The principal-agent approach to testing experts*

Wojciech Olszewski[†] and Marcin Pęski[‡]

Abstract

Recent literature on testing experts shows that it is difficult, and often impossible, to determine whether an expert knows the stochastic process that generates data. Despite this negative result, we show that often exist contracts that allow a decision maker to attain the first-best payoff in the following sense: in the case in which the expert knows the stochastic process, the decision maker achieves the payoff she would obtain if there were no incentive problems; while in the case in which the expert does not know the stochastic process, she achieves the payoff she would obtain in the absence of any expert.

More precisely, this kind of full-surplus extraction is always possible in infinite-horizon models in which future payoffs are not discounted. If future payoffs are discounted (but the discount factor tends to 1), the possibility of full-surplus extraction depends on a constraint involving the forecasting technology.

*We would like to thank the seminar participants at Arizona State University, Johns Hopkins, Northwestern, Ohio State University, University of Pittsburg, and the audiences of the Canadian Economic Theory Conference in Toronto and the First Transatlantic Theory Workshop in Paris, and the European Econometric Society Meeting in Barcelona for helpful comments and conversations. Wojciech Olszewski thanks the National Science Foundation for research support (CAREER award SES-0644930).

[†]Department of Economics, Northwestern University, 2001 Sheridan Road, Evanston IL 60208

[‡]Economics Department, University of Texas at Austin, 1 University Station, Austin TX 78712

1. Introduction

A number of recent papers show that if an empirical test can be passed by an expert who knows a stochastic process that generates states, and reports this knowledge truthfully, then the test can also be passed an expert who knows nothing about the stochastic process, but delivers forecasts strategically in order to pass this particular test. That is, empirical tests cannot distinguish between these two types of experts. (See Foster and Vohra (1998), Fudenberg and Levine (1999), Lehrer (2001), Sandroni (2003), Sandroni, Smorodinsky, and Vohra (2003), Vovk and Shafer (2005), Olszewski and Sandroni (2008) and (2009), and Shmaya (2008).)

The literature on testing experts does not focus on and hence does not specify in which way one benefits either from learning the expert's type or the stochastic process itself. Statistical decision theory views information as a tool for making better decisions. For example, Wald (1949) and (1950) sees hypothesis testing as a decision-maker's strategy in a game with uncertainty. Of course, there is no conflict between this view and the literature on testing experts. Even if the decision problem is not explicitly modelled, one may argue that when we know the expert's type, we are able to make better decisions.

However, the impossibility of screening the expert's type may depend critically on whether the expert's forecasts play the role of advice concerning a specific decision problem, or alternatively, whether one has no specific decision problem in mind, but simply wishes to learn the expert's type. Indeed, there are several reasons for thinking that this difference may be crucial. In practice, a decision maker must take some default action even in the absence of any expert, and may not appreciate forecasts that suggest the same (or similar) actions. On the other hand, if forecasts lead to better decisions, the decision maker may appreciate them, no matter what type of the expert is who provides them. Thus, in an analysis of forecasting in the context of a single decision problem, it seems legitimate to relax the requirement that a "good" test should always pass informed experts, and fail uninformed ones.

In addition, the famous Hannan's Theorem (see Hannan (1957)) says that if a decision maker receives forecasts from a finite number of sources, then there exists a decision scheme that enables her to achieve (in the long run) an average payoff as high as the maximum of the average payoffs she would achieve by receiving forecasts from single sources. Hannan's result seems quite robust to perturbations of the original model, e.g., to introducing some type of strategic forecasters, or a cost of switching between receiving forecasts from distinct sources. Some of these results are reviewed in Foster and Vohra (1999), and in Cesa-Bianchi and Lugosi (2006).[1] The variety of extensions of Hannan's Theorem suggest that the decision maker should be able to achieve a payoff as high as she would achieve when she knew the type of the expert. We must note, however, that there exist important differences between the literature on Hannan's theorem and the present setting; one consequence of this is that we obtain some impossibility results as well.

We study the relationship between a tester and an expert within an infinite-horizon principal-agent model. In a single period, the principal (henceforth, "the decision maker") takes an action, and her utility depends on the action and the unknown state of the world. The decision maker's knowledge regarding the states cannot be summarized by a probability distribution. The agent (henceforth "the expert") can be either informed, in which case he knows the stochastic process that generates the states; or uninformed, in which case he knows no more than the decision maker. The decision maker offers the expert a menu of contracts. According to each contract, the expert is supposed to provide forecasts in selected periods (i.e., the expert delivers a probability distribution over states in each selected period); the decision maker then chooses an optimal action given the expert's forecast. Each menu contains a default contract under which the expert never provides any forecasts. In periods in which the expert provides no forecast - in particular, in each period under the default contract - the decision maker takes a default action, and the expert exerts his outside option. This outside option can be interpreted as the utility of leisure, or avoiding the cost of providing a forecast.[2] The expert's compensation depends on his forecasts and the

[1] We are grateful to Rekesh Vohra for a discussion on the relation between our results and extensions of Hannan's Theorem.

[2] All results of the present paper hold under the assumption that only the informed expert has

realized states. He is free to choose any contract in the menu.

We consider two criteria for evaluating infinite sequences of payoffs. If the parties do not discount future payoffs, i.e., they evaluate sequences of payoffs by taking the limits of averages, then the principal can approximate her first-best payoff. That is, there exists a menu of contracts that satisfies these two conditions: If the expert is informed, the decision maker's payoff is almost equal to the payoff she would get if the expert honestly revealed the process that generates states, and if the decision maker had to compensate the expert for his forecasts only when the expert's forecasts are useful (i.e., the payoff under the contract exceeds the payoff from taking the default action by the value of the expert's outside option). If the expert is uninformed, then even if he accepts a contract, the decision maker's payoff can fall only marginally below the payoff she would obtain in the absence of any expert.

Under our contracts, the expert can choose a number of grace periods during which he cannot be dismissed. After the grace periods are over, the expert is dismissed (once and for all) when the decision maker's average payoff (including transfers to the expert) falls below the payoff to taking default actions. If the expert is not dismissed, he is compensated for his outside option, and receives a little bonus, proportional to the decision maker's payoff.

For a given discount factor $\delta < 1$, the first-best menu (in the above sense) does not exist. Moreover, it does not exist even when $\delta \to 1$ if contracts are *irreversible*, i.e., in a world in which once the expert does not deliver a forecast, he cannot deliver a forecast in any future period. In contrast, if reversible contracts are allowed, first-best menus do exist as $\delta \to 1$.

The distinction between reversible and irreversible contracts reflects the constraints that the contracting parties face in practice. If forecasting requires a relation-specific capital that disappears when the expert is dismissed, the contracts must be irreversible. On the other hand, if the expert can be hired or fired at will, the reversible contracts are appropriate.

The present analysis contrasts with the message of the literature on testing ex-

an outside option.

perts. For example, Olszewski and Sandroni (2007) study a similar principal-agent model, in which parties discount future payoffs. They show that if the informed expert accepts a contract, the uninformed expert also prefers to accept the contract; thus, the decision maker cannot learn the expert's type.[3] In their model, default actions are not specified. But if they were, the assumption that the informed expert finds a contract profitable would mean that the contract must be profitable even for the expert whose forecasts induce the same actions as (or similar actions to) those that the decision maker would take in the absence of any expert. We postpone a detailed discussion of the relation to the literature on testing experts to Section 4.1.

The idea that the expert's forecasts should be studied in the context of a specific decision problem was examined in a recent paper by Echenique and Shmaya (2008). In their paper, a decision maker also takes a sequence of actions, and discounts future payoffs by δ. The decision maker is equipped in a default stochastic process π, and before taking any action, receives from an expert an alternative stochastic process ν. Echenique and Shmaya study empirical tests that give verdicts at infinity. Using Lebesgue's decomposition theorem, they prove the existence of a test which is passed by the informed and honest expert with probability 1; in addition, contingent on the test being passed, the expected payoff of the decision maker from taking the optimal actions according to ν would be at least as high as the expected payoff from taking the optimal actions according to π.[4]

2. Model

At the beginning of each period $t = 1, 2, ...$, a state $s_t \in S$ is realized. Next, a decision maker (female) takes an action $a_t \in A$. We assume that A is a compact metric space, and S (equipped with a $\sigma-$algebra) is a measurable space. At the end of period t, the state s_t is revealed. State s_t is generated according to a probability distribution

[3] Note also that the decision problem studied in their paper is only one specific example, whereas we study a general decision problem in the present paper.

[4] The expectations of the decision maker's payoffs are taken with respect to π, and the expected payoffs are compared in the limit as $\delta \to 1$.

$p_t = p_t(s^t)$ which may depend on the history of past states $s^t = (s_1, ..., s_{t-1})$; let $p_t[s_t]$ denote the probability of state s_t in period t according to distribution p_t. We emphasize that $p_t = p_t(s^t)$ is an arbitrary probability distribution. We make no assumptions regarding the properties of this distribution, e.g., that the probability distributions are i..i.d., Markovian, or exchangeable.

The decision maker does not know state s_t when she takes action a_t. Moreover, she is completely ignorant, and does not even know $p_t(s^t)$. A potential expert (male) may know p_t, $t = 1, 2, ...$ However, the "expert" - like the decision maker - may also not know anything about p_t. In other words, there is a continuum of informed types of the expert (one for every sequence of history-dependent distributions p_t, $t = 1, 2, ...$), and one uninformed type.

A forecast in period t is a probability distribution $f_t = f_t(s^t)$ over states $s_t \in S$, which may depend on the history of states s^t; let $f_t[s_t]$ denote the probability of state s_t in period t according to forecast f_t.

Contracts

In period 0, the decision maker offers the expert a menu of contracts $\mathcal{C}$. A contract c specifies periods in which the expert is supposed to provide a forecast. Let $e_t = 1$ if the expert is supposed to provide a forecast in period t, and let $e_t = 0$ otherwise; e_t may depend on the states, forecasts, and also the values of all other variables observed before period t. Let $e = (e_t)_{t=1}^{\infty}$. A contract also specifies payments w_t; each w_t is a function of the expert's forecast f_t and the realized state s_t, but may also depend on the values of all variables observed before or in period t. We will assume that $w_t = 0$ in periods in which $e_t = 0$. Let $w = (w_t)_{t=1}^{\infty}$. That is, $c = (w, e)$.

We assume that in periods in which $e_t = 0$, the expert receives a payment $\overline{w} \geq 0$ from an external source. One can think about $\overline{w}$ as an outside option; or, assuming that the informed expert may not know $p_t(s^t)$, but can only learn about this probability distribution, the opportunity cost of providing a forecast. It would not affect the results if we assumed that the outside option is lower (or equal to zero) for the uninformed expert. We also assume that each menu contains a *default contract*

$c^0 = (w^0, e^0)$, where $w_t^0 \equiv \overline{w}$ and $e_t^0 \equiv 0$; that is, the expert has the option of rejecting the entire menu of contracts.

A contract is called *irreversible* if $e_t = 0$ implies $e_m = 0$ in every future period $m \geq t$. Contracts which are not irreversible will be called reversible. If forecasting requires a relation-specific capital that disappears when the expert is dismissed, the contract must be irreversible. Otherwise, reversible contracts might be feasible.

Payoffs

Agents discount future payoffs with a common discount factor $\delta \leq 1$. The utility of the decision maker in period t is $u_t(a_t, s_t) - w_t$; it depends on action a_t, state s_t, and payment w_t, and is quasi-linear with respect to this payment. We assume that u_t is a continuous function of a_t and a measurable function of s_t. We might also assume that u_t is a function of history s^t. We dropped s^t from the set of variables in order to simplify notation. If $\delta < 1$, the decision maker's total utility is defined as the (normalized) present value of utilities in periods $t = 1, 2, ...$; more precisely, this utility is

$$U(s, a, w) := (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \left[u_t(a_t, s_t) - w_t \right],$$

where $s = (s_t)_{t=1}^{\infty}$, $a = (a_t)_{t=1}^{\infty}$, and $w = (w_t)_{t=1}^{\infty}$.

The expert's utility is equal to the (normalized) present value of the payments he will receive in periods $t = 1, 2, ...$:

$$W(w, e) := (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} [e_t w_t + (1 - e_t)\overline{w}].$$

Denote by $W^0 \equiv \overline{w}$ the expert's default payoff.

The separability of the decision maker's payoff from actions and payments is important; however, none of the results would change if she had a convex disutility from payment, and the expert had a concave utility in w_t.

The informed expert, who knows the probability distribution over future states, selects a contract and delivers forecasts in order to maximize his expected payoff. The decision maker's and the uninformed expert's preferences cannot be described in a similar manner, since none of them knows the probability distribution over future states. Referring to the decision-theory terminology, we would say that they face Knightian uncertainty, or are not probabilistically sophisticated. The existing literature offers several distinct models of decision-making under ambiguity. However, we need not assume any particular preference representation in the present analysis.

If $\delta = 1$, parties evaluate flows of payoffs according to the limits of averages, instead of normalized weighted averages. This criterion is traditionally interpreted as corresponding to no discounting, or representing preferences of infinitely patient individuals. The decision maker's total utility is equal to

$$\lim_{T\to\infty}\inf U^T(s,a,w),$$

where

$$U^T(s,a,w) = \frac{1}{T}\sum_{t=1}^{T}\left[u_t(a_t,s_t) - w_t\right].$$

Similarly, the expert's utility is equal to the limit inferior of the average payments he will receive.

Action Choice

If a forecast is provided (i.e., in periods in which $e_t = 1$), the decision maker takes an action $a_t^{f_t}$ that maximizes her expected utility, computed under the assumption that the states are generated by the forecasted probability distribution f_t, i.e.,

$$a_t^{f_t} \in \arg\max E^{f_t} u_t(a_t, s_t).$$

When the forecast is not provided, i.e., $e_t = 0$, the decision maker takes a *default action* $a_t^0(s^t)$. This action may depend on the history of states s^t. The default actions are known to the expert.

One can think about the default actions as those that would be taken in the absence of any expert. The decision maker could use her favorite learning strategy in order to use the data on past states to predict which action will fare best in the current period. It is convenient to normalize the decision maker's utility so that

$$u_t\left(a^0\left(s^t\right), s_t\right) \equiv 0 \text{ for each } t\text{, past history } s^t\text{, and state } s_t.$$

Because the utility u_t may depend on the history of past states s^t, the normalization does not impose any additional constraints on the decision maker's payoffs. Due to this normalization, the decision maker's payoff under the default contract is equal to 0 at any history of states. An action $a \in A$ is *undominated* (in period t) if for every other action $a \neq a^{'} \in A$, there exists a state $s \in S$ such that $u_t(a, s) > u_t(a^{'}, s)$, or $u_t(a, s) = u_t(a^{'}, s)$ for every state $s \in S$. We assume that every default action $a_t^0(s^t)$ is undominated.

We also assume that the decision maker is committed from period 0 on to taking the actions prescribed in the previous paragraph. This assumption makes the analysis tractable. Without commitment, the information revealed by the expert's choices could affect the decision maker's actions. The selection of a particular contract or the provision of a particular forecast may contain some information about probability distributions p_t. However, if we allowed the decision maker to take optimal actions (given information available to her), we would have a dynamic choice problem under ambiguous information; in addition, agents would interact while making their choices. Given the current state of the literature on ambiguity, we find the no-commitment case intractable.

So far, all payoffs have been defined as functions of states, actions, and the contractual arrangements. However, since the actions are completely determined by the contract and the forecasts provided by the expert, we will now stop referring to actions, and instead writing the payoffs of the decision maker and the expert as functions of s, f, c.

3. Problem

In the present paper, we assume that the decision maker has two goals: When she faces an informed expert, she compares her payoff under incentive compatible contracts to her payoff under the scenario in which there are no incentive problems. She seeks menus of contracts under which the two payoffs are equal. Second, when she faces an uninformed expert, she compares her payoff to the default payoff obtained in the absence of any expert. She wants to ensure a minimal safety condition such that her payoffs will not fall below the default payoff.

Define

$$U^*(p) := \sup_c \sup_f E^p U(s, f, c) \text{ s.t. } E^p W(s, f, c) \geq W^0. \tag{3.1}$$

This is the highest possible payoff attained when the expert's information is equal to p and when the decision maker can choose (for the expert) the contract and the forecasts, but the individual rationality constraint is satisfied. Define $U^{*,ir}(p)$ similarly, except that the supremum over all contracts c is replaced by the supremum over irreversible contracts $c \in \mathcal{C}^{ir}$.

In the reversible case, it is easy to show that the supremum is attained by truthful forecasts, i.e., $f_t = p_t$, and contract $c^* = (w^*, e^*)$ such that $w_t^* \equiv w^0$ and $e_t^* = 1$ in periods t such that there is an action a_t with

$$E^{p_t} u_t(a_t, s_t) \geq \overline{w}. \tag{3.2}$$

Under such a contract, the expert provides a forecast in periods in which the forecasts are useful for the decision maker, i.e., when the expected benefit from the forecast is no smaller than the cost. If the opposite inequality holds for all actions a_t, then $e_t^* = 0$, and the decision maker takes the default action $a_t^* := a_t^0(s^t)$; recall that this default action yields a payoff equal to 0.

To compute the first-best payoff in the irreversible case, imagine that the expert truthfully reveals the entire stochastic process p up front in period 0; in addition, the decision maker decides when to stop taking optimal actions for the stochastic process p and to begin taking default actions instead. Up to that moment, she has to pay $\overline{w}$ in every single period. The optimal stopping rule is determined by solving a dynamic programming problem in which, contingent on each history s^t, the decision maker compares (1) her continuation payoff from taking default actions, and (2) her continuation payoff from taking the optimal action according to p in period t and continuing to apply her stopping rule from period $t+1$ on. It is easy to see that for any $\delta < 1$, optimal stopping rules exist for all $p = (p_t)_{t=1}^{\infty}$, and $U^{*,ir}(p)$ is equal to the payoff that is achieved by using an optimal stopping rule. It is also easy to see that optimal stopping rules may not exist for $\delta = 1$. In this case, we denote by $U^{*,ir}(p)$ the supremum of the decision maker's payoffs across all stopping rules. The exact expressions can be found in the proof of Proposition 1.

Take a $\lambda > 0$, and consider a menu of contracts $\mathcal{C}$. The expert makes several choices. He first chooses a contract (or rejects the entire menu); next, if he decided to accept a contract, he chooses forecasts to be provided in periods t in which $e_t = 1$. Because the expert maximizes his own payoffs, he might not have any incentive to choose either a contract or forecasts that maximize the expression (3.1). The informed expert's choices $c^* \in \mathcal{C}$ and f^* are called λ incentive compatible if no other contract in the menu and no other forecasts yield him a payoff higher than c^* and f^* by more than λ, i.e.,

$$E^p W(s, f^*, c^*) \geq \sup_{c \in \mathcal{C}} \sup_f E^p W(s, f, c) - \lambda.$$

We study λ incentive compatible choices of the expert, because for $\delta = 1$, optimal choices may not exist, whereas λ incentive compatible choices do exist for any positive λ. Of course, if optimal choices do exist, as they do when $\delta < 1$, then they are λ incentive compatible for any $\lambda > 0$.[5]

Contract c is $\varepsilon-$*safe* if for each realization of states s, each forecast f, the decision maker's payoff never falls more than ε below the default payoff, that is,

$$U(s, a, c) \geq -\varepsilon.$$

Of course, the default contract is $0-$safe. If the expert is uninformed, then even if he accepts a safe contract, the decision maker's payoff can fall only marginally below the payoff she would obtain in the absence of any expert.

Now, take an $\varepsilon > 0$. A menu of contracts $\mathcal{C}$ is ε first best if (1) it consists entirely of ε-safe contracts, and (2) there exists a $\lambda > 0$ such that *for any informed expert* p, if his choices c and f are λ incentive compatible, then the decision maker's payoff

[5] The nonexistence of optimal choices is typical for the literature on testing experts. This literature assumes that the informed expert reveals forecasts truthfully. However, it is often not optimal for the expert to reveal his information honestly; moreover, there often exist no forecasts that maximize the chance of passing a test.

falls by no more than ε below $U^*(p)$, that is,

$$E^p U(s, f, c) \geq U^*(p) - \varepsilon.$$

Irreversibly ε first best menus of irreversible contracts are defined similarly, except that $U^*(p)$ has to be replaced with $U^{*,ir}(p)$; of course, such a menu must consist of irreversible contracts.

Almost first best menus guarantee almost first best payoffs in a quite strong sense. The decision maker does not know whether the expert is informed or not; nor when he *is* informed what his information is. Nevertheless, if the expert is informed, the decision maker's payoff is close to what she would get if the expert honestly revealed his information, and if he got paid $\overline{w}$ only in periods in which the decision maker would like to rely on the expert's information.

If the expert is uninformed, i.e., he has no additional information about the process that generates states, the decision maker cannot expect to accomplish more than preventing the possibility that this type of expert will benefit at her own expense due to a contract she has offered. Because each contract in an ε first best menu is ε-safe, this indeed cannot happen (up to ε) at any single sequence of states. Equivalently, this cannot happen (up to ε) for any single process that generates the states.

Although we focus on the first best contracts, other constraints might be important in practice for the contracting parties. Two important ones are the following:

- In principal-agent problems, it is common to assume a limit on liability. Say that contract $c = (w, e)$ has *limited liability* if for each t, $w_t \geq 0$.

- Second, the contracts should not force the agent to work for the principal against his will. Let $\phi_t \in \{0, 1\}$ be any history-dependent function, which will be interpreted as the expert's decision as to whether to provide forecasts to the decision maker. For any contract $c = (w, e)$, define $e_t^\phi = \phi_t e_t$ and contract $c^\phi = \left(w, e^\phi\right)$. Say that a menu of contracts $\mathcal{C}$ allows for *free exit*, if for any ϕ_t and $c \in \mathcal{C}$, we have that $c^\phi \in \mathcal{C}$. (In the irreversible case, we assume that $\phi_t = 0$ implies $\phi_m = 0$ for any $m \geq t$.)

It turns out that none of these constraints affects our positive results, i.e., the results that establish the existence of first-best menus. All contracts used in our proofs, have limited liability; and we can ensure free exit by including, together with every contract c, all induced contracts c^{ϕ}.

4. No discounting

Our first result refers to the case in which the parties do not discount future payoffs. Consider the following countable family of (irreversible) contracts $c_n^{ir} = (w, e^{n,ir})$, $n = 1, 2, ...$: Under contract c_n^{ir}, $e_t^{n,ir} \equiv 1$ for $t = 1, ..., n$; the payments in the first n periods are irrelevant, since the parties evaluate sequences of payments by limits of averages. For $t > n$, $e_t^{n,ir} = 1$ if

$$\frac{1}{k}\sum_{m=1}^{k} u_m(a_m^{f_m}, s_m) > \overline{w}; \tag{4.1}$$

recall that $a_m^{f_m}$ is an optimal action of the decision maker, given the expert's forecast f_m. If (4.1) is not satisfied, then $e_t = 0$. The payment is defined by

$$w_t = \overline{w} + \theta\left[u_t(a_m^{f_m}, s_t) - \overline{w}\right],$$

where θ is a sufficiently small positive number.

The number n can be interpreted as the length of a "grace period"; if the expert selects contract c_n^{ir}, he cannot be dismissed in the first n periods. After the grace period is over, the expert is compensated for the outside option he has forgone, and receives a bonus, proportional to the utility generated by his forecast. This payment is received as long as his forecasts generate a positive surplus on average. Once the average surplus drops below zero, the expert is dismissed and no longer receives any payment.

Assumption 1. The values $u_t(a, s)$ are uniformly bounded, i.e., there exists a constant M such that

$$\forall_{t,a,s}\ |u_t(a, s)| \leq M.$$

Proposition 1. *Suppose that* $\delta = 1$, *and that Assumption 1 is satisfied. For each* $\varepsilon > 0$ *and* $\theta < \frac{\varepsilon}{4M}$, *the family* $\{c_n^{ir} : n = 1, 2, ...\}$ *is an irreversibly* ε *first best menu of contracts.*

The formal proof of Proposition 1 is relegated to the appendix. Informally, the argument can be explained as follows: Condition (4.1) guarantees that the uninformed expert will not benefit at the decision maker's expense, i.e., that each contract c_n^{ir} is 0-safe. On the other hand, since contracts with an arbitrarily long grace period are available, the patient expert whose forecasts generate a surplus on average is unlikely to be dismissed. Because of the incentives provided by the bonus, the expert's information is with high probability honestly revealed to the decision maker. As long as the bonus is sufficiently small, the expert takes only a small share of the surplus generated by his forecasts.

A simple modification of the above contracts leads to an ε first best menu of contracts for every $\varepsilon > 0$. The decision maker requires the expert to provide forecasts only in periods in which (3.2) holds. More precisely, let $\phi_t = 1$ if condition (3.2) is satisfied, and $\phi_t = 0$ otherwise. For each n, define $e_t^{n,\phi} := \phi_t e_t^{n,ir}$. Then, define $c_{n,\phi} = \left(w, e^{n,\phi}\right)$.

Proposition 2. *Suppose that* $\delta = 1$, *and that Assumption 1 is satisfied. For each* $\varepsilon > 0$ *and* $\theta < \frac{\varepsilon}{4M}$, *the family* $\{c_{n,\phi} : n = 1, 2, ...\}$ *is an* ε *first best menu of contracts.*

We omit the argument since ti follows very closely the proof of Proposition 1.

4.1. A comparison with the literature on testing strategic experts

We view Proposition 2 as a response to the recent literature on testing experts. Several papers show that an uniformed expert, who knows nothing about the stochastic process that generates states, can forecast strategically to pass empirical tests. (The most general results have been obtained by Olszewski and Sandroni (2008) and Shmaya (2008).)

An empirical test T is defined as an arbitrary function that takes as input forecasts $f = (f_t)_{t=1}^{\infty}$ and states $s = (s_t)_{t=1}^{\infty}$, and returns a verdict that is 0 or 1. When the test returns a 1, the expert (or his forecasts) pass the test. When a 0 is returned, the expert (or his forecasts) fail the test. Unlike the current paper, the literature on testing experts does not specify how the tester benefits from the expert's forecasts. Instead, it seeks tests such that: (a) if the informed expert truthfully reports $(p_t)_{t=1}^{\infty}$, he passes the test with high probability; and (b) no matter what the forecasts of the uninformed expert are, he fails the test with high probability (at least) on some sequences of states s.[6]

Olszewski and Sandroni (2008) study tests such that if the expert fails, he fails after a finite number of periods, i.e., $T(f, s) = 0$ implies that there exists an m such that $T(f, s^{'}) = 0$ if the first m states of s and $s^{'}$ coincide. They call a test *future-independent* if for any f and $f^{'}$ that coincide on histories of length m or shorter, a failure of f at a history s^m implies a failure of $f^{'}$ at the history s^m. Olszewski and Sandroni (2008) show that no future-independent test satisfies (a) and (b) simultaneously. This result implies that no test satisfying *the prequential principle*, i.e., no test such that the expert is required to provide forecasts only from period to period, satisfies (a) and (b) simultaneously. Shmaya (2008) shows that tests satisfying the prequential principle must violate either (a) or (b) even if they allow for failing the expert at infinity without failing him at any finite number of periods.

Tests that pass informed and truthful experts but fail uninformed experts do exist (see Dekel and Feinberg (2006), Olszewski and Sandroni (2009)), but they must be future-dependent and violate the prequential principle.

The first-best menus of contracts fulfill two goals. First, they lead to payoffs that approximate the first best payoffs. This corresponds to the requirement that the informed experts pass the test with high probability. Notice that in the present paper, the informed experts may not report forecasts truthfully, as is assumed in the existing literature. Second, we require that the uninformed expert cannot (too) negatively affect the decision maker's payoff. This corresponds to, but conceptually differs from, the requirement that the uninformed expert fails the test. The contracts

[6]The uninformed expert is allowed to produce forecasts randomly, and high probability refers to his random device.

provided in Proposition 1 require the expert to provide forecasts only from period to period, and so our contracts have the properties corresponding to future-independence and the prequential principle.

There is, however, a little caveat.[7] In our model, we assume that the expert knows enough about the future to evaluate the payoffs from his strategy, including the choice of a contracts and recommendations. In particular, when the informed expert chooses a contract from our first-best menu (i.e., he chooses the length of the grace period), this choice will depend on what he knows about the future. Therefore, our menu of contracts does not correspond to a test satisfying the prequential principle, although to a very mild extent (the expert chooses contract by declaring the length of the grace period that will be needed), and every single contract in the menu is "prequential."

Menus of contracts which extract full surplus and correspond to fully future-independent tests do exist, but in a slightly different version of the model. Suppose that the expert makes recommendations in each period, the decision maker decides whether to follow the expert's recommendations or take the status quo, and the expert receives payment $\overline{w}$ only if his recommendation is used. A simple argument based on the Hannan's Theorem shows that there exists (possibly, nondeterministic) decision maker's strategy that ensures that her payoff is no smaller than (a) the payoff from following the expert's recommendations and paying $\overline{w}$ in every period, and (b) the status quo payoff. Because the decision maker's payoff cannot be higher than (a) or (b), this strategy extracts full surplus. Note that this alternative framework limits the role of the expert. Indeed, because he is unable to predict the future, he cannot choose a contract, nor he can shape recommendations in order to maximize the future payoffs.[8]

Finally, although the model with discounting seems more appropriate for studying principal-agent relationships, no discounting is appropriate for the sake of comparing

[7] We are grateful to Eran Shmaya for pointing this out.

[8] The decision maker's strategy derived from the Hannan's Theorem will depend on the forecasts made in the periods in which the decision maker takes the status quo action. Thus, the Hannan's Theorem cannot be used to establish Proposition 2. Indeed, we assume that the expert provides recommendations only in periods in which he receives payments.

the results of the present paper to the literature on testing experts. Indeed, tests studied in the existing literature typically give a verdict only at infinity, and even if restrictions (such as "if the expert fails, he fails after a finite number of periods") are imposed, different finite numbers are treated equally.

4.2. Brier score

Meteorologists are required to meet a sufficiently low Brier score for their forecasts (see Brier (1950)). The Brier score of forecasts $(f_t)_{t=1}^{\infty}$ at period k is defined as

$$B_k := \frac{1}{k}\sum_{m=1}^{k}(f_m[s_m]-1)^2,$$

where $f_m[s_m]$ stands for the probability assigned by forecast f_m to state s_m realized in period m. Once the threshold of the Brier score is met, their compensation is partly based on other performance measures (e.g., calibration).

The contracts described in Section 3 mimic this idea. Indeed, after the grace period, the expert is required to satisfy condition (4.1) in order not to be dismissed. Assume that S is finite, $A = \Delta S$, and the decision maker has the quadratic utility:

$$u(a_m, s_m) = -(a_m[s_m]-1)^2 - \left(\frac{|S|-1}{|S|}\right)^2,$$

where $a_m[s_m]$ stands for the probability assigned by (mixed) action a_m to (pure) action s_m. Then, action $a_m = f_m$ will maximize $E^{f_m}u(a_m, s_m)$. Let the default action be a probability distribution that assigns equal probability to each state, that is, for any s^m and s_m,

$$a_m^0[s_m] = \frac{1}{|S|},$$

where $a_m^0(s^m) \equiv a_m^0$. After checking that the payoff normalization holds, we obtain that the inequality from condition (4.1) is equivalent to

$$B_k > \overline{w}.$$

According to our contracts, once the threshold of the Brier score is met, the expert is compensated by a bonus, proportional to the decision maker's payoff. This simple performance measure is sufficient for achieving our objectives (i.e., proving Proposition 1). However, we conjecture that in the proof of Proposition 1, the bonus can be replaced with several other performance measures, including calibration.

5. Discounting

Suppose now that the parties discount future payoffs by a common discount factor $\delta < 1$. Our first result shows that for any given discount factor, there exist no first-best menus of contracts. An action a is *rationalized by distribution* p_t if

$$E^{p_t} u_t(a, s_t) \geq E^{p_t} u_t(a^{'}, s_t) \text{ for any } a' \in A.$$

Because we assume that the default actions are undominated, they are rationalized by some distribution p.

Assumption 2. **(a)** There exist: a period t, a history $s^t = (s_1, ..., s_{t-1})$, a state s_t, and an action a_t such that

$$u_t(a_t, s_t) > \overline{w}.$$

(b) For every $m < t$, and for $s^m = (s_1, ..., s_{m-1})$, there exists probability distributions $p_m^0(s^m)$ rationalizing actions $a_t^0(s^m)$ such that $p_m^0(s^m)[s_m] > 0$.

Part (a) of the assumption guarantees that the expert's forecasts are minimally useful from the decision maker's perspective, that is, the expert's forecast generates a surplus at least at one history. Part (b) says that this history has a positive probability according to some distributions rationalizing default actions. The assumption is satisfied if the distribution that rationalizes the default actions has full support.

Proposition 3. *Under Assumption 2, for every $\delta < 1$, there exists an $\varepsilon > 0$ such that no menu of contracts is ε first best.*

The formal proof is relegated to the appendix. To provide a intuition for this result, assume that $\overline{w} = 0$, and that there is only one period. Consider a contract (a menu consisting of one contract) according to which the expert delivers a forecast, and receives payment

$$w_t := \theta u_t(a_t^{f_t}, s_t),$$

where $a_t^{f_t}$ is an optimal action given the forecast f_t. That is, the expert receives a bonus proportional to the surplus created by his forecast.

In this way, the incentives of the expert are allied with these of the decision maker. The expert accepts the contract only if his forecast is indeed useful, and delivers a truthful forecast. The expected utility of action $a_t^{f_t}$ when $f_t = p_t$ is at least as high as that from taking the default action a_t^0. Thus, payoff $U^*(p)$ can be $\varepsilon-$approximated provided that θ is sufficiently small. If θ is large, the informed expert retains more than an $\varepsilon-$share of the surplus, and the decision maker's payoffs will be bounded away from $U^*(p)$.

However, if θ is small, the expert bears only a small share of the decision maker's payoff from taking actions which are optimal according to his forecasts, and therefore the contract is typically not ε-safe.

Thus, if $\varepsilon > 0$ is sufficiently small, the one-contract menu is not ε first best for any θ. This argument generalizes to all menus of contracts.

Remark 1. *An analogous result holds for menus of irreversible contracts. However, Assumption 2 is insufficient. One must make the assumption that the expert's forecasts are minimally useful in the world in which dismissed experts cannot be hired again. Part (b) of Assumption 2 must also be properly modified. We will omit the version of Proposition 1 for menus of irreversible contracts. However, we show in Proposition 4 that in the world in which dismissed experts cannot be hired again, irreversibly first-best menus may not exist not only for a given discount factor, but even in the limit when* $\delta \to 1$.

For a given discount factor $\delta < 1$, the decision maker cannot attain payoffs $U^*(p)$ and $U^{*,ir}(p)$. However, by Propositions 1 and 2, she can approximate the first best outcome arbitrarily closely when $\delta = 1$. It is therefore relevant to see whether the payoff can approximate the first best in the limit case in which the contracting parties discount future payoffs but the discount factor tends to 1. Throughout the rest of this section, we will assume that $\overline{w} > 0$. We will show that for sufficiently large discount factors, the decision maker can attain payoffs $U^* (p)$ if she is allowed to offer reversible contracts; in contrast, payoffs $U^{*,ir} (p)$ cannot be attained in the world of irreversible contracts.

Assumption 3. **(a)** There exist a constant $\eta > 0$ such that for any t and history s^t, there exist a state s_t and an action a_t such that

$$u_t(a_t, s_t) - \overline{w} \geq \eta.$$

(b) There exists a constant $\mu > 0$, and a probability distribution $p_t^0(s^t)$ rationalizing actions $a_t^0(s^t)$ such that for all t and s^t, there exists a state s_t' such that

$$\mu \leq p_t^0(s^t)[s_t'] \leq 1 - \mu.$$

Part (a) of Assumption 2 says that in all periods and at any history, it is possible that the expert's forecast will be useful; moreover, this "usefulness" is uniformly bounded from below (by a constant η). This assumption reduces to part (a) of Assumption 2 for all time-independent utility functions, i.e., whenever $u_t(a_t, s_t) = u(a_t, s_t)$ for all t. Part (b) requires that the probability distributions rationalizing default actions are "uniformly" nondegenerated; if the decision maker's utility function is time-independent, it reduces to the requirement that distributions p_t^0 are simply nondegenerated.

Proposition 4. *Suppose that Assumption 3 is satisfied and $\overline{w} > 0$. There exist an $\varepsilon > 0$ and a $\overline{\delta} < 1$ such that for every $\delta \in (\overline{\delta}, 1)$, no menu of irreversible contracts is irreversibly ε first best.*

Proposition 3 shows that irreversibly ε first best menus of irreversible contracts may not exist for large discount factors in a broad range of circumstances. In contrast, an ε first best menu of contracts typically does exist if the discount factor is sufficiently large.

Proposition 5. *Suppose that Assumption 1 is satisfied. For every $\varepsilon > 0$, there exists a $\overline{\delta} < 1$ such that for every $\delta > \overline{\delta}$, there exists an ε first best menu of contracts.*

The formal proofs of Propositions 4 and 5 can be found in the appendix. The intuition for the two propositions can be easily explained by means of the following example: Suppose that there are only two periods $t = 1, 2$, two states $S = \{-1, 1\}$, and three actions $A = \{-1, 0, 1\}$. For all histories s^t, $t = 0, 1$, action $a_t^0(s^t) = 0$ is rationalized by the fifty-fifty probability distribution $p_t^0(s^t)$, and action $a = s$ is optimal in each of the two states s. Say that

$$u_t(s, s) = 6,\ u_t(-s, s) = -12,\ u_t(0, s) = 0$$

for $t = 1, 2$ and $s \in S$. Finally, let $\overline{w} = 1$.

Consider the informed expert who knows that p_1 is fifty–fifty; furthermore, $s_2 = s_1$ with probability 1 if $s_1 = 1$, and $p_2 = p_1$ if $s_1 = -1$. Denote this type of informed expert by I_1. To guarantee the first best payoff (against expert of type I_1), a contract c must have $e_2(-1) = 0$ and $e_2(1) = 1$, i.e., the expert must be dismissed in period 2 contingent on state -1 in period 1, but must not be dismissed contingent on state 1.

To ensure ε-safety, the payment of the decision maker to the expert contingent on state -1 must be close to 0. Otherwise, the uninformed expert could accept contract c, and the decision maker would make a (too) negative payoff contingent on any history s such that $s_1 = -1$. Thus, the expected payment to the expert contingent on state 1 must be close to $3\overline{w}$. Indeed, the expert of type I_1 receives this payment only with probability $\frac{1}{2}$; therefore, in order for that expert to prefer accepting the contract to rejecting the entire menu, he must receive at least $2\overline{w}$ in addition to $\overline{w}$ from the external source received contingent on state -1 in period 1.

Consider now the informed expert who knows that $s_1 = s_2 = 1$ with probability 1. Denote this type of informed expert by I_2. This expert can also accept contract c and provide forecasts identical to those of the expert of type I_1. Then, he receives payment $3\overline{w}$ with probability 1, and the decision maker's payoff is bounded away from the first best.

Notice now that the problem just described disappears in the world of reversible contracts. Indeed, if there were only three types of expert - I_1, I_2, and the uninformed expert - a contract (a menu consisting of one contract) that attains payoffs $U^*(p)$ could be defined as follows: $e_1 = 0$, $e_2(-1) = 0$, and $e_2(1) = 1$; the expert is paid slightly more than $\overline{w}$ contingent on $s_2 = 1$, and is paid nothing contingent on $s_2 = -1$. In general, the proof of Proposition 5 is more complicated, but involves arguments somewhat similar to those used in the proof of Proposition 1.

6. Appendix

6.1. Proof of Proposition 1

If the expert (informed or uninformed) accepts a contract, and is fired in a period t, the decision maker takes default actions, and makes no payment from period $t+1$ on. In such a case, the decision maker's payoff is equal to the default payoff, which is normalized to 0. If the expert is not fired at any t, (4.1) holds for every $T = 1, 2, ..,$ and

$$U^T\left(s, f, c_n^{ir}\right) = (1-\theta)\left(\frac{1}{T}\sum_{t=1}^{T} u_t(a_t^{f_t}, s_t) - \overline{w}\right) \geq 0,$$

where $a_t^{f_t}$ is the best response to forecast f_t. Thus, the long-run payoff never falls below 0, and so contract c_n^{ir} is $0-$safe.

The rest of the proof consists of two parts. First, we show that there exists an n such that if the expert selects contract c_n^{ir}, and reports forecasts truthfully, then the decision maker's payoff falls below $U^{*,ir}(p)$ by no more than $\frac{\varepsilon}{2}$. Recall that $U^{*,ir}(p)$ is defined as the first-best payoff attained by the decision maker, if individual-rationality constraints, but not incentive-compatibility constraints, are satisfied. The exact value of $U^{*,ir}(p)$ is computed in the course of the proof. Second, we show that any almost incentive-compatible choice of the contract approximates the first-best payoff.

Given $p = (p_t)_{t=1}^{\infty}$, let $\Sigma(p)$ be the set of all sequences s such that

$$\liminf_T \frac{1}{T} \sum_{t=1}^{T} u_t(a_t^{p_t}, s_t) > \overline{w},$$

Set $\Sigma(p)$ consists of all sequences of states along which it is never beneficial to fire the truthful expert. For each sequence s, define

$$U^{ir}(s,p) := \begin{cases} \liminf_T \frac{1}{T} \sum_{t=1}^{T} u_t(a_t^{p_t}, s_t) - \overline{w}, & \text{if } s \in \Sigma(p); \\ 0, & \text{otherwise.} \end{cases}$$

Because the decision maker benefits from the expert's truthful forecasts only if $s \in \Sigma(p)$, it follows that

$$E^p U^{ir}(s,p) \geq U^{*,ir}(p). \tag{6.1}$$

For all n, define the set $\Sigma_n(p)$ of all sequences such that

$$\frac{1}{T} \sum_{t=1}^{T} u_t(a_t^{p_t}, s_t) > \overline{w} \text{ for all } T > n.$$

Set $\Sigma_n(p)$ consists of sequences along which the truthful expert is not fired by contract c_n^{ir}.

For every sequence of states s such that $s \notin \Sigma(p)$, we have that $U(s,p,c_n^{ir}) = U^{ir}(s,p) = 0$. On the other hand, for every sequence of states s such that $s \in \Sigma(p)$, there exists a natural number $n(s)$ such that $s \in \Sigma_n(p)$ for all $n \geq n(s)$. This implies that the expert is never fired under contract c_n^{ir} for all $n \geq n(s)$; therefore, $U(s,p,c_n^{ir}) = (1-\theta) U^{ir}(s,p)$.

Given a positive number $\mu > 0$, there exists a natural number n^* such that the probability, according to probability distribution p, of the set of all sequences of states s such that $n(s) < n^*$ is no lower than $1-\mu$. Thus, if the expert selects this contract $c_{n^*}^{ir}$, and reports his forecasts truthfully, the probability of $\Sigma(p) \backslash \Sigma_{n^*}(p)$ is no higher than μ. That is, the probability that the expert will be fired, although the decision maker's payoff would be higher if he were not, is bounded by μ. This shows that

$$(1-\theta) E^p U^{ir}(s,p) - \mu M \leq E^p U\left(s,p,c_{n^*}^{ir}\right) \leq E^p U^{ir}(s,p).$$

Since the inequalities hold for any θ and μ (and an appropriately chosen n^*), we have that

$$E^p U^{ir}(s,p) \leq \sup_n U\left(s,p,c_n^{ir}\right) \leq U^{*,ir}(p),$$

which, taken together with (6.1), yields that $U^{*,ir}(p) = E^p U^{ir}(s,p)$.

Since $\theta < \frac{\varepsilon}{4M}$, if $\mu \leq \frac{\varepsilon}{4M}$, then the loss in the decision maker's payoff under contract $c_{n^*}^{ir}$ (compared to $U^{*,ir}(p)$) is no higher than $\mu M + \theta M$, which is no higher than $\frac{\varepsilon}{2}$.

Of course, the expert need not report forecasts truthfully. Notice, however, that the bonus aligns the expert's and the decision maker's payoffs: under any contract c_n^{ir}, and for any sequence of reported forecasts f,

$$\begin{aligned}
W\left(s,f,c_n^{ir}\right) &= \overline{w} + \theta \liminf \frac{1}{T}\sum_{t=1}^{T} e_t\left(u_t\left(a_t,s_t\right) - \overline{w}\right) \\
&= \overline{w} + \frac{\theta}{1-\theta}\liminf \frac{1}{T}\sum_{t=1}^{T}\left(u_t\left(a_t,s_t\right) - e_t\left[\overline{w} + \theta\left(u_t\left(a_t,s_t\right) - \overline{w}\right)\right]\right) \\
&= \overline{w} + \frac{\theta}{1-\theta} U\left(s,f,c_n^{ir}\right).
\end{aligned}$$

The second equality follows from the fact that $u_t(a_t,s_t) = 0$ whenever $e_t = 0$. Since the expert makes λ incentive compatible choices, his expected payoff can be no lower than

$$\overline{w} + \frac{\theta}{1-\theta}\left(U^{*,ir}(p) - \frac{\varepsilon}{2}\right) - \lambda;$$

therefore, the decision maker's payoff can be no lower than

$$U\left(s,f,c_n^{ir}\right) = \frac{1-\theta}{\theta}\left(W\left(s,f,c_n^{ir}\right) - \overline{w}\right) \geq U^{*,ir}(p) - \frac{\varepsilon}{2} - \frac{1-\theta}{\theta}\lambda.$$

If we take $\lambda = \frac{\varepsilon\theta}{2(1-\theta)} > 0$, then $U\left(s,f,c_n^{ir}\right) \geq U^{*,ir}(p) - \varepsilon$.

6.2. Proof of Proposition 3

Suppose to the contrary that there exists an ε first best menu for every $\varepsilon > 0$. Let t, s^t, and a_t have the properties described in part (a) of Assumption 2. Since action

$a_t^0(s^t)$ is undominated, there exists a state $s_t^{'}$ such that

$$u_t(a_t, s_t^{'}) < 0. \tag{6.2}$$

Perturbing (if necessary) the degenerated probability distribution that assigns measure 1 to s_t, we obtain a probability distribution $p_t^{'}$ which assigns a positive measure to $s_t^{'}$, and such that

$$E^{p_t} u_t(a_t, s_t) > \overline{w}.$$

For simplicity, assume that p_t has two atoms, s_t and $s_t^{'}$, and assigns measure 0 to all sets containing neither of the two atoms. Replacing (again, if necessary) a_t with another action, we may in addition assume that $a_t = a_t^{p_t}$, i.e., a_t maximizes $E^{p_t} u_t(a_t, s_t)$.

Let p_m^0 be a sequence of history-dependent probability distributions rationalizing default actions, which have the properties described in part (b) of Assumption 2 along history s^t.

Consider the informed expert of type $p = (p_m)_{m=1}^{\infty}$ such that $p_m = p_m^0$ in all states and at all histories, except history s^t; and $p_t = p_t^{'}$ contingent on history s^t. Denote this type of expert by I_1. If $\varepsilon < (1-\delta)\delta^{t-1}\,\mathrm{Prob}(s^t)\,(E^{p_t} u_t(a_t, s_t) - \overline{w})$, where $\mathrm{Prob}(s^t) > 0$ is calculated according to p, then any ε first best menu has to contain a contract c which is accepted by this type of informed expert. Under this contract, e_t has to be equal to 1 after history s^t. Otherwise, the decision maker would not benefit from the expert's advice at history s^t, and then his advice is useful. Consequently, payoff $U^*(p)$ would not be attained.

Since the expert of type I_1 can choose the default contract, his total payoff when he accepts contract c cannot be lower than $\overline{w} - \lambda$; otherwise, his choice of contract would not be λ incentive compatible. However, the expected (discounted and normalized) payment of the decision maker to the informed expert of type I_1 may not be higher than $(1-\delta)\delta^{t-1}\overline{w} + \varepsilon$; otherwise, the decision maker could not attain (up to ε) payoff $U^*(p)$. Thus, the expert must receive $\overline{w} - (1-\delta)\delta^{t-1}\overline{w} - \varepsilon$ from the external source.

Consider now the uninformed expert who predicts that states are generated according to p. It follows from $\varepsilon-$safety that if this expert accepts contract c, then for any sequence of states such that $s_t = s_t^{'}$, the (discounted and normalized) payment to the expert does not exceed

$$(1-\delta)\delta^{t-1}u_t(a_t, s_t^{'}) + \varepsilon.$$

This number is negative, and if $\varepsilon < -\frac{1}{2}(1-\delta)\delta^{t-1}u_t(a_t, s_t^{'})$, it is lower than $(1-\delta)\delta^{t-1}\overline{w}$ by at least $\beta := -\frac{1}{2}(1-\delta)\delta^{t-1}u_t(a_t, s_t^{'})$.

Therefore, since the probability distribution p_t assigns positive probability only to sets containing s_t or $s_t^{'}$, the expected (discounted and normalized) payment to the expert contingent on history s^t, s_t must exceed $(1-\delta)\delta^{t-1}\overline{w}$ by $\beta - \lambda$.

Consider now the informed expert of type $q = (q_m)_{m=1}^{\infty}$, where the only difference between p and q is that q_t contingent on history s^t is the probability distribution with an atom of mass 1 at s_t. Denote this type of the expert by I_2. The informed expert of type I_2 can choose contract c, which guarantees him the total expected payoff that exceeds $\overline{w}$ by $\beta - \lambda - \varepsilon$. Indeed, he receives at least $(1-\delta)\delta^{t-1}\overline{w} + (\beta - \lambda)$ with probability 1 from the decision maker and $\overline{w} - (1-\delta)\delta^{t-1}\overline{w} - \varepsilon$ from the external source. Thus, the menu cannot be ε first best, if λ is sufficiently small.■

6.3. Proof of Proposition 4

Take a number $\alpha \in (0, 1)$ such that

$$\alpha\,(1-\mu)\,\eta > (1-\alpha)\overline{w},$$

where η is the positive number from Assumption 3. Then, there exists a number $\kappa > 0$, such that for all sufficiently large δ, there exist a t such that

$$\delta^{t+1}\,(1-\mu)\,\eta > (1-\delta^{t+1})\overline{w} \geq \kappa. \tag{6.3}$$

Indeed, it follows from the fact that for every interval $\alpha \in (\underline{\alpha}, \overline{\alpha})$, if δ is sufficiently large, then there exists a t such that $\delta^{t+1} \in (\underline{\alpha}, \overline{\alpha})$.

By part (a) of Assumption 3, for each period and past history s^t, there exists an $\eta > 0$, an action $a_t^*(s^t)$, and a state $s_t^*(s^t)$ such that $u_t\left(a_t^*(s^t), s_t^*(s^t)\right) - \overline{w} > \eta$. Let $p^* = (p_t^*)_{t=1}^{\infty}$ denote the sequence of history-dependent probability distributions such that $p_t^*(s^t)\left[s_t^*(s^t)\right] = 1$ for all histories s^t.

Let $s_t^{'}(s^t)$ denote the state such that $\mu < p_t^0(s^t)[s_t^{'}(s^t)] \leq 1-\mu$; the existence of state $s_t^{'}(s^t)$ with the required property is guaranteed by part (b) of Assumption 3.

Consider an informed expert of type $p = (p_m)_{m=1}^{\infty}$ such that $p_t = p_t^0$ up to and including period t. If $s_t = s_t^{'}(s^t)$, then $p_t = p_t^0$ also in periods $t+1, t+2, \ldots$. If $s_t \neq s_t^{'}(s^t)$, then beginning in period $t+1$, $p_t = p_t^*$. Notice that we can assume without loss of generality that $a_t^*(s^t)$ is rationalized by $p_t^*(s^t)$; indeed, an action $a_t^{p_t}$ that maximizes $E^{p_t^*}u_t(a_t, s_t)$ also has the required property. Denote this type of informed expert by I_1.

The first part of (6.3) guarantees that the advice of this type of expert is valuable, if this advice is elicited in periods $1, \ldots, t$, and then in periods $t+1, t+2, \ldots$, contingent on $s_t \neq s_t^{'}(s^t)$, but the expert is dismissed contingent on $s_t = s_t^{'}(s^t)$. Suppose that $\varepsilon < \delta^{t+1}(1-\mu)\,\eta - (1-\delta^{t+1})\overline{w}$. To be irreversibly ε first best, a menu of irreversible contracts must contain a contract c such that $e_m := 1$ for all $m \leq t$.

To satisfy $\varepsilon-$safety, the (discounted and normalized) payment from the decision maker to the expert, contingent on sequences s such that $s_t = s_t^{'}(s^t)$, may not exceed ε. Otherwise, the uninformed expert could accept contract c, and predict p_t^0 for all t and s^t; the decision maker's payoffs would then fall below $-\varepsilon$ contingent on some history s such that $s_t = s_t^{'}(s^t)$. Let x be the expected (discounted and normalized) payment to the expert contingent on sequences s such that $s_t \neq s_t^{'}(s^t)$. Together with at most ε from the decision maker and $\delta^{t+1}\overline{w}$ from the external source (both received contingent histories s such that $s_t = s_t^{'}(s^t)$), the expert must receive at least $\overline{w}-\lambda$ in order for his choice to be λ incentive compatible. Thus,

$$\overline{w} - \lambda \leq \mathrm{Prob}[s_t = s_t^{'}\left(s^t\right)](\delta^{t+1}\overline{w} + \varepsilon) + \mathrm{Prob}[s_t \neq s_t^{'}\left(s^t\right)]x,$$

where $\mathrm{Prob}[s_t = s_t^{'}\left(s^t\right)]$ and $\mathrm{Prob}[s_t \neq s_t^{'}\left(s^t\right)]$ are calculated according to p^0.

Since $p_t^0[s_t = s_t^{'}\left(s^t\right)] \leq 1-\mu$, if ε and λ are sufficiently small, then the inequality from the last display can be satisfied only if x exceeds $\overline{w}$ by a positive number which does not vanish as the discount factor δ tends to 1. Consider now the informed expert I_2 who knows that $s_t \neq s_t^{'}(s^t)$ with probability 1. The expert of type I_2 can also accept contract c and provide forecasts identical to those of expert I_1. He then receives payment $x > \overline{w}$ with probability 1. This means that, for sufficiently small $\varepsilon > 0$, the decision maker does not attain payoff $U^{*,ir}(p)$ against the expert of type I_2.■

6.4. Proof of Proposition 5

In order to prove Proposition 5, we will need a lemma, which is stated below, and which follows directly from Azuma's inequality. A sequence of random variables $(X_k)_{k=1}^{\infty}$ is *uniformly bounded* by a constant C if

$$|X_k| \leq C$$

for every $k = 1, 2, ...$ Let

$$S_m := \sum_{k=1}^{m} X_k \ .$$

Sequence $(S_k)_{k=1}^{\infty}$ is called a *supermartingale* if

$$E[S_{m+1} \mid S_1, ..., S_m] \geq S_m,$$

i.e., if

$$E[X_{m+1} \mid X_1, ..., X_m] \geq 0$$

for $k = 1, 2, ...$

Lemma 1. *If $(X_k)_{k=1}^{\infty}$ is a sequence of random variables uniformly bounded by a constant C, $(S_k)_{k=1}^{\infty}$ is a supermartingale, and d is a positive constant, then*

$$\text{Prob}\left(\sum_{k=1}^{m} X_k \leq -d\right) \leq 2\exp\left(-\frac{d^2}{32mC^2}\right). \tag{6.4}$$

We can now prove Proposition 5. Let

$$\overline{C} = M + \overline{w}$$

where M is the number from Assumption 1. Assume that

$$\varepsilon \leq \frac{3}{4}\overline{C}. \tag{6.5}$$

We first describe a family of menus of contracts $\mathcal{C}_T$, $T = 1, 2, ...$ Fix any $T = 1, 2, ...$ The payment w under each contract c in menu $\mathcal{C}_T$ is the same as that used to define ε first best contracts in Section 4, i.e.,

$$w_t = \overline{w} + \theta \left[u_t(a_t, s_t) - \overline{w} \right]. \tag{6.6}$$

Assume that

$$\theta > 0 \text{ and } \theta < \frac{\varepsilon}{6M}.$$

Let e be any function which has the following property: in periods $t = T, 2T, ...$, the expert's forecasts are reviewed by the decision maker. Say that the expert *fails the* $k-$*th review* if

$$\sum_{t=(k-1)T+1}^{kT} e_t \left(u_t(a_t^{f_t}, s_t) - \overline{w} \right) < -\frac{\varepsilon}{3} T \tag{6.7}$$

where $a_t^{f_t}$ maximizes $E^{f_t} u_t(a_t, s_t)$. Under each contract c in menu $\mathcal{C}_T$, $e_t = 0$ if the expert failed a $k-$th review such that $kT < t$. That is, if the expert has passed all previous reviews, he is free to choose in which periods of the current review he wants to deliver a forecast. Once the expert fails a review, he no longer delivers any forecast.

We shall now show that if

$$\delta^T \geq 1 - \frac{\varepsilon}{3\overline{C}}, \tag{6.8}$$

then every contract $c \in \mathcal{C}_T$ is ε-safe.

Indeed, if the expert passes a $k-$th review, then the difference between the decision maker's payoff under the contractual arrangement and her payoff to taking default actions cannot on average - across periods $(k-1)\,T+1, ..., kT$ - be larger than $\frac{\varepsilon}{3} + \theta M < \frac{\varepsilon}{2}$; the expression θM estimates the bonus $\theta \left[u_t(a_t, s_t) - \overline{w} \right]$ received by the expert. So, the discounted and normalized loss in the decision maker's payoff caused between periods $(k-1)\,T$ and kT by contract c cannot be larger than $(1-\delta)\delta^{(k-1)T}\frac{\varepsilon}{2}T$.

If the expert fails a $k-$th review, the decision maker's loss (during that review) is not larger than $\delta^{(k-1)T}\overline{C}\left(1-\delta^{T}\right)$, and the expert is never hired again. Thus, the decision maker's loss in all periods cannot be larger than

$$\overline{C}\left(1-\delta^{T}\right)+(1-\delta)\frac{1}{1-\delta^{T}}\frac{\varepsilon}{2}T\leq\overline{C}\left(1-\delta^{T}\right)+\frac{\varepsilon}{2\delta^{T}},$$

and this number is no larger than ε by (6.8) and (6.5).

We shall now show that there exists a $\overline{T}$ such that if $T\geq\overline{T}$ and

$$\delta^{T}\leq 1-\frac{\varepsilon}{6\overline{C}}, \tag{6.9}$$

then for every $p=(p_t)_{t=1}^{\infty}$, menu $\mathcal{C}_T$ attains $U^*(p)$ provided the expert makes λ incentive compatible choices for sufficiently small λ.

Given a p, for all t and s^t, define

$$a_t^* := \begin{cases} a_t^{p_t}\in\arg\max_a E^{p_t}u_t\left(s_t,a\right) & \text{if} \quad E^{p_t}u_t\left(s_t,a\right)>\overline{w}; \\ \\ a_t^0 & \text{otherwise;}\end{cases}$$

and let

$$w_t^* := \begin{cases} \overline{w} & \text{if} \quad E^{p_t}u_t\left(s_t,a\right)>\overline{w}; \\ \\ 0 & \text{otherwise.}\end{cases}$$

Then,

$$U^*\left(p\right)=(1-\delta)\,E\sum\nolimits_{t=1}^{\infty}\delta^t\left[u_t\left(a_t^*,s_t\right)-w_t^*\right].$$

We prove first that there exists a contract $c(p)\in\mathcal{C}_T$ such that

$$E^pU\left(s,p,c\left(p\right)\right)\geq U^*\left(p\right)-\frac{\varepsilon}{2}. \tag{6.10}$$

That is, if the informed expert accepts this contract and reports his forecasts truthfully, then the decision maker obtains utility $U^*\left(p\right)-\frac{\varepsilon}{2}$.

Consider the contract $c\left(p\right)$ such that if the expert has passed all previous reviews, $e_t=1$ if and only if $E^{p_t}u_t(a_t^{p_t},s_t)\geq\overline{w}$. Then, the actions taken under this contract (if the expert's forecasts are truthful) may differ from a_t^* only after the expert fails a review. This implies that $E^p U(s, p, c(p))$ cannot be smaller than $U^*(p)$ by more than

$$M \sum_{k=1}^{\infty} \delta^{kT} \operatorname{Prob}\left(\text{failure in } k-\text{th review} \mid \text{success in } k'-\text{th review for all } k' < k\right). \tag{6.11}$$

And by (6.4),

$$\operatorname{Prob}\left(\text{failure in } k-\text{th review} \mid \text{success in } k'-\text{th review for all } k' < k\right) =$$

$$\operatorname{Prob}\left(\begin{array}{c} \sum_{t=(k-1)T+1}^{kT} e_t\left(u_t(a_t^{f_t}, s_t) - \overline{w}\right) \leq -\frac{\varepsilon}{3} T \\ \mid \text{he did not fail any } k'-\text{th review for } k' < k \end{array}\right) \leq$$

$$2\exp\left(-\frac{1}{32\overline{C}^2}\left(\frac{\varepsilon}{3}\right)^2 T\right).$$

Now, let $\overline{T}$ be such that

$$2\exp\left(-\frac{1}{32\overline{C}^2}\left(\frac{\varepsilon}{3}\right)^2 \overline{T}\right) \leq \frac{\varepsilon^2}{24M\overline{C}};$$

Then, if $T \geq \overline{T}$, expression (6.11) can be estimated by

$$2M\frac{1}{1-\delta^T}\frac{\varepsilon^2}{24M\overline{C}} \leq \frac{\varepsilon}{2}.$$

Observe now that under any contract $c \in \mathcal{C}_T$, the normalized present value of the payments received by the expert is equal to

$$\begin{aligned}
& W(s, f, c) \\
= \quad & (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\left[e_t\overline{w} + \theta\left(u_t\left(a_t^{f_t}, s_t\right) - \overline{w}\right) + (1-e_t)\overline{w}\right] \\
= \quad & \overline{w} + (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}e_t\theta\left(u_t\left(a_t^{f_t}, s_t\right) - \overline{w}\right) \\
= \quad & \overline{w} + \frac{\theta}{1-\theta}U(s, f, c)
\end{aligned}$$

where the last equality follows from the fact that $u_t(a_t^*, s_t) = 0$ whenever $e_t = 0$.

That is, the informed expert maximizes the decision maker's utility. In particular, if c and f are λ incentive compatible, then, because of (6.10), his expected payoff can be no lower than

$$\overline{w} + \frac{\theta}{1-\theta}\left(U^*(p) - \frac{\varepsilon}{2}\right) - \lambda,$$

and the decision maker's payoff can be no lower than

$$U(s,f,c) = \frac{1-\theta}{\theta}\left(W(s,f,c) - \overline{w}\right) \geq U^*(p) - \frac{\varepsilon}{2} - \frac{1-\theta}{\theta}\lambda.$$

If $\lambda \leq \frac{\theta\varepsilon}{2(1-\theta)}$, then $U(s,f,c) \geq U^*(p) - \varepsilon$.

To complete the proof, we will show that there exists a $\overline{\delta} < 1$ such that for every $\delta \geq \overline{\delta}$, there exists a $T \geq \overline{T}$ such that conditions (6.8) and (6.9) are satisfied. Let

$$\overline{\delta} = \max\left\{\overline{\delta}_1, \overline{\delta}_2\right\},$$

where

$$\overline{\delta}_1 = \frac{1-\frac{\varepsilon}{3\overline{C}}}{1-\frac{\varepsilon}{6\overline{C}}}, \text{ and } \overline{\delta}_2^{\overline{T}} = 1 - \frac{\varepsilon}{3\overline{C}}.$$

For a given $\delta \geq \overline{\delta}$, take the largest T such that condition (6.8) is satisfied. Since $\delta \geq \overline{\delta}_2$, this T must be at least as large as $\overline{T}$. If condition (6.9) were not satisfied, then we would have that

$$\delta^{T+1} \geq \delta^T \cdot \overline{\delta}_1 \geq 1 - \frac{\varepsilon}{3\overline{C}},$$

which would contradict the definition of T.■

References

[1] Brier, G. (1950), "Verification of Forecasts Expressed in Terms of Probability," *Monthly Weather Review* **78**, 1-3.

[2] Cesa-Bianchi, N. and G. Lugosi (2006), *Prediction, Learning and Games*, Cambridge University Press.

[3] Dekel, E. and Y. Feinberg (2006), "Non-Bayesian Testing of a Stochastic Prediction," *Review of Economic Studies*, **73**, 893 - 906.

[4] Echenique F. and E. Shmaya (2008), "You won't harm me if you fool me," mimeo.

[5] Foster, D. and R. Vohra (1998), "Asymptotic Calibration," *Biometrika*, **85**, 379 - 390.

[6] Foster, D. and R. Vohra (1999), "Regret in the On-Line Decision Problems," *Games and Economic Behavior*, **29**, 7-35.

[7] Fudenberg, D. and D. Levine (1999), "An Easier Way to Calibrate," *Games and Economic Behavior*, **29**, 131 - 137.

[8] Hannan, J. (1957), "Approximation to Bayes Risk in Repeated Plays" in *Contributions to the Theory of Games* (M. Dresher, A.W. Tucker, and P. Wolfe, Eds.), Princeton University Press, Princeton, NJ.

[9] Lehrer, E. (2001), "Any Inspection Rule is Manipulable," *Econometrica*, **69**, 1333 - 1347.

[10] Olszewski, W. and A. Sandroni (2007), "Contracts and Uncertainty," *Theoretical Economics*, **2**, 1 - 13.

[11] Olszewski, W. and A. Sandroni (2008), "Manipulability of Future-Independent Tests," *Econometrica*, **76**, 1437-1466.

[12] Olszewski, W. and A. Sandroni (2009), "A Nonmanipulable Test," *Annals of Statistics*, **37**, 1013-1039.

[13] Sandroni, A. (2003), "The Reproducible Properties of Correct Forecasts," *International Journal of Game Theory*, **32**, 151 - 159.

[14] Sandroni, A., R. Smorodinsky, and R. Vohra (2003), "Calibration with Many Checking Rules," *Mathematics of Operations Research,* **28***, 141 - 153.*

[15] Shmaya, E. (2008), "Many Inspections are Manipulable," *Theoretical Economics*, **3**, 367-382.

[16] Vovk, V. and G. Shafer (2005), "Good Randomized Sequential Probability Forecasting is Always Possible," *Journal of the Royal Statistical Society Series B*, **67**, 747 - 763.

[17] Wald, A. (1949), "Statistical Decision Functions,"*Annals of Mathematical Statistics*, **20**, 165–205.

[18] Wald, A. (1950), *Statistical Decision Functions*, John Wiley and Sons, London.