

Learnability and Stochastic Choice

Pathikrit Basu*

September 9, 2019

Abstract

In this paper, we study a non-parametric approach to prediction in stochastic choice models in economics. We apply techniques from statistical learning theory to study the problem of learning choice probabilities. A model of stochastic choice is said to be learnable if there exist learning rules defined on choice data that are uniformly consistent. We construct learning rules via the procedure of empirical risk minimization, where risk is defined in terms of incentive compatible scoring rules. This approach involves mild distributional assumptions on the model, with the main requirement being a constraint on the complexity of the stochastic choice model. Our main result shows that in the context of stochastic preferences, when the VC dimension of a preference model is finite, then we can recover the distribution over preferences (heterogeneity) in the uniform sense. We provide several applications and obtain general rates of convergence based on sample complexity. Under additional conditions, we establish a converse result for the non-learnability of models. An implication of our non-learnability result is that for any learning rule and sample size, the worst case expected distance between true and estimated choice probabilities is at least half times the maximum possible distance. In binary choice settings, this expected error lower bound is above the error from always predicting that each alternative will be chosen with equal probability. Lastly, we also address recoverability of cognitive heterogeneity and non-uniform learning.

1 Introduction

The study of stochastic choice models in economics seeks to understand aggregate demand behaviour and preference heterogeneity in markets ([McFadden \[1978\]](#), [Berry et al.](#)

*E-mail : pathikritbasu@gmail.com. I would like to thank Kalyan Chatterjee, Federico Echenique, Bob Sherman, Matt Shum and Adam Wierman for comments and suggestions.

[1995]). A typical model of stochastic choice involves a set of alternatives from which an agent makes a choice and based on a set of underlying characteristics (\mathcal{X}), the model prescribes the probabilities with which various items (A) would be chosen. This is defined in terms of a stochastic choice function $\sigma : \mathcal{X} \rightarrow \Delta(A)$. The probabilities $\sigma(x)$ can be interpreted as the proportion of times a particular alternative is chosen compared to other feasible alternatives. Yet another interpretation involves the choice probabilities resulting from randomisation on part of the agent. The nature of this randomisation depends on the context and details of the decision making process (Manzini and Mariotti [2007], Fudenberg et al. [2015]). For example, in random utility models, the choice probabilities depend on structural parameters of the agents' preferences and idiosyncratic factors, in the form of random utility shocks, that may cause preferences to alter. In this paper, we study a novel approach to estimating choice probabilities, which correspond to a model of stochastic choice, based on data available on agents' choices among alternatives.

We study the problem of learning the map σ from finite data on choices. Hence, each data point consists of a pair (x_i, a_i) where a_i is the alternative chosen when the characteristics were given by the vector x_i . We are interested in learning rules defined on choice data which lead to accurate estimation of choice probabilities and the criterion we require is *uniform consistency* (Valiant [1984], Vapnik [1998], Kearns and Schapire [1994], Yamanishi [1992]). This requires that for any given precision parameters $\varepsilon, \delta > 0$, there exists a fixed data size $N(\varepsilon, \delta)$, such that if the analyst were to apply the learning rule to a data set of size above $N(\varepsilon, \delta)$, the probability would be at least $1 - \delta$ that the stochastic choice function conjectured by the learning rule would be close to the true stochastic choice function by at most ε . Moreover, this holds true irrespective of the process that generates choice problems and the true function σ_0 . Hence, working with amount of data at least $N(\varepsilon, \delta)$ allows for robust estimates of the choice probabilities, when learning rules are uniformly consistent. This is a desirable feature of the notion of consistency considered here. In many contexts, it is natural to assume that the analyst knows neither the distribution over characteristics nor the true stochastic choice function, and may wish to work with data sizes that guarantee, in a robust manner, a certain degree of precision in estimation. The minimum number of samples $N(\varepsilon, \delta)$, which provides such a guarantee is called the *sample complexity*¹ of the model of stochastic choice and is indeed a central object of study in the present paper.

¹The terminology is borrowed from the learning theory literature in computer science. See, for example, Shalev-Shwartz and Ben-David [2014].

The estimation strategy developed in this paper relies on the principle of empirical risk minimization (Vapnik [1998]). This involves setting up a risk function to evaluate the goodness of fit of the model. Ideally, these would be constructed in a way that the true choice probabilities would minimise expected risk. Then, since the data generating process is unknown, we approximate the expected risk with the empirical risk on the sample. For large enough samples, this would lead to good estimates for choice probabilities. We define risk in terms of incentive compatible scoring rules which are, in principle, reward schemes to incentivise truthful report of subjective probability judgements (Savage [1971], Gneiting and Raftery [2007]). Hence, when applied as a risk function, it ensures that true choice probabilities minimise expected risk. Furthermore, each scoring rule gives rise to a divergence function which allows us to quantify expected regret² in terms of divergence in predictions between the true and estimated choice probabilities. The criterion for consistent learning, in turn, relies on this divergence converging uniformly to zero. Hence, scoring rules serve as a natural candidate for defining risk and a key result in the paper establishes that they admit uniform learning with respect to their associated divergence functions (Proposition 1). Finally, we obtain sample complexity bounds by applying complexity measures such as Rademacher complexity and Pollard dimension. The details of the approach can be found in Section 2.

There are several reasons as to why the above problem would be of relevance to economists. We discuss here a few points, highlighting how the present work contributes to the empirical and theoretical literature. The estimation of choice probabilities in discrete choice models has for a long time been of interest to empiricists and econometricians interested in studying choice behaviour in markets (Manski [1975], McFadden [1978], Berry et al. [1995]). The novelty of the present approach lies in the introduction and assessment of sample complexity. This has two main advantages. Firstly, the sample size $N(\varepsilon, \delta)$ guarantees robust estimates which are independent of the random process that generates the data. Hence, the analyst/econometrician can safely rely on such a sample size to achieve a predetermined level of accuracy in estimates $\varepsilon, \delta > 0$. This is the key consequence of uniform consistency. Indeed, it also implies the usual consistency notion adopted in econometrics, where only convergence to the true parameter is required without a robustness guarantee in terms of sample sizes. Hence, in the usual definition of consistency, depending on the data generating process, more or less number of samples may be needed for accurate estimation and a priori, the analyst may not know how much data would be

²Regret is defined simply as the difference between the expected risk of the estimated choice probabilities and the expected risk of the true choice probabilities.

needed, for which the present approach would provide guarantees. Sample complexity yields uniform rates of convergence. Secondly, sample complexity also acts as a measure of complexity for stochastic choice models. Different models would have different sample complexity and the models with higher sample complexity, would need more samples for good estimation. Hence, in a sense, we can compare different models based on such measures (say Logit v/s Probit) and this allows the analyst or econometrician to make a formal judgement as to which stochastic choice models are simple and which ones are complex.

We now discuss how the present work pertains to the theoretical literature on stochastic choice. Typically, in theoretical work, the stochastic choice function is treated as the primitive and is itself interpreted as the data (Luce [1959], Luce [1977]). However, real data sets only involve finitely many data points and typically, each data point represents a choice made by an individual. This suggests a gap between the assumption and the nature of real data sets and the logic that seems to justify the gap is that there is perhaps a "law of large numbers" argument that one could rely on. However, from a formal standpoint, it is not clear whether this logic can always be implemented. For example, it may be that for some instances (such as observing a choice from a menu) there may be enough data points to evoke a law of large numbers argument but not for other instances. Hence, estimation of the entire map σ from finite data on choices seems far from obvious especially when keeping in mind the wide variety of decision procedures that the theoretical literature considers and which lead to very sophisticated stochastic choice models. Do all such models admit accurate estimation or learnability of choice probabilities?

This last question also brings us to a point of contrast between the theoretical and empirical literature. On the one hand, the theoretical models investigate varied contexts and decision procedures. If we assume these models as plausible in understanding choice behaviour, then perhaps a framework such as the present one could provide us a way to reason, in a simple manner, about the estimation and identification of all these models, allowing us to make good predictions about choice. However, on the other hand, empirical work still largely deals with more classical utility models and often for more complicated models, makes several distributional assumptions (on the data generating process) which are needed to obtain consistency results. In the present approach, we show that mild assumptions are needed for estimation. The main result of the paper is that in the context of stochastic preferences, if the preference class has finite VC dimension, then the corresponding stochastic choice model is learnable and allows us to recover

the underlying distribution over preferences in the population, uniformly over the data generating process. Under some conditions, with a modified definition of VC dimension with strict preferences, we also obtain a converse result which says that a stochastic preference model is non-learnable if the dimension of the underlying preference class is infinite. This allows us to conclude that for any learning rule and any sample size, the worst case expected distance between true and estimated choice probabilities (in $\Delta(A)$) is at least half times the maximum distance i.e $1/\sqrt{2}$. In situations involving choice between exactly two alternatives, this lower bound expected error is above the error level guaranteed by a rule which always makes the prediction that each alternative will be chosen with equal probability.

We provide several applications of our results for specific preference models commonly encountered in Industrial organisation, Decisions under Risk/Uncertainty and Spatial voting. We derive general bounds on rates of convergence based on Rademacher complexity. In particular, when the VC dimension is d (finite) and there are $|A|$ many alternatives to choose from, then the rademacher complexity associated with the stochastic choice model is at most $O(\sqrt{n^{-1}|A|\log(|A|)d\log(d)})$, which gives us an upper bound on the uniform rate of convergence of the estimator, which is defined via empirical risk minimisation with the quadratic scoring rule. We also show that for homothetic preferences specifically, convergence takes place in a stronger sense, the consistent estimates of the stochastic choice function converge in the sup norm, with a proportional rate of convergence. Lastly, we consider non-uniform learning and also address recoverability of cognitive heterogeneity in the setting of the bounded rationality model of choice with consideration sets.

Related Literature We discuss here, in some more detail, the relationship with prior work in the related literature. Typically, estimation of discrete choice models involves parametric or semi-parametric assumptions (McFadden [1984], Manski [1975], Manski [1985]). The most widely used method for estimation is maximum likelihood. In our setting, maximum likelihood arises as a special case when the scoring rule applied is the log rule. We also show that a variant of Manski’s maximum score estimation method (Manski [1975]) corresponds to use of an incentive compatible scoring rule. Non-parametric approaches have also been studied in the economics literature (Matzkin [1992], Briesch et al. [2010]). The present paper has a non-parametric approach as well. However, the estimation/learning problem considered here is different in that we consider estimation of choice probabilities and our approach involves techniques and ideas from statistical

learning theory at the core. In addition, as noted earlier, the consideration of sample complexity is one aspect that is novel in the present analysis. It is also interesting to compare the results of this paper with [McFadden and Train \[2000\]](#). The paper establishes that choice probabilities of a random utility model can be approximated by choice probabilities corresponding to a mixed logit specification. In contrast, in the present setting, mixed logit being a subset of linear preferences, would be uniformly learnable (with linear complexity), while all continuous preferences would be non-learnable (with lower bounded predictive error). Hence, the dimensions of preferences (VC and VC^+ in this paper) play a critical role and its connection to learnability exhibits a perhaps interesting difference between the mixed logit model and stochastic choice from more general continuous preferences. Relatedly, [De Blasi et al. \[2010\]](#) consider a Bayesian non-parametric approach to the problem of recovering preference heterogeneity in the mixed logit framework via estimation of choice probabilities.

Scoring rules have been applied in prior work to study certain binary classification problems ([Buja et al. \[2005\]](#),[Parry et al. \[2016\]](#)) in other settings with different objectives. In contrast, the problem considered here is of consistent uniform learning for stochastic functions in the sense of PAC (probably approximately correct) learning ([Valiant \[1984\]](#)). This paper is hence, most closely related to prior work in probabilistic concept learning ([Kearns and Schapire \[1994\]](#),[Abe et al. \[1991\]](#)) and learning stochastic rules ([Yamanishi \[1992\]](#),[Abe et al. \[2001\]](#)). While the nature of the learning problem is similar, the present paper is different in that we apply scoring rules to study the learning problem. For instance, one interesting feature is that the related papers mention learnability with respect to specific divergences, which has a very natural connection to scoring rules. We show in this paper how the present approach exploits this connection and this is indeed a key aspect of our estimation strategy. Our approach is also more general and subsumes some of the learning problems studied in previous work. We should also emphasise that we are primarily interested here in stochastic choice models in economics and are motivated by problems in demand estimation.

Lastly, the application of PAC learning in economics has been studied in different choice contexts in prior work ([Kalai \[2003\]](#),[Beigman and Vohra \[2006\]](#),[Zadimoghaddam and Roth \[2012\]](#),[Basu and Echenique \[2018\]](#),[Chase and Prasad \[2018\]](#)). However, the literature has only considered deterministic choice models. The distinguishing feature of the present work in relation to this strand of the literature is the consideration of stochastic choice models. We should also say that the present paper contributes to the growing lit-

erature on machine learning methods in economics ([Athey and Imbens \[2019\]](#)).

The outline of the paper is as follows. In section 2, we present the model and the learning problem. In section 3, we discuss consistent learning rules, which we construct on the basis of the principle of empirical risk minimisation. Finally, in section 4, we apply our approach to a variety of stochastic choice models in economics.

2 Model

Let $\mathcal{X} \subseteq \mathbb{R}^K$ be a compact set of *characteristics* and $A = \{a_1, a_2, \dots, a_m\}$ denote a finite set of *alternatives*. In certain contexts, $x = (x_I, x_A) \in \mathbb{R}^K$ shall be a vector including both individual characteristics (x_I) and product characteristics ($x_A = (x_a)_{a \in A}$). In other contexts, involving choice from menus, $\mathcal{X} \subseteq 2^A = \{0, 1\}^A$. Hence, at each $x \in \mathcal{X}$ which is a menu, a choice of an alternative is made, $a \in x$. We define $\mathcal{Z} = \mathcal{X} \times A$ and will denote a typical element of \mathcal{Z} as z .

A *stochastic choice function* is a map $\sigma : \mathcal{X} \rightarrow \Delta(A)$. The interpretation of σ is that at each x , $\sigma(x) = \{\sigma_a(x)\}_{a \in A}$ is a probability vector where $\sigma_a(x)$ denotes the probability that alternative $a \in A$ will be chosen when the underlying characteristics are given by x .

2.1 Data Generating Process

The analyst has access to a finite data set of choices from a population. Each data point consists of an individual's characteristic and the alternative chosen i.e $(x_i, a_i) \in \mathcal{Z}$. Hence, a *data set* is a finite sequence

$$z^n = \{(x_1, a_1), (x_2, a_2), \dots, (x_n, a_n)\}. \quad (1)$$

Hence, a data set is a element $z^n = (z_1, \dots, z_n)$ of \mathcal{Z}^n .

We now describe the data generating process. There is a probability measure $\pi \in \Delta(\mathcal{X})$ ³ which defines the distribution of characteristics in the population. Given π and a stochastic choice map σ , the data is generated as follows. Independent across i , x_i is drawn according to π and then a_i is drawn according to the choice probabilities given by $\sigma(x_i)$.

³Throughout the paper, for any metric space Y , we will denote as $\Delta(Y)$, the set of all Borel probability measures on Y . For any $\nu \in \Delta(Y)$, we shall denote as ν^n , the n -fold product measure on Y^n defined by $\nu^n := \underbrace{\nu \times \nu \times \dots \times \nu}_n$.

Note here that the analyst only observes the characteristic and the alternative chosen i.e. (x_i, a_i) through the data in 1. The analyst knows neither the distribution π nor the stochastic choice function σ , but assumes that it satisfies certain properties. We denote as $\pi \otimes \sigma$, the probability measure induced by π, σ together on the set $X \times A$. This represents the joint distribution from which the data is generated, by taking n i.i.d. samples from $\pi \otimes \sigma$. We shall denote as $\pi \otimes \sigma^n$, the n -fold product measure induced by $\pi \otimes \sigma$ on \mathcal{Z}^n . This is essentially the distribution of the data z^n .

2.2 Learning

A *model* is any family of stochastic choice functions Σ . The objective of the analyst is to learn the true choice probabilities based on the data and a model Σ represents the analyst's hypothesis. Formally, a *learning rule* is a map

$$\hat{\sigma} : \bigcup_{n \geq 1} (\mathcal{X} \times A)^n \rightarrow \Sigma. \quad (2)$$

Suppose now that the true choice probabilities are given by σ_0 and suppose π_0 governs the distribution over characteristics. We shall require that a learning map $\hat{\sigma}$ be so that with enough data, $z^n = \{(x_i, a_i)\}_{i=1}^n$, the estimate of the learning rule $\hat{\sigma}(z^n)$ would be close to the true choice probabilities σ_0 . Here, our notion of closeness between two choice probability maps σ, σ' will be given by

$$d_{\pi_0}(\sigma, \sigma') = \int_{\mathcal{X}} d(\sigma(x), \sigma'(x)) d\pi_0(x), \quad (3)$$

where $d : \Delta(A) \times \Delta(A) \rightarrow \mathbb{R}$ denotes a divergence function or metric on the space of all choice probability vectors on A i.e $\Delta(A)$. For example, d could be squared Euclidean distance, KL divergence or the total variation distance. This leads us to the following definition of consistency for learning rules.

Definition 1. A learning rule $\hat{\sigma}$ is consistent (with respect to d and Σ) if for all $0 < \varepsilon, \delta < 1$, there exists $N(\varepsilon, \delta)$ such that for all $n \geq N(\varepsilon, \delta)$,

$$(\forall \pi_0 \in \Delta(\mathcal{X}))(\forall \sigma_0 \in \Sigma) \left(\pi_0 \otimes \sigma_0^n \left(\{z^n : d_{\pi_0}(\hat{\sigma}(z^n), \sigma_0) < \varepsilon\} \right) \geq 1 - \delta \right). \quad (4)$$

We say that a model Σ is learnable with respect to d if there exists a learning rule $\hat{\sigma}$, which is consistent with respect to d and Σ . Finally, for a given $\varepsilon, \delta > 0$, we denote as $N(\varepsilon, \delta)$, the smallest n for which 4 holds. The function $N : (0, 1)^2 \rightarrow \mathbb{N}$ is called the sample complexity of

Σ (with respect to $\hat{\sigma}$).

The above definition of learnability is based on PAC learnability (see, for example, Valiant [1984], Yamanishi [1992]). In what follows, we shall discuss consistent learning rules for various stochastic choice models.

3 Consistent learning rules

3.1 Empirical Risk Minimization

We shall construct consistent learning rules based on the principle of *empirical risk minimization*. For a detailed treatment, see Vapnik [1998]. Much of this section contains standard ideas from machine learning. However, one should acknowledge Proposition 1, which provides a novel way to solve the problem of PAC learning stochastic functions w.r.t divergences via the use of scoring rules (see Yamanishi [1992], Kearns and Schapire [1994], Abe et al. [1991] Abe et al. [2001]).

For the learning problem, we first define a loss function $V : \Sigma \times \mathcal{X} \times A \rightarrow \mathbb{R}$. Suppose the true distribution and choice probabilities are given by π_0, σ_0 . Then, the *expected risk* corresponding to $\sigma \in \Sigma$ is defined as

$$\bar{V}(\sigma) = \int_{\mathcal{X} \times A} V(\sigma, x, a) d\pi_0 \otimes \sigma_0(x, a). \quad (5)$$

A minimizer of expected risk σ^* is defined as

$$\sigma^* \in \arg \min_{\sigma \in \Sigma} \bar{V}(\sigma). \quad (6)$$

We shall consider loss functions V for which it will hold that $\sigma^* = \sigma_0$ i.e the true choice probabilities would minimise the expected risk (a property also known as Fisher consistency). The principle of empirical risk minimization involves estimating the expected risk by the empirical risk on the sample $z^n = \{(x_i, a_i)\}_{i=1}^n$. For each $\sigma \in \Sigma$, the empirical risk is given by

$$\hat{V}(\sigma) = \frac{1}{n} \sum_{i=1}^n V(\sigma, x_i, a_i) \quad (7)$$

The learning rule that corresponds to empirical risk minimization, denoted by $\hat{\sigma}_E$, is defined as

$$\hat{\sigma}_E(z^n) \in \arg \min_{\sigma \in \Sigma} \hat{V}(\sigma) \quad (8)$$

The minimum in 8 need not always exist. However, if it holds for some M that $V(\sigma, x, a) \geq M$ for all $\sigma \in \Sigma$ and $(x, a) \in \mathcal{Z}$, then the infimum exists and we can define an *almost-ERM* learning rule as follows. We have $\{\varepsilon_n\}_n$ such that $\varepsilon_n > 0$ and $\lim_{n \rightarrow \infty} \varepsilon_n = 0$. The almost-ERM rule selects $\hat{\sigma}_E(z^n)$ such that

$$\hat{V}(\hat{\sigma}_E(z^n)) \leq \inf_{\sigma \in \Sigma_M} \hat{V}(\sigma) + \varepsilon_n \quad (9)$$

In this context, consistency relies on $\hat{\sigma}_E(z^n)$ being close to the true choice probability function σ^* as the function \hat{V} approximates \bar{V} , for large enough n . Consistency here is defined as follows, in terms of V .

Definition 2. A learning rule $\hat{\sigma}$ is said to be consistent (with respect to V and Σ) if for all $0 < \epsilon, \delta < 1$, there exist an $N(\epsilon, \delta)$ such that for all $n \geq N(\epsilon, \delta)$, it holds that

$$(\forall \pi_0 \in \Delta(\mathcal{X}))(\forall \sigma_0 \in \Sigma) \left(\pi_0 \otimes \sigma_0^n \left(\{z^n : \bar{V}(\hat{\sigma}(z^n)) < \inf_{\sigma \in \Sigma} \bar{V}(\sigma) + \epsilon\} \right) > 1 - \delta \right). \quad (10)$$

We say that a model Σ is learnable with respect to V if there exists a learning rule $\hat{\sigma}$, which is consistent with respect to V and Σ .

It turns out that a model Σ is learnable if the family of real-valued functions $V \circ \Sigma = \{V(\sigma, \dots) : \sigma \in \Sigma\}$ is a uniform Glivenko-Cantelli class of functions (see, for example, [Vapnik \[1998\]](#); [Shalev-Shwartz and Ben-David \[2014\]](#)). Each $f \in V \circ \Sigma$ is a function of the form $f : \mathcal{X} \times A \rightarrow \mathbb{R}$. We provide a definition below (see also [Dudley \[2014\]](#)).

Definition 3. A class of real valued functions \mathcal{F} on \mathcal{Z} is said to be a uniform Glivenko-Cantelli class of functions if for all $\epsilon, \delta > 0$, there exists $N(\epsilon, \delta)$ such that for all $n \geq N(\epsilon, \delta)$, for all $\nu \in \Delta(\mathcal{Z})$,

$$\nu^n \left(z^n : \sup_{f \in \mathcal{F}} |1/n \sum_{i=1}^n f(z_i) - \mathbb{E}_\nu(f)| < \epsilon \right) \geq 1 - \delta, \quad (11)$$

where $\mathbb{E}_\nu(f)$ denotes the expectation of f under ν .

A necessary and sufficient condition for a class of real-valued functions to be a uniform Glivenko-Cantelli class of functions is that it have finite V_γ -dimension for each

$\gamma > 0$ (see, for example, [Alon et al. \[1997\]](#)). Other combinatorial measures and notions of dimensions and capacity also guarantee the uniform Glivenko-Cantelli property for a class of functions for. eg. Vapnik's V -dimension, Pollard's P -dimension, P_γ -dimension, Covering numbers and Rademacher Complexity. We shall define and apply these notions and their study implications for sample complexity bounds in the subsequent sections.

3.1.1 Scoring Rules

In this section, we will define loss functions based on scoring rules. A *scoring rule* is a function $S : \Delta(A) \rightarrow \mathbb{R}^A$ (see, for example, [Savage \[1971\]](#), [Selten \[1998\]](#)). The interpretation of S is that it is a mechanism for eliciting subjective probability judgements. If p is a probabilistic prediction about the alternative to be chosen in A and suppose a is the alternative chosen, then $S_a(p)$ is the reward obtained. For each $p, q \in \Delta(A)$, we can define $S(p, q) := \sum_{a \in A} S_a(p)q_a$ to be the expected score when q is the true distribution over A . A scoring rule is said to be *incentive compatible* if

$$S(q, q) \geq S(p, q) \text{ for all } p, q \in \Delta(A) \quad (12)$$

i.e $p = q$ maximises the function $S(., q)$. We say that S is *strongly incentive compatible*⁴ if additionally, $p = q$ is the unique maximizer for each $q \in \Delta(A)$.

We now define a loss function based on S .

$$V^S(\sigma, x, a) := -S_a(\sigma(x)) \quad (13)$$

We can now prove the following lemma.

Lemma 1. *Let $\pi_0 \in \Delta(\mathcal{X})$ be the true distribution over characteristics and let $\sigma_0 \in \Sigma$ be the true choice probability function. Suppose S is incentive compatible. Then, σ_0 minimises $\bar{V}^S(\sigma)$. Furthermore, if S is strongly incentive compatible and $\sigma^* \in \Sigma$ minimizes $\bar{V}^S(.,)$, then $\sigma^*(x) = \sigma_0(x)$ with probability one according to π_0 .*

Proof. For any $\sigma' \in \Sigma$, we have

$$\begin{aligned} \bar{V}^S(\sigma') &= \int_{\mathcal{X} \times A} V^S(\sigma', x, a) d\pi_0 \otimes \sigma_0(x, a) \\ &= - \int_{\mathcal{X}} \left[\sum_{a \in A} S_a(\sigma'(x)) \sigma_{0,a}(x) \right] d\pi_0(x) \end{aligned}$$

⁴We use the terminology of [Selten \[1998\]](#). Incentive compatible (strongly) scoring rules are also referred to as proper (strictly) scoring rules. See, for example, [Gneiting and Raftery \[2007\]](#)

$$\begin{aligned}
&= - \int_{\mathcal{X}} S(\sigma'(x), \sigma_0(x)) d\pi_0(x) \\
&\geq - \int_{\mathcal{X}} S(\sigma_0(x), \sigma_0(x)) d\pi_0(x) \\
&= \bar{V}^S(\sigma_0),
\end{aligned} \tag{14}$$

where the inequality 14 follows from the fact that S is an incentive compatible scoring rule. Now, suppose σ^* minimizes $\bar{V}^S(\cdot)$ and let $E = \{x : \sigma^*(x) \neq \sigma_0(x)\}$. Since S is strongly incentive compatible, note that $S(\sigma_0(x), \sigma_0(x)) > S(\sigma^*(x), \sigma_0(x))$ for all $x \in E$. We already have that $S(\sigma_0(x), \sigma_0(x)) \geq S(\sigma^*(x), \sigma_0(x))$ for all $x \in X$. Hence, if $\pi_0(E) > 0$, then we will have that $\bar{V}^S(\sigma_0) < \bar{V}^S(\sigma^*)$. This contradicts that σ^* minimizes $\bar{V}^S(\cdot)$. \square

The above result shows that σ_0 is the minimiser of expected risk. Each incentive compatible scoring rule leads to a divergence function

$$d_S(p, q) = S(q, q) - S(p, q) \text{ for all } p, q \in \Delta(A). \tag{15}$$

Hence, from Lemma 1, it follows that

$$\begin{aligned}
\bar{V}_S(\sigma) - \inf_{\sigma \in \Sigma_{\mathcal{M}}} \bar{V}_S(\sigma) &= \bar{V}_S(\sigma) - \bar{V}_S(\sigma_0) \\
&= \int_{\mathcal{X}} d_S(\sigma(x), \sigma_0(x)) d\pi_0(x).
\end{aligned}$$

Hence, if we can construct consistent learning rules with respect to \bar{V}_S , then we can construct consistent learning rules with respect to the divergence function d_S .

The learning rule based on empirical risk minimisation corresponds to maximising the empirical score based on the data $D = \{(x_i, a_i)\}_{i=1}^n$. Hence, $\hat{\sigma}_E$ maximises

$$-\hat{V}^S(\sigma) = \frac{1}{n} \sum_{i=1}^n S_{a_i}(\sigma(x_i)), \tag{16}$$

which can be thought of as the empirical score. For a scoring rule S and stochastic choice function σ , we define the function $S \circ \sigma(x, a) := S_a(\sigma(x))$. Consistency of the learning rule $\hat{\sigma}_E$ with respect to d_S relies on the nature of the real-valued function class

$$S \circ \Sigma = \{S \circ \sigma | \sigma \in \Sigma\}.$$

The following is a key result.

Proposition 1. Let Σ be a model of stochastic choice and let S be an incentive compatible scoring rule. Suppose the class of functions $S \circ \Sigma$ is bounded above i.e. there exists an M such that $f(z) \leq M$ for all $f \in S \circ \Sigma$ and $z \in \mathcal{Z}$. If $S \circ \Sigma$ is a uniform Glivenko-Cantelli class, then the model Σ is learnable with respect to the divergence function d_S .

Proof. The proof is in the appendix. It follows from the fact that the ERM rule yields consistency (see Vapnik [1998]) in the sense of Definition 2 and from Lemma 1. \square

We now give some examples of incentive compatible scoring rules and their associated divergence functions.

1. (*Log Rule*): $S(p, a) = \ln(p_a)$. The log scoring rule is incentive compatible and its divergence function corresponds to KL divergence $S(p, q) = d_{KL}(p||q) = \sum_{a \in A} \frac{\ln(p_a)}{\ln(q_a)} p_a$. Note that maximisation of the empirical score with respect to the log rule corresponds to choosing choice probability functions according to the conditional maximum likelihood procedure .
2. (*Brier/Quadratic Scoring Rule*): $S(p, a) = 2p_a - \sum_b p_b^2$. The scoring rule is called the Brier or quadratic scoring rule and is strongly incentive compatible. The divergence function associated with S is the square of the euclidean distance between p and q i.e. $d_S(p, q) = \|p - q\|_2^2$.

Since all L^p metrics on \mathbb{R}^d are equivalent (as all norms on a finite dimensional vector space are equivalent), it follows that consistency of a learning rule with respect to d_S implies consistency with respect to any other L^p metric. Further, from Pinsker's inequality, we have that

$$\|p - q\|_{TV} \leq \sqrt{\frac{1}{2} d_{KL}(p||q)}.$$

Recall the total variational norm is equal to half times the L^1 distance i.e. $\|p - q\|_{TV} = \frac{1}{2} \|p - q\|_1$. Hence, the above inequality implies that empirical score maximisation using the log scoring rule, which corresponds to conditional maximum likelihood, leads to consistency with respect to all L^p metrics.

3. (*Manski's score with tie breaking*): Let $>$ be a complete strict order on A and let $p \in \Delta(A)$. Now, define another strict order $>_p$ on A as follows : $a >_p b$ if either $p_a > p_b$ or it is the case that $p_a = p_b$ and $a > b$. Finally, we let $W(1) \leq W(2) \leq \dots \leq W(|A| - 1)$.

The Manski scoring rule with tie breaking is defined as $S(p, a) = W(\{|c|a \succ_p c\})$. Note that $a \succ_p b$ implies both $p_a \geq p_b$ and $S(p, a) \geq S(p, b)$. Now, by the rearrangement inequality, this means that that $S(q, q) \geq S(p, q)$ for all $p, q \in \Delta(A)$ i.e. S is an incentive compatible scoring rule.

4. (*Subgradient of a convex function*) : Suppose we have a convex function $f : \Delta(A) \rightarrow \mathbb{R}$. Let $\nabla f(x) : \Delta(A) \rightarrow \mathbb{R}^A$ be a subgradient for the function f . We can now define a scoring rule based on f as follows.

$$S_a(p) = f(p) + \nabla f(p)(\delta_a - p).$$

S is incentive compatible and its associated divergence function corresponds to the Bregman divergence of f . Moreover, any incentive compatible scoring rule corresponds to some convex function. For example, $f(p) = \|p\|^2$ leads to the Brier score.

3.2 Dimension and Sample Complexity

In this section, we discuss sufficient conditions for a class of real valued functions to be uniform Glivenko-Cantelli. These place bounds on the capacity or complexity of the function class. There exist several alternative measures to quantify such complexity and we will discuss the central notions here. In particular, we discuss two notions we will use extensively : Rademacher complexity and Pollard Dimension.

Rademacher Complexity. Consider a real-valued function class \mathcal{F} defined on the set \mathcal{Z} . For the sample z^n , the empirical Rademacher complexity is defined as

$$R_{z^n}(\mathcal{F}) = \mathbb{E}_\eta \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \eta_i f(z_i) \right].$$

Here, each η_i is an independent bernoulli random variable which takes value 1 with 0.5 probability and -1 with 0.5 probability. Rademacher complexity of the function class \mathcal{F} (with respect to a distribution $\nu \in \Delta(\mathcal{Z})$) is defined as

$$R_n(\mathcal{F}) = \mathbb{E}_{\nu^n} [R_{z^n}(\mathcal{F})].$$

Now, it turns out that the Rademacher complexity of a function class is closely related to its Glivenko-Cantelli properties. Suppose \mathcal{F} is uniformly bounded i.e. there exists a real number $M > 0$ such that $|f(z)| \leq M$ for all $z \in \mathcal{Z}$ and $f \in \mathcal{F}$. Then, we have that (see, for

example, [Shalev-Shwartz and Ben-David \[2014\]](#)), for each $\nu \in \Delta(\mathcal{Z})$

$$\nu^n \left(z^n : \sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(z_i) - \mathbb{E}_\nu(f) \right| < 2R_{z^n}(\mathcal{F}) + 6M \sqrt{\frac{\log 4/\delta}{2n}} \right) \geq 1 - \delta,$$

Hence, if the empirical Rademacher complexity $R_{z^n}(\mathcal{F})$ converges to zero (say irrespective of the sample z^n), then we obtain the uniform Glivenko-Cantelli property for the function class \mathcal{F} . For example, if one shows that $R_{z^n}(\mathcal{F}) \leq C/\sqrt{n}$, where C is an absolute constant, then we get that the above inequality is satisfied for sample size n above

$$\frac{1}{\varepsilon^2} \left[2C + 6M \sqrt{\frac{\log 4/\delta}{2}} \right]^2.$$

This has implications for the sample complexity of stochastic choice models. Indeed, the proof of [Proposition 1](#) shows that for a stochastic choice model, the sample complexity $N(\varepsilon, \delta)$, is at most $\max\{3/\varepsilon, N'(\varepsilon/3, \delta)\}$, where $N'(\varepsilon/3, \delta)$ is the threshold sample size that guarantees [inequality 11](#) for values $(\varepsilon/3, \delta)$ in the Glivenko-Cantelli definition. We next discuss Pollard dimension, which is related to Rademacher complexity.

Pollard Dimension. We say that a real-valued function class \mathcal{F} *shatters* a set of points (z_1, \dots, z_n) if there exists a vector $(t_1, \dots, t_n) \in \mathbb{R}^n$ such that for all $B \subseteq \{1, \dots, n\}$, there exists $f_B \in \mathcal{F}$ such that

$$\begin{aligned} f_B(z_i) &> t_i \text{ for all } i \in B \\ f_B(z_i) &\leq t_i \text{ for all } i \notin B \end{aligned}$$

The Pollard dimension of \mathcal{F} , denoted as $P(\mathcal{F})$, is defined as

$$P(\mathcal{F}) = \max\{n \mid \exists (z_1, \dots, z_n) \text{ which can be shattered by } \mathcal{F}\}.$$

This complexity measure was introduced by [Pollard \[1984\]](#). It is a generalisation of the VC dimension (see [Vapnik and Chervonenkis \[1968\]](#), [Vapnik \[1998\]](#)), which is defined for sets. Indeed, for a class of 0-1 functions, the Pollard and VC dimensions coincide. Hence, for a collection of sets \mathcal{P} , we have that $VC(\mathcal{P}) = P(\mathcal{P})$. The more standard (equivalent) definition would involve the notion of shattering for sets. A class of sets \mathcal{P} shatters a set of points (z_1, \dots, z_n) if for each $B \subseteq \{1, \dots, n\}$, there exists $P \in \mathcal{P}$ such that $z_i \in P$ if and only

if $i \in B$. Then, the VC dimension of \mathcal{P} is defined as

$$VC(\mathcal{P}) = \max\{n \mid \exists (z_i)_{i=1}^n \text{ which can be shattered by } \mathcal{P}\}.$$

For the function class \mathcal{F} with Pollard dimension d , the empirical Rademacher complexity is

$$R_{z^n}(\mathcal{F}) \leq C \sqrt{\frac{d}{n}},$$

where $C > 0$ is an absolute constant (see [Pollard \[1984\]](#)). Hence, an immediate implication of this is that finite Pollard dimension implies the uniform Glivenko-Cantelli property. In view of [Proposition 1](#), the sample complexity of a stochastic choice model Σ , to be learned with respect to d_S , is upper bounded by a linear function of the Pollard dimension of $S \circ \Sigma$.

4 Stochastic Choice Models

In this section, we will consider a variety of stochastic choice models. In particular, we are interested in the recoverability of heterogeneity. By heterogeneity, we mean the distribution over preferences or parameters of the model, that leads to the stochasticity in choice. The main result of the paper is that when we consider stochastic preferences, we can recover preference heterogeneity if the support of the distribution has finite VC dimension. We discuss several familiar examples for which the result applies. We also establish a converse result on the non-learnability of such models. Lastly, we also consider choice with consideration sets and the recoverability of cognitive heterogeneity.

4.1 Stochastic Preferences

Let $X \subseteq \mathbb{R}^k$ be a set of alternatives. Suppose \mathcal{P} is a class of continuous preference relations (complete and transitive) on X . We are interested in defining stochastic choice functions where preferences are random in \mathcal{P} . Formally, we have a probability measure μ on \mathcal{P} i.e. $\mu \in \Delta(\mathcal{P})$ ⁵.

In this setting, the characteristic vector shall be a vector of alternatives. Hence, $\mathcal{X} \subseteq X^A$, as in earlier. We shall assume that for each $x = (x_a)_a \in \mathcal{X}$ we have $x_a \neq x_b$ for all $a \neq b$.

⁵The set $\Delta(\mathcal{P})$ denotes the set of all Borel probability measures on \mathcal{P} , which has the subspace topology generated by the closed convergence topology on the closed subsets of $X \times X$. See also the appendix.

Further,

$$\mu\left(\bigcup_{a \in A} \{\succ \in \mathcal{P} : x_a > x_b \text{ for all } b \neq a\}\right) = 1. \quad (17)$$

These assumptions ensure that in the random choice process, there are no ties with probability one. This is natural and simplifies the analysis. We now define the stochastic choice function corresponding to μ as

$$\sigma_a(x; \mu) = \mu(\{\succ \in \mathcal{P} : x_a > x_b \text{ for all } b \neq a\}),$$

which is the probability of choosing a given characteristics x and the distribution μ over preferences. The stochastic choice model corresponding to a random preference model on the class of preferences \mathcal{P} is defined as follows.

$$\Sigma^{\mathcal{P}} = \{\sigma(\cdot; \mu) \mid \mu \in \Delta(\mathcal{P}) \text{ satisfies (17) for all } x \in \mathcal{X}\}.$$

We show the following result.

Proposition 2. *Suppose \mathcal{P} is a class of continuous preferences with finite VC dimension d . Then, $\Sigma^{\mathcal{P}}$ is learnable with respect to any L^p metric. Moreover, for the Quadratic scoring rule S^{br} , the Rademacher complexity of $S^{br} \circ \Sigma^{\mathcal{P}}$ is at most $O(\sqrt{n^{-1}|A|\log(|A|d\log(d))})$.*

Proof. Firstly, we shall define a class of deterministic choice rules based on \mathcal{P} . Then, we will show that $\Sigma^{\mathcal{P}}$ is derived by taking continuous convex combinations of the deterministic class. We will then bound the Rademacher complexity corresponding to $S^{br} \circ \Sigma^{\mathcal{P}}$.

For each $\succ \in \mathcal{P}$ and $a \in A$, define the 0-1 valued function

$$\sigma_a(x; \succ) = \begin{cases} 1 & \text{if } x_a > x_b \text{ for all } b \neq a \\ 0 & \text{otherwise} \end{cases}$$

By assumption (17), we get that for all μ , and $a \in A$,

$$\sigma_a(x; \mu) = \int_{\mathcal{P}} \sigma_a(x, \succ) d\mu(\succ).$$

We will now show that for each a , the class of functions $S_a^{\mathcal{P}} = \{\sigma(\cdot; \succ) \mid \succ \in \mathcal{P}\}$ also has finite VC dimension.

Consider a finite set of n points $D = \{x^1, \dots, x^n\} \subseteq \mathcal{X}$. Suppose that D can be shattered

by $S_a^{\mathcal{P}}$. Then, all 2^n possible 0-1 valued labellings of D can be generated by \mathcal{P} . Now, consider the following set of points in $X \times X$.

$$D' = \{(x_a^i, x_b^i)\}_{i \in \{1, \dots, n\}, b \in A \setminus a}.$$

Note that D' is a set of $n(|A|-1)$ points. Let d be the VC dimension of \mathcal{P} . On the one hand, from Sauer's Lemma (see [Kearns and Vazirani \[1994\]](#)), we have that at most $(\frac{en(|A|-1)}{d})^d$ many labellings in D' can be generated by \mathcal{P} . On the other hand, we have that since $S_a^{\mathcal{P}}$ shatters D , at least 2^n labellings are obtained by \mathcal{P} in D' . For large n , we obtain a contradiction as

$$\left(\frac{en(|A|-1)}{d}\right)^d < 2^n. \quad (18)$$

Hence, we establish that $S_a^{\mathcal{P}}$ has finite VC dimension, say d_a . Now, from a well known result (see [Pollard \[1984\]](#)), it follows that there is an absolute constant C_a such that the Rademacher complexity of $S_a^{\mathcal{P}}$, say $\mathcal{R}_{x^n}(S_a^{\mathcal{P}})$, is at most $C_a \sqrt{d_a/n}$ for all x^n . Now, consider the probability of a in the stochastic choice model i.e. the function class $\Sigma_a^{\mathcal{P}} = \{\sigma_a(\cdot; \mu) \mid \mu \text{ satisfies (17)}\}$. Since, $\Sigma_a^{\mathcal{P}}$ is essentially derived by taking convex combinations (including continuous combinations) from $S_a^{\mathcal{P}}$ (see Lemma 4 in the appendix), we obtain that the Rademacher complexity of $\Sigma_a^{\mathcal{P}}$ is at most equal to that of $S_a^{\mathcal{P}}$. Hence, we have $\mathcal{R}_{x^n}(\Sigma_a^{\mathcal{P}}) \leq C_a \sqrt{d_a/n}$.

We now complete the proof by applying the Quadratic scoring rule and bounding the Rademacher complexity of the class $\mathcal{F} = S^{br} \circ \Sigma^{\mathcal{P}}$. Firstly, consider the function class $\mathcal{F}_a = \{f(\cdot, a) \mid f \in S^{br} \circ \Sigma^{\mathcal{P}}\}$. Note that

$$f(x, a) = 2\sigma_a(x; \mu) - \sum_{b \in A} \sigma_b(x; \mu)^2$$

for some μ . Since Rademacher complexity is additive (applied to $\{\Sigma_a^{\mathcal{P}}\}_a$) and has nice Lipschitz composition properties (in this case y^2), we get that $\mathcal{R}_{z^n}(\mathcal{F}_a) \leq 2C_a \sqrt{d_a/n} + \sum_b 2C_b \sqrt{d_b/n}$ (see [Bartlett and Mendelson \[2002\]](#) and also [Shalev-Shwartz and Ben-David \[2014\]](#)). Define $C' = \max_{a \in A} 2C_a \sqrt{d_a} + \sum_b 2C_b \sqrt{d_b}$. Now, for a given $z^n = (x_i, a_i)_{i=1}^n$, define $x^{n_a} = (x_i)_{i \in \{i \mid a_i = a\}}$ and $n_a = |\{i \mid a_i = a\}|$. Using Lemma 5 in the appendix, we get

$$\mathcal{R}_{z^n}(\mathcal{F}) \leq \sum_{a \in A} \frac{n_a}{n} \mathcal{R}_{x^{n_a}}(\mathcal{F}_a)$$

$$\begin{aligned}
&\leq \sum_{a \in A} \frac{n_a}{n} (C'/\sqrt{n_a}) \\
&\leq \sum_{a \in A} \sqrt{\frac{n_a}{n}} (C'/\sqrt{n}) \\
&\leq (C'/\sqrt{n}) \sum_{a \in A} \sqrt{\frac{n_a}{n}} \\
&\leq (C'/\sqrt{n}) \sqrt{|A|}.
\end{aligned} \tag{19}$$

The last inequality follows from the fact that \sqrt{y} is concave. Hence, we establish that $\Sigma^{\mathcal{P}}$ is learnable with respect to any L_p metric.

Finally to obtain the sample complexity bound, consider the inequality 18. Based on the inequality, we can in fact say that the following holds.

$$d_a \leq 4 \frac{d}{\log 2} \log\left(\frac{2d}{\log 2}\right) + 2d \frac{\log e(|A| - 1)}{\log 2}. \tag{20}$$

The above follows from a fact of arithmetic that $x \geq 4a \log(2a) + 2b$ implies $x \geq a \log(x) + b$ (see, for example, [Shalev-Shwartz and Ben-David \[2014\]](#)) whenever $a > 1$ and $b > 0$. Hence, the derived inequality in 20, together with the conclusion in 19 implies that the Rademacher complexity of \mathcal{F} is at most $O(\sqrt{n^{-1}|A|\log(|A|)d \log(d)})$. \square

The above result also allows us to recover the distribution over preferences. For this, we will define the following notion of distance between distributions.

$$d_{\Delta}(\mu, \mu') := \int_{\mathcal{X}} d(\sigma(x; \mu), \sigma(x; \mu')) d\pi_0(x).$$

Then, the following result holds.

Corollary 1. *(Recovering heterogeneity) Suppose \mathcal{P} is a class of continuous preferences with finite VC dimension. Then, there exists an estimator $\hat{\mu}$ (based on a learning rule $\hat{\sigma}$ derived from almost-ERM with the Quadratic scoring rule) of the true distribution μ_0 over preferences such that*

$$d_{\Delta}(\hat{\mu}, \mu_0) \rightarrow_p 0,$$

uniformly over π_0 and μ_0 . Moreover, the rate of convergence is given by the sample complexity bound derived from Proposition 2.

For homothetic preferences, one can say more. Homotheticity of a preference relation \succeq requires that for $x_a, x_b \in X$ and $\alpha > 0$, one has $x_a \succeq x_b$ if and only if $\alpha x_a \succeq \alpha x_b$. This means that if we knew the stochastic choice function σ for such preferences locally, then we can scale alternatives appropriately so that we could elicit the function σ globally. Indeed, for homothetic preferences, the equality $\sigma(x) = \sigma(\alpha x)$ holds. This property implies that we can get consistency in terms of convergence in the sup norm (over \mathcal{X}). We state the result formally below.

Proposition 3. *Assume \mathcal{X} is compact. Also, assume that for each $x \in \mathcal{X}$ and $\alpha \in (0, 1)$, $\alpha x \in \mathcal{X}$. For continuous homothetic preferences, the estimator $\hat{\sigma}$ for σ_0 converges in the sup norm i.e.*

$$\sup_{x \in \mathcal{X}} \|\hat{\sigma}(x) - \sigma_0(x)\| \rightarrow_p 0$$

Moreover, the rate of convergence is given by the sample complexity bound derived in Proposition 2.

Proof. We show that if for $\varepsilon \in (0, 1)$, we have $d_{\pi_0}(\hat{\sigma}, \sigma_0) = \mathbb{E}_{\pi_0} \left[\|\hat{\sigma}(x) - \sigma_0(x)\|^2 \right] < \varepsilon^2/4$, then we have that $\sup_{x \in \mathcal{X}} \|\hat{\sigma}(x) - \sigma_0(x)\| < \varepsilon$. To show the latter, note that it suffices to show that there exists an open ball containing 0 in $\mathbb{R}^{k \times |A|}$, say \mathcal{O} , such that $\sup_{x \in \mathcal{O} \cap \mathcal{X}} \|\hat{\sigma}(x) - \sigma_0(x)\| < \varepsilon$. This is because for any $x \in \mathcal{X}$, there exists a scalar $\alpha > 0$ such that $\alpha x \in \mathcal{O} \cap \mathcal{X}$ and from homotheticity, we have that $\hat{\sigma}(x) = \hat{\sigma}(\alpha x)$ and $\sigma_0(x) = \sigma_0(\alpha x)$. Firstly, note that if $d_{\pi_0}(\hat{\sigma}, \sigma_0) < \varepsilon^2/4$, then there exists $y \in \mathcal{X}$ such that $\|\hat{\sigma}(y) - \sigma_0(y)\| < \varepsilon/2$. From Lemma 2 in the Appendix, it follows that any stochastic choice function σ derived from continuous preferences is continuous in x and hence $h(x) = \|\hat{\sigma}(x) - \sigma_0(x)\|$ is continuous in x . As \mathcal{X} is compact, it follows that $h(x)$ is uniformly continuous. Hence, there exists a $\delta > 0$ such that whenever $\|x - x'\| < \delta$, we have $|h(x) - h(x')| < \varepsilon/2$.

Now, let $\alpha > 0$ be small enough such that $\|\alpha y\| < \delta/2$. By homotheticity it follows that $\hat{\sigma}(y) = \hat{\sigma}(\alpha y)$ and $\sigma_0(y) = \sigma_0(\alpha y)$. Hence, this means that $h(\alpha y) < \varepsilon/2$. Now, let us consider the $\delta/2$ -open ball around αy , say \mathcal{O} in $\mathbb{R}^{k \times |A|}$. Note that as $\|\alpha y\| < \delta/2$, we have $0 \in \mathcal{O}$. Hence, it suffices to show that $\sup_{x \in \mathcal{O} \cap \mathcal{X}} \|\hat{\sigma}(x) - \sigma_0(x)\| < \varepsilon$. Hence, let $z \in \mathcal{O} \cap \mathcal{X}$. This means that $\|z - \alpha y\| < \delta/2 < \delta$. Hence, $|h(z) - h(\alpha y)| < \varepsilon/2$. Finally, the latter, together with the fact that $h(\alpha y) < \varepsilon/2$, implies that $h(z) < \varepsilon$. This proves the result. \square

In this setting, the class of preferences \mathcal{P} would depend on the context. Stochastic preferences supported on linear utility would, for example, lead to the random coefficients model. Several canonical utility based models such as cobb-douglas or CES preferences may be considered. If preferences are over risky prospects, then we may

have stochastic preferences over expected utility models and models exhibiting ambiguity aversion. We would bound the VC dimension of the preference class (see also [Basu and Echenique \[2018\]](#)) and then apply the above result to obtain learnability of the corresponding stochastic choice model supported on the preferences.

Applications We discuss some applications of the above result below.

1. (*Linear Preferences*) Consider the model of preferences \mathcal{P}_{LIN} defined as follows. For each $\succsim \in \mathcal{P}_{LIN}$ there exists a vector $v \in \mathbb{R}^k$ such that

$$x \succsim y \text{ if and only if } v \cdot x \geq v \cdot y.$$

Such preferences satisfy the well-studied independence axiom from decision theory. This means that for all $x, y, z \in X$ and $\lambda \in [0, 1]$,

$$x \succsim y \text{ if and only if } \lambda x + (1 - \lambda)z \succsim \lambda y + (1 - \lambda)z.$$

From a result in [Basu and Echenique \[2018\]](#), it follows that the VC dimension of the set of preferences that satisfy the independence axiom is at most $k + 1$. Hence, it follows that $VC(\mathcal{P}_{LIN}) \leq k + 1$. From [Proposition 2](#), it follows that the corresponding stochastic choice model, in which the vector of coefficients is random, would be learnable.

Consider the following widely used econometric specification. The utility from an alternative a is

$$u(x; \theta, \epsilon) = \theta \cdot x_a + \epsilon_a,$$

where $\theta \in \mathbb{R}^k$ and $\epsilon = (\epsilon_a)_{a \in A} \in \mathbb{R}^A$ are vectors that are randomly drawn according to the probability measure μ . For example, for the Logit model, θ is deterministically fixed and each ϵ_a independently follows a logistic distribution (similarly for Probit). For Mixed Logit, we would have that θ would follow an unknown distribution for eg. multivariate normal. Here we have $v = (\theta, \epsilon)$ and one may rewrite alternatives as $x'_a = (x_a, e_a)$, where e_a is the unit vector in \mathbb{R}^A where the value corresponding to a is one. This also ensures that $x'_a \neq x'_b$ for all $a \neq b$. Note that our formulation here essentially pertains to the case where x is independent of ϵ .

2. (*Random Utility*) Suppose now that preferences are defined by an underlying utility

function. Hence, suppose we have a class of utility functions $\mathcal{U} \subseteq \{u \mid u : X \rightarrow \mathbb{R}\}$. For an given utility function $u \in \mathcal{U}$, the corresponding preference relation \succeq_u is defined as

$$x \succeq_u y \text{ if and only if } u(x) \geq u(y).$$

The model of preferences generated by \mathcal{U} is then defined as $\mathcal{P}_{\mathcal{U}} := \{\succeq_u \mid u \in \mathcal{U}\}$. It turns out that if the class of utility functions has finite Pollard dimension, say d , then the VC dimension of $\mathcal{P}_{\mathcal{U}}$ is also finite, and is of the order $O(d)$. It can be shown that utility classes corresponding to canonical models such as Cobb-Douglas and CES preferences have finite Pollard dimension.

3. (*Decisions under Risk/Uncertainty*) Suppose that the set of alternatives X corresponds to a set of monetary acts over a finite state space (uncertainty) i.e $X \subseteq \mathbb{R}^{\Omega}$ or the lotteries over a finite set of consequences Z i.e. $X \subseteq \Delta(Z)$ (risk). Such contexts have been studied extensively in decision theory (see, for example, [Kreps \[1988\]](#) or [Gilboa \[2009\]](#)). Several studies have investigated different models of preferences based on different notions of expected utility. As was shown in [Basu and Echenique \[2018\]](#), the standard expected utility and choquet expected utility models had finite VC dimension, where the latter has VC dimension exponential in $|\Omega|$. In the context of risk, one may show that expected utility and Quadratic expected utility (see [Chew et al. \[1991\]](#)) preferences have finite VC dimension. Indeed, the study of [Gul and Pesendorfer \[2006\]](#) consider random preferences with expected utility. The related model of random choice with private information due to [Lu \[2016\]](#) can be studied within our framework of linear preferences as well.

Finally, one should note that the nature of axioms of the model and VC dimension bounds are closely related. As was noted earlier, in [Basu and Echenique \[2018\]](#), we show the preferences satisfying the independence axiom have finite VC dimension, linear in $|\Omega|$. In that paper, it was also argued that if preferences satisfy comonotonic independence and admit existence of certainty equivalents, then the VC dimension would be finite and upper bounded by a function exponential in $|\Omega|$.

4. (*Euclidean Preferences*) Consider the context of spatial voting (see [Downs et al. \[1957\]](#)) where $X \subseteq \mathbb{R}^k$ is an ideological space and each candidate in an election is represented by a point in X (alternative). In this voting setting, preferences are defined through an ideal point x^* such that candidates closer to the ideal point are preferred

by the voter. Hence, we have that

$$x \succ_{x^*} y \text{ if and only if } \|x - x^*\| \leq \|y - x^*\|.$$

Hence $\mathcal{P}_{EUC} = \{\succ_{x^*} \mid x^* \in X\}$. In this case, the probability measure μ essentially gives us a distribution over ideal points, which is also referred to as voter heterogeneity. Since euclidean preferences can be defined by finitely many arithmetic computations (see Proposition 8 in the appendix), it follows that the VC dimension is finite and of the order $O(k)$. Our estimation results provide us a way to recover voter heterogeneity in these settings.

4.2 Non-learnability

We now establish a result on the non-learnability of $\Sigma^{\mathcal{P}}$. With some additional conditions, one can show that a converse result obtains. This requires defining a notion of shattering based on strict preferences. We say that a model of preferences \mathcal{P} *strictly shatters* a set of pairs of alternatives $((x_i, y_i))_{i=1}^n$ if for any labelling $(a_i)_{i=1}^n \in \{0, 1\}^n$ of the pairs, there exists a preference relation $\succ \in \mathcal{P}$ such that

$$\begin{aligned} x_i > y_i & \text{ whenever } a_i = 1 \text{ and;} \\ y_i > x_i & \text{ whenever } a_i = 0 \end{aligned}$$

We now consider the following modified definition of VC dimension, which we will call $VC^+(\mathcal{P})$. It is defined as follows.

$$VC^+(\mathcal{P}) = \max\{n \mid \exists ((x_i, y_i))_{i=1}^n \text{ which can be strictly shattered by } \mathcal{P}\}.$$

We will assume here for simplicity that $X = \mathbb{R}^k$ and that $\mathcal{X} = \{x \in X^A \mid x_a \neq x_b \text{ for all } a \neq b\}$ ⁶. A preference relation \succ is said to be monotonic if $x \gg y$ implies that $x > y$. We now state and prove the result.

Proposition 4. *Let \mathcal{P} be a class of monotone continuous preferences such that $VC^+(\mathcal{P}) = +\infty$. Furthermore, suppose that there exists a probability measure satisfying 17 that has full support in \mathcal{P} . Then, the following hold.*

1. *The stochastic choice model $\Sigma^{\mathcal{P}}$ is not learnable with respect to any L^p metric.*

⁶This assumption is not crucial for the result.

2. Suppose we fix the euclidean distance on $\Delta(A)$. Then, for any learning rule $\hat{\sigma}$ and sample size n , the worst case expected distance (error) between true and estimated choice probabilities is at least $1/\sqrt{2}$, which is half of the maximum possible distance. For $|A| = 2$, this lower bound for error is above what can be achieved by predicting that each alternative will be chosen with probability $1/2$.

Proof. The proof applies the probabilistic method where we randomly choose a stochastic choice function from $\Sigma^{\mathcal{P}}$ and show that expected error is bounded away from zero. Then, we argue that there exists some function in $\Sigma^{\mathcal{P}}$ that generates error above a threshold level.

Fix any learning rule $\hat{\sigma}$ and suppose that $0 < \varepsilon < \sqrt{2} - 1$ and $0 < \delta < 1 - 1/\sqrt{2}$. Also, let $n \in \mathbb{N}$ be a sample size. Now, further fix $k \in \mathbb{N}$ such that $k \geq 2$. Since $VC^+(\mathcal{P}) = +\infty$, there exists a strictly shattered set $((x_i, y_i))_{i=1}^{kn}$. Finally, let $t \in (0, 1)$ such that $t^n \geq 0.5$. We will show then, that there exists a set of points $(x^i)_{i=1}^{kn} \subseteq \mathcal{X}$ such that for all $B \subseteq \{1, 2, \dots, kn\}$, there exists $\sigma^B \in \Sigma^{\mathcal{P}}$ such that

$$\begin{aligned} \sigma_a^B(x^i) &\geq t \text{ for all } i \in B \text{ and ;} \\ \sigma_a^B(x^i) &\leq 1 - t \text{ for all } i \notin B, \end{aligned} \tag{21}$$

where $a \in A$ is fixed beforehand. We now show how to construct the set of points $(x^i)_{i=1}^{kn} \subseteq \mathcal{X}$. Let $b \in A \setminus \{a\}$. We first define $x_a^i := x_i$ and $x_b^i := y_i$. Further, for each i we let $\{x_c^i\}_{c \in A \setminus \{a, b\}} \subseteq \mathbb{R}^k$ be such that all the vectors in $\{x_c^i\}_{c \in A}$ are distinct and also that $x_a^i \gg x_c^i$ and $x_b^i \gg x_c^i$ for all $c \in A \setminus \{a, b\}$. Due to monotonicity, this ensures that from the set of alternatives in x^i either x_a^i or x_b^i is maximal.

Since \mathcal{P} strictly shatters $((x_a^i, x_b^i))_{i=1}^{kn}$, for each $B \subseteq \{1, 2, \dots, kn\}$, there exists a preference relation \succsim such that

$$\begin{aligned} x_a^i &> x_b^i \text{ whenever } i \in B \text{ and ;} \\ x_b^i &> x_a^i \text{ whenever } i \notin B. \end{aligned}$$

Now, consider the set of preferences $\mathcal{Q} = \{\succsim \mid x_a^i > x_b^i \text{ for all } i \in B \text{ and } x_b^i > x_a^i \text{ for all } i \notin B\}$. It is open in \mathcal{P} under the closed convergence topology. Hence, it is measurable in the corresponding Borel sigma algebra. Now, consider a probability measure $\mu \in \Delta(\mathcal{P})$ that satisfies the regularity condition in 17 and has full support in \mathcal{P} . Then, since \mathcal{Q} is open, we must have that $\mu(\mathcal{Q}) > 0$. We will now construct another probability measure μ_B satisfying

17 such that its corresponding stochastic choice function σ^B satisfies 21. Firstly, let $\beta > 0$ be such that $\beta\mu(\mathcal{Q}^c) \leq 1 - t$. Then, let $\alpha := \frac{1 - \beta\mu(\mathcal{Q}^c)}{\mu(\mathcal{Q})}$. This implies that $\alpha\mu(\mathcal{Q}) + \beta\mu(\mathcal{Q}^c) = 1$ and $\alpha\mu(\mathcal{Q}) \geq t$. Now, define the probability measure $\mu_B(\mathcal{E}) := \alpha\mu(\mathcal{Q} \cap \mathcal{E}) + \beta\mu(\mathcal{Q}^c \cap \mathcal{E})$. Note that $\mu_B(\mathcal{E}) = 0$ if and only if $\mu(\mathcal{E}) = 0$, which implies that μ_B also satisfies 17. By definition, it now follows that the corresponding stochastic choice function σ^B satisfies 21.

We now proceed with the proof of non-learnability. Since all L^p metrics are equivalent on $\Delta(A)$, it suffices to show the claim for euclidean distance. Let π be the uniform distribution over the set of points $\{x^i\}_{i=1}^{kn}$. Also, let ζ be the uniform distribution over all stochastic choice functions in $\{\sigma^B\}_{B \subseteq \{1, 2, \dots, kn\}}$. We will lower bound the expected error after n samples.

$$\mathbb{E}_{\zeta} \mathbb{E}_{\pi \otimes \sigma^n} \mathbb{E}_{\pi} \left[\|\hat{\sigma}(z^n)(x) - \sigma(x)\| \right].$$

Now, from Fubini's theorem, we will change the order of expectations. The above expectation then becomes

$$\begin{aligned} \mathbb{E}_{\pi^n} \mathbb{E}_{\pi} \mathbb{E}_{\zeta} \left[\|\hat{\sigma}(z^n)(x) - \sigma(x)\| \right] &= \mathbb{E}_{\pi^{n+1}} \left[\mathbb{E}_{\zeta} \|\hat{\sigma}(z^n)(x) - \sigma(x)\| \mid x \in x^n \right] \mathbb{P}_{\pi^{n+1}}(x \in x^n) \\ &+ \mathbb{E}_{\pi^{n+1}} \left[\mathbb{E}_{\zeta} \|\hat{\sigma}(z^n)(x) - \sigma(x)\| \mid x \notin x^n \right] \mathbb{P}_{\pi^{n+1}}(x \notin x^n). \end{aligned}$$

Since π is uniform, it follows that $\mathbb{P}_{\pi^{n+1}}(x \notin x^n) \geq \frac{k-1}{k}$. We will now simply lower bound the conditional expectation $\mathbb{E}_{\pi^{n+1}} \left[\mathbb{E}_{\zeta} \|\hat{\sigma}(z^n)(x) - \sigma(x)\| \mid x \notin x^n \right]$. Consider some fixed $x \notin x^n$. Now, let us write the expectation $\mathbb{E}_{\zeta} \left[\|\hat{\sigma}(z^n)(x) - \sigma(x)\| \right]$ more explicitly as

$$\mathbb{E}_{\sigma \sim \zeta} \mathbb{E}_{c^n \sim \sigma(x^n)} \left[\|\hat{\sigma}(x^n, c^n)(x) - \sigma(x)\| \right].$$

where $\sigma(x^n)$ is the joint distribution induced on $\{a, b\}^n$ by the product of the choice probabilities in $(\sigma(y))_{y \in x^n}$. Consider some fixed σ^B for some $B \subseteq \{1, 2, \dots, 2n\}$. Now, let $d(B, x^n) \in \{a, b\}^n$ be such that for each $j \in \{1, \dots, n\}$ we define $d_j(B, x^n) = a$ if and only if $x_j \in (x_1, \dots, x_n) = x^n$ is such that $x_j = x^i$ for some $i \in B$. Now, recall that for $i \in B$, we have $\sigma_a(x^i) \geq t$ and for $i \in \{1, 2, \dots, 2n\} \setminus B$, we have $\sigma_a^B(x^i) \leq 1 - t$. Hence, by definition, it follows that $\sigma^B(x^n)$ draws the vector $d(B, x^n)$ with probability at least t^n . This gives us

$$\mathbb{E}_{c^n \sim \sigma_B(x^n)} \left[\|\hat{\sigma}(x^n, c^n)(x) - \sigma_B(x)\| \right] \geq t^n \left[\|\hat{\sigma}(x^n, d^n(B, x^n))(x) - \sigma^B(x)\| \right],$$

which implies

$$\mathbb{E}_{\sigma \sim \zeta} \mathbb{E}_{c^n \sim \sigma(x^n)} \left[\|\hat{\sigma}(x^n, c^n)(x) - \sigma(x)\| \right] \geq t^n \mathbb{E}_{\sigma_B \sim \zeta} \left[\|\hat{\sigma}(x^n, d^n(B, x^n))(x) - \sigma^B(x)\| \right].$$

Now, note in the latter expectation, since ζ is uniform and $x \notin x^n$, it follows that the random variable $\hat{\sigma}(x^n, d^n(B, x^n))(x)$ is independent of the event " $\sigma_a^B(x) \geq t$ ", with the latter event having probability 0.5. This is because $d^n(B, x^n)$ only depends on the set of indexes $B(x^n) = \{i \mid x^i \in x^n\}$. Hence, conditional on $\hat{\sigma}(x^n, d^n(B, x^n))(x)$, the expectation

$$\mathbb{E}_{\sigma_B \sim \zeta} \left[\|\hat{\sigma}(x^n, d^n(B, x^n))(x) - \sigma^B(x)\| \mid \hat{\sigma}(x^n, d^n(B, x^n))(x) \right]$$

is such that with probability 0.5 we have $\|\hat{\sigma}(x^n, d^n(B, x^n))(x) - \sigma^B(x)\|$ with $\sigma_a^B(x) \geq t$ and with probability 0.5 we have $\|\hat{\sigma}(x^n, d^n(B, x^n))(x) - \sigma^B(x)\|$ such that $\sigma_a^B(x) \leq 1 - t$. This implies that value of the above conditional expectation is at least $0.5\|(t, 1 - t) - (1 - t, t)\| = 0.5\sqrt{2}|1 - 2t|$. Combining these, we get the lower bound

$$\mathbb{E}_{\pi^{n+1}} \left[\mathbb{E}_{\zeta} \|\hat{\sigma}(z^n)(x) - \sigma(x)\| \mid x \notin x^n \right] \geq 0.5t^n \sqrt{2}|1 - 2t|.$$

Finally, since, $\mathbb{P}_{\pi^{n+1}}(x \notin x^n) \geq \frac{k-1}{k}$, we get that

$$\mathbb{E}_{\zeta} \mathbb{E}_{\pi \otimes \sigma^n} \mathbb{E}_{\pi} \left[\|\hat{\sigma}(z^n)(x) - \sigma(x)\| \right] \geq \frac{k-1}{k} 0.5t^n \sqrt{2}|1 - 2t|.$$

This means that there exists $\sigma \in \{\sigma_B\}_{B \subseteq \{1, 2, \dots, 2n\}}$ such that

$$\mathbb{E}_{\pi \otimes \sigma^n} \mathbb{E}_{\pi} \left[\|\hat{\sigma}(z^n)(x) - \sigma(x)\| \right] \geq \frac{k-1}{k} 0.5t^n \sqrt{2}|1 - 2t|. \quad (22)$$

Taking $k \rightarrow +\infty$ and $t \rightarrow 1$, we get that the RHS converges to $0.5\sqrt{2} = 1/\sqrt{2}$. Since $\frac{1}{\sqrt{2}}\|\hat{\sigma}(z^n)(x) - \sigma(x)\| \in [0, 1]$, by an application of Markov's inequality⁷, it follows that for any $0 < \varepsilon < \sqrt{2}(1 - 1/\sqrt{2}) = \sqrt{2} - 1$ and $0 < \delta < 1 - 1/\sqrt{2}$, we can choose $k \rightarrow +\infty$ and $t \rightarrow 1$ so that

$$\mathbb{P}_{\pi \otimes \sigma^n} \left(\mathbb{E}_{\pi} \left[\|\hat{\sigma}(z^n)(x) - \sigma(x)\| \right] > \varepsilon \right) \geq \delta.$$

Hence, the result obtains. From the lower bound derived in the RHS of 22, we can say

⁷From Markov's inequality, one can show that for a random variable $X \in [0, 1]$ and $\varepsilon \in (0, 1)$, we have that if $\mathbb{E}[X] \geq 1 - \varepsilon^2$, then $\mathbb{P}[X > 1 - \varepsilon] \geq 1 - \varepsilon$.

that

$$\inf_{\hat{\sigma}} \sup_{\pi_0 \in \Delta(\mathcal{X}), \sigma_0 \in \Sigma^{\mathcal{P}}} \mathbb{E}_{\pi_0 \otimes \sigma_0^n} \mathbb{E}_{\pi_0} \left[\|\hat{\sigma}(z^n)(x) - \sigma_0(x)\| \right] \geq 1/\sqrt{2},$$

where the infimum is taken over all possible learning rules $\hat{\sigma}$. Finally, note that when $|A| = 2$, a rule which always predicts probabilities $\sigma_a(x) = 1/2$ on each alternative a , achieves expected error at most $1/\sqrt{2}$. \square

One can argue that the class of all monotone, continuous and convex preferences \mathcal{P} has $VC^+(\mathcal{P}) = +\infty$. In the context of uncertainty from a result in [Basu and Echenique \[2018\]](#), it follows that Max-min preferences (MEU) have infinite VC dimension whenever $|\Omega| \geq 3$ ⁸. However, the arguments can be modified to show that $VC^+(\mathcal{P}_{MEU}) = +\infty$ as well. Of course, non-monotonic preferences can be non-learnable as well. Consider again the following econometric specification : an alternative has utility $u(x_a, \theta) + \epsilon_a$, where each ϵ_a independently follows a normal distribution with (μ_a, σ_a^2) as mean and variance. Suppose that for the utility class $\mathcal{U} = \{u(\cdot; \theta)\}_{\theta \in \Theta}$, the corresponding preference class has $VC^+(\mathcal{P}_{\mathcal{U}}) = +\infty$. We define now preferences of the form $x'_a = (x_a, e_a) \succeq_{\epsilon, \theta} (y_b, e_b) = y'_b$ if and only if $u(x_a, \theta) + \epsilon_a \geq u(x_b, \theta) + \epsilon_b$. Then , it follows that $VC^+(\{\succeq_{\epsilon, \theta}\}_{(\epsilon, \theta) \in \mathbb{R}^A \times \Theta}) = +\infty$. Preference heterogeneity could be specified by a product probability measure $\zeta_{\theta} \otimes \zeta_{\epsilon} \in \Delta(\Theta \times \mathbb{R}^A)$, where ζ_{ϵ} follows a multivariate normal distribution as a product of independent normal distributions with parameters $\{(\mu_a, \sigma_a^2)\}_{a \in A}$. Then, without any restrictions on ζ_{θ} and not restricting $\{(\mu_a, \sigma_a^2)\}_{a \in A}$, it obtains that the implied stochastic choice model would be non-learnable. This follows from the same arguments as in the proof of [Proposition 4](#).

For these models, the above result establishes non-learnability or irrecoverability of heterogeneity in a certain sense (uniform). The dimension being infinite has the consequence that the preference class is complex enough to significantly affect the predictability of the model. The variability in choice behaviour permitted by the model is large enough to deter good worst case guarantees on out-of-sample estimation error.

⁸When $|\Omega| = 2$, the Max-min model has VC dimension equal to 2.

5 Discussion

5.1 Cognitive Heterogeneity

In this section, we will study stochastic choice in the setting where there is data on choice from menus of alternatives. As before, A is a finite set of alternatives. Hence, we will have $\mathcal{X} = 2^A \setminus \{\emptyset\}$. For a choice probability function $\sigma : \mathcal{X} \rightarrow \Delta(A)$, for each menu of alternatives $x = A \in \mathcal{X}$, the probability vector $\sigma(A)$ has full support in A . For each $a \in A$, the quantity $\sigma_a(A)$ denotes the probability of alternative a being chosen from the menu A . We will assume the following holds for σ .

Definition 4. (*Positivity*) A stochastic choice function σ satisfies positivity if for each menu A and $a \in A$, we have $\sigma_a(A) > 0$.

In this context, the choice probabilities from the Logit or the Luce model are as follows. There exists a function $w : A \rightarrow \mathbb{R}_{++}$, which assigns weights to the various alternatives. For each menu A , the alternative $a \in A$ is chosen with probability

$$\sigma_a(A) = \frac{w(a)}{\sum_{b \in A} w(b)}.$$

It is well-known that under positivity, a stochastic choice map σ has the Luce representation if and only if σ satisfies the Independence of Irrelevant Alternatives axiom (IIA). The IIA axiom says that if A and B are two menus and we have two alternatives $a, b \in A \cap B$. Then,

$$\frac{\sigma_a(A)}{\sigma_b(A)} = \frac{\sigma_a(B)}{\sigma_b(B)}.$$

Another example is the stochastic choice model involving consideration sets ([Manzini and Mariotti \[2007\]](#)). The model is defined as follows. There is a function $\gamma : A \rightarrow (0, 1)$ and a strict order $>$ on A . The choice probabilities are given as follows.

$$\sigma_a^{\gamma, >}(A) := (1 - \gamma(a)) \prod_{b \in A: b > a} \gamma(b).$$

The interpretation is that $1 - \gamma(a)$ is the probability with which alternative a will be considered in choosing from a given menu A . Once the considered alternatives are determined, the agent chooses according to the strict order $>$. The following result obtains in this setting.

Proposition 5. *Suppose Σ is either the Logit/Luce model or the model with consideration sets. Then, under the log rule S^{log} , the Pollard dimension of $S^{\text{log}} \circ \Sigma$ is at most $O(|A|^2)$.*

Proof. The proof is in the appendix and applies a result on the VC dimension of neural networks (see Proposition 8). \square

The techniques from the previous section can be extended to allow recoverability of cognitive heterogeneity. This type of heterogeneity corresponds to a distribution μ over $(\gamma, >)$. By arguments similar to the proof of Proposition 2, one shows that finite pollard dimension of $S^{\text{log}} \circ \Sigma$ allows for consistent estimation of the true distribution μ_0 , if we place a restriction that the choice probabilities are bounded below. Since we are applying the log scoring rule, the estimator is essentially based on the conditional maximum likelihood procedure.

5.2 Non-uniform Learning

The notion of consistency considered above requires that convergence should take place uniformly over both the distribution over characteristics π_0 and the true stochastic choice function σ_0 . The main advantage of such a requirement is that it leads to estimates that are robust with respect to the data generating process. Suppose, we instead consider non-uniform consistency which only requires that the divergence between true and estimated choice probabilities converges to zero. We provide a definition.

Definition 5. *A learning rule $\hat{\sigma}$ is pointwise consistent (with respect to d and Σ) if for all $0 < \varepsilon, \delta < 1$, and for all $\pi_0 \in \Delta(\mathcal{X})$, $\sigma_0 \in \Sigma$, there exists $N(\varepsilon, \delta, \pi_0, \sigma_0)$ such that for all $n \geq N(\varepsilon, \delta, \pi_0, \sigma_0)$,*

$$\pi_0 \otimes \sigma_0^n \left(\{z^n : d_{\pi_0}(\hat{\sigma}(z^n), \sigma_0) < \varepsilon\} \right) \geq 1 - \delta.$$

We say that a model Σ is pointwise learnable with respect to d if there exists a learning rule $\hat{\sigma}$, which is pointwise consistent with respect to d and Σ .

The above is also a kind of PAC criterion but only requires that the random variable $d_{\pi_0}(\hat{\sigma}(z^n), \sigma_0)$ converge in probability to zero. Note that we had a stronger notion of consistency before. Hence, if a stochastic choice model is learnable according to Definition 1, then it is also pointwise learnable. It turns out that one can study pointwise learnability via a weaker definition of Glivenko-Cantelli classes called universal Glivenko-Cantelli classes (see, for example, [Dudley et al. \[1991\]](#), [van Handel \[2013\]](#)). This requires that for a real-valued function class \mathcal{F} (defined on a set \mathcal{Z}) we have that

$\sup_{f \in \mathcal{F}} |1/n \sum_{i=1}^n f(z_i) - \mathbb{E}_\nu[f]|$ converges to zero for each probability measure $\nu \in \Delta(\mathcal{Z})$. In contrast to uniform Glivenko-Cantelli classes, the convergence need not be uniform over ν . Note, however, we still have uniformity over \mathcal{F} as we have convergence of the supremum. This would also be closer to the uniformity criterion typically required in extremum estimation (see Amemiya [1985]). Similar to the proof of Proposition 1, one can argue that for a stochastic choice model Σ and scoring rule S , if $S \circ \Sigma$ is universal Glivenko-Cantelli, then Σ is pointwise learnable w.r.t d_S .

A useful characterisation of universal Glivenko-Cantelli classes is given by van Handel [2013] in terms of Boolean σ -independence. A sequence of functions $\{f_j\}_{j \in \mathbb{N}} \subseteq \mathcal{F}$ is Boolean σ -independent at levels $\beta > \alpha > 0$ if for every subset $F \subseteq \mathbb{N}$,

$$\bigcap_{j \in F} \{z \mid f_j(z) < \alpha\} \cap \bigcap_{j \notin F} \{z \mid f_j(z) > \beta\} \neq \emptyset.$$

It then holds that a function class is universal Glivenko-Cantelli if and only if there exists no Boolean σ -independent sequence in \mathcal{F} for any α, β . From this, it follows that a parametrised function class $\mathcal{F} = \{f(\cdot, \theta) \mid \theta \in \Theta\}$, where Θ is a compact metric space and $f(z, \cdot)$ is continuous in θ for each z , is a universal Glivenko-Cantelli class. We can then show that under some conditions, a stochastic preference model supported on continuous preferences would be pointwise learnable. As before, it will be convenient to work with the closed convergence topology on continuous preferences (see appendix). We state the result below.

Proposition 6. *Let $\mathcal{P}_{\text{CONT}}$ be the set of all continuous preferences over the set of alternatives X , equipped with the closed convergence topology. Further, let $\mathcal{M} \subseteq \Delta(\mathcal{P}_{\text{CONT}})$ be any set of probability measures satisfying the condition 17, that is closed in the topology of weak convergence. Then, $\Sigma^{\mathcal{P}_{\text{CONT}}}(\mathcal{M}) := \{\sigma(\cdot; \mu \mid \mu \in \mathcal{M})\}$ is pointwise learnable with respect to any L^p metric.*

Proof. Can be found in the appendix. □

It is interesting to compare the above result with the results in Proposition 2 and Proposition 4. Firstly, there may be continuous preferences with infinite VC dimension (for eg. all continuous convex preferences) and there may be discontinuous preferences with finite VC dimension (for eg. lexicographic preferences). Secondly, conditional on considering only continuous preferences, which allows pointwise learnability, robust estimation is characterised by the dimension of the preference classes (VC and VC^+). If

the dimension is finite, we get uniform consistency but for infinite dimension, the estimates are non-robust, to an extent that the same estimation procedure that would guarantee pointwise learning, would lead to erroneous out-of-sample predictions in the worst case. Lastly, the techniques for obtaining rates of convergence for pointwise learnability, would involve deriving bounds on the bracketing or covering numbers of the function class $S \circ \Sigma$ (see [Dudley et al. \[1991\]](#)). This would also be similar to PAC learning problems in the deterministic case with a non-uniform learnability criterion (see [Blumer et al. \[1987\]](#), [Benedek and Itai \[1994\]](#), [Shalev-Shwartz and Ben-David \[2014\]](#)) and also structural risk minimisation (see [Vapnik and Chervonenkis \[1974\]](#) and [Vapnik \[1998\]](#)).

References

- Naoki Abe, Manfred K Warmuth, and Jun-ichi Takeuchi. Polynomial learnability of probabilistic concepts with respect to the kullback-leibler divergence. In *Proceedings of the fourth annual workshop on Computational learning theory*, pages 277–289. Morgan Kaufmann Publishers Inc., 1991.
- Naoki Abe, Jun-ichi Takeuchi, and Manfred K Warmuth. Polynomial learnability of stochastic rules with respect to the kl-divergence and quadratic distance. *IEICE Transactions on Information and Systems*, 84(3):299–316, 2001.
- Charalambos D Aliprantis and Kim C Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer Science & Business Media, 2006.
- Noga Alon, Shai Ben-David, Nicolo Cesa-Bianchi, and David Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. *Journal of the ACM (JACM)*, 44(4): 615–631, 1997.
- Takeshi Amemiya. *Advanced econometrics*. Harvard university press, 1985.
- Martin Anthony and Peter L Bartlett. *Neural network learning: Theoretical foundations*. cambridge university press, 2009.
- Susan Athey and Guido W Imbens. Machine learning methods that economists should know about. *Annual Review of Economics*, 11, 2019.
- Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.

- Pathikrit Basu and Federico Echenique. Learnability and models of decision making under uncertainty. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 53–53. ACM, 2018.
- Eyal Beigman and Rakesh Vohra. Learning from revealed preference. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 36–42. ACM, 2006.
- Gyora M Benedek and Alon Itai. Nonuniform learnability. *Journal of Computer and System Sciences*, 48(2):311–323, 1994.
- Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890, 1995.
- Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. Occam’s razor. *Information processing letters*, 24(6):377–380, 1987.
- Kim C Border and Uzi Segal. Dynamic consistency implies approximately expected utility preferences. *Journal of Economic Theory*, 63(2):170–188, 1994.
- Richard A Briesch, Pradeep K Chintagunta, and Rosa L Matzkin. Nonparametric discrete choice models with unobserved heterogeneity. *Journal of Business & Economic Statistics*, 28(2):291–307, 2010.
- Andreas Buja, Werner Stuetzle, and Yi Shen. Loss functions for binary class probability estimation and classification: Structure and applications. *Working draft, November, 3, 2005*.
- Zachary Chase and Siddharth Prasad. Learning Time Dependent Choice. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*, volume 124 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 62:1–62:19. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018. ISBN 978-3-95977-095-8.
- Soo Hong Chew, Larry G Epstein, Uzi Segal, et al. Mixture symmetry and quadratic utility. *Econometrica*, 59(1):139–63, 1991.
- Pierpaolo De Blasi, Lancelot F James, John W Lau, et al. Bayesian nonparametric estimation and consistency of mixed multinomial logit choice models. *Bernoulli*, 16(3): 679–704, 2010.
- Anthony Downs et al. *An economic theory of democracy*. 1957.

- Richard M Dudley. *Uniform central limit theorems*, volume 142. Cambridge university press, 2014.
- Richard M Dudley, Evarist Giné, and Joel Zinn. Uniform and universal glivenko-cantelli classes. *Journal of Theoretical Probability*, 4(3):485–510, 1991.
- Drew Fudenberg, Ryota Iijima, and Tomasz Strzalecki. Stochastic choice and revealed perturbed utility. *Econometrica*, 83(6):2371–2409, 2015.
- Itzhak Gilboa. *Theory of decision under uncertainty*, volume 45. Cambridge university press, 2009.
- Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Faruk Gul and Wolfgang Pesendorfer. Random expected utility. *Econometrica*, 74(1):121–146, 2006.
- Werner Hildenbrand. On economies with many agents. *Journal of economic theory*, 2(2): 161–188, 1970.
- Werner Hildenbrand. *Core and Equilibria of a Large Economy.(PSME-5)*. Princeton university press, 2015.
- Gil Kalai. Learnability and rationality of choice. *Journal of Economic theory*, 113(1):104–117, 2003.
- Yakar Kannai. Continuity properties of the core of a market. *Econometrica (pre-1986)*, 38 (6):791, 1970.
- Michael J Kearns and Robert E Schapire. Efficient distribution-free learning of probabilistic concepts. *Journal of Computer and System Sciences*, 48(3):464–497, 1994.
- Michael J Kearns and Umesh Vazirani. *An introduction to computational learning theory*. MIT press, 1994.
- Alexander Kechris. *Classical descriptive set theory*, volume 156. Springer Science & Business Media, 2012.
- David Kreps. *Notes on the Theory of Choice*. Routledge, 1988.
- Jay Lu. Random choice and private information. *Econometrica*, 84(6):1983–2027, 2016.

- R Duncan Luce. *Individual choice behavior: A theoretical analysis*. John Wiley, 1959.
- R Duncan Luce. The choice axiom after twenty years. *Journal of mathematical psychology*, 15(3):215–233, 1977.
- Charles F Manski. Maximum score estimation of the stochastic utility model of choice. *Journal of econometrics*, 3(3):205–228, 1975.
- Charles F Manski. Semiparametric analysis of discrete response: Asymptotic properties of the maximum score estimator. *Journal of econometrics*, 27(3):313–333, 1985.
- Paola Manzini and Marco Mariotti. Sequentially rationalizable choice. *American Economic Review*, 97(5):1824–1839, 2007.
- Andreu Mas-Colell. On the continuous representation of preorders. *International Economic Review*, pages 509–513, 1977.
- Rosa L Matzkin. Nonparametric and distribution-free estimation of the binary threshold crossing and the binary choice models. *Econometrica: Journal of the Econometric Society*, pages 239–270, 1992.
- Daniel McFadden. Modeling the choice of residential location. *Transportation Research Record (673)*, 1978.
- Daniel McFadden and Kenneth Train. Mixed mnl models for discrete response. *Journal of applied Econometrics*, 15(5):447–470, 2000.
- Daniel L McFadden. Econometric analysis of qualitative response models. *Handbook of econometrics*, 2:1395–1457, 1984.
- Ilya Molchanov. *Theory of random sets*, volume 19. Springer, 2005.
- Matthew Parry et al. Linear scoring rules for probabilistic binary classification. *Electronic Journal of Statistics*, 10(1):1596–1607, 2016.
- David Pollard. *Convergence of stochastic processes*. Springer Science & Business Media, 1984.
- Leonard J Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- Reinhard Selten. Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, 1(1):43–61, 1998.

- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Michel Talagrand et al. The glivenko-cantelli problem. *The Annals of Probability*, 15(3): 837–870, 1987.
- Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- Aad W Van Der Vaart and Jon A Wellner. Weak convergence. In *Weak convergence and empirical processes*, pages 16–28. Springer, 1996.
- Ramon van Handel. The universal glivenko–cantelli property. *Probability Theory and Related Fields*, 155(3-4):911–934, 2013.
- Vladimir Vapnik. *Statistical learning theory. 1998*, volume 3. Wiley, New York, 1998.
- Vladimir Vapnik and Alexey Chervonenkis. Theory of pattern recognition, 1974.
- Vladimir N Vapnik and Aleksei Yakovlevich Chervonenkis. The uniform convergence of frequencies of the appearance of events to their probabilities. In *Doklady Akademii Nauk*, volume 181, pages 781–783. Russian Academy of Sciences, 1968.
- Kenji Yamanishi. A learning criterion for stochastic rules. *Machine Learning*, 9(2-3):165–203, 1992.
- Morteza Zadimoghaddam and Aaron Roth. Efficiently learning from revealed preference. In *International Workshop on Internet and Network Economics*, pages 114–127. Springer, 2012.

6 Appendix

6.1 Technical

When dealing with preferences in economic contexts, a popular topology on the space of preference relations is that given by the topology of closed convergence (Kannai [1970], Hildenbrand [1970]). We provide some general definitions in this section and note some important results from the literature that will be useful in our setting. Let (Y, τ_Y) be a topological space. Let \mathcal{C} denote the set of all closed subsets of Y . The closed convergence

topology on \mathcal{C} is defined as follows (see also [Aliprantis and Border \[2006\]](#)). For any compact $K \subseteq Y$ and finitely many open sets $\{U_i\}_{i \in I}$ in Y , we define the following subset of \mathcal{C} .

$$\mathcal{O}(K, \{U_i\}_i) = \{F \in \mathcal{C} \mid F \cap K = \emptyset \text{ and } F \cap U_i \neq \emptyset \text{ for each } i \in I\}.$$

All sets of the form $\mathcal{O}(K, \{U_i\}_i)$ constitute a basis for the closed convergence topology on \mathcal{C} . We will denote this topology as $\tau_{\mathcal{C}}$. When Y is a locally compact and separable metric space, then it turns out that $(\mathcal{C}, \tau_{\mathcal{C}})$ is a compact metrizable space (see [Hildenbrand \[2015\]](#)). Further, it also holds that $(\mathcal{C}, \tau_{\mathcal{C}})$ is second countable (see [Molchanov \[2005\]](#)).

We now note how the closed convergence topology provides us a way to topologize the space of continuous preferences. Let (X, τ_X) be a topological space of alternatives. A preference relation on X is a binary relation $\succeq \subseteq X \times X$ which is complete and transitive. A continuous preference relation \succeq is a preference relation that is closed in $X \times X$ (under the product topology). Hence, a class of continuous preferences $\mathcal{P} \subseteq 2^{X \times X}$ can be endowed with the subspace topology, say $\tau_{\mathcal{P}}$, generated by the topology of closed convergence on closed subsets of $2^{X \times X}$.

The following theorem, due to [Mas-Colell \[1977\]](#), provides conditions under which a utility representation exists that is jointly continuous in the alternatives and the preferences.

Proposition 7. ([Mas-Colell \[1977\]](#)) *Let X be a locally compact and second countable space of alternatives. Further, let \mathcal{P}_{CONT} denote the set of all continuous preference relations on X , equipped with the subspace topology induced by the closed convergence topology. Then, there exists a utility function $U : X \times \mathcal{P}_{CONT} \rightarrow \mathbb{R}$, which is jointly continuous and each $\succeq \in \mathcal{P}_{CONT}$ is represented by $U(\cdot, \succeq)$.*

An immediate corollary of the above result is the following. For any finite set of alternatives $(x_a)_{a \in A}$, the set of preferences $\{\succeq \in \mathcal{P}_{CONT} \mid x_a \succeq x_b \text{ for all } b \neq a\}$ is closed and $\{\succeq \in \mathcal{P}_{CONT} \mid x_a \succ x_b \text{ for all } b \neq a\}$ is open with respect to the closed convergence topology.

Now, consider a subspace of preferences $\mathcal{P} \subseteq \mathcal{P}_{CONT}$. We denote as $\Delta(\mathcal{P})$, the set of all Borel probability measures defined on \mathcal{P} . We will endow $\Delta(\mathcal{P})$ with the topology of weak convergence (see [Aliprantis and Border \[2006\]](#)). Hence, for a sequence $\{\mu_n\}_n$ and probability measure μ in $\Delta(\mathcal{P})$ we have that $\mu_n \rightarrow \mu$ in the weak convergence topology if and only if for every bounded continuous function $f : \mathcal{P} \rightarrow \mathbb{R}$, the convergence $\mathbb{E}_{\mu_n}[f(\succeq)] \rightarrow \mathbb{E}_{\mu}[f(\succeq)]$

holds. Now, if the space of alternatives X is a locally compact and separable metric space, then it is also the case that $X \times X$ is a locally compact and separable metric space. Hence, the subspace $\mathcal{P}_{CONT} \subseteq 2^{X \times X}$ is a second countable compact metric space. This means that for any subspace $\mathcal{P} \subseteq \mathcal{P}_{CONT}$, the set of probability measures $\Delta(\mathcal{P})$ with the weak convergence topology is a separable metrizable space (see [Aliprantis and Border \[2006\]](#)). We now show some measurability and continuity properties of the stochastic choice probabilities $\sigma_a(x; \mu)$.

We will first define the regularity condition that we assume on $\mu \in \Delta(\mathcal{P})$. Consider the set of preferences where x_a is strictly preferred from a menu of alternatives $x = (x_b)_{b \in A}$ i.e. $\mathcal{Q}(x, a) = \{\succ \in \mathcal{P} \mid x_a > x_b \text{ for all } b \neq a\}$. The regularity assumption imposes on μ that we have $\mu(\cup_{a \in A} \mathcal{Q}(x, a)) = 1$. Also, by the definition of choice probabilities, $\sigma_a(x, \mu) = \mu(\mathcal{Q}(x, a))$. The following lemma is useful.

Lemma 2. *Let X be a locally compact, separable metric space of alternatives and let $\mathcal{X} \subseteq X^A$ be such that $x_a \neq x_b$ for all $x = (x_c)_{c \in A} \in \mathcal{X}$. Suppose $\mathcal{P} \subseteq \mathcal{P}_{CONT}$ and let $\Delta^r(\mathcal{P}) \subseteq \Delta(\mathcal{P})$ be the set of all probability measures on the preference class \mathcal{P} satisfying the regularity condition. The stochastic choice function $\sigma_a(x; \mu)$ is continuous in μ for each $x \in \mathcal{X}$; continuous in x for each $\mu \in \Delta^r(\mathcal{P})$. Hence, $\sigma_a(x; \mu)$ is jointly measurable under the Borel sigma-algebra induced by the product topology on $\mathcal{X} \times \Delta^r(\mathcal{P})$.*

Proof. Consider a fixed $x \in \mathcal{X}$ and let $\{\mu_n\}_n$ be a sequence and let μ be a probability measure in $\Delta^r(\mathcal{P})$ such that $\mu_n \rightarrow \mu$ in the weak convergence topology. It suffices to argue that $\mathcal{Q}(x, a)$ is a continuity set according to the measure μ i.e. $\mu(\partial \mathcal{Q}(x, a)) = 0$, where the set $\partial \mathcal{Q}(x, a)$ is the boundary of the set $\mathcal{Q}(x, a)$. Suppose we have a convergent sequence $\{\succ_n\}_n$ in $\mathcal{Q}(x, a)$ with limit \succ . As μ satisfies the regularity condition $\mu(\cup_{a \in A} \mathcal{Q}(x, a)) = 1$, it suffices to show that $\succ \in \{\succ' \in \mathcal{P}_{CONT} \mid x_a \succ' x_b \text{ for all } b \neq a\}$. But this is true since the latter set is closed and contains the set $\mathcal{Q}(x, a)$ as noted above. It now follows that $\sigma_a(x; \cdot)$ is continuous.

Now, let $\mu \in \Delta^r(\mathcal{P})$ and suppose $\{x_n\} \subseteq \mathcal{X}$ is a convergent sequence with limit $x \in \mathcal{X}$. Note that since preferences are continuous, we have that $\mathcal{Q}(x, a) \subseteq \liminf_n \mathcal{Q}(x_n, a)$. This implies that $\mu(\mathcal{Q}(x, a)) \leq \mu(\liminf_n \mathcal{Q}(x_n, a)) \leq \liminf_n \mu(\mathcal{Q}(x_n, a))$. Hence, $\sigma_a(\cdot; \mu)$ is lower semicontinuous. We now show upper-semicontinuity. Define the set $\mathcal{Q}(x, A \setminus \{a\}) := \cup_{b \in A \setminus \{a\}} \mathcal{Q}(x, b)$. Note that since μ satisfies regularity, we have $\mu(\mathcal{Q}(y, a)) = 1 - \mu(\mathcal{Q}(y, A \setminus \{a\}))$ for all $y \in \mathcal{X}$. Now, let $x_n \rightarrow x$. Again, as preferences are continuous, we obtain that $\mathcal{Q}(x, A \setminus \{a\}) \subseteq \liminf_n \mathcal{Q}(x_n, A \setminus \{a\})$ which means $\mu(\mathcal{Q}(x, A \setminus \{a\})) \leq \mu(\liminf_n \mathcal{Q}(x_n, A \setminus \{a\})) \leq \liminf_n \mu(\mathcal{Q}(x_n, A \setminus \{a\}))$.

But this implies that $\mu(\mathcal{Q}(x, a)) \geq \limsup_n \mu(\mathcal{Q}(x_n, a))$. Hence, the function $\sigma_a(\cdot; \mu)$ is continuous as it is both lower and upper semi-continuous.

The two facts established above imply that $\sigma_a(x; \mu)$ is a Caratheodory function (see, for example, [Aliprantis and Border \[2006\]](#), [Kechris \[2012\]](#)). Hence, $\sigma_a(x; \mu)$ is jointly measurable in $(x; \mu)$. \square

Now it follows that under the conditions satisfied by the above lemma, we obtain appropriate measurability of the function class $\mathcal{F} = S^{br} \circ \Sigma^{\mathcal{P}}$ for the Glivenko-Cantelli theorems to apply. From [Van Der Vaart and Wellner \[1996\]](#), a sufficient condition is a seperability condition called pointwise measurability. It requires that there exist a countable subset $\mathcal{F}' \subseteq \mathcal{F}$ such that for any $f \in \mathcal{F}$, there exists a sequence in \mathcal{F}' which converges to f pointwise. Since $\Delta(\mathcal{P})$ is a seperable metric space, it follows that $\Delta^r(\mathcal{P})$ is a seperable metric space as well. Let $\Pi \subseteq \Delta^r(\mathcal{P})$ be a countable dense subset. Since from Lemma 2 above, the choice probability $\sigma_a(x, \mu)$ is continuous in μ and from the definition of the Quadratic scoring rule S^{br} , it follows that the countable set $\mathcal{F}' = \{S^{br} \circ \sigma(\cdot; \mu) \mid \mu \in \Pi\}$ makes \mathcal{F} pointwise measurable. Alternatively, [Talagrand et al. \[1987\]](#) requires that \mathcal{F} belong to $L^1(\nu)$, for each $\nu \in \Delta(X \times A)$.

We will now discuss the existence of probability measures μ which satisfy the regularity condition. The following result applies when preferences are defined through a parametrised utility function $u : X \times \Phi \rightarrow \mathbb{R}$. For any such function u and $\phi \in \Phi$, we will denote as \succeq_ϕ , the preference relation u induces on X . Hence, we define $x \succeq_\phi y$ if and only if $u(x, \phi) \geq u(y, \phi)$. We will say that a probability measure $\xi \in \Delta(\Phi)$ satisfies the regularity condition for u if for all $x = (x_b)_{b \in A} \in X^A$ with $x_a \neq x_b$ for all $a \neq b$, we have $\xi(\cup_{a \in A} \{\phi \mid u(x_a, \phi) > \max_{b \neq a} u(x_b, \phi)\}) = 1$.

The following result holds. We will assume that preferences are locally strict (see [Border and Segal \[1994\]](#)).

Lemma 3. *Let X be a locally compact and seperable metric space of alternatives, \mathcal{P}_{CONT} be the set of all continuous preference relations on X (endowed with the topology induced by the closed convergence topology) and Φ be a metric space of parameters. Let $u : X \times \Phi \rightarrow \mathbb{R}$ be a jointly continuous utility function. Further, suppose that u is locally strict i.e. for all ϕ , if $u(x, \phi) \geq u(y, \phi)$, then every neighbourhood of (x, y) contains a pair (x', y') such that $u(x', \phi) > u(y', \phi)$. Then, the function $\zeta : \Phi \rightarrow \mathcal{P}_{CONT}$ defined as*

$$\zeta(\phi) = \succeq_\phi,$$

is a continuous function. Now, suppose $\xi \in \Delta(\Phi)$ is a Borel probability measure on the space Φ which satisfies the regularity condition for u . Also suppose that the class of preferences $\mathcal{P}_u = \{\succ_{\phi}\}_{\phi \in \Phi}$ is Borel. Then, the induced probability measure μ_{ξ} on the Borel subsets of \mathcal{P}_{CONT} , defined as

$$\mu_{\xi}(\mathcal{B}) = \xi(\zeta^{-1}(\mathcal{B})),$$

satisfies the regularity condition on \mathcal{P}_u . If \mathcal{P}_u is also closed in \mathcal{P}_{CONT} and ξ has full support on Φ , then μ_{ξ} has full support on \mathcal{P}_u .

Proof. To show that ζ is continuous, it will suffice to show that the pre-image of any basic open set is open in Φ . Let $\mathcal{O}(K, \{U_i\}_{i \in I})$ be a basic open set. Now, each U_i is open in $X \times X$. Hence, from local strictness of u it follows that for any ϕ , if $u(x, \phi) \geq u(y, \phi)$, where $(x, y) \in U_i$, there exists another pair $(x', y') \in U_i$ such that $u(x', \phi) > u(y', \phi)$. For convenience, define the function $\Delta_u(x, y, \phi) = u(x, \phi) - u(y, \phi)$. Then, we get that the pre-image of $\mathcal{O}(K, \{U_i\}_{i \in I})$ can be written as

$$\zeta^{-1}(\mathcal{O}(K, \{U_i\}_{i \in I})) = \{\phi \in \Phi \mid \max_{(x,y) \in K} \Delta_u(x, y, \phi) < 0 \text{ and } \sup_{(x,y) \in U_i} \Delta_u(x, y, \phi) > 0 \text{ for each } i \in I\}.$$

Now, as K is compact, from the theorem of the maximum, it follows that $\max_{(x,y) \in K} \Delta_u(x, y, \phi)$ is continuous in ϕ . Hence, the set $\Phi_K = \{\phi \in \Phi \mid \max_{(x,y) \in K} \Delta_u(x, y, \phi) < 0\}$ is open in Φ . Now, note that $\sup_{(x,y) \in U_i} \Delta_u(x, y, \phi)$ is a supremum of lower semi-continuous functions and is hence, lower semi-continuous in ϕ . Hence, the set $\Phi_i = \{\phi \mid \sup_{(x,y) \in U_i} \Delta_u(x, y, \phi) > 0\}$ is open, for each $i \in I$. To conclude, observe that $\zeta^{-1}(\mathcal{O}(K, \{U_i\}_{i \in I})) = \Phi_K \cap (\bigcap_{i \in I} \Phi_i)$ is open as an intersection of finitely many open sets. Hence, ζ is continuous.

We now argue that μ_{ξ} satisfies the regularity condition if ξ satisfies the regularity condition for u . Let $\mathcal{B} = \bigcup_{a \in A} \mathcal{Q}(x, a)$. Then $\mu_{\xi}(\mathcal{B}) = \xi(\bigcup_{a \in A} \{\phi \mid u(x_a, \phi) > \max_{b \neq a} u(x_b, \phi)\}) = 1$. Now, further assume that ξ has full support on Φ and \mathcal{P}_u is closed. Suppose for contradiction that there exists a closed strict subset $\mathcal{P} \subsetneq \mathcal{P}_u$ such that $\mu_{\xi}(\mathcal{P}) = 1$. But this means that $\xi(\zeta^{-1}(\mathcal{P})) = 1$ where the set $\zeta^{-1}(\mathcal{P})$ is closed as ζ is continuous and \mathcal{P} is closed in \mathcal{P}_{CONT} . But $\zeta^{-1}(\mathcal{P}) \subsetneq \zeta^{-1}(\mathcal{P}_u) = \Phi$, where both sets are closed but $\xi(\zeta^{-1}(\mathcal{P})) = 1$ meaning that ξ does not have full support, which is a contradiction. \square

6.2 Proofs from Section 3

The following is the proof for Proposition 1.

Proof. We show that Σ is learnable with respect to d_S . We shall construct an almost-ERM learning $\hat{\sigma}_E$, which would be consistent with respect to d and Σ . For each $n \in \mathbb{N}$, define $\varepsilon_n := 1/n$. Now, for $z^n = ((x_i, a_i))_{i=1}^n \in \mathcal{Z}^n$, let $\hat{\sigma}_E(z^n)$ be such that

$$\hat{V}_S(\hat{\sigma}_E(z^n)) \leq \inf_{\sigma \in \Sigma} \hat{V}_S(\sigma) + \varepsilon_n,$$

where recall that for any $\sigma \in \Sigma$, $\hat{V}(\sigma) := 1/n \sum_{i=1}^n S_{a_i}(\sigma(x_i))$. Note that the infimum exists in the above definition since $S \circ \Sigma$ is bounded above by M . Hence, the learning rule $\hat{\sigma}_E$ is well-defined. We now show it is consistent.

Let $(\varepsilon, \delta) \in (0, 1)^2$. Since $S \circ \Sigma$ is a Glivenko-Cantelli class of functions, for $(\varepsilon/3, \delta) \in (0, 1)^2$, there exists $N'(\varepsilon/3, \delta)$ such that for all $n \geq N'(\varepsilon/3, \delta)$, for all $\nu \in \Delta(\mathcal{Z})$

$$\nu^n \left(z^n : \sup_{f \in S \circ \Sigma} |1/n \sum_{i=1}^n f(z_i) - \mathbb{E}_\mu(f)| < \varepsilon/3 \right) \geq 1 - \delta. \quad (23)$$

Now, let $N(\varepsilon, \delta) := \max\{3/\varepsilon, N'(\varepsilon, \delta)\}$. Now, suppose $n \geq N(\varepsilon, \delta)$. Let $\pi_0 \in \Delta(\mathcal{X})$ be a distribution over characteristics and suppose the true choice probabilities are given by $\sigma_0 \in \Sigma$. For convenience, denote $\mathbb{E}_{z^n}(f) = 1/n \sum_{i=1}^n f(z_i)$.

Now suppose z^n is such that $\sup_{f \in S \circ \Sigma} |\mathbb{E}_{z^n}(f) - \mathbb{E}_{\pi_0 \otimes \sigma_0}(f)| < \varepsilon/3$. Note that this happens with probability at least $1 - \delta$ under $\pi_0 \otimes \sigma_0^n$. Consider the almost-ERM rule $\hat{\sigma}_E$. Define also $f_{E, z^n}(x, a) := S(\hat{\sigma}_E(z^n)(x), a)$, $f_0(x, a) := S(\sigma_0(x), a)$ and $\nu_0 = \pi_0 \otimes \sigma_0$. It follows that

$$|\mathbb{E}_{\nu_0}(f_{E, z^n}) - \mathbb{E}_{\nu_0}(f_0)| = \mathbb{E}_{\nu_0}(f_0) - \mathbb{E}_{\nu_0}(f_{E, z^n}) \quad (24)$$

$$\begin{aligned} &= \mathbb{E}_{\nu_0}(f_0) - \mathbb{E}_{z^n}(f_{E, z^n}) + \mathbb{E}_{z^n}(f_{E, z^n}) - \mathbb{E}_{\nu_0}(f_{E, z^n}) \\ &\leq \mathbb{E}_{\nu_0}(f_0) - \mathbb{E}_{z^n}(f_{E, z^n}) + \varepsilon/3 \end{aligned} \quad (25)$$

$$\leq \mathbb{E}_{\nu_0}(f_0) - \mathbb{E}_{z^n}(f_0) + 1/n + \varepsilon/3 \quad (26)$$

$$\leq \varepsilon/3 + \varepsilon/3 + \varepsilon/3 \quad (27)$$

$$= \varepsilon.$$

Here, 24 follows since σ_0 minimises expected risk as S is incentive compatible (Lemma 1). The inequality 25 follows by assumption that $\sup_{f \in S \circ \Sigma} |\mathbb{E}_{z^n}(f) - \mathbb{E}_{\nu_0}(f)| < \varepsilon/3$. This yields $\mathbb{E}_{z^n}(f_{E, z^n}) - \mathbb{E}_{\nu_0}(f_{E, z^n}) < \varepsilon/3$. Also, 26 follows since $\hat{\sigma}_E$ is almost-ERM. Lastly, 27 follows since $n \geq 3/\varepsilon$ and again from our assumption that $\sup_{f \in S \circ \Sigma} |\mathbb{E}_{z^n}(f) - \mathbb{E}_{\nu_0}(f)| < \varepsilon/3$, which yields $\mathbb{E}_{z^n}(f_0) - \mathbb{E}_{\nu_0}(f_0) < \varepsilon/3$.

Now, since we have

$$\mathbb{E}_{\nu_0}(f_0) - \mathbb{E}_{\nu_0}(f_{E,z^n}) = \int_{\mathcal{X}} d_S(\hat{\sigma}_E(z^n)(x), \sigma_0(x)) d\pi_0(x),$$

it follows that $\hat{\sigma}_E$ is consistent with respect to d_S and Σ . □

6.3 Proofs from Section 4

The following two lemmas will be useful.

Lemma 4. *Suppose \mathcal{F} is a uniformly bounded class of real-valued functions defined on the set \mathcal{Z} and let $\mathcal{M} \subseteq \Delta(\mathcal{F})$.⁹ Define the function f_μ as*

$$f_\mu(z) = \int_{\mathcal{F}} f(z) d\mu(f).$$

Consider the function class $\mathcal{F}_{\mathcal{M}} = \{f_\mu \mid \mu \in \mathcal{M}\}$. Then,

$$R_{z^n}(\mathcal{F}_{\mathcal{M}}) \leq R_{z^n}(\mathcal{F}).$$

for all $z^n \in \mathcal{Z}^n$.

Proof. We first develop some notation. For z^n and a real-valued function f , define $f(z^n) = (f(z_1), \dots, f(z_n)) \in \mathbb{R}^n$. From the definition of Rademacher Complexity, it suffices to show that for any real vector $\alpha \in \mathbb{R}^n$,

$$\sup_{\mu \in \mathcal{M}} \int_{\mathcal{F}} \alpha \cdot f(z^n) d\mu(f) \leq \sup_{f \in \mathcal{F}} \alpha \cdot f(z^n). \quad (28)$$

Suppose for contradiction that 28 does not hold. Then, $\sup_{\mu \in \mathcal{M}} \int_{\mathcal{F}} \alpha \cdot f(z^n) d\mu(f) > \sup_{f \in \mathcal{F}} \alpha \cdot f(z^n)$. Then, there exists $\mu^* \in \mathcal{M}$ such that $\int_{\mathcal{F}} \alpha \cdot f(z^n) d\mu^*(f) > \sup_{f \in \mathcal{F}} \alpha \cdot f(z^n)$. But this means that there exists f^* such that $\alpha \cdot f^*(z^n) > \sup_{f \in \mathcal{F}} \alpha \cdot f(z^n)$. This is a contradiction and hence, the result follows. □

The following is another useful lemma.

Lemma 5. *Suppose, as in the present context, $\mathcal{Z} = \mathcal{X} \times A$, where $|A| < \infty$. Further, let \mathcal{F} be a real-valued function class defined on \mathcal{Z} . For each $a \in A$, define the following function class on*

⁹This is the set of all probability measures on \mathcal{F} .

\mathcal{X} :

$$\mathcal{F}_a = \{f(\cdot, a) \mid f \in \mathcal{F}\}.$$

Then, for all $z^n = (x_i, a_i)_{i=1}^n \in \mathcal{Z}^n$,

$$\mathcal{R}_{z^n}(\mathcal{F}) \leq \sum_{a \in A} \frac{n_a}{n} \mathcal{R}_{x^{n_a}}(\mathcal{F}_a),$$

where $x^{n_a} = (x_i)_{i \in \{i \mid a_i = a\}}$ and $n_a = |\{i \mid a_i = a\}|$.

Proof. Let $z^n \in \mathcal{Z}^n$. Define $N_a = \{i \mid a_i = a\}$ and then $n_a = |N_a|$. Then,

$$\begin{aligned} \mathcal{R}_{z^n}(\mathcal{F}) &= \mathbb{E}_\eta \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \eta_i f(x_i, a_i) \right] \\ &= \mathbb{E}_\eta \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{a \in A} \sum_{i \in N_a} \eta_i f(x_i, a) \right] \\ &\leq \mathbb{E}_\eta \left[\sum_{a \in A} \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i \in N_a} \eta_i f(x_i, a) \right] \\ &= \sum_{a \in A} \mathbb{E}_\eta \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i \in N_a} \eta_i f(x_i, a) \right] \\ &= \sum_{a \in A} \frac{n_a}{n} \mathbb{E}_\eta \left[\sup_{f \in \mathcal{F}} \frac{1}{n_a} \sum_{i \in N_a} \eta_i f(x_i, a) \right] \\ &= \sum_{a \in A} \frac{n_a}{n} \mathcal{R}_{x^{n_a}}(\mathcal{F}_a). \end{aligned}$$

□

6.4 Proofs from Section 5

Proposition 5 follows straightforwardly from the following result on the VC dimension of neural networks.

Proposition 8. (*Anthony and Bartlett [2009]*) Let $\Theta \subseteq \mathbb{R}^p$. Suppose $\mathcal{F} = \{f(w, \theta) : \theta \in \Theta\}$ is a parametrized class of 0-1 valued functions defined on a set $\mathcal{W} \subseteq \mathbb{R}^k$. Further, suppose, for each $f \in \mathcal{F}$ and each $(w, \theta) \in \Theta$, computing the value of $f(w, \theta)$ takes no more than t many operations from the following

1. The arithmetic operations $+$, $-$, \times and \setminus defined on real numbers.

2. Pairwise comparisons of two real numbers involving the relations $\geq, >, \leq, <$ and $=, \neq$.

3. Output 0 or 1.

Then, the VC dimension of \mathcal{F} is at most $4p(t+2)$.

We now prove Proposition 6.

Proof. We will first show that if we have a parametrised function class $\mathcal{F} = \{f(z, \theta) \mid \theta \in \Theta\}$, where Θ is a compact metric space and $f(z, \cdot)$ is continuous in θ for z , then \mathcal{F} contains no Boolean σ -independent sequence for any α, β . Suppose not. Then, there exists a sequence of parameters $\{\theta_n\}_{n=1}^{\infty} \subseteq \Theta$ such that $\{f(\cdot, \theta_n)\}_n$ is a Boolean σ -independent sequence in \mathcal{F} for some $0 < \alpha < \beta$. Since Θ is compact, there exists a convergent subsequence $\{\theta_{n_k}\}_{k=1}^{\infty}$. It then follows that $\{f(\cdot, \theta_{n_k})\}_k$ is also a Boolean σ -independent sequence at α, β . Now consider the sets $K_O = \{k \mid k \text{ odd}\}$ and $K_E = \{k \mid k \text{ even}\}$. The following set is non-empty.

$$\bigcap_{k \in K_O} \{z \mid f(z, \theta_{n_k}) < \alpha\} \cap \bigcap_{k \in K_E} \{z \mid f(z, \theta_{n_k}) > \beta\} \neq \emptyset.$$

Consider an element z in the above intersection. Further, let $\theta^* \in \Theta$ be the limit of $\{\theta_{n_k}\}_{k=1}^{\infty}$. Then, $\lim_{k \rightarrow \infty} f(z, \theta_{n_k}) = f(z, \theta^*)$. However, from the above intersection, it follows that $f(z, \theta_{n_k}) < \alpha$ for k odd and $f(z, \theta_{n_k}) > \beta$ for k even. Hence, we have two subsequences $\{f(z, \theta_{n_k})\}_{k \in K_O}$ and $\{f(z, \theta_{n_k})\}_{k \in K_E}$ with different limits. But this is a contradiction as all subsequences of a convergent sequence must have the same limit.

We now show the pointwise learnability of continuous preferences. Note that $\mathcal{P}_{CONT} = \{\succ \subseteq X \times X \mid \succ \text{ closed in } X \times X\}$. With the the topology of closed convergence on \mathcal{P}_{CONT} , which is also the topology induced by the Hausdorff metric on \mathcal{P}_{CONT} (see, for example, [Aliprantis and Border \[2006\]](#)), it follows that \mathcal{P}_{CONT} is compact. Hence, $\Delta(\mathcal{P}_{CONT})$ is compact in the weak convergence topology which can metrized by the Prokhorov metric. Since \mathcal{M} is now a closed subset of a compact metric space, it also compact. Now, we will define $\Theta := \mathcal{M}$.

Consider the Quadratic scoring rule S^{br} and the stochastic choice model $\Sigma^{\mathcal{P}_{CONT}}(\mathcal{M})$. Then, for any $f \in S^{br} \circ \Sigma^{\mathcal{P}_{CONT}}(\mathcal{M})$, we have

$$f(x, a; \mu) = 2\sigma_a(x; \mu) - \sum_{b \in A} \sigma_b(x; \mu)^2.$$

Now we need only show that if $\mu_n \rightarrow \mu$, then $f(x, a; \mu_n) \rightarrow f(x, a; \mu)$. This follows from Lemma 2.

□