

RECOVERING PREFERENCES FROM FINITE DATA

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

ABSTRACT. We study preferences estimated 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 and menu choice. Our framework unifies the revealed preference tradition with models that allow for errors.

1. INTRODUCTION

This paper provides conditions that guarantee asymptotic, large-sample, nonparametric recoverability of preferences from binary 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 approximate his preference over the entire set of alternatives to an arbitrary degree of precision. We study when this approximation is possible. 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

(Chambers) DEPARTMENT OF ECONOMICS, GEORGETOWN UNIVERSITY

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

(Lambert) DEPARTMENT OF ECONOMICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Date: First version: April 10, 2017; This version: March 23, 2021.

An early version circulated under the title “Preference Identification.” We are particularly grateful to Jeroen Swinkels, the coeditor and six anonymous referees for insightful comments and suggestions. We are thankful to Pathikrit Basu for finding a mistake in a prior draft of the paper and for valuable suggestions. We thank seminar participants at many institutions, the audiences of conferences and workshops at HEC Paris, Paris School of Economics, UC Berkeley, University of Pennsylvania, University of Warwick, University of York, Oxford University, and Virginia Tech. Echenique thanks the National Science Foundation for financial support (Grants SES-1558757 and CNS-518941), and the Simons Institute at UC Berkeley for its hospitality. Lambert thanks Microsoft Research and the Cowles Foundation at Yale University for their hospitality and financial support.

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 recovered in the limit as the available data grow large. These conditions concern both 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 goes to infinity. 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 and preferences over menus. In all these applications, we show how large finite experiments can approximate a preference of the kind that is appropriate for the application.

Our approach works with 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 Partial observability generally leads to partial identification and hinders estimation. In our framework, preferences can still be fully learned in spite of their partial observability.

We 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

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

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: Alice can be agnostic about how the alternatives presented to subjects are sampled, and how subjects are assumed to make mistakes. Specifically, the estimator does not require knowledge of sampling or error probabilities.

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 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; Varian, 1982), works with the classical model of consumer demand with linear budgets and finite 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.

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,

²For example, 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).

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. More recently, Kübler, Malhotra, and Polemarchakis (2020) consider identification of preferences from finite data, with emphasis on applications to demand and aggregate demand. As in the above works, our paper concerns the convergence of preferences that rationalize finite but unbounded data to the underlying data-generating preference, following large sample theory. We differ in three ways: first, we focus on data from pairwise choices instead of choice from budgets; second, we work with general environments beyond choice over commodities; third, we provide a broad sufficient condition on the class of preferences under consideration. Because we abstract away from specific economic environments, our results are applicable across different domains, such as choice of menus, under uncertainty, of intertemporal streams, lotteries, or consumption bundles.

Experimentalists and decision theorists also have an obvious interest in preference recovery from pairwise choices, but little is known about the behavior of preference estimates from finite data. Decision theory works often include a discussion of identification, but presume access to the agent's full preference relation. In contrast, we are interested in knowing if and when the preference relation can be inferred from the data. 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. Our research diverges from his, and from other studies of identification and integrability, in that we focus on large-sample estimation from finite data.

Another 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). However, these results on partial identification do not apply to the general environments we consider. Instead, we work with the classical large-sample theory for M-estimators (Amemiya, 1985; Newey and McFadden, 1994) but derive conditions for consistency without making additional compactness and equicontinuity assumptions (these assumptions are particularly substantial in a nonparametric estimation problem like ours, see Section 6.2 for a detailed discussion). Halevy, Persitz, and Zrill (2018) develop a method for estimating parametric models by minimizing the incompatibility of choice behavior with the proposed model, in the same spirit as our Kemeny-distance estimator. In contrast to our work, their methodology assumes data on choices from linear budgets, and adopts a money-metric version of Afriat’s and Varian’s “critical cost efficiency index” as a measure of distance. A crucial component of their analysis is to decompose measures in loss due to parametric misspecification, and loss due to inconsistency with rationality. 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). Recently, Basu and Echenique (2018) investigate the learnability of four standard models of choice under uncertainty using the notion of Probably Approximately Correct (PAC) learning from computational learning theory. Basu (2019) applies several other measures of model complexity to a study of stochastic choice.

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.

2. MODEL

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

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).

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 , denoted by $x \succeq y$. Associated to any given preference \succeq are its asymmetric part \succ (strict 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 $y \succeq x$. We do not require that preferences be transitive.

If Alice's goal is to infer Bob's preference from his behavior, it is clear that she must somehow discipline the set of preferences being considered. With partial observability, indifferences are unobservable, and 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 encompasses 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. Note that binary choices are used for simplicity: one could have more than two choices instead.

To study the large-sample properties of estimated preferences, 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.

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.

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. Such a model is random, with randomness stemming from sampling and subjects' mistakes. 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 observational data in which she has no control over the problems the subject faces. In the statistical tradition, Alice continues to assume that Bob has an underlying preference, but she allows for his behavior to deviate from what his preference dictates. Alice looks for a preference that fits the observed behavior the best.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

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.

$\{x, y\} \in \Sigma_n$, $c(\{x, y\}) \succeq x$ and $c(\{x, y\}) \succeq y$. Similarly, we say that \succeq *rationalizes* the choice function 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 Bob to choose in contradiction with his preference, but allows for Bob 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 also use λ to denote the product measure on $X \times X$. Deviations from the IID paradigm do not affect the results substantially.⁸

Bob's behavior is guided by his preference, but in every choice problem where Bob is not indifferent, he may make a mistake by choosing an alternative that is *not* preferred. 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. Assume that, for all $\succeq \in \mathcal{P}$, $q(\succeq; \cdot, \cdot)$ is measurable on $X \times X$.

When a subject with preference \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 make no particular assumption.

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, and Bob could in principle face the same choice problem more than once. By convention, we assume that Bob's choices remain the same over such repetitions, so that it is the choice function c that is random and not the choices themselves. However our assumptions guarantee that repetitions occur with probability zero, so that what is assumed about Bob's behavior under repetition is irrelevant to our results.

⁸Specifically, our results on estimation can be extended to more general probability measures on pairs of alternatives within each choice problem to account for correlations. In addition, these results do not depend on the order of choice problems in an experiment. Therefore, as experiments grow large, an experiment in which the empirical distribution of choice problems converges to a well-defined distribution becomes approximately equivalent to an experiment in which the choice problems are drawn IID.

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 6.1.

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.⁹ We connect the topology on preferences with familiar topologies on spaces of utility functions in Section 4, and, for the case of parameterized classes of preferences, with metrics on the parameter space.

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.

⁹Our 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. 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.

Our third assumption constrains the set of experiments. 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, starting from any 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 remain 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.

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, whose proof is in Appendix A.

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 $\succeq_A = \succeq_B$.*

Therefore, local strictness makes it possible to identify uniquely the subject's preference from an infinite set of observations. It is, however, not enough to estimate the subject's preference from finitely many observations: the requirement that \mathcal{P} be closed then becomes essential. To illustrate, suppose that $X = [0, 1]$, that \mathcal{P} is the set of all locally strict preferences, and that Bob prefers larger numbers, so that $x \succeq^* y$ if and only if $x \geq y$. Let \succsim_k be the preference defined by the piecewise linear utility function u_k with $u_k(0) = u_k(1) = 1$ and $u_k(1/k) = 0$; so, $u_k(x) = 1 - kx$ for $x \leq 1/k$ and $u_k(x) = -1/(k-1) + k/(k-1)x$ for $x > 1/k$. Note that \mathcal{P} includes \succeq^* and all \succsim_k for $k \geq 2$. Consider any set of growing experiments $\{\Sigma_n\}$ in which the choice problems use nonzero alternatives. In every experiment Σ_n , Bob's choices, if Bob chooses according to his preference, can be rationalized with \succsim_k for k large enough. As k grows large, \succsim_k does converge, but to a preference distinct from \succeq^* . So there is information regarding Bob's preference that is never fully learned. This problem owes to the fact that the set of locally strict preferences on $[0, 1]$ is not closed: the limiting preference \succsim_∞ is such that $x \succsim_\infty y$ if $x \geq y$ or if $x = 0$; it is thus not locally strict. More generally, no matter the subject's underlying preference and the choice of growing experiments, one can always find locally strict preferences that perfectly rationalize the subject's choices, and yet, as the size of the experiment grows large, converge to total indifference among all alternatives.¹⁰ The rationalizing preferences then convey no information about features of the subject's preference that are not directly observed. Closedness is thus a crucial assumption. As we show in Sections 4 and 5, closedness is not to be taken for granted, but it is fortunately satisfied in a range of important applications.

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

¹⁰See Section S.1 in the Online Appendix.

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.¹¹ 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

¹¹Recall that, in our model, Bob’s choices are consistent when facing the same choice problem more than once, but this assumption is without loss of generality as repetitions will occur with probability zero.

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, a rank distance measure, 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*.¹²

For a statistical preference model $(X, \mathcal{P}, \lambda, q)$, we shall need three assumptions. Assumptions (1) and (2) on the preference environment (X, \mathcal{P}) 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 import is that we almost never draw a decision problem that makes the subject indifferent. It also imposes that the sampling distribution has full support.¹³ Analogously to Assumption (3), this ensures that enough binary choice problems are probed. 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'. λ has full support on X and, for all $\succeq \in \mathcal{P}$, $\{(x, y) : x \sim y\}$ has λ -probability 0.

¹² $|R_n \setminus \succeq|$ 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,

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

with probability one, which justifies our Kemeny distance terminology.

¹³Full support means that there is no proper closed subset of the sample space that has probability 1.

Therefore, under Assumption (3'), repetitions of a binary choice problem occur with probability zero. Together with Assumption (2), Assumption (3') guarantees identification in the usual sense: the distribution of the data under the true preference is different than that at any other preference.¹⁴ This assumption allows us to control the problem of partial observability. If instead indifference were to occur with positive probability, then the model would only be partially identified. Note that Assumption (3') implies that λ is nonatomic, making it impossible to have finite sets of alternatives.

Under the above assumptions, the Kemeny-minimizing estimator is consistent. Recall that ρ denotes the 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, $\{\succeq_n\}_{n \in \mathbb{N}}$ 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. Requiring that the class of preferences be closed allows us 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.

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

¹⁴Identification in the usual sense is implied by the identification condition for consistency, proved in Lemma 7 of Appendix C. Although Assumption (3') is crucial, it is by itself not sufficient; for example, two preferences that differ at one point only cannot be distinguished when λ has full support. This case is ruled out by local strictness.

within a given approximation error. Our third result 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

$$(1) \quad r(\eta) = \inf \{ \mu(\succeq^*, \succeq^*) - \mu(\succeq, \succeq^*) : \succeq, \succeq^* \in \mathcal{P}, \rho(\succeq^*, \succeq) \geq \eta \}.$$

Roughly, the value of $r(\eta)$ captures the smallest possible probability that a subject makes 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 choice 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, then¹⁵*

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

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. This section handles preference environments derived from collections of utility functions. Specifically, we show how 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 . We are interested in sets of preferences \mathcal{P} that correspond to sets of utility functions. A *utility function* is any function $u : X \rightarrow \mathbf{R}$. We endow the space of utility functions with the topology of compact convergence.¹⁶ 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} .

The following proposition states conditions on the set of utility functions being considered under which our main results apply.

¹⁵The big O notation refers to the usual asymptotic upper bound.

¹⁶A 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).

Proposition 1. *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 1 is an immediate implication of Theorem 8 of Border and Segal (1994), who establish the continuity of Φ (see Appendix E).

4.1. Expected utility preferences. Expected utility preferences are an important case of preferences derived from utility functions. 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 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 indices $v \in \mathbf{R}^d$, with the property that $p \succeq p'$ if and only if $v \cdot p \geq v \cdot p'$, where $v \cdot p = \sum_{i=1}^d v_i p_i$ is the expected utility of lottery p . This preference is *nonconstant* if there is at least one pair $p, p' \in \Delta^{d-1}$ for which $p \succ p'$, or equivalently, if the vector of utility indices satisfies $v_i \neq v_j$ for some i, j . Of course, the vector of utility indices that defines a nonconstant expected utility preference is not unique. To ensure uniqueness, we normalize the utility indices by imposing that indices sum to zero and v be on the unit sphere. That is, each preference of \mathcal{P} is uniquely associated to a normalized vector of utility indices in the set¹⁷

$$\left\{ v \in \mathbf{R}^d : \sum_{i=1}^d v_i = 0, \|v\| = 1 \right\}.$$

The following result asserts that the standard expected utility model is situated within our framework.

Corollary 2. *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 Corollary 2 is in Appendix F.

Note that, although we gather data over the entire set of relevant alternatives—here the whole simplex—we could use a smaller set, because expected utility preferences are identified on a small set of lotteries. By learning the preference on this smaller

¹⁷ $\|\cdot\|$ denotes the Euclidean norm $\|x - y\| \equiv \sqrt{\sum_{i=1}^d (x_i - y_i)^2}$.

set, we can infer uniquely the preference on the full set. For example, suppose the alternatives that make the set of binary choice problems are from a subset of the simplex Δ^{d-1} that is convex and compact with nonempty interior, and that the family of binary choice problems are exhaustive relative to that subset. Then, Theorem 1 continues to hold even though preferences continue to be defined over Δ^{d-1} , that is, the rationalizing preference estimates continue to converge to a unique limiting preference. A similar observation applies to Theorems 2 and 3.

In the expected-utility model, we can further refine the convergence rates of Theorem 3, provided that we restrict attention to error probability functions q that are polynomially bounded in the following sense: there exists $C > 0$ and $k > 0$ such that, if lottery p is strictly preferred to lottery p' according to preference \succeq , then

$$(2) \quad q(\succeq; p, p') \geq \frac{1}{2} + C|u \cdot p - u \cdot p'|^k,$$

where u is the normalized vector of utility indices associated with preference \succeq . Equation (2) is satisfied, for example, if $q(\succeq; p, p')$ is lower-bounded strictly above $1/2$ over all \succeq and all p, p' such that $p \succ p'$. It is also satisfied if the probability of making an error decreases with the difference of utilities between the two lotteries, for example if, for a continuously differentiable and nondecreasing function $f : \mathbf{R}_+ \rightarrow \mathbf{R}_+$ with $f'(0) > 0$, we can write $q(\succeq; p, p') = 1/2 + f(u \cdot p - u \cdot p')$ for all \succeq and all p, p' such that $p \succ p'$ (u continues to denote the utility indices associated \succeq). In the former case, we can use $k = 0$, and in the latter case, $k = 1$.

Applying Theorem 3 also requires one to specify a compatible metric on preferences. Because the set of lotteries is compact and metrizable, one could measure the distance between two preferences by the Hausdorff metric. Here however, each preference is naturally represented by a finite-dimensional vector of utility indices, and it can be more straightforward to set, as distance between preferences, the distance between their associated utility indices. So, let $\rho(\succeq, \succe')$ be the Euclidean distance between the two normalized vectors of utility indices of \succeq and \succe' . It can be seen that ρ is a compatible metric (this fact holds quite generally, see Section 6.3).

The focus on error probability functions of the above form then allows for explicit convergence rates for the Kemeny-minimizing estimator, as below.

Corollary 3. *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 Equation (2), the Kemeny-minimizing estimator is consistent*

and, as $\eta \rightarrow 0$ and $\delta \rightarrow 0$,

$$N(\eta, \delta) = O\left(\frac{1}{\eta^{8(d-1)+4k}} \ln \frac{1}{\delta}\right).$$

One can rewrite Corollary 3 to provide a convergence rate of the form $O_p((1/n)^{1/d})$.¹⁸ The proof of Corollary 3 is in Appendix G.

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

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 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 a *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 $X \times X$.¹⁹ 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 $x, y, z, w \in X$,

¹⁸The O_p notation refers to the stochastic boundedness notation.

¹⁹That 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 $(x', y') \in V$, $x' \triangleright y'$.

$x \succeq y \succ z \succeq w$ 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 H.

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 preferences that fail 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 Grodal-transitive 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 I.

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.

Our discussion is summed up in the following result, which we shall see has several applications.

Proposition 2. *Suppose that X satisfies Assumption (1), and is endowed with an open dominance relation \triangleright . Suppose that \mathcal{P} consists of the preferences that are Grodal-transitive and strictly monotone with respect to \triangleright . Suppose further that for each $x \in X$, there exists $y, z \in X$ arbitrarily close to x such that $y \triangleright x$ and $x \triangleright z$. Then \mathcal{P} meets Assumption (2).*

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 defined 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 Corollary 4.²⁰

Corollary 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 (known) 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.

A similar remark holds for the case of intertemporal consumption, in which there is only one good to be consumed at different dates, and the alternative now gives the amount consumed at each date. If we hypothesize that individuals prefer to consume more at each date, the dominance relation \gg continues to apply to this environment. For example, one possible use of our theory is to recover discount factors from the

²⁰While 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.

data, or to check for distortions with respect to standard models such as exponential discounting.²¹

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 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 result establishes that the revealed preference framework applies to the menu preference environment.

Corollary 5. *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).*

²¹See the working paper version of our paper for details, Chambers et al. (2019).

The proof of Corollary 5 is in Appendix J.

5.3. 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 $p'_d = p_d - (d - 1)\varepsilon$, $p >_{\text{FSD}} p'$. As a consequence of Proposition 2, we obtain the following result.

Corollary 6. *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 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 a 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 .²²

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 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.

6.2. 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

²²To 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).

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 7. 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.

6.3. On parameterized sets of preferences. The metric ρ used in our results can be any metric on the space of preferences that is compatible with the topology of closed convergence. For instance, for compact sets of alternatives that satisfy Assumption (1), we can use for ρ the Hausdorff distance between subsets of $X \times X$.

When the space of preferences \mathcal{P} is parameterized, one may wish to work with a metric on preference parameters, such as the usual Euclidean distance when parameters are finite-dimensional vectors. Doing so is possible when the space of preferences and the space of parameters are homeomorphic, as formalized by the following lemma.

Lemma 4. *Let Θ be a compact space of parameters endowed with a metric D , and ϕ the mapping that sends every parameter in Θ to a preference in \mathcal{P} . Suppose that ϕ is one-to-one and onto, and is continuous. If $\rho(\succeq, \succeq')$ is defined as $D(\theta, \theta')$, with θ, θ' the parameters respectively associated with \succeq, \succeq' , then ρ is a compatible metric on \mathcal{P} .*

Lemma 4 follows directly from the observation that, under the above conditions, the inverse of ϕ is continuous by compactness of Θ .²³

To illustrate this result in a concrete case, let us return to the expected-utility preference environment of Section 4.1. Recall that the set of alternatives X is the set of lotteries over a finite collection of prizes $\{\pi_1, \dots, \pi_d\}$, represented as the $(d - 1)$ -dimensional simplex Δ^{d-1} , and \mathcal{P} is the set of nonconstant expected utility preferences. Each preference of \mathcal{P} is naturally parameterized by its normalized vector of utilities: using the notation in Lemma 4, let the space of parameters be

$$\Theta = \left\{ v \in \mathbf{R}^d : \sum_{i=1}^d v_i = 0, \|v\| = 1 \right\},$$

and D be the Euclidean distance.

Let ϕ the mapping that, to each $v \in \Theta$, associates the expected utility preference relation \succeq defined as $p \succeq p'$ if and only if $v \cdot p \geq v \cdot p'$. Let us write $\phi(v) = \Phi(U_v)$, where U_v is the utility function on X defined by $U_v(p) = v \cdot p$, and, as in Section 4, $\Phi(u)$ denotes the preference induced by utility function $u : X \mapsto \mathbf{R}$. First, we observe that $v \mapsto U_v$ is continuous, when, as in Section 4, the space of utility functions is endowed with the topology of compact convergence. Second, we observe that Φ is continuous by Theorem 8 of Border and Segal (1994) (see Appendix E), since by Corollary 2, $\Phi(u)$ is locally strict when u is defined as $u(p) = v \cdot p$ for $v \in \Theta$. Therefore Lemma 4 applies: if $\rho(\succeq, \succeq')$ is defined as the Euclidean distance between the utility indices of \succeq and \succeq' respectively, then ρ is a compatible metric on \mathcal{P} .

We use this fact in Proposition 3 to derive convergence rates of estimated preferences in terms of the Euclidean distance on utility indexes.

6.4. 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

²³Hence, any set of preferences that is open in the topology of closed convergence contains an open ball under the metric ρ , and conversely, any open ball of preferences under the metric ρ contains an open set in the topology of closed convergence. Therefore, the topology of closed convergence and the metric topology induced by ρ are identical.

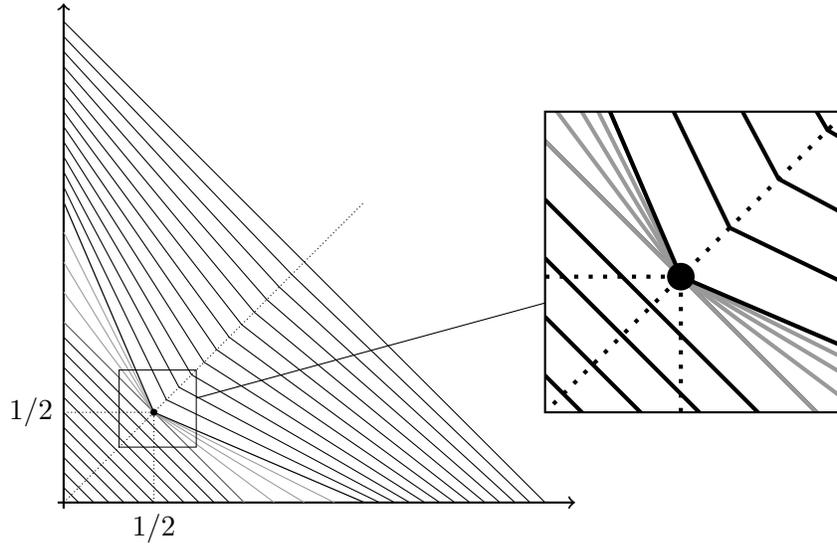


FIGURE 1. A nontransitive preference.

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 gray 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 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

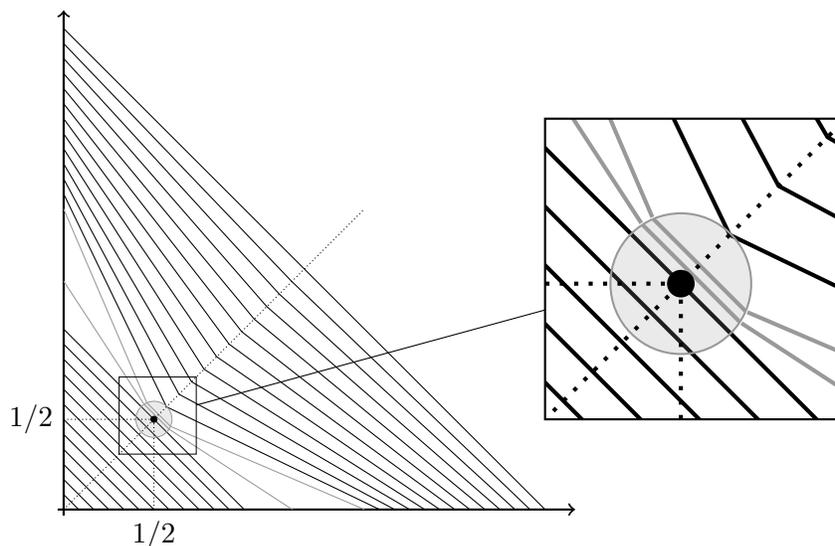


FIGURE 2. A transitive preference.

strictly monotone preferences, \succsim_n , $n \in \mathbf{N}$, where each \succsim_n is transitive, and $\succsim_n \rightarrow \succsim^*$, but \succsim^* is not transitive.

It is generally true that if each \succsim_n (the strict part of \succsim_n) is transitive, then \succsim^* will be transitive as well (see Grodal, 1974), but in some cases one may desire full transitivity of \succsim^* .²⁴ The central element of the example above is that the indifference curves get “squeezed” together too rapidly.

There are several ways out of Grodal’s example, if we wish to obtain a transitive limiting preference \succsim^* . 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. 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 $u : \mathbf{R}_+^d \rightarrow \mathbf{R}$ with the normalization $u(0) = 0$ and 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

²⁴Relations 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).

preference, because a is positive. And by the Arzela-Ascoli Theorem (Theorem 6.4, p. 267, of Dugundji, 1966), \mathcal{U} is compact. Therefore one can appeal to Proposition 1, 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 $a \succ_A b$ and $b \succ_B a \implies a \succeq_B b$, or $b \succeq a$, contradicting the implication $b \succ_B a \implies b \succeq_A a$.

APPENDIX B. PROOF OF THEOREM 1

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

Lemma 5. *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 6. *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 6.

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

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

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 7. *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 that $\succ^* \setminus \succ$ contains a nonempty open set. Indeed, suppose first 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 $x' \succ y'$. By continuity again there are neighborhoods $U' \ni x'$ and $V' \ni y'$ with $U' \succ V'$. Thus $(y', x') \in (V \cap V') \times (U \cap U') \subseteq \succ^* \setminus \succ$. Second, consider the case where there is $x, y \in X$ with $x \succ^* y$ but $(x, y) \notin \succ$. This means that $y \succ x$, so there are neighborhoods $U \ni x$ and $V \ni y$ with $V \succ U$. By local strictness, there is $(x', y') \in U \times V$ with $x' \succ^* y'$. Thus $(y', x') \in \succ \setminus \succ^*$ and we are in the situation above.

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 $(x, y) \in \succ^*$, and that there is an open set on which $\mathbf{1}_{\succ^* \setminus \succ} = 1$. \square

Lemma 8. *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 $x \succ y$. 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,

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

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

Lemma 9. *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 $\{\succeq_i\}_{i \in \mathbf{N}}$ be a sequence of preferences in \mathcal{P} with $\succeq_i \rightarrow \succeq$ as $i \rightarrow \infty$. Suppose that, for all decision problems $\{x_k, y_k\} \in \Sigma_n$, $x_k \sim y_k$. By the same argument as in Lemma 8, either $x_k \succeq_i y_k$ for i large enough and then $x_k \succ y_k$, or $y_k \succeq_i x_k$ for i large enough and then $y_k \succ x_k$. Hence, $\bar{d}_n(\succeq_i, R_n) \rightarrow \bar{d}_n(\succeq, R_n)$ as $i \rightarrow \infty$. And by assumption (3), with probability 1, it is the case that for all decision problems $\{x_k, y_k\} \in \Sigma_n$, $x_k \sim y_k$, 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 6 in Appendix B. That \mathcal{P} is compact and the continuity statement in Lemma 9 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 8, 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),

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

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

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

as long as $n \geq C$.

Hence on the event $\{Z_n - \mathbb{E}Z_n \leq t \text{ and } Y_n - \mathbb{E}Y_n \leq t\}$, 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 $\sqrt{n} \mathbb{E}Z_n$, the third from the definition of \succeq_n , the fourth from the definition of Y_n , and the final inequality from the bound on $\sqrt{n} \mathbb{E}Y_n$.

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 PROPOSITION 1

Proposition 1 follows from the following result by Border and Segal (1994).

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

Under the conditions stated in Proposition 1, Φ is continuous on \mathcal{U} by Theorem 4, so $\Phi(\mathcal{U})$ is compact as the continuous image of a compact set, and hence meets Assumption (2).

APPENDIX F. PROOF OF COROLLARY 2

As explained in Section 4.1, 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 in the Euclidean topology (recall that $\|\cdot\|$ refers to the Euclidean norm).

First, observe that the set of functions $\{U_v : v \in \mathcal{V}\}$, where U_v is defined as $U_v(p) = v \cdot p$, is sequentially compact in the topology of compact convergence. Indeed, 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. Because Δ^{d-1} is compact, in the present case the topology of compact convergence is a metric topology, and the notions of compactness and sequential compactness are equivalent, so $\{U_v : v \in \mathcal{V}\}$ is compact.

Second, each nonconstant expected utility preference is locally strict. Indeed, take $p, p' \in \Delta^{d-1}$ and suppose that $v \cdot p \geq v \cdot p'$. Let $p_*, p'_* \in \Delta^{d-1}$ for which $v \cdot p_* > v \cdot p'_*$

(such a pair exists because v represents a nonconstant preference). Then, for any $\alpha \in (0, 1)$, $v \cdot (\alpha p_* + (1 - \alpha)p) > v \cdot (\alpha p'_* + (1 - \alpha)p')$. Local strictness follows by choosing α arbitrarily small.

Corollary 2 then follows from applying Proposition 1.

APPENDIX G. PROOF OF COROLLARY 3

The proof proceeds by computing an asymptotic lower bound on $r(\eta)$ defined in Equation (1) in Section 3.2 and then applying Theorem 3.

Recall that X is the simplex Δ^{d-1} , and \mathcal{P} is the set of nonconstant expected utility preferences. 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' . Throughout, each preference of \mathcal{P} is associated to a unique normalized vector of utility indices in the set $\{v \in \mathbf{R}^d : \sum_{i=1}^d v_i = 0, \|v\| = 1\}$. For $x \in \Delta^{d-1}$ and $\varepsilon > 0$, 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} .

Let us start with the following lemma.

Lemma 10. *Let $0 < \eta < 1$ and $\succeq_u, \succeq_v \in \mathcal{P}$ with $\rho(\succeq_u, \succeq_v) \geq \eta$. There exists $x_0, y_0 \in X$ such that, for all $x \in \mathcal{B}_{\eta'}(x_0)$ and $y \in \mathcal{B}_{\eta'}(y_0)$,*

$$\begin{aligned} u \cdot x &\geq u \cdot y + \frac{\eta^2}{4d}, \\ v \cdot y &\geq v \cdot x + \frac{\eta^2}{4d}. \end{aligned}$$

where u and v are the normalized utility indices associated respectively with \succeq_u and \succeq_v , and $\eta' \equiv \eta^2/(10d)$. In addition, $\mathcal{B}_{\eta'}(x_0) \times \mathcal{B}_{\eta'}(y_0) \subset X \times X$.

Proof. Let

$$x_0 = \frac{1}{d}\mathbf{1} + \left(\frac{1}{d} - \eta'\right)u, \quad \text{and} \quad y_0 = \frac{1}{d}\mathbf{1} + \left(\frac{1}{d} - \eta'\right)v.$$

The inequality $\rho(\succeq_u, \succeq_v) \geq \eta$ is equivalent to $\|u - v\| \geq \eta$, or $u \cdot v \leq 1 - \eta^2/2$. Let $x \in \mathcal{B}_{\eta'}(x_0)$ and $y \in \mathcal{B}_{\eta'}(y_0)$. The following sequence of inequalities obtains:

$$\begin{aligned}
u \cdot x &= u \cdot (x - x_0) + u \cdot x_0 \\
&\geq u \cdot x_0 - \eta' \\
&\geq u \cdot y_0 + \left(\frac{1}{d} - \eta'\right) \frac{\eta^2}{2} - \eta' \\
&= u \cdot (y_0 - y) + u \cdot y + \left(\frac{1}{d} - \eta'\right) \frac{\eta^2}{2} - \eta' \\
&\geq u \cdot y + \left(\frac{1}{d} - \eta'\right) \frac{\eta^2}{2} - 2\eta' \\
&\geq u \cdot y + \frac{\eta^2}{4d}.
\end{aligned}$$

The first inequality owes to the fact that $|u \cdot (x - x_0)| \leq \|u\| \|x - x_0\| \leq \eta'$. Similarly, $|u \cdot (y_0 - y)| \leq \eta'$ which yields the third inequality. The second inequality comes from $u \cdot x_0 = 1/d - \eta'$ and

$$u \cdot y_0 = \left(\frac{1}{d} - \eta'\right) u \cdot v \leq \left(\frac{1}{d} - \eta'\right) - \left(\frac{1}{d} - \eta'\right) \frac{\eta^2}{2}.$$

The fourth inequality owes to the fact that $\eta < 1$ and

$$\left(\frac{1}{d} - \eta'\right) \frac{\eta^2}{2} - 2\eta' = \frac{\eta^2}{4d} + \frac{\eta^2 - \eta^4}{20d} \geq \frac{\eta^2}{4d}.$$

By a symmetric argument, we also have

$$v \cdot y \geq v \cdot x + \frac{\eta^2}{4d}.$$

Finally, observe that η' is chosen small enough to ensure that both $\mathcal{B}_{\eta'}(x_0)$ and $\mathcal{B}_{\eta'}(y_0)$ are included in the simplex Δ^{d-1} . \square

We now return to the main proof. Let us fix the subject's preference \succeq^* , and let \succeq be any preference of \mathcal{P} with $\rho(\succeq^*, \succeq) \geq \eta$, with $0 < \eta < 1$. As in the proofs 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 $\mathbf{1}_R(x, y) = 1$ if and only if $(x, y) \in R$.

We established in the proof of Theorem 2 that

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

By Lemma 10, there exists $x_0, y_0 \in X$ such that $\mathcal{B}_{\eta'}(x_0) \times \mathcal{B}_{\eta'}(y_0) \subset X \times X$, and if $(x, y) \in \mathcal{B}_{\eta'}(x_0) \times \mathcal{B}_{\eta'}(y_0)$ (with $\eta' \equiv \eta^2/(10d)$) then $x \succ^* y$ while $y \succ x$. Also, if $x \succ^* y$, then $q(x, y) - q(y, x) \geq 0$. Hence,

$$\begin{aligned} \mu(\succeq^*) - \mu(\succeq) &= \int_{\succeq^* \setminus \succeq} [q(x, y) - q(y, x)] d\lambda(x, y) \\ &\geq \int_{\mathcal{B}_{\eta'}(x_0) \times \mathcal{B}_{\eta'}(y_0)} [q(x, y) - q(y, x)] d\lambda(x, y) \\ &\geq \inf \{q(x, y) - q(y, x) : (x, y) \in \mathcal{B}_{\eta'}(x_0) \times \mathcal{B}_{\eta'}(y_0)\} \\ &\quad \times \lambda(\mathcal{B}_{\eta'}(x_0)) \times \lambda(\mathcal{B}_{\eta'}(y_0)). \end{aligned}$$

The Lebesgue measure of each of the $(d-1)$ -dimensional balls $\mathcal{B}_{\eta'}(x_0)$ and $\mathcal{B}_{\eta'}(y_0)$ is

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

where Γ is the Gamma function, so

$$\lambda(\mathcal{B}_{\eta'}(x_0)) \times \lambda(\mathcal{B}_{\eta'}(y_0)) = \Omega(\eta^{4(d-1)})$$

as $\eta \rightarrow 0$, where the big Omega notation refers to the asymptotic lower bound.

Since $x \in \mathcal{B}_{\eta'}(x_0)$ and $y \in \mathcal{B}_{\eta'}(y_0)$ implies $x \succ^* y$, by Equation (2),

$$q(x, y) - q(y, x) \geq 2C|u \cdot x - u \cdot y|^k,$$

where u is the normalized vector of utility indices associated with \succeq^* . By Lemma 10, we have

$$u \cdot x - u \cdot y \geq \frac{\eta^2}{4d}$$

and hence,

$$\inf \{q(x, y) - q(y, x) : (x, y) \in \mathcal{B}_{\eta'}(x_0) \times \mathcal{B}_{\eta'}(y_0)\} \geq 2C \left(\frac{\eta^2}{4d}\right)^k = \Omega(\eta^{2k})$$

as $\eta \rightarrow 0$.

Overall, we get $\mu(\succeq^*) - \mu(\succeq) = \Omega(\eta^{4(d-1)+2k})$, and thus $r(\eta) = \Omega(\eta^{4(d-1)+2k})$. 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), we have

$$N(\eta, \delta) = O\left(\frac{1}{\eta^{8(d-1)+4k}} \ln \frac{1}{\delta}\right).$$

APPENDIX H. 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 6 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 I. 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 J. PROOF OF COROLLARY 5

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 \sqsupset 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.²⁵ 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 $u \in \mathcal{U}$, $h_{m'_A}(u) > h_A(u) - \frac{\varepsilon}{3}$ and $h_{m'_B}(u) < h_B(u) + \frac{\varepsilon}{3}$. Then $h_{m'_A}(u) - h_{m'_B}(u) > h_A(u) - h_B(u) - \frac{2\varepsilon}{3} > \frac{\varepsilon}{3} > 0$, 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

²⁵So that the distance between two functions f, g is given by $\sup_{u \in \mathcal{U}} |f(u) - g(u)|$.

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,” California Institute of Technology Working Paper.
- 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.
- CARVALHO, L., S. MEIER, AND S. W. WANG (2016): “Poverty and Economic Decision Making: Evidence from Changes in Financial Resources at Payday,” *The 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 N. LAMBERT (2019): “Recovering preferences from finite data,” arXiv:1909.05457.

- CHAMBERS, C. P., F. ECHENIQUE, AND E. SHMAYA (2014): “The Axiomatic Structure of Empirical Content,” *The 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?” *The American Economic Review*, 104, 1518–1550.
- CLINTON, J., S. JACKMAN, AND D. RIVERS (2004): “The Statistical Analysis of Roll Call Data,” *The American Political Science Review*, 98, 355–370.
- 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.
- 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.
- 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, 227–241.
- 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., R. MALHOTRA, AND H. POLEMARCHAKIS (2020): “Identification of Preferences, Demand and Equilibrium with Finite Data,” University of Warwick Economics Research Paper No. 1290.
- 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.
- 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,” *The 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.
- 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,” *The American Economic Review*, 101, 664–94.
- WAINWRIGHT, M. J. (2019): *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, Cambridge University Press.
- WILLARD, S. (2004): *General Topology*, Courier Corporation.