

RECOVERING PREFERENCES FROM FINITE DATA

CHRISTOPHER P. CHAMBERS, FEDERICO ECHENIQUE,
AND NICOLAS S. LAMBERT

ABSTRACT. We study preferences recovered from finite choice experiments and provide sufficient conditions for convergence to a unique underlying “true” preference. Our conditions are weak, and therefore valid in a wide range of economic environments. We develop applications to expected utility theory, choice over consumption bundles, menu choice and intertemporal consumption. Our framework unifies the revealed preference tradition with models that allow for errors.

1. INTRODUCTION

This paper concerns the nonparametric recoverability of preferences from finite choice data.

We imagine an experimenter, Alice, offering a sequence of binary choice problems to a subject, Bob. For each choice problem, Bob is presented with a pair of alternatives and is asked to choose one (for example, alternatives could be lotteries over a collection of prizes). Alice wants to ensure that, if she observes Bob on sufficiently many choice problems, then she can recover his preference over the entire set of alternatives to an arbitrary degree of precision.

(Chambers) DEPARTMENT OF ECONOMICS, GEORGETOWN UNIVERSITY

(Echenique) DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES, CALIFORNIA INSTITUTE OF TECHNOLOGY

(Lambert) GRADUATE SCHOOL OF BUSINESS, STANFORD UNIVERSITY

Date: First version: April 10, 2017; This version: November 14, 2019.

An early version circulated under the title “Preference Identification.” We are grateful to Pathikrit Basu for finding a mistake in a prior draft of the paper, and for his insightful suggestions. We thank seminar participants at many institutions, and the audiences of conferences and workshops at Paris School of Economics, UC Berkeley, University of Pennsylvania, University of Warwick and University of York. Echenique thanks the National Science Foundation for its support through the grants SES 1558757 and CNS 1518941, and the Simons Institute at UC Berkeley for its hospitality while part of the paper was written. Lambert gratefully acknowledges the financial support and hospitality of Microsoft Research New York and the Yale University Cowles Foundation.

In this paper, we provide general conditions under which finite choice data can be used to approximate a subject’s “true” preference. The goal is to have conditions that are easy to check and apply broadly.

Our approach is two-pronged. Our first model is anchored in the classical revealed preference tradition, whereby Alice seeks to rationalize the choice data exactly. The model is deterministic. Alice designs a fixed experiment, and hypothesizes that Bob chooses perfectly in accordance with his preference. Our second model is statistical. The selection of experiments is random, and Alice supposes that Bob’s choices are either observed with some error, or made with error. In both models, we provide conditions for the underlying data-generating preference to be learned in the limit. These conditions concern the experimental design, and the preference environment being considered.

The main substantive condition is the *local strictness* of preferences, a property first described by Border and Segal (1994). Local strictness generalizes the familiar notion of local nonsatiation. A locally strict preference means that whenever x is at least as good as y , there are alternatives x' and y' , near x and y respectively, and for which x' is strictly better than y' . We prove that, together with technical conditions, local strictness ensures the convergence of any sequence of rationalizing preferences to the unique underlying preference governing the subject’s choices as the number of observations grows. In the statistical model, we introduce an estimator based on minimizing the Kemeny distance to the observed choices. Again imposing local strictness, we prove that this estimator is consistent, and provide general convergence rates.

The usefulness of our results is illustrated with applications to expected-utility theory and other environments where preferences are defined from utility functions; to monotone preferences in consumption theory, preferences over menus, and exponential discounting in intertemporal choice. In all these applications, we show how large finite experiments can approximate a preference of the appropriate kind. In our statistical model, we provide explicit convergence rates for some of the applications.

Our approach is based on the choice-theoretic notion of partial observability, as in Afriat (1967): when Bob is presented with a choice, he must select no more than one alternative. Therefore, choosing one alternative over another

does not preclude the possibility that Bob would have been equally happy with the other alternative. In Afriat’s case, partial observability has wide-ranging implications, famously rendering the concavity of utility nontestable.¹ In our framework, preferences can still be fully learned in spite of their partial observability.

The results in this paper allow for very general sets of preferences, which translate into a “model-free” approach. If, say, Alice is interested in exponential discounting, then she can estimate a preference without the need to impose this assumption on the data. If Bob is indeed discounting exponentially, then the preference estimates are guaranteed to converge to a preference that follows exponential discounting. And if the preference estimates do not converge to such a preference, then Alice may conclude that the exponential discounting hypothesis is incorrect. She can then evaluate the degree to which Bob’s preference diverges from exponential discounting. The model-free aspect is also present in the statistical model, by allowing Alice to be relatively agnostic about how the alternatives presented to subjects are sampled, and how subjects are assumed to make mistakes.

Overall, our framework combines the elements of different traditions in economic modeling: the nonparametric approach and the finite amount of data in revealed preference analysis, the pairwise comparisons in decision theory and experiments (both online and laboratory), and the source of random errors in empirical research and econometrics.

The paper proceeds as follows. The remainder of this section reviews related works. Section 2 describes the model. Section 3 provides the main results. Sections 4 and 5 put these results to work in several economic environments. Section 6 concludes with a discussion. Proofs are relegated to the appendices.

1.1. Related literature. The rich literature on revealed-preference theory has been primarily concerned with the question of whether observed behavior conforms with standard models in economic theory. The workhorse of this literature, Afriat’s theorem (Afriat, 1967; Diewert, 1973; Varian, 1982), works with the classical model of consumer demand with linear budgets and finite

¹Chambers, Echenique, and Shmaya (2014) provide a discussion of what partial observability entails.

data, and has been expanded in many directions.² This line of research focuses for the most part on constructing revealed preference tests, and discussions of preference recoverability (as, for instance, in Varian, 1982, or more recently Cherchye, De Rock, and Vermeulen, 2011) deal with bounding the sets of possible rationalizing preferences, not on uniquely recovering preferences.

Closer to our work, in the context of consumer demand with linear budgets, Mas-Colell (1978) introduces an “income-Lipschitz” condition and shows that, under this and a boundary condition, any sequence of preferences that rationalizes a sufficiently rich sequence of observations converges to the unique preference that rationalizes the entire demand function. Forges and Minelli (2009) derive the analog of Mas-Colell’s results for nonlinear budget sets. In a model of dynamic asset markets, Kübler and Polemarchakis (2017) derive conditions that permit the identification of utilities and beliefs of a subjective expected utility maximizer, and show that preference estimates from finite data converge to the unique underlying preference as the number of observations grows large. Polemarchakis, Selden, and Song (2017) give conditions under which the identification of preferences is possible, and demonstrate the convergence of preferences estimated from finite data to the unique underlying preference. As in these works, our paper provides conditions for the convergence of preferences that rationalize finite data to the underlying data-generating preference; where we differ is that we focus on data from pairwise choices instead of choice from budgets, allowing us to consider environments beyond choice over commodities, and that we provide a general sufficient condition on the class of preferences under consideration. Because we abstract away from specific economic environments, our results are applicable across different domains: choice of menus, under uncertainty, of intertemporal streams, lotteries, or consumption bundles.

Experimentalists and decision theorists have an obvious interest in preference recovery from pairwise choices, but little is known about the behavior

²For example, Reny (2015) extends Afriat’s theorem to arbitrary, infinite data, Matzkin (1991) and Forges and Minelli (2009) work with nonlinear budget sets, Chavas and Cox (1993) and Nishimura, Ok, and Quah (2017) work with general choice problems. Some extensions were also developed for multiperson equilibrium models, as in Brown and Matzkin (1996) and Carvajal, Deb, Fenske, and Quah (2013), for example.

of preference estimates from finite data. Decision theory papers often include a discussion of identification, but the exercise presumes access to very rich data, in the form of the agent’s full preference relation. In demand theory, there are many studies devoted to the problem of identification—known as the integrability problem—assuming access to a demand function defined on all prices. Matzkin (2006) considers economy-wide data, and uses equilibrium as a means to identify consumers’ utilities. In recent work, Gorno (2019) studies the general problem of identification under partial observability. Gorno provides conditions on decision problems and sets of admissible preferences to ensure identification. We share with Gorno a concern for partial observability, but our research diverges from his, and from other studies of identification and integrability, in that we work with finite data and their large-sample behavior.

One stream of literature combines nonparametric econometric methods with revealed preference theory. In demand analysis, Blundell, Browning, and Crawford (2008) design a statistical test for the revealed preference conditions to be satisfied. Observing that demand responses to price changes can be represented by a set of moment inequalities, they appeal to results on moment inequality estimators by Manski (2003), Chernozhukov, Hong, and Tamer (2007) and Andrews and Guggenberger (2009). In our environments, however, these results on partial identification do not apply. Halevy, Persitz, and Zrill (2018) develop a method for estimating parametric models by minimizing the incompatibility of choice behavior with the proposed model, this is in the same spirit as our Kemeny-distance estimator, but the analysis is quite different. More closely related to our paper, Matzkin (2003) and Blundell, Kristensen, and Matzkin (2010) consider identification in an econometric model of stochastic demand data (see Matzkin, 2007, for a general discussion). We differ from all these papers in our focus on binary choice, in that we look for sufficient conditions on general classes preferences, trying to unify the revealed preference and statistical models, and in how errors are introduced. More recently, Basu and Echenique (2018) investigate the learnability of four standard models of choice under uncertainty using the notion of PAC learning from computational learning theory. Basu (2019) applies several other measures of model complexity to a study of stochastic choice.

Theorem 2 in Section 3 establishes the consistency of what we call the Kemeny-minimizing estimator, which is an M-estimator (Amemiya, 1985; Newey and McFadden, 1994), while Theorem 3 establishes convergence rates. Our result is obtained as a consequence of the canonical consistency of M -estimators: the structure in our model allows us to naturally derive the conditions for consistency without having to make additional compactness and equicontinuity assumptions. See Section 6.3 for a detailed discussion. Such assumptions are particularly challenging in nonparametric estimation like ours.

Finally, a literature in political science (Poole and Rosenthal, 1985, Jackman, 2001, Clinton, Jackman, and Rivers, 2004, are seminal) focuses on binary choice data (roll-call votes), but considers specific parametric models of spatial voting, and uses Bayesian methods for the most part. Our results are broadly applicable to the same data as in this literature, but using very different methods.

2. MODEL

The model features an experimenter, Alice, and a subject, Bob. Bob has preferences over a set of alternatives X , which is a topological space. Alice would like to learn, or recover, Bob's preferences through the device of a choice experiment.³

By *preference relation* or simply *preference* we mean a binary relation \succeq over X that is continuous and complete.⁴ In formal terms, \succeq is the set of pairs $(x, y) \in X \times X$ such that x is at least weakly preferred to y , written $x \succeq y$. Associated to any given preference \succeq are its asymmetric part \succ (strict

³Such experiments, done on a large scale, with large sample size, include for example von Gaudecker, van Soest, and Wengstrom (2011), Chapman, Dean, Ortoleva, Snowberg, and Camerer (2018) and Falk, Becker, Dohmen, Enke, Huffman, and Sunde (2018). Alternatively, we may think of Alice as a researcher, and Bob an individual she has observed in the field. For example, Bob could be a congressman who votes among pairs of competing bills (Poole and Rosenthal, 1985).

⁴Completeness means that for all pairs of alternatives (x, y) , $x \succeq y$ or $y \succeq x$. Continuity means that \succeq as a subset of the product space $X \times X$ is closed; more intuitively, if x is not preferred to y , then x' is also not preferred to y' for all pairs (x', y') in the vicinity of (x, y) . Completeness is standard and continuity is a necessary regularity condition, without it, no meaningful inferences can be made with any finite amount of data.

preference) and its symmetric part \sim (indifference), that is, $x \succ y$ means that $x \succeq y$ but $y \not\succeq x$, while $x \sim y$ indicates that both $x \succeq y$ and $y \succeq x$.

Alice's goal is to infer Bob's preference from his behavior. It is clear that she must somehow constrain, or discipline, the set of preferences being considered. With partial observability, it is very easy to find a preference that rationalizes empirical data. For example, complete indifference rationalizes any observed behavior. Throughout the paper, \mathcal{P} denotes the class of preferences being considered; we think of \mathcal{P} as a set that embraces the possible preferences the subject may have. We refer to a pair (X, \mathcal{P}) as a *preference environment*.

Alice collects information about Bob through a finite experiment, in which Bob confronts a fixed number of *binary choice problems*. In each binary choice problem, Bob is presented with an unordered pair of alternatives, and is asked to choose exactly one of the two alternatives. An *experiment of size n* is represented by a collection $\Sigma_n = \{B_1, \dots, B_n\}$, where $B_k = \{x_k, y_k\}$ is an unordered pair of alternatives that captures a binary choice problem. We stress that an experiment only generates a finite amount of data.

To study the limiting properties of estimated preferences in large samples, we consider not just one experiment, but a set of growing experiments indexed by their size, of the form $\{\Sigma_1, \Sigma_2, \dots\}$, where Σ_n is an experiment of size n and $\Sigma_n \subset \Sigma_{n+1}$. In the sequel, Σ_n always denotes an experiment of size n , and the inclusion property $\Sigma_n \subset \Sigma_{n+1}$ is implicitly assumed. We use the abbreviated notation $\{\Sigma_n\}$ to denote a set of (growing) experiments.

The behavior of a subject who decides over binary choice problems is encoded in a single-valued choice function c that maps unordered pairs of alternatives to alternatives. It records, for every possible binary choice problem $\{x, y\} \subset X$, the alternative $c(\{x, y\}) \in \{x, y\}$ that is chosen. We refer to c as the *choice function*, and impose no a priori restrictions on choice functions.

We follow two traditions in economic modeling. The first tradition is classical revealed preference theory, in which the choice problems of an experiment are selected arbitrarily, and the experimenter seeks to exactly rationalize observed behavior, as in the classical works of Afriat (1967), Mas-Colell (1978) and Varian (1982). In this theory, Bob is assumed to possess a preference and to make choices that comply perfectly with this preference. Alice looks for a

preference that fits exactly the empirical observations. However, this theory does not account for errors, while empirical work often tries to accommodate errors.

The second tradition tackles this problem by imposing a statistical model on the subject's choices. The subject is presented with choices drawn at random, either because the experimental design is explicitly random (as, for example, in Ahn, Choi, Gale, and Kariv, 2014, Choi, Kariv, Müller, and Silverman, 2014, Carvalho, Meier, and Wang, 2016, or Carvalho and Silverman, 2019), or because the experimenter uses field data in which she has no control over the problems the subject faces. In this theory, Alice continues to assume that Bob has an underlying preference, but she allows for his behavior to deviate from his preferences. Alice looks for a preference that fits the best the observed behavior.⁵

2.1. Revealed preference models. In a revealed preference model, experiments are designed arbitrarily by the experimenter. The primitives are the preference environment (X, \mathcal{P}) , and the set of experiments $\{\Sigma_n\}$. We refer to this model by the triple $(X, \mathcal{P}, \{\Sigma_n\})$.

Recall that when presented with a pair of alternatives $\{x, y\}$, Bob is asked to choose *between* x and y —he cannot choose both. In the language of Chambers, Echenique, and Shmaya (2014), our model of choice features partial observability, as in the original work of Afriat (1967).⁶ With partial observability, the appropriate concept of rationalization is weak rationalization. Given a choice function c describing the subject's behavior, and given an experiment Σ_n , we say that a preference \succeq *weakly rationalizes* the observed choices on Σ_n , or simply *rationalizes* the observed choices on Σ_n , if the experiment outcomes are compatible with the subject's preference: for every $\{x, y\} \in \Sigma_n$, $c(\{x, y\}) \succeq x$ and $c(\{x, y\}) \succeq y$. Similarly, we say that \succeq *rationalizes* the choice function

⁵See also Grant, Kline, Meneghel, Quiggin, and Tourky (2016) for a general study of experimental designs that are tolerant to small deviations in the subject's perception of the experiments.

⁶The tradition in revealed preference theory (and in studies of integrability) prior to Afriat was to exactly rationalize a demand function. In Afriat's model, the observed choices are contained in the rationalizing demand, and in consequence concavity of utility is not testable. See Chambers, Echenique, and Shmaya (2014) for a detailed discussion and exploration of the consequences of partial observability.

c if, for every $x, y \in X$, $c(\{x, y\}) \succeq x$ and $c(\{x, y\}) \succeq y$. Hence, weak rationalization does not allow for the subject to choose in contradiction with his preference, but allows for the subject not to reveal the totality of what his preference implies.

2.2. Statistical preference models. In a statistical preference model, experiments are composed of randomly-selected choice problems. More precisely, the choice problems $B_1 = \{x_1, y_1\}, \dots, B_n = \{x_n, y_n\}$ that make up the experiment Σ_n are generated by drawing the alternatives x_k, y_k in each B_k at random from X , independently and identically according to some probability measure λ (X is endowed with the usual Borel σ -algebra).⁷ We abuse notation and also use λ to denote the product measure on $X \times X$.

Bob's behavior is guided by his preference, but does not flawlessly obey his preference. Instead, the model integrates statistical errors: in every choice problem where the subject is not indifferent, he may make a mistake by choosing an alternative that is *not* preferred (alternatively, one can attribute these random mistakes to measurement errors). The corresponding choice function is therefore random. It is determined by an *error probability function* $q : \mathcal{P} \times X \times X \rightarrow [0, 1]$ that quantifies the extent to which a subject is prone to making errors.

When a subject whose preference is \succeq confronts the binary choice problem $\{x, y\}$, he chooses x over y with probability $q(\succeq; x, y)$, and chooses y over x with the complementary probability. We assume that if $x \succ y$ then x is more likely to be chosen, that is, $q(\succeq; x, y) > 1/2$. When $x \sim y$ we assume that the subject is equally likely to choose x or y , but under our assumptions the case in which $x \sim y$ will not matter because it will occur with probability zero. We assume that q is measurable in $X \times X$ for a fixed \succeq .

The primitives of a statistical preference model are the preference environment given by X and \mathcal{P} , the probability measure λ according to which alternatives are drawn, and the error probability function q . We refer to this model by the tuple $(X, \mathcal{P}, \lambda, q)$.

⁷Strictly speaking, an experiment Σ_n is now a multiset, to account for the fact that the same binary decision problem may be drawn more than once, even though typically such repetition occurs with probability zero.

3. MAIN RESULTS

This section provides general results on the convergence of preferences. Throughout, we use the following notion of convergence: a sequence of preferences $\{\succeq_n\}_{n \in \mathbf{N}}$ converges to a preference \succeq^* , written $\succeq_n \rightarrow \succeq^*$ for short, when the following two conditions are satisfied:

- (1) For all alternatives x^*, y^* with $x^* \succeq^* y^*$, there exists a sequence of pairs of alternatives $\{(x_n, y_n)\}_{n \in \mathbf{N}}$ converging to (x^*, y^*) such that $x_n \succeq_n y_n$ for all $n \in \mathbf{N}$.
- (2) For all subsequences $\{\succeq_{n_k}\}_{k \in \mathbf{N}}$, and all pairs of alternatives (x^*, y^*) that are the limit of a sequence $\{(x_{n_k}, y_{n_k})\}_{k \in \mathbf{N}}$ satisfying $x_{n_k} \succeq_{n_k} y_{n_k}$ for all k , we have $x^* \succeq^* y^*$.

Under the assumptions we shall impose, these conditions define convergence in the *closed convergence topology*. Throughout, we endow the space of preferences and binary relations with this topology. The closed convergence topology is a common topology for spaces of sets, such as binary relations, and is the standard topology used for spaces of preferences (Kannai, 1970; Hildenbrand, 1970). It is particularly well suited to the concept of partial observability; we discuss the choice of topology in Section 6.2.

Under conditions that are satisfied in our model, the closed convergence topology on the space of preferences is metrizable, making it possible to quantify approximations and speak of convergence rates. *We fix, and denote by ρ , one compatible metric.* In particular, if X is compact and metrizable, then we can choose as ρ the usual Hausdorff metric for the product space $X \times X$. The notion of closed convergence then coincides with the notion of Hausdorff convergence.⁸

3.1. Convergence in revealed preference models. Our first main result states that convergence of rationalizing preference obtains under certain assumptions on the model primitives. Given a revealed preference model $(X, \mathcal{P}, \{\Sigma_n\})$, consider the following assumptions.

⁸Our results allow for X to be only locally compact. In this case, ρ may still be chosen to coincide with the Hausdorff metric on subsets of the product space $X_\infty \times X_\infty$, where X_∞ is the one-point compactification of X together with some metric generating X_∞ . See Aliprantis and Border (2006) for details.

ASSUMPTION 1. X is a locally compact, separable, and completely metrizable space.

Assumption (1) puts a necessary structure on the set of alternatives. It is satisfied in many common economic environments, as we show in Sections 4 and 5.

The next assumption disciplines the class of the preferences being considered. The central property that allows for meaningful preference recovery is local strictness. This property rules out “thick” indifference curves, in the spirit of the local nonsatiation property of consumer theory. Formally, a preference \succeq is *locally strict* if for every $x, y \in X$ with $x \succeq y$, and every neighborhood V of (x, y) in $X \times X$ there exists $(x', y') \in V$ with $x' \succ y'$ (Border and Segal, 1994).

ASSUMPTION 2. \mathcal{P} is a closed set of locally strict preferences.

The requirement that \mathcal{P} be closed may be seen as a minor technical condition, but it is essential. Perhaps surprisingly, without closedness, it is possible to have locally strict preferences that perfectly rationalize the observations of a subject who chooses exactly according to a locally strict preference \succeq^* , and yet convey no information on the unobserved features of \succeq^* —no matter the number of observations, the experimental design, or the underlying preference \succeq^* . We discuss this point in Section 6.4.

Finally, the choice problems in the experiments must be sufficiently many, and sufficiently diverse, so that observed behavior on all these choice problems can effectively probe the subject’s preference. A set of experiments $\{\Sigma_n\}$, with $\Sigma_n = \{B_1, \dots, B_n\}$, is called *exhaustive* when it satisfies the following two properties:

- (1) $\bigcup_{k=1}^{\infty} B_k$ is dense in X .
- (2) For all $x, y \in \bigcup_{k=1}^{\infty} B_k$ with $x \neq y$, there exists k such that $B_k = \{x, y\}$.

The first property imposes that the alternatives that are used in the set of experiments sample the space of alternatives appropriately. The second property states that the experimenter should be able to elicit the subject’s choices over all alternatives used in her experiments. Note that denseness is the only

real constraint: starting from a countable dense set of alternatives, one can always construct an exhaustive set of experiments via routine diagonalization arguments.

ASSUMPTION 3. $\{\Sigma_n\}$ is exhaustive.

The importance of having a dense set of alternatives is clear: without it, the characteristics of the preference remains unobservable on an open set, and for general classes of preferences, knowledge of the preference outside this set does not suffice to infer those unobservable characteristics. With additional discipline on \mathcal{P} , Assumption (3) can be weakened, as we argue in Section 6.1.

The importance of local strictness for our results hinges on the fact that, for an exhaustive set of experiments, and any two distinct locally strict preferences \succeq_A and \succeq_B of two subjects A and B respectively, there always is at least one experiment for which subject A behaves differently from subject B , thereby allowing the experimenter to distinguish between these two preferences. Thus, with local strictness, a false hypothesis will eventually be demonstrated to be false, whereas without it, too many preferences can be consistent with the data. This fact is stated formally in Lemma 1.

Lemma 1. *Consider an exhaustive set of experiments with binary choice problems $\{x_k, y_k\}$, $k \in \mathbf{N}$. Let \succeq be any complete binary relation, and \succeq_A and \succeq_B be locally strict preferences. If, for all k , $x_k \succeq_A y_k$ and $x_k \succeq_B y_k$ whenever $x_k \succeq y_k$, then $\succeq_A = \succeq_B$.*

The proof of Lemma 1 is in Appendix A.

Under the above assumptions, we establish the convergence of rationalizing preference estimates.

Theorem 1. *Suppose the revealed preference model $(X, \mathcal{P}, \{\Sigma_n\})$ meets Assumptions (1)–(3) and c is an arbitrary choice function. If, for every n , the preference $\succeq_n \in \mathcal{P}$ rationalizes the observed choices on Σ_n , then there exists a preference $\succeq^* \in \mathcal{P}$ such that $\succeq_n \rightarrow \succeq^*$. Moreover, the limiting preference is unique: if, for every n , $\succeq'_n \in \mathcal{P}$ rationalizes the observed choices on Σ_n , then the same limit $\succeq'_n \rightarrow \succeq^*$ obtains.*

The proof of Theorem 1 is in Appendix B.

Theorem 1 asserts that if, in each experiment, the data can be rationalized by some preference in the class \mathcal{P} , then there always exists one preference in \mathcal{P} that rationalizes the choices made over *all* the experiments, and most importantly, there exists only one such preference. The observations are exactly as if the subject's choices were guided by this particular preference, which can be obtained as the limit of the rationalizations as experiments grow in size.

In particular, if we postulate the existence of a preference $\succeq^* \in \mathcal{P}$ according to which the subject chooses on any given decision problem, then no matter the selection of the rationalizing preferences, they always converge to \succeq^* .

Corollary 1. *Suppose the revealed preference model $(X, \mathcal{P}, \{\Sigma_n\})$ meets Assumptions (1)–(3) and c is an arbitrary choice function. If the preference $\succeq^* \in \mathcal{P}$ rationalizes c and if, for every n , the preference $\succeq_n \in \mathcal{P}$ rationalizes the observed choices on Σ_n , then $\succeq_n \rightarrow \succeq^*$.*

3.2. Convergence in statistical preference models. When the subject makes mistakes, looking for a preference that perfectly rationalizes his behavior is moot—a rationalizing preference in the class \mathcal{P} may not exist. Instead, we introduce a simple estimator that approximately rationalizes the observed data, based on minimization of the Kemeny distance (Kendall, 1938; Kemeny, 1959).

The estimator results from a two-step procedure. Let us look at an experiment of size n , Σ_n , drawn at random according to the experimental design. Let c be the choice function that captures the choices of the subject, which is also random. First, from the choices observed on Σ_n , a revealed preference relation is constructed that captures these choices exactly. This revealed preference relation, denoted R_n , is defined by $x R_n y$ for all $\{x, y\} \in \Sigma_n$ such that $c(\{x, y\}) = x$ —that is, $x R_n y$ when the subject chooses x in the choice problem $\{x, y\}$. Note that R_n is sparse, as it only conveys information on the alternatives used in Σ_n . Secondly, the estimated preference \succeq_n is chosen to minimize the distance $d_n(\succeq, R_n)$ between the revealed preference relation just defined, and a preference in $\succeq \in \mathcal{P}$;

$$\succeq_n \in \arg \min \{d_n(\succeq, R_n) : \succeq \in \mathcal{P}\}.$$

Distance d_n is taken to be a version of the Kemeny distance, defined by

$$d_n(\succeq, R_n) = \frac{1}{n} |R_n \setminus \succeq|.$$

In words, $d_n(\succeq, R_n)$ averages the number of mistakes made by the subject on Σ_n if his underlying preference is \succeq . We refer to this estimator as the *Kemeny-minimizing estimator*.⁹

Given a statistical preference model $(X, \mathcal{P}, \lambda, q)$, consider the following assumptions.

Assumptions (1) and (2) on the preference environment remain unchanged. In particular, we continue to assume that the preferences under consideration are locally strict, local strictness being the key unifying property between the revealed and statistical preference models.

We think of Assumption (3'), below, as the analog of Assumption (3) for randomized experiments. Its main requirement is that the sampling distribution have full support, so as to sample thoroughly the space of binary decision problems.¹⁰ Full support can be relaxed for preferences that are identified on a proper subset of X (as we do in Section 5.1). Assumption (3') also requires that we almost never draw a decision problem that makes the subject indifferent, to help prevent incorrect inferences under indifference. With locally strict preferences, this property is commonly satisfied by diffuse distributions.

ASSUMPTION 3'. λ has full support and for all $\succeq \in \mathcal{P}$, $\{(x, y) : x \sim y\}$ has λ -probability 0.

⁹ $|R_n \setminus \succeq|$ denotes the number of elements in $R_n \setminus \succeq$. The Kemeny distance between two finite binary relations R and R' is usually defined as $|R \Delta R'|$, where Δ is the symmetric difference. Note that if $(x, y) \in \succeq \setminus R_n$ and $\{x, y\} \in \Sigma_k$, Then $(y, x) \in R_n \setminus \succ$. In our model, alternatives are strictly ranked by R_n and by \succeq with probability one. Hence,

$$\begin{aligned} \sum_{\{x,y\} \in \Sigma_n} (\mathbf{1}_{(x,y) \in \succeq \setminus R_n} + \mathbf{1}_{(x,y) \in R_n \setminus \succeq} + \mathbf{1}_{(y,x) \in \succeq \setminus R_n} + \mathbf{1}_{(y,x) \in R_n \setminus \succeq}) \\ = 2 \sum_{\{x,y\} \in \Sigma_n} (\mathbf{1}_{(x,y) \in R_n \setminus \succeq} + \mathbf{1}_{(y,x) \in R_n \setminus \succeq}) \end{aligned}$$

with probability one, which justifies our Kemeny distance terminology.

¹⁰Full support means that there is no proper closed subset of the sample space that has probability 1.

Under the above assumptions, the Kemeny-minimizing estimator is consistent. Recall that ρ denotes any compatible metric on the space of preferences.

Theorem 2. *Suppose the statistical preference model $(X, \mathcal{P}, \lambda, q)$ meets Assumptions (1), (2) and (3'), and suppose the subject's preference is $\succeq^* \in \mathcal{P}$. Let \succeq_n denote the Kemeny-minimizing estimator for the n -th experiment Σ_n . Then, $\{\succeq_n\}_{n \in \mathbf{N}}$ converges to \succeq^* in probability; that is, for any $\eta > 0$,*

$$\lim_{n \rightarrow \infty} \Pr(\rho(\succeq_n, \succeq^*) < \eta) = 1.$$

The proof of Theorem 2 is in Appendix C. One challenge is that the class of preferences may be very rich, which increases the potential for overfitting: the noise in the data may be misinterpreted as being part of the subject's true preference. And indeed, for many common spaces of alternatives, it can be shown that one can rationalize perfectly any finite set of observations by a locally strict preference. Imposing that the class of preferences be closed allows to overcome this difficulty.

We stress that Theorem 2 requires *no assumption* on the error probability function q , other than measurability and asking that the subject be more likely to follow his preference than to make a mistake. Alice may remain agnostic about the dependence of q on the underlying preference \succeq and the alternatives to choose from x and y . Moreover, aside from the independence of the draws of alternatives, the requisites on λ , stated in Assumption (3'), are minimal. In particular, calculating the Kemeny estimator does not require any assumptions on q and λ . It requires Alice to specify \mathcal{P} , but allows her to be largely agnostic about the rest of the model.

Concerning the class \mathcal{P} , two remarks are in order. First, Assumption (2) does not require that the preferences in \mathcal{P} be transitive. Thus, our framework can handle theories that contain intransitive choice. Secondly, Assumption (2) demands that \mathcal{P} be closed. The full set of locally strict preferences is not closed in general (Section 6.4 discusses the key issues). When this is the case, working with a smaller set of preferences makes the class closed. Sections 4 and 5 provide several examples.

Having the guarantee that preference estimates converge accurately, Alice may want to know how large of a sample is needed to estimate the subject's

preference within a given approximation error. Our third result, a quantitative version of Theorem 2, establishes lower bounds on the rate of convergence.

To state the result, we introduce some terminology. For any $\eta > 0$ and any $\delta \in (0, 1)$, let $N(\eta, \delta)$ be the smallest value of N such that for all $n \geq N$, and all underlying subject preferences $\succeq^* \in \mathcal{P}$,

$$\Pr(\rho(\succeq^n, \succeq^*) < \eta) \geq 1 - \delta.$$

(By convention, we let $N(\eta, \delta) = \infty$ if no finite value of N exists.) That is, $N(\eta, \delta)$ is the size of the smallest experiment such that the probability that the preference estimate is η -close to the true subject preference is guaranteed to be at least $1 - \delta$, no matter the true subject preference.

In addition, $\mu(\cdot, \succeq^*)$ denotes the probability measure induced on the product space $X \times X$ by \succeq^* , q and λ , as follows:

$$\mu(A, \succeq^*) = \int_A q(\succeq^*; x, y) d\lambda(x, y).$$

Loosely speaking, $\mu(x, y, \succeq^*)$ represents how likely a subject with preference \succeq^* is to choose x over y in a decision problem randomly drawn. In particular, the value of $\mu(\succeq, \succeq^*)$ represents the probability that the choice of a subject with preference \succeq^* over a randomly drawn decision problem is consistent with the preference \succeq .

Finally, for $\eta > 0$, we let

$$r(\eta) = \inf \{ \mu(\succeq^*, \succeq^*) - \mu(\succeq, \succeq^*) : \succeq, \succeq^* \in \mathcal{P}, \rho(\succeq^*, \succeq) \geq \eta \}.$$

Roughly, the value of $r(\eta)$ captures the smallest possible probability that a subject make a choice that is perfectly consistent with his own preference but is inconsistent with a preference at least η -distant.

To obtain convergence rates, we appeal to Vapnik-Chervonenkis theory. Given a preference environment (X, \mathcal{P}) , for $n \in \mathbf{N}$, let S_n be the largest size of the sets

$$\{(\mathbf{1}_{x_1 \succeq y_1}, \dots, \mathbf{1}_{x_n \succeq y_n}) \in \{0, 1\}^n : \succeq \in \mathcal{P}\}$$

over all binary decision problems $\{x_k, y_k\} \subset X$ for $k = 1, \dots, n$. We always have $S_n \leq 2^n$, and, if \mathcal{P} is rich enough, we may have $S_n = 2^n$. The VC dimension of \mathcal{P} (abbreviation for Vapnik-Chervonenkis dimension) is then defined as the maximum value of n such that $S_n = 2^n$, and is infinite if no such maximum exists.

The consistency of the Kemeny-minimizing estimator applies no matter the VC dimension of \mathcal{P} , but when \mathcal{P} has a finite VC dimension, and so is not too rich, we can, in addition, obtain uniform bounds on the convergence rates.

Theorem 3. *If the statistical preference model $(X, \mathcal{P}, \lambda, q)$ meets Assumptions (1), (2), and (3'), and if \mathcal{P} has a finite VC dimension, then¹¹*

$$N(\eta, \delta) = O\left(\frac{1}{r(\eta)^2} \ln \frac{1}{\delta}\right).$$

The proof of Theorem 3 is in Appendix D, in which we also provide a more refined nonasymptotic bound. Of course, the value of $r(\eta)$ depends on the specific preference environment being considered. Below in Sections 4 and 5, we apply Theorem 3 in different environments. Basu and Echenique (2018) compute the VC dimension of some common models of choice, they show, in particular, that the class of expected utility, Choquet expected utility, and two-state max-min preferences have finite VC dimension.

4. PREFERENCES FROM UTILITIES

In this section and the next, we show that the assumptions of our general framework are valid in a variety of important preference environments. In some cases, we compute explicit convergence rates for the Kemeny-minimizing estimator. This section handles preferences based on utility functions, while the next section deals with monotone preferences.

4.1. Expected utility preferences. The standard expected utility model is situated within our framework, and our results can be used to establish convergence to expected utility preferences. Let $\Pi \equiv \{\pi_1, \dots, \pi_d\}$ be a collection of $d \geq 2$ prizes, and let Δ^{d-1} denote the $(d-1)$ -dimensional simplex $\{p \in \mathbf{R}_+^d : p_1 + \dots + p_d = 1\}$. Think of each element p of the simplex as a

¹¹The big O notation refers to the usual asymptotic upper bound.

lottery over the prizes in Π , with p_i the probability of getting π_i . The set of alternatives is Δ^{d-1} , endowed with the Euclidean topology.

An *expected utility preference* stands for any preference \succeq on Δ^{d-1} defined by a vector of utility indexes $v \in \mathbf{R}^d$, with the property that $p \succeq p'$ if and only if $v \cdot p \geq v \cdot p'$. It is *nonconstant* if there is at least one pair $p, p' \in \Delta^{d-1}$ for which $p \succ p'$.

The next proposition makes Theorems 1, 2 and 3 immediately applicable.

Proposition 1. *In the expected-utility environment just described, the set of alternatives $X \equiv \Delta^{d-1}$ endowed with the Euclidean topology meets Assumption (1), and the class \mathcal{P} of all nonconstant expected utility preferences meets Assumption (2).*

The proof of Proposition 1 is in Appendix E.

Moreover, we can refine the convergence rates of Theorem 3, provided that we restrict attention to error probability functions q defined by means of a real function f as follows: when $x \succ y$,

$$(1) \quad q(\succeq; x, y) = 1 - \frac{1}{2}f(\|x - y\|),$$

$$(2) \quad q(\succeq; y, x) = \frac{1}{2}f(\|x - y\|),$$

where $f : \mathbf{R}_+ \rightarrow [0, 1]$ is a strictly decreasing function that is continuously differentiable at 0 and for which $f(0) = 1$ and $f'(0) < 0$.¹² The focus on this class of error probability functions allows for explicit convergence rates of the Kemeny-minimizing estimator, as below.

Proposition 2. *For the statistical preference model $(X, \mathcal{P}, \lambda, q)$, where $X \equiv \Delta^{d-1}$, \mathcal{P} is the set of all nonconstant expected utility preferences, λ is the uniform distribution on Δ^{d-1} and q satisfies Equations (1) and (2), the Kemeny-minimizing estimator is consistent and, as $\eta \rightarrow 0$ and $\delta \rightarrow 0$,*

$$N(\eta, \delta) = O\left(\frac{1}{\eta^{4d-2}} \ln \frac{1}{\delta}\right).$$

¹² $\|\cdot\|$ denotes the Euclidean norm $\|x - y\| \equiv \sqrt{\sum_{i=1}^d (x_i - y_i)^2}$.

One can rewrite the theorem statement to provide an $O_p((1/n)^{1/d})$ convergence rate.¹³ The proof of Proposition 2 is in Appendix F.

Note that the uniform distribution is not at all essential for Theorem 2. It just yields a particularly simple closed form for the bound on $r(\eta)$ to be used in Theorem 3.

4.2. Preferences induced by utility functions. Our framework encompasses preference environments derived from collections of utility functions. Specifically, our assumptions may be derived from conditions on the utility representations under consideration, rather than directly imposing the assumptions on a family of preferences.

Consider a set of alternatives X . A *utility function* is any function $u : X \rightarrow \mathbf{R}$. We endow the space of utility functions with the topology of compact convergence.¹⁴ For any utility function $u : X \rightarrow \mathbf{R}$, let $\Phi(u)$ denote the preference induced by u , that is, the binary relation defined by $x \Phi(u) y$ if and only if $u(x) \geq u(y)$. And for any set of utility functions \mathcal{U} , let $\Phi(\mathcal{U}) = \{\Phi(u) : u \in \mathcal{U}\}$ denote the image of \mathcal{U} .

Proposition 3. *Suppose X satisfies Assumption (1), \mathcal{U} is a compact set of continuous utility functions, and every preference in $\Phi(\mathcal{U})$ is locally strict. Then the class of preferences $\mathcal{P} \equiv \Phi(\mathcal{U})$ meets Assumption (2).*

Proposition 3 is an immediate implication of Theorem 8 of Border and Segal (1994), who establish the continuity of Φ (see Appendix E).

To put Proposition 3 to work in a concrete example, let us look at the case of intertemporal choice. Suppose a good can be consumed at $d \geq 2$ different dates $t_1 < \dots < t_d$. In this environment, an alternative is a vector of the Euclidean space \mathbf{R}_+^d whose i -th entry indicates the amount consumed at the date t_i .

¹³The O_p notation refers to the stochastic boundedness notation.

¹⁴A sequence of functions $f_n : X \rightarrow \mathbf{R}$, $n \in \mathbf{N}$, converges compactly to a function f if and only if it converges uniformly to f on every compact set $K \subseteq X$. This topology on utilities is commonly used in the literature; see for example Mas-Colell (1974) and Border and Segal (1994).

Fix $a, b \in \mathbf{R}_{++}$ with $a < b$ and call \mathcal{V} the set of continuous functions $v : \mathbf{R}_+ \rightarrow \mathbf{R}$ that satisfy, for all $x < y$,

$$a \cdot (y - x) \leq v(y) - v(x) \leq b \cdot (y - x).$$

We interpret $v(x)$ as the utility for an immediate consumption of quantity x of the good. The above inequality constrains marginal utilities to be positive and bounded above and below.

Denote by \mathcal{U} the set of the utility functions u over \mathbf{R}_+^d that are written

$$u(x_1, \dots, x_d) = \sum_{i=1}^d \delta_i v(x_i),$$

where $v \in \mathcal{V}$, and $\delta = (\delta_1, \dots, \delta_d) \in [\varepsilon, 1]^d$ is a vector of discount factors, with ε an arbitrarily small positive lower bound. So, the set \mathcal{U} captures discounted utility preferences with general discount factors.

Compactness of \mathcal{U} follows the Arzelà-Ascoli theorem (for example, Theorem 6.4 of Dugundji, 1966). And clearly each utility function describes a locally strict preference, because if an individual with utility $u \in \mathcal{U}$ prefers the consumption vector $x \in \mathbf{R}_+^d$ to $y \in \mathbf{R}_+^d$, then he strictly prefers the consumption vector $x + \eta \mathbf{1}$ to y , for any $\eta > 0$.

Therefore, Proposition 3 applies, and so does the framework of Section 3. For example, one possible use of our theory is to recover discount factors from the data, or to check for distortions with respect to standard models such as exponential discounting. It is also worth noting that Proposition 1 can be viewed as a consequence of Proposition 3 (see Appendix E for details).

5. APPLICATION TO MONOTONE PREFERENCES

In many economic settings, it is safe to posit the existence of a universal ordering, by which some alternatives are ranked above others by all the individuals of the relevant population: preferences are monotone with respect to this ordering. For example, for the classical consumption environment in which individuals choose bundles of goods, it is usually assumed that individuals strictly prefer to have more of each good. In laboratory experiments, it is common to assume some form of objective ranking, for instance when enforcing

single-switching in multiple-price lists, or when using randomization devices to enforce incentives. Monotonicity with respect to such universal orderings turns out to be a very useful discipline on preferences.

In this section, we adopt the following terminology. Fix a set of alternatives X . We call *dominance relation* any binary relation \triangleright on X that is not reflexive, that is, for each $x \in X$, $x \not\triangleright x$. The relation \triangleright is said to be *open* when \triangleright is an open set in the product space $X \times X$.¹⁵ Being open for a dominance relation can be interpreted as a continuity property, saying that if x dominates y then this domination extends locally around the alternatives x and y .

Given a dominance relation \triangleright , a preference relation \succeq is *strictly monotone* with respect to \triangleright if, for each $x, y \in X$, $x \triangleright y$ implies $x \succ y$. Having a class of strictly monotone preferences captures the above idea that some alternatives are universally preferred to some others in accordance to the dominance relation.

Usually, strict monotonicity alone does not suffice to ensure that the preference is locally strict, the first crucial condition in our framework. It helps to add a notion of transitivity. We call a preference relation \succeq *Grodal-transitive* if for all $x, y, z, w \in X$, $x \succeq y \succ z \succeq w$ implies $x \succeq w$. Named after Grodal (1974), Grodal-transitivity is weaker, and so more permissive, than the usual notion of transitivity. Importantly, together with strict monotonicity, Grodal-transitivity makes the class of preferences closed, the second crucial condition of our framework.

Lemma 2. *Suppose Assumption (1) is met and \triangleright is a dominance relation on X . If \triangleright is open, then the class of preferences that are Grodal-transitive and strictly monotone with respect to \triangleright is closed.*

The proof of Lemma 2 is in Appendix G.

It is worth noting that, in general, closedness is not achieved under the usual notion of transitivity and strict monotonicity. If one wishes to impose transitivity, the class of preferences must be reduced further to obtain a closed set (as we did in the example of Section 4.1). Of course, there is no harm in having a more generous class of preferences. Even if the class includes

¹⁵That is, for each x, y with $x \triangleright y$, there exists a neighborhood V of (x, y) in $X \times X$ such that for all $(x', y') \in V$, $x' \triangleright y'$.

preferences that fails desirable properties such as classical transitivity—and so may include irrelevant preferences—preference estimates are guaranteed to converge to the correct underlying preference, so that any violation by the preference estimates eventually gets corrected in the limit.

The main benefit of Grodal-transitivity is that it is enough to make strictly monotone preferences locally strict under relatively mild conditions.

Lemma 3. *Suppose Assumption (1) is satisfied and \succeq is a preference strictly monotone with respect to the dominance relation \triangleright . If, for each $x \in X$, there exists $y, z \in X$ arbitrarily close to x and such that $y \triangleright x$ and $x \triangleright z$, then \succeq is locally strict.*

The proof of Lemma 3 is in Appendix H.

Therefore, Assumption (2) on the class of preferences is valid as long as X is well behaved, the dominance relation is open, and the preferences considered are Grodal-transitive and strictly monotone. The examples below show that these properties are satisfied in many common preference environments.

5.1. Commodity spaces. The classical setup of consumer demand analysis features a commodity space over $d \geq 2$ goods, where consumers get to choose over bundles of goods (as in Afriat, 1967, Mas-Colell, 1978, or Varian, 1982). The set of alternatives is the Euclidean space \mathbf{R}_{++}^d , the i -th entry of vector (x_1, \dots, x_d) is interpreted as the consumed quantity of the i -th good. This environment is part of our framework when preferences are asked to satisfy a monotonicity condition.

Consider the dominance relation \gg on \mathbf{R}_{++}^d by $x \gg y$ exactly when $x_i > y_i$ for all $i = 1, \dots, d$. An individual whose preference is strictly monotone with respect to \gg means that this individual strictly prefers to have more of every good, a postulate that appears reasonable in many situations, and that is common in economic models. It is evident that the relation \gg is open, and for any $x \in \mathbf{R}_{++}^d$, $x + \varepsilon \mathbf{1} \gg x$ for all $\varepsilon > 0$ while $x \gg x - \varepsilon \mathbf{1} \in \mathbf{R}_{++}^d$ for all small enough $\varepsilon > 0$. Hence, Lemmas 2 and 3 apply, and we get Proposition 4.¹⁶

¹⁶While the set \mathbf{R}_{++}^d is not complete under the Euclidean metric, there exists a compatible complete metric by Alexandroff's Theorem (Theorem 24.12 of Willard, 2004). Of course, \mathbf{R}_{++}^d is also locally compact and separable, and hence Assumption (1) is satisfied.

Proposition 4. *In the commodity-space environment just described, the set of alternatives $X \equiv \mathbf{R}_{++}^d$ endowed with the Euclidean topology meets Assumption (1), and the class \mathcal{P} of all preferences that are Grodal-transitive and strictly monotone with respect to \gg meets Assumption (2).*

The same set of alternatives can be used to describe state-contingent payments, with an objective public distribution over states, and where the i -th entry of a vector encodes the payment received in the state i . In such an environment, one may want to test the validity of the hypothesis that individuals maximize an expected utility function (as, for example, in Green and Srivastava, 1986), or maximize a utility function that is monotone with respect to first-order stochastic dominance (as in Nishimura, Ok, and Quah, 2017). Because both classes of preferences are more restrictive than the class \mathcal{P} considered here, our convergence results continue to apply, which means that we can fully recover preferences and examine precisely the validity of these hypotheses.

We may also work with other dominance relations. In particular, the “smaller” the dominance relation \triangleright (as a subset of $X \times X$), the “larger” the set of strictly monotone preferences, and so the more general the class \mathcal{P} . For example, the relation \gg_α defined by

$$x \gg_\alpha y \quad \text{exactly when for all } i \in 1, \dots, d, \quad x_i - y_i > \frac{\alpha}{d-1} \sum_{j \neq i} (x_j - y_j),$$

is an instance of dominance relation that is open if $0 \leq \alpha < 1$. The relation \gg_0 coincides with the relation \gg just described, and as α grows to 1, \gg_α becomes a relation with empty interior, precisely, the relation defined by $x \gg_1 y$ exactly when $x - y = \beta \mathbf{1}$ for some $\beta > 0$. Hence, as α increases, it becomes less demanding to require that preferences be strictly monotone with respect to \gg_α .

This preference environment allows us to compute explicit convergence rates for the statistical preference model. Given a set of alternatives K , we say that the class \mathcal{P} is *identified on K* if, whenever two preferences coincide on K , they must be identical everywhere. We require that \mathcal{P} be identified on a compact set, and that it has finite VC dimension. This requirement is satisfied by many models of interest; for example, the class of preferences with a constant

elasticity of substitution (CES) utility representation, or, when $\{1, \dots, d\}$ is interpreted as a state space, and the set of alternatives X is interpreted as a space of monetary acts, preferences with a CARA subjective expected utility representation.¹⁷

In our next result, let K^ε denote the set of all points within distance ε of a set K with respect to the Euclidean metric.¹⁸

Proposition 5. *Let K be a compact set in $X \equiv \mathbf{R}_{++}^d$, and fix $\theta > 0$. Suppose the statistical preference model for commodity spaces $(X, \mathcal{P}, \lambda, q)$ is such that \mathcal{P} has a finite VC-dimension and is identified on K , λ is the uniform probability measure on $K^{\theta/2}$, and q satisfies Equations (1) and (2) of Section 4. Then the Kemeny-minimizing estimator is consistent and, as $\eta \rightarrow 0$ and $\delta \rightarrow 0$,*

$$N(\eta, \delta) = O\left(\frac{1}{\eta^{2d+2}} \ln \frac{1}{\delta}\right).$$

The proof of Proposition 5 is in Appendix I.¹⁹

The notion of a dominance relation is fairly general. In consequence, our results extend to environments beyond commodity spaces. We proceed with applications to choice over menus and intertemporal choice.

5.2. Choice over menus. Our next application deals with recovering preferences over menus, following Kreps (1979) and Dekel, Lipman, and Rustichini (2001). Let $\Pi = \{\pi_1, \dots, \pi_d\}$ be a collection of prizes and let Δ_{++}^{d-1} be the interior of the $(d-1)$ -dimensional simplex, interpreted as the set of full-support

¹⁷See Basu and Echenique (2018) for a discussion of other uses of the VC dimension for choice under uncertainty.

¹⁸ $K^\varepsilon = \bigcup_{x \in K} \{y \in X : \|x - y\| \leq \varepsilon\}$ is the generalized ball of radius ε around K .

¹⁹For technical reasons, the bound $N(\eta, \delta)$ in Proposition 5 uses the metric ρ^{K, K^θ} instead of ρ in our earlier definition of $N(\eta, \delta)$. The parameter θ is a ‘‘fudge factor,’’ and the corresponding fudged metric is defined as

$$\rho^{K, K^\theta}(\succeq, \succeq') = \max \left\{ \sup \{ \rho((x, y), \succeq' \cap (K \times K)^\theta) : x \succeq|_K y \}, \right. \\ \left. \sup \{ \rho((x, y), \succeq \cap (K \times K)^\theta) : x \succeq'|_K y \} \right\},$$

where $\succeq|_K$ is the restriction of \succeq to K and, for $A \subseteq X \times X$,

$$\rho((x, y), A) = \inf \{ \| (x, y) - (x', y') \| : (x', y') \in A \}.$$

Note that as θ vanishes to zero, ρ^{K, K^θ} becomes equal to the Hausdorff distance restricted to $K \times K$. For $\theta > 0$, $\rho^{K, K^\theta}(\succeq, \succeq') \leq \rho^{K, K^0}(\succeq, \succeq')$.

distributions over the elements of Π . We endow Δ_{++}^{d-1} with the Euclidean metric.

Let \mathcal{M} denote the set of closed convex subsets of Δ_{++}^{d-1} with nonempty interior. We interpret \mathcal{M} as a set of menus of lotteries. A subject who possesses a menu $m \in \mathcal{M}$ gets to choose a lottery in m , and subsequently receives a prize drawn according to this lottery. The convexity of menus that is assumed here is also implied by the axiom of indifference to randomization introduced by Dekel et al. (2001). We endow \mathcal{M} with the Hausdorff topology, as is standard in menu theory.

We define the dominance relation \sqsupseteq as follows: for two menus m_A and m_B , $m_A \sqsupseteq m_B$ if every expected-utility decision maker with full knowledge of her utility when making menu choices would strictly prefer m_A to m_B . More precisely, we write

$$\mathcal{U} \equiv \left\{ u \in \mathbf{R}^d : \sum_{i=1}^d u_i = 0, \|u\| = 1 \right\}$$

the set of all utility indexes over prizes, up to a normalization (we rule out the trivial preference that is indifferent between any two lotteries). Then we write $m_A \sqsupseteq m_B$ if and only if for every $u \in \mathcal{U}$,

$$\sup_{p \in m_A} u \cdot p > \sup_{p \in m_B} u \cdot p.$$

Since we restrict attention to convex menus, $A \sqsupseteq B$ implies $A \supset B$. Hence, the dominance relation \sqsupseteq is similar to, but weaker than, the set-containment relation traditionally used in models of choice over menus. In particular, monotonicity with respect to \sqsupseteq is less demanding than monotonicity with respect to \supset .

Our next proposition establishes that the revealed preference framework applies to the menu preference environment.

Proposition 6. *In the menu environment just described, the set of alternatives $X \equiv \mathcal{M}$ endowed with the Hausdorff topology meets Assumption (1), and the class \mathcal{P} of all preferences that are Grodal-transitive and strictly monotone with respect to \sqsupseteq meets Assumption (2).*

The proof of Proposition 6 is in Appendix J.

5.3. Dated rewards. We apply our results to intertemporal choice, using the environment introduced by Fishburn and Rubinstein (1982). In this environment, the set of alternatives is \mathbf{R}_{++}^2 , endowed with the Euclidean topology. An element $(t, x) \in \mathbf{R}_{++}^2$ represents a monetary payment x delivered on date t .

Economists usually assume that individuals prefer more money over less, and to be paid earlier rather than later. Hence, the relevant dominance relation, written as $>_\tau$, is defined by $(t, x) >_\tau (t', x')$ if and only if $t < t'$ and $x > x'$. The relation $>_\tau$ is open, and, for any $(t, x) \in \mathbf{R}_{++}^2$ and $\varepsilon > 0$ small enough, $(t - \varepsilon, x + \varepsilon) >_\tau (t, x)$ and $(t, x) >_\tau (t + \varepsilon, x - \varepsilon)$. In consequence, Lemmas 2 and 3 deliver our next result.

Proposition 7. *In the dated rewards environment just described, the set of alternatives $X \equiv \mathbf{R}_{++}^2$ endowed with the Euclidean topology meets Assumption (1), and the class \mathcal{P} of all preferences that are Grodal-transitive and strictly monotone with respect to $>_\tau$ meets Assumption (2).*

Just as in the example in Section 4.2, the dated rewards environment can be used to evaluate the exponential discounting model. Exponential discounting means that (t, x) is preferred to (t', x') if and only if $\delta^t u(x) \geq \delta^{t'} u(x')$, for a given discount factor $\delta \in (0, 1)$ and a given strictly increasing function $u : \mathbf{R}_{++} \rightarrow \mathbf{R}$.

In spite of the class \mathcal{P} being much more general than those representable by exponential discounting, as long as the subject complies with the exponential discounting model, the estimated preferences are guaranteed to converge to the underlying exponential discounting preferences. They would then give an approximation of the discount factor and underlying utility for money. Conversely, if the experimenter were to observe that her preference estimates violate the exponential discounting model, she may conclude that her hypothesis that the subject complies to exponential discounting is false.

5.4. Intertemporal consumption. Continuing with the theme of intertemporal choice, we revisit the example in Section 4.2. There is a good to be consumed at a sequence of $d \geq 2$ increasing dates t_1, \dots, t_d . The set of alternatives is \mathbf{R}_{++}^d , and each element (x_1, \dots, x_d) gives the amount consumed at each date.

As before, we can hypothesize that individuals prefer more of the good early over less later. In the present environment, this postulate is captured by the dominance relation \ggg on \mathbf{R}_{++}^d whereby $x \ggg y$ if and only if for every k ,

$$\sum_{i=1}^k x_i > \sum_{i=1}^k y_i.$$

In the same environment, Nishimura, Ok, and Quah (2017) suggest another postulate: that individuals are neutral to time while they still prefer to get more of the good. The associated dominance relation $>_{\text{sym}}$ is defined as $x >_{\text{sym}} y$ if and only if there exists a permutation σ over $\{1, \dots, d\}$ such that for every i , $x_{\sigma(i)} > y_{\sigma(i)}$.

It is immediately seen that \ggg is open, and that, when $x \in \mathbf{R}_{++}^d$ and $\varepsilon > 0$ is small enough, $x \ggg x - \varepsilon \mathbf{1}$ and $x + \varepsilon \mathbf{1} \ggg x$. The very same observations apply to the relation $>_{\text{sym}}$. By a logic that is now routine, we get the following proposition.

Proposition 8. *In the intertemporal consumption environment just described, the set of alternatives $X \equiv \mathbf{R}_{++}^d$ endowed with the Euclidean topology meets Assumption (1), and the class \mathcal{P} of all preferences that are Grodal-transitive and strictly monotone with respect to \ggg (or with respect to $>_{\text{sym}}$) meets Assumption (2).*

5.5. Choice over lotteries. Consider a set of $d \geq 2$ monetary rewards $\Pi \equiv \{\pi_1, \dots, \pi_d\}$, and let the interior of the $(d-1)$ -simplex, Δ_{++}^{d-1} , be the set of alternatives endowed with the Euclidean topology. We interpret an element $p \in \Delta_{++}^{d-1}$ as a full-support lottery over monetary rewards. This choice domain is an instance of the domain studied in Section 4.1.

Suppose that the elements of Π are ordered as $\pi_1 < \dots < \pi_d$. A natural dominance relation is strict first-order stochastic dominance, noted $>_{\text{FSD}}$, where $p >_{\text{FSD}} p'$ if and only if for all $k = 1, \dots, d-1$,

$$\sum_{i=1}^k p_i < \sum_{i=1}^k p'_i.$$

The reason for using *strict* first-order stochastic dominance is that this relation is open, as opposed to first-order stochastic dominance. Now, let $p \in \Delta_{++}^{d-1}$.

For a small enough positive ε , we can define $p' \in \Delta_{++}^{d-1}$ by $p'_i = p_i - \varepsilon$ for all $i \leq d-1$ and $p'_d = p_d + (d-1)\varepsilon$. Then $p' >_{\text{FSD}} p$. And similarly, when instead $p'_i = p_i + \varepsilon$ for all $i \leq d-1$ and $p'_d = p_d - (d-1)\varepsilon$, $p >_{\text{FSD}} p'$. Hence Lemmas 2 and 3 apply, and we get the following result.

Proposition 9. *In the lottery environment just described, the set of alternatives $X \equiv \Delta_{++}^{d-1}$ endowed with the Euclidean topology meets Assumption (1), and the class \mathcal{P} of all preferences that are Grodal-transitive and strictly monotone with respect to $>_{\text{FSD}}$ meets Assumption (2).*

Note that the class of nonconstant expected-utility preferences studied in Section 4.1 and the class considered in the present lottery choice environment are distinct, and neither one is a refinement of the other.

6. CONCLUDING DISCUSSION

This paper deals with the question of recovering individual preferences from observed choice data, when the data consist of a finite number of binary comparisons. Decision theorists often consider the question of “backing out” a model from data on pairwise choices, motivated by laboratory experiments, in which pairwise choices are common. They assume, however, rich and infinite data sets. Econometricians study the convergence of preference estimates, but their models usually differ from the pairwise choice paradigm. Moreover, the conditions needed for consistency of their estimates are imposed as added-on assumptions, instead of being derived from the properties of the economic model under consideration.

We provide a common unifying framework. We show that the class of preferences considered should be locally strict. Under local strictness, and some regularity conditions, any preference that rationalizes the observed pairwise choices converges to the correct data-generating preference. In the statistical counterpart to our model, the Kemeny-minimizing estimator, which outputs the preferences that best fit the data, is consistent. In addition, convergence rates can be obtained while remaining largely agnostic on the sampling method and the error probability function.

Our results require weak assumptions and apply to a broad range of standard preference environments. We conclude with a discussion on a few aspects of our model.

6.1. On the exhaustiveness of experiments. Because we deal with non-parametric estimation, it is important that the alternatives used through the sequence of growing experiments form a dense set. The reason is that, if the experimenter was to leave an open set of alternatives outside of her experimental design, the subject's preferences over alternatives in that set would be very hard to gauge. Therefore, this denseness is the key condition of the exhaustiveness of experiments.

In practice however, one may want to restrict attention to classes of preferences that are small enough so that there is no need to elicit choices over a set that is dense in X . This is true, for example, for expected utility preferences over lotteries, or homothetic preferences in \mathbf{R}^d . In those cases, the experimenter only cares to infer a single indifference curve, from which she can uncover the entire preference. It is then fine to focus on a small set of alternatives. But one would still need to work with a subset of alternatives dense within a relevant subset of X . An example along these lines is presented in Section 5.1. In this environment, preferences are fully identified on a strict subset of X .

6.2. On the convergence of preferences. At a general level, the topology of closed convergence is defined by the property that individuals with comparable preferences behave similarly on closely related decision problems. Such continuity property appears natural, and even necessary to be able to learn from finite data. In formal terms, if $x \succ y$ for some alternatives (x, y) and a preference \succeq , and if (x', y') are alternatives in a neighborhood of (x, y) , then the topology of closed convergence is defined exactly so that $x' \succ' y'$ for any preference \succeq' close enough to \succeq .²⁰

This topology also has the property that, if the experimenter learns that the subject prefers (at least weakly) x to y through her experiments—because the

²⁰To be even more formal, under the assumptions of our results, the closed convergence topology is the smallest topology for which the set $\{(x, y, \succeq) : x \succ y\}$ is open in the product topology; see Theorem 3.1 of Kannai (1970).

subject chooses x when presented with the pair $\{x, y\}$ —this is also reflected in the limiting preference, when it exists: if for some $N \in \mathbf{N}$ and $x, y \in X$, we have $x \succeq_n y$ for all $n \geq N$, and if $\succeq_n \rightarrow \succeq^*$, then $x \succeq^* y$. And our model allows for the possibility that certain parts of the subject’s preference remain unobserved. For example, the experimenter cannot learn, from a finite number of observations, that x is strictly preferred to y , although she may learn that x is weakly preferred. In this case, if we can still ensure a unique limiting preference \succeq^* , then it means that correct inferences about missing observations have been made. The closed convergence topology is therefore well suited to the concept of weak rationalization.

6.3. On the connection with econometrics. Theorem 2 is an instance of the consistency of M-estimators. Consistency is well known to rely on three properties of the econometric environment (see, for example, Theorem 4.1.1 of Amemiya, 1985, or Theorems 2.1 and 3.1 of Newey and McFadden, 1994). The first is that the true parameter has to be a unique extremum of the population version of the objective function. We prove this property in Lemma 6. The second is the uniform convergence of the sample objective function to the population version. We do not need to assume this property in an ad-hoc fashion; instead, we are able to derive it from our model primitives. Finally, the canonical results on M-estimators need that the parameter space is compact. The topology we use on the space of preferences guarantees its compactness. So, even though our estimation problem is fully nonparametric, we are able to work with a compact space of parameters and we do not need to assume uniform convergence or stochastic equicontinuity.

The most related work in the econometrics literature is Chernozhukov, Hong, and Tamer (2007), who present consistent estimators for partially identified models from moment conditions. We differ from their work in that they provide a general methodology for parametric estimation, while we are specifically interested in the approximation of preferences from pairwise comparisons, and our problem is nonparametric. Because their methodology aims at being general, their main consistency result (Theorem 3.1) assumes the uniform convergence property (part of condition C1). It is not derived from the revealed preference questions that motivate their study. Our consistency result also

depends on an analogous uniform convergence property (as mentioned above, this is true generally of the consistency of M-estimators), but it is obtained as a consequence of the primitives of our model. Obtaining consistency directly from the model primitives is the key contribution of Theorem 2. In certain environments, revealed preference conditions are representable by means of moment inequalities, and the results of Chernozhukov et al. can be applied, Blundell, Browning, and Crawford (2008) is a notable instance of such an application in the context of demand analysis.

6.4. On the closedness of the class of preferences. Our results require that the class \mathcal{P} of preferences considered be closed. For general sets of alternatives, the set of locally strict preferences is, however, not closed. Lemma 1 implies that, if a sequence of locally strict rationalizations along an exhaustive set of experiments has a locally strict limiting preference, then there can be no other limit. Yet, without the assumption that \mathcal{P} is closed, it is possible that two distinct sequences of locally strict rationalizations converge to two different limits, one not being locally strict. Hence, one must impose that the class \mathcal{P} be closed.

To build intuition, consider a simple environment with the real line $X \equiv \mathbf{R}$ as the set of alternatives, and suppose that a locally strict preference \succeq^* generates the data. The argument is general, but to fix ideas, take for \succeq^* the “greater than or equal to” relation. Note that there is a unique choice function c^* generated by \succeq^* . Now, let $\{\Sigma_n\}$ be any exhaustive set of experiments, and let us write \widehat{B}_n for the set of all the alternatives used over the experiments $\Sigma_1, \dots, \Sigma_n$. Of course, $|\widehat{B}_n| < +\infty$ for all n .

Given the choice function c^* , for every n , we can use $\succeq_n = \succeq^*$ as a rationalizing preference for the observed behavior on Σ_n . Then we evidently get $\succeq_n \rightarrow \succeq^*$. However, there are other sequences of locally strict rationalizations. In particular, we can have a sequence which converges to \succeq^I , the relation which ranks any pair of alternatives as indifferent.

The construction is simple by means of utility functions. For each n , we will make a utility function u_n whose induced preference is locally strict and rationalizes the choice data observed on Σ_n . Fix n and let us write \widehat{B}_n as an ordered set $\{b_n^1, \dots, b_n^m\}$ with $b_n^1 < \dots < b_n^m$. For any $x \notin (b_n^1, b_n^m)$, define

$u_n(x) = \arctan(x)$. Likewise, for any $x \in \widehat{B}_n$, define $u_n(x) = \arctan(x)$. Finally, for any $i = 1, \dots, m - 1$, define

$$u_n \left(b_n^i + \frac{b_n^{i+1} - b_n^i}{3} \right) = 1 \quad \text{and} \quad u_n \left(b_n^i + \frac{2(b_n^{i+1} - b_n^i)}{3} \right) = 0,$$

and we then extend u_n piecewise linearly, so that for any x in the open interval $(b_n^i, b_n^i + (b_n^{i+1} - b_n^i)/3)$ for instance, we have

$$u_n(x) = \arctan(b_n^i) + 3 \frac{x - b_n^i}{b_n^{i+1} - b_n^i} (1 - \arctan(b_n^i)),$$

and so on. Importantly, it is immediately seen that the preference \succeq'_n that is induced by u_n is locally strict. Moreover, the limit $\succeq'_n \rightarrow \succeq^I$ holds. The reason is that, for each $x, y \in \mathbf{R}$, there is, for all $\varepsilon > 0$, an integer n large enough and, $x_n, y_n \in \mathbf{R}$ with $|x_n - x| < \varepsilon$, $|y_n - y| < \varepsilon$, such that $u_n(x_n) \geq u_n(y_n)$.

This example illustrates the importance of the assumption that \mathcal{P} be closed to rule out sequences of rationalizations that behave increasingly erratically as experiments grow in size. The ideas behind the examples are general, and other examples can easily be found.

6.5. On the transitivity of preferences. Aside from the assumption of local strictness, our framework applies to very general classes of preferences. In particular, it applies to preferences without classical rationality hypotheses, such as transitivity. Still, one may wish to look for rationalizing preferences that are transitive. While it is perfectly reasonable to focus on transitive preferences, one must interpret Theorem 1 with care. Even when all the preferences that rationalize the observed behavior for a set of experiments can be chosen to be transitive, there is no guarantee that the limiting preference is transitive, even if Assumption (1) is satisfied. This fact owes to an example of Grodal (1974).

Adapted to our context, Grodal's example proceeds as follows. Figure 1 exhibits a nontransitive relation borrowed from Grodal (1974), with $X = \mathbf{R}_{++}^2$ (say, X is a commodity space with two goods). The lines depict indifference curves. All the green indifference curves intersect at one point: $(1/2, 1/2)$. Aside from the point $(1/2, 1/2)$, x is at least as good as y if and only if it lies on a (weakly) higher indifference curve. But, all bundles on an indifference

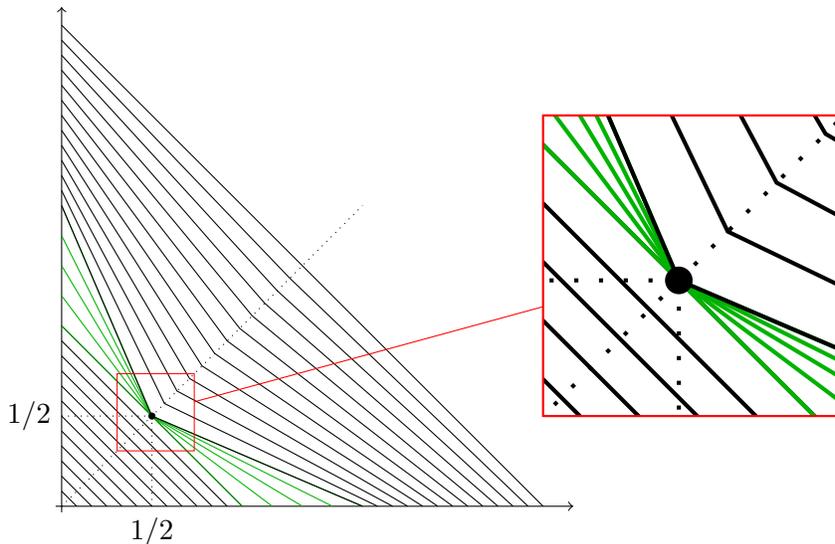


FIGURE 1. A non-transitive preference

curve passing through $(1/2, 1/2)$ are indifferent to $(1/2, 1/2)$. This feature makes the preference nontransitive; specifically, the indifference part of the preference is intransitive here. Let \succeq^* denote this preference.

Imagine an ordered collection of binary choice problems that do not include the alternative $(1/2, 1/2)$. Suppose that this collection is either finite or infinite but countable, as the set used in our growing experiments. Then for every n there is a ball around $(1/2, 1/2)$ that does not contain any alternative in the first n binary choice problems. Consider the preferences pictured in the diagram of Figure 2. Compared to the relation depicted in Figure 1, the preferences of Figure 2 have been modified close to $(1/2, 1/2)$ so that transitivity holds. Thus, one can construct a sequence of strictly monotone preferences, \succeq_n , $n \in \mathbf{N}$, where each \succeq_n is transitive, and $\succeq_n \rightarrow \succeq^*$, but \succeq^* is not transitive.

It is generally true that if each \succ_n (the strict part of \succeq_n) is transitive, then \succ^* will be transitive as well (see Grodal, 1974), but in some cases one may desire full transitivity of \succ^* .²¹ The central element of the example above is that the indifference curves get “squeezed” together too rapidly.

²¹Relations for which the strict part is transitive are usually called quasitransitive, and they possess many of the useful properties possessed by transitive relations. For example, continuous quasitransitive relations possess maximums on compact sets (Bergstrom, 1975).

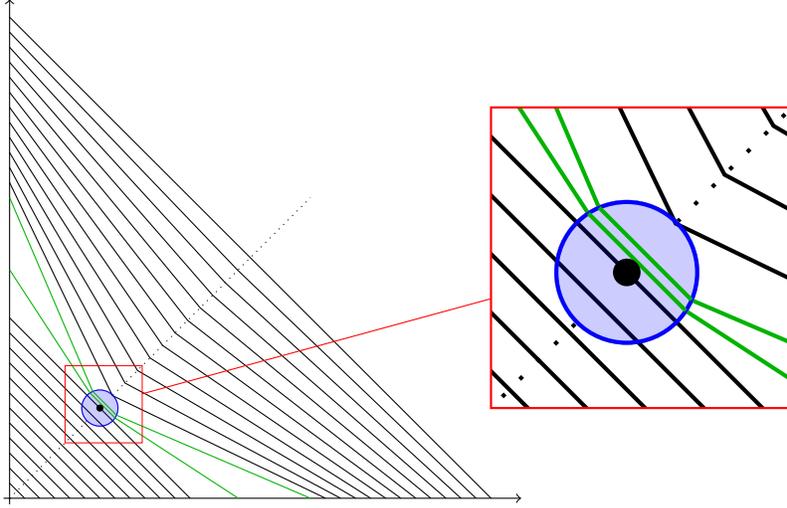


FIGURE 2. A transitive preference

There are several ways out of Grodal's example, if we wish to obtain a transitive limiting preference \succeq^* . We have seen a few in Sections 4 and 5. A rather general approach involves Lipschitz conditions on the class of preferences being considered.

Let us apply the Lipschitz approach to the environment described in Section 4.2. Fix $X = \mathbf{R}_+^d$ as the set of alternatives, such as a commodity space with d goods. Also fix $a, b \in \mathbf{R}_{++}$ with $a < b$, and consider the class of utility functions \mathcal{U} defined as the set of all the continuous utility functions $u : \mathbf{R}_+^d \rightarrow \mathbf{R}$ such that, for all i , and all $x_i, y_i \in \mathbf{R}_+$ with $x_i < y_i$,

$$a \cdot (y_i - x_i) \leq u(y_i, x_{-i}) - u(x_i, x_{-i}) \leq b \cdot (y_i - x_i)$$

for all $x_{-i} \in \mathbf{R}_+^{d-1}$. Hence, each utility function in \mathcal{U} is Lipschitz-bounded above and below. Clearly, every such utility function also describes a transitive, locally strict preference, because a is positive. And by the Arzela-Ascoli Theorem (Theorem 6.4 of Dugundji, 1966), \mathcal{U} is compact. Therefore one can appeal to Proposition 3, and if each rationalizing preference \succeq_n of the n -th experiment of an exhaustive sequence is included in $\Phi(\mathcal{U})$, then the limiting preference exists and is a member of $\Phi(\mathcal{U})$, and so is transitive.

APPENDIX A. PROOF OF LEMMA 1

Suppose, by means of contradiction, that there exist $x, y \in X$ for which $x \succeq_A y$ and $y \succ_B x$. By continuity of \succeq_B and local strictness of \succeq_A , we can assume without loss that $x \succ_A y$ and $y \succ_B x$. Then, by continuity of \succeq_A and \succeq_B , and by denseness of the collection $\{x_k, y_k : k \in \mathbf{N}\}$, there exists $a, b \in \{x_k, y_k : k \in \mathbf{N}\}$ such that $a \succ_A b$ and $b \succ_B a$. However, by completeness of \succeq , either $a \succeq b$, contradicting the implication $a \succeq b \implies a \succeq_B b$, or $b \succeq a$, contradicting the implication $b \succeq a \implies b \succeq_A a$.

APPENDIX B. PROOF OF THEOREM 1

The proof utilizes Lemma 1 and the following two elementary lemmas.

Lemma 4. *If $A \subseteq X \times X$, then $\{\succeq \in X \times X : A \subseteq \succeq\}$ is closed.*

Proof. Let $\{\succeq_n\}_{n \in \mathbf{N}}$ be a sequence in the set $\{\succeq \in X \times X : A \subseteq \succeq\}$ such that $\succeq_n \rightarrow \succeq$. Then for all $(x, y) \in A$, we have $x \succeq_n y$, hence $x \succeq y$. So $(x, y) \in \succeq$. \square

Lemma 5. *The set of all continuous binary relations on X is a compact metrizable space.*

Proof. See Theorem 2 in Chapter B of Hildenbrand (2015), or Corollary 3.95 of Aliprantis and Border (2006). \square

We now return to the main proof of Theorem 1. By assumption, \mathcal{P} is closed, and hence compact as a closed subset of a compact space by Lemma 5.

Let \succeq' be any complete binary relation (not necessarily in \mathcal{P}) such that for all n and all $\{x, y\} \in \Sigma_n$, $x = c(\{x, y\})$ if and only if $x \succeq' y$ (\succeq' is guaranteed to exist because the experiments are nested, $\Sigma_n \subseteq \Sigma_{n+1}$ for all n). Similarly, let \succeq'_n be the revealed preference relation that captures the observations made on the experiment Σ_n , that is, $x \succeq'_n y$ if and only if there is $\{x, y\} \in \Sigma_n$ with $x = c(x, y)$.

For every $n \in \mathbf{N}$, let $P_n = \{\succeq \in \mathcal{P} : \succeq'_n \subseteq \succeq\}$ be the set of relations in \mathcal{P} that rationalize the observed choices on Σ_n . Lemma 4 implies that P_n is closed, and hence compact. Thus, $\{P_n\}_{n \in \mathbf{N}}$ constitutes a decreasing sequence

of closed sets lying in the compact set P , and by the finite intersection property, $\bigcap_{n \in \mathbf{N}} P_n \neq \emptyset$.

So let $\succ^* \in \bigcap_{n \in \mathbf{N}} P_n$. We claim that $\bigcap_{n \in \mathbf{N}} P_n = \{\succ^*\}$. Take any $\succ \in \bigcap_{n \in \mathbf{N}} P_n$. By definition, for any binary decision problem $\{x, y\} \in \bigcup_{n \in \mathbf{N}} \Sigma_n$, if $x \succ' y$ then $x \succ^* y$ and $x \succ y$. Hence Lemma 1 implies $\succ = \succ^*$.

The result now follows as for each $n \in \mathbf{N}$, P_n is compact, and $\bigcap_{n \in \mathbf{N}} P_n = \{\succ^*\}$.

APPENDIX C. PROOF OF THEOREM 2

Throughout this proof, \succ^* is the subject's preference, \succ_n is the Kemeny-minimizing estimator for the n -th experiment Σ_n , and R_n is the revealed preference relation for that n -th experiment as described in Section 3. Since \succ^* is fixed throughout, we write $\mu(\succ, \succ^*)$ as $\mu(\succ)$.

To simplify notation, we also let $\bar{d}_n(\succ, R_n) = \frac{1}{n} |\succ \cap R_n|$, noting that

$$d_n(\succ, R_n) = \frac{1}{n} |R_n \setminus \succ| = 1 - \bar{d}_n(\succ, R_n).$$

In particular, \succ_n maximizes $\succ \mapsto \bar{d}_n(\succ, R_n)$.

The proof makes use of the following three lemmas.

Lemma 6. *For any preference \succ in the class \mathcal{P} , if \succ and \succ^* are distinct, then $\mu(\succ^*) > \mu(\succ)$.*

Proof. First, we observe that $\succ \neq \succ^*$ implies $\succ \neq \succ^*$. Indeed, suppose that $x, y \in X$ satisfies $x \succ y$ but $(x, y) \notin \succ^*$. By completeness, $y \succ^* x$. Then, by continuity, there are neighborhoods U and V of x and y respectively with $V \succ^* U$ (that is, every alternative in V is ranked strictly above every alternative in U according to \succ^*). Because $x \succ y$, by local strictness, there exists $(x', y') \in U \times V$ with $x' \succ y'$. Hence $x' \succ y'$ but $y' \succ^* x'$, so $\succ \neq \succ^*$.

Let $q(x, y)$ be a short notation for $q(\succ^*; x, y)$, and for a binary relation R , let $\mathbf{1}_R(x, y)$ if and only if $(x, y) \in R$. We show that $\mu(\succ^*) - \mu(\succ) > 0$ by the

sequence of inequalities below:

$$\begin{aligned}
\mu(\succ^*) - \mu(\succ) &= \int_{X \times X} [\mathbf{1}_{\succ^*}(x, y)q(x, y) - \mathbf{1}_{\succ}(x, y)q(x, y)] d\lambda(x, y) \\
&= \int_{X \times X} [\mathbf{1}_{\succ^* \setminus \succ}(x, y)q(x, y) - \mathbf{1}_{\succ \setminus \succ^*}(x, y)q(x, y)] d\lambda(x, y) \\
&= \int_{X \times X} [\mathbf{1}_{\succ^* \setminus \succ}(x, y)q(x, y) - \mathbf{1}_{\succ \setminus \succ^*}(y, x)q(y, x)] d\lambda(x) d\lambda(y) \\
&= \int_{X \times X} [\mathbf{1}_{\succ^* \setminus \succ}(x, y)q(x, y) - \mathbf{1}_{\succ^* \setminus \succ}(x, y)q(y, x)] d\lambda(x) d\lambda(y) \\
&= \int_{X \times X} \mathbf{1}_{\succ^* \setminus \succ}(x, y) [q(x, y) - q(y, x)] d\lambda(x) d\lambda(y) \\
&> 0.
\end{aligned}$$

The third equality obtains with a change of variables. The fourth equality uses the completeness of \succ^* and \succ , which means that $(y, x) \in \succ \setminus \succ^*$ if and only if $x \succ^* y$ and $x \not\succeq y$. The fifth equality follows as $\mathbf{1}_{\succ^* \setminus \succ}(x, y)$ and $\mathbf{1}_{\succ^* \setminus \succ}(x, y)$ are equal λ -almost surely. The final inequality owes to the fact that $q(x, y) > 1/2 > q(y, x)$ if $(x, y) \in \succ^*$. \square

Lemma 7. *The mapping $\succ \mapsto \mu(\succ)$ is continuous on \mathcal{P} .*

Proof. Let $\succ \in \mathcal{P}$ and $\{\succ_i\}_{i \in \mathbf{N}}$ be a sequence of preferences in \mathcal{P} with $\succ_i \rightarrow \succ$. If $x \succ_i y$ for infinitely many values of i , then $x \succeq y$. Hence, by completeness of the preference \succ , if the sequence of binary numbers $\{\mathbf{1}_{(x, y) \in \succ_i}\}_{i \in \mathbf{N}}$ diverges, then $x \sim y$.

Next, suppose $x \not\succeq y$. Then, there are two possibilities. Either $\mathbf{1}_{(x, y) \in \succ_i} = 1$ for all i large enough, then $x \succ y$. Or $\mathbf{1}_{(x, y) \in \succ_i} = 0$ for all i large enough, then, by completeness, $\mathbf{1}_{(y, x) \in \succ_i} = 1$ for all i large enough, and so $y \succ x$.

Finally, recall that, by Assumption (3), the set $\{(x, y) : x \sim y\}$ has λ -probability 0. Hence,

$$\begin{aligned}
\mu(\succeq) &= \int_{\{(x,y):x\sim y\}} \mathbf{1}_{(x,y)\in \succeq} q(\succeq^*; x, y) d\lambda(x, y) \\
&= \int_{\{(x,y):x\sim y\}} \lim_{n \rightarrow \infty} \mathbf{1}_{(x,y)\in \succsim_i} q(\succeq^*; x, y) d\lambda(x, y) \\
&= \lim_{i \rightarrow \infty} \int_{\{(x,y):x\sim y\}} \mathbf{1}_{(x,y)\in \succsim_i} q(\succeq^*; x, y) d\lambda(x, y) \\
&= \lim_{i \rightarrow \infty} \int \mathbf{1}_{\succsim_i} d\mu = \lim_{i \rightarrow \infty} \mu(\succsim_i),
\end{aligned}$$

where the interchange between the limit and integration operators follows from Lebesgue dominated convergence. \square

Lemma 8. *For all $\succeq^* \in \mathcal{P}$ and all $n \in \mathbf{N}$, the mapping $\succeq \mapsto \bar{d}_n(\succeq, R_n)$ defined on \mathcal{P} is almost surely continuous at \succeq^* .*

Proof. Let $\succeq \in \mathcal{P}$ and $\{\succsim_i\}_{i \in \mathbf{N}}$ be a sequence of preferences in \mathcal{P} with $\succsim_i \rightarrow \succeq$ as $i \rightarrow \infty$. Suppose that, for all decision problems $\{x_k, y_k\} \in \Sigma_n$, $x_k \approx y_k$. By the same argument as in Lemma 7, either $x_k \succsim_i y_k$ for i large enough and then $x_k \succ y_k$, or $y_k \succsim_i x_k$ for i large enough and then $y_k \succ x_k$. Hence, $\bar{d}_n(\succsim_i, R_n) \rightarrow \bar{d}_n(\succeq, R_n)$ as $i \rightarrow \infty$. And by assumption (3), with probability 1, it is the case that for all decision problems $\{x_k, y_k\} \in \Sigma_n$, $x_k \approx y_k$, which concludes the proof. \square

Recall that, by Assumption (2), \mathcal{P} is closed, and so compact as a closed subset of the set of all continuous binary relations on X , which itself is compact by Lemma 5 in Appendix B. That \mathcal{P} is compact and the continuity statement in Lemma 8 imply that $\bar{d}_n(\succeq, R_n)$ converges uniformly in probability to $\mu(\succeq)$ over the domain \mathcal{P} by uniform laws of large numbers (for example, Lemma 2.4 of Newey and McFadden (1994)). Finally, the mapping $\succeq \mapsto \mu(\succeq)$ is continuous on \mathcal{P} by Lemma 7, and is uniquely maximized at \succeq^* . Under the above conditions, standard consistency theorems apply, such as Theorem 2.4 of Newey and McFadden (1994), and thus $\{\succeq_n\}_{n \in \mathbf{N}}$ converges to \succeq^* in probability.

APPENDIX D. PROOF OF THEOREM 3

As in the proof of Theorem 2, throughout this proof, we fix the subject's preference \succeq^* , \succeq_n denotes the Kemeny-minimizing estimator for the size- n experiment Σ_n , and R_n is the revealed preference relation for Σ_n . We continue to write $\mu(\succeq, \succeq^*)$ as $\mu(\succeq)$, and we also let $\bar{d}_n(\succeq, R_n) = \frac{1}{n} |\succeq \cap R_n|$. Let $V_{\mathcal{P}}$ be the VC dimension of \mathcal{P} , let

$$C = 24^2 \frac{V_{\mathcal{P}} \log(4e^2)}{5n},$$

and, for each n , let

$$\begin{aligned} Z_n &= \sup \{ \bar{d}_n(\succeq, R_n) - \mu(\succeq) : \succeq \in \mathcal{P} \}, \\ Y_n &= \sup \{ \mu(\succeq) - \bar{d}_n(\succeq, R_n) : \succeq \in \mathcal{P} \}. \end{aligned}$$

By the bounded differences inequality (for example, Theorem 6.2 of Boucheron, Lugosi, and Massart, 2013),

$$\Pr(Z_n - \mathbf{E}Z_n > t \text{ or } Y_n - \mathbf{E}Y_n > t) \leq 2e^{-2t^2n}.$$

Moreover, by Theorem 13.7 of Boucheron, Lugosi, and Massart (2013),

$$\max \{ \sqrt{n} \mathbf{E}Z_n, \sqrt{n} \mathbf{E}Y_n \} \leq 72\sqrt{V_{\mathcal{P}} \log(4e^2)},$$

as long as $n \geq C$.

Hence on the event $\{Z_n - \mathbf{E}Z_n \leq t \text{ and } Y_n - \mathbf{E}Y_n \leq t\}$, we have

$$\begin{aligned} \mu(\succeq_n) &\geq \bar{d}_n(\succeq_n, R_n) - Z_n \\ &\geq \bar{d}_n(\succeq^*, R_n) - \left(72\sqrt{\frac{V_{\mathcal{P}} \log(4e^2)}{n}} + t \right) \\ &\geq \mu(\succeq^*) - Y_n - \left(72\sqrt{\frac{V_{\mathcal{P}} \log(4e^2)}{n}} + t \right) \\ &\geq \mu(\succeq^*) - 2 \left(72\sqrt{\frac{V_{\mathcal{P}} \log(4e^2)}{n}} + t \right). \end{aligned}$$

The first inequality follows from the definition of Z_n , the second from the definition of the event we are in, and the bound on $\sqrt{n} \mathbf{E}Z_n$, the third from the

definition of \succeq_n , the fourth from the definition of Y_n , and the final inequality from the bound on $\sqrt{n} \mathbf{E}Y_n$.

Thus, as long as

$$(3) \quad 2 \left(72 \sqrt{\frac{V_{\mathcal{P}} \log(4e^2)}{n}} + t \right) < r = r(\eta),$$

we know that $\rho(\succeq_n, \succeq^*) < \eta$.

Setting $\delta = 2e^{-2t^2n}$ so that $t^2 = \ln(2/\delta)/(2n)$, together with Equation (3), we get

$$144 \sqrt{\frac{V_{\mathcal{P}} \log(4e^2)}{n}} + 2 \sqrt{\frac{\ln(2/\delta)}{2n}} < r.$$

Hence,

$$N(\eta, \delta) \leq \max \left\{ r^{-2} (144 \sqrt{V_{\mathcal{P}} \log(4e^2)} + \sqrt{2 \ln(2/\delta)})^2, C \right\},$$

where the constant C does not depend on δ nor η , and hence, as $\delta \rightarrow 0$ and $\eta \rightarrow 0$,

$$N(\eta, \delta) = O \left(\frac{1}{r(\eta)^2} \ln \frac{1}{\delta} \right).$$

APPENDIX E. PROOF OF PROPOSITIONS 1 AND 3

Proposition 3 follows from the following result from Border and Segal (1994).

Theorem 4 (Theorem 8 of Border and Segal, 1994). *Let (X, d) be a locally compact and separable metric space and \mathcal{R} be the space of continuous preference relations on X , endowed with the topology of closed convergence. Let $C(X, d)$ be endowed with the topology of uniform convergence on compacta. If $\Phi(u)$ is locally strict, then Φ is continuous at u .*

We now prove Proposition 1.

Any nonconstant expected utility preference can be represented by a member of $\mathcal{V} \equiv \{u \in \mathbf{R}^d : \|u\| = 1 \text{ and } \sum_i u_i = 0\}$, which is a compact set. Thus, the set of functions $\{U_v : v \in \mathcal{V}\}$, where U_v is defined as $U_v(p) = v \cdot p$, is compact in the topology of compact convergence.²² Also, each nonconstant

²²Because Δ^{d-1} is compact, in the present case the topology of compact convergence a metric topology. Let $\{U_{v^n}\}_{n \in \mathbf{N}}$ be a sequence of functions where $v^n \in \mathcal{V}$, and let $\{v^{n_k}\}_{k \in \mathbf{N}}$ be a convergent subsequence that converges to v^* . Then $|U_{v^{n_k}}(p) - U_{v^*}(p)| = |(v^{n_k} - v^*) \cdot p| \leq \sqrt{\|v^{n_k} - v^*\| \|p\|} \leq \sqrt{\|v^{n_k} - v^*\|}$, where the first inequality is a Cauchy-Schwarz inequality.

expected utility preference is locally strict. To see this, take $p, q \in \Delta^{d-1}$ and suppose that $v \cdot p \geq v \cdot q$. Let $p^*, q^* \in \Delta^{d-1}$ for which $v \cdot p^* > v \cdot q^*$ (such a pair exists because v represents a nonconstant preference). Then, for any $\alpha \in (0, 1)$, $v \cdot (\alpha p^* + (1 - \alpha)p) > v \cdot (\alpha q^* + (1 - \alpha)q)$. Local strictness follows by choosing α arbitrarily small.

APPENDIX F. PROOF OF PROPOSITION 2

Recall that the set of alternatives X is the simplex Δ^{d-1} , and that we use for ρ the Hausdorff metric. It will be convenient to refer to the elements of the simplex by the generic symbols for alternatives x and y , as opposed to p and p' . The preference \succeq^* continues to denote the true preference of the subject. As in the proof of Theorems 2 and 3, we use $q(x, y)$ as a short notation for $q(\succeq^*; x, y)$, and for a binary relation R , we let $\mathbf{1}_R(x, y)$ if and only if $(x, y) \in R$.

Before we turn to the main proof, we perform preliminary computations with Lemma 9 and 10 below.

For these computations, let \succeq_A and \succeq_B be preferences in \mathcal{P} , and let $\eta > 0$. Suppose that there exists $(\tilde{x}, \tilde{y}) \in \succeq_A$ such that, if $x, y \in \Delta^{d-1}$ and $\|(\tilde{x}, \tilde{y}) - (x, y)\| < \eta$, then $(x, y) \notin \succeq_B$. We define

$$\eta' = \min \left\{ \frac{\eta}{4\sqrt{d(d-1)}}, \frac{1}{d} \right\},$$

and $\alpha = d\eta' \in (0, 1]$. Also, let

$$x^\alpha = \frac{\alpha}{d}\mathbf{1} + (1 - \alpha)\tilde{x} \quad \text{and} \quad y^\alpha = \frac{\alpha}{d}\mathbf{1} + (1 - \alpha)\tilde{y}.$$

Observe that $x(\alpha) \succeq_A y(\alpha)$.

Lemma 9. *The following inequality holds: $\|x^\alpha - \tilde{x}\| < \eta/4$.*

Proof. Since $\tilde{x} \in \Delta^{d-1}$, we have $\|(1/d)\mathbf{1} - \tilde{x}\| \leq \|(1/d)\mathbf{1} - (1, 0, \dots, 0)\|$. Therefore,

$$\begin{aligned} \|x^\alpha - \tilde{x}\| &= \alpha \|(1/d)\mathbf{1} - \tilde{x}\| \\ &\leq \alpha \sqrt{(1 - 1/d)^2 + (d-1)(1/d)^2} \\ &= \alpha \sqrt{1 - 2/d + d(1/d)^2} \\ &= \alpha \sqrt{1 - 1/d} = \eta' \sqrt{d(d-1)} \leq \eta/4. \end{aligned}$$

□

Next, let $v \in \mathbf{R}^d$ be such that $x \succeq_A y$ if and only if $v \cdot x \geq v \cdot y$, for all $x, y \in \Delta^{d-1}$. Such element v exists because $\succeq_A \in \mathcal{P}$. For $x \in \Delta^{d-1}$, we let $\mathcal{B}_\varepsilon(x)$ be the open ball of radius ε and center x in $\{z \in \mathbf{R}^d : \sum_{i=1}^d z_i = 1\}$, which is the affine span of the simplex Δ^{d-1} . For $x, y \in \Delta^{d-1}$, and $\varepsilon > 0$, let

$$\begin{aligned} \mathcal{B}_\varepsilon^+(x) &= \mathcal{B}_\varepsilon(x) \cap \{z \in \mathbf{R}^d : v \cdot z > v \cdot x\}, \text{ and} \\ \mathcal{B}_\varepsilon^-(y) &= \mathcal{B}_\varepsilon(y) \cap \{z \in \mathbf{R}^d : v \cdot z < v \cdot y\}. \end{aligned}$$

Lemma 10. *If $x \in \mathcal{B}_{\eta'}^+(x^\alpha)$ and $y \in \mathcal{B}_{\eta'}^-(y^\alpha)$, then $(x, y) \in \succ_A \setminus \succeq_B$.*

Proof. Clearly, if $x \in \mathcal{B}_{\eta'}^+(x^\alpha) \cap \Delta^{d-1}$ and $y \in \mathcal{B}_{\eta'}^-(y^\alpha) \cap \Delta^{d-1}$, then $v \cdot x > v \cdot x(\alpha) \geq v \cdot y(\alpha) > v \cdot y$, so that $x \succ_A y$. It remains to prove that

$$\mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha) \subseteq \Delta^{d-1} \times \Delta^{d-1},$$

and also that

$$[\mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)] \cap \succeq_B = \emptyset.$$

For the first claim, we show that if $z \in \mathbf{R}^d$ with $z \cdot \mathbf{1} = 1$ and $\|z - x^\alpha\| < \eta'$, then $z \in \Delta^{d-1}$. The argument is by contradiction, as having $z_i < 0$ for some index i would imply that

$$|\alpha/d| < |\alpha/d - z_i| \leq |x_i^\alpha - z_i| \leq \|z - x^\alpha\| < \eta',$$

contradicting the definition of α . A similar observation holds for y^α , so we obtain the inclusion

$$\mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha) \subseteq \mathcal{B}_{\eta'}(x^\alpha) \times \mathcal{B}_{\eta'}(y^\alpha) \subseteq \Delta^{d-1} \times \Delta^{d-1}.$$

For the second claim, note that if $x \in B_{\eta'}(x^\alpha)$ then $\|x - \tilde{x}\| \leq \|x - x^\alpha\| + \|x^\alpha - \tilde{x}\| \leq \eta/4 + \eta' < \eta/2$, where the second inequality owes to Lemma 9. Similarly, if $y \in B_{\eta'}(y^\alpha)$ then $\|y - \tilde{y}\| < \eta/2$. Thus, since we have assumed that, if $\|(\tilde{x}, \tilde{y}) - (x, y)\| < \eta$ then $(x, y) \notin \succeq_B$, we get

$$[\mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)] \cap \succeq_B \subseteq \mathcal{B}_{\eta'}(x^\alpha) \times \mathcal{B}_{\eta'}(y^\alpha) \cap \succeq_B = \emptyset.$$

□

Let us now return to the main proof of Proposition 2. Most of the proof concerns the estimation of $r(\eta)$, for η small. Let $\succeq' \in \mathcal{P}$, and let the Hausdorff distance between \succeq^* and \succeq' be η . There are two cases to consider: either there is an element in \succeq^* that is distance η to \succeq' , or there is an element in \succeq' that is distance η to \succeq^* .

Let us start with the first case, and apply the computations of Lemmas 9 and 10. Choose $(\tilde{x}, \tilde{y}) \in \succeq^*$ such that if, for $x, y \in \Delta^{d-1}$, $\|(\tilde{x}, \tilde{y}) - (x, y)\| < \eta$, then $(x, y) \notin \succeq'$. Define η' , α , x^α , y^α , \mathcal{B}^+ and \mathcal{B}^- as above, but using \succeq^* in place of \succeq_A , and using \succeq' in place of \succeq_B .

By the same argument as in the proof of Lemma 6, we get

$$\begin{aligned} \mu(\succeq^*) - \mu(\succeq') &= \int_{X \times X} \mathbf{1}_{\succeq^* \setminus \succeq'}(x, y) [q(x, y) - q(y, x)] \, d\lambda(x) \, d\lambda(y) \\ &= \int_{\succeq^* \setminus \succeq'} \underbrace{[q(x, y) - q(y, x)]}_{\geq 0} \, d\lambda(x) \, d\lambda(y) \\ &\geq \int_{\mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)} [q(x, y) - q(y, x)] \, d\lambda(x) \, d\lambda(y) \\ &\geq \inf \{q(x, y) - q(y, x) : (x, y) \in \mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)\} \\ &\quad \times \lambda(\mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)), \end{aligned}$$

where the third inequality owes to Lemma 10 and the fact that $q(x, y) \geq q(y, x)$ when $(x, y) \in \succeq^*$.

Note that $\mathcal{B}_{\eta'}^+(x^\alpha)$ and $\mathcal{B}_{\eta'}^-(y^\alpha)$ are $(d-1)$ -dimensional half balls of radius η' , and so the Lebesgue measure of $\mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)$ is

$$\left(\frac{\pi^{(d-1)/2}}{2\Gamma((d-1)/2 + 1)} (\eta')^{d-1} \right)^2,$$

where Γ is the Gamma function. Hence, if η is small enough to have $\eta' < 1/d$, we get

$$(4) \quad \lambda(\mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)) \propto (\eta')^{2(d-1)} = \left(\frac{\eta}{4\sqrt{d(d-1)}} \right)^{2(d-1)}.$$

We now estimate

$$\inf \{q(x, y) - q(y, x) : (x, y) \in \mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)\}.$$

To that effect, observe that when, by Lemma 10,

$$(x, y) \in \mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha) \subseteq \succ^* \setminus \succeq',$$

we have $(y, x) \in \succeq'$ by completeness of \succeq' . So given that the closest element in \succeq' to (x, y) is at distance η , the distance along the path (\tilde{x}, \tilde{y}) to (x^α, y^α) to (x, y) to (y, x) cannot be shorter. That is,

$$\begin{aligned} \eta = \rho((\tilde{x}, \tilde{y}), \succeq') &\leq \|(\tilde{x}, \tilde{y}) - (x^\alpha, y^\alpha)\| \\ &\quad + \|(x^\alpha, y^\alpha) - (x, y)\| \\ &\quad + \|(x, y) - (y, x)\| \\ &< \sqrt{2}\eta/4 + \sqrt{2}\eta' + \|(x, y) - (y, x)\|, \end{aligned}$$

where we have used Lemma 9 and the Pythagorean theorem to conclude that $\|(\tilde{x}, \tilde{y}) - (x^\alpha, y^\alpha)\| < \sqrt{2}\eta/4$. Hence, we get

$$\begin{aligned} \|(x, y) - (y, x)\| &> \eta - \frac{\eta}{2\sqrt{2}} - \frac{\eta}{2\sqrt{2}\sqrt{d(d-1)}} \\ &\geq \frac{(2\sqrt{2}-1)\sqrt{d(d-1)}-1}{2\sqrt{2}\sqrt{d(d-1)}}\eta \\ &\geq \frac{2\sqrt{2}-2}{2\sqrt{2}}\eta = \frac{\sqrt{2}-1}{\sqrt{2}}\eta, \end{aligned}$$

and we have

$$\|(x, y) - (y, x)\| = \sqrt{2}\|x - y\|.$$

Thus, by definition of the error probability function q and the monotonicity of f , for all $(x, y) \in \mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)$,

$$q(x, y) - q(y, x) \geq 1 - f\left(\frac{\sqrt{2}-1}{2}\eta\right).$$

Combining this estimation with Equation (4), and since $f(0) = 1$ and f is continuously differentiable at 0 with $f'(0) < 0$, we get that, as $\eta \rightarrow 0$,

$$\mu(\succeq^*) - \mu(\succeq') = \Omega\left(\eta \frac{\sqrt{2}-1}{2} \left(\frac{\eta}{4\sqrt{d(d-1)}}\right)^{2(d-1)}\right) = \Omega(\eta^{2d-1})$$

where the big Omega notation refers to the usual asymptotic lower bound notation.

We now turn to the other possibility, that the Hausdorff distance between \succeq^* and \succeq' is η because of an element in \succeq' that is distance η to \succeq^* . As before, choose $(\tilde{x}, \tilde{y}) \in \succeq'$ such that if, for $x, y \in \Delta^{d-1}$, $\|(\tilde{x}, \tilde{y}) - (x, y)\| < \eta$, then $(x, y) \notin \succeq^*$. Consider the definitions aforementioned for η' , α , x^α , y^α , \mathcal{B}^+ and \mathcal{B}^- , but with \succeq' in place of \succeq_A and \succeq^* in place of \succeq_B . By Lemma 10 we obtain that $\mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha) \subseteq \succeq' \setminus \succeq^*$.

Then, by a similar logic as above and again using the same argument as in the proof of Lemma 6, we get that

$$\begin{aligned} \mu(\succeq') - \mu(\succeq^*) &= \int_{X \times X} \mathbf{1}_{\succeq' \setminus \succeq^*}(x, y) [q(x, y) - q(y, x)] d\lambda(x) d\lambda(y) \\ &= \int_{\succeq' \setminus \succeq^*} [q(x, y) - q(y, x)] d\lambda(x) d\lambda(y) \\ &\leq \int_{\mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)} [q(x, y) - q(y, x)] d\lambda(x) d\lambda(y) \\ &\leq \sup \{q(x, y) - q(y, x) : (x, y) \in \mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)\} \\ &\quad \times \lambda(\mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)), \end{aligned}$$

where we note that $q(x, y) \leq q(y, x)$ when $(x, y) \notin \succeq^*$.

By a symmetric argument as above, we get that for all $(x, y) \in \mathcal{B}_{\eta'}^+(x^\alpha) \times \mathcal{B}_{\eta'}^-(y^\alpha)$,

$$q(x, y) - q(y, x) \leq f\left(\frac{\sqrt{2}-1}{2}\eta\right) - 1.$$

So, still by a symmetric argument, we get that, as $\eta \rightarrow 0$,

$$\mu(\succeq^*) - \mu(\succeq') = \Omega(\eta^{2d-1}).$$

The previous asymptotic bounds continue to hold for any $\succeq' \in \mathcal{P}$ with Hausdorff distance at least η to \succeq . Hence, as $\eta \rightarrow 0$,

$$\inf\{\mu(\succeq^*) - \mu(\succeq') : \rho(\succeq^*, \succeq') \geq \eta\} = \Omega(\eta^{2d-1}).$$

Finally, the proposition follows from applying Theorem 3 and observing that the VC dimension of \mathcal{P} is no greater than $d+1$ (and so finite) by Proposition 4.20 of Wainwright (2019).

APPENDIX G. PROOF OF LEMMA 2

Let $\{\succeq_n\}_{n \in \mathbf{N}}$ be a converging sequence of Grodal-transitive preferences that are strictly monotone with respect to \triangleright , and let \succeq^* be the limiting binary relation.

Recall that, by Lemma 5 in Appendix B, the set of continuous binary relations is closed, and so \succeq^* is continuous. Also, for each $x, y \in X$, either $x \succeq_n y$ or $y \succeq_n x$, so there is a subsequence $\{\succeq_{n_k}\}_{k \in \mathbf{N}}$ for which either $x \succeq_{n_k} y$ for all k , or for which $y \succeq_{n_k} x$ for all k . Hence, either $x \succeq y$ or $y \succeq x$, which makes \succeq^* complete. Hence, \succeq^* is a preference.

Suppose by means of contradiction that \succeq^* is not strictly monotone with respect to \triangleright . In that case, there are $x, y \in X$ for which $x \triangleright y$ and yet $y \succeq^* x$. Let $\{x_n\}_{n \in \mathbf{N}}$, $\{y_n\}_{n \in \mathbf{N}}$ be any sequences of alternatives in X that converge to x and y and respectively and for which $y_n \succeq_n x_n$ for all n (existence of sequences satisfying this property follows from the definition of closed convergence in Section 3). Because \triangleright is open, for n large enough, $x_n \triangleright y_n$, which contradicts the fact that \succeq_n is strictly monotone with respect to \triangleright . Hence, \succeq^* is strictly monotone.

Finally, we show that \succeq^* is Grodal-transitive. Suppose $x, y, z, w \in X$ satisfy $x \succeq^* y \succ^* z \succeq^* w$. Let $\{x_n\}_{n \in \mathbf{N}}, \{y_n\}_{n \in \mathbf{N}}, \{z_n\}_{n \in \mathbf{N}}, \{w_n\}_{n \in \mathbf{N}}$ be sequences of alternatives in X that converge to x, y, z, w respectively, and for which $x_n \succeq_n y_n$ and $z_n \succeq_n w_n$ (which, again, exist by the definition of closed convergence). Since $y \succ^* z$, for n large enough, $y_n \succ_n z_n$. Consequently, for n large, $x_n \succeq_n y_n \succ_n z_n \succeq_n w_n$, which, by Grodal-transitivity, implies $x_n \succeq_n w_n$, and so $x \succeq^* w$.

APPENDIX H. PROOF OF LEMMA 3

For any $x \in X$, let U_x be the set $\{y : y \succ x\}$. Let us show that for each $x, y \in X$, either $U_x \subseteq U_y$, or $U_y \subseteq U_x$. To see this, suppose by means of contradiction that there is $z \in U_x \setminus U_y$ and $w \in U_y \setminus U_x$. Then we have $y \succeq z \succ x$ and $x \succeq w \succ y$. Therefore, $x \succeq w \succ y \succeq z$, which implies $x \succeq z$ by Grodal-transitivity. This contradicts $z \succ x$.

Now, fix $x, y \in X$ such that $x \succeq y$, and fix a neighborhood V of (x, y) in the product space $X \times X$. By the lemma hypotheses, there exists $(x', y') \in V$ such that $x' \triangleright x$ and $y \triangleright y'$. Then, there are two possibilities: either either $U_{y'} \subseteq U_x$, or $U_x \subseteq U_{y'}$. In the first case, by monotonicity, $y' \succ y$, which implies $y \succ x$, contradicting $x \succeq y$. So, we must have $U_x \subseteq U_{y'}$. Then $x' \in U_x$, so $x' \in U_{y'}$, which implies $x' \succ y'$. Hence, \succeq is locally strict.

APPENDIX I. PROOF OF PROPOSITION 5

Here, the set of alternatives X is the positive orthant \mathbf{R}_{++}^d . As before, the preference \succeq^* refers to the true preference of the subject. Moreover, as in the proof of Theorems 2 and 3, we use $q(x, y)$ as a short notation for $q(\succeq^*; x, y)$, and for a binary relation R , we let $\mathbf{1}_R(x, y)$ if and only if $(x, y) \in R$.

For any $x \in \mathbf{R}_{++}^d$ and $\varepsilon > 0$, let $\mathcal{B}_\varepsilon(x)$ denote the open ball in \mathbf{R}_{++}^d of radius ε and center x . Also let

$$\mathcal{B}_\varepsilon^+(x) = \{z \in \mathcal{B}_\varepsilon(x) : z \gg x\},$$

and

$$\mathcal{B}_\varepsilon^-(x) = \{z \in \mathcal{B}_\varepsilon(x) : x \gg z\}.$$

Fix any two given preferences \succeq_A and \succeq_B in the class \mathcal{P} .

Suppose that there is $(\tilde{x}, \tilde{y}) \in K$ with $\tilde{x} \succeq_A \tilde{y}$ such that if, for any $x, y \in \mathbf{R}_{++}^d$, $\|(\tilde{x}, \tilde{y}) - (x, y)\| < \eta$, then $(x, y) \notin \succeq_B \cap (K \times K)^\theta$.

We first observe that, if $(x, y) \in \mathcal{B}_{\eta/2}^+(\tilde{x}) \times \mathcal{B}_{\eta/2}^-(\tilde{y})$, then $x \succ \tilde{x} \succeq \tilde{y} \succ y$, by monotonicity of the preference \succeq_A . Then, we remark that if $(x, y) \in \mathcal{B}_{\eta/2}^+(\tilde{x}) \times \mathcal{B}_{\eta/2}^-(\tilde{y})$, then

$$\|(\tilde{x}, \tilde{y}) - (x, y)\| < \sqrt{2}\eta/2 < \eta \leq \theta,$$

so $(x, y) \in (K \times K)^\theta$, while $(x, y) \notin \succeq_B$. Therefore, we obtain the inclusion

$$(5) \quad \mathcal{B}_{\eta/2}^+(\tilde{x}) \times \mathcal{B}_{\eta/2}^-(\tilde{y}) \subseteq (\succ_A \setminus \succeq_B) \cap (K \times K)^\theta.$$

Second, if $(x, y) \in \mathcal{B}_{\eta/2}^+(\tilde{x}) \times \mathcal{B}_{\eta/2}^-(\tilde{y})$, then

$$(6) \quad \|(x, y) - (y, x)\| \geq \frac{\sqrt{2} - 1}{\sqrt{2}}\eta.$$

To prove this inequality, note that the completeness of \succeq_B and Equation (5) implies $y \succeq_B x$, as $(x, y) \in (K \times K)^\theta$ ensures that $(x, y) \notin \succeq_B$. There are then two possibilities.

One possibility is that $(x, y) \notin (K \times K)^\theta$, which implies

$$\|(x, y) - (y, x)\| \geq \frac{\sqrt{2} - 1}{\sqrt{2}}\eta$$

as $\|(x, y) - (\tilde{x}, \tilde{y})\| < \sqrt{2}\eta/2$ by the Pythagorean theorem, while $\|(\tilde{x}, \tilde{y}) - (y, x)\| \geq \theta \geq \eta$.

The other possibility is that $(x, y) \in (K \times K)^\theta$, in which case

$$\eta < \|(\tilde{x}, \tilde{y}) - (x, y)\| + \|(x, y) - (y, x)\|,$$

as the closest element in $\succeq_B \cap (K \times K)^\theta$ to (\tilde{x}, \tilde{y}) is at distance at least η from (x, y) . Then $\|(\tilde{x}, \tilde{y}) - (x, y)\| < \eta/2$ implies that

$$\|(x, y) - (y, x)\| > \frac{\eta}{2} > \frac{\sqrt{2} - 1}{\sqrt{2}}\eta.$$

To complete the proof, fix $\succeq^* \in \mathcal{P}$ and consider $\succeq' \in \mathcal{P}$ with $\rho^{K, K^\theta}(\succeq^*, \succeq') \geq \eta$.

First, suppose that there exists $(\tilde{x}, \tilde{y}) \in K$ with $\tilde{x} \succeq^* \tilde{y}$ such that, whenever $\|(\tilde{x}, \tilde{y}) - (x, y)\| < \eta$ for any $x, y \in \mathbf{R}_{++}^d$, then $(x, y) \notin \succeq' \cap (K \times K)^\theta$. Recall that λ is the uniform probability measure on $K^{\theta/2}$, and that $\eta \leq \theta$. Then,

using Equations (5) and (6) together with the same argument as in the proof of Proposition 2, we get

$$\begin{aligned} \mu(\succeq^*) - \mu(\succeq') &= \int_{K^{\theta/2} \times K^{\theta/2}} \mathbf{1}_{\succeq^* \setminus \succeq'}(x, y) [q(x, y) - q(y, x)] d\lambda(x, y) \\ &\geq \int_{\mathcal{B}_{\eta/2}^+(\tilde{x}) \times \mathcal{B}_{\eta/2}^-(\tilde{y})} [q(x, y) - q(x, y)] d\lambda(x) d\lambda(y) \\ &\geq \lambda(\mathcal{B}_{\eta/2}^+(\tilde{x}) \times \mathcal{B}_{\eta/2}^-(\tilde{y})) \times \left[1 - f\left(\frac{\sqrt{2}-1}{2}\eta\right) \right]. \end{aligned}$$

Second, suppose on the other hand, that there exists $(\tilde{x}, \tilde{y}) \in K$ with $\tilde{x} \succeq' \tilde{y}$ such that, whenever $\|(\tilde{x}, \tilde{y}) - (x, y)\| < \eta$ for any $x, y \in \mathbf{R}_{++}^d$, then $(x, y) \notin \succeq^* \cap (K \times K)^\theta$. Then, by symmetric argument as in the first case above, we get

$$\begin{aligned} \mu(\succeq') - \mu(\succeq^*) &= \int_{K^{\theta/2} \times K^{\theta/2}} \mathbf{1}_{\succeq^* \setminus \succeq'}(x, y) [q(x, y) - q(y, x)] d\lambda(x, y) \\ &\leq \int_{\mathcal{B}_{\eta/2}^+(\tilde{x}) \times \mathcal{B}_{\eta/2}^-(\tilde{y})} [q(x, y) - q(x, y)] d\lambda(x) d\lambda(y) \\ &\leq \lambda(\mathcal{B}_{\eta/2}^+(\tilde{x}) \times \mathcal{B}_{\eta/2}^-(\tilde{y})) \times \left[f\left(\frac{\sqrt{2}-1}{2}\eta\right) - 1 \right]. \end{aligned}$$

So, in both cases, we obtain the inequality

$$\mu(\succeq^*) - \mu(\succeq') \geq \lambda(\mathcal{B}_{\eta/2}^+(\tilde{x}) \times \mathcal{B}_{\eta/2}^-(\tilde{y})) \cdot \left[1 - f\left(\frac{\sqrt{2}-1}{2}\eta\right) \right].$$

Finally, recall that $\mathcal{B}_{\eta/2}(\tilde{x})$ and $\mathcal{B}_{\eta/2}(\tilde{y})$ are both d -dimensional balls of radius $\eta/2$, and so each of $\mathcal{B}_{\eta/2}^+(\tilde{x})$ and $\mathcal{B}_{\eta/2}^-(\tilde{y})$ has Lebesgue measure equal to the volume of a d -ball of radius $\eta/2$ divided by 2^d , which is equal to

$$\frac{\pi^{d/2}}{4^d \cdot \Gamma\left(\frac{d}{2} + 1\right)} \eta^d.$$

Since λ is the uniform probability measure on $(K \times K)^{\theta/2}$, we get, as $\eta \rightarrow 0$,

$$\begin{aligned} \lambda(\mathcal{B}_{\eta/2}^+(\tilde{x}) \times \mathcal{B}_{\eta/2}^-(\tilde{y})) &\geq \frac{\pi^d}{16^d \cdot \Gamma\left(\frac{d}{2} + 1\right)^2 \left(\frac{\theta}{2} + 1\right)^{2d} \text{leb}(K)^2} \eta^{2d} \\ &= \Omega(\eta^{2d}), \end{aligned}$$

where Ω stands for the asymptotic lower bound. Then, using that f is differentiable at 0, that $f'(0) < 0$, and that $f(0) = 1$, we get, as $\eta \rightarrow 0$,

$$\inf \{ \mu(\succeq^*) - \mu(\succeq) : \rho^{K, K^0}(\succeq^*, \succeq) \geq \eta, \succeq \in \mathcal{P} \} = \Omega(\eta^{2d+1}),$$

which means that, as $\eta \rightarrow 0$,

$$r(\eta) = \Omega(\eta^{2d+1}).$$

The proposition then follows from a direct application of Theorem 3.

APPENDIX J. PROOF OF PROPOSITION 6

First, we prove that the set of alternatives \mathcal{M} meets Assumption (1).

That \mathcal{M} is locally compact follows from Theorem 1.8.3 of Schneider (1993), which demonstrates that the set of nonempty, convex subsets of the simplex is compact, and Theorem 18.4 of Willard (2004). The fact that \mathcal{M} is separable obtains from Theorem 3.85(3) of Aliprantis and Border (2006), together with the fact that a subset of a separable metric space is itself separable (Problem 16G Part 1 of Willard, 2004). Next, we show that \mathcal{M} is completely metrizable. This space is, by definition, metrizable. The Hausdorff metric is complete on the set of nonempty, closed, convex subsets of Δ^{d-1} ; let us call this set \mathcal{M}^* . This fact owes to a straightforward adaptation of Theorem 1.8.2 of Schneider (1993), together with the fact that \mathcal{M}^* is a Hausdorff-closed set in the space of compact, convex, nonempty subsets of $\{x \in R^d : \sum_i x_i = 1\}$, because a closed subset of a metric space is complete (Theorem 24.10 of Willard, 2004). Further, \mathcal{M} is a Hausdorff (relatively) open subset of \mathcal{M}^* . It then follows by Alexandroff's Theorem (Theorem 24.12 of Willard, 2004) that \mathcal{M} is completely metrizable.

Secondly, we show that the hypothesis of Lemma 2 is satisfied, that is, that the dominance relation \sqsupseteq is open.

This result can be obtained by means of a standard isometry between the set \mathcal{M} endowed with the Hausdorff metric, and the set of support functions of members of \mathcal{M} defined on \mathcal{U} (Lemma 8 of Dekel, Lipman, and Rustichini, 2001; p. 594 of Dekel, Lipman, Rustichini, and Sarver, 2007; Theorem 1.8.11 of Schneider, 1993). For a member m of \mathcal{M} , such a support function is written

$h_m(u) = \sup_{p \in m} u \cdot p$. We endow the set of these support functions with the sup-norm metric.²³ Observe that $m_A \sqsupset m_B$ if and only if for every $u \in \mathcal{U}$, $h_{m_A}(u) > h_{m_B}(u)$. In particular, since \mathcal{U} is compact, $m_A \sqsupset m_B$ if and only if there is $\varepsilon > 0$ for which $h_{m_A}(u) - h_{m_B}(u) > \varepsilon$.

For an element m of \mathcal{M} , let $\mathcal{B}_\delta(m)$ denote the open ball of radius δ centered on m . Suppose $m_A, m_B \in \mathcal{M}$ satisfy $m_A \sqsupset m_B$. Let $m'_A \in \mathcal{B}_{\frac{\varepsilon}{3}}(m_A)$ and $m'_B \in \mathcal{B}_{\frac{\varepsilon}{3}}(m_B)$. By the isometry aforementioned, for each $u \in \mathcal{U}$, $h_{m'_A}(u) > h_{m_A}(u) - \frac{\varepsilon}{3}$ and $h_{m'_B}(u) < h_{m_B}(u) + \frac{\varepsilon}{3}$. Then $h_{m'_A}(u) - h_{m'_B}(u) > h_{m_A}(u) - h_{m_B}(u) - \frac{2\varepsilon}{3} > \frac{\varepsilon}{3} > 0$, so that $m'_A \sqsupset m'_B$. Hence \sqsupset is open.

Thirdly, we show that the hypothesis of Lemma 3 is satisfied. Let $m_A \in \mathcal{M}$. For any $\alpha \in (0, 1)$, let m_B be the Minkowski sum of m_A and Δ^{d-1} , weighted by α and $1 - \alpha$ respectively, i.e., $m_B = \{\alpha p + (1 - \alpha)p' : p \in m_A, p' \in \Delta(X)\}$. Then, m_B is arbitrarily close to m_A by choosing α small enough, and $m_B \sqsupset m_A$. Similarly, fix p_0 in the interior of m_A , and for any $\alpha \in (0, 1)$, let $m_C = \alpha p_0 + (1 - \alpha)m_A$. Then, m_C is arbitrarily close to m_A by choosing α small enough, and $m_A \sqsupset m_C$.

Hence, \mathcal{M} meets Assumption (1), and by Lemmas 2 and 3, the class \mathcal{P} meets Assumption (2).

REFERENCES

- AFRIAT, S. N. (1967): “The Construction of Utility Functions from Expenditure Data,” *International Economic Review*, 8, 67–77.
- AHN, D. S., S. CHOI, D. GALE, AND S. KARIV (2014): “Estimating Ambiguity Aversion in a Portfolio Choice Experiment,” *Quantitative Economics*, 5, 195–223.
- ALIPRANTIS, C. D. AND K. BORDER (2006): *Infinite Dimensional Analysis: A Hitchhiker’s Guide*, Springer, 3rd ed.
- AMEMIYA, T. (1985): *Advanced Econometrics*, Harvard University Press.
- ANDREWS, D. W. AND P. GUGGENBERGER (2009): “Validity of Subsampling and “Plug-in Asymptotic” Inference for Parameters Defined by Moment Inequalities,” *Econometric Theory*, 25, 669–709.
- BASU, P. (2019): “Learnability and Stochastic Choice,” SSRN Working Paper No 3338991.
- BASU, P. AND F. ECHENIQUE (2018): “Learnability and Models of Decision Making under Uncertainty,” Working paper.

²³So that the distance between two functions f, g is given by $\sup_{u \in \mathcal{U}} |f(u) - g(u)|$.

- BERGSTROM, T. C. (1975): “Maximal Elements of Acyclic Relations on Compact Sets,” *Journal of Economic Theory*, 10, 403–404.
- BLUNDELL, R., M. BROWNING, AND I. CRAWFORD (2008): “Best Nonparametric Bounds on Demand Responses,” *Econometrica*, 76, 1227–1262.
- BLUNDELL, R., D. KRISTENSEN, AND R. L. MATZKIN (2010): “Stochastic Demand and Revealed Preference,” Working paper.
- BORDER, K. C. AND U. SEGAL (1994): “Dynamic Consistency Implies Approximately Expected Utility Preferences,” *Journal of Economic Theory*, 63, 170–188.
- BOUCHERON, S., G. LUGOSI, AND P. MASSART (2013): *Concentration Inequalities: A Nonasymptotic Theory of Independence*, Oxford University Press.
- BROWN, D. J. AND R. L. MATZKIN (1996): “Testable Restrictions on the Equilibrium Manifold,” *Econometrica*, 64, 1249–1262.
- CARVAJAL, A., R. DEB, J. FENSKE, AND J. K.-H. QUAH (2013): “Revealed Preference Tests of the Cournot Model,” *Econometrica*, 81, 2351–2379.
- CARVALHO, L., S. MEIER, AND S. W. WANG (2016): “Poverty and Economic Decision Making: Evidence from Changes in Financial Resources at Payday,” *American Economic Review*, 106, 260–284.
- CARVALHO, L. AND D. SILVERMAN (2019): “Complexity and Sophistication,” NBER working paper No. 26036.
- CHAMBERS, C. P., F. ECHENIQUE, AND E. SHMAYA (2014): “The Axiomatic Structure of Empirical Content,” *American Economic Review*, 104, 2303–19.
- CHAPMAN, J., M. DEAN, P. ORTOLEVA, E. SNOWBERG, AND C. CAMERER (2018): “Econographics,” NBER Working Paper No. w24931.
- CHAVAS, J.-P. AND T. L. COX (1993): “On Generalized Revealed Preference Analysis,” *The Quarterly Journal of Economics*, 108, 493–506.
- CHERCHYE, L., B. DE ROCK, AND F. VERMEULEN (2011): “The Revealed Preference Approach to Collective Consumption Behaviour: Testing and Sharing Rule Recovery,” *The Review of Economic Studies*, 78, 176–198.
- CHERNOZHUKOV, V., H. HONG, AND E. TAMER (2007): “Estimation and Confidence Regions for Parameter Sets in Econometric Models,” *Econometrica*, 75, 1243–1284.
- CHOI, S., S. KARIV, W. MÜLLER, AND D. SILVERMAN (2014): “Who Is (More) Rational?” *American Economic Review*, 104, 1518–1550.
- CLINTON, J., S. JACKMAN, AND D. RIVERS (2004): “The Statistical Analysis of Roll Call Data,” *American Political Science Review*, 98, 355370.
- DEKEL, E., B. L. LIPMAN, AND A. RUSTICHINI (2001): “Representing Preferences with a Unique Subjective State Space,” *Econometrica*, 69, 891–934.
- DEKEL, E., B. L. LIPMAN, A. RUSTICHINI, AND T. SARVER (2007): “Representing Preferences with a Unique Subjective State Space: A Corrigendum,” *Econometrica*, 75, 591–600.

- DI EWERT, W. E. (1973): “Afriat and Revealed Preference Theory,” *The Review of Economic Studies*, 40, 419–425.
- DUGUNDJI, J. (1966): *Topology*, Bacon.
- FALK, A., A. BECKER, T. DOHMEN, B. ENKE, D. HUFFMAN, AND U. SUNDE (2018): “Global Evidence on Economic Preferences,” *The Quarterly Journal of Economics*, 133, 1645–1692.
- FISHBURN, P. C. AND A. RUBINSTEIN (1982): “Time Preference,” *International Economic Review*, 23, 677–694.
- FORGES, F. AND E. MINELLI (2009): “Afriat’s Theorem for General Budget Sets,” *Journal of Economic Theory*, 144, 135–145.
- GORNO, L. (2019): “Revealed Preference and Identification,” *Journal of Economic Theory*, 183, 698–739.
- GRANT, S., J. KLINE, I. MENEGHEL, J. QUIGGIN, AND R. TOURKY (2016): “A Theory of Robust Experiments for Choice under Uncertainty,” *Journal of Economic Theory*, 165, 124–151.
- GREEN, R. C. AND S. SRIVASTAVA (1986): “Expected Utility Maximization and Demand Behavior,” *Journal of Economic Theory*, 38, 313–323.
- GRODAL, B. (1974): “A Note on the Space of Preference Relations,” *Journal of Mathematical Economics*, 1, 279–294.
- HALEVY, Y., D. PERSITZ, AND L. ZRILL (2018): “Parametric Recoverability of Preferences,” *Journal of Political Economy*, 126, 1558–1593.
- HILDENBRAND, W. (1970): “On Economies with Many Agents,” *Journal of Economic Theory*, 2, 161–188.
- (2015): *Core and Equilibria of a Large Economy*, Princeton University Press.
- JACKMAN, S. (2001): “Multidimensional Analysis of Roll Call Data via Bayesian Simulation: Identification, Estimation, Inference, and Model Checking,” *Political Analysis*, 9, 227241.
- KANNAI, Y. (1970): “Continuity Properties of the Core of a Market,” *Econometrica*, 38, 791–815.
- KEMENY, J. G. (1959): “Mathematics without Numbers,” *Daedalus*, 88, 577–591.
- KENDALL, M. G. (1938): “A New Measure of Rank Correlation,” *Biometrika*, 30, 81–93.
- KREPS, D. M. (1979): “A Representation Theorem for ‘Preference for Flexibility’,” *Econometrica*, 47, 565–577.
- KÜBLER, F. AND H. POLEMARCHAKIS (2017): “The Identification of Beliefs from Asset Demand,” *Econometrica*, 85, 1219–1238.
- MANSKI, C. F. (2003): *Partial Identification of Probability Distributions*, Springer Science & Business Media.
- MAS-COLELL, A. (1974): “Continuous and Smooth Consumers: Approximation Theorems,” *Journal of Economic Theory*, 8, 305–336.

- (1978): “On Revealed Preference Analysis,” *The Review of Economic Studies*, 45, 121–131.
- MATZKIN, R. L. (1991): “Axioms of Revealed Preference for Nonlinear Choice Sets,” *Econometrica*, 59, 1779–1786.
- (2003): “Nonparametric Estimation of Nonadditive Random Functions,” *Econometrica*, 71, 1339–1375.
- (2006): “Identification of Consumers Preferences When Their Choices Are Unobservable,” in *Rationality and Equilibrium: A Symposium in Honor of Marcel K. Richter*, ed. by C. D. Aliprantis, R. L. Matzkin, D. McFadden, J. C. Moore, and N. Yannelis, Springer, 195–215.
- (2007): “Heterogeneous Choice,” in *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress of the Econometric Society*, ed. by R. Blundell, W. Newey, and T. Persson, Cambridge University Press, vol. 43 of *Econometric Society Monographs*, chap. 4, 75–110.
- NEWAY, W. K. AND D. MCFADDEN (1994): “Large Sample Estimation and Hypothesis Testing,” in *Handbook of Econometrics*, ed. by R. Engle and D. McFadden, North Holland, vol. 4, chap. 36, 2111–2245.
- NISHIMURA, H., E. A. OK, AND J. K.-H. QUAH (2017): “A Comprehensive Approach to Revealed Preference Theory,” *American Economic Review*, 107, 1239–1263.
- POLEMARCHAKIS, H., L. SELDEN, AND X. SONG (2017): “The Identification of Attitudes Towards Ambiguity and Risk from Asset Demand,” Columbia Business School Research Paper No. 17–43.
- POOLE, K. T. AND H. ROSENTHAL (1985): “A Spatial Model for Legislative Roll Call Analysis,” *American Journal of Political Science*, 29, 357–384.
- RENY, P. J. (2015): “A Characterization of Rationalizable Consumer Behavior,” *Econometrica*, 83, 175–192.
- SCHNEIDER, R. (1993): *Convex Bodies: the Brunn-Minkowski Theory*, Cambridge University Press, 1st ed.
- VARIAN, H. R. (1982): “The Nonparametric Approach to Demand Analysis,” *Econometrica*, 50, 945–973.
- VON GAUDECKER, H.-M., A. VAN SOEST, AND E. WENGSTROM (2011): “Heterogeneity in Risky Choice Behavior in a Broad Population,” *American Economic Review*, 101, 664–94.
- WAINWRIGHT, M. J. (2019): *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, vol. 48, Cambridge University Press.
- WILLARD, S. (2004): *General Topology*, Courier Corporation.