Theoretical Economics 3 (2008), 367–382 1555-7561/20080367

Many inspections are manipulable

Eran Shmaya
Kellogg School of Management, Northwestern University

A self-proclaimed expert uses past observations of a stochastic process to make probabilistic predictions about the process. An inspector applies a test function to the infinite sequence of predictions provided by the expert and the observed realization of the process in order to check the expert's reliability. If the test function is Borel and the inspection is such that a true expert always passes it, then it is also manipulable by an ignorant expert. The proof uses Martin's theorem about the determinacy of Blackwell games. Under the axiom of choice, there exist non-Borel test functions that are not manipulable.

Keywords. Forecasting, calibration, zero-sum games.

JEL classification. C72, C73.

1. Introduction

At every period $n = 0, 1, 2, \ldots$ nature chooses an outcome s_n from a finite set S. An expert claims to know the underlying distribution behind nature's choices. To prove his claim, at each period n the expert provides a probabilistic prediction p_n about s_n before s_n is realized. The question addressed in this paper is whether the expert's reliability can be tested from the infinite sequence of predictions $(p_0, p_1, \ldots)$ provided by the expert and the actual observed sequence $(s_0, s_1, \ldots)$.

Assume that an inspector decides on the reliability of the expert by applying some test function—a function whose arguments are the predictions $(p_0, p_1, \ldots)$ made by the expert and the actual realization $(s_0, s_1, \ldots)$, and whose value is either 'pass' if the predictions fit the realization or 'fail' otherwise. The calibration test, which checks whether the predicted frequencies of events equal their observed frequencies, is a well known example of such a test function. Several versions of the calibration test are studied in the literature (Foster and Vohra 1997, Lehrer 2001, Sandroni et al. 2003), and all of them turn out to be manipulable: An ignorant expert, who does not know the distribution of the process, can strategically generate predictions so that he matches the performance of a true expert, who predicts according to the correct distribution. In this paper I prove that every Borel test function is manipulable, without making any assumption about its form.

Eran Shmaya: e-shmaya@kellogg.northwestern.edu

This paper was written while I was a research fellow in the Social and Information Sciences Laboratory at California Institute of Technology. I am grateful to Federico Echenique for detailed comments and suggestions that substantially improved the paper, and also to Chris Chambers, Eddie Dekel, Yossi Feinberg, Nabil Al-Najjar, Wojciech Olszewski, Alvaro Sandroni, and Leeat Yariv.

Olszewski and Sandroni (2008a), extending a previous theorem of Sandroni (2003), have already proved a general manipulability result of the type sought in this paper. They consider an inspection in which the expert can be rejected only at some finite period: if he is not rejected at any finite period then he passes the test. This assumption is natural from an economic perspective, since a real world inspection is based on a finite data sequence. Topologically, this assumption translates to semi-continuity of the test function. Olszewski and Sandroni prove that such a test function is always manipulable.[1] My contribution is twofold: I extend Olszewski and Sandroni's result from semi-continuous test functions to arbitrary Borel functions, thus dispensing with the assumption that rejection is determined at a finite period, and I give an example of a non-Borel test function that is not manipulable. My proof uses Martin's theorem about the determinacy of Blackwell games, which is a new tool in this literature.

A comparison between the results of this paper and those Dekel and Feinberg (2006) and Olszewski and Sandroni (2008b) shows that the scope of the prediction that the expert provides is crucial for the existence of a non-manipulable inspection. These authors prove the existence of non-manipulable inspections that are based on a prediction about the entire infinite realization of the process, which the expert announces before any data is realized. The manipulability theorem of this paper, on the other hand, relies on the fact that at every period the expert provides predictions about the current period, or, more generally, about a finite number of future periods, but not about events that are only determined at infinity. To emphasize this point, I give an example of a non-manipulable sequential inspection, in which at every period the expert makes a prediction about a single event that is only determined at infinity.

Theorems 1 and 2 in Section 2 are the main results of this paper—every Borel test function that does not reject the truth with high probability is manipulable, and, under the axiom of choice, there exists a non-Borel test function that is not manipulable. Section 3 discusses related literature. Section 4 presents Martin's theorem. The proofs of the theorems are in Sections 5 and 6. In Section 7 I give an example of a Borel non-manipulable inspection that is based on repeated predictions about a single event. Section 8 concludes.

2. Manipulable and non-manipulable tests

Let S be a finite set. Elements of $\Delta(S)$, the simplex of probability distributions over S, are called *predictions*. At every period $n = 0,1,2,\dots$ an *outcome* $s_n \in S$ is realized. At every period, before s_n is realized, an expert declares a prediction $p_n \in \Delta(S)$ about s_n, based on the past outcomes $s_0,\dots,s_{n-1}$. A *realization* is given by an infinite sequence $s \in S^{\mathbb{N}}$ of outcomes, where $\mathbb{N} = \{0,1,2,\dots\}$.

Let $S^{<\mathbb{N}} = \bigcup_{n\geq 0} S^n$ be the set of all finite sequences of elements of S, including the empty sequence e. For every realization $s = (s_0, s_1, \dots) \in S^{\mathbb{N}}$ and every $n \in \mathbb{N}$ let $s|_n = (s_0,\dots,s_{n-1})$ be the initial segment of s of length n. In particular, $s|_0 = e$.

[1] In fact, their result is formulated in a wider framework, without the assumption that the predictions are only about the current period. See Section 3.3.

DEFINITION 1. A *prediction rule* is a function $f : S^{<\mathbb{N}} \to \Delta(S)$.

If the expert uses a prediction rule f then his prediction about s_n after observing $(s_0, \dots, s_{n-1})$ is $f(s_0, \dots, s_{n-1})$.

DEFINITION 2. A *test function* is a function $T : \Delta(S)^{\mathbb{N}} \times S^{\mathbb{N}} \to \{0,1\}$.

A test function T dictates, for every infinite sequence of predictions and every realization, whether or not the expert passes the inspection:

DEFINITION 3. Let T be a test function and $s \in S^{\mathbb{N}}$ a realization. A sequence $p \in \Delta(S)^{\mathbb{N}}$ of predictions *passes T over s* if $T(p,s) = 1$. A prediction rule f *passes T over s* if p passes T over s, where $p = (p_0, p_1, \dots)$ is the sequence of predictions of f along the realization, $p_n = f(s_0, \dots, s_{n-1})$.

Ideally, an expert who predicts according to the distribution of the process passes the test with high probability, and an ignorant expert who does not know the distribution of the process is not able to pass the test. Definitions 4 and 5 below formalize these notions. These definitions are special cases of the definitions in Olszewski and Sandroni (2008a), which pertain to a more general notion of test function (see Section 3.2).

DEFINITION 4. A test function T *does not reject the truth with probability* $1-\epsilon$ if

$$\mathbb{P}\big(f \text{ passes } T \text{ over } \Theta_0, \Theta_1, \dots\big) \geq 1 - \epsilon \tag{1}$$

for every sequence of random variables $\Theta_0, \Theta_1, \dots$ with values in S, where the prediction rule $f : S^{<\mathbb{N}} \to \Delta(S)$ is given by

$$f(s_0, \dots, s_{n-1})[s_n] = \mathbb{P}(\Theta_n = s_n \mid \Theta_0 = s_0, \dots, \Theta_{n-1} = s_{n-1}) \tag{2}$$

for every $s_0, \dots, s_n \in S$.

REMARK 1. The random variables $\Theta_0, \Theta_1, \dots$ are defined over some probability space $(\Omega, \mathscr{A}, \mathbb{P})$. The event $\{f \text{ passes } T \text{ over } \Theta_0, \Theta_1, \dots\}$ is the subset of Ω that is given by $\{\omega \in \Omega \mid f \text{ passes } T \text{ over } \Theta_0(\omega), \Theta_1(\omega), \dots\}$. Note that if T is a Borel function then this set is in $\mathscr{A}$. If T is universally measurable[2] then the left side of (1) is still well defined. The inequality (1) is meaningful even when the set is not measurable, in which case the meaning is that there exists some $A \in \mathscr{A}$ such $\mathbb{P}(A) \geq 1 - \epsilon$ and f passes T over $\Theta_0(\omega), \Theta_1(\omega), \dots$ for every $\omega \in A$.

REMARK 2. Instead of quantifying over all sequences of random variables with values in S, I could write Definition 4 by quantifying over all distributions in $S^{\mathbb{N}}$, as in Olszewski and Sandroni (2008a, Definition 1): A test function T does not reject the truth with probability $1-\epsilon$ if

$$\mu_f\big(x \in S^{\mathbb{N}} \mid f \text{ passes } T \text{ over } x\big) \geq 1 - \epsilon$$

[2] A subset M of a standard Borel space X is *universally measurable* if it is μ-measurable for every probability measure μ over X.

for every prediction rule f where μ_f is the distribution induced by f over $S^{\mathbb{N}}$. I prefer the terminology of random variables, which is more suitable to the probabilistic arguments of this paper and which renders the proof simpler.

DEFINITION 5. A test function T is ϵ-manipulable if there exists some probability measure ξ over prediction rules such that

$$\xi\left(\{f \mid f \text{ passes } T \text{ over } s\}\right) \geq 1-\epsilon, \tag{3}$$

for every $s \in S^{\mathbb{N}}$.

If a test function is ϵ-manipulable then an ignorant expert can randomize his prediction rule according to ξ and pass the test with high probability, regardless of the actual realization.

REMARK 3. The set of prediction rules $f : S^{<\mathbb{N}} \to \Delta(S)$ is naturally endowed with a standard Borel structure, as a countable product of copies of a simplex. If T is a Borel function then the set $\{f \mid f \text{ passes } T \text{ over } s\}$ is a Borel set. If T is not a Borel function then the inequality (3) means that there exists some Borel set B of prediction rules such that $\xi(B) \geq 1-\epsilon$ and f passes T over s for every $f \in B$.

The following theorem is the negative result of this paper: Under the Borel assumption, a test function that does not reject a true expert is manipulable by an ignorant expert.

THEOREM 1. *Let $T : \Delta(S)^{\mathbb{N}} \times S^{\mathbb{N}} \to \{0,1\}$ be a Borel test function. If T does not reject the truth with probability $1-\epsilon$, then T is $\epsilon+\delta$-manipulable for every $\delta > 0$.*

REMARK 4. The test function T is a Borel function if the set $\{(p,s) \mid T(p,s)=1\}$ is a Borel subset of $\Delta(S)^{\mathbb{N}} \times S^{\mathbb{N}}$. Theorem 1 and its proof remain valid if the space $\Delta(S)^{\mathbb{N}} \times S^{\mathbb{N}}$ is equipped with the product of discrete topologies over $\Delta(S)$ and S, which gives rise to a larger class of Borel functions than the standard Borel structure over $\Delta(S)^{\mathbb{N}} \times S^{\mathbb{N}}$, since the discrete topology over $\Delta(S)$ is stronger than the Euclidean topology.

Can a test function $T : \Delta(S)^{\mathbb{N}} \times S^{\mathbb{N}} \to \{0,1\}$ that is not Borel be manipulable? It follows from the proof of Theorem 1 that there is a model of set theory without the axiom of choice in which Theorem 1 is valid for an arbitrary test function (Remark 5). The next theorem shows that in ZFC there exists a test function that is not manipulable.

THEOREM 2. *Let $S = \{0,1\}$. Under the axiom of choice, there exists a universally measurable test function $T : \Delta(S)^{\mathbb{N}} \times S^{\mathbb{N}} \to \{0,1\}$ such that*

(i) T does not reject the truth with probability 1.

(ii) For every probability distribution ξ over prediction rules there exists $s \in S^{\mathbb{N}}$ such that

$$\xi\left(\{f \mid f \text{ passes } T \text{ over } s\}\right) = 0.$$

In particular, T is not ϵ-manipulable for any $\epsilon > 0$.

3. Related literature

3.1 *Calibration tests*

Calibration tests (Dawid 1982, Foster and Vohra 1997, Kalai et al. 1999, Fudenberg and Levine 1999, Lehrer 2001, Sandroni et al. 2003) compare the observed frequencies of events over sets of periods with the average predictions over the same sets of periods. I follow Lehrer (2001). For $p \in \Delta(S)$ and a subset V of S let $P[V] = \sum_{s\in V} p[s]$ and let $\mathbf{1}_V : S \to \{0,1\}$ be the indicator function of V. A *simple calibration test* is given by a pair (U,C), where U and C are functions that assign to every observation $(s_0,\dots,s_n) \in S^{<\mathbb{N}}$ subsets $C(s_0,\dots,s_n)$ and $U(s_0,\dots,s_n)$ of S such that $C(s_0,\dots,s_n) \subseteq U(s_0,\dots,s_n)$. The interpretation is that $U(s_0,\dots,s_n)$ is the local universe considered after $s_0,\dots,s_n$ and, within this universe the event $C(s_0,\dots,s_n)$ is checked. The simple calibration test $T^{U,C} : \Delta(S)^{\mathbb{N}} \times S^{\mathbb{N}} \to \{0,1\}$ induced by (U,C) checks whether the conditional probability attached by the expert to the events $C(s_0,\dots,s_n)$ given $U(s_0,\dots,s_n)$ matches the empirical relative frequency. Formally,

$$T^{U,C}(p,s) = 1 \text{ if } \sum_{n=0}^{\infty} \mathbf{1}_{U(s|_n)}(s_n) = \infty \text{ implies}$$
$$\lim_{n\to\infty} \frac{\sum_{i=0}^{n} \mathbf{1}_{U(s|_i)}(s_i)\cdot\left(p_i[U(s|_i)]\mathbf{1}_{C(s|_i)}(s_i) - p_i[C(s|_i)]\right)}{\sum_{i=0}^{n} \mathbf{1}_{U(s|_i)}(s_i)} = 0$$

for every realization $s = (s_0, s_1, \dots) \in S^{\mathbb{N}}$ and every infinite sequence of predictions $p = (p_0, p_1, \dots) \in \Delta(S)^{\mathbb{N}}$. Then $T^{U,C}$ is a Borel test function that does not reject the truth with probability 1.

A more general calibration test is given by a mixture of simple calibration tests. Let $\mathscr{S}$ be the set of simple calibration tests and let λ be a probability distribution over $\mathscr{S}$. Assume the inspector first chooses a simple calibration test $T \in \mathscr{S}$ using λ and then applies this test to the expert's predictions. For every such λ, Lehrer constructs a prediction rule that passes the inspection λ-almost surely. Lehrer's result relates to Theorem 1 of this paper in the following way. Let $\Lambda : \Delta(S)^{\mathbb{N}} \times S^{\mathbb{N}} \to \{0,1\}$ be the test function given by

$$\Lambda(p,s) = 1 \text{ if and only if } \int T(p,s)\lambda(\mathrm{d}T) = 1.$$

Then Λ is a Borel test function that does not reject the truth with probability 1. By Theorem 1, it is ϵ-manipulable for every ϵ. Thus, there exists a manipulation scheme that passes the calibration test with high probability. For the case of calibration tests, Lehrer proves a stronger result: He shows that the test is 0-manipulable and that the manipulation scheme is *pure*—there exists a prediction rule f such that f passes the test Λ over s for every s.[3] Most importantly, Lehrer constructs the manipulating rule, while my proof is not constructive.

[3] Other calibration tests do not admit a pure manipulation (Sandroni et al. 2003, Example 2.1).

3.2 *Infinite-horizon predictions*

According to the setup of Section 2, at each period the expert has to provide a prediction about the outcome of that period. Dekel and Feinberg (2006) and Olszewski and Sandroni (2008b) consider a different framework, in which at the start of the inspection the expert must inform the inspector of his prediction about the entire realization of the process. A test function in this context is a function $t : \Delta(S^{\mathbb{N}}) \times S^{\mathbb{N}} \to \{0,1\}$, where $\Delta(S^{\mathbb{N}})$ is the set of all probability measures over the set $S^{\mathbb{N}}$ of realizations. I call elements of $\Delta(S^{\mathbb{N}})$ *infinite-horizon predictions*, denote tests that are based on infinite-horizon predictions by lower case t, and call them *infinite-horizon tests*. When necessary to avoid confusion, I call elements of $\Delta(S)$ *one-period predictions* and tests in the sense of Definition 2 *one-period tests*.

There is a natural correspondence $f \longleftrightarrow \mu_f$ between prediction rules according to Definition 1 and probability measures over S^∞: for every prediction rule f, μ_f is the joint distribution of a sequence $\Theta_0, \Theta_1, \ldots$ of random variables satisfying (2). Because of this correspondence, every one-period test function $T : \Delta(S)^{\mathbb{N}} \times S^{\mathbb{N}} \to \{0,1\}$ naturally induces[4] an infinite-horizon test function t such that $t(\mu_f, s) = 1$ if and only if f passes T over s, but the converse is not true. Dekel and Feinberg (2006) and Olszewski and Sandroni (2008b) prove the existence of infinite-horizon test functions that do not reject the truth with probability 1 and are not manipulable. In light of Theorem 1 it should also be mentioned that the test function constructed by Olszewski and Sandroni is a Borel function when $\Delta(S^\infty)$ is equipped with its standard Borel structure. In Section 7 I discuss in detail the difference between finite-horizon and infinite-horizon predictions.

3.3 *Future-independent tests*

Olszewski and Sandroni (2008a) consider (infinite-horizon) test functions $t : \Delta(S^{\mathbb{N}}) \times S^{\mathbb{N}} \to \{0,1\}$ of the form

$$t(\mu_f, s) = \begin{cases} 0 & \text{if } (f|_n, s|_n) \in R \text{ for some } n \in \mathbb{N} \\ 1 & \text{otherwise} \end{cases} \tag{4}$$

for some $R \subseteq \bigcup_{n \geq 0} \left(\Delta(S)^{\bigcup_{i=0}^{n} S^i} \times S^n\right)$, where $f|_n$ is the restriction of f to $\Delta(S)^{\bigcup_{i=0}^{n-1} S^i}$ for every prediction rule f and $s|_n = (s_0, \ldots, s_{n-1})$ for every realization $s = (s_0, s_1, \ldots) \in S^{\mathbb{N}}$. The underlying assumption is that rejection must occur at some period n, and that the decision of the inspector at that period depends only on the segment of the data $s|_n$ that was realized before that period and on the prediction rule before that period. The set R (the rejection set) consists of all finite segments of realizations and prediction rules that are considered to be inconsistent with each other.

Test functions of the form (4) are called *future-independent*: the decision whether to reject an infinite-horizon prediction μ_f at some period does not depend on the predictions made by f at later periods. For future-independent tests, Olszewski and Sandroni

[4] There is a minor inaccuracy here because μ_f does not determine f uniquely. Cantankerous readers can assume that $T(p,s) = 0$ whenever $p_n[s_n] = 0$ for some $n \in \mathbb{N}$, so that T determines t uniquely.

prove an analogue of Theorem 1: a future-independent test that does not reject the truth with high probability is manipulable.

It is interesting to compare the theorem of Olszewski and Sandroni with Theorem 1. Consider first a one-period test function $T : \Delta(S)^{\mathbb{N}} \times S^{\mathbb{N}} \to \{0,1\}$. If the infinite-horizon test function induced by T is future-independent, then the set $\{(p,s) \mid T(p,s) = 1\}$ is closed (recall Remark 4) and, in particular, Borel. Therefore in the framework of test function studied in this paper, the scope of Theorem 1 is wider than that of Olszewski and Sandroni's theorem. However, there are future-independent, infinite-horizon tests that are not induced by one-period test functions. In fact, neither of the theorems is contained in the other. Olszewski and Sandroni's inspector is more restricted in that he must decide to reject using a finite number of predictions. On the other hand, he can use the entire prediction rule up to the rejection point, including predictions conditioned on observations outside the actual realization.

4. Blackwell games

A *Blackwell game* is a two-player zero-sum game that is given by (A, B, r) where A and B are the sets of *actions* of player 1 (the maximizer) and 2 (the minimizer) respectively and $r : (A \times B)^{\mathbb{N}} \to [0,1]$ is the *payoff function*.

The game is played as follows. At every stage $n = 0, 1, 2, \ldots$ both players, simultaneously and independently, choose an action. At the end of the stage, each player is informed of his opponent's action. Let a_n, b_n be the actions chosen by players 1 and 2 respectively at stage n. The payoff that player 2 pays player 1 is given by $r(a_0, b_0, a_1, b_1, \ldots)$.

Let $\mathcal{H} = (A \times B)^{<\mathbb{N}} = \bigcup_{n \geq 0} (A \times B)^n$ be the set of *finite histories* of the game, including the empty history e. A behavioral strategy σ of player 1 is given by $\sigma : \mathcal{H} \to \Delta(A)$. Behavioral strategies τ of player 2 are defined analogously. Every pair σ, τ of behavioral strategies induces a probability distribution $\mu_{\sigma,\tau}$ over the set $(A \times B)^{\mathbb{N}}$ of *infinite histories* or *plays*. Let $R(\sigma,\tau) = \int r \, d\mu_{\sigma,\tau}$ be the expected payoff in the game if the players play according to σ, τ.

Determinacy

The *upper value* $\overline{V}(G)$ and the *lower value* $\underline{V}(G)$ of G of a Blackwell game G are given by

$$\overline{V}(G) = \inf_{\tau} \sup_{\sigma} R(\sigma,\tau)$$
$$\underline{V}(G) = \sup_{\sigma} \inf_{\tau} R(\sigma,\tau),$$

where the suprema range over all behavioral strategies σ of player 1 and the infima range over all behavioral strategies τ of player 2. A strategy σ of player 1 is *δ-optimal* if $R(\sigma,\tau) \geq \underline{V}(G) - \delta$ for every behavioral strategy τ of player 2. The game G is *determined* if $\underline{V}(G) = \overline{V}(G)$. Blackwell (1969, 1989) proves the determinacy of Blackwell games (which he calls infinite games with imperfect information) with a payoff function that is the indicator function of a G_δ set, and conjectures that every Blackwell game with Borel payoff function is determined. Vervoort (1996) advances higher in the Borel hierarchy, proving

determinacy for indicators of $G_{\delta\sigma}$ sets. The conjecture is proved by Martin (1998) (see also Maitra and Sudderth 1998 for applications to stochastic games).

MARTIN'S THEOREM. *Let A, B be two countable sets, at least one of which is finite, and let $r:(A\times B)^{\mathbb{N}}\to[0,1]$ be a Borel function. Then the Blackwell game (A,B,r) is determined.*[5]

Random plays

Let (σ,τ) be a pair of behavioral strategies in the Blackwell game (A,B,r). A (σ,τ)-*random play* is a sequence

$$\alpha_0,\beta_0,\ldots,\alpha_n,\beta_n,\ldots$$

of random variables over some probability space, where the values of α_n (respectively β_n) are in A (respectively B) such that

$$\mathbb{P}\left(\alpha_n=a,\beta_n=b\mid\alpha_0,\beta_0,\ldots,\alpha_{n-1},\beta_{n-1}\right)=\\ \sigma(\alpha_0,\beta_0,\ldots,\alpha_{n-1},\beta_{n-1})[a]\cdot\tau(\alpha_0,\beta_0,\ldots,\alpha_{n-1},\beta_{n-1})[b]$$

for every $a\in A$ and $b\in B$.

The measure $\mu_{\sigma,\tau}$ that is induced by (σ,τ) over $(A\times B)^{\mathbb{N}}$ is the joint distribution of some (σ,τ)-random play. The payoff function associated with a pair of behavioral strategies (σ,τ) can also be written in terms of random plays: $R(\sigma,\tau)=\mathbb{E}\left(r\left(\alpha_0,\beta_0,\alpha_1,\beta_1,\ldots\right)\right)$.

Pure and mixed strategies

A *pure strategy* of player 1 in the Blackwell game (A,B,r) is a function $f:B^{<\mathbb{N}}\to A$; for every sequence $b_0,\ldots,b_n$ of past actions of player 2, $f(b_0,\ldots,b_n)$ is player 1's action at stage $n+1$. Every pure strategy is in particular a behavioral strategy. A *mixed strategy* of player 1 is a probability distribution over pure strategies. Kuhn's Theorem (Sorin 2002, Theorem D.1) establishes the equivalence between behavioral and mixed strategies. In particular, for every $\delta>0$ player 1 has a δ-optimal mixed strategy in every Blackwell game, i.e. a mixed strategy ξ such that $R(\xi,g)\geq\underline{V}(G)-\delta$ for every pure strategy g of player 2, where $R(\xi,g)=\int R(f,g)\xi(\mathrm{d}f)$ is the expected payoff for player 1 under ξ,g.

5. PROOF OF THEOREM 1

Let $\Delta^{\mathbb{Q}}(S)=\{p\in\Delta(S)\mid p[s]\in\mathbb{Q}\text{ for every }s\in S\}$ be the set of elements of $\Delta(S)$ with rational values. For a test function T let $G(T)$ be the Blackwell game in which the set of actions of player 1 is $\Delta^{\mathbb{Q}}(S)$, the set of actions of player 2 is S, and the payoff function is the restriction of T to $(\Delta^{\mathbb{Q}}(S)\times S)^{\mathbb{N}}$. Note that every pure strategy of player 1 in $G(T)$ is a prediction rule according to Definition 1. Roughly speaking, player 1 represents the expert and player 2 represents nature. However, in the game $G(T)$, player 2 is allowed to condition his actions on past actions of player 1 (as if nature picks the value of s_n

[5]Martin states his theorem for the case when both action sets are finite and points out that the proof works when one of these sets is countably infinite.

depending on previous predictions made by the expert) and player 1 is only allowed to make predictions with rational values.[6]

The game $G(T)$ satisfies the assumptions of Martin's theorem: the action set of player 1 is finite, the action set of player 2 is countable, and the payoff function is Borel. Therefore $\overline{V}(G(T)) = \underline{V}(G(T))$. The following two lemmas complete the proof of Theorem 1.

LEMMA 1. *If T does not reject the truth with probability $1-\epsilon$ then $\overline{V}(G(T)) \geq 1-\epsilon$.*

LEMMA 2. *T is $(1-\underline{V}(G(T))+\delta)$-manipulable for every $\delta > 0$*

The proof of Lemma 1 uses the following lemma. Recall that, for a finite set S and $p, p' \in \Delta(S)$, a *coupling of* (p, p') is a pair (Θ, Θ') of random variables such that $\mathbb{P}(\Theta = s) = p[s]$ and $\mathbb{P}(\Theta' = s) = p'[s]$ for every $s \in S$, i.e. the marginal distributions of Θ and Θ' are p and p' respectively.

COUPLING LEMMA (Lindvall 1992, Chapter 1, Theorem 5.2). *Let S be a finite set and let $p, p' \in \Delta(S)$. Then there exists a coupling (Θ, Θ') of (p, p') such that $\mathbb{P}(\Theta \neq \Theta') = \|p - p'\|_1/2$.*[7]

PROOF OF LEMMA 1. Let τ be a behavioral strategy for player 2 in $G(T)$. We have to construct a good response for player 1 against τ. The strategy is such that at every stage player 1 predicts the action of player 2 for that stage. Note that since τ is given, at every stage player 1 knows the probability distribution according to which player 2 is going to choose an action. However, since in $G(T)$ player 1 is only allowed to make predictions with rational values, his strategy only approximates this distribution.

Let $\delta > 0$ and let $f, f' : S^{<\mathbb{N}} \to \Delta(S)$ be the prediction rules defined inductively as follows. For every $(s_0, \ldots, s_n) \in S^{<\mathbb{N}}$ let

$$f(s_0, \ldots, s_n) = \tau(p_0, \ldots, p_n, s_0, \ldots, s_n), \tag{5}$$

where $p_i = f'(s_0, \ldots, s_{i-1})$, and let $f'(s_0, \ldots, s_n) \in \Delta^Q(S)$ be such that

$$\|f'(s_0, \ldots, s_n) - f(s_0, \ldots, s_n)\|_1 < \delta/2^n. \tag{6}$$

Then f' is a pure strategy of player 1 in $G(T)$. I am going to construct a (f', τ)-random play $(\Pi_0, \Theta_0, \Pi_1, \Theta_1, \ldots)$ and, on the same probability space, a stochastic process $(\Theta'_0, \Theta'_1, \ldots)$ that equals $(\Theta_0, \Theta_1, \ldots)$ with high probability, such that f' is the correct prediction rule for $(\Theta'_0, \Theta'_1, \ldots)$.

[6] Olszewski and Sandroni (2008b) use another game-theoretic representation of the expert's problem. Their game is a normal form one-shot game in which nature chooses a realization and the expert chooses a prediction rule. They use topological properties of the test function to deduce the determinacy of the game using the classical minimax theorem.

[7] I essentially use the coupling lemma to prove that the set of probability measures μ_f that are induced by prediction rules f with rational values is dense in the norm topology. Cf. Lemma 2 in Olszewski and Sandroni (2008a).

Let $\Theta=(\Theta_n)_{n\in\mathbb{N}}$ and $\Theta'=(\Theta'_n)_{n\in\mathbb{N}}$ be random variables over a probability space defined inductively such that, for $s_0,s'_0,\dots,s_{n-1},s'_{n-1}\in S$, the conditional joint distribution of Θ_n,Θ'_n given the event $\{\Theta_i=s_i,\Theta'_i=s'_i \text{ for } 1\le i<n\}$ satisfies

$$\mathbb{P}\left(\Theta_n=s\mid\Theta_i=s_i,\Theta'_i=s'_i \text{ for } 0\le i<n\right)=f(s_0,\dots,s_{n-1})[s] \tag{7}$$

$$\mathbb{P}\left(\Theta'_n=s'\mid\Theta_i=s_i,\Theta'_i=s'_i \text{ for } 0\le i<n\right)=f'(s'_0,\dots,s'_{n-1})[s'], \tag{8}$$

and, if $s_i=s'_i$ for $0\le i<n$, then also

$$\mathbb{P}(\Theta_n\neq\Theta'_n\mid\Theta_i=s_i,\Theta'_i=s'_i \text{ for } 0\le i<n)\le\delta/2^n. \tag{9}$$

If $s_i=s'_i$ for $0\le i<n$ then the existence of a pair of random variables that satisfies (7), (8), and (9) follows from (6) and the Coupling Lemma. If $s_i\neq s'_i$ for some $0\le i<n$ then the conditional joint distribution of Θ_n,Θ'_n given the event $\{\Theta_i=s_i,\Theta'_i=s'_i \text{ for } 0\le i<n\}$ can be chosen arbitrarily with the marginals given by (7) and (8). Note that from (7) and (8) it follows that

$$\mathbb{P}(\Theta_n=s\mid\Theta_0=s_0,\dots,\Theta_{n-1}=s_{n-1})=f(s_0,\dots,s_{n-1})[s] \tag{10}$$

$$\mathbb{P}(\Theta'_n=s'\mid\Theta'_0=s'_0,\dots,\Theta'_{n-1}=s'_{n-1})=f'(s'_0,\dots,s'_{n-1})[s']. \tag{11}$$

Also, from (9) it follows that

$$\mathbb{P}(\Theta_n\neq\Theta'_n\mid\Theta_i=\Theta'_i \text{ for } 0\le i\le n)\le\delta/2^n,$$

and therefore

$$\mathbb{P}(\Theta\neq\Theta')\le\sum_{n\in\mathbb{N}}\mathbb{P}(\Theta_n\neq\Theta'_n\mid\Theta_i=\Theta'_i \text{ for } 0\le i\le n)\le 2\delta. \tag{12}$$

Let $\Pi=(\Pi_n)_{n\in\mathbb{N}}$ be given by

$$\Pi_n=f'(\Theta_0,\dots,\Theta_{n-1}). \tag{13}$$

Then it follows from (5), (10), and (13) that

$$\mathbb{P}(\Theta_n=s\mid\Pi_0,\Theta_0,\dots,\Pi_{n-1},\Theta_{n-1})=\tau(\Pi_0,\Theta_0,\dots,\Pi_{n-1},\Theta_{n-1})[s]. \tag{14}$$

From (13) and (14) it follows that

$$(\Pi_0,\Theta_0,\Pi_1,\Theta_1,\dots) \text{ is a } (f',\tau)\text{-random play in } G(T). \tag{15}$$

Therefore

$$\begin{aligned}
R(f',\tau)&=\mathbb{E}\big(T(\Pi_0,\Theta_0,\Pi_1,\Theta_1,\dots)\big)\\
&=\mathbb{P}\big(\Pi \text{ passes } T \text{ over } \Theta\big)\\
&\ge\mathbb{P}\big(\Pi \text{ passes } T \text{ over } \Theta'\big)-\mathbb{P}(\Theta\neq\Theta')\\
&=\mathbb{P}\big(f' \text{ passes } T \text{ over } \Theta'\big)-\mathbb{P}(\Theta\neq\Theta')\\
&\ge 1-\epsilon-\mathbb{P}(\Theta\neq\Theta')\\
&\ge 1-\epsilon-2\delta,
\end{aligned}$$

where the first equality follows from (15), the second equality from Definition 3, the first inequality from the fact that

$$\{\Pi \text{ passes } T \text{ over } \Theta'\} \subseteq \{\Pi \text{ passes } T \text{ over } \Theta\} \cup \{\Theta' \neq \Theta\},$$

the third equality from (13) and Definition 3, the second inequality from (11) and the fact that T does not reject the truth with probability $1-\epsilon$, and the third inequality from (12). Thus, for every strategy τ of player 2 and every $\delta > 0$ we have built a pure strategy f' of player 1 such that $R(f', \tau) \geq 1 - \epsilon - 2\delta$. Therefore $\overline{V}(G(T)) \geq 1 - \epsilon$ as desired. □

Proof of Lemma 2. Let ξ be a mixed δ-optimal strategy for player 1 in $G(T)$. We claim that ξ, viewed as a distribution over prediction rules, $(1 - \underline{V}(G(T)) + \delta)$-manipulates T. Indeed, let s be a realization and let g be the pure strategy of player 2 in $G(T)$ given by $g(p_0, \ldots, p_{n-1}) = s_n$ for every $p_0, \ldots, p_{n-1} \in \Delta^Q(S)$.

Let f be a pure strategy of player 1 in $G(T)$. Then it follows from Definition 3 that

$$R(f, g) = \begin{cases} 1 & \text{if } f \text{ passes } T \text{ over } s \\ 0 & \text{otherwise.} \end{cases}$$

It follows that

$$\underline{V}(G(T)) - \delta \leq R(\xi, g) = \int R(f, g)\xi(\mathrm{d}f) = \xi\{f \mid f \text{ passes } T \text{ over } s\},$$

as desired. □

Remark 5. There is a model of set theory without the axiom of choice in which every Blackwell game is determined (Martin 1998, Theorem 13). In this model every set is universally measurable, which makes Definitions 4 and 5 meaningful for an arbitrary test function T. It follows from the proof that in such a model Theorem 1 is valid for an arbitrary test function.

6. Proof of Theorem 2

The test is a modification of the non-manipulable (infinite-horizon) test of Dekel and Feinberg (2006, Proposition 2). A subset M of $\{0,1\}^{\mathbb{N}}$ is *universally null* if M is universally measurable and $\mu(M^c) = 1$ for every non-atomic probability measure μ over $\{0,1\}^{\mathbb{N}}$. It follows from the axiom of choice that there exist universally null sets in $\{0,1\}^{\mathbb{N}}$ of cardinality $\aleph_1$ (Miller 1984, Theorem 5.3).[8] Such a set is not a Borel set. Note that for a universally null set M and an arbitrary probability measure μ one has

$$\mu(M^c \cup A(\mu)) = 1, \tag{16}$$

where $A(\mu) = \{s \in \{0,1\}^{\mathbb{N}} \mid \mu(\{s\}) > 0\}$ is the set of atoms of μ.

[8] The existence result in Miller's paper is formulated for the real line. Since by Kuratowski's isomorphism theorem the space $\{0,1\}^{\mathbb{N}}$ and the real line are isomorphic as Borel spaces, the result holds also in $\{0,1\}^{\mathbb{N}}$.

Let $S=\{0,1\}$ and let M be an uncountable universally null subset of $S^{\mathbb{N}}$. Let $T:\Delta(S)^{\mathbb{N}}\times S^{\mathbb{N}}\to\{0,1\}$ be the test function given by

$$T(p,s)=\begin{cases}0 & \text{if } s\in M \text{ and } \prod_{n\in\mathbb{N}} p_n[s_n]=0\\ 1 & \text{otherwise.}\end{cases}$$

Note that T is universally measurable. I claim that T satisfies the requirements of Theorem 2.

Let $f:S^{<\mathbb{N}}\to\Delta(S)$ be a prediction rule and let $s\in S^{\mathbb{N}}$. Let $\Theta_0,\Theta_1,\dots$ be a sequence of random variables satisfying (2) and let $\mu_f\in\Delta(S^{\mathbb{N}})$ be their joint distribution. Then

$$\begin{aligned}\mu_f(\{s\})&=\mathbb{P}\big(\Theta_n=s_n \text{ for every } n\in\mathbb{N}\big)\\ &=\prod_{n\in\mathbb{N}}\mathbb{P}\big(\Theta_n=s_n\mid\Theta_i=s_i \text{ for } 0\le i<n\big)=\prod_{n\in\mathbb{N}}p_n[s_n],\end{aligned}$$

where $p_n=f(s_0,\dots,s_{n-1})$.

By the last equation and the definition of T it follows that f passes T on $s\in M$ if and only if $\mu_f(\{s\})>0$, i.e. $s\in A(\mu_f)$. The rest of the argument is the same as in Dekel and Feinberg's paper: For every prediction rule f

$$\mu_f\big(\{s\mid f \text{ passes } T \text{ over } s\}\big)=\mu_f(M^c\cup A(\mu_f))=1$$

(the last equality follows from (16)), and therefore T does not reject the truth with probability 1.

To prove the second assertion of Theorem 2, let ξ be a probability distribution over prediction rules and let $s\in M$. If

$$\xi\big(\{f\mid f \text{ passes } T \text{ over } s\}\big)=\xi\big(\{f\mid\mu_f(\{s\})>0\}\big)>0,$$

then in particular $\overline{\xi}(\{s\})>0$ where $\overline{\xi}\in\Delta(S^{\mathbb{N}})$ is the barycenter of ξ, given by

$$\overline{\xi}(B)=\int\mu_f(B)\xi(\mathrm{d}f)$$

for every Borel subset B of $\Delta(S^{\mathbb{N}})$. Thus ξ passes the test with some positive probability over $s\in M$ only if $s\in A(\overline{\xi})$. Since M is uncountable and $A(\overline{\xi})$ is countable, it follows that there are some $s\in M$ over which ξ passes the test with probability 0.

Remark 6. The test constructed by Olszewski and Sandroni (2008b) has the stronger property that for every randomly generated prediction rule ξ, the set of all realizations s over which ξ passes tests with positive probability is a set of first Baire category.

7. Finite-horizon predictions and predictions about a finite set

Theorem 1, which deals with one-period tests, can be generalized to tests with a finite horizon. Let $k\ge 1$ be a natural number. Consider an inspection in which, at every

period n, the expert provides a probabilistic prediction about $(s_n, s_{n+1}, \ldots, s_{n+k-1})$. A *$k$-horizon prediction* is an element of $\Delta(S^k)$. A *k-horizon test function* is a function $T : \Delta(S^k)^{\mathbb{N}} \times S^{\mathbb{N}} \to \mathbb{R}$. A *$k$-horizon prediction rule* is a function $f : S^{<\mathbb{N}} \to \Delta(S^k)$. Definitions 4 and 5 extend in an obvious way to the case of k-horizon test functions, and the analogue of Theorem 1 is also true: If a k-horizon test function does not reject the truth with probability $1-\epsilon$ then it is $\epsilon+\delta$-manipulable for every $\delta > 0$. Moreover, the number k need not be constant or bounded, and can depend on past realizations and predictions.

On the other hand, as Dekel and Feinberg and Olszewski and Sandroni show, infinite-horizon tests can be non-manipulable. In order to emphasize that the manipulability result of this paper relies on the fact that the predictions requested from the expert are about events in the finite horizon, and not just on the fact that the set S^k over which predictions are made at each period is finite, consider the following situation. Fix a Borel set $B \subseteq S^{\mathbb{N}}$. Assume that at every period n, given the partial realization $(s_0, \ldots, s_{n-1})$ observed at that period, the expert is asked to make a prediction $q_n \in [0,1]$ about whether the event B will occur, that is whether the infinite realization s is in B. Consider the test function $T_B : [0,1]^{\mathbb{N}} \times S^{\mathbb{N}} \to \{0,1\}$ given by

$$T_B(q,s) = \begin{cases} 1 & \text{if } \lim_{n\to\infty} q_n = \mathbf{1}_B(s) \\ 0 & \text{otherwise.} \end{cases}$$

A *prediction rule about B* is a function $f : S^{<\mathbb{N}} \to [0,1]$. A prediction rule passes the test T_B over realization $s \in S^{\mathbb{N}}$ if $T_B(q,s) = 1$ where $q \in [0,1]^{\mathbb{N}}$ is given by $q_n = f(s_0, \ldots, s_{n-1})$. It follows from the martingale convergence theorem that the test T_B does not reject the truth with probability 1. That is, for every stochastic process $\Theta_0, \Theta_1, \ldots$ with values in S one has

$$\mathbb{P}(f \text{ passes } T_B \text{ on } \Theta_0, \Theta_1, \ldots) = 1$$

where f is the prediction rule given by

$$f(s_0, \ldots, s_{n-1}) = \mathbb{P}\big((\Theta_0, \Theta_1, \ldots) \in B \mid \Theta_0 = s_0, \ldots, \Theta_{n-1} = s_{n-1}\big).$$

Note that in the inspection induced by T_B the expert is always asked to state a prediction about only two possibilities—either B occurs or B does not occur. Still, as I show in the following example, T_B need not be manipulable when B is an event in the infinite horizon.

Example 1. Let $B \subseteq S^{\mathbb{N}}$ be a Borel set that is not an F_σ set. Then the test T_B is not ϵ-manipulable for any $\epsilon < 1/2$. ◊

Indeed, Let ξ be a probability measure over prediction rules. Let

$$\overline{B} = \left\{ s \in S^{\mathbb{N}} \mid \xi\left(\left\{ f \mid \liminf_{n\to\infty} f(s|_n) > 1/2 \right\} \right) > 1/2 \right\}.$$

Since for every $s \in S^{\mathbb{N}}$ one has

$$\left\{ f \mid \liminf_{n\to\infty} f(s|_n) > 1/2 \right\} = \bigcup_{r=1}^{\infty} \bigcup_{n=1}^{\infty} \bigcap_{k=n}^{\infty} \{ f \mid f(s|_k) > 1/2 + 1/r \}$$

it follows that

$$s\in\overline{B}\longleftrightarrow\exists r,n\quad \xi\left(\bigcap_{k=n}^{\infty}\{f\mid f(s|_k)>1/2+1/r\}\right)>1/2$$
$$\longleftrightarrow\exists r,n,t\ \forall m\geq n\quad \xi\left(\bigcap_{k=n}^{m}\{f\mid f(s|_k)>1/2+1/r\}\right)>1/2+1/t.$$

Therefore

$$\overline{B}=\bigcup_{r,n,t=1}^{\infty}\bigcap_{m=n}^{\infty}\overline{B}(r,n,t,m),$$

where

$$\overline{B}(r,n,t,m)=\left\{s\mid\xi\left(\bigcap_{k=n}^{m}\{f\mid f(s|_k)>1/2+1/r\}\right)>1/2+1/t\right\}.$$

Since the sets $\overline{B}(r,n,t,m)$ are clopen (membership in $\overline{B}(r,n,t,m)$ depends on only a finite number of coordinates) it follows that $\overline{B}$ is an F_σ set. By the choice of B it follows that $B\neq\overline{B}$. If s is any element in the symmetric difference of B and $\overline{B}$ then $\xi\left(\{f\mid f \text{ passes } T_B \text{ on } s\}\right)\leq 1/2$. In particular, T_B is not ϵ-manipulable for any $\epsilon<1/2$.

Remark 7. Let $S=\{0,1\}$. As an example of a Borel subset B of $S^{\mathbb{N}}$ that is not an F_σ set one can take $B=\{s\in S^{\mathbb{N}}\mid\sum_{n\in\mathbb{N}}s_n=\infty\}$, the set of realizations with infinitely many 1's.

8. Conclusions

The inspections considered in this paper are *sequential*: they require the expert to announce at every period a probabilistic prediction in $\Delta(P)$ for some finite set P of possibilities. In one-period inspections, $P=S$, the set of possible outcomes in the period. In k-horizon inspections (Section 7), $P=S^k$, the set of possible outcomes in the k next periods. In inspections T_B about an event B (Section 7), P is a set consisting of two possibilities: 'B occurs' and 'B does not occur'. In contrast, the inspections studied by Dekel and Feinberg (2006) and Olszewski and Sandroni (2008b) require the expert to provide one prediction in $\Delta(S^{\mathbb{N}})$ about the entire realization of the process.

The paper bears good news and bad news for the inspector. There exist sequential inspections that are not manipulable (Theorem 2 and Example 1). However, such inspections must be either non-Borel or rely on predictions about events that are not determined at any finite period.

References

Blackwell, David (1969), "Infinite G_δ games with imperfect information." *Zastosowania Matematyki (Applicationes Mathematicae)*, 10, 99–101. [373]

Blackwell, David (1989), "Operator solution of infinite G_δ games of imperfect information." In *Probability, Statistics, and Mathematics* (Theodore W. Anderson, Krishna B.

Athreya, and Donald L. Igelhart, eds.), 83–87, Academic Press, Boston, Massachusetts. [373]

Dawid, A. Philip (1982), "The well-calibrated Bayesian." *Journal of the American Statistical Association*, 77, 605–610. [371]

Dekel, Eddie and Yossi Feinberg (2006), "Non-Bayesian testing of a stochastic prediction." *Review of Economic Studies*, 73, 893–906. [368, 372, 377, 380]

Foster, Dean P. and Rakesh V. Vohra (1997), "Calibrated learning and correlated equilibrium." *Games and Economic Behavior*, 21, 40–55. [367, 371]

Fudenberg, Drew and David K. Levine (1999), "An easier way to calibrate." *Games and Economic Behavior*, 29, 131–137. [371]

Kalai, Ehud, Ehud Lehrer, and Rann Smorodinsky (1999), "Calibrated forecasting and merging." *Games and Economic Behavior*, 29, 151–169. [371]

Lehrer, Ehud (2001), "Any inspection is manipulable." *Econometrica*, 69, 1333–1347. [367, 371]

Lindvall, Torgny (1992), *Lectures on the Coupling Method*. Wiley, New York. [375]

Maitra, Ashok P. and William D. Sudderth (1998), "Finitely additive stochastic games with Borel measurable payoffs." *International Journal of Game Theory*, 27, 257–267. [374]

Martin, Donald A. (1998), "The determinacy of Blackwell games." *Journal of Symbolic Logic*, 63, 1565–1581. [374, 377]

Miller, Arnold W. (1984), "Special subsets of the real line." In *Handbook of Set-Theoretic Topology* (Kenneth Kunen and Jerry E. Vaughan, eds.), 201–233, North-Holland, Amsterdam. [377]

Olszewski, Wojciech and Alvaro Sandroni (2008a), "Manipulability of future-independent tests." *Econometrica*, forthcoming. [367, 369, 372, 375]

Olszewski, Wojciech and Alvaro Sandroni (2008b), "A nonmanipulable test." *Annals of Statistics*, forthcoming. [368, 372, 375, 378, 380]

Sandroni, Alvaro (2003), "The reproducible properties of correct forecasts." *International Journal of Game Theory*, 32, 151–159. [368]

Sandroni, Alvaro, Rann Smorodinsky, and Rakesh V. Vohra (2003), "Calibration with many checking rules." *Mathematics of Operations Research*, 28, 141–153. [367, 371]

Sorin, Sylvain (2002), *A First Course on Zero-Sum Repeated Games*. Springer, Heidelberg. [374]

Vervoort, Marco R. (1996), "Blackwell games." In *Statistics, Probability, and Game Theory* (Thomas S. Fergurson, Lloyd S. Shapley, and James B. MacQueen, eds.), 369–390, Institute of Mathematical Statistics, Hayward, California. [373]

Submitted 2007-10-28. Final version accepted 2008-7-6. Available online 2008-7-6.