

A PREQUENTIAL TEST FOR EXCHANGEABLE THEORIES

ALVARO SANDRONI AND ERAN SHMAYA

ABSTRACT. We construct a prequential test of probabilistic forecasts that does not reject correct forecasts when the data-generating processes is exchangeable and is not manipulable by a false forecaster.

1. INTRODUCTION

A potential expert, named Bob, claims that he is informed about which stochastic process will generate the future data. A tester, named Alice, does not know whether Bob does, in fact, know anything relevant about the process (i.e., something Alice does not know). However, Bob will repeatedly announce a forecast of next period's outcomes. A basic question is whether Alice can eventually determine whether Bob is informed. A growing literature answers this question in the negative (see Olszewski (2011) for a literature review).

A prequential test takes, as input, Bob's forecasts and realized outcomes and returns, as output, a reject or pass verdict. Sandroni (2003) shows that if a perfectly informed expert (i.e., predicts according to the data generating process) is likely to pass Alice's test then a completely uninformed expert (i.e., knows nothing about the data generating process) can strategically concoct forecasts that are likely to pass the test, regardless of the realized outcomes. Thus, Alice's test cannot screen informed and uninformed, but strategic, experts. Olszewski and Sandroni (2008) extended this result to finite, possibly unbounded horizon.

There are no clear benefits of running a test if it is known, at the outset, that it will return a pass verdict. So, the results in this literature cast doubt on the use of empirical tests when experts are potentially strategic. However, two key assumptions underlying the impossibility

2000 *Mathematics Subject Classification*. Primary: 62A01. Secondary: 91A40.

results should be examined. First, at the outset, Alice knows nothing about how the data will be produced. Second, the data available to her is finite.

The first condition (Alice's complete ignorance) imposes no restrictions on the types of stochastic processes that may generate the data. So, the results hold for tests that are unlikely to reject *any* data generating process. In contrast, consider the case in which Bob and Alice assume that the data will be produced within a class of data generating processes, such as exchangeable processes. Then Alice may look for tests that do not reject any exchangeable process and she may require Bob to forecast consistently with some exchangeable process. This allows for empirical testing that relies on a priori assumptions on how the data might be produced.

The second condition (finite data sets) is typically justified based on the practical observation that the data cannot feasibly be infinite. However, the impossibility results in this literature are conceptual: they show the existence of forecasting schemes that can manipulate Alice's test, but they do not determine how to construct such manipulating schemes in practice. The assumption of finite data sets may be significant because the forecasting schemes that manipulate the test may involve complex processes that require large, potentially infinite, streams of data to be effectively tested.

It is tempting to speculate that Alice's ignorance, combined with finite data sets, is central to the impossibility of screening informed and uninformed experts. An ignorant, but strategic, expert may be able to exploit the tester's ignorance, combined with her limited data sets, to manipulate her test. However, the impossibility theorems still hold even if any one of these two conditions is disposed:

Assume that it is known, to both Alice and Bob, that the data generating process is an exchangeable processes. Alice does not know which exchangeable process runs the data and Bob claims that he knows the exact process (before any data is observed). Then, Alice has, a priori, a significant understanding of how the data will be produced. She knows that the data is produced by some exchangeable process and her empirical test may rely on this condition. Yet, with finite data the impossibility results still hold: Any prequential test that

is likely to pass an informed expert can be manipulated by an uninformed expert, and so is equally likely to pass an uninformed expert even if he is restricted to produce forecasts based on an exchangeable process (Proposition 1 below).

Now consider the case of infinite data sets, but with no restrictions on the types of stochastic processes that may generate the data. Shmaya (2008) showed that the impossibility results still hold: Any prequential test that is likely to pass an informed expert is also likely to pass an uninformed, but strategic, expert (Proposition 2 below).

Thus, neither a restriction to exchangeable processes, nor infinite sequences of forecasts and outcomes, can, by itself, rule out the impossibility theorems. The question addressed in this paper is whether the impossibility theorem still hold if it is assumed *simultaneously* that the data generating process is exchangeable and that the tester can obtain large, perhaps infinite, sequences of forecasts and outcomes. Under these conditions, we show that the impossibility theorems do not hold. We construct a prequential test that is likely to pass an informed expert and cannot be manipulated by an uninformed expert; no matter which forecasting scheme he uses, he cannot be assured that he is likely to pass the test.

Here is the basic idea of our test: By De-Finetti's Theorem, every exchangeable process has a representation as a sequence of i.i.d. coin tosses with an unknown parameter q . By the law of large numbers, the infinite realization pins down the parameter q . In addition, we show that after observing Bob's forecasts along the realization Alice can deduce Bob's prior distribution on the parameter. So, the sequential framework reduces to a one-shot framework in which Bob delivers a prior that (he claims) generates a parameter q that is observed by Alice. The existence of non-manipulable test can now be obtained by adapting a result in Olszewski and Sandroni (2009b) to this one-shot framework.

Our argument reveals a non-intuitive distinction between the infinite and finite horizon setups. With infinite streams of forecasts and outcomes, it is possible to determine q and, hence, to screen informed and uninformed experts. With finite (perhaps arbitrarily long) streams of forecasts and outcomes, it is possible to make arbitrarily accurate statistical inferences about q and yet this is insufficient to screen informed and uninformed experts.

2. SETUP

Let S be a finite set of *outcomes*. Every period $n = 0, 1, \dots$ an outcome is realized. Let $\Omega = S^{\mathbb{N}}$ be the set of *realizations* equipped with the product topology. A *theory* is a probability measure over Ω , representing the distribution of some stochastic process. The set $\Delta(\Omega)$ of theories is a convex and compact metrizable subspace of the topological vector space of all signed measures equipped with the weak-* topology. A *paradigm* is a set of theories.

Let $S^{<\mathbb{N}} = \bigcup_{n \geq 0} S^n$ be the set of *partial realizations* including the empty sequence e . For every partial realization $\sigma = (s_0, \dots, s_{n-1})$ let $N(\sigma) \subseteq S^{\mathbb{N}}$ be the set of realizations $x \in S^{\mathbb{N}}$ that extends σ , i.e. such that σ is the initial segment of x of length n . In particular $N(e) = S^{\mathbb{N}}$.

Definition 1. Let μ be a theory and $x \in S^{\mathbb{N}} = (s_0, s_1, \dots)$ a realization. The *forecast of μ over x at period n* is the element $p_n \in \Delta(S)$ that is given by

$$(1) \quad p_n[s] = \mu(s | s_0, \dots, s_{n-1}) = \frac{\mu(N(s_0, \dots, s_{n-1}, s))}{\mu(N(s_0, \dots, s_{n-1}))}$$

for every outcome $s \in S$. That is, the forecast at period n is the forecast made at period $n-1$ (based on all information available at period $n-1$) on the outcome at period n . Here and later, when $n = 0$ the sequence (s_0, \dots, s_{n-1}) refers to the empty sequence e . The forecast p_n is undefined if $\mu(N(s_0, \dots, s_{n-1})) = 0$.

A *prequential test* is given by a Borel function $T : (\Delta(S) \times S)^{\mathbb{N}} \rightarrow \{\text{FAIL}, \text{PASS}\}$: Bob passes the test if $T(p_0, s_0, p_1, s_1, \dots) = \text{PASS}$, where, for every $k \in \mathbb{N}$, p_k is the forecast on period k according to Bob's theory and s_k is the outcome of period k . This definition of test is based on Dawid's (1984) *prequential principle* that assessment of a theory μ should depend on μ only through the sequence of forecasts that μ made over the realized sequence of outcomes, and not on forecasts conditional on data that has not been realized. More general tests allow the tester (Alice) to use forecasts of Bob's theory after any finite history, including the histories that did not occur. This paper focuses on prequential tests.

The test is *finite* if, for some $n \in \mathbb{N}$, T depends only on the forecasts and outcomes of the first n periods.

Definition 2. Fix a prequential test T and let μ be a theory.

- (1) μ passes the test over a realization $x = (s_0, s_1, \dots) \in S^{\mathbb{N}}$ if $\mu(N(s_0, \dots, s_{n-1})) > 0$ for every n and $T(p_0, s_0, p_1, s_1, \dots) = \text{PASS}$, where p_k is the forecast of μ over x at period k .
- (2) The *acceptance set of μ under T* , denoted $A_\mu^{(T)}$ is the set of all realizations $x \in S^{\mathbb{N}}$ such that μ passes the test over x .

Definition 3. The prequential test T *accepts the data generating process in a paradigm Γ with probability $1 - \epsilon$* if $\mu(A_\mu^{(T)}) \geq 1 - \epsilon$ for every $\mu \in \Gamma$.

If the test accepts the data generating process with high probability and Bob's theory μ is the data generating process then Bob is assured that he will pass the test with high probability.

Definition 4. Let Γ be a paradigm. A test T is ϵ -manipulable w.r.t Γ if there exists some distribution ζ over Γ such that

$$\zeta \otimes \mu(\{(\nu, x) : x \in A_\nu^{(T)}\}) > 1 - \epsilon$$

for every theory $\mu \in \Gamma$.

If a test is ϵ -manipulable then Bob can randomize a theory according to ζ and be sure that, with high probability, he passes the test regardless of the data generating process μ .

If a test is not manipulable then every strategy to fake theories fails on at least one data-generating process. Following Dekel and Feinberg (2006) and Olszewski and Sandroni (2009b), we use in this paper a stronger condition, which entails that every strategy employed by an ignorant expert will fail on a topologically large set of data generating processes. Recall that a set M of a separable metric space Z is co-meager if it contains a countable intersection

of dense open sets of Z . We say that a property is satisfied by co-meager many elements in Z if the set of elements that satisfy this property is co-meager.

Definition 5. Let Γ be a paradigm. A test T is *strongly non-manipulable w.r.t* Γ if, for every distribution ζ over Γ ,

$$(2) \quad \zeta \otimes \mu \left(\{(\nu, x) : x \in A_\nu^{(T)}\} \right) = 0$$

for co-meager many μ -s in Γ .

Note that in Definition 5 we require that the set of μ on which the expert fail be co-meager relative to the topological space Γ . We do not assume that Γ itself is a co-meager subset of $\Delta(\Omega)$. Indeed, most paradigms used in modeling, i.e., most interesting classes of stochastic processes (stationary processes, markov processes, mixing processes) are meager as subsets of $\Delta(\Omega)$.

3. RESULTS

3.1. Previous results. The following result extends an argument of Sandroni (2003) to convex and compact paradigms. The proof relies on Fan's Minimax Theorem. See also Vovk, V. and G. Shafer (2005) for a related result. For completeness, we reproduce the proof in the appendix.

Proposition 1. *Fix $\epsilon > 0$ and let Γ be a convex and compact paradigm. Let T be a finite test. If T accepts the data generating process in the paradigm Γ with probability $1 - \epsilon$ then T is $(\epsilon + \delta)$ -manipulable w.r.t Γ for every $\delta > 0$.*

The following result was proved by Shmaya (2008). The proof relies on Martin's Theorem about determinacy of Blackwell games (Martin (1998)).

Proposition 2. *Fix $\epsilon > 0$ and let $\Gamma = \Delta(\Omega)$ be the paradigm of all theories. Let T be a prequential test. If T accepts the data generating process in the paradigm Γ with probability $1 - \epsilon$ then T is $(\epsilon + \delta)$ -manipulable w.r.t Γ for every $\delta > 0$.*

3.2. New result. Our new result is that, for the compact and convex paradigm of exchangeable processes, there exists an infinite prequential test which is non-manipulable. Therefore, Proposition 2 about manipulability of infinite tests does not hold for convex and compact paradigms. In contrast, Proposition 1 about manipulability of finite tests holds for every compact and convex paradigm.

For the rest of the paper we fix the set of outcome $S = \{0, 1\}$. Theorem 1 below and its proof applies with minor changes to the case of arbitrary finite set of outcomes, however since we are mainly interested in a counter-example the restriction to a two-elements outcome set is worth the notational convenience it provides. We identify the set forecast $p \in \Delta(S)$ with $p[1] \in [0, 1]$.

Before stating our result, we give two simple examples of paradigms for which Proposition 2 does not hold. In Example 1 the paradigm is compact but not convex and in Example 2 the paradigm is convex but not compact.

Example 1. Let Γ be the paradigm of all i.i.d distributions. Let T be the test that passes the expert if and only if

$$(3) \quad p_n = \limsup_{k \rightarrow \infty} (s_0 + \dots + s_{k-1})/k$$

for every $n \geq 0$ (so that the expert passes if he makes the same prediction in all days and this prediction matches the empirical frequency of outcomes). Then T accepts the data-generating process in Γ with probability 1 and is strongly non-manipulable w.r.t. Γ .

Example 2. Let Γ be the convex hull of the paradigm of all i.i.d. distributions. Let T be the test that passes the expert if and only if (3) is satisfied for all but finitely many n -s. Then T accepts the data-generating process in Γ with probability 1 and is strongly non-manipulable w.r.t. Γ .

A theory $\mu \in \Delta(S^{\mathbb{N}})$ is *exchangeable* if $\mu(N(s_0, \dots, s_{n-1})) = \mu(N(s_{\pi(0)}, \dots, s_{\pi(n-1)}))$ for every partial realization $(s_0, s_1, \dots, s_{n-1}) \in S^{<\mathbb{N}}$ and every permutation π over $\{0, 1, \dots, n-1\}$

1}. Clearly, the paradigm Γ of exchangeable processes is compact and convex. Moreover, it follows from De-Finetti's Theorem (see below) that Γ is the closed convex hull of the paradigm of all i.i.d. distributions.

Theorem 1. *Let Γ be the paradigm of exchangeable processes. Then there exists a prequential test T such that T accepts the data-generating process in Γ with probability 1 and T is strongly non-manipulable w.r.t Γ .*

3.3. Additional comments.

Non-prequential tests. Proposition 2 relies on Dawid's prequential principle. If one abandons this principle, allowing Alice to use forecasts of Bob's theory after all finite histories, including those that did not occur, then some non-manipulable tests do exist. In particular, Olszewski and Sandroni (2009b) construct a strongly non-manipulable non-prequential test for the paradigm of all theories. Shmaya (2008) (Section 7) also have non-prequential non-manipulable tests for the paradigm of all theories. In addition, Al-Najjar et al. (2010) construct a (non-prequential) non-manipulable test for the paradigm of theories which are learnable and sufficient for predictions, which is convex but not compact.

Non Borel tests. Another assumption of Proposition 2 is that the test is a Borel function (This assumption is part of our definition of test in Section 2). Shmaya (2008) shows that under the axiom of choice, there exist non-manipulable prequential tests which are not Borel.

Future independent tests. Olszewski and Sandroni (2008) introduce the class of future independent tests. For prequential tests, the future independence property can be seen as weakening of the finite horizon property: instead of requiring that the test terminates at finite time, it only requires that if the expert fail the test then this failure will be demonstrated in finite (possibly unbounded) time. Olszewski and Sandroni (2009a) show that many of their manipulability results for future independent tests can also be proved for convex and compact paradigms. The test we construct is not future independent.

Strong non-manipulability. The definition of strong non-manipulability appeals to a topological notion of ‘large’ set of distribution as a co-meager set. The goal of the definition is to capture the idea that, regardless of the strategy he uses to create forecasts, the uninformed expert will fail on a large set of stochastic processes. This definition has been used in the expert testing literature, but we recognize that such topological notion has well-known drawbacks. In particular, it does not indicate odds of discrediting the uninformed expert. We view the main contribution of our theorem to be the very fact that there exists a test which is not manipulable. The strong non-manipulability property is a bonus.

4. PROOF OF THEOREM 1

4.1. Preliminaries.

De-Finetti’s Theorem. For every $\lambda \in \Delta([0, 1])$ let $\varepsilon_\lambda \in \Delta(\{0, 1\}^\mathbb{N})$ be the distribution of infinite sequence of i.i.d coins with probability q of success, where q is drawn from λ :

$$(4) \quad \varepsilon_\lambda(N(s_0, \dots, s_{n-1})) = \int \prod_{i=0}^{n-1} q^{s_i} (1-q)^{1-s_i} \lambda(dq).$$

Clearly the distribution ε_λ is exchangeable. De-Finetti’s Theorem states that the map $\lambda \in \Delta([0, 1]) \mapsto \varepsilon_\lambda$ is one-to-one and onto the set Γ of exchangeable distributions over $\{0, 1\}^\mathbb{N}$, and its inverse is given by $\mu \in \Gamma \mapsto \bar{\mu} \in \Delta([0, 1])$ where $\bar{\mu} \in \Delta([0, 1])$ is the push-forward of μ under L (i.e., $\bar{\mu}(B) = \mu(L^{-1}(B))$ for every Borel subset B of $[0, 1]$) and $L : \{0, 1\}^\mathbb{N} \rightarrow [0, 1]$ is the limit average

$$(5) \quad L(s_0, s_1, \dots) = \limsup_{n \rightarrow \infty} (s_0 + \dots + s_{n-1})/n.$$

Since the map $\lambda \mapsto \varepsilon_\lambda$ is continuous and its domain $\Delta([0, 1])$ is compact it follows that the maps $\lambda \mapsto \varepsilon_\lambda$ and $\mu \mapsto \bar{\mu}$ are homeomorphisms.

A non-manipulable test. Our proof uses the following proposition.

Proposition 3. *Let Q be a compact metric space and let $\Delta(Q)$ be equipped with the weak-* topology. There exists a Borel function $t : \Delta(Q) \times Q \rightarrow \{\text{FAIL}, \text{PASS}\}$ such that*

(1) For every $\bar{\mu} \in \Delta(Q)$, $\bar{\mu}(\{q \in Q | t(\bar{\mu}, q) = PASS\}) = 1$.

(2) For every $\bar{\zeta} \in \Delta(\Delta(Q))$, one has

$$\bar{\zeta} \otimes \bar{\mu}(\{(\bar{\nu}, q) \in \Delta(Q) \times Q : t(\bar{\nu}, q) = PASS\}) = 0$$

for co-meager many $\bar{\mu} \in \Delta(Q)$.

This proposition has the following nice interpretation: Assume that an expert claims to know the distribution from which an element $q \in Q$ is drawn. Say that the expert passes the test if $t(\bar{\mu}, q) = PASS$, where $\bar{\mu}$ is the distribution proclaimed by the expert and q is the actual realization. The first property says that if the expert provides the correct distribution of q then he passes with probability 1. The second property ensures that, if the expert randomizes his theory according to some distribution $\bar{\zeta}$, then he will fail the test with probability 1 for a large (co-meager) set of truths.

Proposition 3 was proved by Olszewski and Sandroni (2009b). Their theorem is formulated for the underlying space of realization $Q = \{0, 1\}^{\mathbb{N}}$ but their theorem and proof carry through for every compact metric space. We need the case $Q = [0, 1]$.

4.2. Proof. We first prove that when forecasts are made using an exchangeable theory, then the forecasts over a single realization already determines the entire theory.

Lemma 1. *Let $x = (s_0, s_1, \dots) \in \{0, 1\}^{\mathbb{N}}$ be a realization and let $\mu' = \varepsilon_{\lambda'}$ and $\mu'' = \varepsilon_{\lambda''}$ be two exchangeable theories. If $\mu'(1|s_0, \dots, s_{n-1}) = \mu''(1|s_0, \dots, s_{n-1})$ for every $n \in \mathbb{N}$ then $\lambda' = \lambda''$ (and, in particular, $\mu' = \mu''$).*

Proof. Under the assumption of the lemma, it follows by induction from (1) that

$$(6) \quad \mu'(N(s_0, \dots, s_{n-1})) = \mu''(N(s_0, \dots, s_{n-1}))$$

for every $n \in \mathbb{N}$. From (6) and (4) it follows that

$$\int \prod_{i=0}^{n-1} q^{s_i} (1-q)^{1-s_i} \lambda'(dq) = \int \prod_{i=0}^{n-1} q^{s_i} (1-q)^{1-s_i} \lambda''(dq)$$

for every n . Since $\prod_{i=0}^{n-1} q^{s_i}(1-q)^{1-s_i}$ is a polynomial in q of degree n , it follows again by induction that

$$\int q^n \lambda'(dq) = \int q^n \lambda''(dq)$$

for every n . Since a probability measure with bounded support is determined by its moments it follows that $\lambda' = \lambda''$. \square

Corollary 1. *There exists a Borel subset B of $[0, 1]^{\mathbb{N}} \times \{0, 1\}^{\mathbb{N}}$ and a Borel map $f : B \rightarrow \Delta([0, 1])$ such that for every exchangeable theory μ and for every realization $x = (s_0, s_1, \dots)$ one has*

$$(p_0, p_1, \dots, s_0, s_1, \dots) \in B, \text{ and}$$

$$f(p_0, p_1, \dots, s_0, s_1, \dots) = \bar{\mu}.$$

where $p_n = \mu(1|s_0, \dots, s_{n-1})$ and $\bar{\mu}$ is the push-forward of μ under the limit average $L : \{0, 1\}^{\mathbb{N}} \rightarrow [0, 1]$ defined in (5).

Proof. Consider the Borel map $g : \Delta([0, 1]) \times \{0, 1\}^{\mathbb{N}} \rightarrow [0, 1]^{\mathbb{N}} \times \{0, 1\}^{\mathbb{N}}$ that is given by $g(\lambda, s_0, s_1, \dots) = (p_0, p_1, \dots, s_0, s_1, \dots)$ where $p_n = \mu(1|s_0, \dots, s_{n-1})$ for $\mu = \varepsilon_\lambda$. By the previous lemma this map is one-to-one. Therefore its image $B = g(\Delta([0, 1]) \times \{0, 1\}^{\mathbb{N}})$ is a Borel set and the inverse map $g^{-1} : B \rightarrow \Delta([0, 1]) \times \{0, 1\}^{\mathbb{N}}$ is a Borel map (See Kechris (1994), Corollary 15.2). Take $f = \pi \circ g^{-1}$ where $\pi : \Delta([0, 1]) \times \{0, 1\}^{\mathbb{N}} \rightarrow \Delta([0, 1])$ is the projection on the first coordinate. The assertion in the corollary follows from the fact that if μ is exchangeable then $\mu = \varepsilon_\lambda$ when $\lambda = \bar{\mu}$. \square

Proof of Theorem 1. Consider the test $T : [0, 1]^{\mathbb{N}} \times \{0, 1\}^{\mathbb{N}} \rightarrow \{0, 1\}$ which is given by

$$T(p, x) = \begin{cases} t(f(p, x), L(x)), & \text{if } (p, x) \in B \\ \text{FAIL}, & \text{otherwise.} \end{cases}$$

for every realization $x = (s_0, s_1, \dots) \in \{0, 1\}^{\mathbb{N}}$ and every sequence $p = (p_0, p_1, \dots) \in [0, 1]^{\mathbb{N}}$ of forecasts, where $B \subseteq [0, 1]^{\mathbb{N}} \times \{0, 1\}^{\mathbb{N}}$ and $f : B \rightarrow \Delta([0, 1])$ are as in Corollary 1,

$L : \{0, 1\}^{\mathbb{N}} \rightarrow [0, 1]$ is the limit average (5), and t is the test defined in Proposition 3 for the interval $Q = [0, 1]$. From Corollary 1 and the definition of T it follows that

$$(7) \quad \mu \text{ passes } T \text{ over } x \iff t(\bar{\mu}, L(x)) = \text{PASS}$$

for every realization $x \in \{0, 1\}^{\mathbb{N}}$ and every exchangeable distribution μ . We now claim that the desired properties of T follow from the corresponding properties of t . Indeed, for every exchangeable μ

$$\mu(A_{\mu}^{(T)}) = \mu(\{x | t(\bar{\mu}, L(x)) = \text{PASS}\}) = \bar{\mu}(\{q | t(\bar{\mu}, q) = \text{PASS}\}) = 1$$

where the first equality follows from (7), the second from the definition of $\bar{\mu}$ and the third from Proposition 3. Thus T accepts exchangeable data generating processes with probability 1.

Let $\zeta \in \Delta(\Gamma)$ and let $\bar{\zeta} \in \Delta(\Delta([0, 1]))$ be the push-forward of ζ under the map $\nu \mapsto \bar{\nu}$ where $\bar{\nu} \in \Delta([0, 1])$ is the pushforward of ν under L for every $\nu \in \Gamma$. Fix a theory $\mu \in \Gamma$ and let $\bar{\mu}$ be the pushforward of μ under L . Then

$$\begin{aligned} \zeta \otimes \mu(\{(\nu, x) : x \in A_{\nu}^{(T)}\}) &= \\ \zeta \otimes \mu(\{(\nu, x) : t(\bar{\nu}, L(x)) = \text{PASS}\}) &= \bar{\zeta} \otimes \bar{\mu}(\{(\bar{\nu}, q) | t(\bar{\nu}, q) = \text{PASS}\}) \end{aligned}$$

where the first equality follows from (7) and the second from the definitions of $\bar{\zeta}$ and $\bar{\mu}$. Therefore the set of all μ -s such that $\zeta \otimes \mu(\{(\nu, x) : T(\nu, x) = \text{PASS}\}) = 0$ is the preimage of the set of all $\bar{\mu}$ such that $\bar{\zeta} \otimes \bar{\mu}(\{(\bar{\nu}, q) : t(\bar{\nu}, q) = \text{PASS}\}) = 0$ under the map $\mu \mapsto \bar{\mu}$. Since the later set is co-meager by Theorem 3 and the map $\mu \mapsto \bar{\mu}$ is homeomorphism, it follows that the former set is co-meager, as desired. \square

5. TESTING AND MERGING

The expert testing literature asks whether an uninformed expert can produce forecasts that appear to be as good as the forecasts of the informed given the observed outcomes.

A related question is whether an uninformed expert can produce forecasts that are close to the forecasts of the informed expert. In this section we show that the answer to these two questions might be different.

Recall that a distribution $\nu \in \Delta(S^{\mathbb{N}})$ *merges* with a distribution $\mu \in \Delta(S^{\mathbb{N}})$ if

$$\nu(s|s_0, \dots, s_{n-1}) - \mu(s|s_0, \dots, s_{n-1}) \xrightarrow[n \rightarrow \infty]{} 0$$

for μ -almost every realization (s_0, s_1, s_2, \dots) . This notion of merging was introduced by Kalai and Lehrer (1994) as a weakening of Blackwell and Dubins merging (1962). Sorin (1999) relates merging with several game theoretic models. Kalai et al. (1999) relate merging with certain classes of tests.

Say that a paradigm Γ is *learnable* if there exists $\nu \in \Gamma$ such that ν merges with every $\mu \in \Gamma$. If Γ is learnable then an uninformed expert can learn to predict, i.e., to provide forecasts that becomes arbitrary close to the forecasts of the informed expert.

The paradigm of all exchangeable theories is learnable: If $\Omega = \{0, 1\}$ then the theory $\nu = \varepsilon_\lambda$ given in (4) where λ is the uniform distribution on $[0, 1]$ merges with all exchangeable theories. However, by our Theorem 1, there exists a test that screens an uninformed expert from the informed expert based on the outcomes of the process.

The paradigm of all theories is not learnable, so that an uninformed expert cannot produce forecasts that are close to the informed expert's forecasts for every process, but, by Proposition 2, he can do as good as the informed expert in every prequential test.

APPENDIX A. PROOF OF PROPOSITION 1

Let $F : \Delta(\Gamma) \otimes \Gamma \rightarrow [0, 1]$ be given by $F(\zeta, \mu) = \zeta \otimes \mu \left(\left\{ (\nu, x) \in \Gamma \times \Omega : x \in A_\nu^{(T)} \right\} \right)$. Then F is bilinear and, since T is finite, continuous in its second argument. It follows from Fan's Minimax Theorem (1953) that

$$(8) \quad \sup_{\zeta} \min_{\mu} F(\zeta, \mu) = \min_{\mu} \sup_{\zeta} F(\zeta, \mu),$$

where the suprema are over all $\zeta \in \Delta(\Gamma)$ and the minima are over all $\mu \in \Gamma$.

For every $\mu \in \Gamma$, if ζ is Dirac atomic measure over μ then $F(\zeta, \mu) \geq 1 - \epsilon$ since T accepts data generating process in Γ with probability $1 - \epsilon$. Therefore it follows from (8) that for every $\delta > 0$ there exists $\zeta \in \Delta(\Gamma)$ such that $\min_{\mu} F(\zeta, \mu) > 1 - (\epsilon + \delta)$, i.e., that T is $(\epsilon + \delta)$ -manipulable.

REFERENCES

- [1] N. Al-Najjar, R. Smorodinsky, A. Sandroni, and J. Weinstein, Testing Theories with Learnable and Predictive Representations *Journal of Economic Theory*, **145** (2010).
- [2] D. Blackwell and L. Dubins, Merging of opinions with increasing information, *Annals of Mathematical Statistics*, **33** (1962).
- [3] E. Dekel and Y. Feinberg Non-Bayesian Testing of a Stochastic Prediction, *Review of Economic Studies*, **73** (2006).
- [4] A. P. Dawid, Statistical Theory: The Prequential Approach, *J.R.Statist. Soc. A*, **147** (1984).
- [5] K. Fan, Minimax Theorems, *Proceedings of the National Academy of Science U.S.A.*, **39**, 42-47.
- [6] Kechris, A, *Classical Descriptive Set Theory*, Springer Verlag (1994).
- [7] E. Kalai and E. Lehrer, Weak and strong merging of opinions, *J. Math. Econom.*, **23** (1994).
- [8] E. Kalai, E. Lehrer and R. Smorodinsky, Calibrated forecasting and merging, *Games Econom. Behav.*, **29** (1999).
- [9] D. Martin, The Determinacy of Blackwell Games *Journal of Symbolic Logic.*, **63** (1998).
- [10] W. Olszewski Calibration and Expert Testing, *Handbook of Game Theory, volume IV* (2011)
- [11] W. Olszewski and A. Sandroni, Manipulability of Future-Independent Tests *Econometrica*, **76** (2008).
- [12] W. Olszewski, and A. Sandroni Strategic Manipulation of Empirical Tests, *Mathematics of Operations Research*, **34** (2009a).
- [13] W. Olszewski, and A. Sandroni A Nonmanipulable Test, *Annals of Statistics*, **37** (2009b).
- [14] A. Sandroni, The Reproducible Properties of Correct Forecasts, *International Journal of Game Theory*, **32** (2003).
- [15] E. Shmaya, Many inspections are manipulable, *Theoretical Economics*, **3** (2008).
- [16] S. Sorin, Merging, reputation, and repeated games with incomplete information, *Games Econom. Behav.*, **29** (1999).
- [17] V. Vovk and G. Shafer, Good Randomized Sequential Probability Forecasting is Always Possible, *Journal of the Royal Statistical Society Series B*, **67** (2005).

E-mail address: sandroni@kellogg.northwestern.edu

KELLOGG SCHOOL OF MANAGEMENT, NORTHWESTERN UNIVERSITY

E-mail address: e-shmaya@kellogg.northwestern.edu

KELLOGG SCHOOL OF MANAGEMENT, NORTHWESTERN UNIVERSITY