A Least Squares Theory of Revealed Preference*†

Peter P. Caradonna‡

August 16, 2024

Abstract

We develop a least squares theory for the empirical study of models of preference and individual decision-making. Our approach utilizes common invariance properties of various models to obtain cardinal measurements of preference intensity. Our theory is widely applicable, and provides richer, more granular insights into the drivers of a model's predictive success or failure than traditional revealed preference methods, while simultaneously remaining computationally simple. We illustrate our methodology on common models of preferences over consumption bundles, dated rewards, lotteries, consumption streams, and Anscombe-Aumann acts.

1 Introduction

Dating back to at least Samuelson (1938), the study of the testable implications of models of individual preference and decision making has occupied a central

*This paper was previously circulated under the title 'Preference Regression.'

†I am indebted to Christopher Chambers for his encouragement and innumerable helpful discussions over the course of this project. I would also like to thank Marina Agranov, Axel Anderson, Simone Cerreia-Vioglio, Thomas Demuynck, Federico Echenique, Mira Frick, Ivana Komunjer, P.J. Healy, Roger Lagunoff, Jonathan Libgober, Yusufcan Masatlioglu, Kirby Nielsen, Efe Ok, Alexandre Poirier, Luciano Pomatto, John Quah, Debraj Ray, John Rehbeck, Margit Reischer, Mauricio Ribeiro, Philipp Strack, Omer Tamuz, Christopher Turansick, as well as seminar audiences at Caltech, Georgetown, OSU, NYU, Stanford GSB, UCLA, UC Riverside, UCSD, Yale, the 9th Annual ASU Economic Theory Conference, CETC, D-TEA, EC '24, METC, NASMES, and SAET for helpful comments.

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

position within both empirical and theoretical economics. In behavioral and decision theory, experimental falsification, or the discovery of 'paradoxes' documenting widespread empirical inconsistency with respect to various axioms, has been a long-standing driver of progress.[1] Related puzzles in macroeconomics and finance (e.g. the equity premium puzzle, see Mehra and Prescott 1985) have similarly led to the creation of new theories of individual behavior (Epstein and Zin 1989; Constantinides 1990).

A core underlying question in each of these instances is how to evaluate whether the predictions of a theory (e.g. the expected utility hypothesis or constant relative risk aversion) are sufficiently consistent with the observed data. Any such theory necessarily only describes individual behavior in a stylized fashion; because of this, we expect that no theory will perfectly explain any sufficiently rich data set. Thus it is critical to understand not only whether a model is consistent with the data, but rather how best to quantify the *degree* of any observed inconsistency.

For consumption data, a seemingly natural approach is to use standard econometric notions of loss (e.g. mean squared error) to quantify the magnitude of the deviation between observed and predicted demands generated by some theory. However, Varian (1990) argues that this approach (i.e. quantifying inconsistency via deviations between *choices*) despite its tractability, reflects the statistical, rather than economic, significance of violations, which may be unrelated. Instead, Varian argues, one should rely on revealed preference type inconsistency indices, which admit a more natural economic interpretation.[2,3] Nonetheless, these are not without their own disadvantages: such measures are at best imperfect proxies for economically meaningful quantities, tend to be computationally difficult and, in the presence of noise or error, present

[1] E.g., Allais (1953); Ellsberg (1961); Kahneman and Tversky (1979); Rabin (2000).

[2] For examples of such indices, see Afriat (1972); Houtman and Maks (1985); Varian (1990); Echenique et al. (2011); Dean and Martin (2016).

[3] However, stronger interpretations than warranted are often attributed to these indices; see Echenique (2021) for discussion of this point.

statistical challenges.[4,5]

This paper develops a new and widely applicable methodology for evaluating the predictive accuracy of models of preference and individual decision-making. Our approach applies to any theory which satisfies a common form of 'invariance' property. It hinges on a novel generalization of the concept of numeraire commodity, which may be defined for any such family. Using willingness-to-pay measurements denominated in such a 'virtual numeraire,' we are able to quantify the extent to which the data deviate from model predictions in a principled and transparent way.

Our approach enjoys a number of advantages over traditional revealed preference methods. We show that any data vector in our setting can always be *uniquely* decomposed into a rationalizable component, and a sum of (cardinal) revealed preference cycles. Many ordinal methods for trying to extract consistent rankings from choice data require strong completeness assumptions to guarantee the existence or reasonableness of predictions.[6] In contrast, our methodology provides a natural means of distinguishing inconsistency from underlying rationalizable content for general data sets, and without further assumptions.

Our approach is also computationally simple, particularly in comparison to many ordinal inconsistency indices.[7] Quantifying deviation requires evaluating a standard least squares program, subject to finitely many linear inequality constraints. Using our notion of virtual numeraire, we show that this minimization may be regarded as taking place over an appropriate space of utilities. In particular, the value of our least-squares program corresponds to the minimal (quadratically-weighted) *utility loss* that would be suffered by

[4]There is a growing literature on the computational difficulty of computing various inconsistency indices for revealed preference data, see Cherchye et al. (2015); Dean and Martin (2016); Smeulders et al. (2013, 2014, 2021).

[5]E.g., Echenique et al. (2011).

[6]See, for example, Bernheim and Rangel (2009); Nishimura (2018).

[7]E.g., Echenique et al. (2011); Dean and Martin (2016); Smeulders et al. (2013, 2014, 2021).

any model-consistent agent, were they to generate the observed data.[8] This provides additional nuance for evaluating model fit. For example, it may be the case that violations of some theory are frequently observed in ordinal data, yet when cardinal measurements are obtained, the overall utility loss due to these violations is small. In such cases, it may be reasonable to accept the theory as a good approximation of behavior.

The analytic structure afforded by our cardinal data additionally allows us to provide a suite of tools for analyzing a variety of related questions, often more tractably than existing ordinal methods. The solution to our least squares program provides a best-fit estimator, which allows us to straightforwardly select the 'most consistent' preferences from some theory, even when none are fully compatible with the data. For parametric models, this provides an economically meaningful way of obtaining point estimates of parameters. However our approach retains all its power and simplicity even when applied to more complex, non-parametric theories.

By examining which inequality constraints bind at our solution, we are also often able to provide insight into which individual axioms of a theory are most (or least) well-supported by the data. In Section 5, we illustrate how use this to compute the 'shadow price,' in model fit terms, of Gilboa and Schmeidler (1989)'s ambiguity aversion axiom.

Finally, in the presence of stochastic noise or measurement error, we provide a general means of leveraging cross-sectional observations to construct distribution-free statistical tests of consistency. Our approach again remains equally valid for both parametric and non-parametric theories.

To more concretely illustrate our approach, the following example considers the simplest case of quasilinear preferences, where our notion of 'virtual numeraire' reduces to the numeraire commodity. This allows us to highlight how numeraire-denominated measurements can be used to naturally measure deviations from a theory in an economically meaningful fashion.

[8] While our primary focus is on the least squares theory, it is straightforward to consider other notions of loss within our framework. We discuss this in more detail in Section 4.3.1 and Section 4.3.2.

An illustrative example: Consider a subject who has preferences over bundles of (say) money and apples, modelled as elements of $(m, a) \in \mathbb{R}^2_+$. Suppose we wish to evaluate how consistent the subject's behavior is with the maximization of a quasilinear utility $u : \mathbb{R}^2_+ \to \mathbb{R}$ of the form:

$$u(m, a) = v(a) + m, \tag{1}$$

for some concave and increasing $v : \mathbb{R}_+ \to \mathbb{R}$. We consider data of the following form: there is a collection of three bundles $(m_1, a_1), (m_2, a_2)$, and (m_3, a_3), in general position in $\mathbb{R}^2_+$, and the subject is presented with all three possible pairs of bundles from this collection. For each pair of bundles, we observe both (i) which bundle is preferred, and (ii) what quantity α_{ij} of numeraire (here, money) must be added to the less-preferred bundle i to make the subject indifferent between it and the more-preferred bundle j.

Our goal is to test the extent to which the data vector α is consistent with the model. To be consistent with the representation (1), there must exist utilities u_i for each bundle such that $u_i = u(m_i, a_i) = v(a_i) + m_i$ for some fixed choice of v. The set of (u_1, u_2, u_3) that arise in this fashion is defined by a simple, finite system of linear inequalities.

As the vector α is denominated in numeraire, by (1), it may be identified with a vector of differences in *utilities*, under the null hypothesis of consistency. Hence it is natural to quantify deviations from the predictions of our model by measuring the distance from our vector α to the set of differences generated by vectors (u_1, u_2, u_3) consistent with the model. This corresponds to minimizing the mean squared error between the observed and predicted *utility differences*, and requires only solving a simple least squares problem, subject to a finite collection of linear inequality constraints.

The solution to this program characterizes the set of convex, increasing, and quasilinear preferences which best fit the vector of observations. This allows us, for example, to straightforwardly select a best-fitting preference from our model, even when our vector of observations is inconsistent with every such preference.

In fact, a more granular analysis is possible. Our decomposition result implies that any observed error can be uniquely decomposed as error attributable to (i) the subject exhibiting money-pump like behavior (hence failing to treat money as a numeraire) and (ii) violations of the monotonicity and concavity axioms. Indeed, under the null of consistency, the mean squared error obtained from our fitting problem precisely equals the sum of (squared) utility losses arising from money pump violations and violations of our other model axioms.[9]

Moreover, by examining which constraints bind at our solution, we can quantify how much fit would improve by relaxing the assumptions of monotonicity or convexity (or both). In economic terms, this allows us to quantify how much *less* utility a subject would need to lose, under a weakened null hypothesis, in order to explain the data. This allows for far more refined feedback than is typically available for, e.g., model selection exercises. ■

Perhaps surprisingly, quasilinearity turns out to be far from necessary to achieve these results. Suppose instead we are interested in preferences on some abstract consumption space X. We first introduce the notion of a 'virtual commodity,' formally a collection $\{\phi_\alpha\}_{\alpha\geq 0}$ of transformations, each mapping $X \to X$. For any alternative x, we interpret $\phi_\alpha(x)$ as representing x, plus α additional units of the virtual commodity. To ensure consistency between these transformations, we require that $\phi_\beta\big(\phi_\alpha(x)\big) = \phi_{\alpha+\beta}(x)$. This simply says adding β units of the virtual commodity to the alternative which already consists of x plus α units, must yield the same outcome as adding $\alpha+\beta$ units to x all at once.

Given such a commodity, our first main result provides necessary and sufficient conditions for each preference in some family to admit a representation satisfying the system of simultaneous functional equations:

$$u\big(\phi_\alpha(x)\big) = u(x) + \alpha, \tag{2}$$

for all $x \in X$ and each $\alpha \geq 0$. If such a representation obtains for every preference, we say that $\{\phi_\alpha\}_{\alpha\geq 0}$ defines a virtual *numeraire* commodity for the family. When X consisted of bundles $(m, a) \in \mathbb{R}^2_+$ and each ϕ_α was the

[9]For a discussion of other loss functions, see Section 4.3.1 and Section 4.3.2.

transformation that added α extra dollars to any bundle, satisfaction of (2) was equivalent to quasilinearity in money. However, many more general families of preference admit utilities satisfying (2), for appropriately chosen families $\{\phi_\alpha\}_{\alpha\geq 0}$.

For example, suppose instead that X consists of all lotteries over prizes $\{\$1, \$5, \$10\}$ except $\delta_{\$10}$, the lottery that pays \$10 with probability one. For each $\alpha \geq 0$, define the transformation $\phi_\alpha(p) = e^{-\alpha}p + (1 - e^{-\alpha})\,\delta_{\$10}$. Here, 'adding' virtual commodity to a lottery p corresponds to mixing p with $\delta_{\$10}$. Given an increasing, expected utility preference, its standard representation(s):

$$U(p) = \sum_{i\in\{\$1,\$5,\$10\}} p_i\, v(i)$$

will not satisfy (2). However, the monotone transformation $u(p) = -\ln\big[\bar{c} - U(p)\big]$, where $\bar{c} \equiv v(\$10)$ is a normalizing constant, does.

This provides a very general approach to obtaining measurements of preference intensity for a variety of theories or models. We first select a family of transformations $\{\phi_\alpha\}_{\alpha\geq 0}$ that form a virtual numeraire for our theory. Using this choice, we proceed analogously to the quasilinear case: for some collection of binary subsets of X, we elicit (i) which alternative in the pair is preferred (e.g. $x' \succsim x$), and (ii) what quantity of additional, virtual numeraire must be added to the less-preferred alternative to render the subject indifferent between it and the more-preferred (e.g. for which $\alpha^* \geq 0$ is $\phi_{\alpha^*}(x) \sim x'$).[10] Just as in the quasilinear case, compensation measurements denominated in virtual numeraire form an exact proxy for the utility difference between x' and x, but now this difference is measured under any utility satisfying (2).

The paper proceeds as follows. In Section 2 we review related work. Section 3 characterizes the existence of representations satisfying (2). We also provide explicit examples of virtual numeraires $\{\phi_\alpha\}_{\alpha\geq 0}$ for many common classes of preferences over commodity bundles, dated rewards, lotteries, consumption streams, and Anscombe-Aumann acts. Section 4 develops our least-squares theory, illustrated above in the context of quasilinear preferences, for

[10] In Online Appendix E we provide a dominant strategy incentive-compatible mechanism for truthfully eliciting such measurements, for general X and choice of $\{\phi_\alpha\}_{\alpha\geq 0}$.

general virtual numeraire-denominated data. This requires a novel continuous extension theorem for invariant preferences, which is our main workhorse result. Section 5 illustrates the power of our approach in the context of the maxmin expected utility model, and Section 6 considers extensions to the case of stochastic data.

2 Related Literature

The revealed preference literature is too large to adequately survey here; for an excellent overview, see Chambers and Echenique (2016). Nishimura et al. (2017) prove an elegant continuous rationalization result for choice data. Our results differ from theirs at a technical level in two respects: first, our data is cardinal, rather than ordinal, and second, our continuous rationalizing preference is also required to satisfy certain additional *invariance* properties (cf. Ok and Riella 2014, 2021). Similar to us, Chambers et al. (2021) consider revealed preference data drawn from binary choice sets, however in our setting, the ordinal choice data is complemented with additional willingness-to-pay observations.

This paper also contributes to the literature on inconsistency measures for revealed preference data. Among the first to consider this was Afriat (1972) in the context of price-consumption data (see also Polisson and Quah 2022). Echenique et al. (2011) study a money pump index for price-consumption data. We obtain a virtual numeraire-valued analogue of their money pump index in our setting.[11] These papers provide economically natural measures of deviation from rational behavior; in contrast, this paper is concerned with quantifying deviations from the predictions of specific models.

Our concept of virtual numeraire relates to earlier work in measurement theory, e.g. Krantz et al. (2007). Historically, this focused on the measurement of intensity (including preference intensity) across various factors of a product space, and studied which means of evaluating trade-offs across these factors

[11] Houtman and Maks (1985); Varian (1990) are also classical contributions. See also Fudenberg and Liang (2020) for a different approach.

led to additive representations (e.g. Debreu 1959). Our methodology, based around finding an appropriate, exogenous scale $\{\phi_\alpha\}_{\alpha \geq 0}$ to obtain consistent measures of intensity, may be viewed as a 'coordinate-free' generalization of this approach.[12]

In the context of stochastic data, a number of papers study the problem of statistically testing rationalizability for various models. Deb et al. (2023) consider statistical tests of models of price preference. Kitamura and Stoye (2018) provide a non-parametric test for the random utility framework.[13] Fudenberg et al. (2020) provide a test of the drift-diffusion model. Blundell et al. (2008) provide a test of the classical revealed preference axioms for demand data. They observe that rationalizable demand responses to price changes can be characterized by a system of moment inequalities. In Section 6 we exploit similar structure to derive nonparametric statistical tests of rationalizability for a wide variety of models.

On risk domains Smith (1961) and Savage (1971) propose using the *probability* of a subject winning a given prize itself as a utility-linear unit of compensation in incentive schemes for expected utility maximizers.[14] Similarly, Roth and Malouf (1981, 1982) feature an experimental design in which subjects bargain over probability units to ensure constant marginal utility. Our usage of suitable families of transformations $\{\phi_\alpha\}_{\alpha \geq 0}$ to obtain utility-linear measurements generalizes these approaches, and suggests natural applications beyond the problems considered here.

3 A Measure of Intensity of Preference

Our objective in this section is to introduce a method of obtaining consistent, cardinal measurements of preference intensity for a given family of preferences.

[12] For a formal statement of this interpretation, see Theorem 3.

[13] Deb et al. (2018) use a similar approach to develop tests for a model in which consumers also exhibit price preference.

[14] Smith cites Savage (1954) for the origin of this idea.

3.1 Virtual Commodities

Let X be a metric space of alternatives, or consumption space. A preference $\succsim$ is a complete and transitive binary relation on X. Given a preference $\succsim$, we use $\succ$ and $\sim$ to denote the asymmetric and symmetric components, respectively. A preference is continuous if $\{y : y \succsim x\}$ and $\{y : x \succsim y\}$ are closed for all $x \in X$.

A family of transformations $\{\phi_\alpha\}_{\alpha\geq 0}$, where $\phi_\alpha : X \to X$ for all $\alpha \geq 0$, defines a **virtual commodity** if (i) for all $x \in X$, $\phi_0(x) = x$, and (ii) for all $\alpha, \beta \geq 0$ and $x \in X$, $\phi_\beta\big(\phi_\alpha(x)\big) = \phi_{\alpha+\beta}(x)$. We will always assume that any virtual commodity is jointly continuous in α and x.[15] We interpret the alternative $\phi_\alpha(x)$ as x plus α additional units of the virtual commodity. Property (i) requires that adding no units of commodity does not alter any alternative. Property (ii) is a path-independence condition that requires adding β units of commodity, to the alternative consisting of x plus α units, be equivalent to adding $\alpha + \beta$ units to x at once.

Let $\mathcal{M}$ denote a family of continuous preference relations on X, or **model**. We say that $\{\phi_\alpha\}_{\alpha\geq 0}$ is a **virtual numeraire** commodity for $\mathcal{M}$ if, for each $\succsim \in \mathcal{M}$ the following conditions are satisfied:

(N.1) **Invariance**: For all $\alpha \geq 0$, $x, x' \in X$:

$$x \succsim x' \iff \phi_\alpha(x) \succsim \phi_\alpha(x').$$

(N.2) **Monotonicity**: For all $\alpha \geq 0$, $x \in X$:

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

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

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

$$x' \succsim x \implies \exists\, \alpha \geq 0 \text{ s.t. } \phi_\alpha(x) \sim x'.$$

[15]That is, the map $(\alpha, x) \mapsto \phi_\alpha(x)$ is continuous in the product topology on $\mathbb{R}_+ \times X$.

Invariance says that adding the same quantity of the commodity to two alternatives does not affect the preference between them. It rules out cases where adding some common quantity of the commodity to two alternatives causes the preference between them to reverse. Monotonicity says the virtual commodity is a good. Compensability is a richness condition; it rules out lexicographic-like behavior where no amount of additional commodity could compensate an agent for receiving a less-preferred alternative.

3.2 Consistent Measurement of Intensity

Given a virtual numereraire $\{\phi_\alpha\}_{\alpha\geq 0}$ for some model $\mathcal{M}$, we wish to use these transformations as a system of rulers, to obtain a systematic measure of preference intensity across pairs of alternatives. Suppose, for some $\succsim\in\mathcal{M}$, that $x'\succsim x$. We define the **compensation difference** between the less-preferred x and more-preferred x' to be the $\alpha_{xx'}\geq 0$ such that $\phi_{\alpha_{xx'}}(x)\sim x'$. Such a quantity $\alpha_{xx'}\geq 0$ is guaranteed to exist by (N.3) and is unique by (N.2). Symmetrically, the compensation difference from x' to x is defined as $-\alpha_{xx'}$.

To derive testable implications for compensation differences data, we will rely on the following theorem, which is a generalization of standard results on quasilinear representation.

Theorem 1. *A virtual commodity $\{\phi_\alpha\}_{\alpha\geq 0}$ is a virtual numeraire for the model $\mathcal{M}$ if and only if every preference $\succsim\in\mathcal{M}$ admits a continuous utility representation $u: X\to\mathbb{R}$ such that, for all $x\in X$ and all $\alpha\geq 0$:*

$$u\big(\phi_\alpha(x)\big)=u(x)+\alpha. \tag{3}$$

Such a representation is unique up to an additive constant; in particular its utility differences are identified.

We term any utility satisfying (3) a **ϕ-additive** representation. Theorem 1 justifies our use of the term 'numeraire' to describe any virtual commodity satisfying (N.1) - (N.3). It guarantees that each transformation ϕ_α may be viewed as adding α utils of benefit, for *any* preference in $\mathcal{M}$, independently of the alternative to which it is applied. This in turn ensures that *observable* willingness-to-pay measurements denominated in ϕ provide an exact proxy for *unobservable* utility values.

This identification yields testable implications for data. Suppose we observe compensation differences data, measured using $\{\phi_\alpha\}_{\alpha\geq 0}$, over some collection of pairs of alternatives. If a subject's behavior is consistent with any preference satisfying (N.1) - (N.3), Theorem 1 implies that for any subcollection of pairs $\{x_0, x_1\}$, $\{x_1, x_2\}$, ..., $\{x_{L-1}, x_0\}$, we must have:

$$\sum_{l=0}^{L-1} \alpha_{x_l x_{l+1}} = \sum_{l=0}^{L-1} u(x_{l+1}) - u(x_l) = 0, \tag{4}$$

where $x_L \equiv x_0$. Equation (4) is an 'adding-up' condition: it simply says that the compensation difference between x_0 and x_{L-1} must equal the sum of the compensation differences $\alpha_{x_0 x_1} + \alpha_{x_1 x_2} + \cdots + \alpha_{x_{L-2} x_{L-1}}$. In Section 4, we show that for any experiment, these adding-up conditions in fact characterize consistency, and we leverage this to construct our regression theory.

Theorem 1 also provides guidance for the process of finding virtual numeraires. It highlights that the crucial property needed is that the preferences of our model $\mathcal{M}$ be suitably *invariant* under the transformations $\{\phi_\alpha\}_{\alpha\geq 0}$. Such families are often discernible either by inspection, or from existing axiomatic work. In Section 3.3, we provide a number of examples of virtual numeraires for many common classes of preference, across a variety of domains. Taken together, these examples illustrate that while quasilinearity is commonly viewed as a simple, limited class of preferences, analogous 'quasilinear structure,' as captured by Theorem 1, actually obtains in a far wider range of settings and models.

3.3 Examples

In this section, we turn to a number of examples of virtual numeraires, for a variety of models of economic interest. Our objective is to not only provide a ready-made list of such numeraires for applied work, but also highlight that in practice, the problem of finding a virtual numeraire for models is often straightforward.

Quasilinear Preferences

Let $X = \mathbb{R}^L_+$, and let $\mathcal{M}$ denote the collection of all continuous preferences that are quasilinear with respect to the first commodity. Here, the numeraire commodity itself may be represented as a virtual numeraire for $\mathcal{M}$, by defining $\phi_\alpha(x) = x + (\alpha, 0, \ldots, 0)$. The utility representation:

$$u(x) = v(x_2, \ldots, x_L) + x_1,$$

for any $\succsim \in \mathcal{M}$ is ϕ-additive, and by Theorem 1, every ϕ-additive utility is of this form.

Stationary Preferences for Dated Rewards

Consider a decision maker who has preferences over prizes $z \in Z$, delivered at some time $t \in \mathbb{R}_+$ in the future. Let $X = \mathbb{R}_+ \times Z$. An element $(t, z) \in X$ corresponds to the 'dated reward' featuring the prize z being delivered to the decision-maker at time t.[16] Following Fishburn and Rubinstein (1982), let $\mathcal{M}$ denote those preferences which admit an exponentially discounted utility:

$$\hat{u}(t, z) = \rho^t v(z),$$

where $0 < \rho < 1$ and $v : Z \to \mathbb{R}_{++}$. To such a decision-maker, time is a 'bad,' as it delays receipt of the desirable prize z. If we instead ask that (N.2) and (N.3) hold with the opposite relations, to reflect impatience, and adapt the definition of ϕ-additivity correspondingly, $\phi_\alpha(t, z) = (t + \alpha, z)$ defines a virtual numeraire for $\mathcal{M}$.[17] Compensation differences here reflect how long the receipt

[16] See also Ok and Masatlioglu (2007) for a complementary interpretation of preferences over X as the commitment preferences of an agent.

[17] Formally, we ask that for all $x \in X$, and $\alpha > 0$, $\phi_\alpha(x) \prec x$, and if $y \succ x$, there exists some $\alpha > 0$ such that $\phi_\alpha(y) \sim x$. Correspondingly, we understand ϕ-additivity in this context to mean $u\big(\phi_\alpha(x)\big) = u(x) - \alpha$.

of a more-desirable dated prize must be delayed to make the decision-maker indifferent with a less-preferred. While the standard representation $\rho^t v(z)$ is not ϕ-additive, a monotone transformation is:

$$u(t,z) = \frac{-1}{\ln\,\rho}\ln\left[\rho^t v(z)\right] = \frac{\ln v(z)}{-\ln\rho} - t.$$

Homothetic Preferences

Let $X = \mathbb{R}^L_+ \setminus \{0\}$ and let $\mathcal{M}$ denote the collection of all preferences on X which admit a continuous, strictly increasing, and positively homogeneous utility function.[18] Then $\phi_\alpha(x) = e^\alpha\, x$ defines a virtual numeraire for $\mathcal{M}$. For this choice of $\{\phi_\alpha\}_{\alpha\geq 0}$, (N.1) is equivalent to each $\succsim \in \mathcal{M}$ being homothetic.[19] If $\hat{u}$ is a strictly increasing and positively homogeneous representation for a preference in $\mathcal{M}$, then the monotone transformation:

$$u(x) = \ln \hat{u}(x)$$

is easily seen to be ϕ-additive. Moreover, by Theorem 1, every ϕ-additive representation of preferences in $\mathcal{M}$ is of this form, up to an additive constant. Compensation differences here may be interpreted as the amount a less-preferred bundle must be proportionally scaled to achieve indifference with a more-preferred.

Expected Utility Preferences

Let $X = \Delta\big([0,1]\big) \setminus \{\delta_1\}$ denote the space of all monetary lotteries over $[0,1]$ less δ_1, the point-mass at 1. Let $\mathcal{M}$ denote the space of all strictly increasing expected utility (EU) preferences on X. Then $\phi_\alpha(p) = e^{-\alpha}p + \big(1 - e^{-\alpha}\big)\delta_1$ defines a virtual numeraire for $\mathcal{M}$. For a given $\succsim \in \mathcal{M}$, let:

$$\hat{u}(p) = \int v\, dp$$

[18] A utility $\hat{u}$ is positively homogeneous if, for all $\lambda > 0$ and $x \in X$, $\hat{u}(\lambda x) = \lambda \hat{u}(x)$. If $\hat{u}$ is additionally strictly increasing, it must necessarily take strictly positive values on X.

[19] A preference on X is homothetic if, for all $x, y \in X$ and all $\lambda > 0$, $x \succsim y \iff \lambda x \succsim \lambda y$.

be a normalized expected utility representation, with $v(1) = 0$. Then:

$$u(p) = -\ln\big[-\hat{u}(p)\big]$$

represents the same preference and may be seen to be ϕ-additive. By Theorem 1, all ϕ-additive representations for preferences in $\mathcal{M}$ are of this form, up to an additive constant. Here, compensation differences measure how much a less-preferred lottery must be mixed with δ_1 to achieve indifference with a more-preferred.[20]

Many common classes of EU preferences admit other natural choices of virtual numeraire. If $\mathcal{M}$ consists of the collection of all constant absolute risk aversion (CARA) EU preferences, then defining $\phi_\alpha(p)$ as providing p plus α units of additional guaranteed wealth yields a virtual numeraire.[21] Here, the ϕ-additive representations are equal to the certainty equivalent, up to an additive constant. In fact, $\{\phi_\alpha\}_{\alpha\geq 0}$ is a virtual numeraire for a larger class of (not necessarily EU) preferences; see e.g. Safra and Segal (1998); Mu et al. (2021). Analogous results obtain for the class of constant relative risk aversion preferences.

Geometric & Quasi-hyperbolic Discounting

Let X denote the space of all bounded, infinite-horizon, discrete-time consumption streams. Let $\mathcal{M}$ denote the collection of all quasi-hyperbolic preferences on X (e.g. Laibson 1997) with fixed continuous, strictly increasing, and unbounded flow utility $v : \mathbb{R}_+ \to \mathbb{R}_+$, i.e. those preferences admitting a utility of the form:

$$u(x) = v(x_0) + \beta \sum_{t=1}^{\infty} \delta^t v(x_t),$$

[20] Note that in normalizing the Bernoulli utility $v(1) = 0$, we have used the additive degree of freedom in the standard expected utility representation. The remaining multiplicative degree of freedom then precisely becomes the additive degree of freedom in the ϕ-additive representation.

[21] Here, let X be e.g. the set of all finitely supported monetary lotteries, to ensure $\{\phi_\alpha\}_{\alpha\geq 0}$ is well defined.

where $0 < \beta \leq 1$ and $0 < \delta < 1$. When $\beta = 1$, this reduces to the standard geometric discounting model, see Koopmans (1960). For all $t > 0$, define $\phi_\alpha(x)_t = x_t$, and for $t = 0$, let $\phi_\alpha(x)_0 = v^{-1}\big(v(x_0) + \alpha\big)$.[22] It follows immediately that every such u is ϕ-additive, as is any utility of the form:

$$u(x) = v(x_0) + w(x_1, x_2, \dots).$$

There are many other closely related virtual numeraires for $\mathcal{M}$. For example, letting $\phi'_\alpha(x)_t = v^{-1}\big(v(x_t) + \alpha\big)$ for all periods in some set T, rather than just $T = \{0\}$, yields a virtual numeraire as well.

Constant Absolute Ambiguity Aversion

Let S denote a finite set of states of the world, and $X = \mathbb{R}^S$ denote the pure-ambiguity domain of (risk-free) monetary acts. Let $\mathcal{M}$ denote the set of continuous preferences on X that admit a utility representation of the form:

$$u(x) = w\big(v(x_1), \dots, v(x_S)\big), \tag{5}$$

where $v : \mathbb{R} \to \mathbb{R}$ is a fixed, strictly increasing, and unbounded-above state-contingent utility, and $w : \mathbb{R}^S \to \mathbb{R}$ is an increasing, translation-invariant utility functional, i.e. satisfying:

$$w\big(y + \alpha \mathbb{1}_S\big) = w(y) + \alpha$$

for all $\alpha \geq 0$, $y \in \mathbb{R}^S$, where $\mathbb{1}_S$ denotes the vector of all ones.[23] $\mathcal{M}$ may be viewed as encapsulating a wide range of ambiguity attitudes, given fixed risk attitude.[24] It includes, for example, the subjective expected utility, Choquet

[22] Here $\phi_\alpha(x)_t$ denotes the t-th component of $\phi_\alpha(x)$.

[23] Grant and Polak (2013) interpret translation-invariance over utility acts as reflecting constant absolute ambiguity aversion.

[24] Viewing X as a subset of the larger Anscombe-Aumann domain featuring both ambiguity and monetary risk, the common state-contingent utility vu can be viewed as being pinned down by the expected utility risk preference of a subject on this larger domain. Thus studying which preference in $\mathcal{M}$ best fits a given set of empirical data may be viewed as studying the subject's ambiguity attitude, given their prescribed risk preference.

expected utility, maxmin expected utility, variational preference, and dual self expected utility models. Define each ϕ_α component-wise via:

$$\phi_\alpha(x)_s = v^{-1}\big(v(x_s) + \alpha\big).$$

Then every utility of the form (5) is ϕ-additive, and by Theorem 1 every ϕ-additive representation of preferences in $\mathcal{M}$ is of this form.

This approach extends naturally to the full Anscombe-Aumann domain featuring both risk and ambiguity. Suppose, for example, that X consists of all acts mapping $S \to \Delta\big([0,1]\big) \setminus \{\delta_1\}$, and that preferences in $\mathcal{M}$ admit a representation satisfying (5), where v is an expected utility representation over lotteries, and w is positive homogeneous. Then $\phi_\alpha(x)_s = e^{-\alpha}x_s + (1-e^{-\alpha})\delta_1$ defines a virtual numeraire.

3.4 Regularity Conditions

We conclude by briefly stating some basic regularity conditions needed for future results. All these conditions are satisfied in every example in Section 3.3.

A virtual commodity $\{\phi_\alpha\}_{\alpha\geq 0}$ is said to be **regular** if is injective in x and α.[25] An alternative x' is **obtainable** from x, denoted $x \trianglelefteq x'$ if there exists $\alpha \geq 0$ such that $x' = \phi_\alpha(x)$. Define $x \sim_{\trianglelefteq} x'$ if either x is obtainable from x' or vice-versa. The following are mild topological conditions on X and $\{\phi_\alpha\}_{\alpha\geq 0}$ which will be used in subsequent results.

(A.1) **Cross Section**: There exists a continuous map $s : X \to X$, such that $x \sim_{\trianglelefteq} x'$ implies $s(x) = s(x')$, and $x \sim_{\trianglelefteq} s(x)$ for all $x \in X$.

(A.2) **No Accumulation**: 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.

[25] That is, if each transformation ϕ_α is an injective map $X \to X$, and the map $\alpha \mapsto \phi_\alpha(x)$ is an injective map from $\mathbb{R}_+ \to X$ for each $x \in X$.

Condition (A.1) is a weak technical requirement. Roughly speaking, it ensures preferences' indifference sets are not substantially 'less connected' than X.[26] Condition (A.2) says that no alternative may be regarded as the limit of adding an infinite amount of virtual commodity to any other.

4 Least Squares Theory

In this section, we develop a least squares theory for evaluating the predictive accuracy of general models of preference. Toward this end, our main theorem establishes a decomposition result. We show that every data set, arising from any experiment, can be *uniquely* decomposed into two orthogonal components: a rationalizable term, and a sum of cardinal revealed preference cycles. Thus, unlike in the case of classical revealed preference data, we obtain a canonical method of separating the data into consistent and inconsistent components.

4.1 Preliminaries

Fix a choice of consumption space X and regular virtual commodity $\{\phi_\alpha\}_{\alpha\geq 0}$. An **experiment** is a finite collection $\mathcal{E}$ of pairs of elements of X. To avoid trivialities, we assume that, for any $\{x, x'\} \in \mathcal{E}$, it is not the case that $\phi_\alpha(x) = x'$ for some $\alpha \geq 0$ or vice-versa.[27] A **data set** consists of a compensation difference measurement for each pair $\{x, x'\} \in \mathcal{E}$ (recall the compensation difference is equal to the unique α^* that solves $x \sim \phi_{\alpha^*}(x')$ if $x \succsim x'$, or $x' \sim \phi_{\alpha^*}(x)$ if $x' \succsim x$).[28] We identify an experiment $\mathcal{E}$ with the undirected graph whose vertex set $\mathcal{V} \equiv \{x_1, \ldots, x_K\}$ consists of those alternatives appearing in some pair in $\mathcal{E}$, and whose edge set is $\mathcal{E}$.

[26]It rules out, for example, pathological fringe cases where every indifference set of every continuous preference satisfying (N.1) - (N.3) is totally disconnected, even while X itself is connected.

[27]This serves only to rule out questions such as 'how many units of ϕ are needed to achieve indifference between x and x plus ten units of ϕ?'

[28]In Online Appendix E, we provide an incentive compatible mechanism for eliciting this data for general choice of X and $\{\phi_\alpha\}_{\alpha\geq 0}$.

Given a graph $(\mathcal{V}, \mathcal{E})$, let $\vec{\mathcal{E}}$ denote the set of oriented edges, $\vec{\mathcal{E}} = \{(x_i, x_j) \in X \times X : \{x_i, x_j\} \in \mathcal{E}\}$. A **flow** is a function $F : \vec{\mathcal{E}} \to \mathbb{R}$ such that $F_{x_i x_j} = -F_{x_j x_i}$ for all $(x_i, x_j) \in \vec{\mathcal{E}}$.[29] Let $\mathcal{F}$ denote the space of all flows on $(\mathcal{V}, \mathcal{E})$; it is a vector space under coordinate-wise addition and scalar multiplication.[30] It is precisely the collection of possible compensation differences data sets that could arise from the experiment $(\mathcal{V}, \mathcal{E})$.

Gradient Flows

Let $\mathcal{U} = \mathbb{R}^{\mathcal{V}}$ denote the space of utility functions on $\mathcal{V}$. By minor abuse of notation we will write i for x_i, $\bar{u}_i$ for $\bar{u}(x_i)$ and so forth. For any utility vector $\bar{u} \in \mathcal{U}$, its **gradient** is the flow whose value on an oriented edge is given by the signed difference of the utility values at its endpoints:

$$(\operatorname{grad} \bar{u})_{ij} = \bar{u}_j - \bar{u}_i,$$

for all $(i, j) \in \vec{\mathcal{E}}$, $i < j$. This defines a linear map $\operatorname{grad} : \mathcal{U} \to \mathcal{F}$. We say a flow is a gradient flow if it is the gradient of some utility vector in $\mathcal{U}$.

4.2 Rationalizability & Inconsistency

Let $\{\phi_\alpha\}_{\alpha \geq 0}$ be a regular virtual commodity, and $Y \in \mathcal{F}$ a ϕ-compensation differences data set for some experiment $(\mathcal{V}, \mathcal{E})$. The data Y are **ϕ-rationalizable** if there exists a continuous preference $\succsim$ for which $\{\phi_\alpha\}_{\alpha \geq 0}$ is a virtual numeraire, and which satisfies:

$$Y_{ij} \geq 0 \quad \iff \quad \phi_{Y_{ij}}(x_i) \sim x_j \tag{6}$$

for all $\{i, j\} \in \mathcal{E}$ (recall Y_{ij} denotes the observed compensation difference from x_i to x_j). If such a preference exists, we say $\succsim$ rationalizes Y. By Theorem 1, this is equivalent to the existence of a ϕ-additive representation u such that:

$$Y_{ij} = u(x_j) - u(x_i). \tag{7}$$

[29] For a pair $\{x_i, x_j\} \in \mathcal{E}$, $F_{x_i x_j}$ denotes the flow from x_i to x_j.

[30] Formally, we endow it with basis $\{\mathbb{1}_{(i,j)}\}_{\{(i,j) \in \vec{\mathcal{E}} : i < j\}}$, where $\mathbb{1}_{(i,j)}$ denotes the flow equal to one on (i, j) (and hence minus one on (j, i)) and zero along every other oriented edge.

Equation (7) implies that if Y is ϕ-rationalizable, it is necessarily a gradient flow, as $Y = \operatorname{grad} u\big|_{\mathcal{V}}$, and hence it must satisfy the adding-up conditions (4). Conversely, (4) are sufficient for the existence of some $\bar{u} \in \mathcal{U}$ such that $Y = \operatorname{grad} \bar{u}$. However, it is unclear whether this suffices for *rationalizability*, as a priori it is not obvious which vectors $\bar{u} \in \mathcal{U}$ are the restrictions of continuous, ϕ-additive representations. Our next theorem is the main structural result of this paper. It establishes that no matter the choice of $\{\phi_\alpha\}_{\alpha \geq 0}$ or $(\mathcal{V}, \mathcal{E})$, *every* vector $\bar{u} \in \mathcal{U}$ is the restriction of some continuous and ϕ-additive utility. Thus the adding-up conditions (4) are not only necessary, but also sufficient for ϕ-rationalizability, no matter how complex the structure of the environment, experiment, or virtual commodity may be.

Theorem 2. *Let $\{\phi_\alpha\}_{\alpha \geq 0}$ be a regular virtual commodity which satisfies (A.1) and (A.2). Then for every experiment $(\mathcal{V}, \mathcal{E})$, for any dataset $Y \in \mathcal{F}$, the following are equivalent:*

(i) For every collection $(i_0, i_1), (i_1, i_2), \ldots, (i_{L-1}, i_0) \in \vec{\mathcal{E}}$,

$$\sum_{l=0}^{L-1} Y_{i_l i_{l+1}} = 0,$$

where $i_L \equiv i_0$.

(ii) The data Y form a gradient flow.

(iii) The data are ϕ-rationalizable by a continuous preference.

Theorem 2 tells us that the rationalizable data sets always form a linear subspace of $\mathcal{F}$. One consequence of this is that every data vector may be *uniquely* written as a sum of a ϕ-rationalizable term, and a uniquely determined 'residual' component that is orthogonal to every ϕ-rationalizable vector. This latter term reflects the inconsistency in the observed data: it is zero if and only if the data are ϕ-rationalizable. This residual inconsistency vector admits a particularly natural economic interpretation.

Call a flow $C \in \mathcal{F}$ a **perfect cycle** if $C_{i_0 i_1} = C_{i_1, i_2} = \ldots = C_{i_{L-1} i_0} = c$ on some collection of oriented edges $(i_0, i_1), (i_1, i_2), \ldots, (i_{L-1}, i_0) \in \vec{\mathcal{E}}$, and C is equal to zero everywhere else. Perfectly cyclic flows are cardinal analogues of revealed preference cycles. If a perfect cycle with $c > 0$ was observed as data, it would mean:

$$\phi_c(x_{i_l}) \sim x_{i_{l+1}},$$

for all $l = 1, \ldots, L-1$, where $x_{i_L} \equiv x_{i_0}$. This implies that for every l, the agent would be willing to exchange x_{i_l}, plus up to c units of numeraire, for $x_{i_{l+1}}$, and hence could be exploited as a 'numeraire pump' by a savvy arbitrageur.[31]

It is a standard result in graph theory that a flow is orthogonal to every gradient vector if and only if is a sum of perfect cycles.[32] We summarize this below.

Proposition 1. *For any flow $R \in \mathcal{F}$, the following are equivalent:*

(i) R is a sum of perfect cycles.

(ii) R is orthogonal to every gradient flow, i.e.:

$$\sum_{\{(i,j)\in\vec{\mathcal{E}}\,:\,i<j\}} \left(\operatorname{grad}\bar{u}\right)_{ij} R_{ij} = 0 \tag{8}$$

for all $\bar{u} \in \mathcal{U}$.

Thus every data set can be uniquely decomposed into two orthogonal terms: a rationalizable component, and a sum of cardinal revealed preference cycles. This highlights a notable advantage of our approach: for ordinal choice data, it is difficult to separate cyclic inconsistency from rationalizable material (e.g. Bernheim and Rangel 2009; Nishimura 2018). On the other hand, Theorem 2 and Proposition 1 show that for *any* type of compensation differences data, revealed preference cycles and rationalizable content occupy complementary, orthogonal subspaces of $\mathcal{F}$.

This suggests a natural notion of 'best fitting' rationalizable preferences given data: a preference belongs to the best-fit set if its ϕ-additive utility

[31] By (N.2) this also implies $x_{i_l} \prec x_{i_{l+1}}$ for all l, yielding an ordinal revealed preference cycle.

[32] This follows, e.g., by combining Theorem 1 in Jiang et al. (2011) and Corollary 14.2.3 of Godsil and Royle (2001).

differences are exactly equal to the *rationalizable component* of the data vector. Our regression theory operationalizes this idea. By minimizing a least-squares objective, we orthogonally project our data onto the rationalizable subspace, guaranteeing we purify away only cyclic inconsistency, and leave behind the rationalizable content.

Rationalizability for General Models

Often, we wish to test whether a rationalizing preference can be found which, in addition, belongs to a particular model $\mathcal{M}$. When such a rationalizing preference exists, we will say the data Y are **$\mathcal{M}$-rationalizable**. Theorem 2 characterizes $\mathcal{M}$-rationalizability of a number of models, such as the classes of quasilinear or homothetic preferences, stationary preferences over dated rewards, or general CARA preferences (see Section 3.3). In each of these instances, the set $\mathcal{M}$ coincides with the set of *all* preferences for which the relevant ϕ is a virtual numeraire. However, many models of interest are not fully characterized by (N.1) - (N.3) alone.

To test for $\mathcal{M}$-rationalizability generally, Theorem 2 suggests finding the subset $\mathcal{K}_{\mathcal{M}} \subseteq \mathcal{U}$ of vectors that are the restrictions to $\mathcal{V}$ of the ϕ-additive representations of $\mathcal{M}$. These sets reflect the additional 'shape' restrictions that characterize which ϕ-additive utilities represent preferences in $\mathcal{M}$. These sets $\mathcal{K}_{\mathcal{M}}$ are frequently straightforward to compute, and often possess a highly tractable structure.

Example 1. Suppose $X = \mathbb{R}^L_{++}$, $\phi_\alpha(x) = e^\alpha x$ and $(\mathcal{V}, \mathcal{E})$ is arbitrary. Let $\mathcal{M}$ denote the collection of all Cobb-Douglas utilities on X. The ϕ-additive representations of Cobb-Douglas preferences are all of the form:

$$u(x) = \sum_{l=1}^{L} \kappa_l \ln x_l + c.$$

Because the utility differences of any such u are independent of c, we may normalize c to zero without loss. Thus a data vector Y is Cobb-Douglas rationalizable if and only if $Y = \operatorname{grad} \bar{u}$ for some vector $\bar{u}$ for which there exists $\kappa \in \mathbb{R}^L$ such that:

$$\begin{aligned} \bar{u}_i &= \big\langle \kappa, \ell(x_i) \big\rangle \quad \forall i = 1, \dots, K \\ \langle \kappa, \mathbb{1} \rangle &= 1 \\ \kappa &\geq 0, \end{aligned} \tag{9}$$

where $\ell : \mathbb{R}^L_{++} \to \mathbb{R}^L$ denotes the component-wise natural logarithm. Note that $\mathcal{K}_\mathcal{M}$, i.e. the set of such vectors $\bar{u}$, is defined by finitely many linear inequalities. ■

Example 2. Suppose now $X = \mathbb{R}^L_+$ and $\phi_\alpha(x) = (x_1 + \alpha, x_2, \dots, x_L)$. Let $\mathcal{M}$ denote the collection of preferences which admit an increasing, quasilinear, and quasiconcave utility.[33] It is straightforward to show $\bar{u} \in \mathcal{K}_\mathcal{M}$ if and only if there exist vectors $\pi_1, \dots, \pi_K \in \mathbb{R}^L$ and scalars $\gamma_1, \dots, \gamma_L \in \mathbb{R}$ such that:

$$\begin{aligned} &\bar{u}_i = \langle \pi_i, x_i \rangle + \gamma_i && \forall i = 1, \dots, K \\ &\langle \pi_i, x_i \rangle + \gamma_i \leq \langle \pi_j, x_i \rangle + \gamma_j && \forall i, j = 1, \dots, K \\ &\pi_{i,1} = 1 && \forall i = 1, \dots, K \\ &\pi_i \geq 0 && \forall i = 1, \dots, K, \end{aligned} \tag{10}$$

where $\pi_{i,1}$ denotes the first component of the vector π_i.[34] In particular, even though the preferences of $\mathcal{M}$ cannot be described by any finite vector of parameters, $\mathcal{K}_\mathcal{M}$ is still defined by a finite system of linear inequalities. ■

4.3 Least Squares Theory

Let us now fix a model $\mathcal{M}$. Going forward, we will assume that (i) $\mathcal{K}_\mathcal{M}$ is convex, and (ii) $\mathcal{K}_\mathcal{M} + \ker(\mathrm{grad})$ is closed.[35] These conditions ensure that $\mathrm{grad}(\mathcal{K}_\mathcal{M}) \subseteq \mathcal{F}$ is closed and convex.[36] Both conditions are automatically

[33] By quasilinear, we mean with respect to the first commodity.

[34] See Online Appendix G for a formal proof.

[35] The kernel of the gradient, ker(grad), consists of vectors that are constant on the vertex sets of each connected component of $(\mathcal{V}, \mathcal{E})$. In particular, when $(\mathcal{V}, \mathcal{E})$ is connected, it consists only of vectors of the form $(c, \dots, c)$, $c \in \mathbb{R}$.

[36] See, e.g., Holmes (2012), Lemma 17.H.

satisfied if $\mathcal{K}_{\mathcal{M}}$ is defined by a finite set of linear inequalities. This stronger requirement is satisfied by both Example 1 and Example 2, and turns out to hold quite broadly; see Online Appendix G.

This enables us to evaluate the predictive accuracy of any such model by solving a simple, constrained least squares problem:

$$\min_{\bar{u} \in \mathcal{K}_{\mathcal{M}}} \left\| \operatorname{grad} \bar{u} - Y \right\|_2^2. \tag{11}$$

Geometrically, solving (11) amounts to projecting the data vector Y onto the subset of $\mathcal{M}$-rationalizable flows, $\operatorname{grad}(\mathcal{K}_{\mathcal{M}})$. Since this set is closed and convex by hypothesis, (11) admits a unique minimizer, which we denote by $Y^*_{\mathcal{M}}$. We term this flow the **best-fit estimator** for the model $\mathcal{M}$. It captures the preference(s) in $\mathcal{M}$ which minimize the mean squared error between the observed compensation differences Y, and those predicted by the model. For parametric models such as in Example 1, $Y^*_{\mathcal{M}}$ will generally identify a unique preference. In such instances, (11) provides a natural method for estimating model parameters from the data Y. When models are non-parametric, as in Example 2, the estimator $Y^*_{\mathcal{M}}$ will select for a set of preferences; see also Section 5.2.

4.3.1 Economic Content of Mean Squared Error Minimization

Consider a subject with preference $\succsim \in \mathcal{M}$. Any such preference yields a tuple of willingness-to-pay measurements $\operatorname{grad} \bar{u} \in \operatorname{grad} \mathcal{K}_{\mathcal{M}}$, where $\bar{u} = u\big|_{\mathcal{V}}$ for some ϕ-additive representation u. However, suppose instead that some vector of compensation differences Y is observed. If, for some (i, j), we see $\bar{u}_j - \bar{u}_i > Y_{ij}$, the subject was willing to exchange alternative j for i plus Y_{ij} units of numeraire. Were such a trade realized, the subject would suffer a loss in utility terms equal to $L_{ij} = \big|(\bar{u}_j - \bar{u}_i) - Y_{ij}\big| > 0$ under any ϕ-additive utility. Conversely, if we observed $Y_{ij} > \bar{u}_j - \bar{u}_i$, this would mean the subject would *reject* any offer to swap for compensation equal to $Y_{ij} - \varepsilon$, for any $\varepsilon > 0$. In this case, the agent would suffer an opportunity cost of up to L_{ij} utils in missed gains from trade, should they adhere to their reported trade strategy.

This means that at any solution to (11), the value of the least squares program may be regarded as the minimal, *quadratically-weighted* (ϕ-additive) utility loss that would be suffered by any $\mathcal{M}$-consistent agent, should they commit to the observed trades Y.[37] Economically, the quadratic weighting reflects a desire to penalize more strongly those exchanges which yield larger losses in the fitting exercise. However, as our next result shows, as a loss function it also provides a means of distinguishing between different *sources* of error.

Let Y^*_{Add} denote the projection of Y onto the subspace of gradient flows.[38] Our next result says that for any $Y \in \mathcal{F}$, solving (11) is equivalent to a two-stage least squares procedure where Y is first projected onto the gradient flows, then the solution from this fitting exercise, Y^*_{Add}, is regressed on $\mathcal{M}$.

Proposition 2. *For any experiment $(\mathcal{V}, \mathcal{E})$, for every $Y \in \mathcal{F}$,*

$$Y^*_{\mathcal{M}} = \arg\min_{\bar{u} \in \mathcal{K}_{\mathcal{M}}} \left\| \mathrm{grad}\, \bar{u} - Y^*_{Add} \right\|_2^2.$$

In particular,

$$\left\| Y - Y^*_{\mathcal{M}} \right\|_2^2 = \left\| Y - Y^*_{Add} \right\|_2^2 + \left\| Y^*_{Add} - Y^*_{\mathcal{M}} \right\|_2^2. \tag{12}$$

Equation (12) says that the mean squared error associated with (11) is simply the sum of the error attributable to the data not being a vector of ϕ-additive utility differences, plus error stemming purely from the best-fit utility differences not satisfying the shape constraints of the model. This allows us to

[37] One could alternatively specify (11) as an L^1 minimization problem. In this case, the value of the (linear) program would have a direct, *unweighted* interpretation. This may be useful in applications where the unique best-fit estimates provided by the (strictly convex) 2-norm are less important. This is analogous to classical regression theory where least absolute deviations and least squares theories have their own context-dependent advantages.

[38] This may be computed as the gradient of any minimizer to (11), when the constraint set is taken to be all of $\mathcal{U}$. This solution is given in closed form by:

$$Y^*_{Add} = \mathrm{grad}\,[\mathrm{grad}^{\intercal}\mathrm{grad}]^{\dagger}\mathrm{grad}^{\intercal}\, Y,$$

where $\dagger$ denotes the the Moore-Penrose pseudoinverse of a matrix.

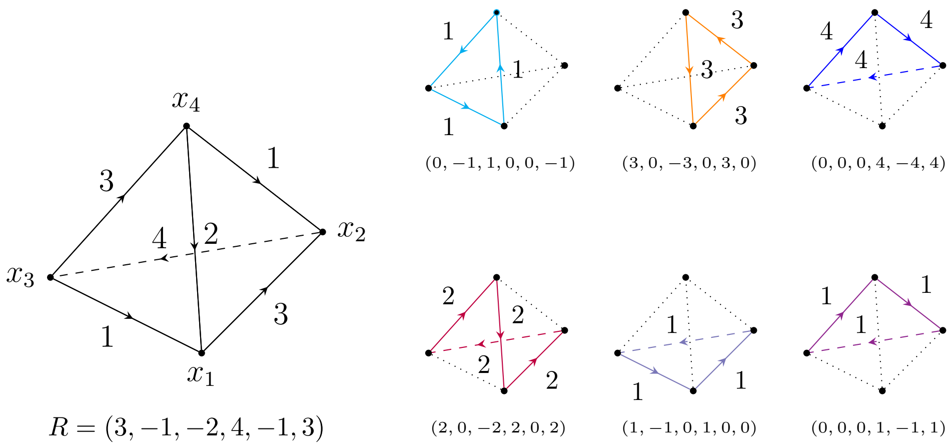

Figure 1: A residual flow $R = (R_{12}, R_{13}, R_{14}, R_{23}, R_{24}, R_{34})$ satisfying (8), along with two decompositions into sums of perfect cycles. The lower bound of $\|R\|_1 = 14$ is attained by the sum of the money pump values of the bottom (though not the top) decomposition.

distinguish between error arising from our identifying assumption that our chosen $\{\phi_\alpha\}_{\alpha \geq 0}$ is indeed a virtual numeraire to the subject, versus model-specific considerations.

4.3.2 An Alternative First-Stage Criterion: The Money Pump

The first stage residual, $Y - Y^*_{Add}$, admits a representation as a sum of cyclic flows, capturing numeraire-valued arbitrage opportunities against the subject. This suggests an analogue of the money pump index of Echenique et al. (2011) as an alternative means of quantifying this deviation.[39] Define the **money pump** value of a perfect cycle C, where $C_{i_0 i_1} = C_{i_1, i_2} = \ldots = C_{i_{L-1} i_0} = c$ on $(i_0, i_1), (i_1, i_2), \ldots, (i_{L-1}, i_0) \in \vec{\mathcal{E}}$ and is zero elsewhere, via $MP(C) = c\,L$. This corresponds to the amount of *numeraire* an arbitrageur could extract from the subject via a cyclic sequence of trades.

It is natural to seek to extend MP from pure cycles to general residuals R linearly, by decomposing R as a sum of pure cycles then summing the

[39]The money pump index was first studied by Echenique et al. (2011), in the context of price-consumption data. Roughly speaking, it reflects the amount of money one could extract from a consumer who violates the generalized axiom of revealed preference.

associated money pump values. However, such decompositions are non-unique; moreover, the sum of the money pump values of different decompositions of the same residual will generally differ, see Figure 1. Instead, we consider the most conservative extension.

Let $\mathfrak{C} \subsetneq \mathcal{F}$ denote the set of pure cycles. For any R satisfying (8), let $\mathfrak{D}(R)$ denote the collection of all finite decompositions of R into pure cycles.[40] We extend $MP : \mathfrak{C} \to \mathbb{R}$ to a function MP^* over all flows satisfying (8) 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 subject as possible, by taking an infimum across all finite decompositions. In spite of its definition as a value function, our next result asserts that MP^* is in fact simply the L^1 norm. This simplicity in our setting is notable, given the computational difficulty of calculating the money pump index for price-consumption data (e.g. Smeulders et al. 2013).

Proposition 3. *For all $R \in \mathcal{F}$ satisfying (8), 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.

While our focus is on L^2 loss, Proposition 3 provides an economically compelling alternative, particularly for models where $\mathcal{K}_{\mathcal{M}} = \mathcal{U}$. It also illustrates the economic naturality of our numeraire-based approach, and provides another example of familiar economic ideas taking on a far simpler, more tractable form in our setting.

5 Application: MEU Preferences

Let $X = \mathbb{R}^S$ denote the pure ambiguity domain of monetary acts for a finite state space S. We will be interested in the model $\mathcal{M}_{MEU}$ consisting of the

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

maxmin expected utility preferences featuring a fixed, known state-contingent utility of consumption $v : \mathbb{R} \to \mathbb{R}$.[41] Such preferences are represented by a utility of the form:

$$u(x) = \min_{\pi \in C} \mathbb{E}_\pi\big(v(x)\big) = \min_{\pi \in C} \sum_{s \in S} \pi_s v(x_s), \tag{13}$$

where $C \subseteq \Delta(S)$ is some compact, convex set of priors over states of the world.[42] The function $\min_{\pi \in C} \mathbb{E}_\pi(\cdot)$ is referred to as the **utility functional**. We assume that v is (i) continuous, (ii) strictly increasing, (iii) unbounded above, and (iv) normalized so that $v(0) = 0$ and $v(1) = 1$. For any act $x \in \mathbb{R}^S$ and $\alpha \geq 0$, we define the virtual commodity $\phi_\alpha(x)_s = v^{-1}\big(v(x_s) + \alpha\big)$ component-wise. Any utility of the form (13) is ϕ-additive. In what follows, see Online Appendix F for omitted derivations.

5.1 Testing Ambiguity Aversion in a Two-State World

Suppose first that there are only two states of the world. Let $x^0 = (0, 0)$ denote the zero act, and $x^1 = (1, 0)$ and $x^2 = (0, 1)$ the Arrow securities for states 1 and 2; we will consider the experiment $\mathcal{E} = \big\{\{0, x^2\}, \{x^1, x^2\}\big\}$.

Figure 2 plots $\text{grad}(\mathcal{K}_{MEU})$, the set of MEU-rationalizable flows for this experiment. The corners of the rationalizable triangle correspond to the preferences in $\mathcal{M}_{MEU}$ whose sets of priors are $\{\delta_1\}$, $\{\delta_2\}$, and $\Delta(S)$ respectively.[43] Given data Y, the best-fit estimator Y^*_{MEU} is obtained by projecting Y onto this triangle.

The faces of the rationalizable triangle have natural axiomatic interpretation. The preferences corresponding to flows along the top edge of $\text{grad}(\mathcal{K}_{MEU})$ are characterized by the property that $x^0 \sim x^1$. This inequality constraint represents the axiomatic requirement that preferences in $\mathcal{M}_{MEU}$ be non-decreasing

[41] Practically speaking, v could be chosen either on the basis of theoretical considerations or first-stage estimation of the subject's risk preference.

[42] Note that there is a one-to-one correspondence between closed, convex sets of priors C and preference in $\mathcal{M}_{MEU}$. We will freely identify such preferences with their sets of priors.

[43] Recall δ_s denotes the probability measuring assigning a mass of one to $\{s\}$.

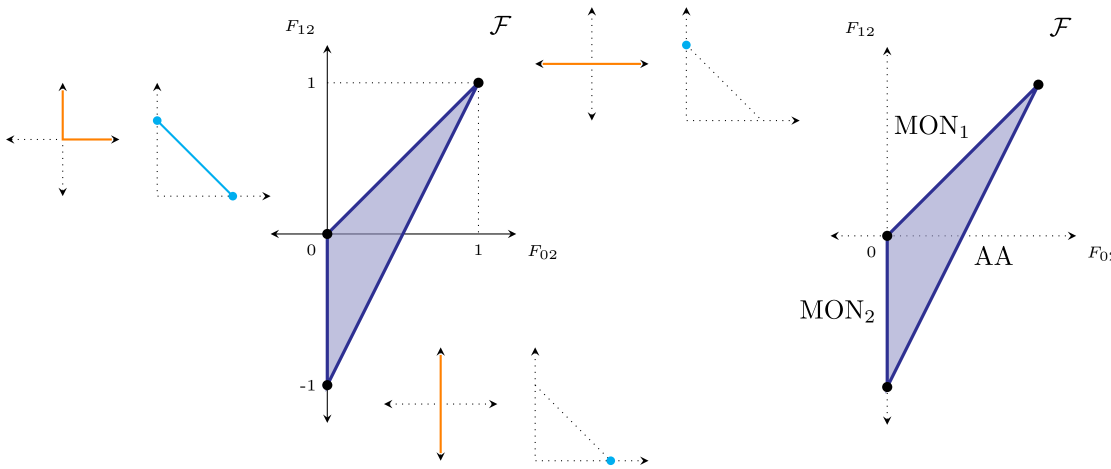

(a) The extremal MEU preferences. The simple nature of $\mathcal{E}$ implies every flow is a gradient.

(b) Model axioms correspond to the inequalities defining the rationalizable set.

Figure 2: The MEU-rationationalizable flows (violet triangle) arising from the experiment $\mathcal{E} = \big\{\{0, x^2\}, \{x^1, x^2\}\big\}$. For each vertex of the triangle, the level set through the origin of the rationalizing MEU functional (orange) and corresponding set of priors (cyan) are shown. Each face of the triangle corresponds to a particular axiomatic constraint of the model: the top and left faces to monotonicity of consumption in state one (resp. two), and the bottom-right to ambiguity aversion.

in state-one consumption. Analogously, the left face is characterized by the property that $x^0 \sim x^2$ and hence reflects monotonicity of consumption in state two. The flows along the bottom-right face correspond to the subjective expected utility preferences. These are precisely the ambiguity-neutral preferences in $\mathcal{M}_{MEU}$, hence this inequality captures the ambiguity aversion axiom of Gilboa and Schmeidler (1989).[44]

Suppose Y^*_{MEU} lies on the relative interior of the lower-right face of ratio-

[44] Any ϕ-additive utility may be written as $u(x) = w\big(v(x_1), \ldots, v(x_S)\big)$, where w is a translation-invariant utility functional (see Section 3.3). This representation is of the form (13) if and only if w is additionally increasing, concave, and positively homogeneous; see, e.g., Ok (2011) H.1.3 Lemma 2. Every data set arising from $\mathcal{E}$ is rationalizable by a ϕ-additive utility where w is additionally positively homogeneous. This leaves monotonicity and the concavity of w (corresponding to ambiguity aversion) as the only falsifiable implications of the model for $\mathcal{E}$.

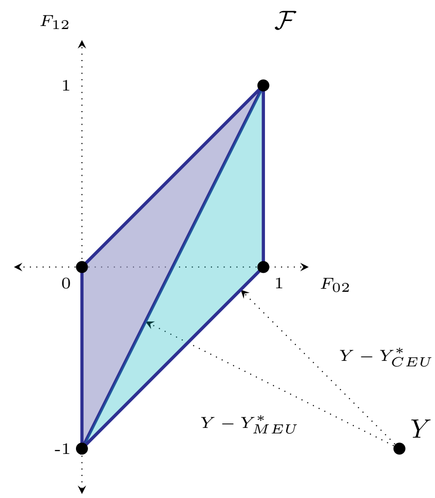

Figure 3: The set of MEU-rationalizable (violet) and CEU-rationalizable (violet or aquamarine) vectors for $\mathcal{E}$. Letting Π_{MEU} and Π_{CEU} denote the respective projections onto these sets, the quantity $\left\|Y - \Pi_{\text{MEU}}Y\right\|_2^2 - \left\|Y - \Pi_{\text{CEU}}Y\right\|_2^2$ reflects the *shadow price*, in mean squared error terms, of imposing ambiguity aversion, conditional upon requiring monotonicity, translation invariance, and homotheticity.

nalizable triangle. Then ambiguity aversion is the sole axiomatic constraint binding at the solution to (11). To quantify the gain in fit from relaxing ambiguity aversion, consider the model obtained by dropping only ambiguity aversion from the MEU axioms. This corresponds to the class of invariant biseparable preferences (e.g. Chandrasekher et al. 2022). However, in our simplified setting, the testable implications of this model coincide with the conceptually simpler Choquet expected utility (CEU) theory of Schmeidler (1989). The CEU-rationalizable set is plotted in Figure 3. By comparing the difference in value of (11) obtained under these two nested constraint sets, one obtains a measure of the shadow price, in model fit terms, of imposing specifically the ambiguity aversion axiom.

5.2 Identification with $|S| > 2$

Suppose now $|S| > 2$. 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 (11) for $\mathcal{M}_{MEU}$ amounts to solving the following constrained least squares problem:

$$\begin{aligned}
\min_{\bar{u}, \pi_1, \ldots, \pi_K} \quad & \left\| \operatorname{grad} \bar{u} - Y \right\|_2^2 \\
\text{subject to} \quad & \bar{u}_i = \langle \pi_i, v_i \rangle \qquad \forall i = 1, \ldots, K \\
& \langle \pi_i, v_i \rangle \leq \langle \pi_j, v_i \rangle \quad \forall i, j = 1, \ldots, K \\
& \langle \pi_i, \mathbb{1}_S \rangle = 1 \qquad \forall i = 1, \ldots, K \\
& \pi_i \geq 0 \qquad \forall i = 1, \ldots, K,
\end{aligned} \tag{14}$$

where $\bar{u} \in \mathcal{U}$, $\pi_1 \ldots, \pi_K \in \mathbb{R}^S$, and $K = |\mathcal{V}|$. Here $v_i = \big(v(x_{i,1}), \ldots v(x_{i,S})\big)$ is the **utility-act** associated with $x_i \in \mathcal{V}$ under v. Once again, in spite of $\mathcal{M}_{MEU}$ being fully non-parametric, the constraint set in (14) is still defined by a finite set of linear inequality constraints.[45]

However, a solution to (14) no longer uniquely identifies a preference in $\mathcal{M}_{MEU}$. Suppose $(\bar{u}, \pi_1, \ldots, \pi_K)$ is a feasible solution to (14). The vector $\bar{u}$, coupled with the utility acts $v_1, \ldots, v_K$, defines a family of hyperplanes $H_{v_i, \bar{u}_i} = \{x \in \mathbb{R}^S : \langle v_i, x \rangle = \bar{u}_i\}$. Let $\bar{H}_{v_i, \bar{u}_i}$ denote the restrictions of these hyperplanes to the affine hull of $\Delta(S)$. The first and second sets of constraints in (14) imply that each $\bar{H}_{v_i, \bar{u}_i}$ supports $\operatorname{co}\{\pi_1, \ldots, \pi_K\}$ at π_i. Thus $\operatorname{co}\{\pi_1, \ldots, \pi_K\}$ is a set of priors defining a preference in $\mathcal{M}_{MEU}$ consistent with $(\bar{u}, \pi_1, \ldots, \pi_K)$. However, many other preferences in $\mathcal{M}_{MEU}$ are also consistent. Let:

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

where $\bar{H}^{+}_{v_i, \bar{u}_i}$ denotes the i-th upper half-space. The following result characterizes the identified set arising from each feasible solution to (14).

[45]In Online Appendix G we characterize the sets of shape constraints for invariant biseparable preferences, as well as SEU, CEU, and variational preferences, allowing for the axiomatic analysis of Section 5.1 to be carried out for general experiments and state spaces, as well as for model selection exercises via Proposition 2.

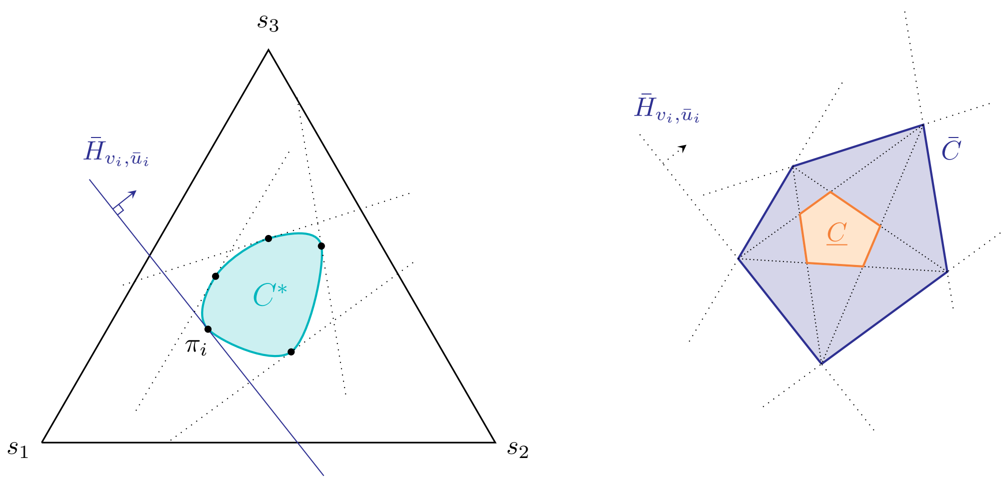

(a) The set of priors C^* associated with a preference in $\mathcal{M}_{MEU}$. The vector $(\bar{u}, \pi_1, \dots, \pi_K)$ is a solution to (14), as for each v_i, the hyperplane $\bar{H}_{v_i,\bar{u}_i}$ supports C^* at π_i.

(b) Every feasible solution to (14) defines a polytope $\bar{C} = \big(\cap_i \bar{H}^+_{v_i,\bar{u}_i}\big) \cap \Delta(S)$. A set of priors $C \subseteq \bar{C}$ defines a rationalizing preference if and only if each facet of $\bar{C}$ contains some extremal point of C.

Figure 4: An experiment with $\mathcal{V} = \{x_1, \dots, x_5\}$ and a rationalizing utility vector $\bar{u}$ define a system of hyperplanes on the simplex. From these hyperplanes, we obtain upper and lower envelope sets of priors $\bar{C}$ and $\underline{C}$. Every set of priors associated with a rationalizing preference in $\mathcal{M}_{MEU}$ is contained within $\bar{C}$ and contains $\underline{C}$.

Proposition 4. *Fix a feasible $(\bar{u}, \pi_1, \dots, \pi_K)$. A closed, convex set of priors $C \subseteq \Delta(S)$ corresponds to a preference in $\mathcal{M}_{MEU}$ consistent with this vector if and only if:*

(i) The set of priors $C \subseteq \bar{C}$, and

(ii) Each hyperplane $H_{v_i,\bar{u}_i}$ contains some extremal point of C.

In particular, the identified set depends only upon $\bar{u}$.

It follows that $\bar{C}$ is the unique largest set of priors consistent with $(\bar{u}, \pi_1, \dots, \pi_K)$.[46] This provides bounds on the priors held by an individual, even absent full iden-

[46] Equivalently, it is the set of priors of the unique, most-ambiguity averse preference in the identified set.

tification: if $Y = \operatorname{grad} \bar{u}$ for $\bar{u} \in \mathcal{K}_{MEU}$, and $\pi \notin \bar{C}$, then π is not held by *any* preferences in $\mathcal{M}_{MEU}$ that rationalize Y.

Figure 4 shows the supporting hyperplanes for the preference in $\mathcal{M}_{MEU}$ whose set of priors is C^*. The upper envelope $\bar{C}$ is simply the intersection of the associated upper half-spaces. It also depicts $\underline{C}$, the set of priors held by *every* rationalizing MEU preference.[47] Thus, while C^* may be unknown, the vector u, along with $\mathcal{V}$, allow one to derive optimal upper and lower bounds, $\underline{C} \subseteq C^* \subseteq \bar{C}$.

These bounds generate further economic predictions. For example, subjects with risk averse MEU preferences 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 sets $\bar{C}$ for two agents are disjoint yields further, testable predictions about trade behavior. Similarly, in an economy of MEU agents without aggregate uncertainty, the Pareto frontier precisely corresponds to the set of full-insurance allocations if and only if the agents share at least one common prior (Billot et al. 2000). Thus observing 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 while the individual preferences themselves may remain unidentified.

6 Statistical Tests of $\mathcal{M}$-Rationalizability

Suppose now we observe cross-sectional data $\{Y^n\}_{n=1}^N$, obtained by repeatedly sampling noisy measurements of an individual's compensation differences. Formally, for all $\{x, x'\} \in \mathcal{E}$ we assume there exists a fixed, non-stochastic 'true' compensation difference $Y^0_{xx'} = -Y^0_{x'x}$. We assume the data $\{Y^n\}_{n=1}^N$ are a random sample of N i.i.d. draws of the random flow $\tilde{Y}$, where for each $(i, j) \in \vec{\mathcal{E}}$ with $i < j$:

$$\tilde{Y}_{ij} = Y^0_{ij} + \epsilon_{ij},$$

[47] Note that generally $\underline{C}$ will not itself define an MEU preference which rationalizes u.

with (i) $\mathbb{E}(\epsilon_{ij}) = 0$, and (ii) $\text{Var}(\epsilon_{ij}) < +\infty$. We do not assume a priori that the the noise shocks $\{\epsilon_{ij}\}$ are uncorrelated or that they are distributed identically across $\mathcal{E}$.

We wish to test whether the vector of true compensation differences, Y^0, is $\mathcal{M}$-rationalizable against the full, multi-sided alternative:

$$H_0 : Y^0 \in \text{grad}(\mathcal{K}_{\mathcal{M}}), \qquad H_1 : Y^0 \notin \text{grad}(\mathcal{K}_{\mathcal{M}}). \tag{15}$$

In other words, we seek to test whether the vector of moments $\mathcal{E}(\tilde{Y})$ belongs to the closed, convex set $\text{grad}(\mathcal{K}_{\mathcal{M}})$. Problems of this form have been well-studied in the econometric literature (e.g. Chernozhukov et al. 2007; Andrews and Guggenberger 2009; Hong and Li 2018; Kitamura and Stoye 2018; Fang and Seo 2019) and off-the-shelf techniques are available for obtaining test statistics and critical values.

Let $\bar{Y} = \frac{1}{N}\sum_n Y^n$ denote the sample average flow, and let $\psi(\tilde{Y})$ denote the distance from a vector $\tilde{Y}$ to $\text{grad}(\mathcal{K}_{\mathcal{M}})$.[48] A generalization of the delta method due to Fang and Santos (2019) guarantees that, under H_0, the quantity $\sqrt{N}\psi(\bar{Y})$ converges in distribution to $\theta\big(N(0,\Sigma)\big)$, where Σ is the covariance matrix of the shock vector ϵ, and θ is a particular, non-linear function related to ψ.[49] The numerical derivative estimator of Hong and Li (2018) provides a convenient method for simulating this distribution, without requiring further analytic calculations.

1. For $b = 1, \ldots, B_N$, let $Z^{*(b)} = \sqrt{B_N}\big(\bar{Y}^{*(b)} - \bar{Y}\big)$, where $\bar{Y}^{*(b)}$ is a draw of the bootstrapped sample mean $\bar{Y}^*$, given the data $\{Y^1, \ldots, Y^N\}$.

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

$$\hat{\theta}_N\big(Z^{*(b)}\big) \equiv \frac{\psi\big(\bar{Y} + \delta_N Z^{*(b)}\big) - \psi\big(\bar{Y}\big)}{\delta_N},$$

[48] That is,

$$\psi(\tilde{Y}) = \min_{F \in \text{grad}(\mathcal{K}_{\mathcal{M}})} \big\|\tilde{Y} - F\big\|_2.$$

[49] Formally, θ is related to the Hadamard directional derivative of ψ evaluated at the true Y_0; see Fang and Santos (2019) for details.

for a choice of sequence of tuning parameters δ_N satisfying $\lim_N \delta_N = 0$, and $\lim_N \delta_N \sqrt{B_N} \to \infty$.

Letting $Z^* = \sqrt{B_N}(\bar{Y}^* - \bar{Y})$ denote the re-scaled bootstrap mean conditional upon the data, Theorem 3.1 of Hong and Li (2018) establishes the consistency of $\hat{\theta}(Z^*)$ for the asymptotic distribution $\theta\big(N(0,\Sigma)\big)$.

Theorem (Hong and Li 2018)**.** *Under the above hypotheses,*

$$\lim_{N\to\infty} \hat{\theta}_N(Z^*) \overset{\mathbb{P}}{\rightsquigarrow} \theta\big(N(0,\Sigma)\big),$$

where $\overset{\mathbb{P}}{\rightsquigarrow}$ *denotes weak convergence in probability conditional upon the data.*

An α-level test of (15) can then be constructed by comparing $\sqrt{N}\psi(\bar{Y})$ to $1-\alpha$ conditional quantile of $\hat{\theta}(Z^*)$. This can be approximated using the empirical distribution of $\{\hat{\theta}_N\big(Z^{*(b)}\big)\}_{b=1}^{B_N}$.

When $\mathcal{K}_{\mathcal{M}}$ is polyhedral, as is often the case in practice, Fang and Santos (2019) provide a consistent, alternative estimator for $\theta\big(N(0,\Sigma)\big)$. Their approach exploits the polyhedral structure of $\mathcal{K}_{\mathcal{M}}$ to directly estimate θ. Their alternative has the added benefit of providing confidence regions for the set of binding constraints at the true Y_0,. This is valuable as these constraints often have an axiomatic interpretation (e.g. Section 5.1). The interested reader is referred to Section 4.2 of Fang and Santos (2019).

7 Conclusion

This paper provides a novel approach to quantifying the predictive accuracy of various models of preference and individual decision-making. Our approach makes particular use of a common, underlying invariance property of various models to obtain *cardinal* measurements of preference intensity. The fine structure of this data forms the basis for the major recurring computational and economic advantages of our approach enjoys over classical revealed preference techniques.

A natural follow-up to the theoretical foundations established in this paper would be to take these techniques to the lab. Alternatively, while as a practical matter it is often straightforward to obtain natural choices of virtual numeraire for a given model (see Section 3.3), it would be of theoretical interest to characterize *which* sets of preferences admit a virtual numeraire. Such results may be of interest beyond decision theory and experiments. For example, in the context of mechanism design, such a result would characterize which type spaces of non-quasilinear preferences could be viewed as quasilinear under an appropriate change of coordinates.

Finally, the computational tractability of our least squares theory makes it a natural framework for considering asymptotic econometric problems. It would be interesting to establish consistency of our best-fit estimators in the presence of noise as $(\mathcal{V}, \mathcal{E})$ gets large, in the vein of Seijo and Sen (2011) or Chambers et al. (2021). Similarly it would be of interest to study the testing problem (15) as $(\mathcal{V}, \mathcal{E})$, rather than the cross-sectional dimension, grows large in an appropriate sense.

References

Afriat, S. N. (1972). Efficiency estimation of production functions. *International economic review*, 568–598.

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica: Journal of the Econometric Society*, 503–546.

Andrews, D. W. and P. Guggenberger (2009). Validity of subsampling and "plug-in asymptotic" inference for parameters defined by moment inequalities. *Econometric Theory 25*(3), 669–709.

Bernheim, B. D. and A. Rangel (2009). Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics. *The Quarterly Journal of Economics 124*(1), 51–104.

Billot, A., A. Chateauneuf, I. Gilboa, and J.-M. Tallon (2000). Sharing beliefs: between agreeing and disagreeing. *Econometrica 68*(3), 685–694.

Blundell, R., M. Browning, and I. Crawford (2008). Best nonparametric bounds on demand responses. *Econometrica 76*(6), 1227–1262.

Chambers, C. P. and F. Echenique (2016). *Revealed preference theory*, Volume 56. Cambridge University Press.

Chambers, C. P., F. Echenique, and N. S. Lambert (2021). Recovering preferences from finite data. *Econometrica 89*(4), 1633–1664.

Chandrasekher, M., M. Frick, R. Iijima, and Y. Le Yaouanq (2022). Dual-self representations of ambiguity preferences. *Econometrica 90*(3), 1029–1061.

Cherchye, L., T. Demuynck, B. De Rock, and P. Hjertstrand (2015). Revealed preference tests for weak separability: an integer programming approach. *Journal of econometrics 186*(1), 129–141.

Chernozhukov, V., H. Hong, and E. Tamer (2007). Estimation and confidence regions for parameter sets in econometric models 1. *Econometrica 75*(5), 1243–1284.

Constantinides, G. M. (1990). Habit formation: A resolution of the equity premium puzzle. *Journal of political Economy 98*(3), 519–543.

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

Deb, R., Y. Kitamura, J. K.-H. Quah, and J. Stoye (2018). Revealed price preference: theory and empirical analysis. *arXiv preprint arXiv:1801.02702*.

Deb, R., Y. Kitamura, J. K.-H. Quah, and J. Stoye (2023). Revealed price preference: Theory and stochastic testing. *Review of Economic Studies 90(2)*, 707–743.

Debreu, G. (1959). Topological methods in cardinal utility theory. *Mathematical Methods in Social Sciences*, 16–26.

Echenique, F. (2021). On the meaning of the critical cost efficiency index. *arXiv preprint arXiv:2109.06354*.

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

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The quarterly journal of economics*, 643–669.

Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption. *Econometrica 57*(4), 937–969.

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

Fang, Z. and J. Seo (2019). A projection framework for testing shape restrictions that form convex cones. *arXiv preprint arXiv:1910.07689*.

Fishburn, P. C. and A. Rubinstein (1982). Time preference. *International Economic Review*, 677–694.

Fudenberg, D. and A. Liang (2020). Machine learning for evaluating and improving theories. *ACM SIGecom Exchanges 18*(1), 4–11.

Fudenberg, D., W. Newey, P. Strack, and T. Strzalecki (2020). Testing the drift-diffusion model. *Proceedings of the National Academy of Sciences 117*(52), 33141–33148.

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

Godsil, C. and G. F. Royle (2001). *Algebraic graph theory*, Volume 207. Springer Science & Business Media.

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

Holmes, R. B. (2012). *Geometric functional analysis and its applications*, Volume 24. Springer Science & Business Media.

Hong, H. and J. Li (2018). The numerical delta method. *Journal of Econometrics 206*(2), 379–394.

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

Jiang, X., L.-H. Lim, Y. Yao, and Y. Ye (2011). Statistical ranking and combinatorial hodge theory. *Mathematical Programming 127*(1), 203–244.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–292.

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

Koopmans, T. C. (1960). Stationary ordinal utility and impatience. *Econometrica*, 287–309.

Krantz, D. H., D. R. Luce, P. Suppes, and A. Tversky (2007). *Foundations of measurement: Additive and polynomial representations*, Volume 1. Courier Corporation.

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

Mehra, R. and E. C. Prescott (1985). The equity premium: A puzzle. *Journal of monetary Economics 15*(2), 145–161.

Mu, X., L. Pomatto, P. Strack, and O. Tamuz (2021). Monotone additive statistics. *arXiv preprint arXiv:2102.00618*.

Nishimura, H. (2018). The transitive core: Inference of welfare from nontransitive preference relations. *Theoretical Economics 13*(2), 579–606.

Nishimura, H., E. A. Ok, and J. K.-H. Quah (2017). A comprehensive approach to revealed preference theory. *American Economic Review 107*(4), 1239–1263.

Ok, E. A. (2011). *Real analysis with economic applications*. Princeton University Press.

Ok, E. A. and Y. Masatlioglu (2007). A theory of (relative) discounting. *Journal of Economic Theory 137*(1), 214–245.

Ok, E. A. and G. Riella (2014). Topological closure of translation invariant preorders. *Mathematics of Operations Research 39*(3), 737–745.

Ok, E. A. and G. Riella (2021). Fully preorderable groups. *Order 38*, 127–142.

Polisson, M. and J. K.-H. Quah (2022). Rationalizability, cost-rationalizability, and afriat's efficiency index. Technical report, IFS Working Paper.

Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica 68*(5), 1281–1292.

Rigotti, L., C. Shannon, and T. Strzalecki (2008). Subjective beliefs and ex ante trade. *Econometrica 76*(5), 1167–1190.

Roth, A. E. and M. W. Malouf (1981). Disagreement in bargaining: An experimental study. *Journal of Conflict Resolution 25*(2), 329–348.

Roth, A. E. and M. W. Malouf (1982). Scale changes and shared information in bargaining: An experimental study. *Mathematical Social Sciences 3*(2), 157–177.

Safra, Z. and U. Segal (1998). Constant risk aversion. *journal of economic theory 83*(1), 19–42.

Samuelson, P. A. (1938). The empirical implications of utility analysis. *Econometrica, Journal of the Econometric Society*, 344–356.

Savage, L. J. (1954). *The foundations of statistics*. John Wiley & Sons.

Savage, L. J. (1971). Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association 66*(336), 783–801.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica: Journal of the Econometric Society*, 571–587.

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

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

Smeulders, B., L. Cherchye, F. C. Spieksma, and B. De Rock (2013). The money pump as a measure of revealed preference violations: A comment. *Journal of Political Economy 121*(6), 1248–1258.

Smeulders, B., F. C. Spieksma, L. Cherchye, and B. De Rock (2014). Goodness-of-fit measures for revealed preference tests: Complexity results and algorithms. *ACM Transactions on Economics and Computation (TEAC) 2*(1), 1–16.

Smith, C. A. (1961). Consistency in statistical inference and decision. *Journal of the Royal Statistical Society: Series B (Methodological) 23*(1), 1–25.

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

Appendices

Appendix A Proof of Theorem 1

Proof. Necessity is trivial. As such, let $\succsim \in \mathcal{M}$ be arbitrary, and suppose $\{\phi_\alpha\}_{\alpha \geq 0}$ satisfies (N.1) - (N.3) for $\succsim$. Fix an alternative $\underline{x} \in X$ and define:

$$c(x) = \begin{cases} \alpha^* & \text{if } x \succsim \underline{x} \text{ and } \phi_{\alpha^*}(\underline{x}) \sim x \\ -\alpha^* & \text{if } \underline{x} \succ x \text{ and } \phi_{\alpha^*}(x) \sim \underline{x}. \end{cases}$$

This is well-defined by (N.2) and (N.3). We first show that c is ϕ-additive. Consider $c\big(\phi_\alpha(x)\big)$ for some $x \in X$, $\alpha \geq 0$. If $x \succsim \underline{x}$, then by (N.3) there exists $\alpha_{\underline{x}x} \geq 0$ such that $x \sim \phi_{\alpha_{\underline{x}x}}(\underline{x})$. By (N.1), $\phi_\alpha(x) \sim \phi_\alpha\big(\phi_{\alpha_{\underline{x}x}}(\underline{x})\big) = \phi_{\alpha_{\underline{x}x}+\alpha}(\underline{x})$. By (N.2), both $\phi_{\alpha_{\underline{x}x}+\alpha}(\underline{x}) \succsim \underline{x}$ and $\phi_{\alpha_{\underline{x}x}}(\underline{x}) \succsim \underline{x}$, hence:

$$c\big(\phi_\alpha(x)\big) = c\big(\phi_{\alpha_{\underline{x}x}+\alpha}(\underline{x})\big) = \alpha_{\underline{x}x} + \alpha = c(x) + \alpha.$$

Suppose instead $\underline{x} \succ x$. If $|c(x)| \geq \alpha$, then $\phi_{|c(x)|-\alpha}\big(\phi_\alpha(x)\big) \sim \underline{x}$, and hence $c\big(\phi_\alpha(x)\big) = -\big(|c(x)| - \alpha\big) = c(x) + \alpha$. If instead $\alpha > |c(x)|$, then by (N.1) $\phi_\alpha(x) = \phi_{\alpha-|c(x)|}\big(\phi_{|c(x)|}(x)\big) \sim \phi_{\alpha-|c(x)|}(\underline{x})$, and thus $c\big(\phi_\alpha(x)\big) = \alpha - |c(x)| = c(x) + \alpha$. Thus for all $x \in X$, $\alpha \geq 0$, $c\big(\phi_\alpha(x)\big) = c(x) + \alpha$ and we conclude c is ϕ-additive.

We now show c represents $\succsim$. Let $x \succsim x'$. By (N.3) there exists $\alpha_{x'x} \geq 0$ such that $\phi_{\alpha_{x'x}}(x') \sim x$. By (N.2) $\alpha_{x'x} > 0$ if and only if $x \succ x'$. But since we have already shown c is ϕ-additive:

$$c(x) = c\big(\phi_{\alpha_{x'x}}(x')\big) = c(x') + \alpha_{x'x}.$$

Thus $c(x) \geq c(x')$, with strict inequality whenever $x \succ x'$, and hence c represents $\succsim$.

We now show c is continuous. As $\succsim$ is continuous and admits a utility representation c, by the Open Gap Lemma (Debreu 1964), we conclude $\succsim$ also admits a continuous utility representation $w : X \to \mathbb{R}$.[50] Suppose $x_n \to x$.

[50] See also Chapter 9 Proposition 5.1 of Ok (2011).

By (N.2) and (N.3), for some choice of $\alpha^* \geq 0$ large enough, we have $\hat{x}_n = \phi_{\alpha^*}(x_n) \succsim \underline{x}$ for all $n \in \mathbb{N}$, and $\hat{x} = \phi_{\alpha^*}(x) \succsim \underline{x}$. By continuity, $\hat{x}_n \to \hat{x}$, and $w(\hat{x}_n) \to w(\hat{x})$. But $w(\hat{x}_n) = w\big(\phi_{c(\hat{x}_n)}(\underline{x})\big)$ and $w(\hat{x}) = w\big(\phi_{c(\hat{x})}(\underline{x})\big)$. By definition of a virtual commodity, the map $(\alpha, x) \mapsto \phi_\alpha(x)$ is jointly continuous, thus so is the map $\theta : \mathbb{R}_+ \to X$ defined by $\theta(\alpha) = \phi_\alpha(\underline{x})$. This implies $w \circ \theta$ is continuous and, by (N.2), strictly increasing. It follows that $(w \circ \theta)\big(c(\hat{x}_n)\big) \to (w \circ \theta)\big(c(\hat{x})\big)$ implies that also $c(\hat{x}_n) \to c(\hat{x})$.[51] Hence $c(x_n) \to c(x)$ too, by ϕ-additivity. As $x_n \to x$ was arbitrary, we conclude that c is in fact continuous.

Finally, it is immediate that for any pair $x \succsim x'$, any two ϕ-additive representations c and c' of $\succsim$ must satisfy:

$$c(x) - c(x') = \alpha_{x'x} = c'(x) - c'(x'),$$

where $\alpha_{x'x} \geq 0$ is the unique non-negative scalar such that $\phi_{\alpha_{x'x}}(x') \sim x$ guaranteed by (N.2) and (N.3). Thus c and c' differ by at most an additive constant. $\square$

Appendix B Proof of Theorem 2

B.1 Preliminaries

Let X be a metric space, and $\{\phi_\alpha\}_{\alpha \geq 0}$ a regular virtual commodity. A **homeomorphism** between topological spaces is a continuous bijection with continuous inverse. A map is an **embedding** if it is a homeomorphism onto its image. Let $H : \mathbb{R}_+ \times Y \to X$ for some metric space Y. We say that H is **equivariant** if:

$$H(\alpha + \beta, y) = \phi_\beta\big(H(\alpha, y)\big)$$

[51] A continuous, strictly increasing function $\mathbb{R}_+ \to \mathbb{R}$ is a homeomorphism onto its image (by invariance of domain; e.g. Munkres 1974) and hence admits a continuous left-inverse.

for all $y \in Y$ and $\alpha, \beta \geq 0$. Note that if $u : X \to \mathbb{R}$ is ϕ-additive, then equivariance of H implies:

$$\begin{aligned} u\big(H(\alpha, y)\big) &= u\big(\phi_\alpha\big(H(0, y)\big)\big) \\ &= u\big(H(0, y)\big) + \alpha \\ &= v(y) + \alpha \end{aligned}$$

where $v(y) \equiv u\big(H(0, y)\big)$. Thus any equivariant homeomorphism H renders every ϕ-additive utility quasilinear.

To state our next result, we first require the following lemma. It says that the binary relation $\sim_\trianglelefteq$, defined via $x \sim_\trianglelefteq y$ if and only if either $x \trianglelefteq y$ or $y \trianglelefteq x$, is an equivalence relation.[52]

Lemma 1. *Let $\{\phi_\alpha\}_{\alpha \geq 0}$ be a regular virtual commodity. Then $\sim_\trianglelefteq$ is an equivalence relation.*

In light of Lemma 1, there is a well-defined quotient space $X/\sim_\trianglelefteq$. We let $q : X \to X/\sim_\trianglelefteq$ denote the associated quotient map, and in all that follows, we will consider $X/\sim_\trianglelefteq$ endowed with its quotient topology; see Munkres (1974) for definitions.

The following result is the central piece of technical machinery needed for the proof of Theorem 2. A proof of this result may be found in Online Appendix D.

Theorem 3 (Embedding Theorem)**.** *Let $\{\phi_\alpha\}_{\alpha \geq 0}$ be a regular virtual commodity. Then X and $\{\phi_\alpha\}_{\alpha \geq 0}$ satisfy (A.1) and (A.2) if and only if there exists an equivariant embedding $H : \mathbb{R}_+ \times X/\sim_\trianglelefteq \to X$ such that*

$$q \circ H\big(\alpha, [x]\big) = [x]$$

for all $[x] \in X/\sim_\trianglelefteq$. Moreover the range of H is closed in X.

[52] Recall, $x \trianglelefteq y$ means that there exists some $\alpha \geq 0$ such that $\phi_\alpha(x) = y$.

B.1.1 Proof of Theorem 2

Proof. It is immediate that (ii) implies (i), and by Theorem 1, (iii) $\implies$ (ii). Thus we first will show that (i) $\implies$ (ii). Without loss of generality, suppose $(\mathcal{V}, \mathcal{E})$ is connected, and let $(\mathcal{V}, \mathcal{E}')$ be a spanning tree.[53] Then for all $x_k \in \mathcal{V}$ there exists a unique sequence $\{x_{j_1}, x_{j_2}\}, \{x_{j_2}, x_{j_3}\}, \ldots, x_{j_{M_k-1}}, x_{j_{M_k}}\} \in \mathcal{E}'$, where $j_1 = 1$ and $j_{M_k} = k$. Define:

$$\bar{u}_k = \sum_{m=1}^{M_k-1} Y_{j_k j_{k+1}}$$

It is straightforward that (i) implies $\bar{u}_k$ does not depend on the choice of spanning tree: if two different choices yielded different values for some k, they must obtain a value for x_k by summing along different paths. But then the two paths from x_1 to x_k would define a loop around which the adding-up condition (i) fails. Thus $\bar{u}$ is well-defined, and it is immediate that $\operatorname{grad} \bar{u} = Y$.

We now show (ii) $\implies$ (iii). By Theorem 3, there exists an equivariant embedding $H : \mathbb{R}_+ \times X/\sim_{\trianglelefteq} \to X$, whose range X_H is closed in X, and intersects every $\sim_{\trianglelefteq}$-equivalence class. Following the notation of the Online Appendix D, let $(t, q_H) : X_H \to \mathbb{R}_+ \times X/\sim_{\trianglelefteq}$ denote the continuous inverse of H. Note that q_H is simply the restriction of the quotient map $q : X \to X/\sim_{\trianglelefteq}$ to X_H.

Suppose now that $Y \in \mathcal{F}$ is a gradient flow. Then there exists a vector $\bar{u} \in \mathcal{U}$ such that $\operatorname{grad} \bar{u} = Y$. Since the gradient of any constant vector is zero, we may assume, without loss of generality, that $\bar{u}$ is component-wise positive. We may similarly assume without loss that $\mathcal{V} \subsetneq X_H$.[54] Define $l : \mathcal{V} \to \mathbb{R}_+$ via $l(x_i) = t(x_i) + \big(\|\bar{u}\|_\infty - \bar{u}_i\big)$. By definition of an experiment, $\mathcal{V}$ and $q_H(\mathcal{V})$ are in one-to-one correspondence, hence we may equivalently regard l as a map from $q_H(\mathcal{V}) \to \mathbb{R}_+$.

[53]If $(\mathcal{V}, \mathcal{E})$ is not connected, our argument is valid applied to each connected component independently.

[54]If it is not, since X_H intersects every $\sim_{\trianglelefteq}$ equivalence class and $\mathcal{V}$ is finite, there exists some $\alpha^* > 0$ such that $\phi_{\alpha^*}(\mathcal{V}) \subseteq X_H$, and we may equivalently just work with this set of 'translates.'

By (A.1), there exists a cross section s for X and $\{\phi_\alpha\}_{\alpha\geq 0}$. By the universal property of the quotient (e.g. Munkres 1974), there is a map $s^* : X/{\sim_\trianglelefteq} \to X$ such that $s = s^* \circ q$. By definition of a cross section, $q \circ s^* = \mathrm{id}_{X/\sim_\trianglelefteq}$, thus q is a left inverse of s^*; since $X/{\sim_\trianglelefteq}$ carries the quotient topology, q is continuous, hence s^* is open. Since s^* is injective by definition of a cross section, it is a homeomorphism from $X/{\sim_\trianglelefteq}$ to a subspace of the metric space X, and therefore $X/{\sim_\trianglelefteq}$ is metrizable, which implies it is normal. The Tietze extension theorem, e.g. Munkres (1974), then guarantees there exists a bounded, continuous function $L : X/{\sim_\trianglelefteq} \to \mathbb{R}_+$ such that $L|_{\bar{q}(\mathcal{V})} = l$.

Let $\mathrm{epi}(L)$ denote the epigraph of L, by minor abuse of notation regarded as a subset of $\mathbb{R}_+ \times X/{\sim_\trianglelefteq}$.[55] We define a binary relation on X in three cases: first, if $x, x' \in H\big(\mathrm{epi}(L)\big) \subseteq X_H$, then let $x \succsim x'$ if and only if $t(x) - t(x') \geq (L \circ q_H)(x) - (L \circ q_H)(x')$. If x but not x' belong to $H\big(\mathrm{epi}(L)\big)$, then let $x \succ x'$. Finally, if neither x nor x' belong to $H\big(\mathrm{epi}(L)\big)$, then let $x \succsim x'$ if and only if $\min\{\alpha \geq 0 : \phi_\alpha(x') \in H\big(\mathrm{epi}(L)\big)\} \geq \min\{\alpha \geq 0 : \phi_\alpha(x) \in H\big(\mathrm{epi}(L)\big)\}$.[56]

Clearly $\succsim$ is complete. We now show that it is transitive and hence a preference relation. Let $x \succsim x'$ and $x' \succsim x''$, and suppose first that $x, x', x'' \in H\big(\mathrm{epi}(L)\big)$. Then:

$$t(x) - t(x') \geq (L \circ q_H)(x) - (L \circ q_H)(x')$$

and

$$t(x') - t(x'') \geq (L \circ q_H)(x') - (L \circ q_H)(x'').$$

By summing, we obtain:

$$t(x) - t(x'') \geq (L \circ q_H)(x) - (L \circ q_H)(x''),$$

and thus $x \succsim x''$. If instead $x, x' \in H\big(\mathrm{epi}(L)\big)$ but x'' is not, then it is immediate that $x \succsim x''$. Moreover, by construction it is impossible that $x', x'' \in H\big(\mathrm{epi}(L)\big)$ but x is not, as $x \succsim x'$ by hypothesis. Thus finally suppose that

[55] The epigraph of L is the set $\big\{([x], \tau) \in X/{\sim_\trianglelefteq} \times \mathbb{R}_+ : \tau \geq L([x])\big\}$. Here, we just reverse the order of the coordinates.

[56] As the range of H intersects every $\sim_\trianglelefteq$ equivalence class, both sets are non-empty, as well as closed and bounded below, which ensures the right-hand inequality is well-defined.

$x, x', x'' \notin H\big(\mathrm{epi}(L)\big)$. But then $x \succsim x''$ 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 $x \in H\big(\mathrm{epi}(L)\big)$. If $x' \succsim x$, then $x' \in H\big(\mathrm{epi}(L)\big) \subseteq X_H$ necessarily, and hence:

$$\begin{aligned}\{x' \in X : x' \succsim x\} &= \big\{x' \in X_H : t(x') - t(x) \geq (L \circ q_H)(x') - (L \circ q_H)(x)\big\} \\ &= H\big(\mathrm{epi}(L + \eta)\big),\end{aligned}$$

where the non-negative scalar $\eta \equiv t(x) - (L \circ q_H)(x) \geq 0$. As L is continuous, $\mathrm{epi}(L+\eta)$ is a closed subset of $\mathbb{R}_+ \times X/\sim_{\trianglelefteq}$. As H is an embedding, $H\big(\mathrm{epi}(L+\eta)\big)$ is closed in X_H; by Theorem 3 X_H is closed in X and hence $H\big(\mathrm{epi}(L+\eta)\big)$ is closed in X as well. Similarly,

$$\begin{aligned}\{x' \in X : x' \succ x\} &= \big\{x' \in X_H : t(x') - t(x) > (L \circ q_H)(x') - (L \circ q_H)(x)\big\} \\ &= H\big(\mathrm{int}\,\mathrm{epi}(L + \eta)\big),\end{aligned}$$

is open in $\bar{X}$. To show this set is open in X, we rely on the following claim.

> **Claim**: The set $\big\{x \in X : x \trianglelefteq H\big(0, q(x)\big)\big\}$ is closed.
>
> *Proof.* Suppose this is not the case. Then there exists $x_n \to x$, where $x_n \trianglelefteq H\big(0, q(x_n)\big)$ for all $n \in \mathbb{N}$, and $x \triangleright H\big(0, q(x)\big)$.[57] Thus for all $n \in \mathbb{N}$ there exists $\alpha_n \geq 0$ such that $\phi_{\alpha_n}(x_n) = H\big(0, q(x_n)\big)$, and $\alpha > 0$ such that $\phi_\alpha\big(H\big(0, q(x)\big)\big) = x$. By equivariance of H, we have that $x = H\big(\alpha, q(x)\big)$ and, for all $n \in \mathbb{N}$, that $\phi_{\alpha+\alpha_n}(x_n) = H\big(\alpha, q(x_n)\big)$. By continuity, $H\big(\alpha, q(x_n)\big) \to H\big(\alpha, q(x)\big)$ and hence $\lim_{n\to\infty} \phi_{\alpha+\alpha_n}(x_n) = x$. If any subsequence of $\alpha + \alpha_n$ converges to some limit $\bar{\alpha} \geq \alpha > 0$, then $\phi_{\bar{\alpha}}(x) = x$, violating regularity of the virtual commodity. Thus $\alpha + \alpha_n \to \infty$, as it is bounded below. Then, for all $\varepsilon > 0$ there exists $N_\varepsilon \in \mathbb{N}$ such that, for all $n \geq N_\varepsilon$ we have that both: (i) $x_n \in B_\varepsilon(x)$ and (ii) $\phi_{\alpha+\alpha_n}(x_n) \in B_\varepsilon(x)$, where $\alpha + \alpha_n$ may be chosen to be arbitrarily large. This violates (A.2), a contradiction. $\square$

[57] This last claim follows via Theorem 3 as it guarantees that $q \circ H\big(\alpha, q(x)\big) = q(x)$ for all $x \in X$, $\alpha \geq 0$ and hence $x \sim_{\trianglelefteq} H\big(0, q(x)\big)$ but $\neg\, x \trianglelefteq H\big(0, q(x)\big)$.

In light of this claim, $H\big(\mathbb{R}_{++} \times X/{\sim_{\trianglelefteq}}\big) = \big\{x \in X : x \trianglelefteq H\big(0, q(x)\big)\big\}^c$ is open in X, and since $H\big(\operatorname{int}\operatorname{epi}(L+\eta)\big) \subseteq H\big(\mathbb{R}_{++} \times X/{\sim_{\trianglelefteq}}\big)$, it too is open in X. Thus its complement, $\{x' \in X : x \succsim x'\}$, is closed.

Suppose now that $x \notin H\big(\operatorname{epi}(L)\big)$, and define $\alpha_x \equiv \min\{\alpha \geq 0 : \phi_\alpha(x) \in H\big(\operatorname{epi}(L)\big)\}$. Then $\phi_{\alpha_x}(x) = H\big((L \circ q)(x), q(x)\big)$, hence:

$$\begin{aligned}
\{x' \in X : x' \succsim x\} &= \{x' \in H\big(\operatorname{epi}(L)\big)^c : x' \succsim x\} \cup H\big(\operatorname{epi}(L)\big) \\
&= \{x' \in H\big(\operatorname{epi}(L)\big)^c : \alpha_x \geq \alpha_{x'}\} \cup H\big(\operatorname{epi}(L)\big) \\
&= \big\{x' \in H\big(\operatorname{epi}(L)\big)^c : x' \in \phi_{\alpha_x}^{-1}\big(H\big(\operatorname{epi}(L)\big)\big)\big\} \cup H\big(\operatorname{epi}(L)\big) \\
&= \phi_{\alpha_x}^{-1}\big(H\big(\operatorname{epi}(L)\big)\big),
\end{aligned}$$

where the third equality follows from the equivariance of H, and the fourth from the fact that $\phi_{\alpha_x}^{-1}\big(H\big(\operatorname{epi}(L)\big)\big)$ is closed under $\{\phi_\alpha\}_{\alpha \geq 0}$ (also by equivariance of H). Since ϕ_{α_x} is continuous and $H\big(\operatorname{epi}(L)\big)$ has already been shown to be closed, we conclude $\{x' \in X : x' \succsim x\}$ is closed. By an analogous argument, we obtain that $\{x' \in X : x' \succ x\}$ is open, and hence $\{x' \in X : x \succsim x'\}$ is closed. Thus $\succsim$ is continuous.

We now verify $\succsim$ obeys (N.1) - (N.3). Suppose then that $x \succsim x'$, and let $\alpha \geq 0$. If $x, x' \in H\big(\operatorname{epi}(L)\big)$, then, as $\phi_\alpha(x) = \phi_{\alpha + t(x)}\big(H(0, q(x))\big)$:

$$\begin{aligned}
(t \circ \phi_\alpha)(x) - (t \circ \phi_\alpha)(x') &= (t \circ \phi_{\alpha+t(x)})\big(H(0, q_H(x))\big) - (t \circ \phi_{\alpha+t(x')})\big(H(0, q_H(x'))\big) \\
&= (t \circ H)(\alpha + t(x), q_H(x)) - (t \circ H)(\alpha + t(x'), q_H(x')) \\
&= t(x) - t(y) \\
&\geq (L \circ q_H)(x) - (L \circ q_H)(y) \\
&= (L \circ q_H \circ \phi_\alpha)(x) - (L \circ q_H \circ \phi_\alpha)(x')
\end{aligned}$$

where the inequality follows from $x \succsim x'$ with $x, x' \in H\big(\operatorname{epi}(L)\big)$. Thus $\phi_\alpha(x) \succsim \phi_\alpha(x')$, as both belong to $H\big(\operatorname{epi}(L)\big)$ by equivariance of H. Suppose now x but not x' belongs to $H\big(\operatorname{epi}(L)\big)$, and thus that $x \succ x'$. Then for all $0 \leq \alpha < \alpha_{x'}$, by definition $\phi_\alpha(x) \succ \phi_\alpha(x')$. Suppose then that $\alpha \geq \alpha_{x'}$. By Lemma 4 $(t \circ \phi_\alpha)(x) = t(x) + \alpha$, where $t(x) \geq (L \circ q_H)(x)$. Similarly, since $x' \notin H\big(\operatorname{epi}(L)\big)$,

$(t \circ \phi_\alpha)(x') < (L \circ q)(x') + \alpha$. Hence:

$$\begin{aligned}(t \circ \phi_\alpha)(x) - (t \circ \phi_\alpha)(x') &= t(x) + \alpha - (t \circ \phi_\alpha)(x') \\ &\geq (L \circ q_H)(x) + \alpha - (t \circ \phi_\alpha)(x') \\ &> (L \circ q_H)(x) - (L \circ q)(x') \\ &= (L \circ q_H \circ \phi_\alpha)(x) - (L \circ q_H \circ \phi_\alpha)(x'),\end{aligned}$$

and hence $\phi_\alpha(x) \succ \phi_\alpha(x')$. Finally, suppose neither x nor x' belong to $H\big(\mathrm{epi}(L)\big)$. Let $x \succsim x'$ and hence $\alpha_{x'} \geq \alpha_x$. For any $\alpha < \alpha_x$, if $\tilde{x} = \phi_\alpha(x)$, we have that $\alpha_{\tilde{x}} = \alpha_x - \alpha$, hence for any such α, it follows that $\phi_{\tilde{\alpha}}(x) \succsim \phi_{\tilde{\alpha}}(x')$. If $\alpha \geq \alpha_x$, then $\phi_\alpha(x) \in H\big(\mathrm{epi}(L)\big)$; if $\phi_\alpha(x')$ is not, then $\phi_\alpha(x) \succ \phi_\alpha(x')$ as desired. If $\phi_\alpha(x') \in H\big(\mathrm{epi}(L)\big)$ too, then:

$$\begin{aligned}(t \circ \phi_\alpha)(x) - (t \circ \phi_\alpha)(x') &\geq (t \circ \phi_\alpha)(x) - (t \circ \phi_\alpha)(x') + \alpha_x - \alpha_{x'} \\ &= (L \circ q_H \circ \phi_\alpha)(x) - (L \circ q_H \circ \phi_\alpha)(x'),\end{aligned}$$

and thus $\succsim$ satsifies (N.1). Property (N.2) holds by definition. Thus now suppose $x' \succsim x$. Then $\phi_{\alpha_x}(x), \phi_{\alpha_x}(x') \in H\big(\mathrm{epi}(L)\big)$, thus, having verified (N.1) it suffices to find some α such that $\phi_{\alpha+\alpha_x}(x) \sim \phi_{\alpha_x}(x')$. Let $\alpha = (t \circ \phi_{\alpha_x})(x') - (L \circ q_H \circ \phi_{\alpha_x})(x')$.[58] Since $(t \circ \phi_{\alpha_x})(x) = (L \circ q_H \circ \phi_{\alpha+\alpha_x})(x)$, it follows that:

$$\begin{aligned}(t \circ \phi_{\alpha+\alpha_x})(x) - (t \circ \phi_{\alpha_x})(x') &= \alpha + (t \circ \phi_{\alpha_x})(x) - (t \circ \phi_{\alpha_x})(x') \\ &= \alpha + (L \circ q_H \circ \phi_{\alpha+\alpha_x})(x) - (t \circ \phi_{\alpha_x})(x') \\ &= (L \circ q_H \circ \phi_{\alpha+\alpha_x})(x) - (L \circ q_H \circ \phi_{\alpha_x})(x'),\end{aligned}$$

thus $\phi_{\alpha+\alpha_x}(x) \sim \phi_{\alpha_x}(x')$ as desired, and we conclude $\succsim$ 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 $Y_{ij} \geq 0$. Suppose first $x_i, x_j \in H\big(\mathrm{epi}(L)\big)$. Since $x_i, x_j \in \mathcal{V}$, by construction $(L \circ q_H)(x_i) = l(x_i) + \big(\|u\|_\infty - u_i\big)$ and likewise x_j. Thus

$$\begin{aligned}t(x_j) - (t \circ \phi_{Y_{ij}})(x_i) &= t(x_j) - t(x_i) - Y_{ij} \\ &= t(x_j) - t(x_i) - (u_j - u_i) \\ &= (L \circ q_H)(x_i) - (L \circ q_H \circ \phi_{Y_{ij}})(x_j)\end{aligned}$$

[58] Note this is well-defined as $\phi_{\alpha_x}(x') \in H\big(\mathrm{epi}(L)\big)$.

and hence $x_i \sim \phi_{Y_{ij}}(x_j)$. If x_i or x_j do not belong to $H\big(\text{epi}(L)\big)$, as $\succsim$ satisfies (N.1), it suffices to verify that:

$$\phi_{Y_{ij}+\max\{\alpha_{x_i},\alpha_{x_j}\}}(x_i) \sim \phi_{\max\{\alpha_{x_i},\alpha_{x_j}\}}(x).$$

However, by construction both these alternatives belong to $H\big(\text{epi}(L)\big)$, and hence the preceding argument applies directly. Thus (i) $\implies$ (ii). The theorem follows. $\square$

Appendix C Proposition Proofs

C.1 Proof of Proposition 2

Proof. By the Pythagorean theorem:

$$\left\|Y - Y^*_{\mathcal{M}}\right\|_2^2 = \left\|Y - Y^*_{Add}\right\|_2^2 + \left\|Y^*_{Add} - Y^*_{\mathcal{M}}\right\|_2^2 \tag{16}$$

and, letting $\Pi_{\text{grad}(\mathcal{K}_{\mathcal{M}})}$ denote L^2 projection onto $\text{grad}(\mathcal{K}_{\mathcal{M}})$,

$$\left\|Y - \Pi_{\text{grad}(\mathcal{K}_{\mathcal{M}})} Y^*_{Add}\right\|_2^2 = \left\|Y - Y^*_{Add}\right\|_2^2 + \left\|Y^*_{Add} - \Pi_{\text{grad}(\mathcal{K}_{\mathcal{M}})} Y^*_{Add}\right\|_2^2. \tag{17}$$

As $Y^*_{\mathcal{M}}$ is the nearest point in $\mathcal{K}_{\mathcal{M}}$ to Y, the fact $\Pi_{\text{grad}(\mathcal{K})} Y^*_{Add} \in \text{grad}(\mathcal{K}_{\mathcal{M}})$ implies:

$$\left\|Y - \Pi_{\text{grad}(\mathcal{K}_{\mathcal{M}})} Y^*_{Add}\right\|_2^2 \geq \left\|Y - Y^*_{\mathcal{M}}\right\|_2^2$$

and hence by (16) and (17),

$$\left\|Y^*_{Add} - \Pi_{\text{grad}(\mathcal{K}_{\mathcal{M}})} Y^*_{Add}\right\|_2^2 \geq \left\|Y^*_{Add} - Y^*_{\mathcal{M}}\right\|_2^2. \tag{18}$$

But analogously, $\Pi_{\text{grad}(\mathcal{K}_{\mathcal{M}})} Y^*_{Add}$ must be the nearest point to Y^*_{Add} in $\text{grad}(\mathcal{K}_{\mathcal{M}})$, hence (18) holds with equality. As the L^2 norm is strictly convex, this means $\Pi_{\text{grad}(\mathcal{K}_{\mathcal{M}})} Y^*_{Add} = Y^*_{\mathcal{M}}$, and substituting into (16) yields the desired equation. $\square$

C.2 Proof of Proposition 3

Proof. 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:

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

By taking infimums across all such decompositions, it follows $\|R\|_1 \leq MP^*(R)$. 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.[59] If $R = 0$ then trivially $MP^*(R) = \|R\|_1 = 0$, hence suppose $R \neq 0$. Let $\mathcal{E}'$ denote the non-empty subset of edges on which $R \neq 0$, and let $\mathcal{V}'$ denote the set of vertices appearing in some edge in $\mathcal{E}'$. Choose $x_{i_0} \in \mathcal{V}'$ arbitrarily. Since $x_{i_0} \in \mathcal{V}'$, by (8) there exists a neighboring vertex x_{i_1} such that $R_{x_{i_0} x_{i_1}} > 0$. Proceeding in this fashion, we may construct a sequence of oriented edges in $\vec{\mathcal{E}}'$ such that $R_{x_{i_l} x_{i_{l+1}}} > 0$. We terminate this process when we choose a vertex that has appeared prior in the sequence.[60] By possibly throwing out some initial segment of this sequence and relabelling indices, we obtain a sequence of oriented edges $(x_{i_0}, x_{i_1}), (x_{i_1}, x_{i_2}), \ldots, (x_{i_{L_1}}, x_{i_0})$ such that $R_{x_{i_l} x_{i_{l+1}}} > 0$, where $i_L \equiv i_0$. Let $c_1 = \min\{R_{x_{i_0} x_{i_1}}, R_{x_{i_1} x_{i_2}}, \ldots, R_{x_{i_{L-1}} x_{i_0}}\}$, and let C_1 denote the pure cycle $\sum_{l=0}^{L} c_1 \mathbb{1}_{(x_{i_l}, x_{i_{l+1}})}$. 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 the subspace satisfying (8); however it is non-zero on a proper subgraph of $(\mathcal{V}', \mathcal{E}')$. Thus repeating this process, we obtain a finite decomposition of R into pure cycles $C_1 + \cdots + C_M$, where for all $m = 1, \ldots, M$, $C_m \geq 0$. This implies:

$$\|R\|_1 = \left\| \sum_{m}^{M} C_m \right\|_1 = \sum_{m=1}^{M} \left\|C_m\right\|_1 = \sum_{m=1}^{M} MP(C_m)$$

[59] 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).

[60] This process necessarily terminates, as $\mathcal{V}'$ is finite.

and hence the lower bound obtains. $\square$

C.3 Proof of Proposition 4

Proof. Any preference in $\mathcal{M}_{MEU}$ may be uniquely identified with its set of priors $C \subseteq \Delta(S)$. Given a closed, convex set of priors C satisfying (i) and (ii) for some $u, \pi^1, \ldots, \pi^K$, let $\hat{\pi}^1, \ldots, \hat{\pi}^K \in C$ denote selections of extremal points from C belonging to the respective hyperplanes $\bar{H}_{\tilde{u}^i, u_i}$. Then trivially $u, \hat{\pi}^1, \ldots, \hat{\pi}^K$ belong to $\mathcal{K}_{MEU}$, and must be consistent with $u, \pi^1, \ldots, \pi^K$ as u is the same in both vectors.

Conversely, given some set of priors C that is consistent with $u, \pi^1, \ldots, \pi^K$, suppose for purposes of contraposition that there exists $\pi^* \in C$, $\pi^* \notin \bar{C}$, then for some i:

$$\langle \pi^*, \tilde{u}^i \rangle < u_i,$$

and hence $u_i \neq U(\tilde{u}^i)$, where U is the MEU functional associated with C. Thus if C is consistent with $u, \pi^1, \ldots, \pi^K$, C must satisfy (i). For (ii), if for all $i = 1, \ldots, K$:

$$u_i = U(\tilde{u}^i) = \min_{\pi \in C} \langle \pi, \tilde{u}^i \rangle,$$

then $\bar{H}_{\tilde{u}^i, u_i}$ must be a supporting hyperplane for C and hence must contain some extremal point. $\square$

Appendix D Proof of Theorem 3

Proof of Lemma 1

Proof. Clearly $\sim_{\trianglelefteq}$ is reflexive and symmetric, hence all that remains is to verify transitivity. Suppose $x \sim_{\trianglelefteq} x'$ and $x' \sim_{\trianglelefteq} x''$. We proceed in three cases: first suppose that only one of x and x'' may be obtained from x'; without loss $x \trianglelefteq x' \trianglelefteq x''$. Then there exists $\alpha_{xx'}, \alpha_{x'x''} \geq 0$ such that $\phi_{\alpha_{xx'}}(x) = x'$ and $\phi_{\alpha_{x'x''}}(x') = x''$ then clearly $\phi_{\alpha_{xx'}+\alpha_{x'x''}}(x) = x''$ and hence $x \trianglelefteq x''$. Thus suppose $x' \trianglelefteq x$ and $x' \trianglelefteq x''$. Then there exists $\alpha_{x'x}, \alpha_{x'x''} \geq 0$ such that $\phi_{\alpha_{x'x}}(x') = x$ and $\phi_{\alpha_{x'x''}}(x') = x''$. Without loss of generality let $\alpha_{x'x} \leq \alpha_{x'x''}$, so $\phi_{\alpha_{x'x''}-\alpha_{x'x}}\big(\phi_{\alpha_{x'x}}(x')\big) = x''$ and thus $\phi_{\alpha_{x'x''}-\alpha_{x'x}}(x) = x''$, and we obtain $x \sim_{\trianglelefteq} x''$. Finally, suppose $x \trianglelefteq x'$ and $x'' \trianglelefteq x'$. Then there exists $\alpha_{xx'}, \alpha_{x''x'} \geq 0$ such that $\phi_{\alpha_{xx'}}(x) = x' = \phi_{\alpha_{x''x'}}(x'')$. Without loss, let $\alpha_{xx'} \leq \alpha_{x''x'}$. Then $x' = \phi_{\alpha_{x''x'}}(x'') = \phi_{\alpha_{xx'}+(\alpha_{x''x'}-\alpha_{xx'})}(x'')$, which in turn equals $\phi_{\alpha_{xx'}}\big(\phi_{\alpha_{x''x'}-\alpha_{xx'}}(x'')\big)$. But, by regularity, the map $\phi_{\alpha_{xx'}}$ is injective hence, $\phi_{\alpha_{x''x'}-\alpha_{xx'}}(x'') = x$ and therefore $x \sim_{\trianglelefteq} x''$. □

In light of Lemma 1, there is a well-defined quotient space $X/{\sim_{\trianglelefteq}}$. We let $q : X \to X/{\sim_{\trianglelefteq}}$ denote the associated quotient map, and in all that follows, we will consider $X/{\sim_{\trianglelefteq}}$ endowed with its quotient topology; see Munkres (1974) for definitions.

Corollary 1. *For all $\alpha \geq 0$, for all $x \in X$, $q(x) = (q \circ \phi_\alpha)(x)$.*

The quotient map q has the property that if $f : X \to Z$ is any map that is constant on each $\sim_{\trianglelefteq}$ equivalence class, then there is a uniquely determined map $f^* : X/{\sim_{\trianglelefteq}} \to Z$ such that $f^*\big(q(x)\big) = f(x)$ for all $x \in X$.[61] In particular, every cross section $s : X \to X$ may equivalently be regarded as a map $s : X/{\sim_{\trianglelefteq}} \to X$. Going forward, whenever we refer to a cross section s we will mean it in this latter sense. To conserve on notation, we will reserve the use of y to denote elements of $X/{\sim_{\trianglelefteq}}$.

We now fix a regular virtual commodity $\{\phi_\alpha\}_{\alpha\geq 0}$ that satisfies (A.1) and (A.2), and a choice of cross section s. Define $H : \mathbb{R}_+ \times X/{\sim_{\trianglelefteq}} \to X$ via $H(\alpha, y) = (\phi_\alpha \circ s)(y)$, and let $X_H = \text{range}(H)$. We wish to show that H is an equivariant embedding. Note that equivariance follows immediately from the definition of H:

$$\begin{aligned}\phi_\beta\big(H(\alpha, y)\big) &= \phi_\beta\big((\phi_\alpha \circ s)(y)\big)\\ &= (\phi_{\beta+\alpha} \circ s)(y)\\ &= H(\beta + \alpha, y).\end{aligned}$$

Thus it remains to verify that H is indeed an embedding.

Lemma 2. *Let $q_H : X_H \to X/{\sim_{\trianglelefteq}}$ be the restriction of q to X_H. Then q_H is an open map.*

[61] See Munkres (1974) for details.

Proof. Let $U \subset X_H$ be open. Then:

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

But, s and $\{\phi_\alpha\}_{\alpha \geq 0}$ are continuous, hence $q_H(U)$ is open. □

Lemma 3. *H is injective.*

Proof.

$$\begin{aligned} H(\alpha, y) &= H(\alpha', y') \\ \implies (\phi_\alpha \circ s)(y) &= (\phi'_\alpha \circ s)(y') \\ \implies (q \circ \phi_\alpha \circ s)(y) &= (q \circ \phi'_\alpha \circ s)(y') \\ \implies (q \circ s)(y) &= (q \circ s)(y') \\ \implies s(y) &= s(y'). \end{aligned}$$

This implies that $y = y'$ as s is a cross section, and that $\alpha = \alpha'$ by regularity.

□

Define $t : X_H \to \mathbb{R}_+$ implicitly, as the unique solution to

$$H\big(t(x), q_H(x)\big) = x.$$

It is well-defined in light of the equivariance of H and regularity of $\{\phi_\alpha\}_{\alpha \geq 0}$. By definition, (t, q_H) is the inverse of H. We now establish the regularity (i.e. continuity) of solutions to the above class of topological implicit function problems (Lemma 5 - Lemma 7).

Lemma 4. *For all $x \in X_H, \alpha \geq 0$, the function t satisfies $(t \circ \phi_\alpha)(x) = t(x) + \alpha$.*

Proof. Let $x \in X_H$. By definition, $H\big(t(x), q_H(x)\big) = x$, and by equivariance, $H\big(t(x) + \alpha, q_H(x)\big) = (\phi_\alpha \circ H)(t(x), q_H(x)) = \phi_\alpha(x)$. Hence $(t \circ \phi_\alpha)(x) = (t \circ H)\big(t(x) + \alpha, q_H(x)\big) = t(x) + \alpha$. □

Lemma 5. *For all $x \in X_H$ there exists a finite open cover $\{N_{\alpha_i}\}_{i=1}^K$ of $H\big([0,t(x)] \times \{q_H(x)\}\big)$ with the following properties:*

1. *For all $i \in \{1,\dots,K\}$, the set $\{\alpha : H(\alpha, \bar{q}(x)) \in N_{\alpha_i}\}$ 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 $i, j = 1,\dots,K$, $\alpha_i < \alpha_j$ 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 accumulation property of (A.2).*

Proof. Fix $x \in X_H$. For all $\alpha \in [0, t(x)]$, define $x_\alpha = H(\alpha, q_H(x)) = (\phi_\alpha \circ s \circ q_H)(x)$. By (A.2), 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 $B_{\varepsilon_\alpha}(x_\alpha) \cap H\big([0,t(x)] \times \{q_H(x)\}\big)$ that contains x_α, and define $N_\alpha = B_{\varepsilon_\alpha}(x_\alpha) \setminus \big[H\big([0,t(x)] \times \{q_H(x)\}\big) \setminus U_\alpha\big]$. As $[0,t(x)] \times \{q_H(x)\}$ is compact in $\mathbb{R}_+ \times X/\sim_{\trianglelefteq}$, by continuity $H\big([0,t(x)] \times \{q_H(x)\}\big)$ is a compact and hence closed subset of X_H. U_α is a relatively open subset of $H\big([0,t(x)] \times \{q_H(x)\}\big)$, hence $H\big([0,t(x)] \times \{q_H(x)\}\big) \setminus U_\alpha$ 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. By Lemma 3, $H\big(\,\cdot\,, q_H(x)\big)$ is injective (and continuous) hence for all α, $\{\alpha' : H\big(\alpha', q_H(x)\big) \in N_\alpha\}$ is an open interval in $[0, t(x)]$.

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

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 $\underline{\alpha}_i < \underline{\alpha}_j$,

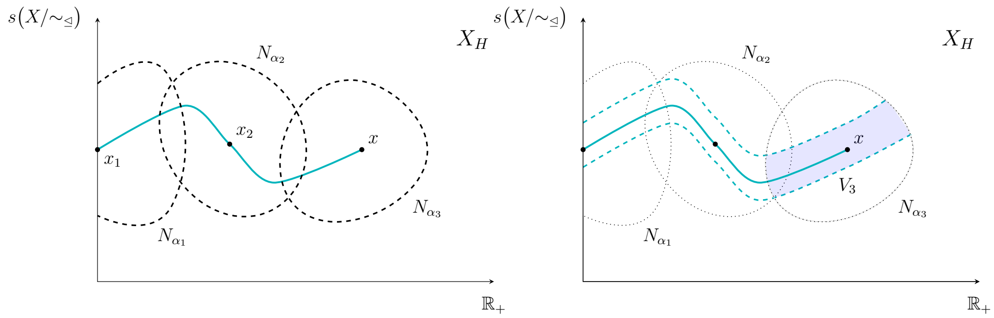

(a) An open cover of the path $H([0, t(x)] \times \{q_H(x)\})$, here in aquamarine. This open cover satisfies all of the properties of Lemma 5.

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

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.[62] 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.). $\square$

Lemma 6. *For all $x \in X_H$ there exists an open neighborhood of x on which t is bounded.*

Proof. Fix $x \in X_H$, and let $\{N_{\alpha_i}\}_{i=1}^{K}$ denote an open cover of $H([0, t(x)] \times \{q_H(x)\})$ of the form guaranteed by Lemma 5. Without loss of generality, suppose that N_{α_1} is the sole cover element to intersect $H(\{0\} \times X/{\sim_{\trianglelefteq}})$.[63]

[62] 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.

[63] 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 $x_n \in \text{range}(s)$,

Define $V_0 = H\big(\{0\} \times X/{\sim_{\trianglelefteq}}\big)$ and, for all $i = 1, \ldots, K$:

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

see Figure 5. We first verify, for all $i = 1, \ldots, K$, that V_i is open. Note that via Lemma 2 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 $V_1 = N_{\alpha_1} \cap \big(q_H^{-1} \circ q_H\big)(V_0 \cap N_{\alpha_1})$, and $V_0 \cap N_{\alpha_1} = N_{\alpha_1} \cap \mathrm{range}(s)$, and hence is relatively open in the range of s. As q_H is a left-inverse of s, $q_H(N_{\alpha_1} \cap V_0)$ is open, and hence so too is V_1.

We now establish that, for all $i = 1, \ldots, K$, $H\big([0, \bar{\alpha}_i) \times \{q_H(x)\}\big) \subseteq \bigcup_{j \le i} V_j$, where we recall that $(\underline{\alpha}_i, \bar{\alpha}_i) = \{\alpha \in [0, t(x)] : H(\alpha, \bar{q}(x)) \in N_{\alpha_i}\}$ for $1 < i < K$, and $[0, \bar{\alpha}_i)$ is the analogue for $i = 1$.[64] For all $i = 1, \ldots, K$, let $x_{\alpha_i} = H(\alpha_i, \bar{q}(x))$ and consider the case of $i = 1$. By hypothesis, $\alpha_1 = 0$, hence $x_{\alpha_1} = (s \circ q_H)(x) \in N_{\alpha_1} \cap V_0$. Then Corollary 1 implies $H\big([0, t(x)] \times \{q_H(x)\}\big) \subseteq (q_H^{-1} \circ q_H)(N_{\alpha_1} \cap V_0)$, and thus $H\big([0, \bar{\alpha}_1) \times \{q_H(x)\}\big) \subseteq V_1$. Suppose now that, for all $1 \le i \le k$, that $H\big([0, \bar{\alpha}_i) \times \{\bar{q}(x)\}\big) \subseteq \bigcup_{j \le i} V_j$, but, for sake of contradiction, that $H\big([0, \bar{\alpha}_{k+1}) \times \{q_H(x)\}\big) \not\subseteq \bigcup_{j \le k+1} V_j$. As $(\underline{\alpha}_{k+1}, \bar{\alpha}_{k+1})$ is an interval, if $\bar{\alpha}_k \in (\underline{\alpha}_{k+1}, \bar{\alpha}_{k+1})$, the contradiction hypothesis would be false, thus it must be that $\bar{\alpha}_k \notin (\underline{\alpha}_{k+1}, \bar{\alpha}_{k+1})$ and hence $H(\bar{\alpha}_k, \bar{q}(x)) \notin N_{\alpha_{k+1}}$. Then $(\underline{\alpha}_k, \bar{\alpha}_k) \cap (\underline{\alpha}_{k+1}, \bar{\alpha}_{k+1}) = \varnothing$. But Lemma 5 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 $H(\bar{\alpha}_k, \bar{q}(x)) \notin \bigcup_{i=1}^K N_{\alpha_i}$, contradicting the fact that $\{N_{\alpha_i}\}_{i=1}^K$ is a cover for $H\big([0, t(x)] \times \{q_H(x)\}\big)$. Thus by induction $H\big([0, \bar{\alpha}_K) \times \{q_H(x)\}\big) \subseteq \bigcup_{j \le K} V_j$, and in particular $x = x_{\alpha_K} \in V_K$.

We now verify that $t|_{V_i}$ is bounded for all $i = 0, \ldots, K$; 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 $x' \in V_1$. By Lemma 4, for any $\underline{x}' \in V_1$, if $\phi_\alpha(\underline{x}') = x'$, then $t(x') = \alpha + t(\underline{x}')$. But since

---

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 $x \in \mathrm{range}(s)$.

[64] This set is indeed an interval by Lemma 5.

N_{α_1} has a no-loitering bound of T_{α_1}, since both $x', \underline{x}' \in V_1 \subseteq N_{\alpha_1}$, we have $t(x') < T_{\alpha_1} + t(\underline{x}')$. However, if $x' \in V_1$, then $(s \circ q_H)(x') \in V_1$, and by definition $(t \circ s \circ q_H)(x') = 0$. Thus for all $x' \in V_1$, $t(x') < T_{\alpha_1}$. Suppose now that, for all $i \leq k$, $t|_{V_i}$ is bounded, and let $x' \in V_{k+1}$. Then, $x' \in N_{\alpha_{k+1}}$ and there exists some $x'' \sim_{\trianglelefteq} x'$, where $x'' \in N_{\alpha_{k+1}} \cap V_j$ where $1 \leq j \leq k$. Suppose $x'' \trianglelefteq x'$. Then:

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

where $T_{\alpha_{k+1}}$ is a (A.2) bound for $N_{\alpha_{k+1}}$, and $\bar{T}_j$ is any upper bound on $t|_{V_j}$ which exists by the induction hypothesis. Note that if $x' \trianglelefteq x''$, then $t(x')$ is bounded above by the same quantity. Thus for all $1 \leq i \leq K$, $t|_{V_i}$ is bounded; as $x \in V_K$ and V_K is open, this establishes the claim. □

Lemma 7. *The map t is continuous.*

Proof. Let $x \in X_H$. By Lemma 6, 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

$$t^*(x') = \arg\min_{\tilde{t} \in [0,K]} d_X\big((\phi_{\tilde{t}} \circ s \circ q_H)(x'), x'\big),$$

for $x' \in B_\varepsilon(x)$. 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 $x \in X_H$ there is a neighborhood of x on which it is continuous, hence it is continuous. □

Corollary 2. *Suppose $\{\phi_\alpha\}_{\alpha \geq 0}$ is a regular virtual commodity satisfying (A.1) and (A.2). Then H is an equivariant embedding.*

Proof of Theorem 3

Proof. By Corollary 2, H is an equivariant embedding and by definition, satisfies the desired identity. Suppose then, for sake of contradiction, that X_H is not closed in X. Then there exists a convergent sequence $x_n \in \text{range}(H)$, $x_n \to x$ with $x \notin \text{range}(H)$. By construction, for every $x \in X$, $x \sim_{\trianglelefteq} H\big(0, q(x)\big) = (s \circ q)(x)$; since $H\big(0, q(x)\big) \trianglelefteq x$ implies $x \in X_H$ by equivariance, it follows that $x \triangleleft H\big(0, q(x)\big)$. Thus there exists $\alpha^* > 0$ such that $\phi_{\alpha^*}(x) = H\big(0, q(x)\big)$. By continuity, $\phi_{\alpha^*}(x_n) \to \phi_{\alpha^*}(x)$, and thus $(t \circ \phi_{\alpha^*})(x_n) \to (t \circ \phi_{\alpha^*})(x) = 0$. By Lemma 4, for all $n \in \mathbb{N}$, the sequence $(t \circ \phi_{\alpha^*})(x_n) = t(x_n) + \alpha^*$, is bounded away from zero by α^*, a contradiction. Thus X_H is closed.

For necessity, suppose an equivariant embedding H with the requisite properties exist. It is immediate that (A.2) holds for $\{\phi_\alpha\}_{\alpha \geq 0}$ and that $H\big(0, q(\cdot)\big)$ defines a cross section. $\square$

Appendix E Dominant Strategy Elicitation of Compensation Data

In this section we 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). For simplicity, we will consider the elicitation problem for a given observation; our results extend to full experiments straightforwardly. Let $\{x, x'\} \in \mathcal{E}$ 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 x'-mechanism. Compensation differences may be elicited by presenting the subject with a choice: they are invited to submit a sell price in either the x- or y-mechanism, but not both. However, in whichever mechanism they do not choose, a sell price of 0 will be submitted on their behalf. After the bids have been submitted, a coin is flipped to select either x or y, and the associated mechanism's reward is allocated to the agent, regardless of which intermediate mechanism they chose to manually submit a sell price for.

We model the agent's decision problem using the states of the world formalism. We do so to highlight that the incentive-compatibility of our mechanism does not depend on the manner in which the subject handles probabilities. Suppose that $\Omega = \mathbb{R}^2_+ \times \{x, x'\}$ denotes the payoff-relevant states of the world; the tuple $(b_x, b_{x'}, z)$ denotes the state in which the computer selects bids b_x in the x-mechanism, $b_{x'}$ in the x'-mechanism, and the payoff-determining mechanism is $z \in \{x, x'\}$. A choice of action for the agent consists of a tuple in $\{x, x'\} \times \mathbb{R}_+$, 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 4. *Let $\{\phi_\alpha\}_{\alpha \geq 0}$ be a virtual commodity, and suppose an agent (i) has preferences $\succsim$ on X that satisfy (N.2) and (N.3), and (ii) 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 $\succsim$-preferred alternative, is $\succsim^*$-optimal.*

Proof. Suppose $x \succsim x'$, with true compensation difference given by $\alpha \geq 0$, $\phi_\alpha(x') \sim x$. Since $\succsim$ satisfies (N.2) and (N.3), this α exists and is unique. Suppose first that the subject chooses to participate in the x'-mechanism and submits a price of s. Then their state-dependent payoff is the act:

$$f_s(b_x, b_{x'}, z) = \begin{cases} \phi_{b_x}(x') & \text{if } z = x \\ \phi_{b_{x'}}(x) & \text{if } z = x',\ b_{x'} \geq s \\ x' & \text{if } z = x',\ s > b_{x'}. \end{cases}$$

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

$$g_s(b_x, b_{x'}, z) = \begin{cases} \phi_{b_{x'}}(x) & \text{if } z = x' \\ \phi_{b_x}(x') & \text{if } z = x,\ b_x \geq s \\ x & \text{if } z = x,\ s > b_x \end{cases}$$

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

$$\phi_{b_x}(x') \succsim x \iff b_x \geq \alpha,$$

hence conditional upon $z = x$, the agent obtains $\max\{\phi_{b_x}(x'), x\}$ from g_α.[65] Now, by (N.2), $\phi_{b_y}(x) \succsim x'$ no matter the value of $b_{x'}$, hence by consistency of $\succsim^*$ the most-preferred f act resulting from a bid in the x'-mechanism is f_0.[66] Thus we wish to show $g_\alpha \succsim^* f_0$. But conditional upon $z = x'$, both g_α and f_0 yield $\phi_{b_{x'}}(x)$, and conditional upon $z = x$, g_α yields $\max\{\phi_{b_x}(x'), x\}$ whereas f_0 yields $\phi_{b_x}(x')$. 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. $\square$

The assumption that $\succsim^*$ was a preference relation is not required for the result. All that was needed was that $\succsim^*$ was consistent with $\succsim$. In principle $\succsim^*$ could be highly incomplete and nontransitive; so long as consistency is satisfied, Theorem 4 remains valid.

Appendix F Proofs Omitted from the Text

F.1 Omitted Arguments from Section 5.1

F.1.1 Positive Homogeneous & Translation Invariant Rationalizations

We first argue that any data set arising from $(\mathcal{V}, \mathcal{E})$ is rationalizable by a utility function of the form:

$$w\big(v(x_1), v(x_2)\big),$$

with w translation-invariant and positive homogeneous. Note that by translation-invariance, it suffices to define the level set through 0 of w, as every other level set must then be determined by this via translation along the diagonal. Firstly,

[65] The max here is understood in the $\succsim$ sense.

[66] That is, from setting $s = 0$.

note that under v the monetary acts $(0, 0)$, $(1, 0)$, and $(0, 1)$ correspond to the identical utility acts. It will be more convenient to work in utility act space. Firstly, Y_{02} yields a utility act $\bar{v}_2 \sim (0, 0)$; if the utility acts $(0, 1) \succsim (0, 0)$ (i.e. $Y_{02} \geq 0$), then the utility act $\bar{v}_2 = (-Y_{02}, 1 - Y_{02}) \sim 0$. If $Y_{02} < 0$, then $\bar{v}_2 = (Y_{02}, 1 + Y_{02}) \sim 0$. Analogously, we can find a utility act $\bar{v}_1 \sim (0, 0)$ by first finding a translation of $(1, 0)$ that is indifferent with $(0, 1)$ and then subtracting Y_{02} from both components to obtain indifference with $(0, 0)$.[67] Define the 0-level set of w to be the union of the rays from 0 through $\bar{v}_1$ and $\bar{v}_2$ respectively, and extend to a functional on $\mathbb{R}^2$ via translation invariance. To see this extension is also positive homogeneous, note that the restriction of w to the half space above (resp. below) the diagonal is linear. Finally, note that any translation-invariant and positive homogeneous rationalization must share the same 0-level set as w, and therefore must arise in this fashion. In particular, this means we do not 'miss' any possible rationalizations by considering this construction.

F.1.2 Relaxing Ambiguity Aversion

As noted above, every translation invariant functional on $\mathbb{R}^2$ may be identified with its level set through 0; if the functional is additionally positive homogeneous, 0-level set must be a union of two rays from 0. If the functional is monotone, these two rays must lie in the second and fourth quadrants respectively; if it is also concave, the upper contour set (i.e. the region bounded between these rays containing the positive 45-degree ray) must be convex. If we drop concavity of the functional, but require monotonicity, translation-invariance and positive homogeneity, then this just requires that $\bar{v}_1$ and $\bar{v}_2$ belong to the fourth and second quadrants of the plane respectively, with no other constraint. Since these utility acts depend on Y, it is straightforward to obtain the constraint rhombus in Figure 3. To see this coincides with CEU, note that for any such pair of rays, there is a vector in $\Delta(S)$ normal to each,

[67] Note that $\bar{v}_1$ must lie below the diagonal of $\mathbb{R}^2$ as it a translation of $(1, 0)$ parallel to the diagonal, and similarly $\bar{v}_2$ must lie above.

and w may be viewed as integrating against the capacity that takes on one of these probability measures depending on whether the integrand lies above or below the diagonal in the space of utility acts.

Appendix G Constraint Set Characterizations

G.1 Quasilinear Increasing Concave Utility

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 (11) with $\mathcal{K} = \mathcal{K}_{QIC}$ is equivalent to solving:

$$\begin{aligned}
\min_{\bar{u} \in \mathcal{U}} \quad & \left\|(\operatorname{grad} \bar{u}) - Y\right\|_2^2 & & \\
\text{subject to} \quad & \bar{u}_i = \langle \pi_i, x_i \rangle + \gamma_i & \forall i = 1, \ldots, K & \\
& \langle \pi_i, x_i \rangle + \gamma_i \leq \langle \pi_j, x_i \rangle + \gamma_j & \forall i, j = 1, \ldots, K & \qquad (19) \\
& \pi_{i,1} = 1 & \forall i = 1, \ldots, K & \\
& \pi_i \geq 0 & \forall i = 1, \ldots, K &
\end{aligned}$$

for $\bar{u} \in \mathcal{U}$ and, for all $i = 1, \ldots, K$, $\pi_i \in \mathbb{R}^L$, $\gamma_i \in \mathbb{R}$ (where $\pi_{i,1}$ denotes the first component of π_i).

Proof. Suppose first that u is a quasilinear (with linear term normalized to identity), increasing, and concave utility. For all $i = 1, \ldots, K$, define $\bar{u}_i = u(x_i)$ and let π_i denote an arbitrary choice of supergradient of u at each x_i. As u is increasing, it follows $\pi_i \geq 0$ for each i. Define $\gamma_i = \bar{u}_i - \langle \pi_i, x_i \rangle$. Then for all $i = 1, \ldots, K$ and all $x \in X$:

$$u(x) \leq u(x_i) + \langle \pi_i, x - x_i \rangle.$$

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

$$u\big(\phi_\alpha(x_i)\big) \leq u(x_i) + \langle \pi_i, (\alpha, 0) \rangle$$

it follows that:

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

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

$$u(\hat{x}) \leq u(x_i) + \langle \pi_i, (-\alpha, 0) \rangle,$$

which yields $-\alpha \leq -\alpha \pi_i^1$ and hence $\pi_i^1 \leq 1$. Thus for all x 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 x on the boundary of X, and hence for all x_i, π_i is of the form $(1, \pi_i^2)$ as claimed.

Conversely, suppose $\bar{u}$, $\{\pi_i\}_{i=1}^K$, $\{\gamma_i\}_{i=1}^K$ is a solution to (19). Define:

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

Then clearly $\hat{u}(x_i) = \bar{u}_i$, and $\hat{u}$ is quasilinear, increasing, and concave. □

G.2 Constant Absolute Ambiguity Aversion Preferences

Throughout, we abuse notation by writing $v(x)$ for the vector $\big(v(x_1), \dots, v(x_S)\big)$, and $\bar{v}_i$ for the utility act $\big(v(x_{i,1}), \dots, v(x_{i,S})\big)$. Finally we will assume that for all $i \neq j$, $\bar{v}_i - \bar{v}_j$ is not a constant vector and that $\phi_\alpha(x)_s = v^{-1}\big(v(x_s + \alpha)\big)$.

G.2.1 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:

$$u(x) = \langle \pi, v(x) \rangle,$$

for some $\pi \in \Delta(S)$. Define $\mathcal{K}_{\text{SEU}}$ as the collection of $\bar{u} \in \mathcal{U}$ that are restrictions of subjective expected utility representations. Then solving (11) with $\mathcal{K} = \mathcal{K}_{\text{SEU}}$ is equivalent to solving:

$$\begin{aligned}
\min_{\bar{u} \in \mathcal{U}} \quad & \big\|(\operatorname{grad} \bar{u}) - Y\big\|_2^2 \\
\text{subject to} \quad & \bar{u}_i = \langle \pi, \bar{v}_i \rangle \qquad \forall i = 1, \dots, K \\
& \langle \pi, \mathbb{1}_S \rangle = 1 \\
& \pi \geq 0.
\end{aligned} \tag{20}$$

Proof. Trivial. $\square$

G.2.2 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 $S = \{1, \ldots, S\}$, and let $\mathfrak{S}_S$ denote the set of permutations on $\{1, \ldots, S\}$. For each $\sigma \in \mathfrak{S}_S$, define:

$$C_\sigma = \{x \in \mathbb{R}^S : x_{\sigma(1)} \geq x_{\sigma(2)} \geq \cdots \geq x_{\sigma(S)}\}. \tag{21}$$

The cones $\{C_\sigma\}_{\sigma \in \mathfrak{S}_S}$ cover $\mathbb{R}^S$. Note that if a functional $w : \mathbb{R}^S \to \mathbb{R}$ 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:

$$w(x) = \int_S x \, dP^\sigma,$$

where, for all $i = 1, \ldots, S$, the probability measure P^σ is defined by:

$$P^\sigma(\sigma(i)) = \nu\big(\{\sigma(1), \sigma(2), \ldots, \sigma(i)\}\big) - \nu\big(\{\sigma(1), \sigma(2), \ldots, \sigma(i-1)\}\big). \tag{22}$$

See Ghirardato et al. (2004) for more discussion. Finally, for notational simplicity, define the shorthand A_i^σ for the set $\{\sigma(1), \sigma(2), \ldots, \sigma(i)\}$.

We say that $u : X \to \mathbb{R}$ is a Choquet expected utility (CEU) representation if:

$$u(x) = \int_S v(x) \, d\nu,$$

where ν is a capacity and the integral denotes Choquet integration. Define $\mathcal{K}_{CEU}$ as the collection of $u \in \mathcal{U}$ that are restrictions of CEU representations. Then solving (11) with $\mathcal{K} = \mathcal{K}_{\mathrm{CEU}}$ is equivalent to solving:

$$
\begin{aligned}
\min_{\bar{u}\in\mathcal{U}} \quad & \left\|(\operatorname{grad}\bar{u}) - Y\right\|_2^2 \\
\text{subject to} \quad & \bar{u}_i = \langle P^\sigma, \bar{v}_i\rangle && \forall \sigma\in\mathfrak{S}_S, \forall i = 1,\ldots,K \text{ s.t. } \tilde{u}^i \in C^\sigma \\
& P^\sigma_{\sigma(j)} = \nu_{A^\sigma_j} - \nu_{A^\sigma_{j-1}} && \forall \sigma\in\mathfrak{S}_S, \forall j = 1,\ldots,S \\
& \nu_A \le \nu_B && \forall A, B \in 2^S \text{ s.t. } A\subseteq B \\
& \nu_\varnothing = 0 \\
& \nu_S = 1
\end{aligned}
\tag{23}
$$

Proof. Suppose u is a CEU representation. Then it corresponds to integration against some capacity ν which by definition then satisfies the last three constraints of (23). From the discussion, e.g., in Ghirardato et al. (2004) (see, in particular, Example 17), each $\bar{v}_i$ belongs to at least one C_σ cone, and restricted to each, u simply amounts to integration (i.e. a dot product) of $\bar{v}_i$ with the measure P^σ. Hence every CEU functional corresponds to a solution to (23). Conversely, it follows trivially that every solution to (23) defines a CEU functional. □

G.2.3 Convex Choquet Expected Utility

A capacity $\nu : 2^S \to \mathbb{R}$ is said to be a convex, if, for all $A, B \subseteq S$:

$$\nu(A) + \nu(B) \le \nu(A\cap B) + \nu(A\cup B).$$

A map $U : X \to \mathbb{R}$ is said to be a convex Choquet expected utility (CCEU) representation if it is of the form:

$$u(x) = \int_S v(x)\, d\nu,$$

for some convex capacity ν. Define $\mathcal{K}_{CCEU}$ as the collection of $u\in\mathcal{U}$ that are restrictions of CCEU representations. Then, solving (11) with $\mathcal{K} = \mathcal{K}_{\text{CCEU}}$ is equivalent to solving:

$$
\begin{aligned}
\min_{\bar{u}\in\mathcal{U}} \quad & \left\|(\operatorname{grad}\bar{u}) - Y\right\|_2^2 \\
\text{subject to} \quad & \bar{u}_i = \langle P^\sigma, \bar{v}_i\rangle && \forall\sigma\in\mathfrak{S}_S, \forall i = 1,\ldots,K \text{ s.t. } \tilde{u}^i \in C^\sigma \\
& P^\sigma_{\sigma(j)} = \nu_{A^\sigma_j} - \nu_{A^\sigma_{j-1}} && \forall\sigma\in\mathfrak{S}_S, \forall j = 1,\ldots,S \\
& \nu_A \le \nu_B && \forall A, B \in 2^S \text{ s.t. } A\subseteq B \\
& \nu_A + \nu_B \le \nu_{A\cup B} + \nu_{A\cap B} && \forall A, B\in 2^S \\
& \nu_\varnothing = 0 \\
& \nu_S = 1
\end{aligned} \tag{24}
$$

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

G.2.4 Maxmin Expected Utility

A map $u : X \to \mathbb{R}$ is said to be a maxmin expected utility (MEU) representation if it is of the form:

$$U(x) = \min_{\pi\in P}\langle\pi, v(x)\rangle,$$

for some compact, convex belief set $P \subseteq \Delta(S)$. Define $\mathcal{K}_{\text{MEU}}$ as the collection of $\bar{u}\in\mathcal{U}$ that are restrictions of MEU representations. Then solving (11) with $\mathcal{K} = \mathcal{K}_{\text{MEU}}$ is equivalent to solving:

$$
\begin{aligned}
\min_{\bar{u}\in\mathcal{U}} \quad & \left\|\operatorname{grad}\bar{u} - Y\right\|_2^2 \\
\text{subject to} \quad & \bar{u}_i = \langle\pi_i, \bar{v}_i\rangle && \forall i = 1,\ldots,K \\
& \langle\pi_i, \bar{v}_i\rangle \le \langle\pi_j, \bar{v}_i\rangle && \forall i, j = 1,\ldots,K \\
& \langle\pi_i, \mathbb{1}_S\rangle = 1 && \forall i = 1,\ldots,K \\
& \pi_i \ge 0 && \forall i = 1,\ldots,K,
\end{aligned} \tag{25}
$$

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

Proof. Suppose first that $\bar{u}\in\mathcal{K}$ is the restriction to $\{\bar{v}_1,\ldots,\bar{v}_K\}$ of some MEU functional w. For $i = 1,\ldots,K$, let $\pi_i \in \partial w(\bar{v}_i)$ denote an arbitrarily selection of supergradients of w. As $w(0) = 0$, by homogeneity, $w(\bar{v}_i) = \langle \pi_i, \bar{v}_i \rangle$ for all $i = 1, \ldots, K$. Furthermore, for all $x \in \mathbb{R}^S$ and all $v_i \in \mathcal{V}$:

$$\begin{aligned} w(x) &\leq w(\bar{v}_i) + \langle \pi_i, x - \bar{v}_i \rangle \\ &= \langle \pi_i, \bar{v}_i \rangle + \langle \pi_i, x - \bar{v}_i \rangle \\ &= \langle \pi_i, x \rangle, \end{aligned}$$

hence for all $\bar{v}_j \in v(\mathcal{V})$, $\langle \pi_j, \bar{v}_j \rangle \leq \langle \pi_i, \bar{v}_j \rangle$. As w is increasing, for each i, $\pi_i \geq 0$. Let $\alpha \in \mathbb{R}$. Since w is translation-invariant, for all $\bar{v}_i$:

$$w(\bar{v}_i + \alpha \mathbb{1}_S) \leq w(\bar{v}_i) + \langle \pi_i, \alpha \mathbb{1}_S \rangle$$

hence

$$w(\bar{v}_i) + \alpha \leq w(\bar{v}_i) + \langle \pi_i, \alpha \mathbb{1}_S \rangle$$

and

$$\alpha \leq \alpha \langle \pi_i, \mathbb{1}_S \rangle. \tag{26}$$

If $\alpha > 0$, $1 \leq \langle \pi, \mathbb{1}_S \rangle$, and if $\alpha < 0$, $1 \geq \langle \pi, \mathbb{1}_S \rangle$. Since (28) holds for all $\alpha \in \mathbb{R}$, we obtain $\langle \pi_i, \mathbb{1}_S \rangle = 1$.

Suppose now that for some collection $\pi_1, \ldots, \pi_K \in \Delta(S)$, we have a vector $\bar{u} \in \mathcal{U}$ satisfying (i) $\bar{u}_i = \langle \pi_i, \bar{v}_i \rangle$ and (ii) $\langle \pi_i, \bar{v}_i \rangle \leq \langle \pi_j, \bar{v}_i \rangle$. Define

$$\hat{u}(x) = \min_{i \in \{1, \ldots, K\}} \langle \pi_i, v(x) \rangle = \min_{\pi \in \mathrm{co}\{\pi_1, \ldots, \pi_K\}} \langle \pi, v(x) \rangle.$$

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

G.2.5 Variational Preferences

A map $u : X \to \mathbb{R}$ is said to be a variational preferences representation if it is of the form:

$$U(x) = \min_{\pi \in \Delta(S)} \langle \pi, v(x) \rangle + c(\pi)$$

for some cost function $c : \Delta(S) \to [0, \infty]$ that is (i) convex, (ii) lower semicontinuous, and (iii) grounded, i.e. attains 0 for some $\pi \in \Delta(S)$. Define $\mathcal{K}_{\mathrm{VAR}}$ as the collection of $\bar{u} \in \mathcal{U}$ that are restrictions of variational utility representations.

We assume that $x_K \in \mathcal{V}$ is the zero act, and hence $\bar{v}_K = 0$. Then solving (11) with $\mathcal{K} = \mathcal{K}_{\mathrm{VAR}}$ is equivalent to solving:

$$\begin{aligned}
\min_{\bar{u}\in\mathcal{U}} \quad & \left\|\operatorname{grad}\bar{u} - Y\right\|_2^2 & \\
\text{subject to} \quad & \bar{u}_i = \gamma_i + \langle \pi_i, \bar{v}_i\rangle & \forall i = 1,\dots,K \\
& \gamma_i + \langle \pi_i, \bar{v}_i\rangle \le \gamma_j + \langle \pi_j, \bar{v}_i\rangle & \forall i,j = 1,\dots,K \\
& \langle \pi_i, \mathbb{1}_S\rangle = 1 & \forall i = 1,\dots,K \\
& \pi_i \ge 0 & \forall i = 1,\dots,K, \\
& \gamma_K = 0, &
\end{aligned} \tag{27}$$

for $\pi_1 \dots, \pi_K \in \mathbb{R}^S$ and $\gamma_1, \dots, \gamma_K \in \mathbb{R}$.

Proof. Suppose first that $\bar{u} \in \mathcal{K}$ is the restriction to $\mathcal{V}$ of some risk-neutral variational utility functional w. For $i = 1,\dots,K$, let $\pi_i \in \partial w(\bar{v}_i)$ be an arbitrary selection of supergradients of w, one at each $\bar{v}_i$. For all $i = 1,\dots K$, let:

$$\gamma_i = \bar{u}_i - \langle \pi_i, \bar{v}_i\rangle.$$

Then, for all i, by construction $\bar{u}_i = \gamma_i + \langle \pi_i, \bar{v}_i\rangle$ and $\gamma_K = 0$ hence so too is $\bar{u}_K$. Moreover, for all $x \in \mathbb{R}^S$ and all $\bar{v}^j$:

$$\begin{aligned}
w(x) &\le w(\bar{v}_j) + \langle \pi_j, x - \bar{v}_j\rangle \\
&= \gamma_j + \langle \pi_j, \bar{v}_j\rangle + \langle \pi_j, x - \bar{v}_j\rangle \\
&= \gamma_j + \langle \pi_j, x\rangle,
\end{aligned}$$

hence in particular, for all $\bar{v}_i$, $\gamma_i + \langle \pi_i, \bar{v}_i\rangle \le \gamma_j + \langle \pi_j, \bar{v}_i\rangle$. As w is increasing, for each i, $\pi_i \ge 0$. Let $\alpha \in \mathbb{R}$. Since w is translation-invariant, for all $\bar{v}_i$:

$$w(\bar{v}_i + \alpha\mathbb{1}_S) \le w(\bar{v}_i) + \langle \pi_i, \alpha\mathbb{1}_S\rangle$$

hence

$$w(\bar{v}_i) + \alpha \le w(\bar{v}_i) + \langle \pi_i, \alpha\mathbb{1}_S\rangle$$

and

$$\alpha \leq \alpha\langle \pi_i, \mathbb{1}_S\rangle. \tag{28}$$

If $\alpha > 0$, $1 \leq \langle \pi, \mathbb{1}_S\rangle$, and if $\alpha < 0$, $1 \geq \langle \pi, \mathbb{1}_S\rangle$. Since (28) holds for all $\alpha \in \mathbb{R}$, we obtain $\langle \pi_i, \mathbb{1}_S\rangle = 1$.

Suppose now that for some collection $\pi_1, \dots, \pi_K \in \Delta(S)$ and $\gamma_1, \dots, \gamma_K \in \mathbb{R}$ with $\gamma_K = 0$, we have a vector $\bar{u} \in \mathcal{U}$ satisfying (i) $\bar{u}_i = \gamma_i + \langle \pi_i, \bar{v}_i\rangle$, and (ii) $\gamma_i + \langle \pi_i, \bar{v}_i\rangle \leq \gamma_j + \langle \pi_j, \bar{v}_i\rangle$. Define

$$\hat{u}(x) = \min_{i\in\{1,\dots,K\}} \gamma_i + \langle \pi_i, v(x)\rangle$$

By construction, $\bar{u}_i = \hat{u}(\bar{v}_i)$ and $\hat{u}$ is a (i) translation invariant, (ii) concave, (iii) increasing, and (iv) normalized hence, by the results of Maccheroni et al. (2006), corresponds to a variational utility representation. □

G.2.6 Dual Self Expected Utility

A map $u : X \to \mathbb{R}$ is said to be a dual-self utility representation if it is of the form:

$$u(x) = \max_{P\in\mathbb{P}^*} \min_{\pi\in P} \langle \pi, v(x)\rangle,$$

where $\mathbb{P}^*$ is a Hausdorff-compact collection of compact, convex subsets of $\Delta(S)$.

Let $(\mathcal{V}, \mathcal{E})$ denote an experiment, where $\bar{v}_K = 0$. Let $\mathcal{K}_{\mathrm{DS}}$ denote the collection of $\bar{u} \in \mathcal{U}$ that are restrictions of dual-self utility representations. Then solving (11) with $\mathcal{K} = \mathcal{K}_{\mathrm{DS}}$ is equivalent to solving:

$$\begin{aligned}
\min_{\bar{u}\in\mathcal{U}} \quad & \left\|\operatorname{grad} \bar{u} - Y\right\|_2^2 \\
\text{subject to} \quad & u_i = \langle \pi_{ii}, \tilde{u}^i\rangle & \forall i = 1, \dots, K \\
& \langle \pi_{ii}, \bar{v}_i\rangle \leq \langle \pi_{ij}, \bar{v}_i\rangle & \forall i, j = 1, \dots, K \\
& \langle \pi_{ji}, \bar{v}_i\rangle \leq \langle \pi_{ii}, \bar{v}_i\rangle & \forall i, j = 1, \dots, K \\
& \langle \pi_{ij}, \mathbb{1}_S\rangle = 1 & \forall i, j = 1, \dots, K \\
& \pi_{ij} \geq 0 & \forall i, j = 1, \dots, K,
\end{aligned} \tag{29}$$

for $\bar{u} \in \mathcal{U}$, $\{\pi_{ij}\}_{i,j=1}^K \in \mathbb{R}^S$.

Proof. Suppose, first, that $\bar{u}, \{\pi_{ij}\}_{i,j=1}^{K}$ is a solution to (29). Define, for each $i = 1, \ldots, K$, the set $P_i = \mathrm{co}\{\pi_{i,1}, \ldots, \pi_{i,K}\}$. Clearly $P_i \subseteq \Delta(S)$ for each i. Let $\mathbb{P}^* = \{P_i\}_{i=1}^{K}$. We claim that:

$$\hat{u}(x) = \max_{P \in \mathbb{P}^*} \min_{\pi \in P} \langle \pi, x \rangle$$

defines a DSEU functional w whose restriction to $v(\mathcal{V})$ is precisely $\bar{u}$. Firstly, as $\langle \pi_{ii}, \bar{v}_i \rangle \leq \langle \pi_{ij}, \bar{v}_i \rangle$ for all $j = 1, \ldots, K$, it follows that:

$$\bar{u}_i = \langle \pi_{ii}, \bar{v}_i \rangle = \min_{\pi \in P_i} \langle \pi, \bar{v}_i \rangle.$$

But, for all $j = 1, \ldots, K$ we have $\langle \pi_{ji}, \bar{v}_i \rangle \leq \bar{u}_i$, hence:

$$\bar{u}_i \geq \langle \pi_{ji}, \bar{v}_i \rangle \geq \min_{\pi \in P_j} \langle \pi, \bar{v}_i \rangle,$$

as $\pi_{ji} \in P_j$. Thus:

$$\begin{aligned} w(\tilde{u}^i) &\equiv \max_{P \in \mathbb{P}^*} \min_{\pi \in P} \langle \pi, \bar{v}_i \rangle \\ &= \min_{\pi \in P_i} \langle \pi, \bar{v}_i \rangle \\ &= \langle \pi_{ii}, \bar{v}_i \rangle \\ &= \bar{u}_i. \end{aligned}$$

Conversely, suppose now that $w(x) = \max_{P \in \mathbb{P}^*} \min_{\pi \in P} \langle \pi, x \rangle$ 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:

$$w(\bar{v}_i) = \min_{\pi \in P_i} \langle \pi, \bar{v}_i \rangle,$$

and let $\pi_{ii} \in P_i$ be any minimizer of the right-hand side.[68] Define, for each $i = 1, \ldots, K$, the utility value $\bar{u}_i = \langle \pi_{ii}, \bar{v}_i \rangle$. Since P_j is an 'active' belief set at $\bar{v}_j$ for each $j \neq i$, there exists, for each j, some $\pi_{ij} \in P_i$ such that $\langle \pi_{ij}, \bar{v}_j \rangle \leq \bar{u}_j$. Since each $\pi_{ij} \in P_i$, then $\bar{u}_i \leq \langle \pi_{ij}, \bar{v}_i \rangle$ for each i. Then, as clearly every $\pi_{ij} \in \Delta(S)$, the collection $\bar{u}, \{\pi_{ij}\}_{i,j=1}^{K}$ is a solution to (29), as required. □

[68] 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.

G.2.7 Dual-Self Variational Utility

A map $w : X \to \mathbb{R}$ 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) ϕ-additive, (ii) monotone, (iii) normalized, i.e. $U(\mathbb{1}_S) = 1$, see Supplementary Appendix to Chandrasekher et al. (2022).

Let $(\mathcal{V}, \mathcal{E})$ denote an experiment. Let $\mathcal{K}_{\text{DSV}}$ denote the collection of $\bar{u} \in \mathcal{U}$ that are restrictions of dual-self variational utility representations. Then solving (11) with $\mathcal{K} = \mathcal{K}_{\text{DSV}}$ is equivalent to solving:

$$\begin{aligned} \min_{\bar{u} \in \mathcal{U}} \quad & \left\| \operatorname{grad} \bar{u} - \bar{Y} \right\|_2^2 \\ \text{subject to} \quad & \bar{u}_i \geq \bar{u}_j \qquad\qquad \forall i, j \text{ s.t. } \bar{v}_i \geq \bar{v}_j \\ & u_K = 0, \end{aligned} \tag{30}$$

where $\bar{v}_i \geq \bar{v}_j$ is understood in the product order on $\mathbb{R}^S$.

Proof. Firstly, suppose w is a dual-self variational functional. Then it clearly is monotone, hence $\bar{v}_i \geq \bar{v}_j$ implies $w(\bar{v}_i) \geq w(\bar{v}_j)$. Moreover,

$$w(\mathbb{1}_S) = w\big(0 + \mathbb{1}_S\big) = w(0) + 1,$$

hence U is normalized if and only if $w(0) = 0$. Thus clearly letting $\bar{u}_i = w(\bar{v}_i)$ satisfies the constraints of (30).

Conversely, suppose $\bar{u}$ is a solution to (30). In light of the characterization provided in Chandrasekher et al. (2022), it suffices to prove there exists an translation-invariant and monotone extension from $\{\bar{v}_1, \ldots, \bar{v}_K\}$ to $\mathbb{R}^S$.[69]

[69]Normalization holds for any ϕ-additive extension, as $\bar{v}_K = 0$.

However, note that by hypothesis, no pair $\bar{v}_i$ and $\bar{v}_j$ lie on the same translate of the diagonal, thus $\bar{u}$ is trivially translation-invariant and by definition monotone on $\{\bar{v}_1, \ldots, \bar{v}_K\}$. Hence by Theorem 1 of Cerreia-Vioglio et al. (2014), there exists a ϕ-additive, monotone, and normalized extension of $\bar{u}$, and hence by Chandrasekher et al. (2022) this corresponds to some dual-self variational utility functional. $\square$

References

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

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

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and A. Rustichini (2014). Niveloids and their extensions: Risk measures on small domains. *Journal of Mathematical Analysis and Applications 413*(1), 343–360.

Chandrasekher, M., M. Frick, R. Iijima, and Y. Le Yaouanq (2022). Dual-self representations of ambiguity preferences. *Econometrica 90*(3), 1029–1061.

Debreu, G. (1964). Continuity properties of paretian utility. *International Economic Review 5*(3), 285–293.

Ghirardato, P., F. Maccheroni, and M. Marinacci (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory 118*(2), 133–173.

Hiriart-Urruty, J.-B. and C. Lemaréchal (2004). *Fundamentals of convex analysis*. Springer Science & Business Media.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica 74*(6), 1447–1498.

Munkres, J. R. (1974). *Topology; a First Course*. Prentice-Hall.

Ok, E. (2011). *Elements of order theory*. Mimeo, New York University.