietze extension theorem, e.g. Munkres (1974), 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.

[76] 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.'

We define a binary relation on X in three cases: first, if $x, y \in \bar{s}\big(\text{epi}(I)\big) \subseteq \bar{X}$, then let $x \succsim y$ if and only if $t(x) - t(y) \geq (I \circ \bar{q})(x) - (I \circ \bar{q})(y)$.[77] If x but not y belong to $\bar{s}\big(\text{epi}(I)\big)$, then let $x \succ y$. Finally, if neither x nor y belong to $\bar{s}\big(\text{epi}(I)\big)$, then let:

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

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, $\succsim$ is complete.[78] 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 Corollary 2, 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.

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

[78]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:

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

where the first equality follows from an argument identical to footnote 75, and the final from Corollary 1. 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 1, so it remains only to prove (iii) $\implies$ (i). Let $x_0, \ldots, x_L \in X$, 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 $l = 0, \ldots L$, 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 3 cannot be significantly relaxed without requiring a completely different proof approach.

Appendix D Proposition Proofs

D.1 Proof of Proposition 1

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, $\operatorname{grad} u = F$, completing the proof. □

D.2 Proof of Proposition 2

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

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

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

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

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 $\mathrm{im}(\mathrm{grad}) \oplus \ker(\mathrm{div})$. □

D.3 Proof of Proposition 3

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.[79] 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 $R \in \ker(\mathrm{div})$, 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.[80] 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 $0 \leq C_1 \leq R$ 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 $\ker(\mathrm{div})$; 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:

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

and hence the lower bound obtains. □

[79] 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).

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

Appendix E Shape Constraint 'Cookbook'

E.1 Quaslinear, Increasing, Concave Utility (Proof of Example 9)

Proof. Suppose first that U is a quasilinear, increasing, and concave utility. For all $i = 1, \dots, K$, 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 $i = 1, \dots, K$ 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 (Hiriart-Urruty and Lemaréchal (2004), 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 (9). Define:

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

Then clearly $U(v_i) = u_i$, and $\tilde{U}$ is quasilinear, increasing, and concave. □

E.2 Cobb-Douglas Preferences (Proof of Example 10)

Proof. Consider $X = \mathbb{R}^L_{++}$, and $\phi(\alpha, x) = e^\alpha x$. Define $H : X \to \mathbb{R}^L$ via:

$$H(x) = \big(\ln x_1, \dots, \ln x_L\big).$$

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 (7) under change of coordinates becomes:

$$\begin{aligned} \min_{u \in \mathcal{U}} \quad & \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2 \\ \text{subject to} \quad & u_i = \langle \beta, H(v_i) \rangle \quad \forall i = 1, \dots, K \\ & \beta \geq 0 \end{aligned} \tag{24}$$

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. □

E.3 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 Remark 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 14.

E.3.1 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 (7) with $\mathcal{K} = \mathcal{K}_{\text{SEU}}$ is equivalent to solving:

$$\begin{aligned} \min_{u \in \mathcal{U}} \quad & \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2 \\ \text{subject to} \quad & u_i = \langle \pi, v_i \rangle \qquad \forall i = 1, \dots, K \\ & \langle \pi, \mathbb{1}_S \rangle = 1 \\ & \pi \geq 0. \end{aligned} \tag{25}$$

Proof. Trivial. □

E.3.2 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:

$$C_\sigma = \{x \in \mathbb{R}^S : x_{\sigma(1)} \geq x_{\sigma(2)} \geq \cdots \geq x_{\sigma(S)}\}. \tag{26}$$

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:

$$P^\sigma(\sigma(i)) = \nu\big(\{\sigma(1), \sigma(2), \dots, \sigma(i)\}\big) - \nu\big(\{\sigma(1), \sigma(2), \dots, \sigma(i-1)\}\big). \tag{27}$$

See Ghirardato et al. (2004) 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 (7) with $\mathcal{K} = \mathcal{K}_{\text{CEU}}$ is equivalent to solving:

$$\begin{aligned}
\min_{u \in \mathcal{U}} \quad & \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2 \\
\text{subject to} \quad & u_i = \langle P^\sigma, v_i \rangle && \forall \sigma \in \mathfrak{S}_S, \forall i = 1, \dots, K \text{ s.t. } v_i \in C^\sigma \\
& P^\sigma_{\sigma(j)} = \nu_{A^\sigma_j} - \nu_{A^\sigma_{j-1}} && \forall \sigma \in \mathfrak{S}_S, \forall j = 1, \dots, S \\
& \nu_A \leq \nu_B && \forall A, B \in 2^S \text{ s.t. } A \subseteq B \\
& \nu_\varnothing = 0 \\
& \nu_S = 1
\end{aligned} \tag{28}$$

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 (28). From the discussion, e.g., in Ghirardato et al. (2004) (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 (28). Conversely, it follows trivially that every solution to (28) defines a CEU functional. □

E.3.3 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 (7) with $\mathcal{K} = \mathcal{K}_{\text{CCEU}}$ is equivalent to solving:

$$
\begin{aligned}
\min_{u\in\mathcal{U}} \quad & \left\|(\operatorname{grad} u) - \bar{Y}\right\|_2^2 \\
\text{subject to} \quad & u_i = \langle P^\sigma, v_i\rangle && \forall \sigma \in \mathfrak{S}_S, \forall i = 1,\dots,K \text{ s.t. } v_i \in C^\sigma \\
& P^\sigma_{\sigma(j)} = \nu_{A^\sigma_j} - \nu_{A^\sigma_{j-1}} && \forall \sigma \in \mathfrak{S}_S, \forall j = 1,\dots,S \\
& \nu_A \le \nu_B && \forall A, B \in 2^S \text{ s.t. } A \subseteq B \\
& \nu_A + \nu_B \le \nu_{A\cup B} + \nu_{A\cap B} && \forall A, B \in 2^S \\
& \nu_\varnothing = 0 \\
& \nu_S = 1
\end{aligned} \tag{29}
$$

Proof. Follows from CEU case, where additionally the supermodularity of ν is enforced. □

E.3.4 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 (7) with $\mathcal{K} = \mathcal{K}_{\text{MEU}}$ is equivalent to solving:

$$
\begin{aligned}
\min_{u\in\mathcal{U}} \quad & \left\|(\operatorname{grad} u) - \bar{Y}\right\|_2^2 \\
\text{subject to} \quad & u_i = \langle \pi_i, v_i\rangle && \forall i = 1,\dots,K \\
& \langle \pi_i, v_i\rangle \le \langle \pi_j, v_i\rangle && \forall i, j = 1,\dots,K \\
& \langle \pi_i, \mathbb{1}_S\rangle = 1 && \forall i = 1,\dots,K \\
& \pi_i \ge 0 && \forall i = 1,\dots,K,
\end{aligned} \tag{30}
$$

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 $v_i \in \mathcal{V}$:

$$
\begin{aligned}
U(x) &\le U(v_i) + \langle \pi_i, x - v_i\rangle \\
&= \langle \pi_i, v_i\rangle + \langle \pi_i, x - v_i\rangle \\
&= \langle \pi_i, x\rangle,
\end{aligned}
$$

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

$$\alpha \leq \alpha \langle \pi_i, \mathbb{1}_S \rangle. \tag{31}$$

If $\alpha > 0$, $1 \leq \langle \pi, \mathbb{1}_S \rangle$, and if $\alpha < 0$, $1 \geq \langle \pi, \mathbb{1}_S \rangle$. Since (32) holds for all $\alpha \in \mathbb{R}$, we obtain $\langle \pi_i, \mathbb{1}_S \rangle = 1$.

Suppose now that for some collection $\pi_1, \dots, \pi_K \in \Delta(S)$, 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

$$\hat{U}(x) = \min_{i \in \{1,\dots,K\}} \langle \pi_i, x \rangle = \min_{\pi \in \text{co}\{\pi_1,\dots,\pi_K\}} \langle \pi, x \rangle.$$

The latter equality follows from standard results on support functions see, e.g., Hiriart-Urruty and Lemaréchal (2004) Theorem 3.3.2. By construction, $u_i = \hat{U}(v_i)$ and $\hat{U}$ is a risk-neutral MEU functional. □

E.3.5 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 $i = 1, \dots, K$, let $\pi_i \in \partial U(v_i)$ be an arbitrary selection of supergradients of U, one at each v_i. For all $i = 1, \dots K$, 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

$$\alpha \leq \alpha \langle \pi_i, \mathbb{1}_S \rangle. \tag{32}$$

If $\alpha > 0$, $1 \leq \langle \pi, \mathbb{1}_S \rangle$, and if $\alpha < 0$, $1 \geq \langle \pi, \mathbb{1}_S \rangle$. Since (32) 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 Maccheroni et al. (2006), corresponds to a variational utility functional. □

E.3.6 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 $\mathcal{K}_{\mathrm{DS}}$ denote the collection of $u \in \mathcal{U}$ that are restrictions of dual-self utility functionals. Then solving (7) with $\mathcal{K} = \mathcal{K}_{\mathrm{DS}}$ is equivalent to solving:

$$
\begin{aligned}
\min_{u \in \mathcal{U}} \quad & \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2 \\
\text{subject to} \quad & u_i = \langle \pi_{ii}, v_i \rangle && \forall i = 1, \ldots, K \\
& \langle \pi_{ii}, v_i \rangle \leq \langle \pi_{ij}, v_i \rangle && \forall i, j = 1, \ldots, K \\
& \langle \pi_{ji}, v_i \rangle \leq \langle \pi_{ii}, v_i \rangle && \forall i, j = 1, \ldots, K \\
& \langle \pi_{ij}, \mathbb{1}_S \rangle = 1 && \forall i, j = 1, \ldots, K \\
& \pi_{ij} \geq 0 && \forall i, j = 1, \ldots, K,
\end{aligned} \tag{33}
$$

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 (33). Define, for each $i = 1, \ldots, K$, the set $P_i = \operatorname{co}\{\pi_{i,1}, \ldots, \pi_{i,K}\}$. 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.[81] 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 (33), as required.

□

E.3.7 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 Chandrasekher et al. (2020).

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 (7) with $\mathcal{K} = \mathcal{K}_{\text{DSV}}$ is equivalent to solving:

$$\begin{aligned} \min_{u \in \mathcal{U}} \quad & \left\| (\operatorname{grad} u) - \bar{Y} \right\|_2^2 \\ \text{subject to} \quad & u_i \geq u_j \qquad \forall i, j \text{ s.t. } v_i \geq v_j \\ & u_K = 0, \end{aligned} \tag{34}$$

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 (34).

[81] 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.

Conversely, suppose u is a solution to (34). In light of the characterization provided in Chandrasekher et al. (2020), it suffices to prove there exists an additive-equivariant and monotone extension from $\mathcal{V}$ to $\mathbb{R}^S$.[82] 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 Cerreia-Vioglio et al. (2014),

$$U(x) = \sup\{u_{v_i} + b : v_i \in \mathcal{V}, b \in \mathbb{R}, \text{ and } v_i + b\mathbb{1}_S \leq x\}$$

defines an additive-equivariant, monotone, and normalized extension of u, and hence by Chandrasekher et al. (2020) this corresponds to some dual-self variational utility functional. □

[82] Normalization holds for any additive-equivariant extension, as $u_K = 0$.

References

Abdellaoui, Mohammed, Aurélien Baillon, Laetitia Placido, and Peter P Wakker, "The rich domain of uncertainty: Source functions and their experimental implementation," *American Economic Review*, 2011, *101* (2), 695–723.

Aczél, János and J. Dhombres, *Functional equations in several variables* number 31, Cambridge university press, 1989.

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

Ahn, David, Syngjoo Choi, Douglas Gale, and Shachar Kariv, "Estimating ambiguity aversion in a portfolio choice experiment," *Quantitative Economics*, 2014, *5* (2), 195–223.

Aliprantis, Charalambos D. and Kim C. Border, *Infinite Dimensional Analysis: a Hitchhiker's Guide*, Berlin; London: Springer, 2006.

Allon, Gad, Michael Beenstock, Steven Hackman, Ury Passy, and Alexander Shapiro, "Nonparametric estimation of concave production technologies by entropic methods," *Journal of Applied Econometrics*, 2007, *22* (4), 795–816.

Andreoni, James and John Miller, "Giving according to GARP: An experimental test of the consistency of preferences for altruism," *Econometrica*, 2002, *70* (2), 737–753.

Anscombe, Francis J and Robert J Aumann, "A definition of subjective probability," *Annals of mathematical statistics*, 1963, *34* (1), 199–205.

Baillon, Aurelien, Yoram Halevy, Chen Li et al., "Experimental elicitation of ambiguity attitude using the random incentive system," *University of British Columbia working paper*, 2021.

Bayer, Ralph C, Subir Bose, Matthew Polisson, and Ludovic Renou, "Ambiguity revealed," *Available at SSRN 2046598*, 2013.

Becker, Gordon M, Morris H DeGroot, and Jacob Marschak, "Measuring utility by a single-response sequential method," *Behavioral science*, 1964, *9* (3), 226–232.

Billot, Antoine, Alain Chateauneuf, Itzhak Gilboa, and Jean-Marc Tallon, "Sharing beliefs: between agreeing and disagreeing," *Econometrica*, 2000, *68* (3), 685–694.

Bossaerts, Peter, Paolo Ghirardato, Serena Guarnaschelli, and William R Zame, "Ambiguity in asset markets: Theory and experiment," *The Review of Financial Studies*, 2010, *23* (4), 1325–1359.

Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper, "Risk and rationality: Uncovering heterogeneity in probability distortion," *Econometrica*, 2010, *78* (4), 1375–1412.

Candogan, Ozan, Ishai Menache, Asuman Ozdaglar, and Pablo A Parrilo, "Flows and decompositions of games: Harmonic and potential games," *Mathematics of Operations Research*, 2011, *36* (3), 474–503.

Caradonna, Peter, "How strong is the weak axiom," 2020.

Cattaneo, Matias D, Xinwei Ma, Yusufcan Masatlioglu, and Elchin Suleymanov, "A random attention model," *Journal of Political Economy*, 2020, *128* (7), 2796–2836.

Cerreia-Vioglio, Simone, Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini, "Niveloids and their extensions: Risk measures on small domains," *Journal of Mathematical Analysis and Applications*, 2014, *413* (1), 343–360.

Chambers, Christopher P, Federico Echenique, and Kota Saito, "Testing theories of financial decision making," *Proceedings of the National Academy of Sciences*, 2016, *113* (15), 4003–4008.

Chandrasekher, Madhav, Mira Frick, Ryota Iijima, and Yves Le Yaouanq, "Dual-self representations of ambiguity preferences," 2020.

Choi, Syngjoo, Raymond Fisman, Douglas Gale, and Shachar Kariv, "Consistency and heterogeneity of individual behavior under uncertainty," *American economic review*, 2007, *97* (5), 1921–1938.

Csató, László, "A graph interpretation of the least squares ranking method," *Social Choice and Welfare*, 2015, *44* (1), 51–69.

Deb, Rahul, Yuichi Kitamura, John K-H Quah, and Jörg Stoye, "Revealed price preference: theory and empirical analysis," *arXiv preprint arXiv:1801.02702*, 2018.

Debreu, Gerard, "Continuity properties of Paretian utility," *International Economic Review*, 1964, *5* (3), 285–293.

Demuynck, Thomas, Christian Seel, Giang Tran et al., "An index of competitiveness and cooperativeness for normal-form games," *American Economic Journal. Microeconomics*, 2020.

Dümbgen, Lutz, "On nondifferentiable functions and the bootstrap," *Probability Theory and Related Fields*, 1993, *95* (1), 125–140.

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

Fang, Zheng and Andres Santos, "Inference on directionally differentiable functions," *The Review of Economic Studies*, 2019, *86* (1), 377–412.

_ **and Juwon Seo**, "A Projection Framework for Testing Shape Restrictions That Form Convex Cones," *arXiv preprint arXiv:1910.07689*, 2019.

Fishburn, Peter C and Ariel Rubinstein, "Time preference," *International Economic Review*, 1982, pp. 677–694.

Fuchs, Laszlo, *Partially ordered algebraic systems*, Vol. 28, Courier Corporation, 2011.

Fudenberg, Drew, Whitney Newey, Philipp Strack, and Tomasz Strzalecki, "Testing the drift-diffusion model," *Proceedings of the National Academy of Sciences*, 2020, *117* (52), 33141–33148.

Ghirardato, Paolo, Fabio Maccheroni, and Massimo Marinacci, "Differentiating ambiguity and ambiguity attitude," *Journal of Economic Theory*, 2004, *118* (2), 133–173.

Gilboa, Itzhak and David Schmeidler, "Maxmin expected utility with non-unique prior," *Journal of mathematical economics*, 1989, *18* (2), 141–153.

Godsil, Chris and Gordon F Royle, *Algebraic graph theory*, Vol. 207, Springer Science & Business Media, 2001.

Grant, Simon and Ben Polak, "Mean-dispersion preferences and constant absolute uncertainty aversion," *Journal of Economic Theory*, 2013, *148* (4), 1361–1398.

Gul, Faruk, "A theory of disappointment aversion," *Econometrica: Journal of the Econometric Society*, 1991, pp. 667–686.

Halevy, Yoram, "Ellsberg revisited: An experimental study," *Econometrica*, 2007, *75* (2), 503–536.

_ , Dotan Persitz, and Lanny Zrill, "Parametric recoverability of preferences," *Journal of Political Economy*, 2018, *126* (4), 1558–1593.

Harville, David, "The use of linear-model methodology to rate high school or college football teams," *Journal of the American Statistical Association*, 1977, *72* (358), 278–289.

Hey, John D and Jinkwon Lee, "Do subjects separate (or are they sophisticated)?," *Experimental Economics*, 2005, *8* (3), 233–265.

_ and Noemi Pace, "The explanatory and predictive power of non two-stage-probability theories of decision making under ambiguity," *Journal of Risk and Uncertainty*, 2014, *49* (1), 1–29.

_ , Gianna Lotito, and Anna Maffioletti, "The descriptive and predictive adequacy of theories of decision making under uncertainty/ambiguity," *Journal of risk and uncertainty*, 2010, *41* (2), 81–111.

Hirani, Anil N, Kaushik Kalyanaraman, and Seth Watts, "Least squares ranking on graphs," *arXiv preprint arXiv:1011.1716*, 2010.

Hiriart-Urruty, Jean-Baptiste and Claude Lemaréchal, *Fundamentals of convex analysis*, Springer Science & Business Media, 2004.

Holmes, Richard B, *Geometric functional analysis and its applications*, Vol. 24, Springer Science & Business Media, 2012.

Holt, Charles A, "Preference reversals and the independence axiom," *The American Economic Review*, 1986, *76* (3), 508–515.

Hong, Han and Jessie Li, "The numerical bootstrap," *The Annals of Statistics*, 2020, *48* (1), 397–412.

Jiang, Xiaoye, Lek-Heng Lim, Yuan Yao, and Yinyu Ye, "Statistical ranking and combinatorial Hodge theory," *Mathematical Programming*, 2011, *127* (1), 203–244.

Kitamura, Yuichi and Jörg Stoye, "Nonparametric analysis of random utility models," *Econometrica*, 2018, *86* (6), 1883–1909.

Koopmans, Tjalling C, "Stationary ordinal utility and impatience," *Econometrica*, 1960, pp. 287–309.

Kuosmanen, Timo, "Representation theorem for convex nonparametric least squares," *The Econometrics Journal*, 2008, *11* (2), 308–325.

Laibson, David, "Golden eggs and hyperbolic discounting," *The Quarterly Journal of Economics*, 1997, *112* (2), 443–478.

Lee, Jinkwon, "The effect of the background risk in a simple chance improving decision model," *Journal of risk and uncertainty*, 2008, *36* (1), 19–41.

Loomes, Graham, "Evidence of a new violation of the independence axiom," *Journal of Risk and uncertainty*, 1991, *4* (1), 91–108.

Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini, "Ambiguity aversion, robustness, and the variational representation of preferences," *Econometrica*, 2006, *74* (6), 1447–1498.

Munkres, James R, *Topology; a First Course*, Prentice-Hall, 1974.

Ok, Efe A, *Real analysis with economic applications*, Princeton University Press, 2011.

_ **and Yusufcan Masatlioglu**, "A theory of (relative) discounting," *Journal of Economic Theory*, 2007, *137* (1), 214–245.

Osting, Braxton, Jérôme Darbon, and Stanley Osher, "Statistical ranking using the L1-norm on graphs," *Inverse Problems & Imaging*, 2013, *7* (3), 907.

Polisson, Matthew, John K-H Quah, and Ludovic Renou, "Revealed preferences over risk and uncertainty," *American Economic Review*, 2020, *110* (6), 1782–1820.

Rigotti, Luca, Chris Shannon, and Tomasz Strzalecki, "Subjective beliefs and ex ante trade," *Econometrica*, 2008, *76* (5), 1167–1190.

Routledge, Bryan R and Stanley E Zin, "Generalized disappointment aversion and asset prices," *The Journal of Finance*, 2010, *65* (4), 1303–1332.

Schmeidler, David, "Subjective probability and expected utility without additivity," *Econometrica: Journal of the Econometric Society*, 1989, pp. 571–587.

Seijo, Emilio and Bodhisattva Sen, "Nonparametric least squares estimation of a multivariate convex regression function," *The Annals of Statistics*, 2011, *39* (3), 1633–1657.

Shapiro, Alexander, "Asymptotic analysis of stochastic programs," *Annals of Operations Research*, 1991, *30* (1), 169–186.

Smeulders, Bart, Laurens Cherchye, and Bram De Rock, "Nonparametric analysis of random utility models: computational tools for statistical testing," *Econometrica*, 2021, *89* (1), 437–455.

Starmer, Chris and Robert Sugden, "Does the random-lottery incentive system elicit true preferences? An experimental investigation," *The American Economic Review*, 1991, *81* (4), 971–978.

Stefani, Raymond T, "Football and basketball predictions using least squares," *IEEE Transactions on systems, man, and cybernetics*, 1977, *7* (2), 117–21.

Tversky, Amos and Daniel Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *Journal of Risk and uncertainty*, 1992, *5* (4), 297–323.

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

Young, H Peyton, "An axiomatization of Borda's rule," *Journal of economic theory*, 1974, *9* (1), 43–52.

_ , "Social choice scoring functions," *SIAM Journal on Applied Mathematics*, 1975, *28* (4), 824–838.