

YOU WON'T HARM ME IF YOU FOOL ME

FEDERICO ECHENIQUE AND ERAN SHMAYA

ABSTRACT. A decision maker faces a new theory of how certain events unfold over time. The theory matters for choices she needs to make, but possibly the theory is a fabrication. We show that there is a test which is guaranteed to pass a true theory, and which is also *conservative*: A false theory will only pass when adopting it over the decision maker's initial theory would not cause substantial harm; if the agent is fooled she will not be harmed.

We also study a society of conservative decision makers with different initial theories. We uncover pathological instances of our test: a society collectively rejects most theories, be they true or false. But we also find well-behaved instances of our test, collectively accepting true theories and rejecting false. Our tests builds on tests studied in the literature in the context of non-strategic inspectors.

1. INTRODUCTION

An agent makes decisions about events that unfold over time: to fix ideas, suppose she is trading in the stock market. The agent has a pre-existing belief about stock-market prices, beliefs which would lead her to adopt a certain contingent plan of actions (trades). Assume that the agent learns about a new theory of stock market prices. If she believes

Date: December, 2007.

Division of the Humanities and Social Sciences, California Institute of Technology.
Emails: fede@caltech.edu and eshmaya@caltech.edu.

the theory, she should adopt a different contingent plan. The agent faces a problem, should she adopt the theory or stick to her original beliefs?

We study agents who are *conservative*, in the sense that they are inclined to distrust the new theory, and follow their pre-existing beliefs. We show that, conservatism notwithstanding, agents can follow the new theory if they at the same time test it. When the theory is correct it is guaranteed to pass the test, and the agent reaps the benefits of following the optimal action. When the theory is incorrect, it will only pass the test when the resulting actions are not too suboptimal under the agent's pre-existing beliefs. We say "not too suboptimal" here to mean that in the limit, for arbitrarily patient agents, there is no loss in following an incorrect theory.

Thus the agent can be true to her conservative inclination, but follow the new theory. Under her pre-existing beliefs, she does not lose too much when the theory passes the test, even when the agent is *fooled*, i.e. when an incorrect theory passed the test. At the same time, she avoids the standard cost of conservatism—the cost of rejecting correct new theories. When the theory fails the test the agent knows that it is false. She can, for example, demand a restitution, or not trust the source of the theory in the future.

Our test has implications for a society of conservative agents with different initial beliefs. While each individual agent is assured that a true theory will pass her test with probability 1, it may be that for any given realization of stock market prices, a true theory would

fail the tests of (in a sense) *most* agents. In fact, for some instances of conservative tests, most theories—whether true or false—fail most agents’ tests at all outcomes.

We then present a conservative test which does not have this problem: A true theory will almost surely pass the tests of most agents. Moreover, There is an instance of conservative test which, in addition, is not manipulable: If the theory is a fabrication of a false expert then such a manipulation will not be successful.

Testing an expert. Our paper is closely related to the literature on testing experts. The expert presents a theory, and claims it is the true probability law governing some events that unfold over time; e.g. the weather or the stock market. The expert may be true, and report the actual law, or false, and be completely ignorant about the actual law. One would like to test the expert’s report using the observed data. In certain environments, one cannot design a test which (a) the true expert is guaranteed to pass and (b) the false expert cannot manipulate (Foster and Vohra, 1998; Lehrer, 2001; Sandroni, 2003; Olszewski and Sandroni, 2007b; Shmaya, 2007).

We depart from the literature by modeling why one would want to test the expert: the expert’s report matters for an agent’s decisions. After seeing the report, the agent has to choose a contingent plan of actions. If she believes the expert, a certain contingent plan is optimal. If she does not believe him, an alternative plan, based on some given beliefs about the outcome, is optimal.

We show that there is a test such that (a) the true expert is guaranteed to pass and (b') the false expert can only pass when following his recommendation would not have lead to a significantly worse decision than ignoring it. So, while our test is manipulable in the sense that a false expert can easily pass it, the agent, if fooled into passing a false expert, is not harmed.

In addition, we consider an expert reporting to a society of agents (a guru). We show that there are tests satisfying (a) and (b') at the level of individual agents, but with very different implications for the society. One test fails most experts in the aggregate, whether they are true or not. Another test passes true experts in the aggregate, while being resistant to manipulation: A false guru can fool each agent separately, but he cannot fool a large set of agents.

Four papers are closely related to ours. First, Al-Najjar and Weinstein (2006) and Feinberg and Stewart (2007) study testing of multiple simultaneous experts. One can think of our agent's pre-existing beliefs as a competing expert's theory. Al-Najjar and Weinstein assume that one of the experts is true, and show that there is a test that only fails to single out the true expert when the different experts' predictions become close over time. This conclusion is similar to ours, in the sense that an agent following the recommendations of two experts who pass Al-Najjar and Weinstein's test would eventually make similar decisions. In fact, we show in Section 4 that one instance of our test is formally similar to Al-Najjar and Weinstein's. One difference between the two

is that, in our context, Al-Najjar and Weinstein’s test is not guaranteed to pass a true expert. This issue is irrelevant in their model, as they want tests which single out a true expert when one is known to be present. Feinberg and Stewart consider a cross-calibration test, which compares the empirical frequencies of events conditioned on the joint predictions made by the two experts. They show that the test always passes the true expert, and fails a false expert against a large set of true experts. We derive a similar result in the context of a false guru – the guru fails the tests of a large set of decision makers. The main difference is that in our framework the guru does not compete against a true expert.

Finally, our tests in Section 4 build directly on the tests in Dekel and Feinberg (2006) and Olszewski and Sandroni (2007a); we show how one can combine the idea in our conservative tests with their tests to obtain tests with certain desirable properties in the society.

2. MODEL AND NOTATION

The primitives of our model are (Z, A, r, λ, π) :

- Z is a finite or countable set of *outcomes* equipped with the discrete topology; let $\Omega = Z^{\mathbb{N}}$ be equipped with the product topology and the induced sigma algebra \mathcal{B} of Borel sets.
- A is a finite set of *actions*.
- $r : Z \times A \rightarrow [0, 1]$ is a *payoff function*.
- $\lambda \in (0, 1)$ is a *discount factor*.

- π is a probability measure over Ω ; π represents given beliefs about Ω .

At every period n a decision maker (DM) chooses an action $a_n \in A$ and receives payoff $r(z_n, a_n)$, where z_n is the outcome of that period. The resulting discounted payoffs are

$$(1 - \lambda) \sum_{n=0}^{\infty} \lambda^n r(z_n, a_n).$$

Let $Z^{<\mathbb{N}} = \bigcup_{n \in \mathbb{N}} Z^n$ be the set of finite sequences of elements of Z , including the empty sequence e . For $\omega = (z_0, z_1, \dots) \in \Omega$ and $n \in \mathbb{N}$, let $\omega|_n = (z_0, \dots, z_{n-1})$ be the initial segment of ω of length n . In particular $\omega|_0 = e$. For $s \in Z^{<\mathbb{N}}$ and $\omega \in \Omega$ we write $s \subseteq \omega$ if $s = \omega|_n$ for some n . For $s \in Z^{<\mathbb{N}}$ let $N_s = \{\omega \in \Omega | s \subseteq \omega\}$.

A (pure) strategy is given by $f : Z^{<\mathbb{N}} \rightarrow A$: $f(z_0, \dots, z_{n-1})$ is the action taken by the DM after observing (z_0, \dots, z_{n-1}) .

For $\omega = (z_0, z_1, \dots) \in \Omega$ and a pure strategy $f : Z^{<\mathbb{N}} \rightarrow A$ let

$$R_\lambda(\omega, f) = (1 - \lambda) \sum_{n \in \mathbb{N}} \lambda^n r(z_n, f(z_0, \dots, z_{n-1}))$$

be the discounted payoff to the DM who uses strategy f when the realization is ω . Say that f is ν -optimal iff

$$f \in \operatorname{argmax} \int R_\lambda(\omega, g) \nu(d\omega),$$

where the maximization above is over strategies g .

We denote the set of all probability measures over Ω by $\Delta(\Omega)$, and endow it with the weak* topology.

A probability measure $\nu \in \Delta(\Omega)$ is a *theory*. A *test function* is a function $T : \Delta(\Omega) \rightarrow 2^\Omega$: a theory ν is accepted if the observed realization of outcomes is in $T(\nu)$. A test function T is *type-I error free* if $\nu(T(\nu)) = 1$ for every $\nu \in \Delta(\Omega)$.

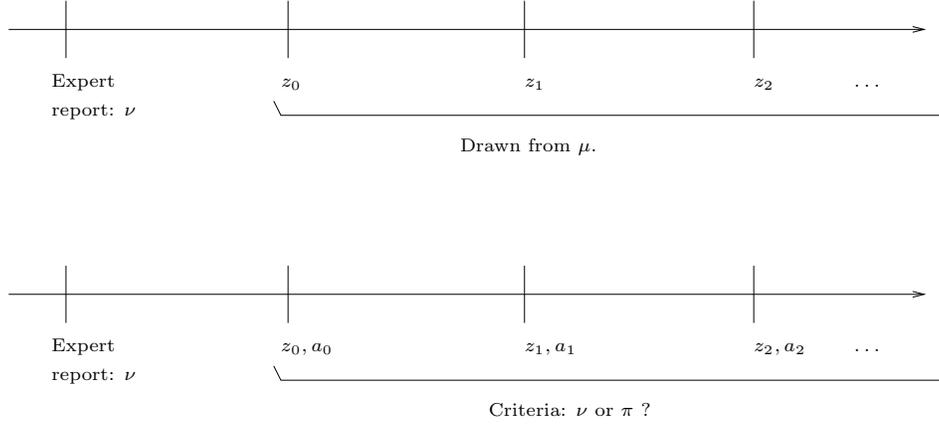
Recall that a set is *meager* if it is contained in an F_σ set with empty interior, and that by the Baire category theorem, a countable union of meager sets is meager. A set is *residual* if its complement is meager. We interpret meager sets as being small, and residual sets as large.

3. INDIVIDUAL CONSERVATISM

We start with an interpretation of conservatism, then present our results.

Consider the first diagram in the figure below. Suppose the outcome of interest is a set Ω of infinite sequences of stock market prices, z_0, z_1, \dots , drawn from a probability law μ . Before any z_n has been realized, an expert claims that the probability law ν governs the realization of prices over time; the expert's report may or may not coincide with μ . A test for the expert prediction is a set of outcomes z_0, z_1, \dots for which one decides that the expert has reported the true law.

We study the situation in the second diagram: A decision maker (DM) sees the expert prediction ν and has to choose, at each n an action a_n . DM has payoffs that depend on the sequences of outcomes, (z_n) and of actions (a_n) . So she cares about the expert's report because it matters for the decision she has to make. If she believes the report is true, she should base her decisions on ν . If she believes it is false, she has some (given) pre-existing beliefs π about the outcomes.



DM is ignorant about μ so she has two criteria on which to base her choices: ν or π . A contingent plan f_ν is optimal if ν were true, and f_π is optimal if DM rejects ν and uses π . Evidently, f_ν is better than f_π if DM knew the expert's report to be true, but f_π is better if she knew it to be false.

We show that there is a test such that, on outcomes for which ν passes the test, f_ν is at least as good as f_π under *either* of the two criteria DM might use.

One way to think of the result is that DM is conservative. She realizes that ν might be false, and is concerned about outcomes where f_ν leads to very different payoffs than f_π . She might want to reject ν on such outcomes. Our test assures DM that she can satisfy this conservative inclination, while getting the benefit of higher payoffs from following the true expert when ν is true.

Formally, we work with the following notion of conservatism. Let $\pi \in \Delta(\Omega)$. A test function T is π -conservative if

$$(1) \quad \lim_{\lambda \rightarrow 1} \int_{T(\nu)} (R_\lambda(\omega, g) - R_\lambda(\omega, f)) \pi(d\omega) \leq 0$$

for every $\nu \in \Delta(\Omega)$ and every ν -optimal strategy f and π -optimal strategy g .

Theorem 1. *Let $\pi \in \Delta(\Omega)$. There exists a π -conservative and type-I error free test function T .*

The proof of Theorem 1 is in Section 5, as are all proofs in the paper. The proof relies on finding a set $A \subseteq \Omega$ such that for π -almost all outcomes on A , the posterior distributions of ν and π converge as n grows, and at the same time $\nu(A) = 1$. Accept ν on A ; i.e. set $T(\nu) = A$. We use Lebesgue's decomposition theorem and Blackwell-Dubins merging theorem to show the existence of A .

4. IMPLICATIONS OF CONSERVATISM FOR THE SOCIETY.

In this section, we consider a society of decision makers, each has her own initial belief, and each uses a conservative test. We identify the decision maker with her initial belief, so the set of decision makers is given by $\Delta(\Omega)$.

We show that the collective response of the society to new theories can depend on the choice of individual tests. Concretely, we present three instances of conservative test. The first instance results in a society which accepts an expert only when the realization is an atom of his forecast: Note that at the individual level, a true theory is accepted with probability one. But in this instance of the test, the true expert is rejected, save for exceptional circumstances. Our second instance is a society which always accepts a true expert with probability 1, so the individual property of being type-I error free aggregates. Our third

instance satisfies all the desirable properties of the second instance and, in addition, cannot be manipulated.

The tests in Instance 2 result from combining the ideas in our conservative tests with the test in Dekel and Feinberg (2006), while the tests in Instance 3 use the test in Olszewski and Sandroni (2007a).

We assume that every DM $\pi \in \Delta(\Omega)$ has a test $T_\pi : \Delta(\Omega) \rightarrow 2^\Omega$. To emphasize the dependence on π we sometimes use the term *individual test* for T_π . A collection $\{T_\pi\}_{\pi \in \Delta(\Omega)}$ of individual tests is called a *collective test*.

Definition. Let $\nu \in \Delta(\Omega)$ be a theory and $\omega \in \Omega$ a realization. Then ν is *collectively accepted over ω* if the set $P_{\nu,\omega} = \{\pi \in \Delta(\Omega) | \omega \in T_\pi(\nu)\}$ is residual. ν is *collectively rejected over ω* if the set $P_{\nu,\omega}$ is meager.

Lemma 1. *Suppose that the individual tests T_π are type-I error free. Then for every theory ν and every atom ω of ν , ν is collectively accepted over ω .*

So an expert is always accepted by the society if the realization is an atom of his forecast. In particular, if ν is atomic then

$$\nu(\{\omega | \nu \text{ is collectively accepted over } \omega\}) = 1,$$

so a true expert that reports an atomic distribution is collectively accepted with probability 1.

Instance 1—The society rejects the true expert when he is not atomic.

Theorem 2. *There exists a collective test $\{T_\pi\}_{\pi \in \Delta(\Omega)}$ such that:*

- (1) *For every $\pi \in \Delta(\Omega)$, the individual test T_π is π -conservative and type-I error free.*
- (2) *A forecast $\nu \in \Delta(\Omega)$ is collectively accepted over a realization ω only if ω is an atom of ν .*

Remark. Recall that by Lemma 1, if the individual tests are type-I error free then a forecast $\nu \in \Delta(\Omega)$ is collectively accepted over a realization ω which is an atom of ν . Thus, among all the collective tests that are individually type-I error free, the test in Theorem 2 is the one that is the least favorable towards gurus, be they true or false.

Remark. The test in Theorem 2 accepts a theory ν on $\{\omega \in \Omega \mid \sup_n L_{\pi,\nu}^n(\omega) < \infty\}$, with $L_{\pi,\nu}^n(\omega) = \frac{\pi(N_{\omega|n})}{\nu(N_{\omega|n})}$. In contrast, the test defined by Al-Najjar and Weinstein (2006) accepts ν on $\{\omega \in \Omega \mid \sup_n L_{\pi,\nu}^n(\omega) < 1\}$.

Note that their test can reject a true expert: If π is absolutely continuous w.r.t ν then the expert is rejected on the points where $1 < \frac{d\pi}{d\nu}$. On the other hand, Al-Najar and Weinstein's test satisfies uniform convergence in ν in (1).

Remark. The test in Theorem 2 depends only on forecasts along $\omega = (z_0, z_1, \dots)$:

$$L_{\pi,\nu}^n(\omega) = \frac{\nu(z_0) \cdot \nu(z_1|z_0) \cdot \dots \cdot \nu(z_n|z_0, \dots, z_{n-1})}{\pi(z_0) \cdot \pi(z_1|z_0) \cdot \dots \cdot \pi(z_n|z_0, \dots, z_{n-1})},$$

where $\nu(z_n|z_0, \dots, z_{n-1})$ and $\pi(z_n|z_0, \dots, z_{n-1})$ are the forecasts made by ν and π about z_n given z_0, \dots, z_{n-1} .

Instance 2 – The society always accepts a true expert.

Theorem 3. *There exists a collective test $\{T_\pi\}_{\pi \in \Delta(\Omega)}$ such that*

- (1) *For every $\pi \in \Delta(\Omega)$ the individual test T_π is conservative and type-I error free.*
- (2) *A true expert is collectively accepted with probability 1:*

$$\nu(\{\omega | \nu \text{ is collectively accepted over } \omega\}) = 1$$

for every $\nu \in \Delta(\Omega)$.

Instance 3 – Resistance to Strategic manipulation. Consider a false guru, who does not know the true distribution of the process. Such an expert can randomize a theory according to some probability distribution $\zeta \in \Delta(\Delta(\Omega))$. In the following theorem we describe a society that is immune to strategic manipulation by a false expert in the sense that over a topologically large set of realization such an expert will be almost surely rejected by the society.

Theorem 4. *There exists a collective test $\{T_\pi\}_{\pi \in \Delta(\Omega)}$ that satisfies the properties of Theorem 3 and, in addition, renders the society immune to strategic manipulation in the following sense: For every $\zeta \in \Delta(\Delta(\Omega))$, the set $\{\omega | \zeta(\{\nu | \nu \text{ is collectively rejected over } \omega\}) = 1\}$ is residual.*

Remark. Note that a given decision maker with initial belief π can be fooled by a false guru reporting theory π (In which case she will be fooled but not harmed). However, the society described in Theorem 4 will collectively reject the false guru. He can fool one individual, but, generically, he cannot fool them all.

5. PROOFS

5.1. Preliminaries.

5.1.1. *Lebesgue's decomposition.* Let π, ν be probability measures over Ω . π is *absolutely continuous w.r.t ν* ($\pi \ll \nu$) if $\nu(A) = 0$ implies $\pi(A) = 0$. π is *orthogonal to ν* ($\pi \perp \nu$) if there exists $B \in \mathcal{B}$ such that $\nu(B) = 1$ and $\pi(B) = 0$.

Let $B \in \mathcal{B}$. Say that π is *absolutely continuous w.r.t. ν on B* (and write $\pi \ll_B \nu$) if $\nu(A) = 0$ implies $\pi(A \cap B) = 0$.

Proposition 5. (*Lebesgue Decomposition*) *Let ν and π be two probability measures. Then $\pi = \pi^{(r)} + \pi^{(s)}$ where $\pi^{(r)}$ and $\pi^{(s)}$ are probability measures such that $\pi^{(r)} \ll \nu$ and $\pi^{(s)} \perp \nu$.*

Corollary 6. *For every ν, π there exists a set B such that $\nu(B) = 1$ and $\pi \ll_B \nu$.*

Proof. Let $\pi = \pi^{(r)} + \pi^{(s)}$ be the Lebesgue decomposition of π w.r.t. ν . Let B be a set such that $\nu(B) = 1$ and $\pi^{(s)}(B) = 0$. For every $A \in \mathcal{B}$, $\nu(A) = 0$ implies

$$\pi(A \cap B) = \pi^{(r)}(A \cap B) + \pi^{(s)}(A \cap B) \leq \pi^{(r)}(A) + \pi^{(s)}(B) = 0,$$

as desired. □

5.1.2. *Merging.* For a probability measure ν over Ω and $s \in Z^{<\mathbb{N}}$ we let ν_s be the probability measure over Ω that is given by

$$\nu_s(A) = \nu(A|N_s) = \nu(A \cap N_s)/\nu(N_s)$$

if $\nu(N_s) > 0$ and defined arbitrarily if $\nu(N_s) = 0$. The law of total expectation says that for every $n \in \mathbb{N}$ and every bounded Borel function R over Ω

$$\int R \, d\nu = \int \left(\int R \, d\nu_{\omega|n} \right) \nu(d\omega).$$

The distance between two probability measures ϕ_1 and ϕ_2 over Ω is given by

$$d(\phi_1, \phi_2) = \sup_{D \in \mathcal{B}} |\phi_1(D) - \phi_2(D)|.$$

Note that

$$(2) \quad \left| \int R \, d\phi_1 - \int R \, d\phi_2 \right| \leq d(\phi_1, \phi_2)$$

for every Borel function $R : \Omega \rightarrow [0, 1]$. Let π and ν be two probability measures over Ω . We say that ν *merges with* π if

$$(3) \quad \lim_{n \rightarrow \infty} d(\pi_{\omega|n}, \nu_{\omega|n}) = 0,$$

for π -almost every $\omega \in \Omega$. The following result was proved by Blackwell and Dubins.

Proposition 7. *Blackwell and Dubins (1962) If $\pi \ll \nu$ then ν merges with π*

Let $B \in \mathcal{B}$. Say that ν *merges with π on B* if (3) is satisfied for π -almost every $\omega \in B$.

Corollary 8. *If $\pi \ll_B \nu$ then ν merges with π on B .*

Proof. If $\pi(B) = 0$ then the corollary is satisfied trivially. Assume $\pi(B) > 0$, and let π' be the probability measure over Ω that is given by $\pi'(D) = \pi(D|B) = \pi(D \cap B)/\pi(B)$ for every $D \in \mathcal{B}$. Then $\pi' \ll \nu$.

Let $s \in Z^{<\mathbb{N}}$ and $D \in \mathcal{B}$. Then it follows from the definition of π' that

$$\pi'(D|N_s) = \pi(D \cap N_s \cap B)/\pi(N_s \cap B) = \pi(D \cap B|N_s)/\pi(B|N_s).$$

It follows that

$$\pi'(D|N_s) \geq \pi(D \cap B|N_s) \geq \pi(D|N_s) - (1 - \pi(B|N_s)),$$

and

$$\pi'(D|N_s) \leq \pi(D|N_s)/\pi(B|N_s) \leq \pi(D|N_s) + (1 - \pi(B|N_s))/\pi(B|N_s).$$

Therefore $d(\pi_s, \pi'_s) \leq (1 - \pi(B|N_s))/\pi(B|N_s)$. It follows that

$$(4) \quad \lim_{n \rightarrow \infty} d(\pi_{\omega|n}, \pi'_{\omega|n}) \leq \lim_{n \rightarrow \infty} (1 - \pi(B|N_{\omega|n}))/\pi(B|N_{\omega|n}) = 0$$

π -almost surely on B since by the martingale convergence theorem $\pi(B|N_{\omega|n}) \rightarrow \mathbf{1}_B(\omega)$.

By Blackwell-Dubins Theorem, since $\pi' \ll \nu$ it follows that

$$(5) \quad \lim_{n \rightarrow \infty} d(\pi'_{\omega|n}, \nu_{\omega|n}) = 0$$

π' -almost surely on B . Since $\pi \ll_B \pi'$ it follows that (4) is satisfied π -almost surely on B . From (4) and (5) it follows that

$$d(\pi_{\omega|n}, \nu_{\omega|n}) \leq d(\pi_{\omega|n}, \pi'_{\omega|n}) + d(\pi'_{\omega|n}, \nu_{\omega|n}) \xrightarrow[n \rightarrow \infty]{} 0$$

π -almost surely on B , as desired. \square

5.2. Proof of Theorem 1. By Bellman's principle of optimality, if f is ν -optimal then for every $s = (z_0, \dots, z_{n-1})$ and every strategy g one has

$$\int_{N_s} R_\lambda^n(y, g) \nu_s(dy) \leq \int_{N_s} R_\lambda^n(y, f) \nu_s(dy)$$

where for every strategy h and every $\omega = (z_0, z_1, \dots) \in \Omega$, $R_\lambda^n(y, h) = (1 - \lambda) \sum_{k=n}^{\infty} r(z_k, h(z_0, \dots, z_{k-1}))$ (an optimal strategy is optimal from every stage onward).

For every $\nu \in \Delta(\Omega)$ let $T(\nu) \in \mathcal{B}$ such that $\nu(T(\nu)) = 1$ and $\pi \ll_{T(\nu)} \nu$. Existence of $T(\nu)$ follows from Corollary 6. By Corollary 8, ν merges with π on $T(\nu)$.

Let f be a ν -optimal strategy and g a π -optimal strategy. We claim that (1) is satisfied. Indeed, let $\epsilon > 0$. Let $n \in \mathbb{N}$ be large enough such that $\pi(T(\nu) \setminus G) < \epsilon$ where $G = T(\nu) \cap \{x \mid d(\pi_{\omega|_n}, \nu_{\omega|_n}) < \epsilon\}$, and let $0 < \Lambda < 1$ be large enough such that $(1 - \Lambda^n)/(1 - \Lambda) < \epsilon$. Let $\lambda > \Lambda$. For $\omega = (z_0, z_1, \dots) \in \Omega$ and a strategy h let

$$R_\lambda^n(x, h) = (1 - \lambda) \sum_{n=N}^{\infty} \lambda^n r(z_n, f(z_0, \dots, z_{n-1})).$$

Then, by the choice of Λ and since $r(z, a) \in [0, 1]$, it follows that

$$(6) \quad 0 \leq R_\lambda^n(\omega, h) \leq R_\lambda(\omega, h) \leq R_\lambda^n(\omega, h) + \epsilon.$$

for every $\omega \in \Omega$ and every strategy h . Now,

$$\begin{aligned} \int_{T(\nu)} R_\lambda(\omega, g) \pi(d\omega) &< \int_G R_\lambda(\omega, g) \pi(d\omega) + \epsilon < \int_G R_\lambda^n(\omega, g) \pi(d\omega) + 2\epsilon = \\ &\int_G \left(\int R_\lambda^n(y, g) \pi_{\omega|N}(dy) \right) \pi(d\omega) + 2\epsilon \leq \int_G \left(\int R_\lambda^n(y, g) \nu_{\omega|N}(dy) \right) \pi(d\omega) + 3\epsilon \leq \\ &\int_G \left(\int R_\lambda^n(y, f) \nu_{\omega|N}(dy) \right) \pi(d\omega) + 3\epsilon \leq \int_G \left(\int R_\lambda^n(y, f) \pi_{\omega|N}(dy) \right) \pi(d\omega) + 4\epsilon = \\ &\int_G R_\lambda^n(\omega, f) \pi(d\omega) + 4\epsilon \leq \int_{T(\nu)} R_\lambda(\omega, f) \pi(d\omega) + 4\epsilon. \end{aligned}$$

The first inequality follows from the choice of n and the fact that $R_\lambda(\omega, g) \leq 1$. The second inequality follows from (6). The first equality follows from the law of total expectation. The third inequality follows from the definition of G and (2). The fourth inequality follows from the fact that f is ν -optimal and dynamic consistency. The fifth inequality follows from the definition of G and (2). The second equality follows from the law of total expectation. The last inequality follows from (6).

We proved that for every $\epsilon > 0$ there exists $0 < \Lambda < 1$ such that

$$\int_{T(\nu)} R_\lambda(\omega, g) \pi(d\omega) < \int_{T(\nu)} R_\lambda(\omega, f) \pi(d\omega) + 4\epsilon$$

for every $\lambda > \Lambda$. This completes the proof of the theorem. \square

Remark. Though our proof is based on Blackwell-Dubins merging theorem, a weaker notion of merging (Kalai and Lehrer, 1994) would have been sufficient for our cause.

5.3. Proofs from Section 4.

Proof of Lemma 1. Let $\nu \in \Delta(\Omega)$ be a forecast and let ω be an atom of ν . Since $\nu(T_\pi(\nu)) = 1$, it follows that $\omega \in T_\pi(\nu)$ for every π . Therefore $P_{\nu,\omega} = \Delta(\Omega)$. \square

Proof of Theorem 2. Let $T_\pi(\nu)$ be given by

$$(7) \quad T_\pi(\nu) = \left\{ \omega \in \Omega \left| \sup_n \frac{\pi(N_{\omega|n})}{\nu(N_{\omega|n})} < \infty \right. \right\}.$$

We claim first that T_π is π -conservative and does not reject the truth with probability 1. Indeed, fix $\pi, \nu \in \Delta(\Omega)$ and let $L^n : \Omega \rightarrow [0, \infty)$ be given by

$$L_{\pi,\nu}^n(\omega) = \frac{\pi(N_{\omega|n})}{\nu(N_{\omega|n})}$$

So $L_{\pi,\nu}^n(\omega)$ is the likelihood ratio of π and ν over the realization $\omega|_n$.

By (Durrett, 1996, Chapter 4.3.c) $L_{\pi,\nu}^n$ is a martingale under ν . Let $L = \sup L_{\pi,\nu}^n$. Then the Lebesgue Decomposition (see Section 5.1.1) of π over ν is given by $d\pi^{(r)} = Ld\nu$ and $\pi^{(s)}(A) = \pi(A \cap \{L = \infty\})$. Therefore T_π is a version of the test we construct in the proof of Theorem 1. In particular, T_π is π -conservative and does not reject the truth with probability 1.

Let $\nu \in \Delta(\Omega)$ be a theory and let $\omega \in \Omega$ be a realization such that $\nu(\{\omega\}) = 0$. We claim that ν is collectively rejected over ω . By (7)

$$P_{\nu,\omega} = \bigcup_{M=1}^{\infty} \bigcap_{n=1}^{\infty} \left\{ \pi \in \Delta(\Omega) \mid \pi(N_{\omega|n}) \leq M \cdot \nu(N_{\omega|n}) \right\}.$$

Since the sets $N_{\omega|n}$ are clopen it follows that the function $\pi \mapsto \pi(N_{\omega|n})$ is continuous in the weak* topology and therefore the sets $\left\{ \pi \in \Delta(\Omega) \mid \pi(N_{\omega|n}) \leq M \cdot \nu(N_{\omega|n}) \right\}$ are closed in the weak* topology and therefore $P_{\nu,\omega}$ is an F_σ set.

We now claim that $P_{\nu,\omega}$ has an empty interior. Indeed, let $\pi \in P_{\nu,\omega}$, and let M be such that $\pi(N_{\omega|n}) \leq M \cdot \nu(N_{\omega|n})$ for every n . Then

$$\pi(\{\omega\}) = \lim_{n \rightarrow \infty} \pi(N_{\omega|n}) \leq M \cdot \lim_{n \rightarrow \infty} \nu(N_{\omega|n}) = M \cdot \nu(\{\omega\}) = 0.$$

Therefore $P_{\nu,\omega} \subseteq \{\pi \in \Delta(\Omega) | \pi(\{\omega\}) = 0\}$. The assertion that $P_{\nu,\omega}$ has empty interior follows from Lemma 2 below. \square

Lemma 2. *For every $\omega \in \Omega$ the set $\{\pi \in \Delta(\Omega) | \pi(\{\omega\}) > 0\}$ is dense.*

Proof of Lemma 2. For every $\mu \in \Delta(\omega)$ let $\pi_n = (1 - 1/n)\mu + 1/n\delta_\omega$ where δ_ω is dirac measure over ω . Then $\pi_n(\{\omega\}) > 0$ and $\pi_n \rightarrow \mu$ in the norm topology, and, in particular in the weak* topology. \square

The proof of Theorem 3 uses the following proposition (Dekel and Feinberg, 2006, Proposition 1)

Proposition 9. (1) *For every $\nu \in \Delta(\Omega)$ there exists a meager subset F of Ω such that $\nu(F) = 1$.*

(2) *For every meager subset F of Ω the set $\{\pi \in \Delta(\Omega) | \pi(F) > 0\}$ is a meager subset of $\Delta(\Omega)$.*

Proof of Theorem 3. For every $\nu \in \Delta(\Omega)$ let $t(\nu)$ be a meager subset of Ω such that $\nu(t(\nu)) = 1$. For every $\pi \in \Delta(\Omega)$ let $T_\pi : \Delta(\Omega) \rightarrow 2^\Omega$ be such that $T_\pi(\nu) = t(\nu)$ whenever $\pi(t(\nu)) = 0$ and $T_\pi(\nu)$ is defined arbitrary such that $\nu(T_\pi(\nu)) = 1$ and $\pi \ll_{T_\pi(\nu)} \nu$ whenever $\pi(t(\nu)) > 0$ (this is possible by using Lebesgue's Decomposition in the same way

as in the proof of Theorem 1). Then T_π is a version of the test we constructed in the proof of Theorem 1 and, in particular, is π -conservative and does not reject the truth with probability 1.

Fix $\nu \in \Delta(\Omega)$ and let $\omega \in t(\nu)$. It follows from the definition of T that $\{\pi|\pi(t(\nu)) = 0\} \subseteq P_{\nu,\omega}$ and it follows from Proposition 9 that the set $\{\pi|\pi(t(\nu)) = 0\}$ is residual. Therefore the set $P_{\nu,\omega}$ is residual. It follows that ν is collectively accepted over ω for every $\omega \in t(\nu)$. Since $\nu(t(\nu)) = 1$ the result follows. \square

The proof of Theorem 4 uses the following proposition proved by Olszewski and Sandroni (2007a).

Proposition 10. *There exists a function $t : \Delta(\Omega) \rightarrow 2^\Omega$ such that*

- (1) *For every $\nu \in \Delta(\Omega)$, $t(\nu)$ is meager and $\nu(t(\nu)) = 1$.*
- (2) *For every $\zeta \in \Delta(\Delta(\Omega))$ the set*

$$\{\omega|\zeta(\{\nu|\omega \notin t(\nu)\}) = 1\}$$

is residual.

Proof of Theorem 4. Let $t : \Delta(\Omega) \rightarrow 2^\Omega$ be as in Proposition 10 and let $\{T_\pi\}_{\pi \in \Delta(\Omega)}$ be such that $T_\pi(\nu) = t(\nu)$ whenever $\pi(t(\nu)) = 0$ and $T_\pi(\nu)$ is defined arbitrary such that $\nu(T_\pi(\nu)) = 1$ and $\pi \ll_{T_\pi(\nu)} \nu$ whenever $\pi(t(\nu)) > 0$. Then T_π is a version of the test constructed in the proof of Theorem 3.

Fix a realization ω and let $\nu \in \Delta(\Omega)$ be such that $\omega \notin t(\nu)$. Then by the definition of T_π , $P_{\omega,\nu} \subseteq \{\pi|\pi(t(\nu)) > 0\}$. By Proposition 9 it

follows that the set $\{\pi|\pi(t(\nu)) > 0\}$ is meager and therefore the set $P_{\omega,\nu}$ is meager. Thus, if $\omega \notin t(\nu)$ then $P_{\omega,\nu}$ is meager, i.e. ν collectively fails over ω . Let $\zeta \in \Delta(\Delta(\Omega))$. It follows from the previous observation that, for every realization ω ,

$$\zeta(\{\nu|\omega \notin t(\nu)\}) = 1 \rightarrow \zeta(\{\nu|\nu \text{ is collectively rejected over } \omega\}) = 1.$$

Since the set $\{\omega|\zeta(\{\nu|\omega \notin t(\nu)\}) = 1\}$ is residual by Proposition 10 it follows that the set $\{\omega|\zeta(\{\nu|\nu \text{ is collectively rejected over } \omega\}) = 1\}$ is residual. \square

6. CONCLUDING REMARKS

We have worked on the case where the test comes after the decisions were made. This is in line with the literature on testing experts because in that literature there is only one realization of ω which plays a role. With more than one realization, the situation is closer to that of statistical sampling, and the problem of testing depends on how one models the dependence between each observation.

REFERENCES

- AL-NAJJAR, N., AND J. WEINSTEIN (2006): “Comparative Testing of Experts,” mimeo, Northwestern University.
- BLACKWELL, D., AND L. DUBINS (1962): “Merging of opinions with increasing information,” *Annals of Mathematical Statistics*, 33, 882–886.
- DEKEL, E., AND Y. FEINBERG (2006): “Non-Bayesian testing of a stochastic prediction,” *Review of Economic Studies*, 73.

- DURRETT, R. (1996): *Probability: Theory and Examples*. Wadsworth Publishing Company.
- FEINBERG, Y., AND C. STEWART (2007): “Testing Multiple Forecasters,” Research Paper 1957, Stanford GSB.
- FOSTER, D. P., AND R. V. VOHRA (1998): “Asymptotic Calibration,” *Biometrika*, 85(2), 379–390.
- KALAI, E., AND E. LEHRER (1994): “Weak and Strong Merging of opinions,” *Journal of Mathematical Economics*, 23, 73–86.
- LEHRER, E. (2001): “Any Inspection Is Manipulable,” *Econometrica*, 69(5), 1333–1347.
- OLSZEWSKI, W., AND A. SANDRONI (2007a): “A NonManipulable Test,” preprint.
- (2007b): “Strategic Manipulation of Empirical Tests,” mimeo, Northwestern University.
- SANDRONI, A. (2003): “The reproducible properties of correct forecasts,” *Internat. J. Game Theory*, 32(1), 151–159.
- SHMAYA, E. (2007): “Many Inspections are Manipulable,” mimeo, Caltech.